

Distinction Holonomy: A Gauge Theory of Persistence, Repair, and Continuation

with a Worked Multiscale Instantiation in Spherepop and TARTAN

Flyxion
Independent Researcher

August 2026

Abstract

A persistent object is usually treated as something that remains identical while time passes. This paper proposes the opposite starting point: persistence is not the preservation of an object but the preservation of a family of admissible distinctions under transformation. An entity persists when successive states remain mutually continuable, even when their material realization, internal configuration, or descriptive representation changes. From this perspective, memory is not a storage mechanism, repair is not merely restoration, and identity is not primitive. All three are manifestations of a single mathematical structure: the transport of distinctions through a changing state space.

Part I formulates this structure as a gauge theory over a bundle of admissible continuations, developed with attention to where the gauge-theoretic language is rigorously earned and where it is only a schematic borrowing to be discharged later. A base manifold represents physical or computational histories, while each fiber contains the distinctions that remain available at that point in history. Continuation is represented by a connection, failure of perfect continuation by curvature, and repair by an optimization over admissible lifts through the distorted bundle. Persistent identity becomes a class of transport representations rather than an invariant label. Technical appendices supply a categorical formulation, a stochastic formulation, a gauge-invariant variational theory of repair, an explicit toy model of holonomy, the coupling to RSVP field dynamics, an account of type theory as an internal language for distinction-preserving computation, and a stack-theoretic treatment of global coherence.

Part II is a worked instantiation, optional on a first reading, in which the abstract apparatus is specialized to two existing formal systems: Spherepop, an event-sourced calculus of four primitive operations (Pop, Refuse, Bind, Collapse), and TARTAN, a recursive multiscale tiling of a continuous field substrate. Spherepop is shown to be a conservative, history-enriched extension of the typed lambda calculus, and separately to generate a categorical, generally noninvertible, connection whose curvature is literal operator non-commutativity. TARTAN supplies the missing multiscale geometry between Spherepop’s discrete events and Part I’s continuous transport, yielding a notion of persistence as a compatible family across scales. A reader interested only in the core theory can stop at the end of Part I; a reader who wants to see the apparatus applied to a concrete computational system should continue into Part II.

Contents

I	A Gauge Theory of Persistence, Repair, and Continuation	2
1	The Reversal of the Ontological Question	2
2	The Distinction Bundle	2
3	Continuation as a Connection	3
4	Curvature versus Holonomy	4
5	Memory as Path-Dependent Transport	4
6	Holonomy and Identity	5
7	Repair as Gauge-Invariant Curvature Compensation	5
8	The Minimal Repair Principle	6
9	Admissibility as a Field	7
10	Distinction Entropy and Normalized Persistence	7
11	Relation to RSVP	8
12	Continuation Capacity and Repair Flux	8
13	Continuation Operators over Discourse	9
14	Semantic Curvature	9
15	Identity as a Restricted Transport Representation	9
16	Ontological Phase Transitions	10
17	A General Persistence Functional	10
18	The Deep Unification	11
19	The Ontological Consequence	11
20	Ontological Death, Relativized	11
21	A Non-Identifiability Theorem	12
22	Toward a General Theory of Persistence	12
23	Conclusion to Part I: The Object Is the Future It Still Permits	13

II	A Worked Multiscale Instantiation: Spherepop and TARTAN	13
24	Spherepop as a Conservative Extension of Typed Lambda Calculus	13
24.1	Configurations and Primitives	14
24.2	Binding as Historical Substitution	14
24.3	Embedding, Simulation, and Erasure	14
25	Spherepop as a Categorical Connection	15
25.1	The Continuation Fiber	16
25.2	The Four Generators	16
25.3	The Categorical Connection and Its Reversible Sector	16
25.4	Curvature as Operational Noncommutativity	17
25.5	Holonomy as Residual Historical Action	17
25.6	Two Pathologies of Bind: The Double Bind and Schismogenesis	18
26	TARTAN: Recursive Multiscale Tiling	19
26.1	Tiles and Coherence	19
26.2	The Four Primitives as Tiling Operations	20
26.3	Annotated Noise and Induced Curvature	20
26.4	Three Scales of Curvature	21
26.5	Recursive Holonomy and Multiscale Persistence	21
26.6	TARTAN Repair	22
27	The Combined Architecture	22
A	Categorical Formulation of Continuation	23
A.1	History Category and Continuation Functor	23
A.2	Identity as Restricted Yoneda Equivalence	23
A.3	Histories as Objects, and Repair as Natural Transformation	24
B	Stochastic Continuation Geometry	24
C	Variational Repair and Optimal Control	25
D	An Explicit Toy Model of Holonomy	25
E	Continuation and RSVP Field Dynamics	26
F	Type Theory as an Internal Language for Distinction-Preserving Computation	26
F.1	Types as Objects, Contracts as Subobjects	26
F.2	Refinement Types and Distinction Contracts	27
F.3	Type Theory as Internal Language, Not Subset	27
G	A Stack-Theoretic Treatment of Global Coherence	28
G.1	Presheaves, Sheaves, and Their Limits	28
G.2	Local Systems versus Connections	28
G.3	Obstruction and Extendability	29

Part I. A Gauge Theory of Persistence, Repair, and Continuation

1. The Reversal of the Ontological Question

The conventional ontological question is: what is the thing that persists? This question presupposes that persistence is a secondary property of an already-defined object. One first identifies an entity x , then asks whether $x(t)$ remains the same entity at later times.

The present framework reverses the order and asks instead what makes a distinction continuable, then reconstructs an entity from the answer. Let $\mathcal{X}$ denote a state space. A state $x \in \mathcal{X}$ is not yet an entity in the ontologically interesting sense; it is merely a configuration. What matters is the set of transformations under which some distinction encoded by x remains admissible. Let $C(x)$ denote the set of admissible continuations of x . We regard $x \rightsquigarrow y$ as a continuation relation if $y \in C(x)$. The primitive object is therefore not x but the pair $(x, C(x))$, and a persistent structure is one for which the continuation relation remains sufficiently nonempty under transformation.

This changes the meaning of destruction. A structure is not destroyed merely because $x_t \neq x_{t+\Delta t}$. Destruction occurs when $C(x_{t+\Delta t})$ ceases to overlap adequately with the continuation structure inherited from x_t . Identity is therefore not state equality:

$$\text{Identity} \neq \text{state equality}, \quad \text{Identity} \sim \text{continuation compatibility}.$$

2. The Distinction Bundle

Let M be a manifold of histories. A point $p \in M$ represents a local historical context: a physical time, computational stage, biological condition, conversational state, or organizational configuration. Above each p , define a fiber E_p containing the distinctions available at that historical location, giving a fiber bundle

$$\pi : E \rightarrow M.$$

To keep the later gauge-theoretic claims rigorous rather than schematic, we fix a minimal model and treat every subsequent generalization as an explicit extension of it, not a silent redefinition. In the minimal model, E_p is a finite-dimensional vector space of distinction observables, $G \subseteq \text{GL}(E_p)$ is a structure group acting on the fibers, and

$$\nabla : \Gamma(E) \rightarrow \Omega^1(M, E)$$

is a connection compatible with G . Everything said in this Part about connections, curvature, and holonomy is rigorous relative to this model. Later generalizations — replacing the vector space fiber with a lattice, a Boolean or effect algebra, a category, or a measurable space — are extensions in the technical sense developed in Appendix A, and the reader should treat expressions borrowed from that broader setting as provisional until they are re-derived there.

A distinction in the minimal model is a linear observable $d \in E_p^*$, or more generally an element of E_p itself under a fixed pairing; two states x, y are equivalent relative to a family $\{d_i\} \subset E_p^*$ when $d_i(x) = d_i(y)$ for every i in the family, even though $x \neq y$ as points of the underlying state space. This gives an explicit version of a familiar but usually informal phenomenon: a system can change continuously while remaining recognizably the same, because the relevant distinctions are held fixed while irrelevant coordinates move.

A remark on the origin of the concept is in order before the formal development proceeds. The idea that not every detectable variation is constitutive – that what matters is a difference capable of altering what happens next – has a direct precursor in Gregory Bateson’s definition of information as a difference which makes a difference. Bateson’s point was that a physical variation only becomes informative relative to a receiving system for which it changes some subsequent state or action; an undetected or inconsequential variation carries no information for that system, however large it may be in absolute terms. The distinction fiber E_p is the same idea made compositional and quantitative: E_p does not contain every measurable difference at p , only those capable of changing which continuations remain admissible.

This can be made fully operational rather than left as a motivating slogan. Let $d : X \rightarrow \mathcal{V}_d$ be a candidate distinction and let $C_d(x)$ denote the continuations available when d is retained. Say d is *consequential at x* when there exist x_1, x_2 with

$$d(x_1) \neq d(x_2) \quad \text{and} \quad C_d(x_1) \not\supset C_d(x_2).$$

A detectable difference that has no effect on any relevant continuation need not belong to the constitutive distinction fiber; conversely, a difference may be physically small yet constitutive if it changes which transformations remain possible. On this reading, d belongs in E_p exactly when it is consequential at some x realizing p , exactly as an undetected or inconsequential difference carries no Batesonian information. This interpretation treats “difference” as a relation among an observation, a set of permitted operations, and a specified future horizon, not as a mystical or universal substance, and motivates treating E_p as a distinction space rather than a raw state space from the outset.

3. Continuation as a Connection

Suppose a history is parameterized by $\gamma : [0, 1] \rightarrow M$. At each point $\gamma(t)$ we have a distinction fiber $E_{\gamma(t)}$. To transport a distinction along the history we require a map

$$\mathcal{P}_{\gamma, t_0 \rightarrow t_1} : E_{\gamma(t_0)} \rightarrow E_{\gamma(t_1)}.$$

These transport operators always satisfy $\mathcal{P}_{t \rightarrow t} = I$ and

$$\mathcal{P}_{t_1 \rightarrow t_2} \mathcal{P}_{t_0 \rightarrow t_1} = \mathcal{P}_{t_0 \rightarrow t_2}.$$

This composition law is the ordinary functoriality of parallel transport along a single, concatenated path; it holds for curved connections exactly as it does for flat ones, and should not itself be read as a statement about “perfect continuation.” The property that does depend on curvature is a different one: whether two *different* paths sharing the same endpoints induce the same transport,

$$\mathcal{P}_{\gamma_1} = \mathcal{P}_{\gamma_2} \quad \text{for } \gamma_1, \gamma_2 : p \rightarrow q.$$

This second property generally fails when curvature, or global topology, is nontrivial, and it is this failure – comparison across paths, not composition along one path – that the rest of this Part is built on.

In differential form, transport is generated by a connection $\nabla = d + A$, where A is a connection one-form valued in the Lie algebra $\mathfrak{g}$ of G . For a curve γ , parallel transport satisfies

$$\frac{Ds}{dt} = \frac{ds}{dt} + A_\mu(\gamma(t)) \dot{\gamma}^\mu(t) s = 0, \quad \text{so} \quad \frac{ds}{dt} = -A_\mu \dot{\gamma}^\mu s.$$

The connection A answers: given what currently exists, what counts as the next admissible state? A conventional dynamical law evolves states. A continuation connection evolves admissibility.

4. Curvature versus Holonomy

The curvature of the connection is $F = dA + A \wedge A$. For two infinitesimal directions X, Y ,

$$[\nabla_X, \nabla_Y] = F(X, Y),$$

so curvature measures the infinitesimal failure of continuation operators to commute along nearby paths. It is essential, however, not to conflate curvature with the broader notion of path-dependent memory. If $F = 0$ everywhere, parallel transport is locally path-independent – but only locally, on a simply connected neighborhood. A flat connection on a base space that is not simply connected can still have nontrivial holonomy around a noncontractible loop:

$$F = 0, \quad \text{Hol}_\gamma \neq I.$$

This is not a corner case; it is the source of a second, purely topological, kind of memory. The theory therefore contains two distinct mechanisms for historical dependence: local, infinitesimal curvature, and global, topological holonomy on an otherwise flat connection. Conflating them would understate the theory's own resources.

