Repair Categories and the Local Geometry of Continuation

Flyxion

July 2026 · Draft 0.1

Abstract

A continuation system carries two independent notions of proximity: a directed *repair contrast* $\mathsf{R}(C, C')$, measuring the cost of recognizing C' as a repaired form of C, and an *executable cost* $\mathsf{E}(C, C')$, the least total cost of an operational path from C to C'. We study the geometry these jointly induce, beginning from a fully explicit twelve-state system in which every quantity is computable by hand. The computation isolates a fork that was not visible at the level of analogy: if R satisfies the triangle inequality — that is, if repair space is a Lawvere-enriched category — then the Pythagorean projection defect onto any admissible subset is necessarily nonpositive, and coherence between comparison and execution is certifiable by a local check on generators. Positive projection interference, the signature phenomenon of dually flat information geometry, requires a divergence-type contrast and admits no local coherence certificate. We then pose and solve the finite characterization problem in the style of Chentsov. Monotonicity alone is maximally underdetermined: natural monotone invariants of the category separate every ordered pair of the model, detecting operational data — promotion versus phase, proximity to the singular stratum — that the enrichment forgets. Adding invariance under repair-sufficient morphisms collapses them exactly: sufficiency, provably scarce inside irreversible systems, can be manufactured between systems by amalgamation, and the invariants so characterized are precisely the monotone functions of the *bidirected* executable cost, the ordered pair of forward and reverse repair cost. Directedness survives characterization; operational-bracket data survives only monotone naturality. The smooth dual-geometric theory is presented as the tangent approximation to the enriched substrate, not its foundation.

Introduction

Most theories of continuation assume a fixed structure of comparison. An object changes; the framework for measuring change stands still. The work collected under the Admissibility Program has repeatedly encountered the opposite situation: admissibility conditions vary across a space of contexts, repair criteria drift, projection systems reorganize, and retained histories differ from point to point. In such a setting the naive difference between a continuation here and a continuation there conflates two things — change in the object and change in the frame — and the problem of separating them is, structurally, the problem that led differential geometry to the connection and the covariant derivative. The tempting conclusion is that a theory of admissibility should import that machinery. This paper argues, and in the finite case proves, that the inheritance runs the other way. The differential-geometric apparatus is a tangent approximation to something more primitive, procedural and enriched-categorical in kind; and when one asks — in the manner of Chentsov — which quantities are actually *characterized* by the invariances a continuation system supports, the answer is not a connection at all.

Every continuation system carries two proximities, and they are this paper's primitives. The *executable cost* $\mathsf{E}(C, C')$ is the least total cost of a sequence of available operations transforming C into C'; it is generated by the system's own operational algebra, and its triangle inequality is nothing but the composition of procedures — in Lawvere's sense, E is the free enrichment of the operational graph. The *repair contrast* $\mathsf{R}(C, C')$ is a directed comparison defined independently of the operations: the cost of recognizing C' as a repaired form of C. Nothing forces these to agree. Some states are easy to compare but difficult to reach; some are easy to reach but expensive to repair. The discrepancy $\mathsf{G} = \mathsf{E} - \mathsf{R}$ between comparison and execution is the central object of study.

Everything in the paper answers to a single worked example: a twelve-state system with two operations, one of them irreversible, small enough that every claim reduces to a table. The example was built to discipline the abstractions, and it did more than that; each of the paper's results was found in its tables before it was proved. Three findings organize the paper. First, a *fork*: if R satisfies the triangle inequality — if repair space is itself an enriched category — then the Pythagorean projection defect onto any admissible subset is necessarily nonpositive, and coherence between comparison and execution ($\mathsf{R} \leq \mathsf{E}$) is certifiable by a local check on generators. Positive projection interference, the signature of dually flat information geometry, is possible only for divergence-type contrasts, which admit no such local certificate. The enriched and divergence regimes of repair are separated by a computable sign. Second, a *profusion*: posing the finite characterization problem — which functionals of pairs of states are monotone under all admissibility morphisms — the answer is that monotonicity alone constrains almost nothing. The natural monotone invariants separate every ordered pair of the model, and what they detect is operational data the enrichment forgets: which generator does the work of a geodesic, and how close it runs to the singular stratum where the operations cease to commute. Third, a *collapse*, which is the paper's main theorem: adding invariance under repair-sufficient morphisms — the analogue of Chentsov's sufficient statistics — cuts the profusion down exactly. Sufficient morphisms are provably scarce inside an irreversible system, but they can be *manufactured* between systems by amalgamation, gluing two systems along a pair of states with matching costs; and the invariants that survive are precisely the monotone functions of the *bidirected* executable cost, the ordered pair $(\mathsf{E}(x, y), \mathsf{E}(y, x))$. What survives every distinction-preserving redescription of a continuation system is two numbers per ordered pair: the cost of getting there, and the cost of getting back.

The theorem stands to Chentsov's as irreversible systems stand to statistical ones, with the hinge inverted: where the classical proof leans on the abundance of sufficient statistics, the present proof leans on a construction, because abundance fails. And the characterized invariant is not a symmetric metric but a directed pair; directedness survives characterization, while the operational-bracket data — the torsion-like, order-sensitive structure — survives only the weaker discipline of monotone naturality. The differential-geometric objects that motivated the inquiry reappear at the end, in their proper place: on the simplex relaxation of a finite system, the metric and dual connections of information geometry arise as the second- and third-order jets of a smooth contrast, and the paper's closing conjecture is that the smooth characterization, when it exists, will pin the germ of the bidirected cost — the dual geometry emerging as the jet expansion of the characterized invariant rather than as a structure selected by fiat.

Section 1 presents the twelve-state system and the fork. Section 2 develops the enriched-categorical foundation. Section 3 poses and solves the finite characterization problem. Section 4 constructs the smooth tangent theory and states what remains open.

1 A Twelve-State Continuation System

We begin from something almost embarrassingly concrete: an example small enough that every claim in this paper reduces to a table.

Definition 1.1 (The system). *Let* $\mathcal{C} = \{0,1,2\} \times \{0,1,2,3\}$*, a state* $\mathsf{C} = (\ell, p)$ *consisting of a* level ℓ *and a cyclic* phase p*. The operational algebra* $\mathcal{O}$ *is generated by two total maps:*

$$a(\ell,p) = (\ell,\, p+1 \bmod 4), \qquad b(\ell,p) = \begin{cases} (\ell+1,\, p) & \ell < 2, \\ (2,\, 0) & \ell = 2, \end{cases}$$

with state-independent costs $c(a) = 1$ *and* $c(b) = 2$*. No operation decreases* ℓ*: promotion is irreversible, and saturation of* b *at the top level resets the phase.*

Definition 1.2 (Executable cost). $\mathsf{E}(C, C')$ *is the infimum of* $\sum_i c(o_i)$ *over finite words* $o_n \cdots o_1$ *in* $\{a, b\}$ *with* $o_n \cdots o_1(C) = C'$*, and* $+\infty$ *if no such word exists.*

Metric geometry begins from distance. Repair geometry begins from reconstruction.

