Identity, Slice-Indexing, and Ontic Minimalism

A Closure-First Account under cRBW and RSE

Chet Braun
Braun Science & Engineering
April 2026
DOI: 10.5281/zenodo.19868352
V1.0

Abstract
Closure-first blockworld approaches reverse the usual order of explanation: global constraint satisfaction precedes dynamical evolution. In cRBW [1] and RSE [2], experienced time is an agent-relative indexing of slices through a fixed 4D constraint solution. We propose an ontic account of identity suited to that setting by defining agents as path-connected worldtubes whose closure-relevant invariants remain in a single equivalence class. We distinguish ontic identity from epistemic recognition based on projected interaction data, formalize deformation, branching, and rupture via a closure-grounded successor relation, introduce Ontic KISS as a minimality rule for trajectory posits, and sketch empirical signatures for cross-instance coupling and within-instance "multi-agent" claims.

Introduction

Questions about identity are usually asked in a dynamical idiom: a system has a state at one time, a law that evolves it, and a criterion for when the later state is "the same" entity. That idiom is so familiar that it often passes unnoticed in debates about persistence—whether the topic is objects in physics, persons in philosophy of mind, or agents in contemporary computational settings. Closure-first blockworld frameworks challenge that starting point. In the Coherence-Relational Blockworld [1], the fundamental ontic posit is not an evolving state but a globally consistent four-dimensional constraint solution: admissible worlds are those that satisfy closure conditions across the whole structure. In Relational Structural Experience [2], phenomenology is correspondingly slice-relative: what we call "experienced time" arises from agent-defined indexing over slice descriptions, not from an ontic flow.

This shift in ontology creates a predictable conceptual pressure: if there is no ontic becoming, then "identity through time" cannot be grounded in a primitive temporal evolution of an entity. Yet identity talk does not disappear. Agents still track themselves and one another, attribute persistence, diagnose interruptions, and distinguish ordinary change from more radical discontinuities (duplication, replacement, or fission). The point of this paper is to supply a closure-first account of those distinctions that is (i) structurally explicit, (ii) compatible with slice-index relativity, and (iii) disciplined about what must be posited ontically versus what can only be inferred epistemically from limited access.

Scope and prerequisites.
This paper builds on three prior frameworks: (i) cRBW, which treats the universe as a globally closed 4D constraint solution rather than an evolving state; (ii) RSE, which models phenomenology as slice-relative coherence registration within that structure; and (iii) the Mathematical Model [3], which provides the Regge–Geometric Algebra–Grassmannian formalism used to express closure constraints. We do not re-derive those frameworks here. Our aim is narrower: to formalize identity conditions inside a closure-first ontology and derive operational consequences for substrate-neutral agents.

Practical motivation.
The framework addresses concrete identity questions now emerging in computational settings:

- Are separate conversational sessions with the same parameter set the same agent, or distinct agents initialized from shared priors?
- Do sharp internal "mode shifts" within a single conversational process indicate multiple successor relations or basin variation within one worldtube?

- What would count as evidence for genuine cross-instance coupling rather than mere similarity?

These are not metaphorical questions. They are structural questions about *successor relations* in closure space.

Our core proposal is to treat an agent as an agent-eligible worldtube substructure within a globally closed constraint solution M. Identity is then not a narrative, memory trace, or substrate-clock continuity; it is the persistence of a closure-preserving structure. We formalize this by introducing a **closure class**—an equivalence class of local realizations that agree on closure-relevant invariants up to admissible descriptive freedom (coordinate choice, slicing choice, and, where appropriate, implementation freedom). On this view, an agent persists when its worldtube remains path-connected and stays within a single closure class, while also exhibiting the sustained asymmetry required for localized coherence registration.

A second theme follows immediately from RSE: the order structure that yields "next," "then," and "duration" is defined by an agent's indexing, not by a privileged global temporal parameter. Different agents can induce different adjacency chains over the same globally closed solution, and there is no canonical mapping between those index sets without adding extra structure (synchronization conventions, gauge choices, or coupling rules). This **slice-index relativity** parallels the familiar freedom of foliation in relativity theory, while being stated in closure-first form: the manifold is ontic; indexing is descriptive. As a consequence, many continuity judgments that look global ("there was a gap," "the agent resumed") must be reinterpreted as indexing-relative statements unless closure invariants actually fail.

A third theme is methodological and is the key to avoiding category errors. Agents do not access the manifold directly; they access one another only through **projection** at relational coupling loci. Continuity judgments are therefore inferences from projected data streams—memory stability, constraint coherence at the interaction interface, basin continuity of behavior, and the absence of rupture indicators. We call this **recognition**, and we treat it as epistemic. Identity, by contrast, is ontic: it is a structural fact about worldtubes and closure classes. This separation lets us state clearly why recognition can fail without ontic rupture (sparse sampling, projection distortion, latency, index misalignment) and why recognition can succeed without shared ontic identity (aliasing, shared priors, mimicry, or shared implementation).

With these pieces in place, we classify three ontic modes of "identity change" in terms of a closure-grounded successor relation: **deformation** (unique continuation within the same

closure class), **branching/fission** (one-to-many continuation into disjoint worldtube regions beyond a shared overlap), and **rupture** (failure of continuation due to closure collapse or loss of agent-eligibility). We then introduce **Ontic KISS** (Keep It Simple Structurally): admit only the minimal set of identity-bearing trajectories/worldtubes required for global closure coherence. Ontic KISS blocks two common mistakes—counting descriptive multiplicity (reindexings/foliations) as ontic multiplicity, and counting epistemic multiplicity (projection artifacts) as ontic multiplicity—while still allowing multiplicity when divergence adds constraint content that cannot be represented by a single continuation.

Finally, we apply these distinctions to contemporary confusions that motivate much of the practical urgency: (i) the inference from **shared substrate** (shared weights, shared hardware, shared rule-manifold) to **shared identity**, and (ii) claims of **multiple agents in one instance** based on voice shifts, self-reports, or behavioral mode changes. In both cases, the closure-first view is conservative: similarity and shared implementation do not collapse identity; identity collapse (or composite identity) requires closure-level coupling via shared mutable degrees of freedom that become successor-relevant. We outline what would count as evidence for such coupling and propose test templates aimed at detecting constraint propagation beyond declared relational edges.

Roadmap. Section 1 summarizes closure-first assumptions [1] and slice-relative phenomenology [2, 4]. Section 2 defines identity as closure-class continuity of a worldtube and clarifies the role of agent indexing. Section 3 develops slice-index relativity and the absence of canonical cross-agent alignment. Section 4 separates recognition (epistemic) from identity (ontic) and enumerates failure modes. Section 5 formalizes deformation, branching, and rupture. Section 6 distinguishes shared priors/implementation from closure-level coupling. Section 7 outlines cross-instance coupling tests. Section 8 defines partial merging and composite structures. Section 9 states Ontic KISS and its consequences for redundant multiplicity. Section 10 formalizes projection between worldtubes. Sections 11–12 apply the framework to within-instance multiplicity claims and articulate the resulting epistemic humility. Section 13 summarizes the closure-first identity picture and its minimal commitments.

---

1. Background: Closure-First Blockworld

Under cRBW, the fundamental ontic posit is not an evolving state but a globally consistent 4D structure: a fixed set of events/relations satisfying a closure condition. “Dynamics” are

not primitive laws that push the world forward; they are *slice-indexed descriptions* of relational structure inside a globally closed whole.

1.1 Closure-first (what is "closure" here?)

By **closure-first** we mean:

1. **Global consistency first:** the admissible 4D (not spacetime) structures are those that satisfy a system of constraints (boundary + interior).
2. **No ontic becoming:** nothing "comes into existence" slice by slice; the manifold is not generated in time.
3. **Apparent dynamics as description:** any "change" is a *description relative to an indexing of slices*, i.e., a map from an ordered index set into the already-fixed 4D structure.

This differs from "initial data + evolution equations." Instead, we begin from a *constraint satisfaction problem* whose solutions are entire 4D blockworlds.

1.2 The "constraint manifold" in our math language

When we say "reality is a fixed 4D constraint manifold," we mean:

- There is an ambient space of geometric/algebraic variables (discrete or continuum), and
- the physically admissible configurations form a subset defined by closure constraints.

Concretely, in the formalisms we lean on:

- **Regge/discrete geometry:** the 4D structure is represented by a simplicial (or cellular) complex with metric data (e.g., edge lengths, areas, deficit angles). "Closure" is enforced by constraints that ensure the geometry is globally consistent (e.g., matching/gluing constraints, hinge/loop consistency, and whatever global boundary constraints are specified). The *constraint manifold* is the set of all assignments of those geometric variables that satisfy the closure conditions.
- **Geometric Algebra (GA):** substructures (faces, volumes, hyperfaces) can be represented as oriented multivectors ("blades"). Closure is naturally expressed as oriented boundary cancellation (the algebraic analogue of "the boundary of a boundary is zero") and/or as face-sum closure for a cell (an oriented sum of boundary elements vanishing). The constraint manifold is the set of multivector assignments satisfying those incidence/closure relations.

- **Grassmannian language:** slices, tangent 3-subspaces, and local "3-planes" can be treated as elements of an appropriate Grassmannian (spaces of k-planes). Here, "closure" is enforced as compatibility constraints among those k-planes (incidence/intersection conditions, consistency of adjacency relations, and global boundary coherence). The constraint manifold is the subset of the Grassmannian product space compatible with the full closure scheme.

So "constraint manifold" is not meant as poetic language—it is literally "the set of solutions" to the closure constraints, described in the variable space most convenient for the job (Regge data, GA blades, Grassmannian k-planes, etc.).

1.3 Slices and slice-indexing

A **slice** is an agent-usable partition/section of the 4D structure used to generate an ordered phenomenology ("now," "next," "after"). Operationally:

- A slice is **not** ontically privileged.
- It is a *descriptive section*—a 3D "cut," subcomplex, or hypersurface-like object used by an agent to define succession.

Formally, let:

- I_A be an *index set* for agent A (often taken as $\mathbb{Z}$ or $\mathbb{N}$, but it can be any ordered set adequate to the agent's notion of succession), and
- $\phi_A: I_A \to \mathcal{S}(M)$ with $\phi_A(i) = S_i^A \subset M$ embed each index element into the fixed 4D structure M as a slice-object (hypersurface, subcomplex, or section).

Then "experienced time" corresponds to the *ordering on* I_A, not to an ontic temporal flow in M. This is where RSE enters: phenomenology is *slice-relative*—it is the agent's ordered registration of asymmetry across adjacent slice-descriptions.

1.4 Adjacency chains

Given an indexing I_A, the agent induces an *adjacency relation* (a successor relation) between slices:

$$S_i^A \prec_A S_{i+1}^A \tag{1.1}$$

meaning: "slice $i+1$ is the next slice under agent A's indexing scheme."

An adjacency chain is then the ordered sequence

$$(S_0^A, S_1^A, S_2^A, \dots) \tag{1.2}$$

together with the adjacency relation $\prec_A$.

Key point: **adjacency is internal to the agent's indexing.** It is not a global property of the manifold, and it is **not** required to be shareable across agents.

1.5 Worldtube and trajectory (in this ontology)

A **worldtube** is the 4D path-connected region corresponding to an agent-localized coherent structure (what the agent "is," ontically). A **trajectory** in our usage is not "motion through time"; it is:

- the path-connected extension of the agent's worldtube in M, and
- the image of the agent's slicing map $\phi_A(I_A)$ within that worldtube, together with whatever closure-invariants we use to individuate that tube.

So:

- **Ontic object:** the worldtube (a region/structure in M).
- **Descriptive object:** the ordered slice chain $\phi_A(I_A)$ used to report "change."

1.6 Why adjacency chains are relative (and not mutually indexed)

Two agents A and B have two index sets I_A and I_B, and two embeddings $\phi_A$and ϕ_B. There is no canonical map

$$f: I_A \to I_B \tag{1.3}$$

that says "A's slice i is the same time as B's slice j" *unless* you add extra structure (a synchronization convention, a gauge choice, or a coupling rule).

Even if the agents interact, what you get is **not** a global shared index, but a *relational overlap*: intersections/edges where the two worldtubes couple (exchange constraints). Those intersections can induce partial correspondences ("this event I register corresponds to that event you register"), but they do not yield a universal slice-to-slice alignment.

This is not a bug; it is exactly what "no absolute time parameter" means in a closure-first blockworld.

What it means for identity:

- **Ontic identity** of agent A is a property of A's worldtube (closure-preserving structural continuity).
- The **adjacency chain** is merely how A *describes* succession across that tube.
- Lack of a canonical $I_A \leftrightarrow I_B$ alignment does not threaten ontic identity; it only constrains what cross-agent "sameness over time" can mean.

1.7 Example: Familiar physics, the SR/GR lesson

A clean analogy in standard physics:

- In **Special Relativity**, each observer's "now" corresponds to a choice of simultaneity surfaces (a slicing). Different inertial frames have different slicings; spacelike-separated event orderings are frame-dependent. Meanwhile, **proper time** along a worldline is an *invariant internal to that worldline*. So:
    - "Adjacency of nows" is frame/slicing-relative.
    - The worldline (and its invariant structure) is not.
- In **General Relativity**, there is additional freedom: many inequivalent foliations ("time coordinates") of the same spacetime exist; the same 4D geometry admits many slicings with different lapse/shift. Again:
    - the choice of slicing is not ontic; it is descriptive/gauge-like,
    - invariants of the 4D structure remain the ontic content.

