

Categorical Locality Before Spacetime: Axioms for Physical Contexts, Composition, and Descent

Matthew Long

The YonedaAI Collaboration

YonedaAI Research Collective

Chicago, IL

matthew@yonedaai.com · <https://yonedaai.com>

23 Jul 2026

Abstract

This paper asks how much locality can be formulated before a manifold, a point set, or a causal order has been chosen. The proposed baseline consists of a category of physical contexts, morphisms called processes, a symmetric monoidal operation for composition, a pullback-stable coverage, and contravariant coefficient assignments. Compatibility is agreement after restriction to overlaps. Effective descent is the assertion that every compatible family comes from a unique global object, or from an object unique up to equivalence in a categorified target. The distinction between compatibility and effectiveness is essential. It prevents a coverage from silently becoming a gluing theorem.

Several modest results are proved in full. Word contexts form a strict monoidal model, while finite multisets form a strict symmetric monoidal model. An effective-descent witness gives a reconstruction map with both inverse identities. Identity coverage is effective for every presheaf, and products of set-valued sheaves remain sheaves. A finite locality reduct admits incompatible order-shaped, number-valued, topology-on-label-set, and pseudometric annotations. This proves non-uniqueness only for those added annotations. It does not construct admissible non-geometric physical models or settle whether the axioms determine realized causal or Lorentzian geometry. The larger goals of a physically selected coverage, higher-categorical coherence, general effective descent, and comparison functors to geometric theories remain **OPEN**. The paper records these limits against project claims P1-C01 through P1-C04 and assumptions P1-A01 through P1-A07.

Contents

1 Research question and scope

4

1.1	The question	4
1.2	Results and their limits	4
1.3	Nonclaims	5
1.4	Relation to the executable artifacts	5
2	Primitive signature	6
2.1	Contexts and processes	6
2.2	Compositional structure	6
2.3	What the signature omits	7
3	Coverage without open sets	8
3.1	Covering families	8
3.2	Pullbacks and overlaps	8
3.3	Refinement	9
4	Compatibility and effective descent	9
4.1	Set-valued coefficients	9
4.2	Categorified coefficients	10
5	Modest local-to-global theorems	11
5.1	Reconstruction from an effective witness	11
5.2	Identity coverage	12
5.3	Products of local coefficients	12
5.4	Invariance under natural isomorphism	13
5.5	What these theorems establish	13
6	A finite diagnostic witness	13
6.1	Diagnostic expansions of a reduct	13
6.2	The locality reduct	13
6.3	Incompatible diagnostic expansions	14
7	Geometric and algebraic comparators	15
7.1	Why comparisons are not proofs	15
7.2	Bordisms and topological field theory	15
7.3	Factorization algebras	16
7.4	Algebraic quantum field theory	16
7.5	Effective and non-effective descent	16
7.6	Higher topoi	16
7.7	Comparison matrix	17
8	Assumption audit	17
8.1	Circularity checks	18

9	Executable evidence and formal boundary	18
9.1	Lean	18
9.2	Haskell	19
9.3	Evidence matrix	19
10	Open reconstruction targets	19
10.1	Higher categorical coherence	19
10.2	Physical selection of coverage	20
10.3	Ambient category for effectiveness	20
10.4	Realization functors	21
10.5	The next-paper interface	21
11	Discussion	21
11.1	What is gained	21
11.2	What is not gained	22
11.3	Novelty boundary	22
11.4	Failure modes	22
12	Conclusion	22
A	Axiom checklist	23
A.1	Context and process data	23
A.2	Tensor data	23
A.3	Coverage data	23
A.4	Coefficient and descent data	24
B	Finite model details	24
B.1	Monoidal table	24
B.2	Coverage table	24
B.3	Coefficient table	25
B.4	Diagnostic annotation table	25
C	Proof dependency map	25
D	Claim and assumption cross-reference	26
D.1	P1-C01	26
D.2	P1-C02	26
D.3	P1-C03	26
D.4	P1-C04	26
E	Review checklist for reuse	27

1 Research question and scope

1.1 The question

The research question is precise:

What categorical data are sufficient to state composition, coverage, compatibility, effective descent, and locality without taking points, open subsets of a space, manifolds, causal cones, or a metric as primitive?

There are two parts to this question. The first is formal. One needs types for contexts, processes, covers, local data, and gluing. The second is physical. One needs a reason to regard the chosen covers as physically local. This paper answers part of the formal question. It does not settle the physical selection problem.

The word “before” in the title refers to dependency order. It does not refer to time. A process arrow is not assumed to be a temporal transition. A tensor product is not assumed to mean spacelike separation. A cover is not assumed to be an open cover. These interpretations may be added by a realization functor in a later paper. They are absent from the primitive signature here.

1.2 Results and their limits

The paper separates project claims from smaller mathematical lemmas. That distinction matters because a valid lemma can coexist with an open reconstruction target. Table 1 gives the current disposition.

Claim	Target	Paper status
P1-C01	The selected contexts and processes form the declared weak symmetric monoidal higher category and satisfy coherence.	OPEN
P1-C02	The physically selected coverage has effective descent in the declared target.	OPEN
P1-C03	Geometric and algebraic comparators instantiate the axioms through explicit realization functors.	OPEN
P1-C04	The locality axioms do not determine causal order, dimension, topology, or metric.	OPEN. Theorem 6.4 proves only a narrower local theorem about diagnostic annotations.

Table 1: Project claim status after the results of this paper. All four rows remain open because their missing physical or realization bridges are not supplied by the modest propositions proved below.

Claim P1-C01 remains open because the paper works at an ordinary one-categorical baseline. It gives a strict monoidal example, but it does not construct the selected weak higher category demanded by assumption P1-A01. Claim P1-C02 remains open because the paper does not derive a physically privileged coverage or prove effectiveness for all intended coefficients. Claim P1-C03 remains open because the comparisons in [Section 7](#) are hypothesis audits, not constructed realization functors.

Claim P1-C04 remains open. Its stable statement requires admissible non-geometric models and concerns causal, dimensional, topological, and metric structure in the intended realized sense. [Theorem 6.4](#) gives the same locality reduct two different diagnostic annotations of each listed kind. That local result shows that arbitrary object-label annotations are not fixed by the reduct. It does not show that either annotation is realized causal or Lorentzian geometry, so it does not discharge P1-C04.

1.3 Nonclaims

This paper does not derive:

- a causal or temporal order;
- dimension or a point-set topology;
- a differentiable structure, metric, connection, or curvature;
- quantum probabilities, states, or a Born rule;
- an equation of motion or any other dynamics.

Categorical composition says how arrows compose. It does not say how a physical state evolves. Even when a later realization reads some arrows dynamically, that reading will require an explicit interface and extra assumptions.

