

HAPI FOUNDATION PAPER

# Governance as the Product of Agency Preservation

*A HAPI Foundation Thesis on Real Governance, Agency Theater, and Institutional Restoration*

Human Agency Preservation Infrastructure (HAPI) Foundation Paper

Author: Michael Bower

Date: May 18, 2026

Status: Working theory manuscript, not peer reviewed

**Core thesis: Governance becomes real when the conditions of meaningful human agency are preserved. When agency is lost, governance becomes theater.**

---

## Abstract

This paper develops the HAPI thesis that governance is the product of agency preservation. It argues that governance does not become real through rules, policies, approvals, dashboards, audits, or compliance artifacts alone. Governance becomes real when the humans inside a system retain meaningful participation at the points where consequence can still be changed. The paper distinguishes real governance from governance theater, defines the agency conditions required for legitimate governance, and explains how clarity, authority, refusal, revision, memory, contestability, capacity, and accountability produce durable governance. It applies the model to institutions and agentic AI systems, showing why human oversight fails when humans remain present but lose the power to understand, refuse, revise, or contest. The conclusion proposes HAPI as a theory, audit method, restoration model, and infrastructure project for rebuilding systems from agency preservation outward.

**Keywords:** human agency, governance, agency preservation, agency restoration, governance theater, institutional design, agentic AI, accountability, human oversight, HAPI

# Table of Contents

1. Introduction
2. The Governance Problem
3. Working Definition
4. Why Governance Cannot Start With Control
5. Governance Theater
6. The Agency Conditions of Real Governance
7. The Governance From Agency Model
8. Governance and Authority
9. Governance and Refusal
10. Governance and Memory
11. Governance and Accountability
12. Governance and Values
13. Governance in Agentic AI
14. Governance in Institutions
15. True Governance Versus False Governance
16. The HAPI Governance Test
17. From Governance Product to HAPI Practice
18. Conclusion

# 1. Introduction

Modern systems often begin governance from the wrong place. They begin with rules, policies, dashboards, approvals, risk scores, audits, escalation paths, and compliance artifacts. These tools can be useful, but they do not create real governance by themselves. A system can have all of them and still leave humans unable to understand, refuse, revise, remember, contest, or meaningfully participate.

HAPI begins deeper. It begins with agency.

The central claim of this paper is that governance is the product of agency preservation. Governance becomes real when the humans inside a system retain the conditions required for meaningful participation. When those conditions are missing, governance becomes theater. People may still be present. They may still sign, approve, attend, review, acknowledge, consent, or be held accountable. But if they cannot meaningfully affect the outcome before consequence, their presence is symbolic rather than governing.

This paper builds on the previous HAPI foundation papers, *The Theory of Agency*, *Agency Loss*, and *Agency Restoration*. It argues that governance should not be understood first as control over humans. Governance should be understood as the structure that protects the conditions of human agency.

The core thesis is simple:

If agency is preserved, governance becomes real. If agency is lost, governance becomes theater.

## 2. The Governance Problem

The word governance is often used to describe the visible control structure of a system. A company has policies. A hospital has consent forms. A school has procedures. A religious institution has discipline. A government has law. An AI system has oversight. A workflow has approval rules. These structures appear to answer the question: who controls what?

But the deeper question is: can the humans who are responsible, affected, or authorized still meaningfully participate?

Many systems satisfy the surface form of governance while failing the agency condition beneath it. They include people but do not empower them. They ask for approval but do not preserve refusal. They assign responsibility but withhold authority. They keep records but not memory that humans can use. They claim transparency while leaving the real operating logic unreadable. They preserve compliance while losing conscience.

This is the governance problem HAPI addresses.

Governance is not real merely because a system has a process. Governance is real when the process preserves human authority, judgment, refusal, revision, memory, accountability, and contestability at the points where consequence can still be changed.

### 3. Working Definition

HAPI defines governance as the durable structure by which agency is preserved, directed, constrained, remembered, corrected, and made accountable across a system.