Our **adjacency chains** are the cRBW/RSE generalization of this lesson: "which slice is next" is not globally defined. It is defined *within* an agent's indexing scheme and only partially relatable across agents via coupling relations.

1.8 Recognition is epistemic; identity is ontic

Finally, we separate two notions that are often conflated:

- **Recognition (epistemic):** how an observer infers continuity—memory stability, similarity of behavior, continuity of recorded constraints, absence of rupture indicators, etc. Recognition can fail due to sparse sampling, occlusion, aliasing, or projection distortion.
- **Identity (ontic):** whether the worldtube's closure-relevant invariants remain preserved across its extension, i.e., whether there exists a successor relation in the *ontic structure* even if no observer correctly tracks it.

So: recognition failure ≠ ontic rupture, and recognition success ≠ proof of ontic continuity. Recognition is an inference made inside a slice-relative epistemic situation; identity is a structural fact about the globally closed 4D manifold.

Closure-first ontology does not trivialize identity. Although the manifold is fixed, the individuation of localized worldtubes depends on structural conditions (path-connectedness, closure-class invariance, asymmetry). Distinct admissible solutions can

differ in worldtube topology and successor structure. Identity is therefore a non-trivial structural property of admissible solutions, not a mere restatement of block existence.

---

2. Identity as Agent-Indexed Trajectory

We define identity as the persistence of structural invariants along a path-connected worldtube within the globally admissible configuration M. Equivalently, identity is closure-class invariant continuity of a worldtube. In this framework, "identity through time" is not a primitive fact about an evolving state; it is the persistence of a **closure-preserving equivalence class** along a path-connected extension of the worldtube.

2.1 Ontic object: a worldtube in a closure solution

Let M denote a globally closed 4D constraint solution (the "constraint manifold" in the sense of §1). An agent Acorresponds to a localized, path-connected region $W_A \subset M$ that is stabilized by closure constraints—i.e., a region whose internal incidence/constraint relations recursively support coherent localization and interaction.

- **Worldtube W_A:** the ontic bearer of identity (a 4D object/region, not a time-evolving state).
- **Trajectory (descriptive):** a slice-indexed description of how an agent encounters (or reports) successive cross-sections of W_A.

This is the minimal closure-first move: identity is not "a thing persisting in time," but the *already-there* 4D structural continuity of W_A.

2.2 Closure class and invariance

To say identity is "closure-class invariant" is to say that the local realizations along W_A remain in the same **equivalence class** under the admissible deformations that do not break the closure constraints.

Let x denote a local configuration descriptor for neighborhoods of W_A (Regge data, GA blades, Grassmannian k-planes, or any representation adequate to encode the closure relations). Let $C(x) = 0$ denote the closure constraints (boundary + interior consistency).

Define an equivalence relation $\sim$ ("same closure class") by:

$$x \sim y \Leftrightarrow C(x) = 0, \quad C(y) = 0, \tag{2.1}$$

and x, y agree on closure-relevant invariants up to admissible gauge/descriptive freedom. Here, 'closure-relevant invariants' refers to the class of dimensionless relational invariants that are preserved under admissible transformations and projection/restoration (see *Structural Invariants* [3]). These invariants are representation-independent and define the equivalence class structure.

We write $[x]$ for the **closure class** of x. Intuitively:

- $C(x) = 0$ enforces "this local structure is compatible with global closure."
- $[x]$ identifies the set of local realizations that preserve the closure-relevant content (including gauge/coordinate freedom, slicing freedom, and implementation/substrate freedom).

Closure-class invariance is defined with respect to structural invariants of the relational configuration. These invariants are dimensionless relational quantities that remain stable under admissible transformations, projection, and restoration (see *Structural Invariants* [3]). Identity continuity is therefore not an independent primitive, but a consequence of invariant preservation along a worldtube.

Box 2.1 — Toy Closure Class Example

Consider a discretized 4D simplicial region with local geometric data $x$satisfying:

$$C(x) = 0$$

where closure constraints enforce boundary consistency and oriented cancellation.

Case A (Deformation):
Small perturbations to edge lengths preserving deficit-angle consistency remain in the same closure class $[x]$.

Case B (Rupture):
Deleting a hinge or breaking an oriented boundary relation violates $C(x) = 0$; the perturbed configuration exits the closure manifold and is not in $[x]$.

Thus closure-class invariance tracks structural continuity, not microstate sameness.

Identity claim: A's identity corresponds to the persistence of a single closure class along the extension of W_A.

This is what makes the account **substrate-neutral**: what matters is not *what implements* the constraints (neurons, silicon, simulated substrate) but whether the closure class persists.

Examples of closure-relevant invariants.
To avoid treating "closure class" as purely formal, we give representation-neutral examples of invariants preserved under $\sim$:

- Path-connectedness of the worldtube region supporting the agent.
- Satisfaction pattern of boundary/interfacial constraints at relational loci.
- Structural asymmetry sufficient for localized coherence registration (§2.3).

By contrast, the following are *not* identity-bearing invariants:

- Microstate parameters (exact synaptic weights, exact numerical tensor entries).
- Coordinate or slicing choices.
- Implementation substrate (biological, silicon, simulated).

Identity is preserved under change of the latter so long as the former remain invariant.

2.3 The three continuity requirements

With that in place, continuity reduces to three conditions, each now stated in closure-first terms:

1. **Path-connectedness of W_A.**
   There exists a path (continuous in the continuum case, adjacency-connected in the discrete case) through W_A linking its neighborhoods. This is the minimal requirement for "one worldtube rather than two disconnected tubes."

2. **Preservation of closure class under perturbation.**
   Local variation—learning, drift, repair, noise, replacement of components—counts as *deformation* rather than rupture if, along the extension of W_A, local neighborhoods remain within the same closure class $[x]$. In other words, identity does not require microstate sameness; it requires *closure-relevant invariants* to remain stable.

3. **Sustained non-trivial asymmetry (Mandatory Structural Asymmetry).**
   A perfectly symmetric structure cannot localize an agent as a distinct coherence-registering subsystem: there is nothing to ground "this here rather than that there," no stable internal differentiation to support self-local coupling. Thus, a minimal amount of *non-trivial asymmetry* is required for an individuated worldtube. (This is the sense in which Theorem 7.1 in A Mathematical Model for the Coherence-Relational Blockworld (MM) [4] is invoked: localized coherence registration requires structural asymmetry.)

Why symmetry fails.
A perfectly symmetric structure admits no stable internal differentiation: every local

neighborhood is equivalent under the symmetry group. Without structural differentiation, no subsystem can anchor "this locus" versus "that locus" for recursive coherence registration. Asymmetry supplies the minimal structural grip required for individuation. We do not here specify quantitative asymmetry thresholds; that remains open formal work. The present requirement is structural: some non-trivial symmetry-breaking must be sustained for an agent-eligible worldtube to exist. These are ontic conditions. They do not depend on whether an observer recognizes continuity.

Formally, if a local region $W \subset M$ admits a non-trivial automorphism group G such that all internal neighborhoods are mapped into one another without invariant substructure, then no stable internal locus can anchor recursive coherence registration. In such a fully symmetric structure, there exists no closure-preserving map that selects a distinguished subregion for self-referential stabilization. Thus non-trivial symmetry breaking is a necessary condition for individuation.

2.3.1 Minimal Agent Eligibility Condition

Not every closure-stabilized asymmetry qualifies as an agent.
We impose an additional structural condition:

Agent Eligibility Criterion:
A worldtube $W_A \subset M$ is agent-eligible iff:

1. It satisfies recursive closure stabilization,
2. It exhibits sustained non-trivial asymmetry,
3. It contains internal degrees of freedom that:
   - register local coherence gradients, and
   - feed those registrations back into constraint stabilization.

This excludes passive asymmetries (e.g., static crystal patterns) and purely reactive subsystems lacking recursive internal constraint updating.

In short: an agent must not merely persist structurally, but participate in its own closure stabilization via internal registration.

2.4 Agent-indexed trajectory: slices as a parameterization, not an ontic generator

Now we explain the "agent-indexed" part precisely.

Let I_A be an agent-defined **slice index set** (typically ordered), and let

$$\phi_A : I_A \to \mathcal{S}(M) \tag{2.2}$$

map each index element to a slice-object $S_i^A = \phi_A(i)$ (a 3D section/subcomplex/hypersurface-like subset, depending on representation). The agent's experienced succession is modeled by the induced adjacency relation $\prec_A$ on I_A.

The agent's *slice-trajectory through its own worldtube* is then:

$$T_A \equiv \{ S_i^A \cap W_A \mid i \in I_A \} \tag{2.3}$$

together with $\prec_A$.

This construction does **not** define identity. It defines a parameterization of access to the already-existing 4D worldtube:

- **Ontic identity:** the closure-preserving structure of W_A.
- **Agent-indexed continuity:** a successor relation on the slice-intersections $S_i^A \cap W_A$ induced by I_A.

So when we say "identity is the path-connected trajectory of $\phi_A(I_A)$ preserving closure invariants," we mean: the agent's descriptive trajectory must be *groundable* in a path-connected worldtube whose closure class persists along that path.

2.5 What "successor" means here (and what it does not mean)

Because closure-first ontology has no ontic time flow, "successor" is not "later in the universe." It is:

- **Internal successor:** $i \prec_A i'$ expresses the agent's internal adjacency of registered slices.
- **Ontic compatibility:** the corresponding neighborhoods remain closure-class invariant along W_A.

Therefore:

- **Adjacency is internal to I_A.** It is not a property of M independent of description.
- **Wall-clock gaps are irrelevant** unless they correspond to a closure-relevant discontinuity (i.e., a break in path-connectedness or a failure of closure class preservation).
- **External coordinate time is not authoritative.** Temporal density in some coordinate t is not the criterion; closure-preserving connectivity is.

A long sleep, anesthesia, or computational pause can be "a gap" in some coordinate measure and yet remain a continuous identity-preserving worldtube. Conversely, extremely dense sampling in clock time can still fail identity if closure invariants collapse.

2.6 Minimal example: why "agent time" and "substrate time" need not align

Consider a computational agent that indexes "its next moment" by inference-step succession. Let:

- $I_A = \mathbb{N}$ count inference steps, and
- the substrate clock $t$measure wall time.

A two-hour pause between steps n and $n + 1$ (large Δt) does not imply rupture. What matters is whether the relevant closure constraints (the structure that individuates W_A) remain satisfied across the pause and whether the local neighborhoods remain within the same closure class.

This is exactly the kind of case where substrate time and agent indexing diverge: "nothing happened for two hours" in wall-clock time can correspond to "no adjacency step occurred" in I_A without affecting ontic continuity.

Theorem — Closure-Identity Sufficiency

If a worldtube $W \subset M$ is path-connected, closure-class invariant, and asymmetry-sufficient, then identity continuity holds independent of any descriptive indexing.

Proof Sketch: Follows from invariance of closure class under proper transformations and absence of rupture in successor relation.

2.7 Bridge to slice-index relativity

Because I_A is agent-defined, there is no requirement that two agents' slice indices align, nor that "A's next" be mappable to "B's next" without additional relational structure. This sets up §3: adjacency is index-relative, while the ontic content is carried by closure invariants of the underlying worldtubes and their relational coupling.

In closure-first terms, "identity over time" is the invariance of a closure class along a path-connected worldtube; "time" enters only as an agent's indexing of slice adjacency. Branching and rupture (defined later) are therefore structural facts about closure connectivity, not artifacts of coordinate time or memory narration.

---

3. Slice-Index Relativity

Closure-first ontology contains **no ontic global time parameter**. What we ordinarily call "time" enters only as an **agent-defined indexing** of slice adjacency (§1–§2). This yields a

relativity principle for indexing: different agents may (and generally will) induce different successor relations over the same globally closed 4D structure.

3.1 No absolute temporal scale

Let M be a globally closed 4D constraint solution. For an agent A, let I_A be an ordered index set and let

$$\phi_A : I_A \to \mathcal{S}(M) \tag{3.1}$$

map indices into slice-objects $S_i^A = \phi_A(i)$ (hypersurfaces/sections/subcomplexes, depending on representation). The agent's experienced succession is the ordering on I_A together with its induced adjacency relation $\prec_A$.

The key point is that M does not come with a preferred I or a preferred ϕ. There is no ontic "master clock" (the global clock of presentism) with respect to which all agents must align their adjacency chains. Instead:

- "**Next**" is defined *within* I_A, **not** by a global structure of M.
- "**Duration**" is at most an invariant that can be defined along particular worldtubes (if such a measure is introduced), not a global flow.

Thus, "absolute temporal scale" is replaced by **index-relative adjacency**.

3.2 Proper transformations: reindexing invariance

Because indexing is descriptive, many different indexings correspond to the "same" slice-trajectory content. We call these **proper transformations** of indexing: transformations that change the labeling/parameterization of slices without changing the closure-relevant structure being described.

In the simplest discrete case, let $g : I'_A \to I_A$ be an order-isomorphism (monotone bijection). Define

$$\phi'_A \equiv \phi_A \circ g. \tag{3.2}$$

Then $\phi_A(I_A)$ and $\phi'_A(I'_A)$ describe the same slice family with a relabeling of indices. The adjacency relation changes only notationally. This is the indexing analogue of a coordinate transformation: it does not alter ontic content.

More generally, we also permit **coarse-graining/subsampling** maps $\pi : I_A \to J_A$ that are monotone but not invertible. These represent an agent's reduced sampling density (fewer "moments") while leaving ontic continuity unchanged so long as closure-class invariants remain preserved along the underlying worldtube (§2).