5. Memory as Path-Dependent Transport

This gives a precise definition of memory that does not require a storage device, and that is correctly calibrated against the curvature/holonomy distinction just drawn. Let $\gamma_1, \gamma_2 : [0, 1] \rightarrow M$ share endpoints $\gamma_1(0) = \gamma_2(0) = p$ and $\gamma_1(1) = \gamma_2(1) = q$. If $\mathcal{P}_{\gamma_1} \neq \mathcal{P}_{\gamma_2}$, the system's present distinction structure depends on its historical path. Memory is characterized by

$$\gamma_1 \sim \gamma_2 \iff \mathcal{P}_{\gamma_1} = \mathcal{P}_{\gamma_2}.$$

Nonzero curvature is sufficient to produce infinitesimal path dependence, but it is not necessary for holonomic memory, by the flat-but-multiply-connected case above. The correct general statement is therefore:

Memory is operationally detectable path-dependent transport.

Curvature measures its local infinitesimal component; holonomy measures the finite, and potentially purely topological, result. A further qualification: a system whose present state differs after different histories does contain a physical record of those histories in its present configuration, even though no separate symbolic storage variable was introduced. Holonomy gives memory without a designated memory register, not memory without any present encoding at all.

6. Holonomy and Identity

Consider a closed loop $\gamma : S^1 \rightarrow M$ based at p . Parallel transport around the loop produces an operator

$$\text{Hol}_\gamma = \mathcal{P} \exp \left(- \oint_\gamma A \right) \in G.$$

Under a gauge transformation $A \mapsto A^g$, holonomy transforms by conjugation,

$$\text{Hol}_\gamma(A^g) = g(p) \text{Hol}_\gamma(A) g(p)^{-1},$$

so the conjugacy class of Hol_γ is gauge-invariant, while Hol_γ itself, in a chosen trivialization, is not. This has a direct consequence for how identity can be defined. Belonging to a single fixed conjugacy class is gauge-invariant but generally too weak to determine identity: many inequivalent connections can share the holonomy of one particular loop, or even the same holonomy group, without being continuation-equivalent in any richer sense.

A better identity invariant is the full transport functor restricted to a family of constitutive paths, rather than one loop or its conjugacy class. Let $\Pi_1^{\text{thin}}(M)$ be the thin path groupoid of M – paths identified only under homotopies sweeping out no area, the natural domain for parallel transport. Define

$$\mathcal{P}_A : \Pi_1^{\text{thin}}(M) \longrightarrow G\text{-}\mathbf{Tor},$$

the transport functor sending each thin-homotopy class of path to the corresponding torsor map between fibers. Two connections are then identity-equivalent, relative to a family Γ_{ess} of constitutive paths, when the restrictions of their transport functors to Γ_{ess} are naturally isomorphic – not merely when one particular holonomy’s conjugacy class agrees.

This also repairs a typing issue in the informal quotient $\Gamma_p / \approx$ used loosely below and in the introduction: the equivalence $\mathcal{P}_{\gamma_1} \sim \mathcal{P}_{\gamma_2}$ is only well-typed when γ_1 and γ_2 share both endpoints, or when a specified identification of their endpoint fibers is supplied. With that qualification, the persistent object is a gauge-invariant feature of the restricted transport functor rather than a single conjugacy class:

Object = orbit, under the transport functor restricted to Γ_{ess} , of distinguishability under continuation.

7. Repair as Gauge-Invariant Curvature Compensation

Ordinary theories of repair begin with a target state x^* and define repair as minimizing distance from it, $x \rightarrow x^*$. This is insufficient for systems whose identity depends on relational structure. Suppose the system suffers a deformation $A \rightarrow A + \delta A$. The repair problem is not necessarily $\min |\delta A|$. Instead we seek a correction R such that the resulting connection

$A' = A + \delta A + R$ restores some required transport structure – but the loss used to measure “restoration” must itself be gauge-invariant, or the theory will judge two physically identical connections as differently repaired merely because they are described in different trivializations.

The naive loss $|\text{Hol}_{A'}(\gamma) - \text{Hol}_A(\gamma)|^2$ fails this test: a holonomy is a group element in a chosen trivialization, and its matrix difference changes under a gauge transformation even when nothing physical has changed. The corrected repair functional uses a class function of the holonomy instead – a Wilson loop observable

$$W_\rho(\gamma; A) = \text{Tr } \rho(\text{Hol}_\gamma(A))$$

for a representation ρ of G , which is invariant under conjugation and hence under gauge transformation. Let Γ_{ess} denote a family of essential loops. The repair functional is

$$\mathcal{L}_{\text{repair}}[R] = \int_{\Gamma_{\text{ess}}} |W_\rho(\gamma; A + \delta A + R) - W_\rho(\gamma; A)|^2 d\mu(\gamma) + \lambda |R|^2,$$

so that repair is $R^* = \arg \min_R \mathcal{L}_{\text{repair}}[R]$. An equivalent gauge-minimized formulation optimizes over the residual gauge freedom directly:

$$\inf_{g \in G} d_G(\text{Hol}_{A+\delta A+R}(\gamma), g \text{Hol}_A(\gamma) g^{-1})^2.$$

This produces a radical reinterpretation: repair does not restore the past state, it restores the continuation equivalence class, measured by quantities that cannot be faked by a mere change of description. A repaired organism, repaired machine, reconstructed institution, or recovered semantic concept may contain entirely different local material while preserving the same global continuation geometry.

This raises a distinction the earlier, unrefined version of this theory elided. A pure gauge transformation $A \mapsto A^g$ changes only the description of a connection, not the connection’s physical content; performing one is *diagnostic gauge alignment*, not repair. Genuine, physical repair changes the gauge-equivalence class of the damaged connection so that it approaches the original class – it is a compensatory intervention, not a redescription. The three should be kept terminologically separate: redescription (pure gauge change, no physical effect), diagnostic gauge alignment (using gauge freedom to reveal that no physical repair was in fact needed), and physical compensatory intervention (an actual change to A that alters gauge-invariant quantities).

8. The Minimal Repair Principle

The functional above generalizes to a broader family of constitutive distinctions. Let $\mathcal{D}$ be a set of distinctions considered constitutive of identity, and let $\Delta_D(A, A')$ measure the distortion of distinction D under replacement of A by A' , using a gauge-invariant comparison as in the previous section. Define the structural repair cost

$$\mathcal{C}(A, A') = \int_{\mathcal{D}} w(D) \Delta_D(A, A') dD.$$

The preferred repair is $A^* = \arg \min_{A'} \mathcal{C}(A, A')$ subject to

$$\exists R \in \mathcal{R} \text{ such that } \text{Hol}_{A_d+R}(\gamma) \in \mathcal{O}_\gamma \quad \forall \gamma \in \Gamma_{\text{ess}},$$

where $\mathcal{R}$ is the admissible space of repair one-forms and $\mathcal{O}_\gamma$ is the target conjugacy class for loop γ ; this replaces an earlier ill-typed intersection condition between a space of one-forms and a set of holonomies. The constraint is not equality with an original configuration but admissibility of the resulting holonomy class. This gives a mathematical basis for the intuition that successful repair often produces something different from the original that is nevertheless “the same thing”: sameness is measured by continuation, not composition.

9. Admissibility as a Field

The continuation connection can be embedded in a more general admissibility field. Let $\mathcal{A}(x, y) \in [0, 1]$ measure the degree to which y is an admissible continuation of x , or, more carefully, let $\mathcal{A}_{\Delta t}(x, x + \dot{x} \Delta t)$ denote the admissibility of an infinitesimal transition. An infinitesimal rate is recovered as

$$\mathcal{A}_{\Delta t}(x, x + \dot{x} \Delta t) \asymp e^{-\Delta t L(x, \dot{x})}, \quad L(x, \dot{x}) = - \lim_{\Delta t \rightarrow 0} \frac{1}{\Delta t} \log \mathcal{A}_{\Delta t}(x, x + \dot{x} \Delta t),$$

which defines a genuine Lagrangian rather than treating $V(x, y) = -\log \mathcal{A}(x, y)$ as directly a function of $(x, \dot{x})$. Continuation then becomes a least-action problem:

$$\gamma^* = \arg \min_{\gamma} \int_0^T L(\gamma(t), \dot{\gamma}(t)) dt.$$

The velocity Hessian of L ,

$$g_{ij}(x, \dot{x}) = \frac{\partial^2 L}{\partial \dot{x}^i \partial \dot{x}^j},$$

defines a genuine Riemannian metric only when it is symmetric, nondegenerate, and positive definite; without those assumptions it gives a Finsler-type, or pseudo-Finsler, structure. Subsequent statements about “the geometry of persistence” should be read with this qualification attached. Physical distance remains one special case of a much larger notion of distance between continuations, but the general case is Finslerian rather than Riemannian.

10. Distinction Entropy and Normalized Persistence

Let a system possess a probability distribution over distinctions $p(d_i)$, with distinction entropy $H_D = -\sum_i p(d_i) \log p(d_i)$, where the enumeration and weighting of $\{d_i\}$ is itself an operational choice – sampled by an observer, activated by the system, or weighted by environmental relevance – that must be fixed before H_D is meaningful. Two systems may share Shannon entropy while possessing radically different continuation structures. Define a conditional distinction entropy

$$H_D(t + \Delta t \mid t) = - \sum_{d, d'} p(d, d') \log p(d' \mid d),$$

and the corresponding persistence information

$$I_D(t; t + \Delta t) = H_D(t + \Delta t) - H_D(t + \Delta t | t).$$

The bare condition $I_D(t; t + \Delta t) > 0$ is too weak to indicate persistence: almost any temporally correlated stochastic process satisfies it. A more useful normalized quantity, comparable across systems and scales, is

$$\mathfrak{P}_D(\tau) = \frac{I(D_t; D_{t+\tau})}{H(D_t)}, \quad H(D_t) > 0,$$

optionally compared against a null model and required to exceed a scale-dependent threshold before being read as evidence of genuine persistence.

11. Relation to RSVP

The framework is particularly informative when coupled to the Relativistic Scalar–Vector Plenum (RSVP). Let the RSVP fields be $\Phi(x, t)$, $\mathbf{v}(x, t)$, $S(x, t)$, with usual dynamics defining a physical field configuration $X_{\text{RSVP}} = (\Phi, \mathbf{v}, S)$. Distinction Holonomy Theory adds a second layer, a distinction functional $\mathfrak{D}[\Phi, \mathbf{v}, S]$, and the connection becomes

$$A_\mu = A_\mu[\Phi, \partial\Phi, \mathbf{v}, \partial\mathbf{v}, S, \partial S].$$

The RSVP fields do not merely constitute a medium in which structures exist; they induce a geometry of admissible continuation. Schematically,

$$A = \alpha_\Phi d\Phi + \alpha_v \iota_{\mathbf{v}} d\mathbf{x} + \alpha_S dS + A_{\text{int}}, \quad F = dA + A \wedge A.$$

A persistent structure is then a region of RSVP configuration space for which the induced holonomy remains within a stable conjugacy class $\mathcal{O} = \{ghg^{-1} \mid g \in G\}$: persistence requires approximately $\text{Hol}_\gamma(A) \in \mathcal{O}$ for a family of admissible loops γ , subject to the caveat of Section 6 that a fixed conjugacy class alone is a necessary, not sufficient, identity criterion. This gives a reading of thermodynamic persistence: a structure does not survive because its matter remains fixed, but because the field configuration continues to support the same family of distinguishable transformations.

12. Continuation Capacity and Repair Flux

An earlier formulation of continuation capacity as $\mathcal{K}(x) = \mu(\mathcal{C}(x))$ is severely measure-dependent for continuous state spaces: the value can change under reparameterization or diverge, and it counts every admissible future equally regardless of probability. A corrected definition uses either a threshold on normalized probability mass,

$$\mathcal{K}_\theta(x) = K(x, \{y : \mathcal{A}(x, y) \geq \theta\}),$$

or an effective support,

$$\mathcal{K}_{\text{eff}}(x) = \exp H(Q(\cdot \mid x)),$$

where $Q(\cdot \mid x)$ is a normalized continuation distribution (Appendix B). Neither definition should be read as measuring successful persistence by itself: a randomizing failure mode can increase either quantity while destroying every constitutive distinction. Capacity must therefore always be indexed by an essential distinction family D_{ess} before it is used as a diagnostic.