1.4 Relation to the executable artifacts

The repository contains a Lean interface and a finite Haskell model. They have deliberately smaller scopes than the paper. The Lean declaration `PhysicalContext` packages an object of an explicitly supplied category with claim metadata. The declaration `EffectiveDescent` contains an equivalence between global data and compatible local data. Its theorem `restrict_reconstruct` checks one inverse law.

The Haskell modules `Foundational.Context` and `Foundational.Descent` test finite word composition, finite covers, and overlap compatibility. Those tests are **COMPUTATIONAL**. They do not prove the project claims. They do provide executable witnesses that composition and coverage are different pieces of structure.

2 Primitive signature

2.1 Contexts and processes

Definition 2.1 (Context-process category). A *context-process category* is a locally small category $\mathbf{Ctx}$. Its objects are called *physical contexts*. For $U, V \in \text{Ob}(\mathbf{Ctx})$, a morphism $f : U \rightarrow V$ is called a *process from U to V*.

The adjective “physical” is a role designation. It adds no mathematical field to an object. In particular, an object of $\mathbf{Ctx}$ does not carry a point set, a topology, coordinates, a dimension, a causal relation, or a metric unless one of these is separately supplied.

The process terminology also carries no orientation convention beyond the source and target of a morphism. A morphism can later represent an embedding, a coarse-graining map, a preparation, a boundary inclusion, or something else. The category axioms alone do not decide among these readings.

Proposition 2.2 (Process coherence at the baseline). *For composable processes $U \xrightarrow{f} V \xrightarrow{g} W \xrightarrow{h} X$,*

$$h \circ (g \circ f) = (h \circ g) \circ f.$$

For every process $f : U \rightarrow V$,

$$f \circ \text{id}_U = f \quad \text{and} \quad \text{id}_V \circ f = f.$$

Proof. These equations are the associativity and unit axioms in [Theorem 2.1](#). No geometric or physical interpretation is used. $\square$

Remark 2.3. [Theorem 2.2](#) is a complete proof at the chosen one-categorical level. It does not settle P1-C01. That project claim asks for the selected weak symmetric monoidal higher category and all of its coherence data.

2.2 Compositional structure

Definition 2.4 (Compositional context system). A *compositional context system* is a symmetric monoidal category

$$(\mathbf{Ctx}, \otimes, I, \alpha, \lambda, \rho, \sigma).$$

The tensor $U \otimes V$ is called the *parallel composite*. The unit is I . The natural isomorphisms $\alpha, \lambda, \rho, \sigma$ are part of the declared structure and satisfy the usual pentagon, triangle, and hexagon equations.

The phrase “parallel composite” is chosen to avoid a stronger physical claim. It does not mean spacelike separation. Nor does it imply statistical independence. It records only the mode of composition supplied by the model.

Assumption P1-A01 asks for a weak higher-categorical version of [Theorem 2.4](#). This paper does not construct that object. Instead it fixes a one-categorical truncation and gives a strict finite model whose coherence can be checked directly.

Construction 2.5 (Word and multiset contexts). Fix a set A of atom labels. Let $\mathbf{Word}(A)$ be the discrete category whose objects are finite words in A . There is only an identity morphism on each word. Define

$$\langle a_1, \dots, a_m \rangle \otimes \langle b_1, \dots, b_n \rangle = \langle a_1, \dots, a_m, b_1, \dots, b_n \rangle.$$

The empty word is I . Let $\mathbf{MSet}(A)$ instead have finite multisets as objects and multiset sum as tensor. It is again discrete. Because multiset sum is commutative on the nose, its symmetry maps are identities.

Proposition 2.6 (Discrete word category is strict monoidal). *The discrete category $\mathbf{Word}(A)$, with concatenation as tensor and the empty word as unit, is a strict monoidal category.*

Proof. Let x, y, z be finite words. Both $(x \otimes y) \otimes z$ and $x \otimes (y \otimes z)$ list the entries of x , followed by the entries of y , followed by the entries of z . They are equal as lists. Concatenating the empty word on either side changes no entry. Thus associativity and both unit laws hold as equalities of objects. Since $\mathbf{Word}(A)$ is discrete, tensor sends the only available pair of morphisms, $(\text{id}_x, \text{id}_y)$, to $\text{id}_{x \otimes y}$ and is a bifunctor. The associator and unitors are identities, and every coherence diagram commutes. $\square$

Corollary 2.7 (Strict symmetric multiset composition). *Finite multisets under multiset sum form a strict symmetric monoidal model.*

Proof. Multiset sum is associative and has the empty multiset as an identity. It is commutative. In the discrete category on finite multisets, the associator, unitors, and symmetry are identity morphisms. Every coherence diagram therefore consists only of identity morphisms and commutes. $\square$

This proposition is mirrored by the Haskell property tests for `composeContexts`. The test evidence ranges over generated finite lists. The mathematical proof above ranges over all finite words. Neither result supplies a physical choice of atoms or processes.

2.3 What the signature omits

Let $\mathcal{L}_{\text{loc}}$ denote the signature consisting of:

1. objects and process hom-sets;
2. categorical identities and composition;
3. a symmetric monoidal product and its coherence data;
4. a coverage, defined in [Section 3](#);
5. coefficient presheaves and descent comparisons, defined in [Section 4](#).

There are no sorts or fields for points, neighborhoods, causal precedence, dimension, lengths, angles, durations, curvature, actions, or equations of motion. This is the signature audit required by assumption P1-A06. Omission is not yet an independence theorem. That stronger conclusion requires the expansion argument in [Section 6](#).

3 Coverage without open sets

3.1 Covering families

Definition 3.1 (Coverage). A *coverage* Cov on Ctx assigns to each object U a class $\text{Cov}(U)$ of families

$$\mathcal{U} = \{u_i : U_i \rightarrow U\}_{i \in J}.$$

Its members are called *covering families*. The assignment satisfies:

1. the singleton $\{\text{id}_U : U \rightarrow U\}$ covers U ;
2. for every cover $\{u_i : U_i \rightarrow U\}$ and morphism $f : V \rightarrow U$, the pullbacks $V \times_U U_i$ exist and the family

$$\{\pi_i : V \times_U U_i \rightarrow V\}_{i \in J}$$

covers V ;

3. if $\{u_i : U_i \rightarrow U\}$ covers U and $\{v_{ik} : V_{ik} \rightarrow U_i\}_{k \in K_i}$ covers every U_i , then $\{u_i \circ v_{ik} : V_{ik} \rightarrow U\}_{i,k}$ covers U .

This is a pretopology-style presentation [5, Chapter III, Section 2, Definition 2]. One may pass to the generated Grothendieck topology. The family formulation makes the later compatibility equations explicit. Higher-categorical versions replace pullbacks and equalities by homotopy pullbacks and coherent equivalences. That replacement belongs to the open extension of P1-C01.