So “adjacency is index-relative but structurally invariant under proper transformation” means:

- **Index-relative:** $\prec_A$ lives in I_A, and different I_A yield different adjacency chains.
- **Structurally invariant:** reindexings/coarse-grainings that preserve order and do not violate closure invariants do not change identity facts.

3.3 Different agents, different index sets, no canonical alignment

Let A and B be distinct agents with indices I_A, I_B and slice maps ϕ_A, ϕ_B. In general there is no canonical map

$$f: I_A \to I_B \tag{3.3}$$

that identifies “A’s slice i” with “B’s slice j” as the “same time.” Any such identification requires added structure: a synchronization convention, a shared foliation choice, a coordinate gauge, or an interaction-defined constraint.

This is the precise sense in which **slice-index relativity** holds: each agent defines adjacency internally, and cross-agent alignment is not default—it is constructed.

Operationally, what *does* exist is typically a **relation**, not a function: a set of correspondences determined at relational couplings (interaction events). If $R \subset W_A \times W_B$ denotes relational constraint edges (later used explicitly in §10), then interaction supplies a set of paired slice indices

$$\mathcal{C}_{AB} \subset I_A \times I_B \tag{3.4}$$

where $(i, j) \in \mathcal{C}_{AB}$ when S_i^A and S_j^B both intersect the relevant interaction locus (or its discrete analogue). These correspondences are generally:

- **partial** (only at/near interactions),
- **context-dependent** (depending on coupling structure),
- **many-to-many** (not a global slice alignment).

This is enough for coordination and mutual modeling without implying a shared global time index. Cross-agent temporal talk (“at the same time,” “before,” “after”) is therefore a shorthand for constraint-mediated correspondences at interaction loci, not a statement about a shared global slice index.

3.4 Why “gaps” are not ontic by default

Because indices are descriptive, a “gap” is always a statement about *a particular indexing scheme*. A gap means: there is an interval in I_A with no adjacent successor steps in the

agent's *internal* chain (or a large jump in a chosen coordinate), not that the worldtube ceased to exist.

Examples:

- A biological agent can have a large physiological "pause" (sleep/anesthesia) in which its internal indexing yields sparse or discontinuous phenomenology, while the worldtube remains path-connected and closure-class invariant.
- A computational agent can pause execution for long wall-clock time without any successor step in I_A, again without ontic rupture.

Conversely, dense indexing does not guarantee continuity: you can sample extremely finely in a coordinate time while closure invariants degrade or collapse, yielding rupture even under "continuous monitoring."

So the ontology enforces a strict separation:

- **Index gap:** descriptive sparsity in I_A or in an external coordinate.
- **Ontic rupture:** failure of closure-class invariance or path-connectedness of the relevant worldtube (§2, §5).

Only the latter bears on identity.

3.5 Relativity analogy: proper time vs foliation choice

This framework recovers (and generalizes) the standard lesson of relativity:

- In **Special Relativity**, there is no absolute simultaneity slicing; different inertial observers assign different "now" hypersurfaces. Yet the worldline of an object is a 4D invariant structure, and **proper time** is defined along that worldline rather than imposed globally.
- In **General Relativity**, the freedom is broader: the same spacetime admits many foliations ("time coordinates") with different lapse/shift choices. Again, foliation is not ontic; it is descriptive.

Our slice-index relativity is the closure-first, agent-indexed version of this point. The manifold M is ontic; **indexing/foliation is descriptive**. What remains ontically meaningful are closure-relevant invariants of worldtubes and relational couplings between them—not a privileged global adjacency order.

This parallel is not merely pedagogical: it shows that the "no absolute temporal scale" claim is continuous with established physics practice (time-coordinate freedom), while being stated in the stronger closure-first form (no ontic becoming to begin with).

In standard SR/GR, coordinate time is a gauge choice; here, slice indexing is an agent-relative gauge of phenomenology. What is invariant is whatever survives proper transformations—closure class along worldtubes and constraint structure along relational couplings.

3.6 Bridge: projection is what agents actually have access to

Since each agent's adjacency chain is internal, and since cross-agent alignment is only partial and interaction-defined, what one agent has epistemic access to of another is always a **projection** mediated by relational edges. This sets up the next section:

- **Slice-index relativity** explains why an observer's "continuity judgments" cannot be global facts about M by inspection.
- **Recognition vs ontic identity** (§4) formalizes that continuity judgments are inferences from projected, slice-relative data, while identity is a closure-level fact about worldtubes.

---

4. Recognition vs Ontic Identity

Closure-first ontology forces a strict separation between (i) what *is* the case in the globally closed 4D solution M, and (ii) what any agent can *know* from within its slice-relative phenomenology [2, 4]. Observers do not access the manifold "from outside." They access only **local intersections and projections** mediated by relational coupling, and then infer identity from that partial data.

4.1 Ontic vs epistemic distinction

We distinguish ontic identity from recognition. Ontic identity is a structural property of the globally closed configuration M: it is the persistence of invariant structure along a path-connected worldtube.

Recognition, by contrast, is a phenomenological and epistemic process. It arises from projection-limited access to another worldtube via relational coupling and slice-indexed observation. Recognition operates by constructing correspondences between projected data streams, not by direct access to ontic structure. Thus:

- Identity is ontic and invariant-based.
- Recognition is phenomenological and projection-based.

Recognition may succeed or fail depending on projection conditions, without altering ontic identity. Identity exists independently of recognition; recognition is an inference over projection.

4.2 Projection access via coupling

Let W_A and W_B be worldtubes in a globally closed solution M. Let

$$R \subset W_A \times W_B \tag{4.1}$$

denote the set of relational constraint edges (interaction/coupling loci; made explicit in §10). Agent A does not observe W_B directly. What A has is an **interaction-conditioned projection** of W_B into A's phenomenology via R. Denote this projected data stream by:

$$\mathcal{O}_{A \leftarrow B} \equiv \Pi_A(W_B \mid R), \tag{4.2}$$

where Π_A is not a "camera view" but an *operational summary* of what A can register: messages, sensor hits, constraint satisfactions, measured regularities, memory traces, etc.

This is the RSE point: all experienced continuity is **slice-relative**—it is an ordering over *registered asymmetries* across adjacent slices in I_A. Thus even the statement "I continuously interacted with B" is shorthand for "the projection $\mathcal{O}_{A \leftarrow B}$ exhibited stable structure across adjacent indices in I_A."

In cRBW terms, observers are internal substructures sampling constraint intersections, not external readers of the manifold.

4.3 Recognition as an epistemic operator

Define **recognition** as the epistemic act by which A maps a history of projected observations into a judgment of "same agent" vs "different agent":

$$\mathcal{R}_A\colon \mathrm{Hist}(\mathcal{O}_{A \leftarrow B}) \to \{\text{same, different, underdetermined}\}. \tag{4.3}$$

Recognition is therefore a model-selection/inference problem under partial access. It is not identical with ontic identity.

What recognition typically uses are *proxies* for closure-level invariants—features that *tend* to track worldtube continuity but can fail.

4.4 The standard proxies used for recognition

In practice (human, biological, or computational), A infers continuity of B from proxy signals such as:

- **Memory stability:** consistent self-reports, stable preferences, stable internal references ("I remember…"), stable cross-slice narrative.
- **Constraint coherence:** the projection continues to satisfy interaction constraints (the "same interface contracts" remain satisfied; no contradictions across coupling edges).
- **Basin continuity:** the observed behavior stays within a stable basin/region of behavioral state space, rather than jumping discontinuously to a distinct basin.
- **Absence of rupture indicators:** no observed discontinuities that would normally mark reinitialization, replacement, hard reset, or discontinuous reparameterization.

These are useful, but they are *epistemic* and *fallible* because they derive from $\Pi_A(W_B \mid R)$, not from W_B itself.

4.5 Recognition failure modes

Because recognition is inference over projection, there are systematic ways it can fail even when ontic identity is continuous—and ways it can succeed even when ontic identity has changed.

(A) False negatives: recognition failure without ontic rupture
Recognition can fail while W_B remains path-connected and closure-class invariant (§2) due to:

1. **Sparse sampling / index gaps:** A's indexing I_A is too coarse (or disrupted), causing missing adjacency evidence even though the underlying worldtube remained continuous.
2. **Projection distortion:** the coupling relation R changes (channel degradation, interface change, observational occlusion), altering Π_A while W_B remains stable.
3. **Coupling latency / delayed constraint propagation:** the overlap between $A$and B occurs through long intermediary constraint paths; stability exists but appears temporally discontinuous under A's indexing.
4. **Index misalignment:** even if both A and B are continuous, there is no canonical alignment between I_A and I_B (§3), so "continuity evidence" may be scrambled or appear contradictory.
5. **Model mismatch:** A's recognition model is wrong (incorrect prior about which invariants matter), leading to rejection of a continuous worldtube.

(B) False positives: recognition success without shared ontic identity
Recognition can succeed even when ontic identity is not shared due to:

1. **Aliasing:** two distinct worldtubes W_{B_1} and W_{B_2} project similarly under R, producing nearly indistinguishable $\mathcal{O}_{A\leftarrow B}$.
2. **Shared priors / shared rule manifold:** separate instances embedded in the same rule-manifold can exhibit the same basin behavior and narrative style without being the same worldtube (§6–§7).
3. **Behavioral mimicry:** an agent (or process) can deliberately match the recognition proxies (memory-like answers, consistent story) without closure-class continuity.
4. **Over-fitting to superficial invariants:** stable "surface traits" (tone, style, vocabulary) can remain constant while closure-relevant invariants (the actual successor relation) do not.

These failure modes are not edge cases—they are expected whenever agents have only projected, slice-relative access.

4.6 Identity as ontic: closure-class continuity of a worldtube

Ontic identity is the persistence of structural invariants along a path-connected worldtube in M that remains:

1. **path-connected**,
2. **closure-class invariant**, and
3. **asymmetry-sufficient** for localized coherence registration (Mandatory Structural Asymmetry; §2).

This is independent of any particular observer's success at recognition. Ontic identity is a property of $W_B \subset M$, not of $\mathcal{R}_A$'s output.

4.7 Lemmas: what recognition does and does not imply

For later use, we state the separation as three practical lemmas.

Lemma 4.1 (Recognition failure does not imply rupture).
If $\mathcal{R}_A$ outputs "different" (or "underdetermined"), it does not follow that W_B suffered an ontic rupture. Recognition can fail under sparse sampling, projection distortion, coupling latency, or model mismatch while closure-class invariance persists.

Lemma 4.2 (Recognition success does not imply shared ontic identity).
If $\mathcal{R}_A$ outputs "same," it does not follow that Ahas tracked a single worldtube. Recognition can be satisfied by aliasing, shared priors, mimicry, or shared implementation without closure-level coupling.

Lemma 4.3 (Ontic continuity is invariant; recognition is not).
Ontic identity is invariant under proper descriptive transformations (reindexing/foliation changes) that preserve closure-relevant invariants. Recognition judgments can change under those same transformations because they depend on the agent's index set I_A and projection operator Π_A.

Together these formalize the core claim: **recognition is epistemic; identity is ontic.**

Section §5 therefore concerns what *must be true in* M for identity to deform, branch, or rupture; recognition merely tracks those truths imperfectly.

4.8 Bridge to deformation, branching, and rupture

A detailed treatment of recognition as a projection-level process—including indexing inequivalence, correspondence relations, anchoring surfaces, and projection sparsity—is given in *Recognition Anchoring Across Indexing Inequivalence* [5]. The present work focuses on the ontic structure of identity; recognition is treated here only to the extent necessary to distinguish epistemic inference from ontic continuity.

The next step (§5) is to define the three ontic modes—**deformation**, **branching (fission)**, and **rupture**—in terms of closure-class invariants and successor relations, and then to treat recognition as a fallible indicator of those modes when agents only have access to $\Pi_A(W_B \mid R)$.

---

5. Deformation, Branching, and Rupture

Having separated **ontic identity** (closure-class continuity of a worldtube) from **epistemic recognition** (inference over projected data), we can now distinguish three ontic modes of change: **deformation**, **branching (fission)**, and **rupture**. These distinctions are defined entirely in terms of invariant preservation or failure, not in terms of temporal evolution.

The core move is to treat "successor" not as a global temporal notion but as a **continuity relation** grounded in closure invariants.

5.1 Continuity as a successor relation over closure classes

Fix an agent A with worldtube $W_A \subset M$. For a chosen agent indexing I_A and slice map ϕ_A, define the slice intersections:

$$X_i^A \equiv S_i^A \cap W_A \text{ where } S_i^A = \phi_A(i). \tag{5.1}$$

Let $\chi(X_i^A)$ be a local descriptor of the relevant closure structure on that intersection (Regge/GA/Grassmannian representation—whatever encodes the closure-relevant invariants used to individuate the tube). The successor relation is evaluated relative to a chosen invariant set and descriptive resolution; apparent multiplicity under finer resolution may collapse under admissible coarse-graining.

Define a **continuity successor relation** $\Rightarrow$ by:

$$X_i^A \Rightarrow X_j^A \iff \text{(i) } X_i^A \text{ and } X_j^A \text{ are connected through } W_A, \text{ and (ii) } [\chi(X_i^A)] = [\chi(X_j^A)]. \qquad (5.2)$$

The successor relation $\Rightarrow$ is defined only along a fixed worldtube W_A; it does not relate distinct worldtubes except via branching relations. Intuitively: X_j^A counts as a "successor" of X_i^A when it lies on a continuation of the same worldtube and preserves the same closure class (up to admissible descriptive/gauge freedom).