Definition 1.3 (Repair contrast). *Independently of* $\mathcal{O}$*, define the directed contrast*

$$\mathsf{R}((\ell,p),(\ell',p')) = \max(\ell'-\ell, 0) + 5\max(\ell-\ell', 0) + \tfrac{1}{2}\, d_{\mathrm{cyc}}(p,p'),$$

where d_{cyc} *is cyclic distance on* $\{0,1,2,3\}$*. Recognizing an upgrade is cheap; recognizing a downgrade is expensive but finite; phase mismatch costs* $\frac{1}{2}$ *per cyclic step. Note that* R *satisfies the triangle inequality: it is a quasi-metric, hence a Lawvere enrichment of* $\mathcal{C}$*.*

Three quantities are computed exhaustively over all 132 ordered pairs (source and full computation in the appendix script).

The comparison–execution gap. Define $\mathsf{G} = \mathsf{E} - \mathsf{R} \in [0, \infty]$. Of the 132 pairs, 48 have $\mathsf{G} = \infty$: every downward repair is *recognizable* at finite contrast (5 per level) but *inexecutable*, since no operation lowers ℓ. Among finite pairs the largest gaps occur where R exploits cyclic symmetry that $\mathcal{O}$ cannot: for instance $\mathsf{R}((0,0),(2,3)) = 2.5$ while $\mathsf{E}((0,0),(2,3)) = 7$, because the phase must be cranked forward three times and two promotions paid. Representational nearness and operational accessibility dissociate, and the dissociation is quantitative.

Lemma 1.4 (Local-to-global coherence). *Suppose* R *satisfies the triangle inequality and is* generator-dominated*:* $\mathsf{R}(C, o(C)) \le c(o; C)$ *for every generator* $o \in \mathcal{O}$ *and every* $C \in \mathcal{C}$*. Then* $\mathsf{R} \le \mathsf{E}$*, i.e.* $\mathsf{G} \ge 0$*, on all of* $\mathcal{C} \times \mathcal{C}$*.*

Proof. For any executable word $o_n \cdots o_1$ from C to C' with intermediate states C_i, the triangle inequality gives $\mathsf{R}(C, C') \le \sum_i \mathsf{R}(C_{i-1}, C_i) \le \sum_i c(o_i; C_{i-1})$. Take the infimum over words. □

In the model, generator domination holds with room to spare (a moves R by at most $\frac{1}{2} \le 1$; b by at most $1 \le 2$), so $\mathsf{G} \ge 0$ is a theorem of the system rather than an empirical observation. The lemma fails without the triangle inequality; we return to this below.

Localized operational torsion. The generators commute everywhere except on the saturation stratum: for every p,

$$a \circ b\,(2,p) = (2,1), \qquad b \circ a\,(2,p) = (2,0),$$

while $a \circ b = b \circ a$ on $\ell < 2$. Order-dependence is not diffuse; it is concentrated exactly where b changes character from injective promotion to non-injective collapse.

Proposition 1.5 (Localized torsion witness)**.** *The commutator defect of $\mathcal{O}$ is supported on the singular stratum $\{\ell = 2\}$ of the operational algebra.*

The projection defect and its sign. For an admissible subset $\mathcal{A} \subseteq \mathcal{C}$ with projection $C^* = \arg\min_{t \in \mathcal{A}} \mathsf{R}(C,t)$, define

$$\Delta_{\mathcal{A}}(C, C') = \mathsf{R}(C, C') - \mathsf{R}(C, C^*) - \mathsf{R}(C^*, C'), \qquad C' \in \mathcal{A}.$$

With the level-aligned set $\mathcal{A} = \{(1,p) : p\}$ we find $\Delta_{\mathcal{A}} \equiv 0$: the contrast is separable and $\mathcal{A}$ is autoparallel with respect to the product structure. With the diagonal set $\mathcal{A} = \{(0,0), (1,2), (2,1)\}$, fourteen of thirty-six pairs give $\Delta_{\mathcal{A}} \neq 0$, with range $[-2, 0]$. The defect never becomes positive, and this is not an accident of the example.

Proposition 1.6 (Sign restriction in enriched repair)**.** *If R satisfies the triangle inequality, then $\Delta_{\mathcal{A}}(C, C') \leq 0$ for every $\mathcal{A}$, every projection point, and every $C' \in \mathcal{A}$. Consequently*

$$\Delta_{\mathcal{A}} > 0 \implies \mathsf{R} \textit{ is not a Lawvere enrichment.}$$

Proof. Immediate from $\mathsf{R}(C, C') \leq \mathsf{R}(C, C^*) + \mathsf{R}(C^*, C')$. □

The fork

Positive Pythagorean interference — the signature of dually flat information geometry, where the generalized Pythagorean theorem yields exact positive decompositions — is *impossible* in enriched repair. At the same location, Lemma 1.4 shows that enriched repair admits a *local certificate* of coherence ($\mathsf{R} \leq \mathsf{E}$ verified generator by generator), whereas a divergence-type contrast, lacking the triangle inequality, admits none: its coherence with execution must be imposed globally. Repair structures therefore bifurcate into an *enriched regime* (directed metric, locally certifiable, projection can only lose) and a *divergence regime* (contrast function, dual connections, positive interference possible, coherence a global axiom). The remainder of this paper develops the first regime rigorously and situates the second as its tangent theory.

2 Repair as a Lawvere-Enriched Category

The twelve-state system used no structure beyond a set of states, a costed operational graph, and a directed contrast. This section fixes that data abstractly and shows that the executable cost E is not an auxiliary construction but the free object of the theory. Throughout, costs are valued in the quantale $\mathbb{V} = ([0,\infty], \geq, +)$, the base of enrichment in Lawvere's generalized metric spaces [1]: objects of a $\mathbb{V}$-category are points, the hom from x to y is a cost $d(x,y) \in [0,\infty]$ with $d(x,x) = 0$, and the composition law *is* the triangle inequality $d(x,z) \leq d(x,y) + d(y,z)$.

Definition 2.1 (Continuation system)**.** *A* finite continuation system *is a triple* $(\mathcal{C}, \mathcal{O}, c)$ *where* $\mathcal{C}$ *is a finite set of states,* $\mathcal{O}$ *is a set of generating operations* $o\colon \mathrm{dom}(o) \subseteq \mathcal{C} \to \mathcal{C}$*, and* $c(o; C) \in (0, \infty)$ *assigns each application its cost. The* operational graph *has vertex set* $\mathcal{C}$ *and a directed edge* $C \to o(C)$ *of weight* $c(o; C)$ *for each applicable generator.*

Definition 2.2 (Executable enrichment)**.** *The* executable cost $\mathsf{E}(C, C')$ *is the infimum of* $\sum_i c(o_i; C_{i-1})$ *over finite operational words from C to* C'*, with* $\mathsf{E}(C, C) = 0$ *(the empty word) and* $\mathsf{E}(C, C') = \infty$ *when no word exists.*

E satisfies the triangle inequality by concatenation of words, so $(\mathcal{C}, \mathsf{E})$ is a $\mathbb{V}$-category. It is, moreover, the *free* one on the operational graph: enrichment is not an assumption about repair but a consequence of composition.