Assumption P1-A02 is exactly the base-change requirement in the second clause. The paper treats it as declared structure. It does not prove that a proposed physical laboratory or algebraic system selects such a coverage.

Warning 3.2 (Coverage is not locality by itself). A coverage lists admissible families. It does not assert that local data agree, and it does not assert that compatible data glue. Calling every coverage “locality” would collapse three distinct questions: which families cover, which restrictions are compatible, and which compatible families are effective.

3.2 Pullbacks and overlaps

Fix a cover $\mathcal{U} = \{u_i : U_i \rightarrow U\}$. For each pair i, j , write

$$U_{ij} = U_i \times_U U_j$$

with projections $p_{ij}^i : U_{ij} \rightarrow U_i$ and $p_{ij}^j : U_{ij} \rightarrow U_j$. These pullbacks are categorical overlaps. They are not intersections of subsets unless a realization functor proves that comparison.

Triple overlaps are iterated pullbacks

$$U_{ijk} = U_i \times_U U_j \times_U U_k.$$

Associativity of pullback is canonical only up to isomorphism in a general category. For set-valued compatibility, the pairwise equations below imply the corresponding equalities after restriction to triple overlaps. For category-valued data, coherent higher comparison cells must be retained.

Lemma 3.3 (Compatibility survives further restriction). *Let $F : \mathbf{Ctx}^{\text{op}} \rightarrow \mathbf{Set}$ be a presheaf. Suppose $x_i \in F(U_i)$ satisfy*

$$F(p_{ij}^i)(x_i) = F(p_{ij}^j)(x_j)$$

for every pair i, j . For any morphism $q : W \rightarrow U_{ij}$, the two restrictions to W agree.

Proof. Apply the function $F(q)$ to the displayed equality. Functoriality gives

$$F(q)F(p_{ij}^i)(x_i) = F(q)F(p_{ij}^j)(x_j).$$

Contravariant composition rewrites the two sides as the restrictions along $p_{ij}^i \circ q$ and $p_{ij}^j \circ q$. $\square$

3.3 Refinement

Definition 3.4 (Refinement). A cover $\{v_a : V_a \rightarrow U\}_{a \in K}$ *refines* $\{u_i : U_i \rightarrow U\}_{i \in J}$ if there is a function $r : K \rightarrow J$ and morphisms $w_a : V_a \rightarrow U_{r(a)}$ such that

$$u_{r(a)} \circ w_a = v_a$$

for every a .

Proposition 3.5 (Compatible families restrict to refinements). *Let $F : \mathbf{Ctx}^{\text{op}} \rightarrow \mathbf{Set}$ be a presheaf. A compatible family on a cover restricts to a compatible family on every refinement for which the required pullbacks exist.*

Proof. Let (x_i) be compatible and define $y_a = F(w_a)(x_{r(a)})$. On $V_a \times_U V_b$, both routes to $U_{r(a)} \times_U U_{r(b)}$ commute with the structure maps to U . The equality for $x_{r(a)}$ and $x_{r(b)}$ therefore pulls back to equality for y_a and y_b . This is an application of [Theorem 3.3](#). $\square$

The proposition is formal. It does not say that every physically meaningful refinement is present in $\mathbf{Cov}$. That is part of the selection problem in open question [0Q-P1-02](#).

4 Compatibility and effective descent

4.1 Set-valued coefficients

Definition 4.1 (Coefficient presheaf). A *coefficient presheaf* is a functor

$$F : \mathbf{Ctx}^{\text{op}} \longrightarrow \mathbf{Set}.$$

For a process $f : V \rightarrow U$, the function $F(f) : F(U) \rightarrow F(V)$ is called restriction along f .

Assumption [P1-A03](#) says that the intended assignment has this functorial structure or a declared higher analogue. The word “coefficient” is neutral. Elements of $F(U)$ need not be observables, states, fields, or measurement outcomes. Those readings require separate definitions.

Definition 4.2 (Compatible family). For a cover $\mathcal{U} = \{u_i : U_i \rightarrow U\}$, a family $(x_i)_{i \in J}$ with $x_i \in F(U_i)$ is *compatible* when

$$F(p_{ij}^i)(x_i) = F(p_{ij}^j)(x_j)$$

in $F(U_{ij})$ for all i, j . Write $\text{Desc}_F(\mathcal{U})$ for the set of compatible families.

Definition 4.3 (Comparison map). The *descent comparison map* is

$$\kappa_{\mathcal{U}} : F(U) \longrightarrow \text{Desc}_F(\mathcal{U}), \quad x \longmapsto \left(F(u_i)(x) \right)_{i \in J}.$$

Lemma 4.4 (The comparison map is well defined). *The restrictions of every $x \in F(U)$ form a compatible family.*

Proof. For each i, j , the pullback square gives $u_i \circ p_{ij}^i = u_j \circ p_{ij}^j$. Contravariant functoriality yields

$$F(p_{ij}^i)F(u_i)(x) = F(u_i \circ p_{ij}^i)(x) = F(u_j \circ p_{ij}^j)(x) = F(p_{ij}^j)F(u_j)(x).$$

Thus $\kappa_{\mathcal{U}}(x)$ satisfies [Theorem 4.2](#). □

Definition 4.5 (Separatedness and effective descent). The presheaf F is *separated for $\mathcal{U}$* if $\kappa_{\mathcal{U}}$ is injective. It has *effective descent for $\mathcal{U}$* if $\kappa_{\mathcal{U}}$ is bijective. It is a *sheaf for $\mathbf{Cov}$* if it has effective descent for every cover in $\mathbf{Cov}$.

Surjectivity means existence of a global glue. Injectivity means uniqueness. Compatibility alone supplies neither. This is the mathematical content behind [Theorem 3.2](#).

Definition 4.6 (Locality datum). A *locality datum* is a tuple

$$(\text{Ctx}, \otimes, I, \mathbf{Cov}, \mathcal{F})$$

where $(\text{Ctx}, \otimes, I)$ is a compositional context system, $\mathbf{Cov}$ is a coverage, and $\mathcal{F}$ is a declared class of coefficient presheaves. It is *effective for $\mathcal{F}$* if every $F \in \mathcal{F}$ is a sheaf for $\mathbf{Cov}$.

The definition makes locality relative to coefficients. A coverage can be effective for one presheaf and fail for another. It can also be effective in one ambient category and fail after the target is restricted. The non-effective projective descent example in the Stacks Project directly shows this dependence [7].

