

# THE HAPI CONSTITUTION

Principles for Preserving Human Agency in High-Pressure Systems

*A Human Agency Preservation Infrastructure Thesis Paper*

---

Author: Michael Bower

Project: Human Agency Preservation Infrastructure (HAPI)

Version: v0.1 Working Thesis Draft

Status: Conceptual foundation paper, not peer reviewed

Date: May 2026

*Core thesis: A system preserves human agency only when it protects understanding, authority, refusal, revision, memory, accountability, and contestability before power becomes consequence.*

---

## Abstract

This paper proposes the HAPI Constitution, a set of first principles for preserving human agency in systems that operate under pressure, scale, automation, institutional authority, or moral consequence. Human Agency Preservation Infrastructure begins from the premise that governance is not real merely because policies, dashboards, procedures, approvals, or leaders are present. Governance becomes real when humans retain the meaningful ability to understand, refuse, revise, approve, remember, contest, and remain accountable. The HAPI Constitution translates this premise into practical principles that can guide institutional design, AI governance, religious and moral authority, workplace systems, healthcare processes, and public infrastructure.

The paper argues that agency loss often occurs through respectable structures: compliance programs, automated workflows, spiritual administration, safety processes, performance metrics, consent forms, escalation paths, and review rituals. These mechanisms can protect people, but they can also become agency theater when they preserve the appearance of participation while removing real power. The HAPI Constitution is therefore framed as a boundary document. It defines what a system must not do if it claims to preserve human agency, and what it must protect if governance is to remain legitimate.

Keywords: human agency, HAPI, agency preservation, governance, agency theater, refusal, accountability, authority, contestability, restoration, AI governance

## Table of Contents

1. Introduction
2. Why HAPI Needs a Constitution
3. Core Definitions
4. The Constitutional Invariant
5. The Ten Agency-Preservation Principles
6. Anti-Principles: What HAPI Rejects
7. The Agency Rights of Participants
8. Governance Duties of Systems
9. Institutional Application
10. Agentic AI Application
11. Religious and Moral Authority Application
12. Healthcare and Crisis Application
13. Evaluation Questions
14. Failure Modes
15. Conclusion

## Appendix A: Principle Checklist

## 1. Introduction

HAPI, Human Agency Preservation Infrastructure, requires more than a theory of agency. It also requires a constitutional layer: a set of stable commitments that define what must be preserved whenever systems gain power over human participation. Without such commitments, agency preservation can become a slogan, a brand, or a compliance label rather than a binding discipline.

A constitution does not solve every case in advance. It establishes the boundary conditions under which cases must be judged. It says what kind of system HAPI is trying to build, what kind of system HAPI refuses to become, and what must remain protected even when efficiency, authority, automation, safety, urgency, or institutional convenience creates pressure to bypass human participation.

*The purpose of the HAPI Constitution is to prevent agency preservation from becoming another form of agency capture.*

This is especially important because the language of agency can be easily absorbed by systems that do not actually preserve it. A company can say employees are empowered while giving them no authority to refuse. A church can say conscience matters while making belonging dependent on institutional approval. A hospital can ask for informed consent while overwhelming the patient with fear and complexity. An AI platform can keep a human in the loop while allowing the system to move too quickly for meaningful intervention.

The HAPI Constitution therefore begins from a simple distinction: a system may include humans without preserving human agency. Presence is not participation. Approval is not authority. Consent is not comprehension. Oversight is not meaningful if it cannot still change the outcome.

### Constitutional thesis

A system preserves agency only when it protects the human ability to understand, refuse, revise, approve, remember, contest, and remain accountable at the point where consequence can still be changed.

## 2. Why HAPI Needs a Constitution

HAPI aims to audit systems and institutions, identify where agency has been weakened or captured, and help rebuild from the core agency-loss problem outward. That mission requires more than diagnostic categories. It requires commitments that constrain HAPI itself.

Any agency-preservation project faces a paradox. To restore agency, it must often introduce structures: audits, gates, reports, authority maps, receipts, review processes, and governance tools. But those same structures can become the thing they were meant to prevent. A gate can become a false gate. A support system can create dependency. A governance tool can become surveillance. A score can become compliance theater. A safety process can become a machine for avoiding responsibility.

The HAPI Constitution exists to keep restoration from becoming capture. It is not merely a public statement of values. It is a design constraint for every future HAPI audit, product, institution, tool, and governance model.

- It defines the agency conditions HAPI must protect.
- It clarifies the difference between support and capture.
- It prevents governance from becoming control for its own sake.
- It keeps human participation from becoming ornamental.
- It gives auditors and builders a repeatable standard for evaluating systems.

