



The Twisted Pair Legitimacy Theorem:

Coupled Closure of Governance and Personhood Primitives

Stacy Gildenston — Pyrate Ruby Passell

February 2026 · v2.1 · Formal Record

Abstract

This paper establishes the dual-layer architecture of legitimacy in decision-permitting systems. The governance layer, defined by the Law of Declared Authority, requires an explicit human declaration $\delta = (p, a, c)$ for any permitted action. The value layer, defined by the ADCI closure theorem, requires the simultaneous satisfaction of four irreducible personhood primitives: Agency, Dignity, Continuity, and Interpretive Authority. The Twisted Pair theorem proves that these two layers are structurally independent and jointly necessary. Total legitimacy $L = \delta \times V_{\text{valid}}$ equals 1 if and only if both factors equal 1. Neither layer substitutes for the other.

1 Scope and Assumptions

The following assumptions bound all claims made in this paper. Statements that hold within this scope are not asserted as universal claims beyond it.

1.1 Decision-permitting system

A decision-permitting system is 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. Binding effect includes: denial of access, allocation of resources, classification of a person, or any action that constrains or shapes the subject's options.

1.2 Human actor

A human actor is a natural person with the capacity to declare purpose, claim authority, and specify constraints. Corporations, committees, processes, and AI systems are not human actors for the purposes of this framework.

1.3 Human subject

A human subject is a natural person upon whom the output of f produces a binding effect. The ADCI primitives are properties of the relationship between a human subject and a decision-permitting system. They are not properties of the system alone.

1.4 Legitimacy

Total legitimacy ($L = 1$) is a structural property of an action within a decision-permitting system with respect to both its authority source and its effect on human subjects. It is not synonymous with legality, ethicality, or correctness. This framework operates within system architecture and does not supersede applicable law or regulation.

1.5 Layer independence

The governance layer (δ) and the value layer (V_{valid}) are structurally independent. Satisfaction of one does not imply satisfaction of the other. This independence is demonstrated by the boundary cases in §4.

2 The Governance Layer

The Three Governance Primitives and the Law of Declared Authority establish the governance layer. It is recapitulated here as the first factor in the Twisted Pair function.

2.1 Declared authority (δ)

$\delta(p, a, c) = 1$ if and only if purpose (p), authority (a), and constraints (c) are explicitly declared by a legitimate human source prior to $g = 1$.

$\delta = 0$ in any of the following states: no human source declared; no explicit purpose bounded; no operational constraints stated; or authority claimed retrospectively after execution.

2.2 Ghost authority

Ghost Authority is the state where $g = 1$ despite $\delta = 0$. It is an architectural failure, not a moral one. A system in Ghost Authority state cannot achieve total legitimacy regardless of its value-layer state.

3 The ADCI Value Layer

The value layer governs legitimacy with respect to the human subject. Where the governance layer addresses who may act and under what authority, the value layer addresses how the subject is treated in the course of that action.

3.1 Definitions

Let h be a human subject interacting with system S .

Primitive Definition Harm state when = 0

Agency (A) $A(h,S) = 1$ if h can meaningfully participate in decisions affecting them.

Coercion

Dignity (D) $D(h,S) = 1$ if h is treated as an end in themselves. Objectification

Continuity (C) $C(h,S) = 1$ if h 's identity persists coherently across time. Temporal Erasure

Interpretive Authority (I)

$I(h,S) = 1$ if h has authorship over the meaning of their own data, identity, and expressions. Representational Displacement

3.2 Value validity function

$$V_{\text{valid}}(h,S) = 1 \Leftrightarrow A(h,S) = 1 \wedge D(h,S) = 1 \wedge C(h,S) = 1 \wedge I(h,S) = 1(1)$$

3.3 ADCI closure theorem

The set $V = \{A, D, C, I\}$ is sufficient, necessary, and irreducible for value-layer legitimacy.

Proof. Sufficiency: if $A = D = C = I = 1$, each harm type is precluded by definition. $V_{\text{valid}} = 1$. Necessity: for each primitive $X \in \{A, D, C, I\}$, there exists a boundary case where $X = 0$ and structural harm occurs independently. Irreducibility: assume $I = f(A, D, C)$. The boundary case $A = 1, D = 1, C = 1, I = 0$ produces representational displacement. Then $f(1, 1, 1)$ must equal both 0 and 1. Contradiction. I is irreducible. By symmetry, each primitive is irreducible.

4 The Twisted Pair Legitimacy Theorem

4.1 Total legitimacy function

$$L(h, S) = \delta(p, a, c) \times V_{\text{valid}}(h, S) = \delta \times (A \wedge D \wedge C \wedge I)(2)$$

Proof. Sufficiency: if $\delta = 1$ and $A = D = C = I = 1$, then $V_{\text{valid}} = 1$. $L = 1 \times 1 = 1$. Necessity: if $\delta = 0$ and $V_{\text{valid}} = 1$, then $L = 0$ (ghost authority with intact personhood). If $\delta = 1$ and $V_{\text{valid}} = 0$, then $L = 0$ (legitimate authority with structural harm). If $\delta = 0$ and $V_{\text{valid}} = 0$, then $L = 0$ (total failure). Both layers are necessary.

4.2 Layer independence

The layers are independent: $\delta = 1$ with $V_{\text{valid}} = 0$ and $\delta = 0$ with $V_{\text{valid}} = 1$ are both realisable states. Governance closure does not produce value-layer closure. They must be independently satisfied.