With that qualification, a system undergoing destructive smoothing may exhibit $d\mathcal{K}_\theta/dt < 0$ restricted to D_{ess} ; a self-organizing structure instead maintains $d\mathcal{K}_\theta/dt \approx 0$ despite substantial microscopic change; repair corresponds to positive continuation-capacity flux, $d\mathcal{K}_\theta/dt > 0$. The distinction between entropy production and structural persistence can be written as $\dot{S} > 0$ while simultaneously $\dot{\mathcal{K}}_\theta \geq 0$ on D_{ess} . There is no contradiction: a system can increase ordinary thermodynamic entropy while preserving or increasing the number of constitutive distinctions it can continue to realize.

13. Continuation Operators over Discourse

The same geometry applies to computational and linguistic systems. Let $\mathcal{S}$ be a semantic state space. An instruction acts as an operator $P : \mathcal{S} \rightarrow \mathcal{S}$, and a discourse is a composition $P_n \circ P_{n-1} \circ \dots \circ P_1$. The relevant object is not the sequence of outputs but the shrinking family of admissible continuations $\mathcal{A}_0 \supseteq \mathcal{A}_1 \supseteq \dots \supseteq \mathcal{A}_n$. Each instruction deforms the continuation connection: $A_{n+1} = A_n + \delta A(P_n)$.

Two instruction sequences can produce the same immediate output while producing different future continuation geometries: $P_2 P_1(s) = Q_2 Q_1(s)$ does not imply $P_3 P_2 P_1 = Q_3 Q_2 Q_1$. The systems can be observationally equivalent at one time while possessing different holonomies. Context, on this reading, is not merely information about the past; it is a deformation of the connection governing the future.

14. Semantic Curvature

Consider two semantic operators P, Q with commutator $[P, Q] = PQ - QP$. For a small rectangular loop generated by P and Q ,

$$e^{\epsilon P} e^{\epsilon Q} e^{-\epsilon P} e^{-\epsilon Q} = I + \epsilon^2 [P, Q] + O(\epsilon^3),$$

so the residual transformation after returning to the original semantic location is $\Delta_{\text{loop}} = \epsilon^2 [P, Q] + O(\epsilon^3)$. This gives a precise criterion for historical dependence: if $[P, Q] = 0$, the order of contextual constraints is locally irrelevant; if $[P, Q] \neq 0$, the semantic system possesses path-dependent memory. The same equation describes a physical system, a neural process, a discourse, a state machine, or an institution. A fully worked instance of this commutator, with an explicit base manifold, bundle, and loop, is given in Appendix D.

15. Identity as a Restricted Transport Representation

Let Γ_p denote the set of histories beginning at p that share a common endpoint q within Γ_p , and define $\gamma_1 \approx \gamma_2$ when $\mathcal{P}_{\gamma_1} \sim \mathcal{P}_{\gamma_2}$ under gauge equivalence, restricted throughout to pairs of paths with compatible endpoints as required in Section 6. An entity is then represented not by one trajectory but by an equivalence class $[\gamma] \in \Gamma_p / \approx$, restricted to the constitutive path

family Γ_{ess} . The object is not the trajectory, nor even the terminal state; it is the equivalence class of histories that preserve the same continuation structure on Γ_{ess} :

$$\text{Entity} = \text{history class modulo continuation equivalence on } \Gamma_{\text{ess}}.$$

The traditional question, is this the same object, becomes $[\gamma_1] = [\gamma_2]$? The answer depends on the chosen distinction algebra, connection, and constitutive path family. Identity is therefore relative but not arbitrary: it is relative to a mathematically specified structure of admissible distinctions and tests.

16. Ontological Phase Transitions

Suppose a control parameter λ changes the connection, $A = A(\lambda)$. There may exist a critical value λ_c at which the holonomy group changes discontinuously,

$$\text{Hol}_p(A_{\lambda < c}) \not\cong \text{Hol}_p(A_{\lambda > c}).$$

The system has then undergone an ontological phase transition, stronger than a phase transition in ordinary state variables: the space of possible continuations itself has changed. An entity can disappear without any single state variable becoming singular. Formally, if $\mathcal{E}_\lambda$ is the set of persistent equivalence classes in the sense of Section 15, an ontological transition occurs when a structural invariant of $\mathcal{E}_\lambda$ changes discontinuously.

17. A General Persistence Functional

A persistence functional should compare transport against a fixed reference rather than measuring a quantity that is constant for unitary holonomy by construction; $\text{Tr}[\text{Hol}_\gamma(A) \text{Hol}_\gamma(A)^\dagger]$ is exactly such a degenerate quantity, since it equals $\dim(\rho)$ for any unitary representation and therefore carries no information. The corrected functional compares against a fixed reference connection A_0 via a fidelity,

$$\mathcal{P}[A] = \int_{\Gamma} w(\gamma) \left| \text{Tr}(U_0(\gamma)^\dagger U_A(\gamma)) \right|^2 d\mu(\gamma) - \lambda \int_M |F|^2 dV,$$

where $U_0(\gamma) = \text{Hol}_\gamma(A_0)$ and $U_A(\gamma) = \text{Hol}_\gamma(A)$. The first term now genuinely measures retained transport structure relative to the reference; the second penalizes uncontrolled curvature. Whether a persistent structure extremizes this functional, $\delta\mathcal{P} = 0$ with $\delta^2\mathcal{P} < 0$, should be treated as a proposed model of self-maintenance, not a derived consequence of the preceding geometry: open dissipative systems generally settle into attractors or nonequilibrium steady states rather than extrema of an ordinary action, and the passage from this functional to a local conservation law

$$D_\mu F^{\mu\nu} = J_{\text{cont}}^\nu, \quad J_{\text{cont}}^\nu = \frac{\delta\mathcal{P}_{\text{persistence}}}{\delta A_\nu},$$

is schematic; a nonlocal integral over holonomies produces distributional loop currents in general, not automatically a smooth local current, and a careful variation would be needed to justify the local form. What appears phenomenologically as self-maintenance can be represented, as a modeling proposal, by the stabilization of a gauge field over distinction space relative to a reference.

18. The Deep Unification

Several apparently unrelated phenomena become instances of the same mathematical object under this reading. A biological organism maintains a continuation bundle. A memory system maintains path-dependent transport. A repaired machine restores transport class. A discourse modifies the semantic connection. An institution persists when its admissible transformations remain coherent. A concept persists when its inferential continuation class survives reformulation. A physical dissipative structure persists when its dynamical environment maintains a stable transport structure. In every case the question is the same: which distinctions remain transportable through transformation, and relative to which family of constitutive tests?

This yields a general hierarchy,

distinction $\rightarrow$ admissibility $\rightarrow$ continuation $\rightarrow$ curvature $\rightarrow$ holonomy $\rightarrow$ persistence,

with repair reversing the destructive direction, curvature $\rightarrow$ diagnosis $\rightarrow$ compensation $\rightarrow$ restored transport $\rightarrow$ renewed persistence, and memory occupying the intermediate position, path $\rightarrow$ noncommutativity or topological holonomy $\rightarrow$ historical dependence. The apparent plurality of phenomena is generated by one underlying mathematical architecture, made precise rather than merely evocative by the corrections of Sections 4 through 13.

19. The Ontological Consequence

The most consequential claim of the theory is that persistence precedes objecthood. If an object is defined independently of its possible transformations, persistence is an additional property. If instead an object is defined as a stable equivalence class under admissible transformations, persistence is constitutive:

$$x \equiv [C(x)].$$

The thing is the continuation structure; matter realizes it; information describes it; memory records its path dependence; repair restores its transport class; dynamics realizes its trajectory; identity identifies its restricted transport representation. This is a different ontology from either substance metaphysics, thing $\rightarrow$ change, or simple process metaphysics, process $\rightarrow$ thing. Distinction Holonomy Theory proposes that continuation structure gives rise to process and thing simultaneously: the object is a stabilized possibility of continuation.

20. Ontological Death, Relativized

Let $\mathcal{C}_t$ denote the continuation space of a structure at time t . An earlier overlap measure,

$$\frac{\mu(\mathcal{C}_t \cap \mathcal{P}_{t_0 \rightarrow t} \mathcal{C}_{t_0})}{\mu(\mathcal{C}_{t_0})},$$

is not generally normalized to $[0, 1]$ unless parallel transport preserves the measure, since the denominator need not bound the numerator when $\mathcal{P}_{t_0 \rightarrow t} \mathcal{C}_{t_0}$ changes measure under transport. A symmetric, better-behaved alternative is the Jaccard overlap,

$$\mathfrak{D}(t, t_0) = \frac{\mu(\mathcal{C}_t \cap \mathcal{P}_{t_0 \rightarrow t} \mathcal{C}_{t_0})}{\mu(\mathcal{C}_t \cup \mathcal{P}_{t_0 \rightarrow t} \mathcal{C}_{t_0})} \in [0, 1].$$

A structure remains persistent when $\mathfrak{D}(t, t_0) > \theta$ for a specified threshold, degrades when $0 < \mathfrak{D}(t, t_0) \leq \theta$, and undergoes ontological death when $\mathfrak{D}(t, t_0) = 0$. This criterion should be understood, throughout, as relative to a specified distinction family, measure, horizon, and threshold – ontological death relative to a stated continuation model, not an absolute condition. This definition deliberately permits complete material replacement: if $x_t \neq x_{t_0}$ but $\mathfrak{D}(t, t_0) > 0$, the entity can remain persistent, while $x_t \approx x_{t_0}$ does not guarantee persistence if the continuation structure has changed. The criterion separates phenomenological similarity from ontological continuity.

21. A Non-Identifiability Theorem

The claim that present-state equivalence does not imply identity equivalence becomes substantive, rather than nearly tautological, when stated as a non-identifiability result about estimation, not merely an observation about definitions.

Theorem 21.1. *Let $O : X \rightarrow Y$ be an observation map. Suppose there exist two latent continuation processes, with transport functors $\mathcal{P}_A$ and $\mathcal{P}_B$ over histories γ_A, γ_B sharing a common base point and inducing the same distribution on $O(X_t)$ at every t , but with*

$$\text{Hol}_{\gamma_A} \neq \text{Hol}_{\gamma_B}$$

in the sense of distinct induced predictive distributions on future observations, or distinct restricted transport functors on Γ_{ess} . Then no estimator that is a function of the instantaneous observation $O(X_t)$ alone can consistently recover the continuation class of the underlying process.

The proof is immediate from the hypothesis: any such estimator factors through $O(X_t)$, on which the two processes are indistinguishable by construction, so it must return the same value on both, contradicting the requirement to recover distinct classes. The content of the theorem lies entirely in exhibiting that such pairs (γ_A, γ_B) exist whenever the continuation connection has nonzero curvature or nontrivial holonomy – which Appendix D does explicitly. Identity is therefore necessarily historical whenever the continuation connection possesses nonzero curvature or nontrivial holonomy: the state can be identical while the future remains different, and no amount of instantaneous observation closes that gap.

22. Toward a General Theory of Persistence

The theory’s most defensible core, stripped of the extensions treated as appendices, is:

distinction bundle + connection + path-groupoid transport + gauge-invariant repair loss.

It can be summarized by five equations: distinction, as a linear functional on the fiber, $d \in E_p^*$; admissibility, as an infinitesimal rate $L(x, \dot{x})$ derived from a transition kernel; continuation, $\nabla = d + A$; curvature, $F = dA + A \wedge A$; and identity,

$$\mathfrak{I} = \Gamma_{\text{ess}} / \approx_{\mathcal{P}}, \quad \gamma_1 \approx_{\mathcal{P}} \gamma_2 \iff \mathcal{P}_{\gamma_1} \sim \mathcal{P}_{\gamma_2} \text{ on shared endpoints.}$$

The theory replaces the primitive concept of persistent substance with a quotient space of continuation-compatible histories, restricted to a stated family of constitutive tests. The decisive quantity is not how much of the original state remains, but how much of the original

future remains possible – a distinction small in language and large in mathematics. The fundamental measure of persistence is therefore not $|x_t - x_0|/|x_0|$ but a quantity of the form

$$\mathfrak{P}(t) = \text{Overlap} \left(C_t, C_0^{\parallel} \right),$$

computed with the Jaccard normalization of Section 20, where $C_0^{\parallel}$ is the original continuation space transported into the present fiber. Persistence is a geometric relation between possibilities, made precise rather than merely asserted by the corrections collected in this Part.