Definition 2.3 (Admissibility morphism)**.** *A map* $F\colon (\mathcal{C}_1, \mathcal{O}_1, c_1) \to (\mathcal{C}_2, \mathcal{O}_2, c_2)$ *between continuation systems is an* admissibility morphism *if there is an assignment* $o \mapsto F_*(o)$ *of generators to (possibly empty) operational words such that*

$$F(o(C)) = F_*(o)(F(C)) \quad \textit{and} \quad c_2(F_*(o); F(C)) \leq c_1(o; C)$$

for every applicable pair, where c_2 *of a word is the sum of its step costs. Intertwining says F respects what can be done; cost-nonincrease says F never makes an operation more expensive. Coarse-grainings, quotients by zero-cost equivalences, and forgetful projections are all of this form.*

Lemma 2.4 (Nonexpansiveness)**.** *Every admissibility morphism satisfies* $\mathsf{E}_2(F(C), F(C')) \leq \mathsf{E}_1(C, C')$*.*

Proof. The F_*-image of an operational word from C to C' is an operational word from $F(C)$ to $F(C')$ of no greater total cost; take infima. □

In enriched language, Lemma 2.4 says every admissibility morphism is a $\mathbb{V}$-functor $(\mathcal{C}_1, \mathsf{E}_1) \to (\mathcal{C}_2, \mathsf{E}_2)$. The next proposition locates E uniquely among all contrasts compatible with the generators; it is the abstract form of Lemma 1.4.

Proposition 2.5 (Maximality of the free enrichment)**.** *Let* $d\colon \mathcal{C} \times \mathcal{C} \to [0, \infty]$ *satisfy* $d(C, C) = 0$*, the triangle inequality, and generator domination* $d(C, o(C)) \leq c(o; C)$*. Then* $d \leq \mathsf{E}$ *pointwise. Hence* E *is the greatest generator-dominated enrichment on* $(\mathcal{C}, \mathcal{O}, c)$*, and* $\mathsf{G} = \mathsf{E} - \mathsf{R} \geq 0$ *holds for* every *enriched, generator-dominated repair contrast* R *— Lemma 1.4 is its instance.*

Proof. Chain d along any word and take infima, as in Lemma 1.4. Freeness in the categorical sense: E is the left Kan extension of the weighted operational graph along the inclusion of generators into words, i.e. the value of the left adjoint to the forgetful functor from $\mathbb{V}$-categories to costed directed graphs. □

Two consequences follow. First, the comparison–execution gap G acquires an intrinsic meaning: for enriched contrasts it measures the failure of R to be *free*, i.e. how far the system's abstract sense of nearness falls below what its own operations generate. Second, the sign restriction of Proposition 1.6 is now visible as a fact about the base of enrichment: any $\mathbb{V}$-category satisfies $\Delta_{\mathcal{A}} \leq 0$, so positive interference marks a contrast that is not a hom-structure over $([0, \infty], \geq, +)$ at all.

Definition 2.6 (Repair sufficiency)**.** *An admissibility morphism F is* repair-sufficient *for a contrast* R *if* $\mathsf{R}_2(F(C), F(C')) = \mathsf{R}_1(C, C')$ *for all* C, C'*.*

Repair-sufficient morphisms coarse-grain without erasing any repair-relevant distinction; they play the role that sufficient statistics play in Chentsov's theorem [2] and that deterministic sufficiency plays in Markov categories [4]. The characterization problem of the next section asks which functionals are monotone under all admissibility morphisms and invariant precisely under the repair-sufficient ones.

Remark 2.7 (Where the divergence regime lives). Nothing in Definitions 2.1–2.6 requires R to be triangular. A divergence-type contrast simply fails to be a $\mathbb{V}$-category structure; it is a two-point function on the objects of one. The smooth theory of Section 4 will treat such contrasts as functions on $\Delta(\mathbb{C})$ whose diagonal jets induce geometry, with the enriched structure persisting underneath as the executable skeleton.

Remark 2.8 (Specialization to executable-cost sufficiency). Definition 2.6 defines repair-sufficiency for an arbitrary contrast R, and nothing said so far ties R to E. From this point on, the paper studies sufficiency for the particular contrast $\mathsf{R} = \mathsf{E}$: condition (I) in Problem 3.1 below, and every theorem that follows from it, characterizes invariance under E-*sufficient* morphisms, not under morphisms sufficient for an independently specified R. This is a deliberate narrowing, not an oversight, but it changes what is being characterized. Theorem 3.10 answers the question "what survives redescription that preserves executable cost exactly," not "what survives redescription that preserves the recognition contrast R." The two questions coincide only when $\mathsf{R} = \mathsf{E}$, which Definition 2.1 does not assume of a general continuation system. The joint invariant theory of the pair (R, E) — sufficiency with respect to both primitives simultaneously — is left open; it is discussed again in Section 5 and in the closing remarks.

3 Characterization in the Finite Case

Chentsov's theorem [2] characterizes the Fisher metric and the α-connections not by construction but by variance: they are the structures monotone under all Markov morphisms, invariant exactly under sufficient statistics. The finite analogue for continuation systems replaces Markov kernels by admissibility morphisms and asks the corresponding question.

Problem 3.1 (Finite characterization)**.** *Determine all assignments Φ, giving each finite continuation system $(\mathbb{C}, \mathbb{O}, c)$ a functional $\Phi_{\mathbb{C}}\colon \mathbb{C} \times \mathbb{C} \to [0, \infty]$, such that*

(*M*) $\Phi_{\mathbb{C}_2}(F(C), F(C')) \leq \Phi_{\mathbb{C}_1}(C, C')$ *for every admissibility morphism F, and*

(*I*) *equality holds in* (*M*) *whenever F is repair-sufficient for* E (*Definition 2.6 with* $\mathsf{R} = \mathsf{E}$).

Two partial results bound the answer from below and constrain it from above.

Proposition 3.2 (The executable cone)**.** *For every nondecreasing $f\colon [0, \infty] \to [0, \infty]$ with $f(0) = 0$, the assignment $\Phi = f \circ \mathsf{E}$ satisfies* (*M*) *and* (*I*)*. The class of solutions is closed under pointwise suprema, infima, and nondecreasing reparametrization, and therefore contains the closed cone generated by* E.

Proof. (M) is Lemma 2.4 composed with monotone f; (I) is immediate from the equality defining E-sufficiency. Closure properties follow pointwise. □

Proposition 3.3 (Rigidity)**.** *Any solution Φ of Problem 3.1 is determined by its values on operational graphs, is invariant under all cost-preserving system isomorphisms, and satisfies $\Phi \leq f \circ \mathsf{E}$ for $f(t) = \sup\{\Phi_{I_t}(x, y)\}$, where I_t is the two-state system with a single operation of cost t. In particular Φ vanishes on the diagonal and is infinite only where E is infinite, provided f is finite on finite arguments.*

Proof sketch. Isomorphisms are invertible admissibility morphisms whose inverses are also admissibility morphisms, so (M) applied both ways gives invariance. For the bound: given states C, C' with $\mathsf{E}(C, C') = t < \infty$, any operational geodesic word induces an admissibility morphism from a chain system into $(\mathcal{C}, \mathcal{O}, c)$ hitting C, C', and a morphism from the chain onto I_t; monotonicity along both pins $\Phi_{\mathcal{C}}(C, C')$ against Φ_{I_t}. Where $\mathsf{E}(C, C') = \infty$ no such comparison exists and Φ is unconstrained from above. □

The gap between Propositions 3.2 and 3.3 is the substance of the problem: whether monotonicity plus invariance force Φ to factor through E *exactly*. On the twelve-state system this can be settled by exhaustive computation, and the answer is negative.