### 3. Core Definitions

Term	Definition
Agency	Meaningful human participation under conditions of capacity, authority, clarity, refusal, memory, and accountability.
Agency preservation	The protection of the conditions that allow humans to understand, refuse, revise, approve, act, remember, contest, and remain accountable.
Agency restoration	The process of rebuilding agency after overload, coercion, dependency, automation bypass, false authority, or institutional fragmentation has weakened participation.
Agency capture	A system pattern where support, authority, belonging, tools, or governance begin replacing the agency they claim to protect.
Agency theater	The appearance of human participation without the conditions needed for real participation.
True gate	A boundary that governs rightful passage while preserving agency, authority, accountability, and restoration.
False gate	A boundary that controls access, identity, belonging, action, or legitimacy without rightful authority or restorative purpose.

### 4. The Constitutional Invariant

The central invariant of HAPI is this:

*Preserve meaningful human agency before, during, and after consequential system action.*

This invariant applies across domains. In AI systems, it means delegated agency must not outrun human judgment. In institutions, it means procedures must not remove meaningful participation. In religion, it means human authority must not replace conscience or direct accountability before God. In healthcare, it means the patient must not become a passive object inside a complex system. In work, it means employees must not be held responsible for outcomes they cannot influence.

The invariant has three temporal requirements:

1. Before consequence: the person must have clarity, context, authority, and the ability to refuse or revise.
2. During consequence: the system must remain bound to the scope, purpose, and authority that made the action legitimate.
3. After consequence: the system must preserve memory, proof, contestability, and accountability.

## 5. The Ten Agency-Preservation Principles

### 1. No responsibility without authority

A person should not be held accountable for outcomes they could not meaningfully understand, direct, refuse, or change.

### 2. No approval without refusal

Approval is meaningful only if the person can say no without being punished, erased, humiliated, or bypassed.

### 3. No consent without comprehension

Consent is not real when the person lacks the context, time, language, or capacity required to understand what is being accepted.

### 4. No automation without memory

Systems that act on behalf of humans must preserve what happened, why it happened, who authorized it, and what evidence supported it.

### 5. No governance without contestability

A governed person must have some path to question, appeal, correct, or contest decisions that affect them.

### 6. No support that creates dependency

Support should restore capacity and participation, not make the person less able to act without the system.

### 7. No human-in-the-loop theater

A human presence is not oversight unless the human can still change the outcome before consequence.

### 8. No hidden transfer of authority

Systems must not quietly transfer decision power from humans to institutions, metrics, automation, or committees while preserving the appearance of human control.

## 9. No moral outsourcing

Humans and organizations cannot delegate away responsibility for consequential action. They can delegate tasks, but they remain accountable for the delegation.

## 10. No restoration without agency

A process that punishes, excludes, or controls without restoring clarity, responsibility, and participation is not restorative governance.

## 6. Anti-Principles: What HAPI Rejects

A constitution is defined not only by what it affirms, but by what it refuses. HAPI rejects patterns that simulate agency while removing its load-bearing conditions.

- Ornamental participation: placing humans in a workflow after the meaningful decision has already been made.
- Compliance substitution: treating rule-following as proof that agency has been preserved.
- Coercive belonging: making social, spiritual, professional, or institutional belonging dependent on surrendering conscience or refusal.
- Metric capture: replacing human judgment with scores that cannot be contested or contextualized.
- Authority laundering: using committees, tools, policies, or agents to hide who actually made a decision.
- Restoration theater: claiming to restore a person while keeping them dependent, humiliated, voiceless, or unable to contest their status.
- Gatekeeper dependency: requiring people to pass through a human or institutional gate that does not exist to restore their agency.

## 7. The Agency Rights of Participants

If a person is affected by a high-pressure system, HAPI treats that person as a participant whose agency must be preserved, not as a passive object to be processed. The following are not legal rights in this draft. They are agency rights: conditions that must be considered if a system claims to preserve human agency.

Agency Right	Meaning
Right to clarity	The person should be able to understand what is happening and why.
Right to refusal	The person should have a meaningful path to say no before consequence.
Right to revision	The person should be able to correct, redirect, or narrow a proposed action.
Right to memory	The person should be able to see what happened, who acted, and why.
Right to contestability	The person should have a path to challenge errors, unfairness, overreach, or misclassification.
Right to proportionality	The system response should not destroy agency beyond what the situation requires.
Right to restoration	Corrective processes should aim to restore agency, not merely enforce exclusion or control.

## 8. Governance Duties of Systems

For every agency right, there is a corresponding system duty. A system that claims to preserve agency must be designed to protect participation under pressure, not only under ideal conditions.

4. The duty to make authority visible.
5. The duty to slow consequential action when human judgment must participate.
6. The duty to preserve proof after action.
7. The duty to distinguish real approval from pressured approval.
8. The duty to expose policy-reality splits.
9. The duty to prevent support from becoming dependency.
10. The duty to make correction restorative rather than merely punitive.

These duties apply whether the system is technical, institutional, religious, medical, educational, or organizational. The form changes, but the agency pattern remains stable.

## 9. Institutional Application

Institutions often erode agency unintentionally. They introduce policies to reduce chaos, metrics to increase accountability, procedures to ensure consistency, and approval chains to prevent harm. These structures can be necessary. But they become agency-eroding when they separate responsibility from authority, turn people into procedural participants, or make contestability impossible.

