NOTE-THEOREM-HABITAT / PROOF-ARCHITECTURE RESEARCH NOTE

A Theorem Needs a Habitat

J. Scott Little

Source: July 2026 source edition Publication treatment: 8 September 2026

Research notes — not peer reviewed. Trust, but verify. Always.

When may a proof obligation move from a thick reconstruction to its thin limit?

A note distinguishing conditional theorem architecture from an inhabited physical model, with explicit rules for transferring proof obligations.

Scope and status

The note does not discharge the model-specific Yang–Mills obligations, construct the thin net, or prove a mass gap. “Paper IV will handle it” is valid only when the transferred task is actually a limiting obligation.

This edition adds a reading guide and publication notices. The source’s mathematical status and historical provenance remain attached to its claims.

A way into the note

The central distinction

A theorem proving that hypotheses imply a conclusion does not also prove that the intended physical system satisfies those hypotheses. The latter requires its own model certificate.

A concrete example

Suppose a limiting theorem assumes a diagram of positive-thickness objects. It can analyze their limit only after those objects and comparison arrows have been supplied; the limit cannot create its missing inputs.

What this approach gains

The dependency framework locates the exact stage that owns each obligation. It helps distinguish a legitimate limiting task from an unresolved finite-thickness construction hidden behind a later paper reference.

Reading route

Read the abstract and Sections 1–2, then follow the dependency-category and cut-criterion discussion. Use the application as an obligation map, with each witness restricted to its certified domain.

Citation: J. Scott Little, A Theorem Needs a Habitat, July 2026 source edition; website production edition, 8 September 2026. Cite theorem and printed body-page numbers from the note itself.

A Theorem Needs a Habitat

Conditional Reconstruction, Model Inhabitation, and the Transfer of Proof Obligations in Positive-Thickness Yang–Mills Theory

J. Scott Little
Research Notes on Yang–Mills Reconstruction

July 2026

Abstract

A long reconstruction argument can be logically complete as an abstract theorem while still failing to establish that the intended physical model exists. The gap is not cosmetic. A conditional architecture specifies which objects, morphisms, descent data, analytic domains, and geometric controls would suffice; an inhabited realization proves that a selected physical system actually supplies them. This distinction becomes unavoidable in a multi-paper positive-thickness Yang–Mills program. Paper III must construct the thick incised realization that Papers IV and V later use. A thin-limit theorem may analyze a verified diagram of positive-thickness objects, but it cannot retroactively create missing objects, arrows, Cauchy anchors, margins, or realization equivalences in that diagram.

This note formalizes that asymmetry. We introduce a dependency category of proof obligations, distinguish architecture certificates from model certificates, and give a cut criterion for deciding whether an obligation belongs to the thick reconstruction paper, may be discharged by an existing witness, is legitimately transferred to the thin-limit paper, or should be removed as an unnecessary claim. The framework is applied to the front end

$$P1 \to P2 \to P3 \to P4 \to P5 \to H1 \to H2 \to H3 \to \mathrm{CIR} \to \mathrm{RID/CP4} \to \mathrm{CPR}$$

of the positive-thickness Yang–Mills reconstruction. We explain why branchwise witnesses cannot be promoted beyond their domains, why complete-ground coercive provenance must be transported rather than merely renamed, and why the sentence "Paper IV will handle it" is valid only for genuinely limiting obligations. The resulting methodology is both a proof-hygiene principle and a structural insight: in a descent-based theory, a theorem architecture is analogous to a site, while an inhabited physical realization is a compatible section over that site.

Claim boundary. This note develops proof architecture and obligation-transfer criteria for a positive-thickness reconstruction program. It does not itself discharge the model-specific Yang–Mills obligations listed below, construct the zero-thickness Haag–Kastler net, prove Poincare covariance of the thin theory, establish a spectrum condition, or prove a Yang–Mills mass gap.