Theorem 3.4 (Endomorphism constraints do not force factorization)**.** *The system of Section 1 admits exactly* 22 *admissibility endomorphisms, of which the identity is the only* E*-sufficient one. The constraint digraph they generate on the* 132 *ordered state pairs has* 132 *singleton equality classes, and there exists an explicit functional — the chain depth of a pair in the condensation order of the constraints — that satisfies* (M) *and* (I) *with respect to every endomorphism yet takes five distinct values on the level set* $\{\mathsf{E} = 2\}$ *alone.*

Proof. Computational; script `characterize.py` in the appendix. Cost-nonincrease bounds each generator image to a word of cost at most the generator's cost, and reachability of every state from $(0, 0)$ makes an endomorphism determined by the triple $(F(0, 0), F_*(a), F_*(b))$; the 96 candidates are checked for loop closure, yielding 22. Sufficiency, the constraint digraph, its strongly connected components, and the verification of (M) and (I) for the chain-depth functional are then finite checks. □

The separating witnesses are structurally meaningful. Within $\{\mathsf{E} = 2\}$ the functional distinguishes a two-step phase adjustment $(0, 0) \to (0, 2)$ from a single promotion $(0, 0) \to (1, 0)$ and from a promotion adjacent to the singular stratum $(1, 0) \to (2, 0)$: it is sensitive to the *operational type* of a geodesic, precisely the data that E forgets and the commutator defect sees.

Theorem 3.4 uses only endomorphisms; Problem 3.1 quantifies over all morphisms between all systems. Two further results show that enlarging the constraint class does not rescue factorization.

Proposition 3.5 (Factorization of forced constraints)**.** *Let X be a finite continuation system. Any chain of monotonicity constraints leading from an ordered pair of X through pairs of other systems and back to an ordered pair of X is implied by the constraint of a single endomorphism of X. Consequently the equalities forced on pairs of X by naturality over the entire category coincide with those forced by* $\mathrm{End}(X)$*; for the twelve-state system, none.*

Proof. Each constraint edge is induced by a morphism applied to a pair, and admissibility morphisms compose (intertwining assignments compose word-wise; cost-nonincrease is transitive). A constraint chain $X \to Y_1 \to \cdots \to X$ is therefore induced by the composite morphism, an endomorphism of X. The claim about the model is Theorem 3.4. □

Theorem 3.6 (A natural separating family: co-tests)**.** *For a finite continuation system T with a distinguished ordered pair (t, t'), define on every system X*

$$\Phi_X^{T,t,t'}(x, y) = \begin{cases} 0 & \text{if some admissibility morphism } F\colon T \to X \text{ has } F(t) = x,\ F(t') = y, \\ 1 & \text{otherwise.} \end{cases}$$

Then $\Phi^{T,t,t'}$ satisfies (M) over the entire category. On the twelve-state system, taking T to be the system itself: the co-test based at $((0,0),(1,0))$ takes the value 0 exactly on the promotion-type and saturation-collapse pairs and 1 on all phase-type pairs, separating the level sets $\{\mathsf{E}=1\}$ and $\{\mathsf{E}=2\}$; it satisfies (I) with respect to every E-sufficient endomorphism; and the 132 co-tests based at the model's own pairs are pairwise distinct as functionals on it, so the natural monotone invariants of the category separate every *ordered pair of the model.*

Proof. Monotonicity: if $F\colon X \to Y$ and $\Phi_X(x,y) = 0$ via $H\colon T \to X$, then $F \circ H$ witnesses $\Phi_Y(F(x),F(y)) = 0$; the case $\Phi_X = 1$ is vacuous. The statements about the model are finite checks (`crosstest.py`): morphisms $T = \mathcal{C} \to \mathcal{C}$ are the 22 endomorphisms, so all values are computable. □

The co-tests show that monotonicity alone separates every ordered pair: the natural (M)-invariants of the category are strictly finer than the enrichment and detect operational data — promotion versus phase, proximity to the singular stratum — that E forgets. The question left open by Theorems 3.4 and 3.6 was whether condition (I), over E-sufficient morphisms *between* systems, cuts this profusion back. It does — completely — and the mechanism is that sufficiency, scarce inside any one irreversible system, can be *manufactured* between systems by gluing.

Lemma 3.7 (Sufficiency is isometric embedding)**.** *Every E-sufficient admissibility morphism is injective and isometric for E. (Costs are strictly positive, so $\mathsf{E}(C,C') = 0$ iff $C = C'$.)*

Lemma 3.8 (Amalgamation)**.** *Let X, X' be finite continuation systems with ordered pairs (x,y), (x',y') such that $\mathsf{E}_X(x,y) = \mathsf{E}_{X'}(x',y')$ and $\mathsf{E}_X(y,x) = \mathsf{E}_{X'}(y',x')$. Let $Y = X \sqcup X'/(x \sim x', y \sim y')$, where the generator sets $\mathcal{O}_X, \mathcal{O}_{X'}$ are taken as disjoint copies before forming the quotient — so that two generators of the same name or the same action on states are still treated as distinct generators of Y unless explicitly identified — with the union of these disjoint copies as the generator set of Y. Then both inclusions $X \hookrightarrow Y$ and $X' \hookrightarrow Y$ are E-sufficient admissibility morphisms. If instead only the inequalities $\mathsf{E}_{X'}(x',y') \geq \mathsf{E}_X(x,y)$ and $\mathsf{E}_{X'}(y',x') \geq \mathsf{E}_X(y,x)$ hold, the inclusion of X alone is E-sufficient (the one-sided form).*

The proof requires knowing exactly how a Y-path can interleave the two components. The next lemma settles this before any cost estimate is made, so that the induction below has no hidden assumptions about how many times a path may cross between X and X'.

Lemma 3.9 (Path decomposition across a two-point gluing)**.** *Write $g \in \{x,y\}$ for the two gate points, identified with $\{x',y'\}$ in Y, and let $w = o_n \cdots o_1$ be any operational word in Y from a to b with $a, b \in \mathcal{C}_X$, acting on states $C_0 = a, C_1, \ldots, C_n = b$. Since $\mathcal{O}_X$ and $\mathcal{O}_{X'}$ are disjoint (the generator sets are defined on the disjoint union before gluing), each step o_i lies in exactly one of the two sets; grouping consecutive steps of the same type partitions $\{1,\ldots,n\}$ into an alternating sequence of maximal X-*segments *and* excursions*. Every excursion begins and ends at a gate point.*

Proof. Consider a maximal run of $\mathcal{O}_{X'}$-steps at positions $i,\ldots,j$. If $i > 1$, the preceding step $o_{i-1} \in \mathcal{O}_X$ places $C_{i-1} \in \mathcal{C}_X$; if $i = 1$, $C_{i-1} = a \in \mathcal{C}_X$ by hypothesis. Either way $C_{i-1} \in \mathcal{C}_X$. But $o_i \in \mathcal{O}_{X'}$ requires $C_{i-1} \in \mathcal{C}_{X'}$, its domain. Since $\mathcal{C}_X \cap \mathcal{C}_{X'} = \{x,y\}$ after gluing, $C_{i-1} \in \{x,y\}$. Symmetrically, C_j is the codomain of $o_j \in \mathcal{O}_{X'}$, so $C_j \in \mathcal{C}_{X'}$; if $j < n$ the following step $o_{j+1} \in \mathcal{O}_X$ forces $C_j \in \mathcal{C}_X$, and if $j = n$ then $C_j = b \in \mathcal{C}_X$ by hypothesis. Either way $C_j \in \{x,y\}$. Hence every excursion runs from a gate point to a gate point, regardless of how many excursions w contains or in what order they occur. □

Proof of Lemma 3.8, two-sided case. Fix $a, b \in \mathcal{C}_X$ and let w be any operational word in Y from a to b. By Lemma 3.9, w decomposes into X-segments separated by $k \geq 0$ excursions $e_1, \ldots, e_k$, each running between gate points. We argue by induction on k that $\mathrm{cost}(w) \geq \mathsf{E}_X(a,b)$.