23. Conclusion to Part I: The Object Is the Future It Still Permits

The usual ontology of persistence asks why a thing remains itself despite change. Distinction Holonomy Theory asks a more primitive question: what changes are still possible without destroying the distinction by which the thing is constituted? A state is an instantaneous realization; a process is a trajectory; a memory is path-dependent transport; a repair is gauge-invariant curvature compensation; an identity is a restricted transport representation; an object is a stable continuation equivalence class relative to a stated family of tests. These are not separate phenomena but different projections of the same structure.

A persistent entity is not what remains unchanged; it is what remains continuable. To destroy an entity is not necessarily to alter its state beyond recognition – it is to eliminate the space of admissible futures through which its distinctions can continue to exist. This makes continuation more fundamental than identity, repair more fundamental than restoration, and distinction more fundamental than representation. The ontology of persistence is a geometry of possible futures: what we call a thing is the region of that geometry that refuses, for a time, to become impossible.

Part II specializes this apparatus to two concrete formal systems. It is not required for the theory stated above, which stands on its own; it is offered as a demonstration that the apparatus can be operated on a real computational substrate rather than only on continuous, idealized fibers.

Part II. A Worked Multiscale Instantiation: Spherepop and TARTAN

This Part is optional. It specializes Part I to two existing formal systems – an event-sourced calculus and a recursive multiscale tiling – and can be skipped by a reader interested only in the abstract theory. Its purpose is to show that the transport/curvature/holonomy apparatus is operable on a concrete computational substrate, not merely evocative of one.

24. Spherepop as a Conservative Extension of Typed Lambda Calculus

Before Spherepop’s four primitives can be read as generators of a continuation connection, it is worth establishing precisely what kind of calculus they generate. The claim is not that

$\{\text{Pop}, \text{Refuse}, \text{Bind}, \text{Collapse}\}$ alone reconstitute the ordinary lambda calculus; on their canonical semantics they form a state-transition algebra over an append-only history and an option space. The provable claim is narrower and stronger:

Spherepop is a conservative, history-enriched extension of typed lambda calculus.

Conservative means every ordinary lambda term is a Spherepop term, every ordinary beta-reduction can be performed inside Spherepop, and Spherepop assigns the same result to a pure lambda term apart from the additional history recording its evaluation.

24.1 Configurations and Primitives

A Spherepop configuration is $C = (\sigma, h)$, where σ is the current live option space and $h = (e_1, \dots, e_n)$ is the ordered append-only history. The primitive operations act on configurations:

$$\text{Pop}(\omega) : (\sigma, h) \rightarrow (\sigma \setminus \{\omega\}, h \cdot e_{\text{pop}}(\omega)), \quad \text{Refuse}(a) : (\sigma, h) \rightarrow (\sigma, h \cdot e_{\text{refuse}}(a)),$$

$$\text{Bind}(x, a) : (\sigma, h) \rightarrow (\sigma[x \mapsto a], h \cdot e_{\text{bind}}(x, a)), \quad \text{Collapse}_c : (\sigma, h) \rightarrow (\pi_c(\sigma), h \cdot e_{\text{collapse}}(c)),$$

where $\pi_c : X \twoheadrightarrow Q$ is an admissible quotient certified by c . Evaluation in Spherepop therefore both produces a value and changes the live continuation space, recording how the result became admissible.

24.2 Binding as Historical Substitution

Ordinary application reduces by substitution, $(\lambda x:A. t) a \rightarrow_{\beta} t[x := a]$, usually treated as a metalevel operation with no persistent object-level trace. Spherepop internalizes the substitution as a binding event, $\text{Bind}(H, x, a) = H \cdot e_{\text{bind}}(x, a)$, extending the environment $\rho' = \rho[x \mapsto a]$. Spherepop beta-reduction is

$$(\sigma, H, \rho, (\lambda x:A. t) a) \rightarrow_{\beta_{\text{sp}}} (\sigma', H \cdot e_{\text{bind}}(x, a), \rho[x \mapsto a], t),$$

so Spherepop preserves the defining computational operation of lambda calculus – application of an abstraction is substitution of an argument – and adds only the requirement that the substitution become an irreversible historical event.

24.3 Embedding, Simulation, and Erasure

Define the embedding $\llbracket - \rrbracket : \Lambda \rightarrow \text{Spherepop}$ recursively by $\llbracket x \rrbracket = x$, $\llbracket \lambda x:A. t \rrbracket = \lambda x:\llbracket A \rrbracket. \llbracket t \rrbracket$, and $\llbracket t u \rrbracket = \llbracket t \rrbracket \llbracket u \rrbracket$.

Proposition 24.1 (Substitution lemma). *For all lambda terms t, u : $\llbracket t[x := u] \rrbracket = \llbracket t \rrbracket[x := \llbracket u \rrbracket]$.*

The proof is by structural induction on t , with the abstraction case using alpha-renaming to avoid variable capture; each case reduces directly to the induction hypothesis applied inside the embedding.

Theorem 24.1 (Simulation). *If $t \rightarrow_{\beta} t'$ in ordinary lambda calculus, then for every initial configuration (σ, H) , $(\sigma, H, \llbracket t \rrbracket) \rightarrow_{\text{sp}}^+ (\sigma', H', \llbracket t' \rrbracket)$ for some σ', H' with $H \preceq H'$ a prefix.*

It suffices to check the beta-redex case $t = (\lambda x:A. v) u, t' = v[x := u]$. Under the embedding, Spheredpop performs the binding event $e_\beta = e_{\text{bind}}(x, \llbracket u \rrbracket)$ and reduces to $\llbracket v \rrbracket[x := \llbracket u \rrbracket]$, which by the substitution lemma equals $\llbracket v[x := u] \rrbracket = \llbracket t' \rrbracket$. Reduction inside a larger term follows by closure under evaluation contexts.

Theorem 24.2 (Typing preservation). *If $\Gamma \vdash_\lambda t : A$, then $H; \llbracket \Gamma \rrbracket \vdash_{\text{SP}} \llbracket t \rrbracket : \llbracket A \rrbracket \triangleright H'$ for some $H' \succeq H$.*

The variable and abstraction rules carry over directly; the application rule becomes

$$\frac{H; \Gamma \vdash f : \Pi x:A. B \quad H; \Gamma \vdash a : A}{H; \Gamma \vdash f a : B[x := a] \triangleright H \cdot e_{\text{bind}}(x, a)},$$

the ordinary application rule augmented by monotone history production; no ordinary derivation is invalidated.

Define the erasure map $|-| : \text{Spheredpop} \rightarrow \Lambda$ on the pure fragment by $|x| = x, |\lambda x:A. t| = \lambda x:A. |t|, |tu| = |t| |u|$, discarding runtime histories: $|(\sigma, H, t)| = |t|$. Then $|\llbracket t \rrbracket| = t$ for every lambda term t , and any Spheredpop step evaluating an embedded pure term either is administrative, with $|u| = t$ unchanged, or corresponds to lambda reduction, $t \rightarrow_\beta |u|$.

Theorem 24.3 (Conservativity). *For any closed pure lambda term t and normal form v : $t \rightarrow_\beta^* v$ if and only if $(\sigma_0, \epsilon, \llbracket t \rrbracket) \rightarrow_{\text{SP}}^* (\sigma_f, H_f, \llbracket v \rrbracket)$ up to administrative history events, and erasing the Spheredpop state recovers v .*

This follows by combining the embedding, simulation, typing preservation, and erasure results above. Spheredpop therefore contains typed lambda calculus conservatively: its distinctive contribution is that substitution is no longer an invisible metalevel event. Selection becomes Pop; inadmissibility becomes Refuse; substitution becomes a recorded Bind; quotienting becomes Collapse. Schematically, ordinary beta-reduction becomes the richer transition

$$\text{propose} \rightarrow \text{check} \rightarrow \text{bind} \rightarrow \text{commit or refuse} \rightarrow \text{record}.$$

Spheredpop's computational type can be written schematically as $\llbracket A \rightarrow B \rrbracket_{\text{SP}} = A \rightarrow \text{State}(\Sigma)(\text{Writer}(H)(B))$, or, with refusal represented explicitly, $\Sigma \times H \rightarrow (B + \text{Refusal}) \times \Sigma \times H$: a standard lambda-calculus architecture enriched by state, writer, and exception-like effects. Spheredpop's originality is not abandoning lambda calculus but making historical commitment and continuation-space mutation constitutive of evaluation.

25. Spheredpop as a Categorical Connection

Having fixed Spheredpop's relation to lambda calculus, we now read its four primitives as generators of continuation transport in the sense of Part I. The essential qualification is that Pop, Refuse, and Collapse are generally noninvertible, so the resulting theory is not an ordinary Lie-group gauge theory; it is a categorical, generally semigroup-valued, connection whose reversible sector reduces to the conventional case of Part I.

25.1 The Continuation Fiber

Over each historical context $p \in M$, define a fiber $E_p = (\Omega_p, \mathcal{R}_p, \mathcal{Q}_p)$, where Ω_p is the live option space, $\mathcal{R}_p$ the current dependency structure, and $\mathcal{Q}_p$ the current quotient structure. A Spherpap world is $\mathcal{W}_p = (H_p, E_p)$ with append-only history H_p . Each primitive operation produces a transport map $T_a : E_p \rightarrow E_q$ and appends the corresponding event: $(H_p, E_p) \xrightarrow{a} (H_p \cdot e_a, E_q)$. A history $\gamma = (a_1, \dots, a_n)$ induces the composite $T_\gamma = T_{a_n} \circ \dots \circ T_{a_1}$: Spherpap's discrete connection.

25.2 The Four Generators

For $\omega \in \Omega$ with committed set K , $T_{\text{Pop}(\omega)}(\Omega, K, \mathcal{R}, \mathcal{Q}) = (\Omega \setminus \{\omega\}, K \cup \{\omega\}, \mathcal{R}, \mathcal{Q})$: Pop moves ω from possible continuations into committed history, and is normally noninvertible, since the post-Pop state does not by itself contain every fact needed to reconstruct the pre-Pop option space:

Pop = selection plus irreversible commitment.

For a candidate a with admissibility $\text{Adm}_E(a) \in \{0, 1\}$, $T_{\text{Refuse}(a)}(H, E) = (H \cdot e_{\text{Refuse}}(a), E)$ in the canonical case where refusal does not consume Ω . The live option space is unchanged, but the historical fiber changes, and future admissibility may depend on the record: $\text{Adm}_{H \cdot e_{\text{Refuse}}(a), E}(b) \neq \text{Adm}_{H, E}(b)$ in general. Refuse is the identity only after forgetting history:

Refuse = exclusion recorded without option consumption.

For a directed dependency relation $\mathcal{R} \subseteq X \times X$, $T_{\text{Bind}(a,b)}(\Omega, \mathcal{R}, \mathcal{Q}) = (\Omega, \mathcal{R} \cup \{(a, b)\}, \mathcal{Q})$; if Bind realizes substitution, $T_{\text{Bind}(x,u)}(\rho) = \rho[x \mapsto u]$. For an admissibility field, $\mathcal{A}_{\text{Bind}(a,b)}(x, y) = \mathcal{A}(x, y) + J_{ab}(x, y)$ for a coupling term J_{ab} :

Bind = deformation of independence into relational continuation,

and is the most natural source of noncommutativity among the four, since introducing one dependency can change the meaning of a later selection, refusal, or collapse.

For a certified equivalence $\sim_c$ on X with quotient $Q_c = X / \sim_c$ and projection $\pi_c : X \twoheadrightarrow Q_c$, $T_{\text{Collapse}_c} = \pi_c$; represented as an endomorphism on a common ambient space it is idempotent, $T_{\text{Collapse}_c}^2 = T_{\text{Collapse}_c}$, hence a projection rather than generally a gauge transformation:

Collapse = certified loss of distinctions through quotienting,

preserving identity when its kernel contains no constitutive distinction, and destroying it otherwise.