This definition differs from governance as mere control. Control can exist without agency. A prison has control. A bureaucracy has control. A social scoring system has control. An automated workflow can have control. But control alone does not produce legitimate governance.

Real governance requires rightful authority connected to meaningful participation.

In HAPI terms, governance has five core functions:

1. It preserves agency by keeping humans meaningfully present at consequential boundaries.
2. It directs agency by clarifying roles, purposes, values, and authority.
3. It constrains agency by preventing harmful, unauthorized, or incoherent action.
4. It remembers agency by preserving proof, context, decisions, refusals, revisions, and outcomes.
5. It restores agency by learning from breakdowns and rebuilding the conditions that were lost.

Governance is therefore not the opposite of agency. Real governance is the structure that lets agency remain coherent under pressure.

### 4. Why Governance Cannot Start With Control

When governance starts with control, humans become objects of management. Their behavior is measured, constrained, routed, scored, and audited. Sometimes this is necessary, especially when systems involve safety, law, finance, medicine, security, or institutional responsibility. But if the system begins and ends with control, it can easily become agency-eroding.

Control-first governance asks:

- How do we make people comply?
- How do we reduce risk exposure?
- How do we document approval?
- How do we protect the institution?
- How do we prove a process existed?

Agency-first governance asks deeper questions:

- Can the person understand what is happening?
- Can the responsible human affect the outcome?
- Can the affected human contest the result?
- Can authority refuse before consequence?
- Can memory preserve what happened and why?
- Can accountability land where control actually existed?

The difference matters. Control-first governance can produce orderly agency loss. Agency-first governance produces responsible order.

## 5. Governance Theater

Governance theater occurs when a system performs the visible signs of governance without preserving the agency conditions that make governance real.

The system may have policies, approvals, dashboards, review committees, consent forms, AI oversight language, audits, compliance reports, or moral vocabulary. But the human role inside the system is hollow. The human cannot understand, refuse, revise, slow, halt, contest, or remember in a way that affects consequence.

Common forms of governance theater include:

- rubber-stamp approval, where humans approve without time, context, or real refusal
- fake consent, where refusal is technically available but practically punished or impossible
- dashboard governance, where visibility replaces authority
- policy theater, where written values do not match operational reality
- audit theater, where records exist but cannot change behavior
- human-in-the-loop theater, where humans are present but too late to affect the outcome
- accountability theater, where responsibility is assigned to people who lacked control

Governance theater is dangerous because it gives systems moral cover. It allows the institution to claim that humans were involved while the real agency path was already captured, bypassed, or predetermined.

## 6. The Agency Conditions of Real Governance

HAPI identifies several agency conditions that must be preserved for governance to be real.

The first is clarity. Humans must be able to understand what is happening and what is being proposed.

The second is authority. The person asked to decide must have rightful power over the decision.

The third is refusal. Approval is not meaningful unless the person can say no before consequence.

The fourth is revision. Human participation is stronger when people can modify the proposal rather than merely accept or reject it.

The fifth is memory. The system must preserve what happened, who decided, what was known, what was refused, what changed, and what consequence followed.

The sixth is accountability. Responsibility must land where authority and control actually existed.

The seventh is contestability. Affected humans must have a way to challenge, appeal, correct, or revisit decisions that affect them.

The eighth is capacity. Humans must have enough time, energy, support, knowledge, and emotional stability to participate meaningfully.

When these conditions are preserved, governance becomes real. When they are absent, governance becomes symbolic.

## Agency Conditions and Governance Functions

Agency condition	Governance function	Failure when absent
Clarity	Makes the decision understandable	Humans approve what they cannot understand
Authority	Connects decision rights to responsibility	People are accountable without control
Refusal	Protects the power to stop consequence	Approval becomes symbolic
Revision	Allows correction before commitment	Participation becomes yes or no theater
Memory	Preserves evidence and continuity	Accountability becomes unstable
Contestability	Allows challenge and repair	Affected people are trapped by outcomes
Capacity	Makes participation practically possible	Humans are overloaded into compliance

## 7. The Governance From Agency Model

HAPI proposes a simple model:

Agency preservation -> agency restoration -> emergent governance -> aligned infrastructure

Agency preservation prevents humans from being reduced to rubber stamps. Agency restoration rebuilds participation where it has already been lost. Governance then emerges as the durable structure that protects the restored agency. Infrastructure operationalizes that governance in tools, workflows, receipts, gates, audits, training, and organizational practice.

