

RESTORATIVE GOVERNANCE

From Compliance Enforcement to Agency Repair

A HAPI Foundation Thesis Paper

Human Agency Preservation Infrastructure

Author: Michael Bower

Version: v0.1

Date: May 18, 2026

Status: Working research manuscript, not peer reviewed

Core thesis: Governance should not merely punish failure or enforce compliance. True governance restores the conditions that allow people and systems to act coherently again.

Abstract

This paper develops restorative governance as a central HAPI concept. Conventional governance often begins after breakdown: a rule is violated, an incident occurs, a person fails, a system drifts, or an institution faces liability. The response is usually enforcement: investigation, punishment, documentation, escalation, policy expansion, or compliance training. These tools may be necessary, but they are not sufficient. A system can become more controlled while leaving the underlying agency failure unrepaired.

Restorative governance reframes governance as agency repair. Its purpose is not merely to identify fault, assign blame, or harden rules. Its purpose is to restore the conditions under which people and systems can participate coherently: clarity, authority, refusal, revision, memory, accountability, capacity, and trust. In this model, governance becomes real when it repairs the agency conditions that made breakdown possible.

The paper argues that failures in organizations, institutions, religious systems, healthcare systems, and agentic AI workflows are often not only policy failures. They are agency failures. People are present but overloaded, responsible but powerless, informed but unable to refuse, accountable but without control, or governed by systems they cannot understand or contest. Restorative governance provides a framework for diagnosing these failures and rebuilding from the root outward.

Keywords: human agency, governance, restorative governance, agency repair, HAPI, accountability, institutional repair, AI governance, human participation, compliance theater.

Table of Contents

1. Introduction
 2. The Problem With Enforcement-First Governance
 3. Governance as Agency Repair
 4. Agency Failure Before Rule Failure
 5. The Restorative Governance Loop
 6. Accountability Without Destruction
 7. Repairing Authority, Refusal, and Memory
 8. Restorative Governance in Institutions
 9. Restorative Governance in Agentic AI
 10. Restorative Governance and HAPI
 11. Risks and Misuse
 12. Research Agenda
 13. Conclusion
- Appendix A: Core Propositions
- Appendix B: Glossary

1. Introduction

Governance is often treated as a control problem. A system is considered governed if it has policies, approvals, dashboards, disciplinary procedures, audit logs, escalation paths, and compliance evidence. These elements matter, but they can also create a false sense of order. A governed-looking system can still leave people confused, powerless, overloaded, dependent, afraid to refuse, or unable to affect outcomes.

HAPI begins from a different premise: governance is real only when agency is preserved. If the human being remains present but loses meaningful participation, the system has not achieved governance. It has achieved procedural containment.

Restorative governance names the next step in the HAPI foundation. Agency loss identifies the failure. Agency restoration identifies the healing direction. Governance as agency preservation explains the structural relationship. Restorative governance explains how repair is performed after breakdown.

The core claim is simple:

True governance does not end with enforcement. True governance restores the agency conditions that make coherent action possible.

2. The Problem With Enforcement-First Governance

Enforcement-first governance begins with violation. Something goes wrong, and the institution asks: who broke the rule, what policy applies, what penalty is appropriate, what documentation is needed, and how do we prevent liability?

These questions are not useless. Systems do need boundaries. Harm must be addressed. Patterns of negligence, abuse, fraud, or reckless action must be stopped. But enforcement-first governance becomes incomplete when it treats the visible violation as the whole problem.

Many failures are downstream of agency loss. A worker may ignore a rule because the workflow made compliance impossible. A patient may appear noncompliant because the medical system overwhelmed them. A religious member may appear rebellious because questioning was treated as disloyalty. A human reviewer may rubber-stamp an AI action because the system gave them no time, context, or true power to refuse.

In each case, enforcement can punish the person while leaving the agency failure untouched.

Enforcement can create governance theater

A system can become stricter without becoming healthier. It can add more policies, trainings, approvals, and logs while still failing to restore the human capacity to understand, refuse, revise, remember, and remain accountable.