4.2 Categorized coefficients

Let $\mathcal{D}$ be a category with the products and further limits used below, and let $F : \text{Ctx}^{\text{op}} \rightarrow \mathcal{D}$ be a functor. For a cover indexed by J , write

$$U_{i_0 \dots i_n} = U_{i_0} \times_U \cdots \times_U U_{i_n}.$$

Define a cosimplicial diagram directly in $\mathcal{D}$ by

$$C_F^n(\mathcal{U}) = \prod_{(i_0, \dots, i_n) \in J^{n+1}} F(U_{i_0 \dots i_n}).$$

The coface and codegeneracy maps come from the pullback projections and diagonals. This construction needs the displayed products in $\mathcal{D}$, but it does not require coproducts of the U_i inside $\mathbf{Ctx}$. Define the descent object, when the indicated limit exists, by

$$\mathrm{Desc}_F(\mathcal{U}) = \lim_{[n] \in \Delta} C_F^n(\mathcal{U}).$$

The restrictions from U to every iterated overlap induce

$$\kappa_{\mathcal{U}} : F(U) \longrightarrow \mathrm{Desc}_F(\mathcal{U}).$$

If coefficients take values in a 2-category, the corresponding assignment is a pseudofunctor and the descent object is a pseudolimit. For an infinity-category, it is a homotopy-coherent functor and the descent object is a homotopy limit.

Definition 4.7 (Categorified effective descent). The cover $\mathcal{U}$ has *effective descent* for F in $\mathcal{D}$ when $\kappa_{\mathcal{U}}$ is an isomorphism. For an infinity-categorical target, replace “isomorphism” by “equivalence” and use the homotopy limit.

Higher topos theory supplies an established setting for such descent machinery [4]. It does not select $\mathbf{Ctx}$, $\mathbf{Cov}$, or the physical interpretation of F . The passage from this paper’s baseline to a chosen higher setting is part of assumptions P1-A01, P1-A03, and P1-A04.

5 Modest local-to-global theorems

5.1 Reconstruction from an effective witness

Definition 5.1 (Effective-descent witness). For a set-valued presheaf and cover, an *effective-descent witness* is a bijection

$$e : F(U) \xrightarrow{\sim} \mathrm{Desc}_F(\mathcal{U})$$

that equals the comparison map $\kappa_{\mathcal{U}}$.

Theorem 5.2 (Reconstruction and uniqueness). *Let e be an effective-descent witness. Define*

$$\mathrm{reconstruct}_e = e^{-1}.$$

Then

$$e(\mathrm{reconstruct}_e(x)) = x$$

for every compatible family x , and

$$\mathrm{reconstruct}_e(e(y)) = y$$

for every global datum $y \in F(U)$. The reconstructed global datum is unique.

Proof. The first two equations are the two inverse identities for the bijection e . For uniqueness, suppose $y, y' \in F(U)$ both restrict to x . Then $e(y) = x = e(y')$. Because e is injective, $y = y'$. $\square$

The Lean theorem `EffectiveDescent.restrict_reconstruct` formalizes the first inverse identity after recording that the equivalence agrees with restriction. The theorem is complete but conditional on the witness. It does not prove that an intended physical coefficient assignment has such a witness.

Corollary 5.3 (No hidden choice). *The reconstruction map in Theorem 5.2 uses no choice beyond the supplied equivalence.*

Proof. The inverse function is a field of the bijection. No selection from nonempty fibers is performed. $\square$

5.2 Identity coverage

Proposition 5.4 (Identity covers are effective). *Let $\mathbf{Cov}_{\text{id}}$ contain only the singleton identity cover of each object. Every presheaf $F : \mathbf{Ctx}^{\text{op}} \rightarrow \mathbf{Set}$ is a sheaf for $\mathbf{Cov}_{\text{id}}$.*

Proof. For the cover $\{\text{id}_U : U \rightarrow U\}$, a family is one element $x \in F(U)$. The only overlap is $U \times_U U \cong U$, and the compatibility equation is $x = x$. Thus $\text{Desc}_F(\{\text{id}_U\})$ is canonically $F(U)$. The comparison map is the identity under this identification, hence a bijection. $\square$

Remark 5.5. The proposition shows that effective descent can be vacuous. It gives no reason to call the identity coverage physically adequate. A useful locality theory needs nontrivial covers and a justification for them.

5.3 Products of local coefficients

Proposition 5.6 (Products preserve set-valued locality). *Let $F, G : \mathbf{Ctx}^{\text{op}} \rightarrow \mathbf{Set}$ be sheaves for $\mathbf{Cov}$. Define H objectwise by*

$$H(U) = F(U) \times G(U).$$

Then H is a sheaf for $\mathbf{Cov}$.

Proof. Fix a cover $\mathcal{U}$. A compatible H -family is a family (x_i, y_i) whose restrictions agree on every overlap. Equality in a product is componentwise. Therefore (x_i) is a compatible F -family and (y_i) is a compatible G -family. Because F and G are sheaves, there are unique glues $x \in F(U)$ and $y \in G(U)$. The pair $(x, y) \in H(U)$ restricts to (x_i, y_i) . If (x', y') is another glue, uniqueness for F gives $x' = x$ and uniqueness for G gives $y' = y$. Thus the comparison map for H is bijective. $\square$

Corollary 5.7 (Finite products). *Every finite objectwise product of set-valued sheaves is a sheaf.*

Proof. Apply Theorem 5.6 inductively. The nullary product is the terminal presheaf, whose unique compatible family has a unique glue. $\square$

5.4 Invariance under natural isomorphism

Proposition 5.8 (Locality is invariant under coefficient isomorphism). *Suppose $F, G : \mathbf{Ctx}^{\text{op}} \rightarrow \mathbf{Set}$ are naturally isomorphic. If F is a sheaf for $\mathbf{Cov}$, then G is a sheaf for $\mathbf{Cov}$.*

Proof. Let $\eta : F \Rightarrow G$ be a natural isomorphism. For each cover, apply η componentwise to global data and to matching families. Naturality makes the square of comparison maps commute. The vertical maps are bijections. If the comparison map for F is a bijection, the comparison map for G is a composite of bijections and is therefore a bijection. $\square$

5.5 What these theorems establish

The results in this section are intentionally modest. They establish formal consequences after the categorical data have been specified. They do not establish that a proposed physical family is a cover. They do not establish that the desired coefficient objects form a sheaf. They do not turn the morphisms of $\mathbf{Ctx}$ into a causal relation.

The distinction determines the status of P1-C02. [Theorem 5.2](#) is **PROVED**. The project claim that the selected physical coefficients actually satisfy effective descent remains **OPEN**. Assumption P1-A04 has not been discharged for a nontrivial physical target.