Contents

1 The hidden edge between architecture and existence **3**

1.1 Why the distinction is especially sharp in Yang–Mills reconstruction 4

2 Conditional architectures and inhabited realizations 4

2.1 The three layers of a trustworthy theorem statement 5

3 The thick–thin asymmetry 6

3.1 The direction of logical time 6

4 A dependency category of proof obligations 7

4.1 Proof-obligation descent 7

5 The four-way obligation classification 8

6 Application to the positive-thickness Yang–Mills front end 9

6.1 P1: regional time-zero extension 9

6.2 P2: selected Cauchy extraction and anchoring 9

6.3 P3 and P4: defect naturality and exactness 10

6.4 P5: fixed-scale and cofinal lifting 10

6.5 H1, H2, and H3: from a formal carrier to usable surgery geometry 10

6.6 CIR: category recovery 11

6.7 RID and CP4: identifying the constructed object 11

6.8 CPR: complete-ground coercive provenance 11

7 Branchwise inhabitation and the hierarchy of witnesses 11

7.1 The branch matrix 12

8 Exact, controlled, and limiting transport 12

9 Failure modes exposed by the obligation audit 13

9.1 Notation identity mistaken for object identity 13

9.2 Definition mistaken for sufficiency 13

9.3 Connected enhancement fibers mistaken for contractible fibers 14

9.4 A small auxiliary radius mistaken for the selected thickness 14

9.5 Objectwise positivity mistaken for cofinal uniformity 14

9.6 Observable strictness mistaken for representation neutrality 14

9.7 A later physical mechanism used to create an earlier habitat 14

10 A theorem-level workflow for obligation closure **14**

10.1 Wave 0: source control 14

10.2 Wave 1: classify before proving 15

10.3 Wave 2: prove in dependency order 15

10.4 Wave 3: transfer only across a proved cut 15

10.5 Wave 4: downgrade honestly when needed 15

11 Why this is more than manuscript hygiene **15**

11.1 Physical content lives at interfaces 15

11.2 The obstruction may lie in simultaneous coexistence 16

11.3 A conditional theorem can still be profound 16

12 Consequences for the division among Papers III, IV, and V **16**

13 A compact obligation-transfer theorem **17**

14 Research directions **17**

15 Conclusion **18**

A Obligation classification template **18**

B A first-pass Paper III versus Paper IV boundary **19**

1 The hidden edge between architecture and existence

A theorem often arrives in two stages that are easy to conflate.

First, one proves an implication of the form

$$\mathsf{H}_1 \wedge \cdots \wedge \mathsf{H}_N \Longrightarrow \mathsf{C}. \tag{1}$$

This is a genuine mathematical theorem. It may contain deep geometry, operator theory, higher descent, and analytic control. It determines exactly what would follow from a specified collection of hypotheses.

Second, one must show that the physical model under discussion satisfies those hypotheses. That is a different theorem:

$$\mathfrak{M}_{\mathrm{phys}} \models \mathsf{H}_1, \ldots, \mathsf{H}_N. \tag{2}$$

Only the combination of (1) and (2) yields

$$\mathfrak{M}_{\mathrm{phys}} \models \mathsf{C}. \tag{3}$$

The first theorem constructs an architecture. The second proves that the architecture is inhabited by the intended model.

A theorem needs a habitat. An abstract reconstruction theorem identifies a class of admissible habitats. A model theorem must prove that the intended physical data actually occupy one of them.

This distinction is familiar in elementary mathematics. A theorem about complete metric spaces does not show that a proposed space is complete. A theorem about self-adjoint operators does not show that a given formal differential expression has a self-adjoint realization on the claimed domain. A descent theorem does not show that a proposed local packet satisfies the descent hypotheses. The distinction becomes harder to see when the architecture itself is several hundred pages long, because the conditional theorem may feel like the whole proof.

In a reconstruction program, however, the missing inhabitation theorem is often exactly where the physics resides.

1.1 Why the distinction is especially sharp in Yang–Mills reconstruction

The positive-thickness Yang–Mills program does not begin with one global Hilbert space and one already defined Hamiltonian. It begins with finite physical packages, regional observable systems, selected scale transport, Cauchy presentations, chart-local representations, complete-ground projections, analytic cores, coefficient lines or correspondences, and higher coherence data. The program must show that these ingredients coexist on a common positive-thickness spacetime habitat.

The front-end dependency chain is schematically

$$\begin{aligned} &\text{finite physical time-zero packages} \longrightarrow P1 \longrightarrow P2 \longrightarrow P3/P4 \longrightarrow P5 \\ &\longrightarrow \Sigma_{\mathrm{obs}}^{\Gamma_{\mathrm{cof}}} \longrightarrow H1/H2/H3 \longrightarrow \mathrm{CIR} \longrightarrow \mathrm{RID/CP4} \longrightarrow \mathrm{CPR}. \end{aligned} \tag{4}$$

The resulting thick realization is then passed to a thin reconstruction theorem.

If any indispensable arrow in (4) is missing for the selected physical branch, the later limit theorem does not receive its input.