Base case $k = 0$. Then w is entirely an X-word from a to b, so $\mathrm{cost}(w) \geq \mathsf{E}_X(a,b)$ directly from Definition 2.2.

Inductive step. Let e_1 run from gate g to gate g' (each in $\{x, y\}$), with cost $c(e_1) \geq \mathsf{E}_{X'}(g, g')$, viewing g, g' as the corresponding points of X'. By hypothesis $\mathsf{E}_{X'}(x', y') = \mathsf{E}_X(x, y)$, $\mathsf{E}_{X'}(y', x') = \mathsf{E}_X(y, x)$, and trivially $\mathsf{E}_{X'}(x', x') = \mathsf{E}_X(x, x) = 0 = \mathsf{E}_{X'}(y', y') = \mathsf{E}_X(y, y)$; in every case $\mathsf{E}_{X'}(g, g') = \mathsf{E}_X(g, g')$. Since $\mathcal{C}$ is finite and every cost is strictly positive, an infimizing X-word from g to g' may be taken simple (non-repeating), hence exists with cost exactly $\mathsf{E}_X(g, g') \leq c(e_1)$; call it u_1. Replace e_1 by u_1 in w to obtain w': the states immediately before and after e_1 are g and g', unchanged by the substitution, so w' is a well-formed word from a to b with $k - 1$ excursions and $\mathrm{cost}(w') \leq \mathrm{cost}(w)$. By the inductive hypothesis, $\mathrm{cost}(w') \geq \mathsf{E}_X(a,b)$, so $\mathrm{cost}(w) \geq \mathrm{cost}(w') \geq \mathsf{E}_X(a,b)$.

Since w was arbitrary, $\mathsf{E}_Y(a,b) = \inf_w \mathrm{cost}(w) \geq \mathsf{E}_X(a,b)$. The reverse inequality $\mathsf{E}_Y(a,b) \leq \mathsf{E}_X(a,b)$ holds because every X-word is a Y-word. Hence $\mathsf{E}_Y(a,b) = \mathsf{E}_X(a,b)$ for *every* pair $a, b \in \mathcal{C}_X$, not merely (x, y): the inclusion $X \hookrightarrow Y$ is E-sufficient. Exchanging the roles of X and X' gives the same conclusion for $X' \hookrightarrow Y$. $\square$

Proof of Lemma 3.8, one-sided case. The induction is unchanged for words with endpoints in X: the excursion bound $c(e_\ell) \geq \mathsf{E}_{X'}(g_\ell, g'_\ell) \geq \mathsf{E}_X(g_\ell, g'_\ell)$ still licenses replacing e_ℓ by an X-word of no greater cost, so $X \hookrightarrow Y$ remains E-sufficient. The symmetric statement for X' need not hold: if X offers a strictly cheaper route between the gate points than X' does internally, an X'-pair may acquire a shortcut through X, and $X' \hookrightarrow Y$ can fail to be sufficient. $\square$

Theorem 3.10 (Characterization by the bidirected cost)**.** *An assignment Φ, giving each finite continuation system a functional on ordered pairs with $\Phi(x, x) = 0$, satisfies (M) and (I) over the category of finite continuation systems if and only if there is a function f on cost pairs, nondecreasing in each argument with $f(0, 0) = 0$, such that*

$$\Phi_X(x, y) \;=\; f(\mathsf{E}_X(x, y),\, \mathsf{E}_X(y, x))$$

for every system X and every pair.

Proof. Sufficiency of the form. Lemma 2.4 applied to both orderings gives (M) for nondecreasing f; an E-sufficient morphism preserves both coordinates, giving (I); pairs collapsing to the diagonal give $\Phi = 0 \leq$ anything.

Necessity. Let $V_{s,t}$ denote the two-state system with a forward generator of cost s and a reverse generator of cost t (each omitted when infinite), so $\mathsf{E}_V(v_0, v_1) = s$, $\mathsf{E}_V(v_1, v_0) = t$. Given any pair (x, y) in any X with cost pair (s, t), amalgamate X with $V_{s,t}$ along $(x, y) \sim (v_0, v_1)$; by Lemma 3.8 both inclusions are sufficient, so (I) twice gives $\Phi_X(x, y) = \Phi_Y(\text{glued pair}) = \Phi_{V_{s,t}}(v_0, v_1) =: f(s, t)$. Thus Φ factors through the cost pair. Monotonicity of f: for $s' \leq s$, $t' \leq t$ the identity on states with generator reassignment defines an admissibility morphism $V_{s,t} \to V_{s',t'}$ (a generator maps to one of no greater cost, or to a generator absent in the source there is nothing to assign), and (M) gives $f(s', t') \leq f(s, t)$. $\square$

Corollary 3.11 (Fate of the co-tests)**.** *No co-test invariant that fails to factor through the bidirected cost satisfies (I) over the category. Explicitly, for $\Phi^{\mathcal{C},(0,0),(1,0)}$: gluing a second copy of the model at $((0,0),(1,0))$ onto the pair $((0,0),(0,2))$ — boundary costs $(2,\infty)$ against $(2,2)$, so the base inclusion* *is sufficient by the one-sided form of Lemma 3.8 — makes the second copy's inclusion a cover of the image pair, forcing the value* 0 *where the base system has* 1 (*verified in* `amalgam.py`).

Theorem 3.10 is a finite, continuation-theoretic analogue of the *structure* of Chentsov's characterization argument, with two departures that keep the analogy from running deeper than it should. What is proved here is a characterization under a chosen notion of admissibility morphism and a specifically constructed notion of sufficiency, on a finite state space; it is Chentsov-like in shape, not a finite instance of the same depth of result, since the classical theorem characterizes the Fisher metric among all smooth monotone metrics on an infinite-dimensional statistical manifold, a generality this paper does not attempt. First, the role that the *abundance* of sufficient statistics plays in the classical proof is played here by a *construction*: sufficient morphisms are rare inside irreversible systems (Theorem 3.4) but can always be manufactured between systems by amalgamation, and the characterization rests entirely on the manufactured ones. Second, the invariant characterized is not a symmetric metric but the *bidirected* executable cost — the ordered pair of forward and reverse repair cost. Directedness survives characterization; operational-bracket data does not. The torsion-sensitive invariants of Theorem 3.6 are thereby given an exact status: they are natural under monotone comparison but not under lossless recoding. What sufficiency-invariance preserves is what survives every distinction-preserving change of description, and the theorem says this is precisely two numbers per ordered pair: how much it costs to get there, and how much it costs to get back. The two-primitive architecture accordingly resolves as a stratification rather than a theorem or a refutation: the (M)-natural theory sees the pair (enrichment, operational structure); the (M)+(I)-natural theory sees the bidirected enrichment alone. The precedent of Markov categories [4] suggests how the argument should be organized synthetically; the infinite-dimensional Chentsov theorem [5] remains the caution for the continuous case.