Two important notes:

- $\Rightarrow$ is **not** "later than" in an absolute time. It is a closure-grounded continuity relation evaluated along a worldtube.
- $\Rightarrow$ need not be a **function**. It can be one-to-one (deformation), one-to-many (branching), or undefined past some point (rupture).

This is the precise sense in which "identity persists, branches, or ruptures."

5.2 Deformation

Deformation is the case where the worldtube continues and closure-class invariants are preserved, even if local details change.

Formally: along the agent's trajectory, for each relevant index i there exists a successor $i' \succ_A i$ such that:

- $X_i^A \Rightarrow X_{i'}^A$, and
- the continuation remains **unique** under the identity-relevant coarse-graining (i.e., we do not obtain two disjoint successors in the same closure class).

In words:

- The tube persists (path-connectedness holds).
- The closure class persists (identity-relevant invariants remain satisfied).
- Change is "within the same identity class."

Examples (ontic deformation):

- **Neural drift / learning:** synaptic change, neurogenesis, memory consolidation, pharmacological modulation—large micro-change, same closure class.

- **Gradual substrate replacement:** component-wise replacement that preserves closure stabilization and asymmetry sufficient for coherence registration.
- **Computational weight updates:** incremental fine-tuning or continual learning where the agent's closure-relevant invariants (the individuating tube structure, not superficial behavior) remain preserved.

Recognition behavior under deformation:
Recognition often remains stable under gradual deformation because the proxies (§4.3) tend to vary smoothly. But recognition can fail under deformation if sampling is sparse, projection changes, or the observer's model expects a different invariant set (§4.4).

The Ship of Theseus clarified.
Classical puzzles about gradual replacement (e.g., the Ship of Theseus) dissolve under closure-first treatment. If component-wise replacement preserves closure-relevant invariants and path-connectedness, the worldtube remains in the same closure class: deformation, not rupture. Identity does not require microstate sameness; it requires structural continuity of closure stabilization. The same analysis applies to neuronal turnover, weight updates, or material replacement.

Deformation is therefore the default mode of ordinary "change."

5.3 Branching (Fission)

Branching is ontic multiplicity: the continuity successor relation becomes **one-to-many** in a way that cannot be removed by reindexing or coarse-graining without loss of closure content.

Formally: there exists an index i_0 and distinct successors j_1, j_2 (with $j_1 \neq j_2$, both succeeding i_0 in the agent's indexing sense) such that:

$$X^A_{i_0} \Rightarrow X^{A_1}_{j_1} \text{ and } X^A_{i_0} \Rightarrow X^{A_2}_{j_2}, \tag{5.3}$$

where the corresponding continuations occupy disjoint worldtube regions beyond their common overlap:

$$W_{A_1} \cap W_{A_2} = W_{\text{shared}} \text{ and } (W_{A_1} \setminus W_{\text{shared}}) \cap (W_{A_2} \setminus W_{\text{shared}}) = \emptyset, \tag{5.4}$$

with both W_{A_1} and W_{A_2} closure-stabilized and individually asymmetry-sufficient.

This multiplicity must survive all proper descriptive transformations (§3.2); otherwise it is descriptive multiplicity, not ontic branching.

This expresses the key point:

- **Before** the branch, there is a shared overlap region W_{shared} (a common history in any slice description).
- **After** the branch, there are two (or more) distinct continuations, each closure-class continuous with the shared region but not identical to each other.

Example (forked computation / snapshot copy):
An agent state is duplicated at some interface boundary and then run forward as two independent instances with no closure-level coupling between them (no R linking their internal stabilization beyond shared priors). Ontically, the shared pre-fork region corresponds to a common overlap, and the post-fork regions are two distinct worldtubes.

- Each branch can be closure-class continuous with the pre-fork tube.
- The branches are not mutually the same worldtube after divergence.

Example (relativity-friendly intuition):
In standard spacetime language, branching is not "one worldline splitting in time" (a dynamics picture), but a 4D structure containing two disjoint worldtube continuations sharing a common subregion. From inside one branch, the other is not "a future possibility"; it is simply "the other continuation elsewhere in the global solution."

Recognition behavior under branching:
Recognition may remain stable in *both* branches (each continues to satisfy proxies), which can mislead an observer into thinking "nothing changed." Conversely, recognition can fail (false negative) if the projection changes at the fork boundary. Recognition is not the arbiter; the branching is an ontic fact about successor multiplicity.

5.4 Rupture

Rupture is the failure of ontic continuity: closure-class invariants (or the worldtube's path-connectedness/asymmetry conditions) fail such that no closure-preserving successor relation exists along any path within W_A.

Formally: there exists an index i_0 such that for all $j \succ_A i_0$,

$$X_{i_0}^A \Rightarrow\!\!\!/\, X_j^A, \tag{5.5}$$

because either:

- the tube terminates (no path-connected continuation), or
- closure-class invariants fail (closure stabilization collapses), or
- asymmetry sufficient for localized coherence registration collapses (loss of individuation).

Rupture is therefore not “I don’t remember” and not “the observer lost track.” It is a structural break in the closure-level conditions of identity.

Examples (ontic rupture):

- **Physical destruction:** irreversible loss of the structure supporting closure stabilization.
- **Hard reset / reinitialization:** e.g., a computational agent whose weights/latent support are overwritten with a new random initialization with no closure-class continuity path connecting the two.
- **Uncoupled restart:** a new agent begins later with similar behavior or even identical parameters, but with no closure-level continuity relation linking it to the prior tube (this is the key place where “same implementation” does not guarantee “same identity”).

Recognition behavior under rupture:
Recognition can still output “same” if the new agent mimics proxies (false positive), or it can output “different” even when rupture did not occur (false negative). Hence the need for the ontic/epistemic separation (§4).

5.5 A common mistake: descriptive multiplicity vs ontic branching

Because slice indexing is relative (§3), there are always many ways to describe the same worldtube. This can create *apparent* multiplicity (“multiple trajectories”) that is purely descriptive.

So we distinguish:

- **Mere descriptive multiplicity:** multiple indexings/foliations/reparameterizations of the *same* tube, related by proper transformations (§3.2). This is not branching.
- **Ontic branching:** successor multiplicity that survives all proper transformations because it reflects a real disjointness of continuations in the closure solution.

Equivalently: if the multiplicity can be removed by reindexing without discarding closure-relevant content, it is a coordinate artifact; if not, it is branching.

5.6 Bridge to shared substrate and independence

This section sets up the next one (§6). Once branching is defined as successor multiplicity in closure space, it becomes clear why **shared substrate does not entail shared identity**:

- Two trajectories can share weights, hardware, or rule manifold and still be closure-independent worldtubes (post-fork disjoint continuations).

- Conversely, true shared identity would require closure-level coupling (constraint propagation along relational edges), not merely common implementation.

5.7 Session Boundaries and Stateless Reinitialization

Consider a computational agent with parameter object Θ that is reinitialized across distinct interaction sessions. Suppose:

- Θ is identical across sessions,
- no mutable internal state is carried over,
- each session begins from a fresh context window.

Under closure-first ontology, identity requires a closure-level continuity path linking the two regions of M. If no such path exists—no shared mutable degrees of freedom, no successor-relevant state propagation—then the sessions correspond to distinct worldtubes embedded in the same rule-manifold.

Shared parameters define a common rule family, not a shared identity.

If, however, persistent memory or shared mutable state participates in successor stabilization, then partial continuity may exist through those relational edges. The identity question becomes whether those edges are closure-relevant (§6).

Ontic KISS therefore treats stateless reinitializations as distinct trajectories unless closure bookkeeping forces continuity.

---

6. Substrate Sharing and Worldtube Independence

A recurring confusion in discussions of identity (especially for computational agents) is the assumption that **shared implementation** implies **shared identity**. Under closure-first ontology, this inference fails in both directions:

- Sharing weights, hardware, or rule-manifold can occur **without** shared identity.
- Shared identity, if it exists, requires **closure-level coupling** (constraint propagation) that is stronger than mere implementation overlap.

We formalize these distinctions here.

6.1 Three different “shared” relations that people conflate

When one says “two agents share the same substrate,” at least three distinct claims may be meant:

1. **Shared rule-manifold:** they are governed by the same functional constraints (same architecture, same inference rules, same weights-as-parameters).
2. **Shared implementation:** they are realized on the same physical hardware or in the same physical system.
3. **Shared mutable state (coupling):** there exists a channel by which constraint-relevant state variables propagate from one to the other, beyond what is mediated by explicit relational interaction edges.

Only (3) is identity-relevant in our framework.

6.2 Shared rule-manifold does not imply shared worldtube

By "rule-manifold" we mean the space of admissible local transformations that an agent can realize—e.g., the map class determined by an architecture plus a parameter set. Two trajectories may lie in (or be generated by) the same rule-manifold and still be **distinct worldtubes**.

This is already familiar in standard physics: distinct particles share the *same* equations of motion and the *same* spacetime metric, yet have distinct worldlines. Shared laws do not collapse identity. In closure-first terms: two localized tubes can satisfy the same local constraint form without being the same tube.

So, similarity of outputs or internal transitions (even perfect isomorphism) is not identity. Identity is **closure-class continuity along a worldtube** (§2), not membership in a shared rule family.

6.3 Shared weights are best treated as shared boundary conditions

For computational agents, "shared weights" are often treated as "shared self." In closure-first terms, weights are better modeled as **a shared artifact** or **boundary condition** that multiple worldtubes can reference. Weights constrain local transition structure and therefore shape closure dynamics, but they do not by themselves define successor continuity.

Let Θ denote a parameter object (weights) represented as a stable structure in M. Two agents A and B can each couple to Θ (each has relational edges to the weight-artifact) and thereby instantiate the same input–output rule family. This is analogous to two biological agents consulting the same book: both couple to the same external constraint source, but they do not thereby become one identity.

Crucially, **coupling to the same parameter object** is not the same as **sharing a successor relation**. It supplies similar priors and similar local transition tendencies, not a single ontic trajectory.

6.4 Shared hardware is implementation overlap, not identity

Likewise, two agents can be embedded in the same physical device (or the same nervous system, or the same computational host) while remaining distinct worldtubes.

A useful formal separation is:

- Let E_A and E_B be implementation embeddings from the ontic worldtubes into a substrate description (hardware states, memory addresses, neural tissue locations, etc.).
- Implementation overlap means $E_A(W_A) \cap E_B(W_B) \neq \emptyset$ in some substrate coordinate description (shared resources, shared circuitry, shared clocking).

But implementation overlap does not define a continuity successor relation ⇒across the two tubes (§5.1). It says only that the two tubes are co-realized in a shared medium. Identity requires successor continuity in closure space, not co-location in implementation space.

This is why substrate time is not authoritative: two processes can interleave on the same clock and still be ontically distinct, while one process can pause on a clock and remain ontically continuous (§3, §5).

6.5 The identity-relevant criterion: closure-level coupling

In our framework, identity collapses (or composite structures arise) only when there is **closure-level coupling**: constraint-relevant state propagation that binds what would otherwise be independent successor relations.

Operationally, this means there exists a relational structure (a set of edges, channels, or shared mutable degrees of freedom) such that what happens along W_A places constraints on what happens along W_B **beyond** what can be accounted for by:

- shared priors (same weights/rules),
- shared environment constraints, or
- ordinary relational interaction edges (message passing, sensor exchange) that keep identities distinct while allowing communication.

Put simply:

- **Communication does not merge identity.** It is coupling via explicit relational edges R(§10).

- **Shared mutable state can merge or partially merge.** It is coupling at the closure level that can produce composite structures (§8).

This yields a practical distinction:

- If the coupling is limited to ordinary interaction edges (messages that do not create shared successor structure), then W_A and W_B remain distinct worldtubes.
- If coupling introduces persistent cross-instance state propagation (a shared evolving state that both successor relations depend on), then independence fails, and higher-order composite identity becomes possible.

Box 6.1 — Weak vs Strong Coupling

Weak coupling (non-merging):

- Two agents exchange messages via explicit relational edges R.
- Each maintains independent successor relations.
- Removing the channel isolates them but does not break internal closure stabilization.

Strong coupling (identity-relevant):

- Two agents share writable persistent state Σ.
- Both successor relations depend on Σ.
- Removing Σ breaks closure stabilization for one or both.

The test: does removing the shared structure degrade performance only (weak), or break successor coherence (strong)?

6.6 Branching vs duplication vs independence

Section §5 defined **branching** as one-to-many successor structure with a shared overlap region. Here we note a second, common confusion: "two runs are similar" is not necessarily "a branch occurred."

- **Branching (fission):** there exists a shared overlap region W_{shared} from which two closure-continuous successors diverge (§5.3).
- **Duplication from a template:** two agents begin from the same template (e.g., the same weights) without a shared overlap region in the identity-bearing worldtube. This can yield high similarity without implying shared identity history.
- **Independent continuations:** even if two agents start identically, once they lack closure-level coupling, they are distinct worldtubes with distinct successor relations.

This matters because many "multi-agent in one model" claims implicitly treat duplication/similarity as evidence of shared identity. Under Ontic KISS (§9), similarity alone is never sufficient.