25.3 The Categorical Connection and Its Reversible Sector

Because Pop, Refuse, and Collapse are generally noninvertible, their transports belong to a category or monoid rather than a Lie group G . Let $\mathcal{H}$ be the category whose objects are Spherpap configurations and whose morphisms are finite sequences of primitive operations, and let $\mathcal{D}$ be a category of distinction structures. A continuation functor

$$\mathsf{T} : \mathcal{H} \longrightarrow \mathcal{D}, \quad \mathsf{T}(a) = T_a : E_p \rightarrow E_q, \quad \mathsf{T}(\gamma_2 \circ \gamma_1) = \mathsf{T}(\gamma_2) \circ \mathsf{T}(\gamma_1)$$

is the categorical form of the connection. Restricting to the maximal reversible subcategory $\mathcal{H}^\times \subseteq \mathcal{H}$ and to isomorphisms in $\mathcal{D}$ recovers a groupoid on which transport may be represented by the conventional group-valued connection of Part I. Outside $\mathcal{H}^\times$, Spherpap requires genuinely noninvertible parallel transport, and the theory of Part I should be read as governing this reversible sector exactly.

25.4 Curvature as Operational Noncommutativity

For two primitive operations a, b , order matters when $T_b T_a \neq T_a T_b$. When the transports act linearly on a common space, define the discrete curvature defect $F(a, b) = T_b T_a - T_a T_b$; in general, where subtraction is unavailable, curvature is represented by the pair $(T_b \circ T_a, T_a \circ T_b)$ or a comparison morphism between them:

Curvature asks: does changing the order of two admissible operations change the future distinction structure?

Concretely, $T_{\text{Pop}(\omega)} T_{\text{Bind}(a, \omega)} \neq T_{\text{Bind}(a, \omega)} T_{\text{Pop}(\omega)}$ when binding to ω is possible before ω is consumed but impossible afterward; $T_{\text{Collapse}_c} T_{\text{Bind}(a, b)} \neq T_{\text{Bind}(\pi_c(a), \pi_c(b))} T_{\text{Collapse}_c}$ when the collapse identifies distinctions needed to establish the binding; and $T_{\text{Pop}(\omega)} T_{\text{Refuse}(a)} \neq T_{\text{Refuse}(a)} T_{\text{Pop}(\omega)}$ whenever the recorded refusal modifies the admissibility of ω . These are not analogies to curvature; they are explicit failures of transport squares to commute, formalized as the obstruction to a coherent identification $E_{ba} \cong E_{ab}$ between the two paths from E to a common prospective target.

A primitive curvature table records which generator pairs commute: $[\text{Pop}, \text{Refuse}] \neq 0$ in general, since exclusion can affect subsequent selection; $[\text{Pop}, \text{Bind}] \neq 0$, since a selected option may participate in a new dependency; $[\text{Pop}, \text{Collapse}] \neq 0$, when several selectable options collapse into one quotient class; $[\text{Refuse}, \text{Bind}] \neq 0$, when the refused relation is the one Bind would establish; $[\text{Refuse}, \text{Collapse}] \neq 0$, when Collapse erases the distinction identifying the refused candidate; and $[\text{Bind}, \text{Collapse}] \neq 0$, when projection identifies endpoints or removes a dependency distinction. The exact vanishing conditions are Spherpap's local flatness laws: for instance $[\text{Bind}(a, b), \text{Collapse}_c] = 0$ exactly when the binding descends through the quotient, i.e. there exists a unique $\overline{\text{Bind}}$ on Q_c with $\pi_c \circ \text{Bind} = \overline{\text{Bind}} \circ \pi_c$ – the standard categorical condition that Bind factor through Collapse.

25.5 Holonomy as Residual Historical Action

For a loop $\gamma : p \rightarrow p$ returning the observable base state to its starting point, $\text{Hol}_\gamma = T_{a_n} \circ \dots \circ T_{a_1}$. Even when the base state returns to p , the distinction structure or history may have changed, $\text{Hol}_\gamma(E_p) \neq E_p$; in Spherpap this holonomy is, in general, a noninvertible endomorphism $\text{Hol}_\gamma \in \text{End}_{\mathcal{D}}(E_p)$ rather than an automorphism, a generalized or monoid-valued holonomy. If the observable option space returns to its original cardinality after selection and replenishment, the final history still contains the committed Pop, so

$$\pi_{\text{obs}}(\text{Hol}_\gamma(E_p)) = \pi_{\text{obs}}(E_p), \quad \text{while} \quad \text{Hol}_\gamma(E_p) \neq E_p.$$

The residual difference is Spherpap memory. The resulting correspondence with Part I is

Pop $\leftrightarrow$ selection and commitment, Refuse $\leftrightarrow$ recorded exclusion, Bind $\leftrightarrow$ coupling, Collapse $\leftrightarrow$ projecti

with composites defining transport, noncommutativity defining curvature, and residual loop action defining holonomy:

Spherepop is an event-generated categorical connection on a bundle of option spaces.

25.6 Two Pathologies of Bind: The Double Bind and Schismogenesis

The categorical connection just constructed gives a precise home for two patterns first identified, without this apparatus, in Gregory Bateson’s work on communication and relationship. These should be retained as interactional examples clarifying the formal machinery, not presented as an established general explanation of any clinical condition; the double bind’s original clinical association with schizophrenia in particular is far stronger a claim than the material here supports, and is not asserted.

In the present formalism, a double bind is a configuration in which a participant is subject to incompatible constraints and is additionally prevented from challenging the level at which those constraints are imposed. Let a, b be candidate actions with admissibility predicates $\mathcal{A}_1(a), \mathcal{A}_2(b)$. A simple contradiction is one in which no action satisfies both constraints, $\{x : \mathcal{A}_1(x) = 1\} \cap \{x : \mathcal{A}_2(x) = 1\} = \emptyset$. A double bind adds a metaconstraint forbidding refusal or reformulation of the constraint system itself: if R denotes the operation of challenging the framing, then

$$\mathcal{A}_{\text{meta}}(R) = 0.$$

The failure is therefore not merely that the object-level option space is empty; the system also excludes the operation that would revise that space. In Spherepop terms, incompatible Binds constrain every available Pop, while Refuse is unavailable at the level where it would be effective: if $\text{Bind}(a, x)$ and $\text{Bind}(b, x)$ jointly constrain x in a way no admissible $\text{Pop}(\omega)$ can satisfy, and no morphism of $\mathcal{H}$ available at the object level targets the binding itself, then the object-level history admits no repair within the existing control space – not, more strongly, “no admissible repair” as such, since a repair may become possible once that control space is expanded. This is the operative distinction from a plain curvature trap: the obstruction is not that two paths disagree, but that the binding structure generating the incompatibility is not itself an admissible target of Pop or Refuse at the level where the victim operates. Repair, correspondingly, requires a metalevel change unavailable within the trapped level: distinguishing previously conflated message levels (a new Pop that was not previously an option), refusing one of the bindings directly (Refuse applied to $\text{Bind}(a, x)$ or $\text{Bind}(b, x)$ themselves, not to x), or retiling the interaction, in the sense of Section 26.6 below, into contexts governed by different constraints.

Schismogenesis – Bateson’s term for the progressive amplification of a pattern of interaction between two parties – is the corresponding pathology of iterated, rather than blocked, Bind. Let x_n, y_n describe two coupled participants, $x_{n+1} = F(x_n, y_n)$, $y_{n+1} = G(y_n, x_{n+1})$, with combined update $T(x_n, y_n) = (F(x_n, y_n), G(y_n, F(x_n, y_n)))$. Divergence occurs when repeated application $T^n(x_0, y_0)$ amplifies a relational difference: in a symmetrical process, similar responses mutually escalate; in a complementary process, unlike responses reinforce an asymmetric relation. This should be described as divergent holonomy only under the precise condition that an identifiable interactional cycle returns the system to the same coarse relational context while changing its finer continuation structure – holonomy is not a synonym for any feedback loop, and applies specifically when a closed or recurrent path has a residual

transport effect, T^n returning the coarse-grained state while $\text{Hol}_{T^n} \neq I$ on the finer distinction structure. Neither party's local state explains the divergence; it is generated entirely by repeated transport around the relational loop, exactly as Part I's discourse-curvature apparatus (Section 14) would predict for a coupling that deforms its own connection at each iteration.

It is worth being explicit about the scope of this comparison. Bateson supplies the phenomena and the essential intuition – that some communicative patterns admit no move within their own level, and that some relational dynamics diverge under their own iteration – well before any gauge-theoretic vocabulary existed to describe them. The formal apparatus above does not retroactively attribute the mathematics to Bateson, nor does it revive the stronger clinical claims associated with the double bind's original context; it shows only that the phenomena he identified are recognizable, in retrospect, as a pathological Bind structure and a divergent holonomy, respectively, within a machinery built for other purposes and only later found to fit.

26. TARTAN: Recursive Multiscale Tiling

Neither Part I's continuous fibers nor Spherpap's discrete events, by themselves, supply the geometry connecting them: a principled account of which regions of a continuous field admit one coherent local continuation law, and how that account changes under refinement. TARTAN – Trajectory-Aware Recursive Tiling with Annotated Noise – provides this multiscale structure, recursively partitioning the history space into coherence tiles, each carrying a local field state, entropy, noise annotation, scale, and trajectory estimate. The methodological precedent is Turing's, in two distinct respects: the theory of computation shows how complex behavior can be generated by composition of elementary operations, and the reaction–diffusion account of morphogenesis shows how local interactions can generate spatial pattern without a central pattern-maker; TARTAN combines these only at the level of formal organization – complex global structure analyzed through local rules, spatial coupling, and the repeated differentiation of initially less differentiated regions – and neither claims equivalence to reaction–diffusion dynamics nor that every tiled system reproduces biological morphogenesis. Spherpap then governs changes to the tiling:

Pop $\leftrightarrow$ selection and refinement, Refuse $\leftrightarrow$ exclusion of an inadmissible adjacency,

Bind $\leftrightarrow$ coupling between tiles, Collapse $\leftrightarrow$ projection or coarse-graining.

Distinction holonomy measures what survives transport through the resulting recursively changing tile system.

26.1 Tiles and Coherence

Let the RSVP field state be $\Psi = (\Phi, \mathbf{v}, S)$ over a domain M . At scale k , a tiling $\mathcal{T}_k = \{T_i^{(k)}\}_{i \in I_k}$ covers M , with each tile

$$T_i^{(k)} = (U_i^{(k)}, \Psi_i^{(k)}, A_i^{(k)}, \eta_i^{(k)}, \tau_i^{(k)}, H_i^{(k)}),$$

recording support, local RSVP state, local continuation connection, annotated noise, predicted trajectory, and local append-only history. Recursive tiling allows a tile to be replaced by children, $T_i^{(k)} \rightsquigarrow \{T_{i1}^{(k+1)}, \dots, T_{im}^{(k+1)}\}$ with $U_i^{(k)} = \bigcup_a U_{ia}^{(k+1)}$, giving a refinement structure $\mathcal{T}_0 \preceq \mathcal{T}_1 \preceq \mathcal{T}_2 \preceq \dots$: not one base geometry, but a recursively changing family of bases.

A tile should be formed where field dynamics admit a sufficiently coherent local continuation model. With a coherence functional

$$\kappa = \alpha_\Phi \|\nabla \Phi\|^2 + \alpha_v \|\nabla \mathbf{v}\|^2 + \alpha_S \|\nabla S\|^2 + \alpha_F \|F\|^2 + \alpha_\eta \text{Var}(\eta),$$

a tile is a connected region $\{x \in M : \kappa(x, t) \leq \varepsilon_i\}$: an improvement over defining coherence from entropy gradient alone, since a region can have small $|\nabla S|$ either because it is genuinely coherent or because it has been destructively homogenized, and the curvature and noise terms help distinguish these. A coherence tile is a region over which transport is approximated by one local connection $\nabla_i = d + A_i$ with curvature $F_i = dA_i + A_i \wedge A_i$; low curvature means nearby alternative paths within the tile transport distinctions approximately equivalently, not that the tile is globally memoryless.

26.2 The Four Primitives as Tiling Operations

Pop becomes trajectory-aware refinement: if a tile T contains several distinguishable continuation regimes, $\text{Pop}_\omega(T) = (T_\omega, T_{-\omega})$ with $T_\omega = \{x \in T : \omega \text{ selected}\}$, and the selection is committed to history. Pop therefore operates at two levels at once – choosing one continuation, and introducing a new tile boundary that is itself a realized distinction:

Pop = commitment that becomes geometry.

Trajectory awareness means refinement can depend on predicted divergence, $d(\tau_\omega(t+\Delta t), \tau_{-\omega}(t+\Delta t)) > \delta$, so states locally similar now may be separated because their admissible futures diverge.