This order matters.

If an organization begins with governance infrastructure before agency preservation, it may build a more sophisticated control system without solving the human problem underneath it. It may add approvals without refusal, audits without memory, policies without lived practice, and AI oversight without human authority.

If the organization begins with agency preservation, the governance layer becomes more truthful. It is no longer a symbolic wrapper. It becomes the system by which meaningful human participation is protected across time.

## 8. Governance and Authority

Authority is one of the central conditions of governance. A system cannot be meaningfully governed if no one knows who has the right to decide, who may refuse, who may revise, who owns risk, and who is accountable after consequence.

Many systems fail because authority is either hidden, fragmented, overstated, or displaced.

Hidden authority means decisions are shaped by people, incentives, models, or workflows that are not visible to the person supposedly responsible.

Fragmented authority means no one can tell who owns the final decision or risk.

Overstated authority means a person appears to have power but cannot actually affect the outcome.

Displaced authority means the person held accountable is not the person who controlled the path of action.

Real governance reconnects authority to accountability. The person or body with authority must be visible. Their scope must be clear. Their decision must be recorded. Their refusal must be protected. Their accountability must match their control.

## 9. Governance and Refusal

Refusal is the boundary that separates real governance from procedural theater.

A human who cannot refuse is not governing. A reviewer who cannot stop the action is not governing. A patient who cannot decline without coercion is not meaningfully consenting. A worker who cannot raise safety concerns without retaliation is not participating in safety governance. A religious member who cannot question without social destruction is not exercising conscience under protected agency. An AI oversight approver who cannot halt an agent before execution is not exercising meaningful oversight.

Refusal is not rebellion. Refusal is the protective edge of legitimate participation.

In HAPI, every governance structure should be tested by asking:

- Who can say no?
- When can they say no?
- What happens when they say no?
- Is refusal recorded?
- Is refusal respected?
- Can refusal redirect the system before consequence?

If refusal is absent, the governance system may still control people, but it does not preserve agency.

## 10. Governance and Memory

Memory is required for accountability. A system that cannot remember cannot govern itself honestly.

Memory is more than data storage. Governance memory preserves the human meaning of a decision. It records what was proposed, who had authority, what context was available, what risks were known, what objections were raised, what was revised, what was approved, what was refused, what was executed, and what happened afterward.

Without memory, a system cannot learn. It repeats errors, hides drift, loses accountability, and forces people to relitigate reality from scratch. Without memory, the powerful can deny what happened and the vulnerable cannot prove what happened.

Receipts matter because they preserve agency after consequence. They make the action path reviewable. They allow the system to learn from patterns instead of relying on memory fragments, institutional narratives, or after-the-fact rationalization.

A governance system without memory becomes unstable. A governance system with memory but no access becomes surveillance. A true governance memory must be usable for accountability, correction, contestability, and restoration.

## 11. Governance and Accountability

Accountability is often misunderstood as blame. HAPI treats accountability as the ability to connect consequence back to rightful authority, decision, evidence, refusal, revision, and responsibility.

Accountability becomes unjust when responsibility is separated from control. This happens when workers are blamed for systems they cannot change, patients are blamed for decisions they could not understand, citizens are blamed for processes they cannot contest, users are blamed for automated outcomes they cannot see, or employees are held accountable for AI actions they could not supervise.

Real accountability asks:

- Who controlled the action path?
- Who had authority to approve or refuse?
- What evidence was available?
- What constraints were ignored?
- What did the system know?
- What did the human know?
- Where did the decision bind?
- Where did agency fail?

Accountability should restore truth. It should not be used to relocate institutional failure onto the nearest human body.