2 Conditional architectures and inhabited realizations

Definition 2.1 (Conditional reconstruction architecture)**.** *A* conditional reconstruction architecture *is a tuple*

$$\mathsf{Arch} = (\mathcal{B}, \mathcal{D}, \mathcal{T}, \mathcal{H}, \mathcal{C}) \tag{5}$$

consisting of:

1. *a base category or site $\mathcal{B}$ of regions, scales, presentations, or backgrounds;*
2. *a declared data type $\mathcal{D}$ over $\mathcal{B}$;*
3. *admissible transport and coherence data $\mathcal{T}$;*
4. *a family of hypotheses $\mathcal{H}$;*
5. *a conclusion constructor $\mathcal{C}$ assigning a reconstructed object whenever $\mathcal{H}$ is satisfied.*

Definition 2.2 (Model inhabitation)**.** *Let $\mathfrak{M}$ be a proposed physical model. An* inhabitation certificate *for $\mathfrak{M}$ in* Arch *is a realization*

$$\rho_{\mathfrak{M}} : \mathfrak{M} \longrightarrow (\mathcal{D}, \mathcal{T}) \tag{6}$$

together with proofs that every hypothesis in $\mathcal{H}$ holds for the realized data, with the declared domains, supports, branches, scales, and comparison maps.

Definition 2.3 (Inhabited reconstruction theorem)**.** *A reconstruction theorem is* inhabited for the branch $\mathfrak{s}$ *when it contains both:*

1. *the conditional architecture theorem;*
2. *a branch-specific inhabitation certificate $\rho_{\mathfrak{s}}$.*

Remark 2.4. *An explicit example may establish nonemptiness of the abstract model class without inhabiting the intended physical branch. Conversely, a physical branch may satisfy the local data conditions without satisfying cofinal or limiting uniformity. "There is an example" and "the intended branch is realized" are logically different statements.*

2.1 The three layers of a trustworthy theorem statement

A useful theorem statement separates three layers.

Layer	**Content**
Architecture	If data of the declared type satisfy the hypotheses, the reconstruction exists and has the stated properties.
Model verification	A named physical branch supplies the data and satisfies the hypotheses.
Downstream consequence	The reconstructed object supports later limiting, covariance, or spectral analysis.

A manuscript becomes difficult to audit when these layers are compressed into one omnibus theorem. In particular, the phrase "under the standing assumptions" can hide whether a condition is an imported theorem, a newly proved result, a model-specific proof obligation, or a genuinely downstream hypothesis.

3 The thick–thin asymmetry

The central logical asymmetry can be stated without any Yang–Mills-specific notation.

Let Γ be a selected cofinal category of positive scales and let

$$F : \Gamma \longrightarrow \mathcal{C} \tag{7}$$

be the diagram of positive-thickness objects and admissible comparison maps. A thin theory is then sought as a limit, colimit, completion, or reconstructed object derived from F, for example

$$F_0 \simeq \operatorname*{holim}_{r\in\Gamma} F(r), \qquad \text{or} \qquad F_0 \simeq \operatorname*{hocolim}_{r\in\Gamma} F(r), \tag{8}$$

possibly after applying observable, representation, state, or form functors.

Theorem 3.1 (No retroactive thin-limit discharge)**.** *A theorem asserting existence or properties of a thin object constructed from the diagram F cannot discharge an obligation required to define an object $F(r)$, a morphism $F(a)$, or a coherence cell of F.*

Proof. The expressions in (8) are defined only after the diagram F has been specified. A missing object, arrow, or coherence cell is not a failed convergence property of a defined diagram; it is a failure to define the input diagram itself. Any construction of the thin object therefore presupposes the relevant thick existence and functoriality data. A later theorem may prove compactness, convergence, independence, or no leakage for the diagram, but it cannot serve as the prior definition of that diagram without circularity. □

Corollary 3.2 (Thick existence belongs before thin reconstruction)**.** *If an obligation is necessary to construct the positive-thickness regional net, its selected Cauchy presentations, its lifting stack, its incised category, its realization equivalence, or its objectwise complete-ground certificate, then that obligation belongs to the positive-thickness paper.*

Remark 3.3. *The theorem does not say that every estimate must be uniform in the thick paper. It says that the thick object and every structure used to define it must exist there. Uniformity needed only for convergence may legitimately be deferred.*

3.1 The direction of logical time

The proof dependency has a temporal orientation:

$$\text{finite input} \longrightarrow \text{thick object} \longrightarrow \text{thin limit} \longrightarrow \text{spectral transfer}. \tag{9}$$

A valid proof may use earlier objects to define later ones. It may not use a property of the later object to create the earlier input unless an independent fixed-point theorem is explicitly formulated and proved. In the present reconstruction, no such reverse fixed-point theorem is claimed.

Warning 3.4 (The phrase "Paper IV will handle it")**.** *The phrase is valid only when "it" denotes convergence, compactness, limiting independence, no leakage, or another property of an already constructed thick diagram. It is invalid when "it" denotes existence of the diagram, its regional arrows, its selected Cauchy anchors, its fixed-scale lifting objects, or the geometric margins required to define them.*

4 A dependency category of proof obligations

The obligation audit can be formalized as a finite directed category.