Refuse becomes prohibited gluing: for adjacent tiles T_i, T_j with a proposed gluing map $g_{ij} : E_i|_{U_i \cap U_j} \rightarrow E_j|_{U_i \cap U_j}$, $\text{Refuse}(g_{ij})$ removes g_{ij} from the admissible edges of the tile-adjacency graph $\mathcal{G}$ while the refusal itself remains in history. Refuse denies a proposed overlap identification (in sheaf language), records that a proposed output fails an interface contract (in type-theoretic language, Appendix F), or removes a candidate transport from the admissible connection (in gauge language).

Bind becomes intertile coupling, $\text{Bind}(T_i, T_j) : T_i \rightarrow T_j$, represented by an operator $B_{ij} : E_i \rightarrow E_j$ or a coupled-field action term $\sum_{i,j} \int_{U_i \cap U_j} \|s_j - B_{ij}s_i\|^2 dV$; the local dynamics becomes $\partial_t \Psi_j = \mathcal{F}_j[\Psi_j] + \lambda_{ij} \mathcal{B}_{ij}[\Psi_i, \Psi_j]$. Because binding changes both the tile graph and the local dynamics, it is again the natural source of curvature: $\text{Pop}_i \circ \text{Bind}_{ij} \neq \text{Bind}_{ij} \circ \text{Pop}_i$ in general, since the two orders generate different descendant couplings.

Collapse becomes recursive coarse-graining: given child tiles $\{T_{i1}^{(k+1)}, \dots, T_{im}^{(k+1)}\}$, a certificate c defines $\pi_c : \prod_a T_{ia}^{(k+1)} \rightarrow T_i^{(k)}$, admissible only when constitutive structure descends through the quotient – a coarse binding $\bar{B}$ with $\pi_c \circ B = \bar{B} \circ \pi_c$, and an effective coarse connection $A_i^{(k)} = \mathcal{R}_c(A_{i1}^{(k+1)}, \dots, A_{im}^{(k+1)})$ for a renormalization operator $\mathcal{R}_c$. Collapse is destructive exactly when $\pi_c(x) = \pi_c(y)$ but $\text{Hol}(x) \not\sim \text{Hol}(y)$ for distinct child histories:

Collapse = coarse-graining certified by continuation equivalence.

26.3 Annotated Noise and Induced Curvature

Each tile carries a noise annotation $\eta_i = (\mu_i, \Sigma_i, \ell_i, \chi_i)$ – local bias, covariance, correlation scale, provenance – rather than one undifferentiated random term. The local connection becomes

stochastic, $A_i = \bar{A}_i + \Xi_i$ with $\mathbb{E}[\Xi_i] = \mu_i$, and the curvature is correspondingly random, with expectation

$$\mathbb{E}[F_i] = d\bar{A}_i + \bar{A}_i \wedge \bar{A}_i + \mathbb{E}[\Xi_i \wedge \Xi_i] + \text{cross terms.}$$

Consequently $\mathbb{E}[\Xi_i] = 0$ does not imply $\mathbb{E}[F_i] = F(\bar{A}_i)$: zero-mean connection noise can generate nonzero effective curvature. Correlated noise need not merely obscure an underlying continuation geometry; it can create systematic holonomy.

26.4 Three Scales of Curvature

Curvature separates into three levels. Within a tile, differential curvature is $F_i = dA_i + A_i \wedge A_i$ as usual. Across a boundary, transition curvature measures incompatibility between local connections: compatible transport requires $A_j = g_{ij}^{-1} A_i g_{ij} + g_{ij}^{-1} dg_{ij}$, and the boundary defect

$$K_{ij} = A_j - (g_{ij}^{-1} A_i g_{ij} + g_{ij}^{-1} dg_{ij})$$

signals that the two tiles cannot be treated as local descriptions of one smoothly transported structure when $K_{ij} \neq 0$. Across scales, recursive curvature measures the failure of refinement and transport to commute: for a coarse-graining map $r_{k+1,k}$, scale compatibility requires $r_{k+1,k} \circ T_\gamma^{(k+1)} = T_\gamma^{(k)} \circ r_{k+1,k}$, and the defect

$$K_\gamma^{(k+1,k)} = r_{k+1,k} \circ T_\gamma^{(k+1)} - T_\gamma^{(k)} \circ r_{k+1,k}$$

being nonzero means a history visible at the fine scale is misrepresented after coarse-graining: the system can appear flat macroscopically while retaining hidden microscopic holonomy.

26.5 Recursive Holonomy and Multiscale Persistence

A trajectory through TARTAN moves among tiles and scales, $\gamma : T_{i_0}^{(k_0)} \rightarrow T_{i_1}^{(k_1)} \rightarrow \dots \rightarrow T_{i_n}^{(k_n)}$, each step an operator $U_m : E_{i_{m-1}}^{(k_{m-1})} \rightarrow E_{i_m}^{(k_m)}$ that may be a continuous transport, a Spherepop event, an intertile binding, a refinement, or a quotient projection. The TARTAN holonomy is the path-ordered composite $\text{Hol}_\gamma^{\text{TARTAN}} = U_n \circ \dots \circ U_1$, generally an endomorphism rather than a group element. Even for a loop returning to the same coarse tile, $T_{i_n}^{(k_n)} = T_{i_0}^{(k_0)}$, one may have $\text{Hol}_\gamma^{\text{TARTAN}} \neq I$: the residual transformation records not only where the trajectory traveled but which distinctions were selected, excluded, coupled, or collapsed, and at which scales:

TARTAN memory = holonomy through space, history, and scale.

At each scale k , let $\mathfrak{J}_k$ denote the continuation identity represented by $\mathcal{T}_k$, with coarse-graining maps $r_{k+1,k} : \mathfrak{J}_{k+1} \rightarrow \mathfrak{J}_k$. A scale-coherent persistent entity is a compatible family $(\mathfrak{J}_0, \mathfrak{J}_1, \dots)$ with $r_{k+1,k}(\mathfrak{J}_{k+1}) = \mathfrak{J}_k$, and the complete entity is the inverse limit

$$\mathfrak{J} = \varprojlim_k \mathfrak{J}_k.$$

An object need not exist at one privileged resolution; it exists as compatibility among its representations across resolutions. A failure confined to one scale need not destroy the entity; ontological failure in the TARTAN sense occurs only when no compatible family remains, $\varprojlim_k \mathfrak{J}_k = \emptyset$ – stronger than disappearance from any single tiling.

26.6 TARTAN Repair

Damage may occur within a tile, at a boundary, or between scales, so repair has three components, $R = (R_{\text{local}}, R_{\text{glue}}, R_{\text{scale}})$, restoring intratile transport, gluing compatibility, and cross-scale agreement respectively, with combined functional

$$\mathcal{L}_{\text{TARTAN}}[R] = \sum_i \int_{U_i} \|F_i^R - F_i^{\text{ref}}\|^2 dV + \lambda_b \sum_{i,j} \|K_{ij}^R\|^2 + \lambda_s \sum_k \|K^{(k+1,k),R}\|^2 + \lambda_R \|R\|^2,$$

minimized subject to preserving the constitutive holonomy observables. The distinctive TARTAN possibility is retiling: if a damaged tile can no longer sustain one coherent continuation law, repair need not restore its original interior. Spherepop can Pop it into several tiles, $T_i \rightarrow \{T_{i1}, T_{i2}\}$, Bind their viable dependencies, Refuse destructive crossings, and Collapse distinctions irrelevant at the repaired scale: repair as adaptive repartitioning of the future, and the operation the double-bind discussion of Section 25 invokes by name.

27. The Combined Architecture

The complete architecture reads

RSVP $\rightarrow$ continuous field dynamics, TARTAN $\rightarrow$ trajectory-aware multiscale decomposition,

Spherepop $\rightarrow$ event grammar transforming that decomposition, DHT (Part I) $\rightarrow$ transport, curvature, holonomy

RSVP supplies the evolving substrate $\Psi = (\Phi, \mathbf{v}, S)$; TARTAN identifies the regions and scales over which local continuation laws are coherent; Spherepop changes those regions through Selection, Exclusion, Coupling, and Projection; and Part I's apparatus measures the historical effect of transporting constitutive distinctions through the resulting structure:

Pop = trajectory-sensitive refinement, Refuse = prohibited gluing,

Bind = intertile coupling, Collapse = continuation-preserving coarse-graining.

Their noncommutativity generates curvature; their path-ordered composition generates holonomy; TARTAN makes both recursive and scale-dependent. The resulting picture departs from Part I's single-tile idealization in one significant respect: an entity is not represented by one distinction fiber over one base manifold. It is represented by the compatibility of its local continuation structures across a recursively changing tiling:

A persistent entity is a multiscale descent object whose constitutive holonomy survives admissible retiling.

This is the sense in which Part II is an instantiation of Part I rather than a departure from it: every object introduced here – tile, gluing map, coarse-graining, retiling – is a concrete realization of a construct Part I already named abstractly (fiber, transition function, coarse-graining functor, admissible lift), specialized to a substrate, RSVP fields decomposed by TARTAN and transformed by Spherepop, on which the abstract claims can be checked rather than merely asserted.

A. Categorical Formulation of Continuation

The bundle formulation of Part I generalizes by replacing the state-space picture with a category of admissible continuations, and this appendix supplies the extension referenced there.

A.1 History Category and Continuation Functor

Let $\mathcal{H}$ be a category of historical contexts: objects are contexts $p, q, \dots$, and a morphism $\gamma : p \rightarrow q$ is an admissible history from p to q , with composition given by temporal concatenation and $1_p : p \rightarrow p$ representing continuation without displacement. The simplest choice is a path category of a history space M ; if curvature is to remain detectable, ordinary homotopy is too coarse, and parallel transport should factor through the thin path groupoid $\Pi_1^{\text{thin}}(M)$, as in Section 6.

Let $\mathcal{D}$ be a category of distinction structures. A continuation theory is a functor $T : \mathcal{H} \rightarrow \mathcal{D}$, with $T(p) = D_p$ the distinction structure at p and $T(\gamma) : D_p \rightarrow D_q$ the transport along γ . Functoriality, $T(1_p) = 1_{D_p}$ and $T(\gamma_2 \circ \gamma_1) = T(\gamma_2) \circ T(\gamma_1)$, is exactly the composition-along-one-path law of Section 3 and does not require flatness. Curvature appears when different histories with the same endpoints induce different morphisms, $\gamma, \gamma' : p \rightarrow q$ with $T(\gamma) \neq T(\gamma')$; holonomy is the restriction of T to loops, $\text{Hol}_p : \text{End}_{\mathcal{H}}(p) \rightarrow \text{Aut}_{\mathcal{D}}(D_p)$. The categorical core is therefore:

A continuation structure is a representation of a history category.

The target category $\mathcal{D}$ fixes what a distinction means: **Vect** or Hilbert spaces for observable-valued distinctions, Boolean, Heyting, or effect algebras for propositional distinctions, partition lattices for partition-valued distinctions, or measurable spaces with Markov kernels for probabilistic distinctions. Part I's minimal model (Section 2) fixes $\mathcal{D} = \mathbf{Vect}$; the theory should not define a distinction simultaneously as a function, a partition, and a fiber element, but rather fix $\mathcal{D}$ and treat the others as separate realizations. If continuation is represented by a state evolution $f_\gamma : X_p \rightarrow X_q$, an observable $d_q : X_q \rightarrow \mathcal{V}$ pulls back contravariantly, $f_\gamma^* d_q = d_q \circ f_\gamma$, with $f_{\gamma_2 \circ \gamma_1}^* = f_{\gamma_1}^* \circ f_{\gamma_2}^*$; a manuscript in this framework must fix once and for all whether continuation transports states forward or distinctions backward.