## 12. Governance and Values

Company values often fail because they remain abstract. They appear on websites, posters, handbooks, onboarding decks, and compliance documents, but they do not bind the action path.

HAPI treats values as governance inputs only when they shape agency-preserving behavior.

A company that values safety must preserve the agency to stop unsafe action. A company that values integrity must preserve the agency to tell the truth without retaliation. A company that values respect must preserve contestability and dignity. A company that values accountability must

connect responsibility to authority. A company that values innovation must preserve revision and learning without punishing honest failure.

Values become real when they are translated into agency conditions.

The question is not only: what values does the company claim?

The question is:

What forms of human agency must be preserved for those values to become operationally true?

This is how HAPI connects agency preservation to company values. Governance emerges from restored agency, then naturally aligns to the values the organization actually wants to embody.

### 13. Governance in Agentic AI

Agentic AI makes the governance problem urgent because action can outrun human judgment. AI systems can propose actions, call tools, update records, send messages, trigger workflows, modify systems, and influence operational reality. This turns governance from a policy problem into an action-path problem.

HAPI treats AI agents as delegated operational agency. They are not moral persons, but they are not passive tools either. They act under borrowed human or organizational authority. Therefore, their actions must remain bound to human agency.

The technical stack developed under HAPI exists for this reason:

PGDL challenges the proposal before it becomes actionable.

AAG authorizes, revises, escalates, or blocks proposed action.

Runtime Binding keeps execution bound to what was authorized.

Receipts preserve proof and accountability.

Governance Reality Reports evaluate whether oversight was real or theater.

Continuity Findings test whether governance stayed coherent over time.

This stack does not replace human agency. Its purpose is to preserve agency at machine speed.

### 14. Governance in Institutions

The same agency pattern appears outside AI. Institutions often fail when they preserve structure while losing meaningful human participation.

In healthcare, a patient may sign forms while lacking clarity, time, emotional capacity, or real understanding.

In workplaces, employees may attend safety meetings while lacking authority to stop unsafe production pressure.

In schools, students may be managed through behavior systems while losing voice, dignity, and meaningful development.

In religious communities, members may be given belonging only while their conscience remains institutionally compliant.

In government, citizens may formally participate while real decisions are too complex, remote, captured, or procedurally insulated to contest.

HAPI does not treat these as unrelated failures. They are agency failures. The surface systems differ, but the underlying loss is similar: humans remain inside the system while meaningful participation is weakened or removed.

Institutional governance becomes real only when the institution protects the agency of the people affected by its power.

## 15. True Governance Versus False Governance

A useful HAPI distinction is true governance versus false governance.

True governance preserves rightful agency. It clarifies authority, protects refusal, supports revision, preserves memory, connects accountability to control, and restores participation when it breaks.

False governance captures agency. It uses procedures to control humans while claiming to protect them. It preserves approval rituals without refusal. It creates policies that do not bind reality. It moves accountability away from power. It makes people dependent on the system for identity, permission, memory, or legitimacy.

True governance says:

You remain a meaningful participant. This structure protects your ability to act responsibly.

False governance says:

You are included, but the system will decide what your participation means.

This distinction matters because false governance can look more orderly than true governance. It may produce cleaner documents, faster approvals, stronger hierarchy, and more complete dashboards. But order without agency is not real governance. It is controlled participation.

### True Governance Versus False Governance

True governance	False governance
Preserves agency at the decision boundary	Preserves the appearance of participation
Connects authority to accountability	Assigns responsibility without control
Protects refusal before consequence	Treats approval as a ritual
Uses memory for repair and learning	Uses records as institutional cover
Aligns values to lived operation	Keeps values in documents only
Restores capacity over time	Creates dependency on the system

## 16. The HAPI Governance Test

A governance system can be tested through a simple agency-preservation audit.