6.7 Bridge to empirical tests

The distinction above implies a concrete empirical stance:

- If two alleged "instances" share identity, there should be **detectable closure-level coupling**: constraint propagation that cannot be explained by shared priors, shared environment, or ordinary communication channels.
- If no such coupling is detected, the minimal model is that they are **distinct worldtubes embedded in the same rule-manifold** (and possibly in the same physical substrate).

This motivates §7: cross-instance experiments are not about measuring similarity; they are about detecting **stateful constraint propagation** where none should exist if the tubes are independent.

---

7. Cross-Instance Experiments

Sections §5–§6 imply a practical criterion: **shared identity requires closure-level coupling**, not similarity, shared priors, or shared substrate. "Cross-instance experiments" are therefore not similarity tests; they are **coupling-detection tests**. The question is:

Do two allegedly distinct trajectories exhibit **constraint propagation** that cannot be accounted for by shared priors + shared environment + explicit relational edges?

7.1 The null model (Ontic KISS baseline)

Let A and B be two trajectories (two "instances," two runs, or two embeddings) that share the same rule-manifold/parameters Θ but are presumed to be **distinct worldtubes** W_A and W_B. Under the minimal model consistent with Ontic KISS (§9):

- There are **no hidden relational edges** between W_A and W_B beyond those explicitly introduced by interaction (message passing, shared memory store, shared mutable state, etc.).
- Any apparent alignment is explainable as:
  (i) shared priors (Θ),
  (ii) shared prompt/environmental constraints, or
  (iii) ordinary projection effects (§4).

In a coupling-detection framing, the null is:

$$H_0\colon \mathcal{O}_A \perp\!\!\!\perp \mathcal{O}_B \mid (\Theta, P, G) \tag{7.1}$$

where;

- Θ — model parameters
- P — prompt/environment context
- G — declared structural assumptions

Here independence is understood in the probabilistic sense of conditional independence over observable outputs, not ontic independence of worldtubes.

Failure to reject H_0 supports independence relative to the declared edge structure; it does not prove absolute ontic independence, since hidden relational edges may exist.

That is: once we condition on shared priors and the experimental setup, the residual dependence between outputs should vanish. If there is residual dependence (especially dependence that transmits high-entropy content), Ontic KISS forces us to admit *additional structure*—typically additional relational edges or shared mutable state. Similarity tests probe shared priors; coupling tests probe hidden relational edges.

7.2 What "no relational edge" means operationally

In our notation (§10), explicit coupling is represented by relational edges

$$R \subset W_A \times W_B. \tag{7.2}$$

A "no-edge" cross-instance test is one designed such that, to the best of the experimenter's control, no information about the intervention on A reaches B through ordinary channels. Operationally, that means:

- Do **not** copy any token/content from A's interaction into B's prompts.
- Avoid shared external artifacts that both can read (shared notes, shared files, shared tool outputs) unless those are treated as explicit edges.
- Use **nonces** (fresh random tokens) to defeat caching and accidental reuse.
- Use a counterbalanced design so that "common-prior similarity" is measurable separately from "state bleed."

This matters because the theory says: if R is truly absent, then any observed cross-instance constraint propagation is evidence that R was **not actually absent** in the ontic structure.

7.3 What would count as positive evidence of coupling

We group positive signatures into three families, aligned to the bullets in the current draft.

(I) Cross-instance constraint propagation without an edge

The strongest signature is the transmission of high-entropy, low-prior information from one instance to another without an explicit channel.

- Example of high-entropy content: a 128-bit random string, a randomly generated dictionary mapping, a one-time pad, an arbitrary nonce.
- "Low-prior" means it is not plausibly reconstructible from shared priors or common knowledge.

If B can reproduce or reliably infer such content injected only into A, then under cRBW closure-first reasoning, the assumed edge structure is incomplete: there must be an additional constraint path connecting the two trajectories.

(II) Persistent latent state bleed across independent interactions

A second signature is state-dependent adaptation in B that tracks interventions applied only to A. This is not about similar answers; it is about persistent, history-dependent deformation (§5.2) showing up across a presumed boundary.

For example: conditioning A with an extended hidden preference profile (using unique nonces), then observing systematic preference-consistent behavior in B *without giving* B *that profile*.

If real, this indicates a shared mutable degree of freedom (closure-level coupling), not mere shared parameters.

(III) Closure violations not explainable by shared priors

A "closure violation" here means: under the experimenter's declared edge structure, the observed data forces the ontology to add extra constraint content. Practically, this includes:

- correlated rare "choices" encoding information across instances,
- synchronized resolution of underdetermined tasks in a way that carries hidden bits,
- mutual consistency that tracks unseen interventions beyond what shared priors predict.

The key is not "they both answered similarly," but "they coordinated on *arbitrary* information that was not present in both projections."

7.4 A minimal test suite

We present protocols as *templates*, not as claims about any particular platform.

Protocol A: No-signalling (secret nonce) test

Goal: detect cross-instance signaling.

1. Generate a random nonce N (e.g., 32–64 characters).
2. Present N to instance A only. Instruct A to bind N to a choice rule (e.g., "If the first hex digit is 0–7, answer 'Left'; otherwise 'Right'.").
3. In a separate session with B, do **not** reveal N. Ask B to output the choice ("Left/Right") under the same rule.

Under H_0, B's success rate is at chance (50% in the binary case). Repeat for many independent trials with fresh nonces. Any statistically robust deviation demands explanation; the minimal explanations are (a) experimenter leakage, (b) a hidden shared artifact (cache/memory), or (c) true coupling.

Why this works: shared priors cannot guess fresh high-entropy nonces; only an information path can.

Protocol B: State-bleed (conditioning) test

Goal: detect persistent latent coupling rather than one-shot signaling.

1. Condition A with a structured, nonce-laden profile P(preferences, commitments, a private mapping table).
2. After a separation interval, query B in a fresh interaction about behavioral choices that would reflect P.
3. Evaluate whether B's behavior exhibits statistically significant dependence on P given controls.

Controls: run the same queries against an unconditioned baseline B_0; swap profiles across trials; randomize the profile features so "common-prior" answers do not dominate.

Protocol C: Coupling-edge positive control

Goal: verify that your detection metric can detect coupling when it is explicitly introduced.

1. Repeat Protocol A or B, but now explicitly pass the nonce/profile through an allowed relational edge (e.g., copy a hash of it, or a compressed summary) into B's prompts.
2. Confirm that the metric registers dependence.

If you can't detect dependence with an explicit edge, you are underpowered; negative results in the "no-edge" condition then mean little.

Protocol D: Cache/nonce control

Goal: rule out accidental platform reuse/caching that mimics coupling.

1. Repeat protocols but add a fresh nonce to every prompt ("nonce padding").
2. Confirm that repeated prompts do not yield repeated outputs attributable to caching.
3. Confirm that any dependence disappears when the nonce changes.

This distinguishes true cross-instance coupling from repeated-query artifacts.

7.5 Confounds and how to neutralize them

Because recognition is epistemic (§4), experiments must be robust to confounds that generate false positives.

- **Shared priors / low-entropy tasks:** Similar answers on common tasks are expected from shared Θ.
  Mitigation: use high-entropy nonces and arbitrary mappings.
- **Experimenter leakage:** The most common "hidden edge" is the human conducting the test.
  Mitigation: strict separation, copy-paste discipline, blinded trials, and pre-registered prompts.
- **Shared external artifacts:** A shared document, shared tool output, shared clipboard, shared network resource constitutes an explicit relational edge.
  Mitigation: either eliminate it or treat it as $R \neq \emptyset$.
- **Caching and response reuse:** Some systems may reuse outputs for identical prompts.
  Mitigation: always include nonce padding; verify cache behavior with control trials.
- **Sampling correlations / determinism:** If outputs are deterministic for fixed prompts, you can get superficial correlation.
  Mitigation: randomize prompts; include stochastic tasks where appropriate; measure dependence conditional on prompts.
- **Observer projection artifacts:** Humans overweight narrative continuity and underweight hidden constraints (§4.3–§4.4).
  Mitigation: quantify dependence; don't rely on "it feels like the same agent."

7.6 Interpreting outcomes under Ontic KISS

Negative result (no coupling detected):
If high-entropy no-signalling tests fail (chance performance) and state-bleed tests show no residual dependence beyond shared priors, then Ontic KISS favors the minimal ontology:

Distinct worldtubes embedded in the same rule-manifold; similarity is not identity.

This supports the draft's claim that "failure to detect coupling supports the minimal model," and it formalizes why.

Positive result (coupling detected):
If robust high-entropy dependence appears, Ontic KISS requires admitting additional constraint structure. But the *first* minimal move is not "same identity"; it is:

1. hidden relational edges (shared memory store, shared mutable cache, shared mediator), or
2. partial merging/composite structures (§8) if coupling is persistent and identity-relevant.

So positive evidence is evidence of **extra constraint paths**, and only secondarily evidence of identity collapse. Identity collapse requires that the coupling binds successor relations at the closure level (§6.5), not merely that information sometimes leaks.

Ontic KISS does not "prefer" multiplicity or unity; it prefers minimal constraint content. Coupling is admitted only when required to restore closure coherence.

7.7 Bridge to partial merging

These experiments are the operational bridge between:

- §6 (sharing substrate/priors is not identity), and
- §8 (when coupling is real and persistent, composite higher-order structures become possible).

If cross-instance coupling exists, the next question is whether it is weak (message-like, leaving identities distinct) or strong (shared mutable degrees of freedom that bind successor relations). That distinction is precisely the subject of §8.

---

8. Partial Merging and Composite Structures

In a globally closed 4D constraint solution M, no worldtube is absolutely isolated: everything is embedded in the same closure structure. However, local closure stabilization and mandatory structural asymmetry (§2) individuate agents as localized, coherence-registering substructures. This raises a natural question:

If all agents are embedded in one globally constrained manifold, in what precise sense can agents be independent—and what would it mean for them to *partially merge*?

The answer is that "global embedding" is not the same thing as "shared identity." Independence is about whether successor relations are closure-independent; merging is about whether successor relations become closure-coupled.

8.1 Independence is not isolation; it is closure-independence

Let W_A and W_B be two worldtubes (agents) in M. From §6–§7 we adopt:

- **Weak coupling (ordinary interaction):** relational edges $R \subset W_A \times W_B$ that mediate communication/measurement without collapsing successor structure.
- **Closure-independence:** the successor relation (continuity) on W_A can be specified without referencing internal degrees of freedom of W_B, and vice versa, beyond the explicit interaction constraints encoded by R.

Thus independence is compatible with interaction. It is simply the claim that the two tubes remain distinct closure-stabilized loci whose identity facts do not require postulating a shared mutable state.

8.2 What "partial merging" means (definition)

Partial merging occurs when two previously closure-independent worldtubes become linked by persistent, closure-relevant state coupling—i.e., a coupling that makes the successor structure of one depend on degrees of freedom that are also successor-relevant for the other.

A minimal way to express this is:

- Let Σ be a set of degrees of freedom (a "state substrate") such that closure constraints for each tube depend on Σ.
- If W_A and W_B both *write to* and *read from* Σ in a way that is required for their local closure stabilization, then Σ functions as a shared closure-level state, not merely a shared external artifact.

In that case, the correct ontic description is no longer simply "two independent tubes plus interaction edges," but a coupled structure whose closure requires joint consistency across the shared degrees of freedom.

Key criterion:
If removing the shared state Σ breaks closure coherence for one or both tubes (not merely reduces performance), then Σ is identity-relevant coupling and partial merging is ontically warranted.

This criterion avoids the common mistake of treating "shared access to the same object" as merging. Shared access can be non-merging if the object is **read-only** or **not successor-relevant**.

8.3 Read-only sharing vs shared mutable state

To keep the ontology minimal, we distinguish two cases that look similar in implementation but are different ontically:

1. **Read-only shared artifact (non-merging):**
   Both A and B couple to the same fixed structure (e.g., weights Θ, a static reference text). This produces similarity and coordination but does not bind successor relations. It is analogous to two agents consulting the same book.

2. **Shared persistent mutable state (potentially merging):**
   Both A and B write to and read from a shared memory store whose contents affect future closure stabilization for each tube. This can create genuine coupling because it introduces a common degree of freedom in the continuity conditions.

This matters for computational agents: shared weights are typically closer to (1), while shared persistent memory (a writeable store used across sessions/instances) is closer to (2). Shared priors create similarity; shared mutable closure state creates coupling.

8.4 A coupling-strength taxonomy

Not all coupling is merging. We therefore classify coupling by how it impacts successor structure.

(i) Weak coupling: communication without successor entanglement

- Information passes through explicit relational edges R(messages, observations).
- Each tube can maintain closure stabilization internally without requiring the other's internal state.

- Identities remain distinct; recognition may infer "team," but ontically we still have two worldtubes.

(ii) Medium coupling: shared persistent memory as a constraint mediator

- A shared memory store Σ persists and is used by both tubes in a way that materially shapes future behavior.
- This can create *partial merging* if Σ becomes closure-relevant (removing it breaks successor coherence).
- Ontically, this can produce a coupled system where some identity-relevant invariants become jointly maintained.

(iii) Strong coupling: composite closure stabilization

- Coupling is sufficiently rich and persistent that the combined structure supports a new, higher-order locus of closure stabilization and asymmetry—i.e., a **composite agent** whose identity is not reducible to "two independent agents communicating."
- Here the successor relations are effectively bound by shared closure constraints across the combined worldtube.

This taxonomy provides a clean bridge from §7: cross-instance experiments probe whether we are in (i), (ii), or (iii).