4 Smooth Limits and Dual Geometry

The enriched theory is native to the finite, procedural setting. The differential-geometric apparatus enters only through a relaxation, and we present it as such: a tangent theory of the regular stratum, not a foundation.

4.1 The simplex relaxation

Embed $\mathbb{C} \hookrightarrow \Delta(\mathbb{C})$, the simplex of probability weights on states, sending C to the vertex δ_C. Relaxed states are mixtures $\pi \in \Delta(\mathbb{C})$; generators extend to affine maps $o_*\pi = \pi \circ o^{-1}$ (pushforward), and a directed contrast extends to a smooth two-point function $\mathcal{R}\colon \Delta(\mathbb{C}) \times \Delta(\mathbb{C}) \to [0, \infty]$ agreeing with R on vertices. This is the same move Chentsov makes: the finite object is primary, the geometry lives on distributions over it.

4.2 Induced dual geometry

Where $\mathcal{R}$ is smooth with nondegenerate diagonal Hessian, the Eguchi construction [3] yields, in local coordinates on $\Delta(\mathbb{C})$,

$$g_{ab} = -\left.\frac{\partial^2 \mathcal{R}(\pi, \pi')}{\partial \pi^a\, \partial \pi'^b}\right|_{\pi'=\pi}, \qquad \Gamma_{abc} = -\left.\frac{\partial^3 \mathcal{R}(\pi, \pi')}{\partial \pi^a \partial \pi^b\, \partial \pi'^c}\right|_{\pi'=\pi},$$

with the dual Γ^* obtained by exchanging primed and unprimed arguments. The metric is symmetric regardless of the asymmetry of $\mathcal{R}$; the directedness of repair survives in the cubic tensor $T = \Gamma^* - \Gamma$ and the dual pair (∇, ∇^*), jointly metric-compatible: $Xg(Y,Z) = g(\nabla_X Y, Z) + g(Y, \nabla^*_X Z)$. Both induced connections are torsion-free by construction. Directed repair asymmetry and procedural order-dependence are therefore distinct departures from Riemannian geometry, and only the first is visible to the contrast function.

4.3 Torsion as representation defect

Order-dependence enters through the operational algebra. Suppose the relaxed generators near the identity are generated by vector fields X_a on $\Delta(\mathcal{C})$, with operational bracket $[X_a, X_b]_{\mathcal{O}}$ determined by the composition law. For any candidate connection ∇, define

$$\mathrm{Tor}(X_a, X_b) = \nabla_{X_a} X_b - \nabla_{X_b} X_a - [X_a, X_b]_{\mathcal{O}}.$$

Torsion so defined does not record noncommutativity as such; it records the failure of geometric transport to represent the operational bracket. Three cases separate: commuting operations with torsion-free geometry; noncommuting operations exactly represented, again torsion-free; and genuine mismatch. The twelve-state system suggests where mismatch lives:

Conjecture 4.1 (Singular support)**.** $\mathrm{supp}(\mathrm{Tor})$ *is contained in the operational singular set — the closure of the locus where some generator fails to be injective or applicable.*

In the model this is the saturation stratum $\{\ell = 2\}$, on which the commutator defect of (a, b) is entirely concentrated.

4.4 Decomposing the projection defect

Section 1 exhibited both behaviours of the defect $\Delta_{\mathcal{A}}$: identically zero for a level-aligned admissible set, strictly negative on a diagonal one. In the smooth regime the generalized Pythagorean theorem holds when the ambient geometry is dually flat *and* the admissible set is autoparallel in the relevant dual connection; a nonzero defect therefore has two independent sources, ambient non-flatness and submanifold curvature, plus their interaction. Any diagnostic use of $\Delta_{\mathcal{A}}$ must decompose it accordingly. The enriched regime adds the sign restriction of Proposition 1.6 on top: for triangular $\mathcal{R}$ the entire positive half of the phenomenon is excluded, so dually flat positive decompositions certify that the relaxed contrast has left the Lawvere regime. The boundary between the regimes is crossed exactly when the relaxation destroys the triangle inequality — as it does, for instance, when $\mathcal{R}$ is taken to be a Bregman or Kullback–Leibler-type divergence extending a vertex-triangular R. Which extensions preserve triangularity, and what geometry is possible under that constraint, appears to be open.

4.5 The continuous problem

Conjecture 4.2 (Operational characterization, smooth case)**.** *For a suitable class of smooth continuation systems, the tuple* $(g, \nabla, \nabla^*, \mathrm{Tor})$ *— with* (g, ∇, ∇^*) *induced by the regular part of the contrast and* Tor *the representation defect of the operational bracket — is, up to normalization, the unique local geometry monotone under all admissibility morphisms and invariant under the repair-sufficient subclass.*

We state this as a target, not a claim, but Theorem 3.10 now constrains its form. In the finite case, sufficiency-invariance characterized the bidirected cost, not the operational bracket: the smooth analogue should therefore expect the germ of the bidirected cost along the diagonal to be the characterized object, with the metric g as its symmetric second-order jet and the cubic tensor T carried by the third-order asymmetry — the dual geometry reappearing as the jet expansion of the characterized invariant rather than as an independently selected structure — while Tor, like the co-tests, is expected to be natural only under monotone comparison. The amalgamation technique should transfer: gluing smooth systems along cost-matched pairs is a boundary construction, and its regularity is where the difficulty will concentrate. The infinite-dimensional Chentsov theorem required four decades and carefully engineered regularity [5]; there is no reason to expect the continuation version to be easier. What this paper establishes is finite and secure: that the discrete theory is enriched-categorical rather than differential-geometric in kind; that the enriched and divergence regimes are separated by a computable sign; that coherence between comparison and execution is a local certificate in the first regime and a global axiom in the second; and that the sufficiency-invariant content of a finite continuation system is exactly its bidirected executable cost, no more and no less.

5 Interpretations within the Admissibility Program

The finite theory just established has consequences beyond the category of continuation systems in which it was proved. This section draws out five of them, in each case deriving a statement from Theorem 3.10 or Proposition 1.6 rather than merely asserting a thematic resemblance.

5.1 Admissibility and reachability

For a state $x \in \mathcal{C}$, define the *forward continuation neighborhood* $V(x) = \{y \in \mathcal{C} : \mathsf{E}(x, y) < \infty\}$, the set of states reachable from x at finite executable cost, and dually the *backward* neighborhood $V^{-1}(x) = \{y \in \mathcal{C} : \mathsf{E}(y, x) < \infty\}$. Admissibility of a continuation is then not a property of a destination state in isolation but of the neighborhood it occupies: two states x, x' with $V(x) = V(x')$ and $V^{-1}(x) = V^{-1}(x')$ are interchangeable with respect to what can follow and what could have preceded them, even if $x \neq x'$ and $\mathsf{R}(x, x') > 0$.

Given a comparison d on subsets of $\mathcal{C}$, define the *admissibility distortion* $D(x, x') = d(V(x), V(x')) + d(V^{-1}(x), V^{-1}(x'))$. This quantity is a functional of $\{\mathsf{E}(x, \cdot), \mathsf{E}(\cdot, x)\}$ alone, since $V(x)$ and $V^{-1}(x)$ are level sets of E. Theorem 3.10 therefore applies to it directly: an admissibility distortion built from reachability sets in this way is automatically expressible as a function of sufficiency-invariant content, because it is a function of the bidirected cost. An admissibility distortion built instead from R-neighborhoods $\{y : \mathsf{R}(x, y) < \theta\}$ carries no such guarantee, by Remark 2.8: it may distinguish states that every E-sufficient redescription identifies. The bidirected cost is, in this sense, exactly the resolution at which reachability-based admissibility judgments remain invariant, and no finer.