A.2 Identity as Restricted Yoneda Equivalence

A naive continuation profile $\mathfrak{C}(x) = \{\text{Hom}_{\mathcal{C}}(x, y)\}_y$ omits how postcomposition acts between these hom-sets; the natural invariant is the representable functor $h_x = \text{Hom}_{\mathcal{C}}(x, -) : \mathcal{C} \rightarrow \mathbf{Set}$. By the Yoneda lemma, $h_x \cong h_y \iff x \cong y$: the full continuation profile recovers categorical isomorphism exactly and produces no identity notion weaker than it. To obtain the relative identity Part I intends (Section 15), restrict to a subcategory of constitutive tests $\mathcal{T} \hookrightarrow \mathcal{C}$ and define $h_x^{\mathcal{T}} = \text{Hom}_{\mathcal{C}}(x, -)|_{\mathcal{T}}$, with

$$x \sim_{\mathcal{T}} y \iff h_x^{\mathcal{T}} \cong h_y^{\mathcal{T}}.$$

This says x and y are identical relative to the futures and distinctions that matter, without being isomorphic under every possible test; the dependence of identity on a selected distinction algebra is captured by this restricted Yoneda construction, which is what Γ_{ess} denotes throughout Part I.

A coarse-graining functor $F : \mathcal{C} \rightarrow \mathcal{C}'$ preserving $x \simeq_{\mathcal{T}} y \Rightarrow F(x) \simeq F(y)$ merely respects an existing equivalence; preserving the morphisms constitutive of persistence requires suitable faithfulness on the designated subcategory, or preservation and reflection of the chosen equivalence relation. Markov kernels are not literally probability-enriched hom-sets in the naive sense, since kernel composition involves integration; the correct setting is the Kleisli category of the Giry monad, a Markov category, or an enriched category over a monoidal category of kernels.

A.3 Histories as Objects, and Repair as Natural Transformation

There are two categorical ontologies available. In the first, states are objects and histories are morphisms, naturally producing holonomy from endomorphisms $p \rightarrow p$. In the second, histories are objects and history comparisons are morphisms; a clean combination is the arrow category $\text{Arr}(\mathcal{H})$, whose objects are histories $\gamma : p \rightarrow q$ and whose morphisms are commuting squares $b \circ \gamma = \gamma' \circ a$.

Suppose undamaged and damaged systems are functors $T_0, T_d : \mathcal{H} \rightarrow \mathcal{D}$. A repair cannot generally be one map between states; it should transform one entire continuation representation into another. A candidate repair is a natural transformation $R : T_d \Rightarrow T_r$ into a repaired theory T_r , naturality requiring $T_r(\gamma) \circ R_p = R_q \circ T_d(\gamma)$ for every $\gamma : p \rightarrow q$: repairing before transport and transporting before repair agree. Exact restoration requires a natural isomorphism $T_r \cong T_0$; the weaker and more usual DHT requirement, given a forgetful functor $Q : \mathcal{D} \rightarrow \mathcal{D}_{\text{ess}}$ onto constitutive distinctions, is $Q \circ T_r \cong Q \circ T_0$: the repaired system need not recover every microscopic distinction, only the transport of the constitutive ones – the categorical form of Section 8’s minimal repair principle.

B. Stochastic Continuation Geometry

Deterministic continuation is insufficient for physical, biological, and cognitive systems. Let the system evolve by $dX_t = b(X_t, t) dt + \sigma(X_t, t) dW_t$ with transition kernel $K_t(x, dy)$. Given an admissibility kernel $\mathcal{A}_t(x, y)$, define the weighted transition kernel $\tilde{K}_t(x, dy) = \mathcal{A}_t(x, y) K_t(x, dy)$, total continuation capacity $\mathcal{K}_t(x) = \int_X \tilde{K}_t(x, dy)$, and the normalized continuation distribution

$$Q_t(dy | x) = \frac{\mathcal{A}_t(x, y) K_t(x, dy)}{\mathcal{K}_t(x)}.$$

The persistence information between t and $t + \tau$ is

$$I_{\text{cont}}(t, \tau) = D_{\text{KL}}(Q_{t, t+\tau} \| Q_{t, t+\tau}^{\text{null}}),$$

using the standard relative-entropy notation, against an explicitly chosen reference process Q^{null} ; calling this quantity persistence information is defensible only once that null model is fixed, since otherwise it measures distinguishability from an unspecified reference, not persistence as such. A structure is persistent when $I_{\text{cont}}(t, \tau)$ remains bounded away from zero over the relevant scale, and, following Section 10, the normalized version $I_{\text{cont}}(t, \tau)/H(Q_{t, t_0})$ is generally the more comparable diagnostic.

Given a damaged kernel Q^{dam} , the optimal repair kernel, as an information projection onto the admissible family, is

$$Q^* = \arg \min_{Q \in \mathcal{Q}_{\text{admissible}}} D_{\text{KL}}(Q \parallel Q^{\text{dam}}) \quad \text{subject to} \quad I_{\text{cont}}(Q) \geq I_{\text{crit}},$$

with $\mathcal{Q}_{\text{admissible}}$ and the relevant absolute-continuity assumptions specified per application. Repair is a constrained projection in probability space: the system need not return to its original distribution, only to a region of probability space in which the necessary distinctions remain continuable.

C. Variational Repair and Optimal Control

Let A_0 be an undamaged connection, $A_d = A_0 + \delta A$ the damaged connection, R a repair field, and $A_r = A_d + R$ the repaired connection. With action

$$S[A_r] = \int_M \left[\frac{1}{4} \text{Tr}(F_r \wedge \star F_r) + V(A_r) \right],$$

the repair problem is

$$R^* = \arg \min_R \{ S[A_d + R] + \lambda |R|^2 \} \quad \text{subject to} \quad \text{Hol}_{A_d+R}(\gamma) \in \mathcal{O}_\gamma \quad \forall \gamma \in \Gamma_{\text{ess}},$$

using the corrected, well-typed admissibility condition of Section 8 rather than a set-intersection between one-forms and holonomies. The Euler–Lagrange equation takes the schematic form $D_\mu F^{\mu\nu} + \delta V / \delta A_\nu + 2\lambda R^\nu = J_\gamma^\nu$, where J_γ^ν is the constraint force associated with preserving holonomy, giving the general form: damage plus constraint yields an optimal compensatory field.

If the admissible repair space is $\mathcal{R} \subseteq \Gamma(T^*M \otimes \mathfrak{g})$, repair exists only if there is some $R \in \mathcal{R}$ with $\text{Hol}_{A_d+R}(\gamma) \in \mathcal{O}_\gamma$ for every constitutive loop, producing a sharp distinction between recoverable damage (such an R exists) and ontological damage (no admissible repair lies within the available control space):

irreversibility = absence of an admissible continuation-preserving repair.

D. An Explicit Toy Model of Holonomy

Consider generators

$$X = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}, \quad Y = \begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix}, \quad [X, Y] = XY - YX = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \neq 0.$$

The group commutator $\gamma_\epsilon = e^{\epsilon X} e^{\epsilon Y} e^{-\epsilon X} e^{-\epsilon Y}$ satisfies, by Baker–Campbell–Hausdorff,

$$\gamma_\epsilon = I + \epsilon^2 [X, Y] + O(\epsilon^3).$$

This algebraic fact alone is not yet a holonomy model, since no base manifold, bundle, connection, or loop in the base has been specified. To make it one: take $M = \mathbb{R}^2$, the trivial rank-two bundle, and the constant connection

$$A = X dx + Y dy.$$

Then $F = dA + A \wedge A = [X, Y] dx \wedge dy \neq 0$, and parallel transport around a small coordinate rectangle of side ϵ reproduces γ_ϵ above, up to a sign fixed by the transport convention and loop orientation. A system transported around this loop returns to its original base point $(0, 0)$ but not to its original fiber state:

$$\Delta_\gamma = \gamma_\epsilon - I = \epsilon^2[X, Y] + O(\epsilon^3) \neq 0.$$

No storage variable and no explicit historical record has been introduced; the memory is encoded entirely in the failure of transport to be path-independent, and the pair of paths realizing $\gamma_A \neq \gamma_B$ required by the non-identifiability theorem of Section 21 is exhibited explicitly by the two orderings $e^{\epsilon X} e^{\epsilon Y}$ versus $e^{\epsilon Y} e^{\epsilon X}$ around this loop.

E. Continuation and RSVP Field Dynamics

Let $\Psi = (\Phi, \mathbf{v}, S)$ evolve by $\partial_t \Psi = \mathcal{F}[\Psi]$, and let the continuation connection be a functional of the configuration, $A_\mu = A_\mu[\Psi, \partial\Psi]$, with curvature $F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu + [A_\mu, A_\nu]$. A coupled action

$$S_{\text{total}} = S_{\text{RSVP}}[\Phi, \mathbf{v}, S] + \eta \int_M \text{Tr}(F \wedge \star F) + \zeta \int_M \mathcal{V}_{\text{cont}}(\Psi, A)$$

couple the RSVP fields to a geometric persistence sector on variation with respect to Ψ . The local persistence density

$$\rho_P(x) = \text{Tr}(F_{\mu\nu} F^{\mu\nu}) + \alpha \mathcal{K}_\theta(x),$$

using the corrected capacity of Section 12, gives a persistent dissipative structure as a dynamically maintained region $\Omega \subset M$ for which $\int_\Omega \rho_P dV$ remains within a bounded attractor despite ongoing entropy production: a possible bridge between RSVP's thermodynamic field picture and the continuation ontology of Part I.

F. Type Theory as an Internal Language for Distinction-Preserving Computation

Type theory is not literally a subset of category theory, but it can be interpreted categorically, and doing so makes precise a distinction implicit throughout this paper between an interface that merely composes and one that transports the distinctions its consumers require.

F.1 Types as Objects, Contracts as Subobjects

Interpret a type A as an object $\llbracket A \rrbracket \in \text{Ob}(\mathcal{C})$ and a typed function $f : A \rightarrow B$ as a morphism $\llbracket f \rrbracket : \llbracket A \rrbracket \rightarrow \llbracket B \rrbracket$, with composite programs interpreted as composite morphisms. Type checking, $\text{cod}(f) = \text{dom}(g)$, is already a categorical continuation condition: the output of f must land in the interface expected by g .

An ordinary type contract enforces only this outer shape. A function $n : \text{Image} \rightarrow \text{Label}$ that ignores its input and returns a random label has exactly the type of a competent classifier $f : \text{Image} \rightarrow \text{Label}$; both are equally composable morphisms with the same domain and codomain. Hence

$$\text{type correctness} \not\equiv \text{behavioral correctness},$$

not because type theory has failed, but because the chosen type is too weak to express the relevant distinction. In a categorical interpretation, a predicate P on A is represented by a subobject $P \hookrightarrow A$; a function $f : A \rightarrow B$ satisfies a precondition/postcondition contract (P, Q) when it restricts to a morphism $\bar{f} : P \rightarrow Q$ making the evident square commute. A noise function can inhabit the coarse hom-set $\text{Hom}_C(A, B)$ while failing to inhabit the contract-respecting subspace $\text{Hom}_{\text{Contract}}((A, P), (B, Q))$: the category of bare functions is too permissive, and the relevant category has objects (A, P) and morphisms required to satisfy $P(x) \Rightarrow Q(f(x))$, or, for a relational postcondition, $P(x) \Rightarrow R(x, f(x))$.

F.2 Refinement Types and Distinction Contracts

Dependent and refinement types make the required relation explicit rather than leaving it to an unenforced comment. Instead of $f : A \rightarrow B$, require

$$f : \prod_{x:A} \sum_{y:B} R(x, y),$$

so that for every input the function must supply both an output and evidence that the constitutive relation R holds; a pure-noise function generally cannot inhabit this stronger type, since it cannot supply the witness. Equivalently, defining $B_x = \{y : B \mid R(x, y)\}$, the strengthened interface is $f : (x : A) \rightarrow B_x$ – for instance, replacing $\text{sort} : \text{List}(A) \rightarrow \text{List}(A)$ with

$$\text{sort} : (xs : \text{List}(A)) \rightarrow \{ys : \text{List}(A) \mid \text{Sorted}(ys) \wedge \text{Permutation}(xs, ys)\}.$$

A function returning random permutations satisfies the weak interface but not the refined one, which records exactly the distinctions – multiplicity and order – that must survive. This is the difference between output compatibility and continuation compatibility in the vocabulary of the rest of this paper.

F.3 Type Theory as Internal Language, Not Subset

Category theory begins with objects and morphisms and asks how they compose; type theory supplies a formal language constructing only the morphisms admitted by its judgments, $\Gamma \vdash t : A$, meaning t is an admissible inhabitant of A in context Γ . A function judgment $\Gamma \vdash f : A \rightarrow B$ says f can participate in the continuation $A \xrightarrow{f} B$; dependent and refinement types let the interface constrain the relation between input and output, $\Gamma, x:A \vdash f(x) : B(x)$. The precise relation is:

Type theory is an internal language for suitable categories.