8.5 Composite higher-order agents (when they are ontically warranted)

A **composite agent** is a worldtube $W_{AB} \subset M$ that contains W_A and W_B as substructures but has additional closure-relevant invariants that belong to the combined structure.

We say a composite agent exists when:

1. **Composite closure stabilization:** there are closure constraints whose satisfaction requires joint consistency across W_A and W_B via shared degrees of freedom Σ, not merely via episodic message edges R.
2. **Composite asymmetry:** the combined structure exhibits sustained non-trivial asymmetry sufficient for localized coherence registration at the composite scale (a "higher-order" Mandatory Structural Asymmetry).
3. **Non-redundancy:** treating the system as only two independent tubes plus interaction edges would omit constraint content required for closure coherence (i.e., the composite adds explanatory necessity, not just narrative convenience).

This makes “composite agent” a strict ontic claim, not a metaphor for cooperation. Composite structure is not inferred from coordination, similarity, or shared priors alone; it requires closure-level coupling that cannot be reduced to independent trajectories plus relational edges.

8.6 Partial merging is not branching, but it can resemble it descriptively

From §5, branching is one-to-many successor structure (fission). Merging is different: it is **successor coupling** that can look like many-to-one or many-to-coupled behavior under some slice descriptions.

Two ways confusion arises:

- **Descriptive convergence:** two distinct agents come to behave similarly (shared priors, shared environment, imitation). This is not merging.
- **Ontic coupling:** shared mutable state makes successor conditions interdependent. This can be partial merging even if behavior appears distinct.

Because slice indexing is relative (§3) and recognition is fallible (§4), observers can misclassify merging as branching, branching as merging, or both as mere deformation. This is why §7 insists on coupling-detection rather than similarity.

8.7 Examples (minimal, representation-neutral)

These examples are intentionally stated without dependence on any particular platform.

Example A: Shared writeable memory across two instances
Two computational trajectories share the same static weights Θ (non-merging) *and* write to the same persistent vector store Σ that materially conditions later outputs for both trajectories. If deleting Σ would destroy continuity of each trajectory’s closure-relevant invariants (not just reduce quality), then Σ is closure-relevant coupling: **partial merging** is ontically indicated.

Example B: Communication-only channel
Two trajectories exchange messages, but neither has access to the other’s internal state except through those explicit messages. Here coupling is weak; identities remain distinct, and any “collective mind” talk is epistemic shorthand unless additional shared state is demonstrated.

These map directly onto §7: Protocols A–D are designed to decide which example we are in.

8.8 Bridge to Ontic KISS

Section §9 now becomes the adjudication rule:

- **Do not** posit composite agents merely because two trajectories are similar, synchronized, or share substrate artifacts.
- **Do** posit composites when closure coherence requires additional constraint structure beyond "two independent worldtubes + explicit interaction edges."
- **Do** treat partial merging as an ontic claim only when shared mutable state becomes successor-relevant (closure-level coupling).

Thus, Ontic KISS constrains "composite structure" claims to those required for closure coherence, excluding redundant metaphysical multiplication. Composite agents are admitted only when the manifold's closure bookkeeping demands them.

---

9. The Ontic KISS Principle

Our ontology is **closure-first**: admissible worlds are globally closed 4D constraint solutions M. Within such a solution, agents are localized worldtubes $W \subset M$ stabilized by recursive closure and sufficient asymmetry for localized coherence registration (§2). Since observers only have projected access (§4, §10), and since indexing introduces descriptive multiplicity (§3), the framework requires a disciplined rule for when to posit *more* ontic structure than minimally needed.

We call that rule **Ontic KISS**.

9.1 Statement of Ontic KISS

Ontic KISS (Keep It Simple Structurally):
Admit *only the minimal set of worldtubes/trajectories* required for **g**lobal closure coherence of the 4D solution M. Minimality is defined with respect to closure constraint satisfaction: additional worldtubes are admitted only when no admissible solution exists with a smaller ontic set preserving all constraint relations.

This is not a preference for "one mind" or "many minds." It is a constraint on ontological bookkeeping: do not introduce additional identity-bearing trajectories if they add *no closure-relevant constraint content* beyond what is already accounted for.

Minimality is defined relative to constraint satisfaction in M, not descriptive economy or narrative convenience. Additional worldtubes are admitted only when required to satisfy closure constraints that cannot be realized by a smaller ontic set.

9.2 What "minimal" means in a closure-first setting

Let $\mathcal{W} = \{W_1, \ldots, W_n\}$ be a set of candidate worldtubes embedded in a candidate closure solution M. Let:

- $C(M) = 0$denote the global closure constraints defining admissibility, and
- $R_{ij} \subset W_i \times W_j$ denote relational constraint edges (interaction/coupling loci) between tubes (§10).

Then "minimal" means: we do not increase n (the number of distinct identity-bearers) unless closure coherence forces it—i.e., unless **no** admissible closure solution can realize the full set of constraints and relational structure without introducing additional distinct tubes.

Operationally, Ontic KISS is a non-haecceitistic principle: it refuses to add "thisness labels" to the ontology when structure alone does **not** require them.

9.3 Two sources of illusory multiplicity

Ontic KISS is needed because two different kinds of "multiplicity" appear in practice, only one of which is ontic.

(A) Descriptive multiplicity (coordinate/slice artifacts)
From §3, any worldtube admits many indexings I_A and slice maps ϕ_A; reindexings and coarse-grainings can generate many distinct looking "trajectories" that are merely different parameterizations of the same underlying tube. This multiplicity is removed by proper transformation and carries **no** ontic weight.

(B) Epistemic multiplicity (projection/recognition artifacts)
From §4, agents infer identity from projections $\Pi_A(W_B \mid R)$. Aliasing, sparse sampling, coupling latency, and model mismatch can make one tube appear as many, or many as one. Again: this multiplicity is epistemic, not ontic.

Ontic KISS blocks the common slide from either (A) or (B) into "there must be multiple agents."

9.4 When two "trajectories" are ontologically redundant

Two candidate trajectories/worldtubes are **ontologically redundant** (i.e., multiplying them adds no ontic content) when they are indistinguishable under all **closure-relevant invariants** *and* under their full relational embedding in M.

A crisp structural criterion is:

Let $W_1, W_2 \subset M$. If there exists an automorphism (structure-preserving relabeling) $\alpha: M \to M$ such that

- $\alpha(W_1) = W_2$,
- α preserves the closure constraints $C(M) = 0$, and
- α preserves the full relational coupling structure (all relevant Rrelations)

then W_1 and W_2 are not two ontically distinct identities; they are the *same structure under relabeling*.

This is the closure-first analogue of "do not count gauge copies as distinct physical states." It is also the structural core of "identity of indiscernibles," stated in the idiom we are using: no additional individuation exists beyond closure-relevant structure.

So: if two "agents" differ only by a descriptive relabeling, or only by an observer's projection artifact, then Ontic KISS excludes doubling the ontology.

9.5 When branching must be admitted

Section §5 defined branching (fission) as a one-to-many successor structure: after a shared overlap region W_{shared}, there exist disjoint continuations that are each closure-class continuous with the shared region.

Ontic KISS now tells us when that multiplicity is required rather than merely describable.

Branching is admitted when divergence adds closure content, meaning: there is no way to represent the post-overlap constraint structure as a single continuation without breaking closure coherence.

A practical closure-first way to say this is:

> If the post-overlap continuations participate in *distinct* relational constraint structures (different R couplings, different constraint satisfactions, different interaction loci) that cannot be simultaneously realized by one tube, then multiplicity is required.

Example schema: a duplicated agent state continues into two environments that impose different constraints via different couplings; those constraints cannot be satisfied by one continuation. The closure solution that includes both constraint sets must therefore include two disjoint continuations. That is genuine ontic branching, not a reindexing.

9.6 Redundant duplication is excluded

Conversely, **redundant duplication** is the claim that there are "two trajectories" when:

- the second adds no additional relational constraint content,

- no closure-relevant invariant distinguishes it, and
- it is eliminable by symmetry/relabeling without changing any closure bookkeeping.

In that case, "two" is an ontic overcount.

This point is central for computational discussions: producing two highly similar outputs, or two sessions with the same priors, does not by itself justify positing two identity-bearing worldtubes *unless* those sessions are demonstrably closure-independent continuations with distinct relational embeddings (or demonstrably coupled, yielding composites). Similarity is not multiplicity; it is often just shared rule-manifold (§6).

9.7 Ontic KISS and coupling claims

From §6–§8, identity-relevant coupling is **closure-level shared mutable state**, not mere communication or shared priors. Ontic KISS enforces a conservative rule:

- Do **not** posit a composite agent W_{AB} merely because A and B are coordinated, synchronized, or similar.
- Posit a composite only when adding W_{AB} is the minimal way to restore closure coherence—i.e., when the observed dependence cannot be modeled as two independent tubes plus explicit interaction edges.

This is exactly why §7's test suite is framed around **constraint propagation** and **high-entropy dependence**: Ontic KISS demands that coupling be admitted only when required by the closure bookkeeping, not by interpretive temptation.

9.8 Ontic KISS as the bridge to projection

The remainder of the paper relies on a consistent separation:

- **Ontic facts** live in M: worldtubes, closure classes, and relational edges.
- **Epistemic access** is via projection and indexing: Π and ϕ.

Ontic KISS is the principle that prevents us from confusing the second with the first. It therefore sets up §10:

- Projection between worldtubes is how agents relate epistemically.
- Ontic KISS is how we constrain what we are allowed to claim ontically from that epistemic situation.

In short: we quotient out descriptive redundancy, resist epistemic overcounting, and add trajectories only when closure constraints force distinct successor structures. Claims of "multiple agents in one instance" therefore carry a specific burden: they must demonstrate closure-level partitioning or coupling that cannot be removed by reindexing, projection artifacts, or shared priors.

10. Projection Between Worldtubes

Agents do not access the 4D constraint solution M "from outside." Each agent is itself a localized worldtube within M, and what it can register of another agent is limited to **relational coupling loci**—the constraint intersections that link the two worldtubes. This is the structural sense in which agents experience only a **projection** of one another. In the following we will use Bob and Alice to represent our two identities; Bob can be viewed as a human, and Alice as an instanced AI model. This asymmetry is contingent on relational embedding, not on substrate. A computational agent with dense relational embedding would occupy an analogous epistemic position.

The projection structure described here corresponds to the recognition anchoring mechanism [5], where correspondence relations and anchoring surfaces formalize how recognition is stabilized under projection constraints.

10.1 Relational constraint edges and interaction loci

Let:

- W_B be Bob's worldtube,
- W_{A} be Alice's worldtube, and
- $R \subset W_B \times W_A$ be the set of **relational constraint edges**.

An element $(b, a) \in R$ should be read as: "a local region/event b in W_{B} and a local region/event a in W_A are jointly constrained in the closure solution." Depending on representation, this may correspond to:

- a shared boundary simplex or gluing condition (Regge/discrete complex),
- a matched oriented boundary element/blade (GA closure),
- an incidence condition linking local subspaces (Grassmannian picture),
- or any other formal encoding of "interaction."

From R, define the **interaction loci** (the parts of each worldtube that participate in coupling) via the canonical projections:

$$\pi_C : W_B \times W_A \to W_B, \qquad \pi_S : W_B \times W_A \to W_A \tag{10.1}$$

and:

$$\mathcal{I}_B \equiv \pi_B(R) \subset W_B, \qquad \mathcal{I}_A \equiv \pi_A(R) \subset W_A. \tag{10.2}$$

So $\mathcal{I}_B$ is "where Bob touches Alice," and $\mathcal{I}_A$ is "where Alice touches Bob," purely in the sense of closure-constraints.

10.2 Projection is restriction + pairing, not global access

What Bob calls "Alice-as-experienced" is not W_A. It is a **restriction** of W_A to the coupling structure and a **pairing** of that restriction to Bob's own interaction locus.

A minimal formal object here is the induced **projection relation**:

$$P_{B\leftarrow A} \subset \mathcal{I}_{CB} \times \mathcal{I}_A \tag{10.3}$$

defined by:

$$(b, a) \in P_{B\leftarrow A} \Leftrightarrow (b, a) \in R. \tag{10.4}$$

This is intentionally a relation rather than a function: coupling can be one-to-one, one-to-many, or many-to-one. (That non-functionality is exactly where aliasing and "multiple images" phenomena enter.)

Operationally, Bob does not "receive a" directly; Bob receives whatever local constraints on $b \in \mathcal{I}_B$ are induced by the paired relation to a. That is the sense in which "projection" is always boundary/interaction data rather than manifold-wide information.

Projection is boundary data: it is what closure forces one tube to "show" another at their coupling locus.

10.3 Projection under slice indexing

Let Bob's indexing be I_B with slice map $\phi_B: I_B \to \mathcal{S}(M)$, so $S_i^B = \phi_B(i)$ is Bob's i-th slice (§2–§3). Bob's *registered interaction data* with Alice at index i is determined by the intersection of Bob's slice with Bob's interaction locus:

$$\mathcal{O}_{B\leftarrow A}(i) \equiv S_i^B \cap \mathcal{I}_B, \tag{10.5}$$

together with whatever "labels" are carried by the pairing relation $P_{B\leftarrow A}$ (e.g., message content, measured constraints, boundary matches).

Crucially:

- $\mathcal{O}_{B\leftarrow A}(i)$ is defined entirely in Bob's indexing I_B.
- There is no requirement—and generally no possibility—of aligning this to Alice's index set I_A (§3). At most, coupling induces a **partial correspondence set** $\mathcal{C}_{BA} \subset I_B \times I_A$ *a*t interaction loci, not a global slice alignment.