6 A finite diagnostic witness

6.1 Diagnostic expansions of a reduct

Definition 6.1 (Diagnostic expansion and local selection). Let $\mathcal{L}_{\text{loc}}$ be the locality signature fixed in [Section 2](#). Let $\mathcal{L}_{\text{ext}}$ add one further structure S . For this section, a *diagnostic S -expansion* of a locality model M is an arbitrary interpretation of S on the object labels of the same underlying model without changing any $\mathcal{L}_{\text{loc}}$ data. We say that the reduct *locally selects S on M* if all such diagnostic expansions are isomorphic by an isomorphism whose reduct is the identity on M .

This definition is deliberately about uniqueness. It does not assert that a diagnostic annotation is an admissible realized geometry, nor does it decide whether an extra structure is first-order definable in another formalization. If one locality reduct has two nonisomorphic diagnostic expansions, the reduct does not choose between those annotations. That statement is weaker than P1-C04.

6.2 The locality reduct

Construction 6.2 (Boolean locality reduct). Let $B = \{0, 1\}$ with tensor

$$x \otimes y = \max(x, y)$$

and unit 0. Regard B as a discrete category. Equip it with identity coverage. Let the coefficient class contain the terminal presheaf $\mathbf{1} : B^{\text{op}} \rightarrow \mathbf{Set}$. Call the resulting locality datum M_B .

Lemma 6.3 (The reduct satisfies the baseline axioms). *The datum M_B is a strict symmetric monoidal context system with a pullback-stable coverage, and its declared coefficient presheaf has effective descent.*

Proof. Maximum on $\{0, 1\}$ is associative, commutative, and has identity 0. Because the category is discrete, this gives a strict symmetric monoidal structure. Every identity pullback exists and is an identity. Identity families remain identity families after base change and composition. Thus the identity coverage satisfies [Theorem 3.1](#). The terminal presheaf is a sheaf by [Theorem 5.4](#). $\square$

6.3 Incompatible diagnostic expansions

We now add diagnostic annotations to the object set $\{0, 1\}$. Each addition leaves the category, tensor, coverage, and coefficient presheaf unchanged.

For an order-shaped annotation, consider:

$$R_{\text{disc}} = \{(0, 0), (1, 1)\}, \quad R_{\text{chain}} = \{(0, 0), (0, 1), (1, 1)\}.$$

Both are partial orders. They are unequal. Any isomorphism whose reduct is the identity fixes both objects, so it cannot carry one relation to the other.

For a number-valued annotation, consider:

$$d_0(0) = d_0(1) = 0, \quad d_1(0) = 0, \quad d_1(1) = 1.$$

Both are functions $B \rightarrow \mathbb{N}$. They are unequal under the identity reduct. No dimension axiom appears in $\mathcal{L}_{\text{loc}}$.

For a topology annotation on the finite object-label set B , consider:

$$\tau_{\text{ind}} = \{\emptyset, B\}, \quad \tau_{\text{disc}} = \mathcal{P}(B).$$

Both satisfy the topology axioms. The identity map cannot identify them.

For a pseudometric annotation on the finite object-label set B , consider m_1 and m_2 with

$$m_r(x, x) = 0, \quad m_r(0, 1) = m_r(1, 0) = r$$

for $r = 1, 2$. Both are pseudometrics. They differ under the identity on B .

Theorem 6.4 (Local non-uniqueness of diagnostic expansions). *In the sense of [Theorem 6.1](#), the Boolean locality reduct M_B does not locally select an order-shaped annotation, a number-valued annotation, a topology on its finite object-label set, or a pseudometric annotation on that set.*

Proof. Use the locality model M_B from [Theorem 6.2](#). [Theorem 6.3](#) verifies the baseline locality axioms. The two relations R_{disc} and R_{chain} are nonisomorphic diagnostic expansions over the identity reduct. The functions d_0 and d_1 give nonisomorphic number-valued annotations. The topologies τ_{ind} and τ_{disc} give nonisomorphic topology-on-label-set annotations. The pseudometrics m_1 and m_2 give nonisomorphic pseudometric annotations. In every case, the underlying category, tensor, coverage, and coefficient presheaf are exactly the same. Therefore the reduct locally selects none of the four added annotations. $\square$

Corollary 6.5 (Arrow direction does not select the order annotation). *The process relation of M_B does not select between the two order annotations above.*

Proof. The category underlying M_B is discrete in both diagnostic expansions. Its morphisms are unchanged while the order annotation changes. $\square$

Corollary 6.6 (Coverage does not select the topology annotation). *The coverage of M_B does not select between the two topologies on its object labels.*

Proof. The identity coverage of M_B is unchanged under the indiscrete and discrete diagnostic expansions. $\square$

Remark 6.7 (What the theorem does not say). The theorem is a local result about arbitrary annotations of the two object labels in M_B . An order on those labels is not a causal relation on events. A number-valued label is not a realized dimension. A topology or pseudometric on those labels is not a spacetime topology or a Lorentzian metric. The theorem therefore neither supplies the admissible non-geometric model required by P1-A07 nor proves P1-C04. It also does not say that a later realization cannot reconstruct geometry from additional hypotheses.

The construction does not supply the independence witness requested by assumption P1-A07. It does place the narrower diagnostic failures behind CE-P1-02 and CE-P1-03 in a common finite setting: composition exists with only trivial coverage, and a site can support descent while its reduct fails to choose between two order annotations.

7 Geometric and algebraic comparators

7.1 Why comparisons are not proofs

Existing geometric theories show that categorical composition and local-to-global structures are useful. They do not show that the present pre-geometric signature has the intended physical meaning. The distinction is recorded by assumptions P1-A05 and P1-A06.

An explicit comparison would require a functor

$$R : \text{Ctx}_{\text{pre}} \longrightarrow \text{Ctx}_{\text{geom}}$$

or a functor in the reverse direction, together with specified preservation data for tensor products, covers, and descent comparisons. It would also require an adequacy statement. For example, one would need to say whether R is faithful, conservative, fully faithful, essentially surjective on a controlled subcategory, or only an interpretation. No such functor is constructed in this paper.

7.2 Bordisms and topological field theory

Atiyah's axioms treat closed manifolds as objects and bordisms as morphisms [1]. Disjoint union supplies a symmetric monoidal operation. A topological field theory sends this

geometric category to algebraic data in a way compatible with composition. This is a concrete comparator for process composition.

It is not a pre-geometric proof. The domain already contains manifolds, boundaries, and gluing of bordisms. Using it as the definition of `Ctx` would import topology and dimension at the first step. The present program instead asks whether a context category can be specified without those fields and later mapped into a bordism category under explicit realization assumptions.

7.3 Factorization algebras

Factorization algebras assign data to opens and provide structure maps for families of disjoint opens [2]. They offer a mature account of observables with local-to-global behavior. They separate prefactorization structure from descent-like conditions, which is the same type boundary used here.

The usual domain is a topological space or manifold. Its opens and disjointness relation are already geometric. Therefore factorization algebras test whether a later realization preserves local operations. They do not derive the pre-geometric coverage `Cov`.