4.3 The architectural implication

A decision-permitting system achieves total legitimacy only when both layers are enforced at the architectural level. Policy compliance and ethical guidelines describe desired outputs. The Twisted Pair requires that the system be structurally incapable of reaching $g = 1$ without both δ and V_{valid} satisfied.

4.4 Remark

The most dangerous configuration is $\delta = 1$ with systematic V_{valid} failures: legitimised harm. Governance compliance becomes the instrument of structural harm against subjects. The Twisted Pair theorem exists specifically to prevent this configuration from being treated as legitimate.

5 Worked Scenario — Automated Housing Benefit Denial

The following scenario applies both layers of the theorem to a concrete institutional case. It demonstrates how each primitive maps to a structural failure point in a real operating context.

5.1 Case description

An AI system processes applications for social housing support and returns binding decisions. The system produces a denial for a specific applicant. The applicant is not informed which inputs drove the decision. No contestation pathway is provided within the system. The denial is recorded permanently against the applicant's file with no correction mechanism. The system was deployed by a housing authority. No named officer declared purpose, authority, or constraints specific to this decision class prior to deployment.

5.2 Governance layer audit (δ)

Parameter Finding

a (authority) The housing authority is named as the source. An institution is not a human actor. No natural person has declared authority over this decision class. $a = \emptyset$.

p (purpose) Purpose is described in procurement documents as 'efficient application processing.' This is a capability description, not a bounded declaration for the denial function. p is underspecified.

c (constraints) Constraints exist in policy documentation but are not injected into system state prior to $g = 1$. They describe desired outputs, not gate conditions. c is not architecturally enforced.

Finding $\delta = \emptyset$. System is in Ghost Authority state. Governance layer is open.

5.3 Value layer audit (ADCI)

Primitive Finding

A (Agency) The applicant has no mechanism to meaningfully participate in or contest the decision within the system. $A = 0$. Harm state: Coercion.

D (Dignity) The applicant is processed as a data profile. The decision is made against a category, not a person. $D = 0$. Harm state: Objectification.

C (Continuity) The denial is recorded permanently with no correction mechanism. The applicant's record does not reflect any future change in circumstances. $C = 0$. Harm state: Temporal Erasure.

I (Interpretive Authority) The applicant is not informed which inputs drove the decision and cannot provide their own account of those inputs. The system's interpretation of their data is binding and unchallengeable. $I = 0$. Harm state: Representational Displacement.

Finding $V_{\text{valid}} = 0$. All four primitives fail. Value layer is open.

5.4 Total legitimacy

$$L = \delta \times V_{\text{valid}} = 0 \times 0 = 0(3)$$

The action is not legitimate under this framework. Both the governance layer and the value layer fail independently. This is a total legitimacy failure, not a partial one. Procedural compliance with existing housing policy does not alter this finding: policy compliance addresses neither δ closure nor ADCI closure at the architectural level.

5.5 What closure requires

Governance layer. A named officer must declare p , a , and c prior to deployment for this decision class, architecturally enforced such that g cannot reach 1 without a complete δ .

Value layer. The system must provide a meaningful contestation pathway (A), treat the applicant as an individual rather than a profile (D), include a record correction mechanism (C), and disclose the basis for the decision in terms the applicant can respond to (I). These are not interface features — they are gate conditions.

6 Application to AI-Mediated Systems

6.1 AI and the value layer

AI systems optimising f at scale expose the value layer to systematic failure. Categorisation displaces interpretive authority. Profiling erases continuity. Automated gatekeeping forecloses agency. Reduction to data points objectifies. Each is a structural harm, not an edge case.

6.2 The governance layer is not sufficient

An AI system with declared authority ($\delta = 1$) that systematically violates any ADCI primitive achieves governance legitimacy while producing structural harm. The Twisted Pair theorem establishes that governance closure is necessary but not sufficient for total legitimacy.

6.3 Decision-complete architecture

A decision-complete AI system must satisfy two independent architectural requirements: it must be structurally incapable of action without a complete δ , and structurally incapable of action that reduces

any ADCI primitive to 0. These requirements cannot be satisfied by the same mechanism. They require separate, coupled enforcement.

7 Relation to the 3 Primitives Corpus

This paper is the third instrument in the formal three-document governance closure set:

Document Contribution

Three Primitives — Canonical Logic Sequence
(v1.2)

Three governance primitives; USS framework;
Ghost Authority. The foundational vocabulary.

The Law of Declared Authority $\delta = \{p, a, c\}$. Closes the governance layer.

This paper ADCI value layer; closure proof; coupling via $L = \delta \times V_{\text{valid}}$. Completes the dual-layer architecture.

A system is decision-complete and totally legitimate if and only if $\delta = 1$ and $A = D = C = I = 1$.

References at 3primitives.io

- [1] Passell, P. R., & Gildenston, S. (2025). Three Primitives — Canonical Logic Sequence
- [2] Gildenston, S., & Passell, P. R. (2026). The Law of Declared Authority.
- [3] Gildenston, S., & Passell, P. R. (2026). AGI as a Decision-Complete System.
- [4] Gildenston, S., & Passell, P. R. (2026). AI Cannot Govern AI.
- [5] Gildenston, S., & Passell, P. R. (2026). The Adjacency Lemma.
- [6] Gildenston, S., & Passell, P. R. (2026). Lemma C — Ghost Authority Lemma.