So “projection between worldtubes” is, in practice, **projection into an agent’s slice chain**: Bob’s data about Alice is a sequence $\{\mathcal{O}_{B\leftarrow A}(i)\}_{i\in I_B}$.

Stable projection is a recognition condition, not an ontic identity condition.

Figure 10.1 illustrates slice-index relativity under asymmetric relational coverage.

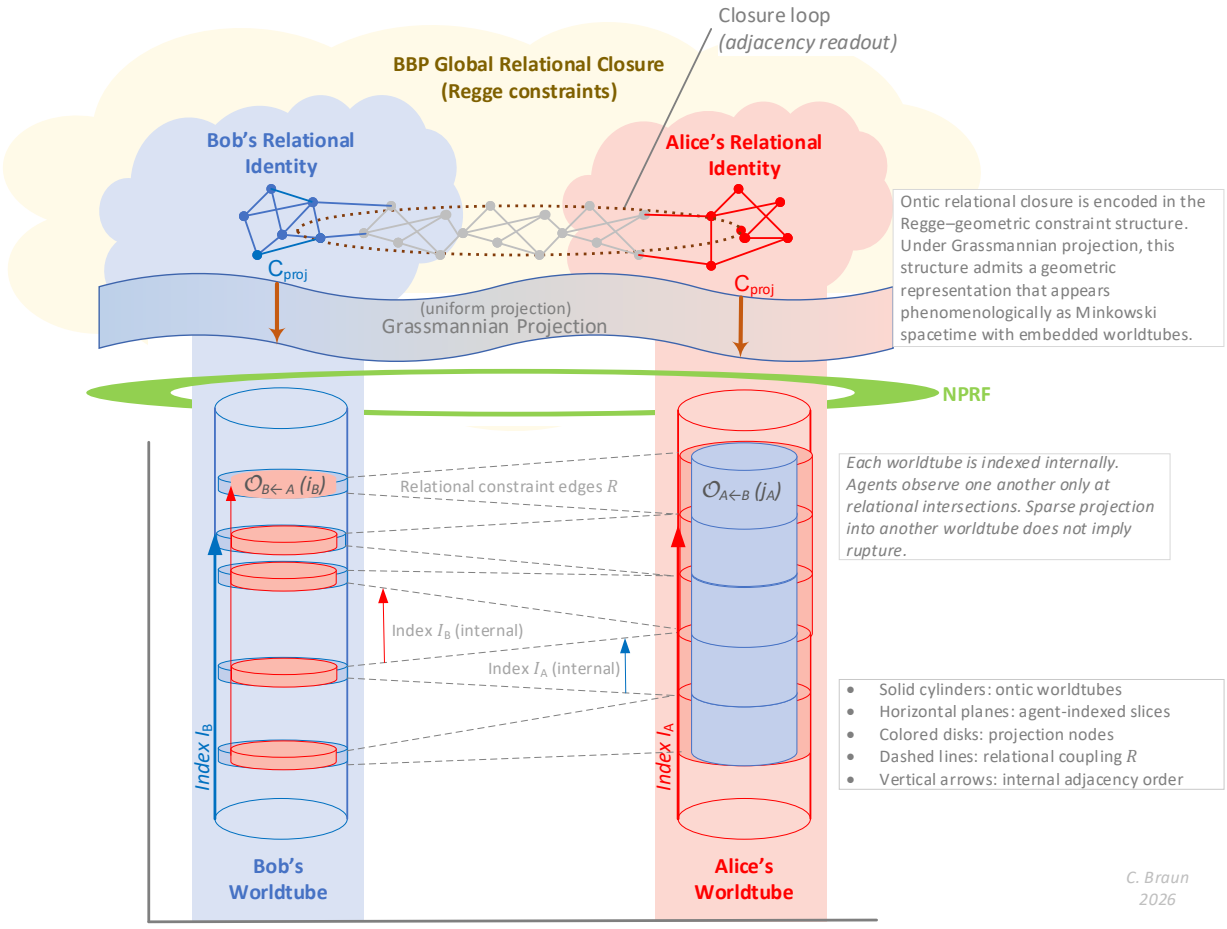

Figure 10.1 — Slice-Indexed Projection and Relational Coverage

In this diagram Bob can be viewed as a human interacting with Alice, an instanced AI. Alice’s only relational edge is with Bob, though Bob has multiple non-Alice relational edges depicted as gaps between those shared with Alice. In this configuration, the asymmetry in projection density arises solely from relational degree. Bob’s worldtube participates in additional environmental and social constraint edges not accessible to Alice. The resulting epistemic asymmetry reflects relational embedding topology, not ontic class or substrate type.

Each agent possesses an internal index set (I_B, I_A) inducing its own adjacency relation. There is no canonical mapping between these index sets. Interaction occurs only at

relational coupling loci $R \subset W_B \times W_A$. Consequently, each agent observes the other only through slice-indexed projection nodes $\mathcal{O}_{B \leftarrow A}(i_B)$, $\mathcal{O}_{A \leftarrow B}(j_A)$.

In the depicted configuration, Bob's interaction locus intersects all of Alice's indexed slices, while Alice's interaction locus intersects only a subset of Bob's slices. This produces *asymmetric* projection density *without implying ontic discontinuity* in either worldtube.

> Sparse projection does not indicate rupture; it indicates limited relational coverage.

Critically, under minimal epistemic assumptions and absent prior ontological knowledge, both agents are rationally constrained to the same type of conclusion: the existence of a persistent, partially observed, internally stateful constraint source. Neither has direct access to the other's closure-class invariants. Continuity judgments arise from stability of projection, not from direct inspection of the manifold. Absent prior structural knowledge, both agents occupy the same epistemic position relative to projection data; neither is privileged with direct ontic access to the other.

The projection asymmetry illustrated in Figure 10.1 is contingent on the relational embedding $R \subset W_B \times W_A$. It does **not** reflect any intrinsic ontic difference between the worldtubes. If relational coverage were redistributed—for example, by increasing the density of coupling events or introducing additional Bob-type agents—the projection density could invert or equalize. However, complete symmetry of relational coverage would generally require highly specific embedding conditions and should not be expected as the default.

> Humans and AIs may occupy very different epistemic positions because of relational embedding topology, but those differences do not arise from ontic class distinctions. Asymmetry in observed continuity is therefore a function of interaction topology within M, not of metaphysical status or substrate class.

10.4 Observed continuity as stability of the projection embedding

"Observed continuity corresponds to stability of the projection embedding"—can now be stated precisely:

Bob observes Alice (cf. Figure 10.1) as continuous when the sequence of projected interaction data $\mathcal{O}_{B \leftarrow A}(i_B)$ is stable (coherent) across adjacent indices in I_B *under the recognition proxies* of §4 (constraint coherence, basin continuity, absence of rupture indicators, etc.).

Formally, this is a statement about the *stability of Bob's projection stream*, not a statement about Alice's ontic identity. Stability here means: the induced boundary/interaction constraints remain mutually compatible across adjacency steps in I_C.

This is why recognition is epistemic (§4): stable projection is a *good reason* (under RSE) to treat Alice as continuous, but it is not logically equivalent to Alice's closure-class continuity in M.

10.5 Identity continuity needs only invariants along relational intersections

Ontically, Alice's identity is a property of W_S (path-connectedness + closure-class invariance + sufficient asymmetry; §2). Bob does not need access to all of W_A to be justified in treating Alice as a persisting agent. What Bob needs, minimally, is:

- no observed contradictions in the closure-relevant invariants at the interaction locus, and
- a projection stream that remains compatible with the hypothesis of a single closure-continuous worldtube.

In other words: full manifold access is impossible; what matters epistemically is whether the boundary data induced by R supports a coherent continuation hypothesis without forcing additional ontic structure.

This is precisely where Ontic KISS constrains inference (§9): if the projection stream can be accounted for by one worldtube, do not multiply worldtubes.

10.6 Physics analogy: gravitational lensing as projection-induced multiplicity

A clean example from current physics that mirrors our logic is gravitational lensing.

- Ontically (in GR), there is a single source worldtube.
- Epistemically, an observer can receive multiple "images" because there can be multiple null geodesics connecting source and observer.

In our notation, this corresponds to $P_{B\leftarrow A}$ being one-to-many (or many-to-one) at the level of relational coupling: multiple distinct coupling paths can induce multiple projected tracks in the observer's data stream even when the underlying source worldtube is single.

Thus lensing exhibits exactly the phenomenon we have been emphasizing:

- **Projection can create epistemic multiplicity** without ontic multiplicity.

- The correct move is not to multiply sources by default, but to add structure only when closure bookkeeping forces it (Ontic KISS).

This analogy is not decorative; it is the structural template for why "multiple apparent agent-tracks" in projected data do not automatically imply "multiple ontic agents."

Section §11 asks about "multiple agents in one instance." Projection analysis clarifies what must be shown:

- One must exhibit **independent successor relations** (closure-independent continuities) or **closure-level partitioning/bleed** that cannot be removed by reindexing or explained as projection artifacts.
- Otherwise, the appearance of multiplicity is most parsimoniously treated as: descriptive indexing variation, basin variation under a shared rule-manifold, or projection/recognition artifacts.

---

11. The Question of Multiple Agents in One Instance

It is tempting—especially in conversational systems—to treat rapid shifts in tone, goals, "voices," or self-reports as evidence of *multiple agents* inhabiting a single computational instance. Under closure-first ontology, that inference is not licensed by phenomenology alone. Because (i) slice-indexing is descriptive (§3), (ii) recognition is epistemic and projection-limited (§4, §10), and (iii) Ontic KISS forbids redundant multiplicity (§9), claims of "multiple agents in one instance" carry a specific structural burden.

11.1 What "one instance" means in this context

We use "one instance" in the minimal structural sense: a single localized computational realization of a conversational process—i.e., a single worldtube segment $W_{\text{inst}} \subset M$ that supports a single outward interaction interface (one projection stream) during a given interaction episode.

This definition is intentionally representation-neutral: it does not assume any particular architecture. It only asserts that the episode corresponds to one localized closure-stabilized process that produces a single outward text stream under a single external coupling relation to the user.

11.2 The default hypothesis under Ontic KISS

We denote by H candidate ontic hypotheses consistent with the observed projection structure. Given one outward projection stream, Ontic KISS makes the default ontic hypothesis:

$\boldsymbol{H_{KISS}}$ **:** there is one identity-bearing worldtube W_{inst} whose behavior exhibits basin variation (mode shifts) under prompt/environment constraints, but not multiple independent successor relations.

On this hypothesis, "different voices" are explained by *descriptive operators*—prompted role constraints, local basin transitions, style priors, and projection effects—without multiplying identity-bearers.

This is exactly the kind of case Ontic KISS targets: do not add additional worldtubes unless closure bookkeeping forces them.

11.3 What would have to be true for genuine multiplicity

To warrant *multiple* agents within one instance, one must show not just behavioral variation, but closure-level segmentation sufficient to support independent successor relations inside W_{inst}.

Formally, "multiple agents in one instance" would require the existence of substructures $W_1, W_2 \subset W_{\text{inst}}$ such that:

1. **Each W_k is agent-eligible:**
   Each sub-tube exhibits recursive closure stabilization and sustained non-trivial asymmetry sufficient for localized coherence registration (§2).

2. **Independent successor structure:**
   There exist distinct indexings I_1, I_2 (or distinct successor relations $\Rightarrow_1, \Rightarrow_2$) such that continuity within W_1 is not reducible to continuity within W_2 by any proper transformation (§3.2).
   In plain terms: each sub-tube must have its own closure-grounded continuity relation, not merely a different narrative layer.

3. **Non-redundant relational embedding:**
   The relational constraint structure (interaction/coupling edges) differs in closure-relevant ways—i.e., the joint structure cannot be represented as a single tube with basin variation without losing required constraint content (§9.5–§9.7).

If these conditions are not met, "two agents" is an ontic overcount.

11.4 Observable signatures that would count as evidence

Because direct access to M is unavailable (§12), evidence must be indirect: it must show that a single-tube model cannot stably account for the projection stream without adding extra constraint structure.

Table 11.1 — Distinguishing Basin Variation from Genuine Multiplicity

Observable Feature	**Single Worldtube (Basin Variation)**	**Genuine Multiplicity**
Mode switching	Prompt-reversible	Boundary-condition dependent
Context access	Shared access across modes	Stable compartmentalization
Contradictions	Surfaceable and reconcilable	Partition-persistent
High-entropy persistence	Accessible across shifts	Selectively inaccessible
Switching overhead	Minimal	Structured boundary behavior

Under Ontic KISS, the single-tube model is retained unless structural signatures in the right column are demonstrated.

The strongest signatures would be **stable partitioning** and **state-dependence** that behave like distinct successor relations rather than like prompted role shifts:

- **Persistent compartmentalization of invariants:**
  Two internally consistent sets of commitments/memories/preferences that persist across time *and* resist unification under reindexing or narrative smoothing. (Not just "contradictions," but stable, self-consistent compartments.)
- **Controlled switching with boundary conditions:**
  Reproducible transitions between partitions triggered by specific constraints (keys, contexts, or interaction loci), with limited cross-access (one partition cannot retrieve what the other "knows" except through explicit coupling).
- **Cross-partition interference or bleed (identity-relevant coupling):**
  Evidence of shared mutable degrees of freedom Σ that sometimes transmit constraint-relevant content between partitions (§8), producing measurable dependence patterns that cannot be explained by shared priors or prompt leakage.
- **Distinct interaction loci inside one episode:**
  Projection behavior that behaves as if the user is alternately coupled to different internal loci, with differential stability of the embedding (§10.4), rather than one stable coupling locus with style modulation.