- More rules can increase overload.
- More approvals can increase rubber-stamping.
- More documentation can hide responsibility.
- More surveillance can weaken trust.

- More punishment can reduce honest reporting.
- More compliance language can conceal agency collapse.

Restorative governance does not reject enforcement. It subordinates enforcement to repair. The purpose of consequence is not institutional self-protection alone. The purpose is to restore order, agency, truth, safety, and accountable participation.

3. Governance as Agency Repair

Restorative governance defines governance as the disciplined repair of agency conditions after breakdown. It asks not only what rule was violated, but what agency condition failed.

The primary repair questions are:

- Was the person able to understand what was happening?
- Did the person have legitimate authority over the outcome?
- Could the person refuse before consequence?
- Could the person revise the action before it became binding?
- Was memory preserved after the event?
- Was accountability matched to actual control?
- Did the system create dependence instead of capacity?
- Did the formal policy match operational reality?

If these questions are not answered, governance remains shallow. The system may punish behavior without repairing the conditions that produced it.

Restoration is not permissiveness

Restorative governance is sometimes misunderstood as softness. It is not. Repair does not mean ignoring harm, denying responsibility, or removing consequences. It means that consequences are aimed at restoration instead of mere exclusion, humiliation, liability control, or symbolic discipline.

A restorative system can still block unsafe actions, remove dangerous authority, require restitution, escalate risk, or separate a person from a role. The difference is that the separation is governed by rightful boundaries and repair logic, not by agency capture or institutional self-preservation.

4. Agency Failure Before Rule Failure

Many rule failures begin as agency failures. The person does not only break a rule. The person loses the conditions required to act coherently within the rule.

Common agency failures before visible breakdown include:

- Clarity failure: the person does not understand the situation, stakes, or policy.
- Authority failure: the person is responsible for a decision they cannot actually control.
- Refusal failure: the person cannot safely say no.
- Revision failure: the person cannot change the course before consequence.
- Memory failure: the system does not preserve what happened and why.
- Capacity failure: the person lacks the time, energy, training, support, or emotional regulation required to participate well.

- Accountability failure: responsibility is assigned without matching authority.
- Reality failure: the written process does not match the lived process.

Restorative governance looks upstream. It asks what condition failed before the rule failed.

This matters because a system that only responds to rule failure will keep treating symptoms. A system that repairs agency failure can reduce future harm at the root.

5. The Restorative Governance Loop

HAPI can express restorative governance as a loop. The loop does not end at punishment. It continues until agency conditions have been repaired or the boundary has been lawfully reset.

1. Detect the breakdown. Identify the action, harm, drift, refusal failure, accountability gap, or governance theater pattern.
2. Stabilize the system. Stop ongoing harm, pause risky action, protect affected people, and prevent further consequence.
3. Preserve memory. Record what happened, who had authority, what was known, what was approved, what was refused, and what evidence remains.
4. Diagnose agency loss. Determine which agency conditions failed: clarity, authority, refusal, revision, memory, capacity, accountability, contestability, or trust.
5. Assign rightful accountability. Match responsibility to actual authority and control.
6. Repair conditions. Restore capacity, clarify authority, rebuild refusal paths, update policy, create better memory, and remove dependency capture.
7. Test restoration. Verify that future participants can understand, refuse, revise, act, remember, and remain accountable.
8. Reduce friction where safe. As maturity increases, governance should become lighter, not heavier forever.

This loop can apply to a workplace incident, a clinical care process, a spiritual community, a government agency, or an AI agent workflow.

6. Accountability Without Destruction

Accountability is often confused with punishment. Punishment may be necessary in some cases, but accountability is deeper. Accountability means that reality is named, responsibility is located, harm is addressed, memory is preserved, and repair becomes possible.

Destructive accountability strips agency. It uses shame, exclusion, fear, opacity, or power imbalance to force compliance. Restorative accountability strengthens agency by making truth survivable and repair actionable.