Simply typed lambda calculus corresponds to cartesian closed categories; dependent type theory to locally cartesian closed categories, categories with families, or comprehension categories; homotopy type theory to higher groupoids and ∞ -categorical structures. Two functions of the same type $f, n : A \rightarrow B$ are observationally equivalent relative to a testing language when $T \circ f = T \circ n$ for every admissible test $T : B \rightarrow O$; if some admissible T separates them, the interface has omitted a distinction relevant to some continuation – the restricted-Yoneda idea of Appendix A applied to software interfaces, where a specification selects a testing category $\mathcal{T}$ and two implementations are contract-equivalent when $T \circ f \sim T \circ g$ for every $T \in \mathcal{T}$.

For stochastic programs represented as Markov kernels $K : A \rightsquigarrow B$, the analogous strengthened type is

$$K : \prod_{x:A} \left\{ \mu \in \mathcal{P}(B) \mid \mu(\{y : R(x, y)\}) \geq 1 - \varepsilon \right\},$$

which a pure-noise kernel with the right stochastic interface but the wrong distribution fails to inhabit. This clarifies the standing of type-shaped interfaces throughout: a raw function is weaker than a well-typed function, which is weaker than a contract-preserving function, which is weaker than a distinction-preserving continuation in the sense of Part I. The governing principle is:

An interface is useful only to the extent that its type exposes the distinctions its consumers require.

Types guarantee that an output can enter the next stage; contracts guarantee it enters with the relevant properties intact; category theory explains how these stages compose; Distinction Holonomy Theory adds that the identity of a computation lies in the continued transport of those properties, not in the mere passage of values through correctly shaped interfaces.

G. A Stack-Theoretic Treatment of Global Coherence

Category theory, in the appendices above, describes continuation along transformation. This appendix addresses a different question: whether locally available continuations assemble into a globally coherent structure.

G.1 Presheaves, Sheaves, and Their Limits

For a topological space, manifold, or other site M , let $\mathcal{C}(U)$ be the set of continuation structures realizable on an open $U \subseteq M$, with restriction maps $\rho_{UV} : \mathcal{C}(U) \rightarrow \mathcal{C}(V)$ for $V \subseteq U$. A genuine sheaf satisfies local identity (agreement on a cover implies equality) and gluing (a compatible family on a cover has a unique amalgamation). It follows that a true sheaf cannot contain a compatible family that fails to glue; a claimed “gluing obstruction” for a sheaf of sets is therefore either a failure of the compatibility condition itself, or a signal that the correct object is not a sheaf of sets but something with weaker equality.

That weaker object is generally required here. Gauge fields are defined locally and related on overlaps by gauge transformations, not literal equality: for local connections A_i on U_i , $A_j = g_{ij}^{-1} A_i g_{ij} + g_{ij}^{-1} dg_{ij}$ on overlaps, with transition functions satisfying the cocycle condition $g_{ij} g_{jk} = g_{ik}$ on triple overlaps. The natural global object assigns a *groupoid* $\mathcal{C}(U)$ to each U – objects are local continuation structures, morphisms are gauge equivalences – and descent asks whether local objects and their overlap isomorphisms assemble into a global object. This is a stack, not a sheaf of sets, and it is the appropriate object for a theory in which identity is already intended to be gauge-relative: equality of local sections is too strict, coherent equivalence is correct.

G.2 Local Systems versus Connections

An ordinary sheaf does not by itself produce curvature or holonomy. A local system – a locally constant sheaf – encodes flat transport and is equivalent, under mild conditions, to a representation $\rho : \pi_1(M, p) \rightarrow \text{Aut}(D_p)$: this gives global holonomy with only flat local geometry,

exactly the topological memory of Section 4, with $F = 0$ but $\rho([\gamma]) \neq I$. To describe nonzero curvature requires a bundle or principal bundle with connection, or an appropriate sheaf or stack of connections, with parallel transport as a functor $\mathsf{T}_A : \Pi_1^{\text{thin}}(M) \rightarrow G\text{-Tor}$. Local system corresponds to flat transport; bundle with connection corresponds to possibly curved transport; a continuation sheaf alone does not generate the full connection-curvature-holonomy sequence without this additional differential structure.

G.3 Obstruction and Extendability

If transition functions $g_{ij} : U_i \cap U_j \rightarrow G$ satisfy $g_{ij}g_{jk}g_{ki} = 1$ on triple overlaps, they define a principal G -bundle; if this fails, the defect $c_{ijk} = g_{ij}g_{jk}g_{ki}$ measures a gluing inconsistency, determining, in the abelian case, a genuine Čech cohomology class, and requiring nonabelian cohomology or stack-theoretic descent more generally. Persistence is then the existence of descent data whose obstruction class vanishes, $[\omega_{\text{cont}}] = 0$; ontological fragmentation occurs when local continuation structures exist but no admissible choice of overlap maps yields coherent descent, and repair, correspondingly, modifies local sections or overlap maps so that $[\omega_{\text{dam}}] \xrightarrow{R} [\omega_{\text{repaired}}] = 0$ – genuine obstruction cancellation, not merely making local sections resemble one another.

A complementary, more local notion of damage asks, for $V \subseteq U$ and a local section $s_V \in \mathcal{C}(V)$, whether s_V lies in the image of $\rho_{UV} : \mathcal{C}(U) \rightarrow \mathcal{C}(V)$ – whether s_V extends through U . With $\mathcal{U}^+(V)$ the family of permitted future extensions, the extension profile

$$\text{Ext}(s_V) = \{U \in \mathcal{U}^+(V) \mid \exists s_U \in \mathcal{C}(U), s_U|_V = s_V\}$$

gives a natural sheaf-theoretic realization of the paper’s central claim:

What persists is the structured capacity for extension.

Repair can be posed as a lifting problem: given $q : \mathcal{E} \rightarrow \mathcal{B}$ mapping fully realized continuation structures to their damaged or partial descriptions, and a damaged section $b \in \mathcal{B}(U)$, repair asks whether a lift $e \in \mathcal{E}(U)$ with $q(e) = b$ exists satisfying the required continuation constraints. Recoverable damage means an admissible lift exists; irrecoverable damage means none does within the permitted repair class – the standard mathematical form of “an admissible lift through a distorted bundle,” and the sheaf-theoretic counterpart of the TARTAN retiling repair of Section 26.6.

G.4 The Combined Object

The categorical and sheaf-theoretic accounts unify in a stack, or fibered category, over historical contexts, $\mathcal{C} \rightarrow \mathcal{H}$, with fiber $\mathcal{C}_p$ the distinction structures at p , transport induced by histories $\gamma : p \rightarrow q$, descent conditions determining whether transported structures assemble globally, and a connection supplying differential transport within this fibered structure. Within this formulation, distinction is an object in a fiber, continuation is transport along a history morphism, memory is path-dependence of transport, identity is equivalence of restricted transport representations, repair is a lift or natural transformation restoring descent, and ontological death is loss of admissible global extendability:

A persistent entity is a coherently extendable transport structure, considered up to the equivalences preserved

References

- [1] V. I. Arnold, *Mathematical Methods of Classical Mechanics*, Springer, New York, 1989.
- [2] M. F. Atiyah, *The Geometry and Physics of Knots*, Cambridge University Press, Cambridge, 1990.
- [3] J. C. Baez and J. Huerta, “An Invitation to Algebraic General Relativity,” *General Relativity and Gravitation*, 43, 2335–2392, 2011.
- [4] J. C. Baez and J. Dolan, “Higher-Dimensional Algebra and Topological Quantum Field Theory,” *Journal of Mathematical Physics*, 36, 6073–6105, 1995.
- [5] G. Bateson, *Steps to an Ecology of Mind*, Chandler Publishing, San Francisco, 1972.
- [6] G. Bateson, D. D. Jackson, J. Haley, and J. Weakland, “Toward a Theory of Schizophrenia,” *Behavioral Science*, 1(4), 251–264, 1956.
- [7] C. H. Bennett, “Logical Reversibility of Computation,” *IBM Journal of Research and Development*, 17, 525–532, 1973.
- [8] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, 2nd ed., Wiley-Interscience, Hoboken, 2006.
- [9] B. A. Dubrovin, A. T. Fomenko, and S. P. Novikov, *Modern Geometry—Methods and Applications*, Springer, New York, 1992.
- [10] S. Eilenberg and N. Steenrod, *Foundations of Algebraic Topology*, Princeton University Press, Princeton, 1952.
- [11] M. Flato, A. Lichnerowicz, and D. Sternheimer, “Deformation Theory and Quantization,” *Annals of Physics*, 111, 67–151, 1978.
- [12] J. W. Gibbs, *Elementary Principles in Statistical Mechanics*, Charles Scribner’s Sons, New York, 1902.
- [13] M. Gromov, *Metric Structures for Riemannian and Non-Riemannian Spaces*, Birkhäuser, Basel, 1999.
- [14] A. Hatcher, *Algebraic Topology*, Cambridge University Press, Cambridge, 2002.
- [15] P. Holmes, J. L. Lumley, G. Berkooz, and C. W. Rowley, *Turbulence, Coherent Structures, Dynamical Systems and Symmetry*, Cambridge University Press, Cambridge, 2012.
- [16] K. Jacobs, *Stochastic Processes for Physicists: Understanding Noisy Systems*, Cambridge University Press, Cambridge, 2010.
- [17] A. N. Kolmogorov, “Three Approaches to the Quantitative Definition of Information,” *Problems of Information Transmission*, 1, 1–7, 1965.
- [18] R. Landauer, “Irreversibility and Heat Generation in the Computing Process,” *IBM Journal of Research and Development*, 5, 183–191, 1961.

- [19] S. Mac Lane, *Categories for the Working Mathematician*, 2nd ed., Springer, New York, 1998.
- [20] B. B. Mandelbrot, *The Fractal Geometry of Nature*, W. H. Freeman, New York, 1982.
- [21] J. Nash, "Equilibrium Points in N-Person Games," *Proceedings of the National Academy of Sciences*, 36, 48–49, 1950.
- [22] E. Noether, "Invariante Variationsprobleme," *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen*, 235–257, 1918.
- [23] J. Pearl, *Causality: Models, Reasoning, and Inference*, 2nd ed., Cambridge University Press, Cambridge, 2009.
- [24] I. Prigogine and G. Nicolis, *Self-Organization in Nonequilibrium Systems*, Wiley, New York, 1977.
- [25] I. Prigogine, *From Being to Becoming: Time and Complexity in the Physical Sciences*, W. H. Freeman, San Francisco, 1980.
- [26] M. R. Rasetti, *Geometry of Nonlinear Differential Equations*, Springer, Berlin, 1986.
- [27] G. Ricci and T. Levi-Civita, "Méthodes de calcul différentiel absolu et leurs applications," *Mathematische Annalen*, 54, 125–201, 1900.
- [28] W. Rudin, *Functional Analysis*, 2nd ed., McGraw-Hill, New York, 1991.
- [29] D. A. Salamon, *Spin Geometry and Seiberg–Witten Invariants*, Springer, Berlin, 1999.
- [30] C. E. Shannon, "A Mathematical Theory of Communication," *Bell System Technical Journal*, 27, 379–423, 623–656, 1948.
- [31] S. Smale, "Differentiable Dynamical Systems," *Bulletin of the American Mathematical Society*, 73, 747–817, 1967.
- [32] S. H. Strogatz, *Nonlinear Dynamics and Chaos*, 2nd ed., Westview Press, Boulder, 2015.
- [33] J. von Neumann, *Mathematical Foundations of Quantum Mechanics*, Princeton University Press, Princeton, 1955.
- [34] R. M. Wald, *General Relativity*, University of Chicago Press, Chicago, 1984.
- [35] E. Witten, "Topological Quantum Field Theory," *Communications in Mathematical Physics*, 117, 353–386, 1988.
- [36] N. M. J. Woodhouse, *Geometric Quantization*, 2nd ed., Oxford University Press, Oxford, 1992.
- [37] N. Yoneda, "On the Homology Theory of Modules," *Journal of the Faculty of Science, University of Tokyo*, 7, 193–227, 1954.