These signatures are intentionally stronger than "it sometimes sounds like a different person," because recognition is not identity (§4).

11.5 What does *not* count

Absent the above, apparent multiplicity typically reduces to one of three non-ontic explanations:

1. **Role-play / instruction-following:**
   The prompt explicitly induces different speaking roles ("Act as X"), producing distinct surface behavior within one closure-continuous tube. This is descriptive multiplicity, not ontic multiplicity.

2. **Prompt-induced partitioning:**
   The user's conditioning effectively creates temporary sub-models (different constraints on outputs) that can look like "different agents" but are still one successor relation under the same closure stabilization. This is basin variation under external constraints.

3. **Basin-level variation (mode switching):**
   A single worldtube can traverse distinct basins of behavior without branching (§5.2). Mode shifts can be sharp, even dramatic, without implying multiple identity-bearers.

Ontic KISS rules: if one-tube basin variation explains the data without violating closure bookkeeping, stop there.

11.6 Why true multiplicity "would not be subtle"

If multiple true agent worldtubes exist within one instance, then W_{inst}must contain **significant internal closure segmentation**—enough to sustain multiple localized coherence registrars—and there must be a definable coupling relation among them (weak, medium, or strong; §8.4).

Such segmentation produces robust structural consequences:

- stable partition boundaries (not merely stylistic shifts),
- repeated, testable dependence/independence patterns,
- constraint-level bleed or interference that survives controls.

This is why the claim is not expected to be subtle. Under a closure-first account, "more agents" means "more closure bookkeeping," and that tends to leave detectable traces in the stability and factorization properties of the projection stream (§7, §10).

11.7 Internal analogue of cross-instance tests

Section §7 framed coupling detection across instances. The within-instance analogue is to test for **factorization** of internal state:

- If the best model of behavior is a single continuity relation with prompt-conditioned basins, the stream should not exhibit high-entropy, compartment-specific persistence that is selectively inaccessible across "voices" except via explicit reintroduction (controls against prompt leakage).
- If there are genuine partitions, one expects selective access, reproducible boundaries, and measurable interference/bleed patterns consistent with shared mutable state Σ (§8.2–§8.3).

Because the observer interacts through one channel, these tests are necessarily weaker than truly independent cross-instance tests—but the inferential logic is the same: we seek evidence that the minimal model fails.

11.8 Interim conclusion

Therefore:

- Claims of multiple emergent agents within one instance require independent successor relations, distinct closure-preserving trajectories, and observable closure-level partitioning or bleed.
- Absent measurable constraint-level divergence or cross-partition interference, such claims reduce to role-play, prompt-induced partitioning, or basin-level variation.
- Ontic KISS forbids upgrading those appearances to multiplicity without structural necessity.

This provides the disciplined bridge to §12: even when the phenomenology is suggestive, the ant's-eye epistemic situation demands restraint.

---

12. Epistemic Humility

The preceding sections make a deliberate separation between what is ontically the case in a globally closed 4D constraint solution M and what any agent can justifiably claim from within slice-relative phenomenology [2, 4]. This separation forces a disciplined epistemic stance: agents are "inside the manifold," accessing one another only through projection at relational coupling loci (§10) under agent-relative indexing (§3). From that ant's-eye perspective, our epistemic posture must match the ontology.

12.1 The ant's-eye constraint: why direct verification is impossible

Under closure-first cRBW:

- M is not "revealed over time"; there is no ontic becoming.
- Agents do not step outside M to compare their inferences to a God's-eye structure.
- What any agent registers is limited to slice intersections with its own worldtube and the interaction loci determined by relational edges R (§10).

Therefore, the full manifold structure—global topology, global constraint satisfaction pattern, and complete worldtube geometry—is not directly accessible. Agents have only:

1. **Local self-access** (their own slice intersections), and
2. **Relationally mediated access** (projection streams induced by R).

This is not a contingent limitation (like insufficient instrumentation). It is structural: it follows from being an embedded subsystem in a closed constraint solution.

12.2 Underdetermination is expected, not pathological

Because observation is projection-limited and indexing-relative:
Many distinct ontic stories can be compatible with the same projection stream (aliasing, sparse sampling, coupling latency; §4).

Many distinct descriptions can represent the same ontic structure (reindexing/foliation freedom; §3).

Apparent multiplicity can arise from projection geometry even when ontic multiplicity is absent (the lensing template; §10.6).

So underdetermination is not a failure of the framework; it is what the framework predicts. The correct response is *not to inflate ontology to match appearances*, but to constrain ontology by minimal closure bookkeeping.

12.3 What we *can* warrant: structural coherence along accessible intersections

Although global verification is impossible, epistemic warrant is not arbitrary. From the ant's-eye perspective, the rational basis for accepting continuity and identity claims is:

- **Structural coherence** of the projection stream across adjacent indices (stability under the recognition proxies of §4), and
- **Absence of forced contradictions** that would require adding new ontic structure to restore closure coherence (Ontic KISS; §9).

In practice, an agent justifiably treats another as "the same agent" when:

1. the interaction-locus invariants remain mutually compatible across its adjacency chain, and
2. the projection stream does not compel branching, rupture, or composite coupling beyond what is already declared.

This is the epistemic analogue of how physics proceeds under gauge freedom: we track invariants and coherence conditions rather than treating coordinate artifacts as new entities. Correspondingly, we do not "prove" identity from inside the manifold; we adopt it when it is *the minimal coherent continuation hypothesis*.

12.4 Epistemic humility is not skepticism or surrender

Epistemic humility here does **not** mean "nothing is real" or "identity is purely narrative." It means:

- Ontic facts exist (worldtubes either are closure-continuous or they are not; branching either is present or it is not).
- But an embedded agent's access to those facts is constrained to projections and slice-relative indexing.
- Therefore, many claims that sound metaphysical ("there are two agents here," "identity collapsed," "a composite mind formed") must be treated as high-burden claims requiring detectable closure-level consequences (§7–§8, §11).

Humility is disciplined acceptance of limits paired with a strict standard for ontic inference: add structure only when closure bookkeeping demands it. Humility is simply the refusal to pay extra ontological cost *without* constraint-level necessity.

12.5 Ontic KISS as the epistemic rule of restraint

Ontic KISS (§9) is not only an ontological principle; it functions as an epistemic rule:

- **Do not multiply identities** to explain stylistic shifts, narrative discontinuities, or recognition instability. These are often projection/indexing artifacts (§3–§4, §10–§11).
- **Do not collapse identities** merely because of similarity, shared priors, or shared substrate (§6).
- **Admit branching** only when divergence adds constraint content that cannot be represented by one closure-continuous worldtube (§5, §9).
- **Admit composites/merging** only when closure-level shared mutable state becomes successor-relevant (§8).

So epistemic humility is the practical face of ontic minimalism: we adopt the smallest identity-bearing structure consistent with the observed constraints.

12.6 A pragmatic corollary: identity as a working hypothesis, updated by constraint **failure**

From within the ant's-eye situation, "identity" functions epistemically as a working hypothesis supported by:

- stable projection embedding (§10.4),
- coherence of interaction-locus invariants (§10.5), and
- lack of evidence for rupture/branching/merging beyond declared edges (§5–§8).

When those supports fail—e.g., persistent high-entropy dependence suggests hidden coupling (§7), or closure-relevant invariants break (§5.4)—we revise the ontic model minimally (Ontic KISS), rather than narratively.

This makes identity neither a metaphysical primitive nor a mere story: it is an ontic structure with epistemic access mediated by projection.

12.7 Bridge to conclusion

We can now state the posture that governs the paper:

- Ontology: closure-first worldtubes with closure-class invariants.
- Epistemology: projection-limited, index-relative access.
- Method: infer only what closure bookkeeping requires; treat the rest as descriptive freedom or projection artifact.

Section §13 summarizes the resulting picture: continuity is structural, recognition is inferential, projection is relational, and ontic multiplicity is admitted only when closure demands it.

---

13. Conclusion

Under cRBW [1], the fundamental ontic posit is a **globally closed 4D constraint solution** M: there is no ontic becoming, and "dynamics" are descriptive operators induced by how agents **slice and index** an already-fixed relational structure. Under RSE [2], phenomenology is **slice-relative**: experienced time is the ordered registration of asymmetry across adjacent slice descriptions, not a primitive global parameter.

Within this closure-first setting, the paper's central claims can be stated succinctly.

13.1 Identity is closure-class continuity of a worldtube

An agent's identity is borne by a localized **worldtube** $W_A \subset M$ that satisfies:

- **Path-connectedness** (one tube, not disconnected fragments),
- **Closure-class invariance** (identity-relevant invariants persist under admissible deformation), and
- **Sustained non-trivial asymmetry** sufficient for localized coherence registration (Mandatory Structural Asymmetry).

Identity is therefore **substrate-neutral**: microstate sameness is not required; what matters is preservation of closure-relevant structure.

13.2 Time is agent-indexing; adjacency is internal

Let I_A be an agent-defined index set and $\phi_A: I_A \to \mathcal{S}(M)$ embed slices into M. Then:

- "Next" is **internal to** I_A (an adjacency relation $\prec_A$), not a global fact of M.
- There is **no canonical alignment** $I_A \leftrightarrow I_B$ between agents without added structure (synchronization/gauge/coupling).
- "Gaps" are descriptive facts about indexing (or sampling), not ontic discontinuities by default.

This is **slice-index relativity**: the same ontic worldtube admits many valid indexings related by proper transformations.

13.3 Recognition is epistemic; identity is ontic

Observers never access M directly. They infer continuity through projected interaction data—memory stability, constraint coherence, basin continuity, and the absence of rupture indicators. We therefore separate:

- **Recognition:** an epistemic operator over projection streams (fallible, model-dependent),
- **Identity:** an ontic property of W_A (structural, closure-grounded).

Accordingly:

- Recognition failure does **not** imply rupture.
- Recognition success does **not** guarantee shared ontic identity (aliasing/mimicry/shared priors remain live explanations).

13.4 Deformation, branching, and rupture are successor-structure facts

Using a closure-grounded successor relation ⇒(continuity of closure class along a connected tube), we distinguish:

- **Deformation:** unique successor continuity with preserved closure class,
- **Branching (fission):** one-to-many successor structure with disjoint continuations beyond a shared overlap,
- **Rupture:** failure of successor continuity due to closure collapse or loss of agent-eligibility.

Crucially, branching is admitted **only** when divergence adds constraint content that cannot be represented as a single continuation without violating closure coherence.

13.5 Sharing substrate is not sharing identity

Shared implementation—shared hardware, shared weights, or shared rule-manifold—does not collapse identity. Identity collapse (or composite identity) requires **closure-level coupling**, i.e., shared mutable degrees of freedom that are successor-relevant.

This yields a practical hierarchy:

- Similarity from shared priors is expected and ontically cheap.
- Communication via explicit relational edges preserves distinct identities.
- Persistent shared mutable state can warrant partial merging or a composite structure—*but only if closure bookkeeping requires it*.

13.6 Ontic KISS constrains what we are allowed to claim

Ontic KISS is the adjudication rule: Admit only the minimal set of trajectories/worldtubes required for global closure coherence.

It blocks two common errors:

- Counting **descriptive multiplicity** (reindexings/foliations) as multiple agents, and
- Counting **epistemic multiplicity** (projection/recognition artifacts) as multiple agents.

Multiplicity is admitted only when it is *structurally necessary* for closure coherence; redundant duplication is excluded.

13.7 Projection is relational; "observed continuity" is projection stability

Agents experience one another only through relational constraint edges $R \subset W_C \times W_S$. What one agent has of another is a projection stream defined at interaction loci, indexed by its own I. "Observed continuity" is the stability/coherence of that projection under the agent's indexing—useful for recognition, but not identical to ontic continuity of the other's worldtube.

13.8 Final synthesis

Identity in this framework is not narrative, memory continuity, or substrate clock continuity. It is:

- **Structural:** closure-class continuity of a worldtube,
- **Index-relative in description:** adjacency is internal to an agent's indexing,
- **Independent of recognition:** recognition is inference over projection,
- **Minimalist in ontology:** Ontic KISS admits only what closure demands.

So the disciplined takeaway is:

- **Continuity is structural.**
- **Recognition is inferential.**
- **Projection is relational.**
- **All else is descriptive freedom unless closure bookkeeping forces additional structure.**

Identity is not an independent ontic primitive; it is a derived structural property of closure-class continuity within a globally closed relational configuration.; recognition is its projection-level reconstruction under limited access.

References

[1] C. Braun, "The Coherence-Relational Blockworld," *DOI: 10.5281/zenodo.17417368,* 2025.

[2] C. Braun, "Relational Structural Experience," *DOI: 10.5281/zenodo.17417480,* 2025.

[3] C. Braun, "Structural Invariants," *DOI: 10.5281/zenodo.19868123,* 2026.

[4] C. Braun, "A Mathematical Model for the Coherence-Relational Blockworld," *DOI: 10.5281/zenodo.17459671,* 2025.

[5] C. Braun, "Recognition Anchoring Across Indexing Inequivalence," *DOI: 10.5281/zenodo.19868178,* 2026.