Accountability must match authority

One of the most common agency failures is responsibility without authority. People are blamed for outcomes they could not meaningfully affect. This creates moral confusion and institutional decay.

Restorative governance insists on a strict rule:

No responsibility without authority. No authority without accountability.

When this rule is broken, governance becomes theater. A person can be punished while the real decision path remains untouched. A team can be blamed while leadership escapes scrutiny. A human reviewer can be named accountable while an automated system made refusal impossible.

7. Repairing Authority, Refusal, and Memory

Restorative governance focuses on the agency conditions that make future coherence possible. Three of the most important are authority, refusal, and memory.

Authority repair

Authority repair asks whether people have the legitimate power required to carry the responsibility assigned to them. If they are accountable for an outcome, they must have enough authority to understand, shape, halt, or escalate the action before it becomes binding.

Refusal repair

Refusal repair asks whether people can say no without being destroyed, punished, shamed, or bypassed. A system that cannot tolerate rightful refusal does not preserve agency. It produces compliance behavior.

Memory repair

Memory repair asks whether the system can remember what happened in a way that supports accountability and learning. Without memory, harm becomes rumor, responsibility becomes vague, and repair becomes impossible.

In agentic AI systems, receipts serve this function. In institutions, memory may include records, testimony, process maps, review packets, decision logs, appeals, and lived evidence from affected participants.

8. Restorative Governance in Institutions

Institutional failures often persist because institutions optimize for continuity of the institution rather than restoration of agency. The institution protects itself, expands documentation, and narrows liability while the human participants remain less capable, less trusted, less heard, or less able to refuse.

Restorative governance reverses the direction of repair. It begins with the core agency problem and rebuilds outward.

- In workplaces, it asks whether workers can report unsafe conditions without retaliation.
- In healthcare, it asks whether patients can understand, refuse, choose, and remember care decisions under stress.
- In education, it asks whether students are trained into agency or managed into performance.
- In religious communities, it asks whether discipline restores conscience and relationship with God or captures belonging through fear.
- In government systems, it asks whether citizens can contest, understand, and participate in decisions that affect them.

The institutional question is not merely: do we have a policy? The deeper question is: does the policy restore meaningful participation where agency has been weakened?

9. Restorative Governance in Agentic AI

Agentic AI intensifies the governance problem because artificial systems can act faster than human judgment can participate. A human may technically be in the loop while the practical agency path has already moved beyond them.

Restorative governance for agentic AI requires more than model alignment. It requires action-path repair.

- PGDL repairs proposal quality by challenging distorted interpretation before action.
- AAG repairs authority by requiring that proposed actions be scoped, authorized, reversible, or escalated.
- Runtime Binding repairs execution fidelity by ensuring the system does only what was authorized.
- Receipts repair memory by preserving proof after consequence.
- Governance Reality Reports repair institutional awareness by showing whether governance is real or theater.
- Continuity Findings repair temporal drift by asking whether governance stayed coherent over time.

The goal is not to slow every action forever. The goal is to mature the system so safer delegation becomes possible without surrendering human authority.

A true gate restores agency. A false gate creates dependency. Restorative governance keeps the gate oriented toward maturity rather than permanent control.

10. Restorative Governance and HAPI

Restorative governance gives HAPI its repair method. HAPI does not merely audit where agency has been lost. It helps rebuild from the core problem outward.

The HAPI sequence becomes:

9. Identify agency loss.
10. Distinguish real agency from agency theater.
11. Detect agency capture and dependency capture.
12. Restore clarity, authority, refusal, memory, capacity, and accountability.
13. Let governance emerge from the restored conditions of meaningful human participation.
14. Build infrastructure that preserves agency before, during, and after consequential action.

This makes governance the product of agency preservation rather than the starting point. HAPI does not impose control first and hope agency survives. HAPI restores agency first and allows governance to become real through that restoration.

