The Law of Declared Authority:

Necessary Closure for Legitimacy in Decision-Permitting Systems

Stacy Gildenston — Pyrate Ruby Passell

March 2026 · v2.1 · Governing Constraint

Abstract

This note formalises the Law of Declared Authority as the necessary closure condition of the Three Primitives framework. It establishes that legitimacy in a decision-permitting system is not an emergent property of computation, but a structural requirement arising from explicit human declaration. The law is derived directly from the E–H–S Nucleus Invariant #1 and associated proofs. A system that permits action without a declaration of purpose, authority, and constraints necessarily enters a state of Ghost Authority and is therefore structurally invalid. This document is a statement record. The foundational proofs on which it depends are archived separately and cited in the references.

1 Formal Statement

Let S be a decision-permitting system in which a permission function g determines whether the output of a decision function f may be executed.

Definition 1 (Declared Authority). A declaration of authority is a tuple δ = (p, a, c), where:

- p Purpose: the bounded scope within which the action is authorised.
- a Authority: the named human source who holds and declares the mandate.
- c Constraints: the operational limits that bound the permitted action.

δ = 0 denotes the absence of any declared purpose, authority, or constraints at the moment permission is evaluated.

Law 1 (Declared Authority). In any decision-permitting system:

$$g = 1 \Rightarrow \exists\, \delta = (p, a, c) \text{ declared by a legitimate human source} \quad (1)$$

$$\delta = 0 \Rightarrow g = 0 \quad (2)$$

Any system state in which g = 1 and δ = 0 constitutes a Ghost Authority state and is structurally invalid.

2 Derivation

This law introduces no new axioms. It is a direct structural consequence of two previously established results within the Three Primitives framework.

2.1 From the E–H–S Nucleus Invariant #1

The Canonical Logic Sequence establishes that for any escalated event x where E(x) = 1, the set of permitted actions A(x) has cardinality greater than 1. The mapping from escalation to action selection is therefore non-unique: no computable function f can determine a uniquely legitimate action from system state alone.

$$E(x) = 1 \Rightarrow |A(x)| > 1 \Rightarrow f \text{ is underdetermined} \quad (3)$$

Because f is underdetermined, a selection among elements of A(x) requires an additional parameter that is not derivable from system state. That parameter is δ. Without δ, any selection made is arbitrary with respect to legitimacy.

2.2 From the No Closed Mathematical Governance Lemma

The AI Cannot Govern AI paper proves via Rice's theorem that no total computable function $h : S \times A \to \{0, 1\}$ can autonomously resolve authorisation without an exogenous declaration. Authorisation is a non-trivial semantic property of system behaviour under context $c \in C$, and contexts are exogenous to system state S.

Proof. Assume δ were derivable from internal system state s alone. Then the permission function g would collapse into the decision function f, eliminating the structural distinction between authorisation and execution. The governance gap Δ(s) would become empty for all s, contradicting the No Closed Mathematical Governance Lemma. Therefore δ cannot be derived from s. It must be declared exogenously by a human source prior to g = 1.

2.3 Implication

The Law of Declared Authority converts human oversight from an ethical preference into a structural requirement. Any system that permits action without explicitly naming its source of authority at the moment of permission is mathematically incomplete and cannot be considered legitimate. This is not a policy claim. It is a consequence of the underdetermination of f and the non-computability of authorisation.

3 Ghost Authority

Definition 2 (Ghost Authority). A system is in Ghost Authority state if and only if g = 1 and δ = 0. In this state, action is executed under apparent permission without a legitimate human authority declaration.

Ghost Authority is an architectural failure, not a moral one. It is the structural void that results when the permission gate fires without a declared authority source. Lemma C — Ghost Authority Lemma establishes that treating a protocol as the source of a decision, rather than as an input to a human decision, necessarily produces this state.

Ghost Authority is identifiable by the Ownership Test: at the point where consequences attach, if no human explicitly declares "I decide this," the decision lacks legitimate authority.

4 Scope

This law applies to any system — computational, institutional, or hybrid — in which a permission gate g determines whether a decision function f is allowed to execute, and where the output of f produces a binding effect on one or more human subjects.

The law does not supersede applicable law or regulation. Satisfaction of the δ closure condition is

necessary but not sufficient for total legitimacy. The Twisted Pair Legitimacy Theorem establishes the additional requirement of ADCI value-layer closure.

5 References at 3primitives.io

[1] Passell, P. R., & Gildenston, S. (2025). Three Primitives — Canonical Logic Sequence (Clean-room v1.2).

[2] Gildenston, S., & Passell, P. R. (2026). AI Cannot Govern AI: A Formal Proof of Structural Openness in Intelligent Systems.

[3] Gildenston, S., & Passell, P. R. (2026). AGI as a Decision-Complete System: A Mathematical Definition.

[4] Gildenston, S., & Passell, P. R. (2026). The Adjacency Lemma: The Interface–Authority Boundary.

[5] Gildenston, S., & Passell, P. R. (2026). Lemma C — Ghost Authority Lemma.

[6] Gildenston, S., & Passell, P. R. (2026). The Twisted Pair Legitimacy Theorem.