This bears on the reconstruction literature in one further respect. A state x is reconstructible from a later state y exactly when $x \in V^{-1}(y)$; the theorem's content, applied here, is that what survives redescription is not the identity of x but the pattern of which states lie in $V^{-1}(y)$ and at what bidirected cost. Reconstruction, on this reading, is a question about continuation structure rather than about recovering a state's identity as such.

The witness adequacy functional of the smooth, continuous-trajectory setting — $\Lambda(w) =$

$\sup_{H \in \mathcal{A}(w)} d(H, \widehat{H}(w))$, the worst-case distance from a witness's admissible continuation set to a model's output — is the natural smooth-case cousin of the reachability neighborhoods $V(x)$ and $V^{-1}(x)$ defined here: both quantify how tightly partial information constrains a continuation set, one by a sup over a metric and the other by a level set of E. Making that correspondence precise is part of the smooth characterization problem of Section 4 and is not attempted here.

5.2 Projection defect and CLIO

The projection-defect construction of Section 1 admits a reading in the vocabulary of CLIO, where a system's *substrate* is the space in which change occurs and an *interface* is a subspace onto which the substrate is represented. Under that correspondence, $\mathcal{C}$ with E is the substrate, an admissible subset $\mathcal{A}$ is an interface, and the projection $C^* = \arg\min_{t \in \mathcal{A}} \mathsf{R}(C, t)$ is the representation of C at that interface. The defect $\Delta_{\mathcal{A}}(C, C') = \mathsf{R}(C, C') - \mathsf{R}(C, C^*) - \mathsf{R}(C^*, C')$ measures representational distortion: the amount by which comparing two substrate points directly disagrees with comparing them through the interface.

Proposition 1.6 restricts which direction that distortion can run. If the substrate's comparison structure is enriched, $\Delta_{\mathcal{A}} \leq 0$ for every interface: representation through an interface can only make two points look *more* alike than direct comparison does, never less. Positive projection interference — an interface making two points look less alike than they are, the signature phenomenon of dually flat information geometry — requires R to fail the triangle inequality, i.e. requires leaving the executable substrate for a divergence-type contrast that is not itself a hom-structure of any enrichment. The enriched/divergence fork of Section 1 is, under this correspondence, a fork in what kinds of representational distortion an interface can produce: monotone-only distortion in the enriched regime, and reversal of apparent similarity only in the divergence regime, where coherence between substrate and interface becomes a global postulate rather than a locally checkable one (Lemma 1.4 has no divergence-regime analogue).

5.3 Distinction geometry and repair

Definition 5.1 (Operational faithfulness)**.** *Say* R *is* operationally faithful *on the ordered pair* (x, y) *if some operational word* w *from* x *to* y *satisfies* $\mathrm{cost}(w) = \mathsf{R}(x, y)$ *— that is, if the recognition cost* $\mathsf{R}(x, y)$ *is actually achieved by an executable path, not merely bounded by one.*

Proposition 5.2. *If* R *is a generator-dominated enrichment, then* $\mathsf{G}(x, y) = 0$ *if and only if* R *is operationally faithful on* (x, y)*.*

Proof. If some word w realizes $\mathsf{R}(x, y)$, then $\mathsf{E}(x, y) \leq \mathrm{cost}(w) = \mathsf{R}(x, y) \leq \mathsf{E}(x, y)$, the second inequality by Lemma 1.4, so $\mathsf{G}(x, y) = \mathsf{E}(x, y) - \mathsf{R}(x, y) = 0$. Conversely, if $\mathsf{G}(x, y) = 0$ then $\mathsf{E}(x, y) = \mathsf{R}(x, y)$; since $\mathcal{C}$ is finite and every cost is strictly positive, the infimum defining $\mathsf{E}(x, y)$ is attained by some word w (Lemma 3.8's proof makes the same finiteness observation), and $\mathrm{cost}(w) = \mathsf{E}(x, y) = \mathsf{R}(x, y)$, so w realizes $\mathsf{R}(x, y)$ and R is operationally faithful on (x, y). $\square$

The gap $\mathsf{G} = \mathsf{E} - \mathsf{R}$ has, to this point, been treated as a computed diagnostic; Proposition 5.2 gives its vanishing a name — recognition and execution agree on a pair exactly when R is operationally faithful there — and turns "recognition agrees with execution," used informally throughout the paper, into a condition with a proof rather than a figure of speech. It bears directly on a question Distinction Geometry and Repair Theory pose separately: the first asks how distinguishable two states are, the second how recoverable one is from the other. The twelve-state system shows these are not one structure under two names. Forty-eight of the 132 ordered pairs have $\mathsf{R} < \infty$ but $\mathsf{E} = \infty$: every downward repair is recognizable, often cheaply, yet inexecutable, since no operation lowers the level ℓ. States can therefore be maximally distinguishable in the operational sense while being cheaply comparable under R. Conversely, among finite pairs, $\mathsf{R}((0,0),(2,3)) = 2.5$ against $\mathsf{E}((0,0),(2,3)) = 7$: states a recognition process treats as nearly identical are separated by substantial executable cost, because R exploits a cyclic symmetry the operational algebra does not respect.

G is best read, on this evidence, as a *distinction–repair mismatch*: the extent to which recognizability and recoverability diverge, rather than a bookkeeping difference. Theorem 3.10 sharpens what the mismatch means for invariant theory. Recognizability, as encoded in R, need not be sufficiency-invariant at all; it sits outside the characterization of Section 3 except where it coincides with E (Remark 2.8). What is characterized is recoverability alone, in both directions. A theory that identifies distinguishability with recoverability is, on this evidence, conflating two independent structures; the present results quantify the conflation without resolving it, and the joint invariant theory of (R, E) remains the natural next problem.

5.4 Continuation geometry

The neighborhood $V(x)$ of Section 5.1 organizes into a continuation field Γ assigning to each state its reachable future, with the bidirected cost $(\mathsf{E}(x,y), \mathsf{E}(y,x))$ as local coordinates on pairs within it. In this language, Theorem 3.10 restates as: what survives every admissibility-preserving redescription of a continuation system is not the identity of its states but the geometry of Γ — the reachable-future structure and its bidirected costs. This is a restatement, not an extension, of the theorem, offered to align vocabulary with related work where the same distinction is drawn without the present finite-case proof.

5.5 Further connections

Two remarks, offered as interpretation rather than as claims proved here. Executable cost has, in RSVP, the shape of a path-integral functional through capacity, transport, and entropy fields, $\mathsf{E}(x,y) = \inf_\gamma \int_\gamma L(\Phi, v, S)\, dt$; the finite system of this paper is then a discrete prototype of that reachability structure, small enough to compute exhaustively, but nothing proved here shows that RSVP's continuous reachability functionals satisfy an analogue of Theorem 3.10 — that remains conjectural, alongside the smooth case of Section 4. In systems with multiple interacting repair modules, distinct modules may agree on recognition (R) while disagreeing on executable repair (E); the present results say that only the latter agreement is what a shared, sufficiency-invariant repair policy can rely on, with G measuring the gap between what a module recognizes as reparable and what it can execute.