7.4 Algebraic quantum field theory

Haag and Kastler associate observable algebras with spacetime regions and impose isotony and locality conditions [3]. This is an algebraic presentation, but its indexing regions retain spacetime content. The fact that the assigned objects are algebras does not make the domain pre-geometric.

For the present program, an AQFT net would be a realization target. One would need to construct both a region-valued realization of contexts and an algebra-valued realization of coefficients. The compatibility of those two maps would be a theorem, not terminology.

7.5 Effective and non-effective descent

Faithfully flat module descent gives a concrete effective theorem: modules with descent data over a faithfully flat map descend, and base extension is an equivalence [6]. This example demonstrates the force of an effective-descent statement. It names the cover class, ambient category, coefficient type, and comparison functor.

The Stacks Project also records projective scheme descent data that are not effective in schemes, although they become effective in algebraic spaces [7]. This obstruction is directly relevant to P1-C02. One cannot infer effectiveness from compatibility alone. Changing the ambient category can change the answer.

7.6 Higher topoi

Higher topos theory gives an infinity-categorical home for homotopy-coherent descent [4]. It supplies tools for a future upgrade of [Theorem 4.7](#). It does not identify which infinity-category should model physical contexts or which covering families should count as physical locality.

7.7 Comparison matrix

Framework	Primitive domain	Local structure	Why it does not prove the present bridge
Atiyah TQFT	Manifolds and bordisms	Composition and disjoint union	Geometry is already in the domain.
Factorization algebras	Opens of a space or manifold	Disjoint-union products and descent	Opens and disjointness are already selected.
Haag and Kastler AQFT	Spacetime regions	Isotony and algebraic locality	The region poset presupposes spacetime localization.
Faithfully flat module descent	Rings and modules	Effective descent along a faithfully flat map	The theorem does not select physical contexts or covers.
Higher topoi	An infinity-site or infinity-topos	Homotopy-coherent descent	The physical site and interpretation remain inputs.

Table 2: Established comparators and the missing pre-geometric bridge.

8 Assumption audit

The assumption ledger prevents a formal definition from being mistaken for a physical derivation. Table 3 records exactly where each Paper 1 assumption enters.

ID	Content	Disposition in this paper
P1-A01	Chosen weak higher-categorical composition	OPEN . Only the one-categorical baseline and a strict word model are given.
P1-A02	Coverage stable under base change	Postulated in Theorem 3.1. No physical selection theorem is given.
P1-A03	Coefficient assignment is a presheaf, stack, or higher analogue	Typed in Theorems 4.1 and 4.7. Existence for intended physics is OPEN .
P1-A04	Chosen descent data are effective in the declared target	Used explicitly in Theorem 5.1. General effectiveness is OPEN .
P1-A05	Explicit structure-preserving comparator functors	OPEN . Section 7 gives only hypothesis audits.

ID	Content	Disposition in this paper
P1-A06	Geometry excluded from the primitive signature	Established by the signature audit in Section 2 . No geometric field occurs in the primitive data.
P1-A07	At least one model of the axioms lacks causal, dimensional, topological, and metric structure	UNVALIDATED . Theorem 6.4 concerns arbitrary diagnostic annotations, not the exact admissible non-geometric model required here.

Table 3: Paper 1 assumption audit. “Typed” means that an assumption is visible in the interface. It does not mean the assumption is physically justified.

8.1 Circularity checks

Three checks are useful when this interface is reused.

First, a cover must not be defined as a family on which the desired coefficient assignment already glues. That definition would make effective descent true by construction while hiding the physical selection problem.

Second, a realization functor must not attach a manifold to each context and then cite the manifold’s open-cover topology as a derivation of pre-geometric coverage. That would reverse the stated dependency order.

Third, a process arrow must not be renamed a causal arrow without checking the axioms required by the later causal reconstruction theorem. [Theorem 6.5](#) gives a direct finite obstruction to that move.

9 Executable evidence and formal boundary

9.1 Lean

The Lean modules are interfaces, not a kernel proof of the whole paper. The relevant declarations are:

- `PhysicalContext`, an object of an explicit category plus a claim record;
- `Cover`, a family of arrows into a target plus assumption IDs;
- `DescentData`, local and global data with a compatibility predicate and a restriction map;
- `EffectiveDescent`, an equivalence that agrees with restriction;
- `EffectiveDescent.restrict_reconstruct`, the checked inverse law.

The declaration `Cover` does not prove that a family is covering. It stores the family and its dependency identifiers. The declaration `EffectiveDescent` does not synthesize an equivalence. It requires one. These choices prevent the implementation from upgrading P1-C02 merely because the types compile.

The registry records all Paper 1 claims as open except where the durable claim ledger is separately updated. This is conservative. A later formalization of [Theorem 6.4](#) could place the finite expansion argument in Lean, but the paper proof does not depend on such a future artifact.

9.2 Haskell

The Haskell model represents a context by a finite list of atom labels. Concatenation implements the tensor-like composition. QuickCheck tests associativity and the empty-context identity. This is finite executable evidence for [Theorem 2.6](#).

A finite cover consists of a universe and a list of patches. The validator checks that:

- the universe and every patch have unique labels;
- patches are nonempty and distinct;
- every patch lies in the universe;
- the union of patches is the universe.

Local sections are checked for one section per patch, exact label coverage, and agreement on every shared label. The test suite includes both an agreeing family and a family that disagrees on an overlap. It also rejects an incomplete cover.

These tests support a narrow **COMPUTATIONAL** claim: the finite implementation distinguishes composition, covering, and compatibility and behaves as specified on generated inputs. They do not establish effective descent. There is no function in the Haskell module that constructs a global section from every compatible family.

9.3 Evidence matrix

10 Open reconstruction targets

10.1 Higher categorical coherence

The ordinary category in [Theorem 2.1](#) suppresses higher identifications between processes. A future version of P1-C01 must specify:

1. the dimension or truncation level of the higher category;
2. the weak composition operations;
3. associators, unitors, braidings, and higher coherence cells;
4. the relation between tensor composition and process composition;

Artifact or result	Status	Exact scope
Theorem 2.6	PROVED	All finite words under concatenation.
Context composition tests	COMPUTATIONAL	Generated finite words in the Haskell model.
Theorem 5.2	PROVED	Reconstruction from a supplied effective witness.
restrict_reconstruct	PROVED	One inverse identity for the Lean witness structure.
Finite covers and compatibility tests	COMPUTATIONAL	Validation and overlap agreement for the declared finite representation.
Theorem 6.4	PROVED	Local non-uniqueness of four kinds of diagnostic annotation over one locality reduct.
P1-C01, P1-C02, P1-C03, P1-C04	OPEN	Higher coherence, physically selected effective coverage, and explicit realization functors, plus an admissible non-geometric model for the stable non-determination claim.