For any consequential system, ask:

- Can affected humans understand what is happening?
- Can responsible humans affect the outcome before consequence?
- Can the correct authority refuse, revise, or escalate?
- Are approvals meaningful or merely procedural?
- Is responsibility connected to actual control?
- Is memory preserved in a way humans can use?
- Can affected people contest or correct outcomes?
- Do policies match operational reality?
- Does support restore capacity or create dependency?
- Does the system become lighter as human maturity increases, or does it capture agency forever?

These questions reveal whether governance is real or theatrical.

The test is not whether governance artifacts exist. The test is whether human agency remains live at the action boundary and reviewable after consequence.

## 17. From Governance Product to HAPI Practice

If governance is the product of agency preservation, then HAPI can become both a theory and a practice.

As a theory, HAPI explains why systems collapse into theater when agency is removed.

As a diagnostic method, HAPI audits where agency has been lost.

As a restoration model, HAPI rebuilds clarity, authority, refusal, revision, memory, accountability, and capacity.

As infrastructure, HAPI builds tools that preserve agency before, during, and after action.

This creates a different kind of governance business. It does not sell control first. It sells agency restoration, then allows governance to emerge from the restored conditions of meaningful participation.

The practical promise is:

We help organizations find where agency has been lost, rebuild the conditions that restore meaningful human participation, and align the resulting governance layer to both company values and human agency.

## 18. Conclusion

Governance is often treated as a system of rules placed over human behavior. HAPI reframes governance as the product of preserved agency.

A system is not truly governed because it has policies, approvals, dashboards, audits, or compliance language. It is governed when humans retain meaningful participation at the points where consequence can still be changed, and when memory remains after consequence so accountability, contestability, and learning remain possible.

This reframing explains why so many modern systems feel governed but unstable. They preserve the appearance of human participation while stripping away the agency conditions that make participation real.

HAPI's claim is therefore both simple and demanding:

Preserve agency first. Let governance emerge from restored agency. Build infrastructure that keeps that governance coherent over time.

Governance is not the starting point. Governance is the product of agency preservation.

## Appendix A: Key Propositions

6. Governance is the product of agency preservation.
7. Governance becomes theater when human presence remains but meaningful participation is removed.
8. Control is not the same as governance. Real governance requires rightful authority connected to meaningful participation.
9. Approval is not meaningful without refusal before consequence.
10. Accountability is unjust when responsibility is separated from authority and control.
11. Memory is required for accountability, contestability, and learning.
12. Values become operationally real only when they preserve concrete forms of human agency.
13. Agentic AI requires action-path governance because delegated operational agency can outrun human judgment.
14. True governance restores and protects agency. False governance captures agency while performing control.
15. HAPI audits agency loss, restores agency conditions, and builds infrastructure that keeps governance coherent over time.

## Appendix B: Glossary

Term	Definition
<b>Agency</b>	The capacity to understand, choose, refuse, revise, act, remember, contest, and remain accountable.
<b>Agency preservation</b>	The protection of meaningful human participation before, during, and after consequential action.
<b>Agency restoration</b>	The rebuilding of conditions that allow meaningful participation after agency has been weakened or removed.
<b>Governance theater</b>	The appearance of governance without the agency conditions that make governance real.
<b>Refusal</b>	The ability to halt, redirect, challenge, or decline a consequential action before it binds.
<b>Governance memory</b>	Records that preserve proposal, authority, approval, refusal, revision, execution, outcome, and learning.
<b>False governance</b>	A control structure that captures agency while claiming to protect or include humans.
<b>True governance</b>	A structure that preserves rightful agency, authority, memory, and accountability.

## References

Bower, Michael. HAPI Theory of Agency v0.1. Human Agency Preservation Infrastructure foundation manuscript, 2026.

Bower, Michael. Agency Loss: How Systems Preserve Human Presence While Removing Human Participation. HAPI foundation manuscript, 2026.

Bower, Michael. Agency Restoration: Rebuilding Capacity, Authority, Refusal, Memory, and Accountability. HAPI foundation manuscript, 2026.

Bower, Michael. The True Gate: Alignment Theory, Human Agency Infrastructure, and Runtime Governance for Agentic AI. Working thesis draft, 2026.