A HAPI-aligned institution does not ask only whether rules were followed. It asks whether people retained meaningful participation in the decisions that affected them. Did employees understand the policy? Could they refuse unsafe work? Could patients contest misclassification? Could students challenge opaque decisions? Could members question authority without losing belonging?

### Institutional test

An institution preserves agency when its procedures protect the person's ability to understand, refuse, revise, remember, and contest, especially when the system is under pressure.

## 10. Agentic AI Application

Agentic AI intensifies the need for a constitutional layer because agents can convert intention into action faster than humans can participate. A chatbot produces text. An agent can call tools, send messages, update records, deploy code, move money, publish content, or alter operational reality. This means agency can be delegated into executable systems.

The HAPI Constitution does not reject delegation. It rejects ungoverned delegation. Humans may use agents to extend capacity, but the system must preserve authority, refusal, memory, and accountability. The agent should amplify human agency, not replace it.

This maps to a runtime governance stack: PGDL scrutinizes proposed meaning before action, AAG authorizes or blocks action, Runtime Binding constrains execution to the authorized permit, Receipts preserve memory, and Governance Reality Reports evaluate whether the system remained real rather than theatrical.

*AI should amplify human agency without outrunning human judgment.*

## 11. Religious and Moral Authority Application

Religious systems can preserve agency by forming conscience, restoring responsibility, protecting community, and pointing people toward truth. They can also capture agency when human authority becomes the gate through which conscience, belonging, repentance, family access, or spiritual legitimacy must pass.

The HAPI Constitution does not deny that communities need boundaries, discipline, doctrine, or accountability. It asks whether those structures restore agency or capture it. A true gate shepherds a person toward truth and restoration. A false gate stands where only God, conscience, or truth belongs.

This distinction matters because spiritual authority can become uniquely destabilizing when it controls not only behavior, but identity, belonging, family, and perceived access to God. A HAPI-aligned spiritual community must preserve conscience, honest questioning, repentance, restoration, and proportional correction.

## 12. Healthcare and Crisis Application

Crisis weakens agency. Illness, fear, fatigue, pain, grief, and information overload can make it difficult for people to understand options, ask questions, refuse interventions, or remember what was explained. Healthcare systems therefore face a special agency duty: they must not treat formal consent as proof of meaningful participation.

Agency-preserving care would support comprehension, decision memory, family participation, second opinions, contestability, and patient dignity. It would not leave the patient responsible for choices made under confusion, panic, or institutional speed. It would treat agency as part of healing, not as a bureaucratic checkbox.

## 13. Evaluation Questions

The HAPI Constitution can be translated into audit questions. These questions are designed to expose the difference between real agency and agency theater.

- Can the affected person understand what is happening?
- Can the person refuse before consequence?
- Can the person revise or narrow the action?
- Can the person contest the decision afterward?
- Is the approving human still able to change the outcome?
- Is authority visible and proportional?
- Are receipts or records preserved?

- Does the support system restore capacity or create dependency?
- Does the system punish refusal, questioning, or conscience?
- Does accountability match actual control?

## 14. Failure Modes

Failure Mode	Constitutional Violation
Rubber-stamp approval	Violates no approval without refusal.
Authority without control	Violates no responsibility without authority.
Opaque consent	Violates no consent without comprehension.
Hidden automation	Violates no hidden transfer of authority.
Unchallengeable scores	Violates no governance without contestability.
Dependency capture	Violates no support that creates dependency.
False gate authority	Violates no restoration without agency.
Receipt gaps	Violates no automation without memory.

## 15. Conclusion

The HAPI Constitution is a boundary document for agency preservation. It defines the commitments that must remain stable as HAPI moves from theory into audits, governance models, infrastructure, software, institutional consulting, healthcare applications, religious analysis, and agentic AI systems.

The central claim is simple: governance is not real unless agency is preserved. A person who cannot understand, refuse, revise, remember, contest, or affect the outcome is not meaningfully participating. A system that keeps such a person present while removing those capacities is not preserving agency. It is performing agency.

HAPI therefore begins with restoration, not control. It audits where agency has been lost, rebuilds the conditions of meaningful participation, and allows governance to emerge from that restored agency. The Constitution exists to keep that mission honest.

*Agency preservation is the mission. Governance is the product. Infrastructure is the implementation.*

## Appendix A: Principle Checklist

A system should be considered constitutionally weaker when it violates any of the following checks:

- The person is accountable but lacks authority.
- The person approves but cannot refuse.
- The person consents but does not understand.
- The system acts but leaves no memory.
- The person is governed but cannot contest.
- The system supports but creates dependency.
- The process includes a human but the human cannot alter the outcome.
- Authority moves to automation, metrics, committees, or institutions without being named.
- The organization delegates action but denies responsibility.
- Correction controls the person without restoring agency.