Table 4: The evidence boundary. Executable examples are not promoted to theorem proofs.

5. a formal or complete mathematical coherence proof.

Assumption P1-A01 cannot be discharged by saying that coherence is standard. The selected model and its precise weak structure must be named. The word model proves only that the baseline is consistent in a strict one-dimensional case.

10.2 Physical selection of coverage

Open questions OQ-P1-02 and OQ-P1-04 ask what makes a cover physical. Several choices are possible. A cover might be generated by operational accessibility, by algebraic subcontext maps, or by a primitive locality relation. These choices need not agree.

The identity coverage shows that formal effectiveness is too weak as a selection principle. A coverage containing every family may be too strong or even ill behaved. The intended coverage needs a criterion that can fail on a concrete family. Without such a criterion, P1-C02 remains open even if many formal sheaf lemmas are available.

10.3 Ambient category for effectiveness

Counterexample CE-P1-01 shows that descent data can be non-effective in one category and effective in a larger one. Thus assumption P1-A04 must name the target. “Objects glue” is incomplete unless “objects” and the equivalence notion are fixed.

For set-valued coefficients, bijectivity is enough. For category-valued coefficients, equivalence rather than equality is appropriate. For infinity-categorical coefficients, homotopy-

coherent descent is required. Moving among these targets changes both the theorem and its proof obligations.

10.4 Realization functors

Claim P1–C03 requires explicit comparison functors. A future theorem should state a diagram such as

$$\begin{array}{ccc} \text{Ctx}_{\text{pre}}^{\text{op}} & \xrightarrow{F_{\text{pre}}} & \mathcal{D}_{\text{pre}} \\ R^{\text{op}} \downarrow & & \downarrow S \\ \text{Ctx}_{\text{geom}}^{\text{op}} & \xrightarrow{F_{\text{geom}}} & \mathcal{D}_{\text{geom}} \end{array}$$

together with a natural isomorphism or comparison transformation. It should say which covers R preserves and whether S preserves the limits used in descent. The diagram is a target specification. No functor in it is constructed here.

10.5 The next-paper interface

Paper 2 may consume only the following:

1. a declared context category and compositional operation;
2. a declared coverage satisfying its formal axioms;
3. presheaf and descent types;
4. the conditional reconstruction theorem from an effective witness;
5. the proved local non-uniqueness theorem for diagnostic annotations.

Paper 2 may not consume a causal order, topology, metric, or dynamics from this paper. It may not assume that every coefficient assignment is a sheaf. If it moves the data into a condensed setting, it must prove that the chosen functor preserves the limits and descent comparisons actually used.

11 Discussion

11.1 What is gained

The framework makes three decisions visible. Composition, coverage, and effective descent occupy different fields in the data. A user can therefore change the coverage without changing process composition, or change the coefficient target without changing the covers. The finite diagnostic witness shows why this separation matters.

The framework also gives later reconstruction papers a clean input boundary. A later geometric realization must add geometric structure through a named map and hypotheses. A later quantum reconstruction must add operational and probabilistic interfaces. A later kinetic model must add dynamics and a limiting regime. None of those additions can be cited as a consequence of locality alone.

11.2 What is not gained

The axioms do not explain why a physical system should furnish Ctx . They do not identify experiments with coefficient objects. They do not choose the coverage. They do not prove nontrivial effectiveness.

There is also no theorem here saying that every useful geometric theory factors through this interface. The comparator section identifies the data such a theorem would have to preserve. Until the functors are constructed, the comparison remains conceptual.

11.3 Novelty boundary

Category theory, sheaves, and descent are established mathematics. The examples reviewed in [Section 7](#) are established frameworks. The contribution here is narrower: an assumption-indexed interface for a reconstruction-first research program, together with a strict separation between formal locality and later geometric realization.

The finite diagnostic non-uniqueness theorem is elementary. Its value is diagnostic, not technical depth. It exposes four unsupported inferences at once but does not establish the corresponding realized-geometric independence claim. The paper does not present standard descent machinery as a new theorem.

11.4 Failure modes

The framework can fail in several concrete ways. The proposed context collection may not close under process composition. Required pullbacks may not exist. A proposed coverage may fail base-change stability. A coefficient assignment may fail functoriality. Compatible data may fail to be effective. A realization may fail to preserve covers or descent limits.

Each failure has a different repair. Enlarging the ambient category can repair effectiveness but may change the meaning of the objects. Restricting the cover class can restore a sheaf condition but may discard physically needed localizations. Adding higher morphisms can repair coherence but introduces new proof obligations. The ledger keeps these repairs attached to their assumptions.

12 Conclusion

Physical locality can be typed without starting from a manifold. The minimal baseline uses contexts, processes, a compositional operation, a coverage, coefficient presheaves, compatibility, and an explicit effective-descent comparison. This statement is about mathematical formulation. It is not yet a derivation of physical locality.

The complete results are modest. Word composition is coherent. Finite multiset composition is symmetric. An effective witness reconstructs compatible data uniquely. Identity coverage is effective. Products and naturally isomorphic copies of sheaves remain local. One finite locality reduct has incompatible order-shaped, number-valued, topology-on-label-set, and pseudometric annotations. They are not realized causal relations, spacetime dimensions, spacetime topologies, or Lorentzian metrics.

The stronger targets remain where the evidence leaves them. The higher-categorical model is **OPEN**. The physical selection of covers is **OPEN**. Nontrivial effectiveness for intended coefficients is **OPEN**. Explicit realization functors to geometric and algebraic comparators are **OPEN**. No causal order, dimension, topology, metric, or dynamics has been recovered.

A Axiom checklist

This appendix restates the baseline as a checklist suitable for later implementations.

A.1 Context and process data

1. A type or universe of contexts.
2. A hom-type $\text{Hom}_{\text{Ctx}}(U, V)$ for each ordered pair.
3. An identity process on every context.
4. A partial composition operation on composable processes.
5. Associativity and unit proofs.

These clauses define the ordinary category. If a higher category is selected, each equality must be replaced or supplemented by the appropriate coherent cells.

A.2 Tensor data

1. A bifunctor $\otimes : \text{Ctx} \times \text{Ctx} \rightarrow \text{Ctx}$.
2. A unit object I .
3. Associator and left and right unitor isomorphisms.
4. A symmetry isomorphism.
5. Pentagon, triangle, and hexagon equations.

No clause interprets $\otimes$ as spatial juxtaposition. That interpretation must be supplied later.

A.3 Coverage data

1. A class $\text{Cov}(U)$ of families into each target U .
2. Identity families.
3. Pullbacks required for base change.
4. Stability under base change.

5. Stability under composition of covering families.

No clause says that a cover is jointly surjective on points. There may be no primitive points.