A caution belongs here. Where repair operators have been defined elsewhere as contractions toward an admissible manifold, that contraction property is an additional hypothesis, not a consequence of anything proved in this paper: neither R nor E is assumed here to shrink under iteration, and Theorem 3.10 places no requirement of that kind on either primitive. Any specialization of the present framework to contraction-type repair operators should state the contraction hypothesis explicitly rather than inherit it by association.

6 What Survives

Theorem 3.10 admits two readings, and they point the program in different directions.

The first is deflationary and should be stated plainly. If one asks what a continuation system *is*, invariantly — what remains when every distinction-preserving redescription has been quotiented away — the answer is austere: a bidirected cost structure, the forward and reverse price of every transition. The rich operational texture that the monotone invariants detect — the difference between a promotion and a phase adjustment of equal cost, the concentration of noncommutativity on a singular stratum, the entire torsion-like layer of structure — none of it is sufficiency-invariant. It is real, and it is natural in the weaker sense, but it is perspective-bound: visible to comparison, erased by lossless recoding. A theory that wants to speak invariantly about continuation systems has exactly two numbers per ordered pair to speak with.

The second reading is the one the Admissibility Program should take. The theorem does not say the operational stratum is illusory; it says the operational stratum is precisely the content of a *description*, over and above the content of the system described. The gap between the (M)-natural theory and the (M)+(I)-natural theory measures how much structure lives in the choice of presentation — in which operations are taken as generators, where their singularities fall, how their order matters — and the theorem locates that structure exactly rather than gesturing at it. In the vocabulary of the surrounding work: the bidirected cost is what is *restorable* from any faithful witness of the system; the bracket data is what a witness additionally knows in virtue of how it witnessed. That the invariant content of a continuation system reduces to reachability cost in both directions — not to a truth-structure, not to a state-structure, but to what can be reached and what can be recovered — is, in the finite case exactly, the admissibility thesis in the form of a theorem.

The question that opened this inquiry — whether an admissibility-theoretic framework could derive the covariant derivative — can now be answered with the precision it originally lacked. In the weak sense, yes, and trivially: any transport of continuations across varying contexts has the shape $\partial + \Omega$. In the strong sense the question was wrongly posed. The covariant derivative of Riemannian geometry is the unique transport compatible with a symmetric quadratic comparison structure; the present setting has a directed, frequently non-quadratic, frequently non-smooth comparison structure, and what is unique given *its* invariances is not a connection but the bidirected cost. Connection-like objects reappear only in the tangent approximation, as jets of the characterized invariant, which is where they belong. What the frameworks of this program derive from first principles is not the Γ of general relativity; it is the thing of which Γ is a smooth, symmetric, reversible shadow.

Beyond the continuous conjecture of Section 4, three problems remain open in the finite theory itself. The characterization concerns functionals of a single ordered pair; the invariant theory of n-point configurations, where amalgamation has less room to move, is untouched, and it is there that bracket data may re-enter the sufficiency-invariant world. The morphism class was fixed by one notion of cost-nonincrease; coarser and finer classes — cost-scaling morphisms, morphisms preserving an admissible cone, morphisms preserving a repair pairing — will characterize different invariants, and the map from morphism class to characterized structure is the proper generalization of the selector problem with which this inquiry began. And the repair contrast R, which the characterization deliberately quotiented out, returns as data the moment one fixes a presentation: the joint invariant theory of the pair (R, E) — of recognition against execution — is the natural next paper.

One sentence of the program survives this inquiry unchanged and acquires a proof schema.

Repair is not something performed inside a geometry; repair is the operation from which the geometry is inferred. What the finite theory adds is the terminus of the inference. Pose it invariantly, and the geometry it infers is the bidirected cost: what it costs to arrive, and what it costs to return.

Appendix: Computation

The complete computation is reproduced by the accompanying scripts: `toymodel.py` (Section 1: Dijkstra over the operational graph for E; exhaustive evaluation of R, G, the commutator defect, and $\Delta_{\mathcal{A}}$ for both admissible sets), `characterize.py` (Theorem 3.4), `crosstest.py` (Theorem 3.6), and `amalgam.py` (Lemma 3.8 instances and Corollary 3.11). The full matrices follow; the diagonal is zero.

	00	01	02	03	10	11	12	13	20	21	22	23
00	0	1	2	3	2	3	4	5	4	5	6	7
01	3	0	1	2	5	2	3	4	6	4	5	6
02	2	3	0	1	4	5	2	3	6	7	4	5
03	1	2	3	0	3	4	5	2	5	6	7	4
10	∞	∞	∞	∞	0	1	2	3	2	3	4	5
11	∞	∞	∞	∞	3	0	1	2	4	2	3	4
12	∞	∞	∞	∞	2	3	0	1	4	5	2	3
13	∞	∞	∞	∞	1	2	3	0	3	4	5	2
20	∞	∞	∞	∞	∞	∞	∞	∞	0	1	2	3
21	∞	∞	∞	∞	∞	∞	∞	∞	2	0	1	2
22	∞	∞	∞	∞	∞	∞	∞	∞	2	3	0	1
23	∞	∞	∞	∞	∞	∞	∞	∞	1	2	3	0

Table 1: Executable cost E. Rows: source state ℓp; columns: target.

	00	01	02	03	10	11	12	13	20	21	22	23
00	0	0.5	1	0.5	1	1.5	2	1.5	2	2.5	3	2.5
01	0.5	0	0.5	1	1.5	1	1.5	2	2.5	2	2.5	3
02	1	0.5	0	0.5	2	1.5	1	1.5	3	2.5	2	2.5
03	0.5	1	0.5	0	1.5	2	1.5	1	2.5	3	2.5	2
10	5	5.5	6	5.5	0	0.5	1	0.5	1	1.5	2	1.5
11	5.5	5	5.5	6	0.5	0	0.5	1	1.5	1	1.5	2
12	6	5.5	5	5.5	1	0.5	0	0.5	2	1.5	1	1.5
13	5.5	6	5.5	5	0.5	1	0.5	0	1.5	2	1.5	1
20	10	10.5	11	10.5	5	5.5	6	5.5	0	0.5	1	0.5
21	10.5	10	10.5	11	5.5	5	5.5	6	0.5	0	0.5	1
22	11	10.5	10	10.5	6	5.5	5	5.5	1	0.5	0	0.5
23	10.5	11	10.5	10	5.5	6	5.5	5	0.5	1	0.5	0

Table 2: Repair contrast R. Rows: source state ℓp; columns: target.

References

[1] F. W. Lawvere, *Metric spaces, generalized logic, and closed categories*, Rend. Sem. Mat. Fis. Milano 43 (1973).

[2] N. N. Chentsov, *Statistical Decision Rules and Optimal Inference*, AMS Translations (1982).

[3] S. Amari and H. Nagaoka, *Methods of Information Geometry*, AMS/Oxford (2000).

[4] T. Fritz, *A synthetic approach to Markov kernels, conditional independence and theorems on sufficient statistics*, Adv. Math. 370 (2020).

[5] N. Ay, J. Jost, H. V. Lê, L. Schwachhöfer, *Information Geometry*, Springer (2017).