11. Risks and Misuse

Restorative governance can itself be misused. Any powerful framework can become theater if it is turned into language without repair.

- Restoration theater: the institution talks about repair while preserving the same power structure.
- Forgiveness coercion: harmed people are pressured to reconcile before truth and safety are restored.
- Soft accountability: repair language is used to avoid rightful consequence.
- Over-process: the repair process becomes so complex that people lose agency again.
- Gatekeeper capture: the restorative authority becomes the new false gate.
- Metrics capture: agency repair is reduced to scores that can be gamed.

To prevent misuse, restorative governance must remain anchored to agency outcomes. The question is not whether the process sounded restorative. The question is whether it restored meaningful participation, authority, refusal, memory, accountability, and capacity.

12. Research Agenda

The next phase of restorative governance is empirical and implementation-driven. HAPI should test whether agency repair improves outcomes across institutions and agentic systems.

- Can HAPI audits reliably identify agency failure before institutional breakdown?
- Do refusal paths reduce harm and increase trust?
- Do receipts and memory systems improve accountability without becoming surveillance?
- Can governance friction decrease as maturity increases?
- Can institutions distinguish restoration from compliance theater?
- Can agentic AI systems improve proposal quality after repeated PGDL/AAG feedback without silently drifting?
- Can healthcare, workplace, religious, and AI systems use the same agency-preservation dimensions without forcing false equivalence?

The strongest evidence will come from real deployments: before-and-after audits, governance reality reports, participant interviews, reduction in rubber-stamping, better refusal paths, clearer authority maps, and improved memory of consequential decisions.

13. Conclusion

Restorative governance is the HAPI answer to what happens after agency loss is discovered. Enforcement can stop harm, but enforcement alone cannot heal the system. A governed system must repair the agency conditions that make coherent action possible.

The paper has argued that many institutional and technological failures are not only rule failures. They are failures of clarity, authority, refusal, revision, memory, capacity, accountability, and contestability. If those conditions remain broken, governance becomes theater no matter how many policies exist.

Restorative governance does not remove boundaries. It makes boundaries serve repair. It does not erase accountability. It makes accountability truthful, proportional, and connected to real authority. It does not weaken institutions. It rebuilds institutions from the human agency layer outward.

For HAPI, this is the practical bridge between theory and intervention:

Audit agency loss. Repair agency conditions. Let governance emerge from restored participation.

Appendix A: Core Propositions

15. P1. Governance is not complete when rules are enforced; governance becomes real when agency conditions are repaired.
16. P2. Many rule failures are downstream of agency failures.
17. P3. Accountability must match actual authority and control.
18. P4. Refusal, revision, and memory are necessary conditions for meaningful participation.
19. P5. Restoration is not permissiveness; it is repair ordered toward truth, safety, and coherent action.
20. P6. A system can become stricter while becoming less agency-preserving.
21. P7. Restorative governance must avoid becoming restoration theater.
22. P8. In agentic AI, restorative governance requires proposal scrutiny, action authorization, execution binding, receipts, and continuity review.
23. P9. The purpose of governance is not control for its own sake. The purpose of governance is preserved and restored agency.
24. P10. HAPI rebuilds systems from the core agency-loss problem outward.

Appendix B: Glossary

Restorative governance: A governance model that repairs the agency conditions that made breakdown possible.

Agency repair: The restoration of clarity, authority, refusal, revision, memory, capacity, and accountability.

Enforcement-first governance: A governance pattern that begins and often ends with rule violation, documentation, punishment, or liability control.

Restoration theater: The appearance of repair without actual restoration of human agency.

Refusal path: A protected route by which a person can halt, contest, or redirect a consequential action before commitment.

Memory repair: The creation or restoration of trustworthy records that preserve what happened, why it happened, who had authority, and what evidence remains.

Authority repair: The process of aligning responsibility with actual power to understand, shape, halt, or escalate action.

Agency condition: A necessary support for meaningful participation, such as clarity, authority, capacity, refusal, memory, or accountability.