A.4 Coefficient and descent data

1. A target category for coefficients.
2. A contravariant assignment on contexts and processes.
3. Categorical overlaps or a Čech object.
4. A compatibility object or category.
5. A comparison map from global data.
6. A stated equivalence notion.
7. An effectiveness theorem or an explicit assumption.

The final clause cannot be replaced by a compatibility check. This is the main type boundary enforced by the paper.

B Finite model details

B.1 Monoidal table

The tensor table for $B = \{0, 1\}$ is

$\otimes$	0	1
0	0	1
1	1	1

Associativity can be checked on eight triples. It also follows because maximum on a total order is associative. Commutativity is immediate from symmetry of maximum. The unit row and column show that 0 is the identity.

B.2 Coverage table

The discrete category has only id_0 and id_1 . The covers are

$$\{\text{id}_0\} \quad \text{and} \quad \{\text{id}_1\}.$$

There is no nonidentity base change to check. Pulling an identity cover back along an identity returns the same cover. Composing identity covers returns an identity cover.

B.3 Coefficient table

The terminal presheaf assigns a singleton to both objects:

$$\mathbf{1}(0) = \{*\}, \quad \mathbf{1}(1) = \{*\}.$$

Every restriction is the identity function on a singleton. For either identity cover, there is one matching family and one global element. The comparison is the unique bijection between singleton sets.

B.4 Diagnostic annotation table

Diagnostic annotation	Annotation A	Annotation B
Order-shaped annotation	Equality relation	Chain $0 \leq 1$
Number-valued annotation	$d(0) = d(1) = 0$	$d(0) = 0, d(1) = 1$
Topology annotation on finite label set	Indiscrete	Discrete
Pseudometric annotation on finite label set	$m(0, 1) = 1$	$m(0, 1) = 2$

Table 5: Pairs of diagnostic annotations on the same finite object-label set over one locality reduct. The identity on the reduct does not identify either pair.

C Proof dependency map

Result	Dependencies	Excluded inference
Theorem 2.2	Category axioms	Does not prove higher coherence.
Theorem 2.6	Finite list concatenation	Does not select physical atoms.
Theorem 2.7	Finite multiset sum	Does not interpret symmetry as spacelike separation.
Theorem 3.3	Presheaf functoriality	Does not produce a global glue.
Theorem 3.5	Compatibility and pullbacks	Does not put the refinement in the coverage.
Theorem 4.4	Pullback square and functoriality	Does not prove injectivity or surjectivity.

Result	Dependencies	Excluded inference
Theorem 5.2	Effective-descent witness	Does not construct the witness.
Theorem 5.4	Identity coverage	Does not justify a nontrivial physical coverage.
Theorem 5.6	Effectiveness of both factors	Does not establish effectiveness for either factor.
Theorem 5.8	Natural isomorphism and effectiveness of one side	Does not infer physical equivalence from formal isomorphism.
Theorem 6.4	Boolean reduct and paired diagnostic expansions	Does not produce realized geometric structures or an admissible non-geometric physical model.

Table 6: Proof dependencies and the stronger conclusions they do not license.

D Claim and assumption cross-reference

D.1 P1-C01

Dependencies: P1-A01 and P1-A02. The ordinary category and strict word model cover only a truncation of the claim. The weak higher category and its physical coverage remain open.

D.2 P1-C02

Dependencies: P1-A02, P1-A03, and P1-A04. The comparison map is typed and reconstruction from a witness is proved. The existence of a witness for the intended nontrivial coverage remains open.

D.3 P1-C03

Dependencies: P1-A05 and P1-A06. The primitive signature passes the no-geometry audit. The comparator functors and preservation theorems remain open.

D.4 P1-C04

Dependency: P1-A07. The Boolean locality reduct and its paired diagnostic annotations prove the local non-uniqueness theorem in Theorem 6.4. They do not realize causal or Lorentzian geometry and do not establish the stable claim. P1-C04 remains open, and P1-A07 is not established.

E Review checklist for reuse

A later paper that consumes this interface should answer all of the following.

1. Which exact category or higher category contains the contexts?
2. Which arrows are processes, and what physical reading is assigned to them?
3. Which tensor is used, and what does it mean?
4. Which families cover?
5. Why are those covers physical rather than merely convenient?
6. Which pullbacks or homotopy pullbacks exist?
7. What is the coefficient target?
8. What is the contravariant assignment?
9. What counts as compatibility?
10. In which category is descent claimed to be effective?
11. Is the comparison an equality, isomorphism, equivalence, or a weaker map?
12. Which theorem supplies effectiveness?
13. Which assumption IDs enter that theorem?
14. Which counterexample tests were attempted?
15. Which geometric or algebraic realization functors are constructed?
16. Which structures do those functors preserve?
17. Which claims remain open after the construction?

If an answer is absent, the later theorem should keep the corresponding bridge **OPEN**. If an answer is a physical postulate, it should be marked as such. If an answer is a mathematical theorem, its hypotheses and source must be given.

References

- [1] Michael F. Atiyah. Topological quantum field theory. *Publications Mathématiques de l'IHÉS*, 68:175–186, 1988. doi: 10.1007/BF02698547. URL <https://doi.org/10.1007/BF02698547>.
- [2] Kevin Costello and Owen Gwilliam. *Factorization Algebras in Quantum Field Theory*, volume 1. Cambridge University Press, 2017. doi: 10.1017/9781316678626. URL <https://doi.org/10.1017/9781316678626>.

- [3] Rudolf Haag and Daniel Kastler. An algebraic approach to quantum field theory. *Journal of Mathematical Physics*, 5(7):848–861, 1964. doi: 10.1063/1.1704187. URL <https://doi.org/10.1063/1.1704187>.
- [4] Jacob Lurie. *Higher Topos Theory*, volume 170 of *Annals of Mathematics Studies*. Princeton University Press, 2009. ISBN 978-0-691-14049-0. URL <https://arxiv.org/abs/math/0608040>.
- [5] Saunders Mac Lane and Ieke Moerdijk. *Sheaves in Geometry and Logic: A First Introduction to Topos Theory*. Universitext. Springer-Verlag, New York, 1992. ISBN 978-0-387-97710-2. doi: 10.1007/978-1-4612-0927-0. URL <https://doi.org/10.1007/978-1-4612-0927-0>.
- [6] The Stacks Project Authors. Descent, proposition 35.3.9, tag 023n. The Stacks Project, 2026. URL <https://stacks.math.columbia.edu/tag/023N>. Accessed 23 July 2026.
- [7] The Stacks Project Authors. Non-effective descent data for projective schemes, tag 08ke. The Stacks Project, 2026. URL <https://stacks.math.columbia.edu/tag/08KE>. Accessed 23 July 2026.