Definition 4.1 (Obligation category)**.** *An* obligation category $\mathcal{O}$ *has:*

- *objects given by proof obligations;*
- *a morphism $A \to B$ whenever a certificate for B uses a certificate for A;*
- *composites given by transitive proof dependence.*

Acyclicity is expected after quotienting mutually equivalent formulations.

For the positive-thickness front end, a coarse skeleton is

$$P1 \to P2 \to P3 \to P4 \to P5 \to H1 \to H2 \to H3 \to \mathrm{CIR} \to \mathrm{RID/CP4} \to \mathrm{CPR}. \tag{10}$$

The actual category is finer. Each node contains subobligations concerning objects, morphisms, 2-cells, domains, supports, branches, uniformity scopes, and comparison maps.

Definition 4.2 (Certificate assignment)**.** *A* certificate assignment *is a functorial rule*

$$\mathsf{Cert} : \mathcal{O} \longrightarrow \mathcal{K} \tag{11}$$

into a category $\mathcal{K}$ of proof objects, source references, computations, witness packages, or explicit constructions. A node is discharged only when its assigned certificate has the exact type required by that node.

The type requirement matters. A finite-dimensional witness cannot be inserted into a node demanding a cofinal family. A line-bundle cocycle computation cannot discharge a graph-domain preservation statement. A normal algebra isomorphism cannot by itself transport an unbounded form domain.

4.1 Proof-obligation descent

There is a useful, though deliberately limited, analogy with descent.

Local proof certificates may exist for separate pieces of a theorem:

$$c_i \in \mathsf{Cert}(U_i). \tag{12}$$

To produce a global model theorem, they must agree on overlaps in the sense relevant to the proof architecture:

$$c_i|_{U_{ij}} \simeq c_j|_{U_{ij}}. \tag{13}$$

If comparison certificates and higher compatibility are required, the proof package itself has a descent problem.

Principle 4.3 (Proof-obligation descent)**.** *A collection of locally valid lemmas is not yet a global model-verification theorem unless the lemmas share the same objects, domains, branches, scales, and comparison maps and their interfaces commute.*

This principle explains why very strong separate results can fail to close a reconstruction theorem. One result may use a strict branch, another a twisted branch; one may use a fixed stratum, another a cofinal tail; one may use a complete-ground projection, another a low-energy island; one may work on a common core, another only after closure. The lemmas do not glue merely because each is true.

Remark 4.4. *The analogy does not assert that every proof ledger carries a canonical Grothendieck topology. It is a structural model: the architecture declares what counts as a cover and what compatibility is required for local certificates to assemble.*

5 The four-way obligation classification

Every explicit obligation should be placed in exactly one of four classes.

Definition 5.1 (Obligation status)**.** *For an obligation $Q \in \mathrm{Ob}(\mathcal{O})$, define*

$$\mathsf{Status}(Q) \in \{\mathrm{III-A}, \mathrm{III-B}, \mathrm{IV}, \mathrm{NC}\}, \tag{14}$$

where:

1. III− A*: must be proved in Paper III for the selected positive-thickness realization;*
2. III− B*: is already discharged by an existing theorem, witness, or imported construction of the correct type;*
3. IV*: is genuinely a thin-limit obligation and may be transferred to Paper IV;*
4. NC*: is unnecessary for the principal theorem or is an unsupported overclaim and should be removed or weakened.*

Theorem 5.2 (Dependency-cut criterion)**.** *An obligation Q is transfer-admissible from Paper III to Paper IV only if both conditions hold:*

1. *no construction or theorem internal to Paper III uses Q to define or prove existence of a positive-thickness object, arrow, domain, category, comparison cell, or branch certificate;*
2. *every downstream use of Q concerns a limiting property of an already defined positive-thickness diagram.*

Proof. If the first condition fails, Definition 3.1 applies. If the second fails, the obligation is not purely limiting and its transfer changes the type or scope of the Paper III output. Conversely, if both conditions hold, Paper III can pass the fully defined thick diagram and record Q as an explicit hypothesis of the thin-limit theorem without circularity. □

Corollary 5.3 (Typical Paper IV obligations)**.** *Uniform tightness on a selected thinning tail, expectation-algebra convergence, admissible thin descent, Mosco convergence, Orlicz or entropy tail control, no leakage, cofinal independence of the thin net, and limiting covariance are transfer candidates, provided the thick objects and maps they concern are already constructed.*

Corollary 5.4 (Typical nontransferable obligations). *Regional extension, selected Cauchy anchoring, local defect exactness, fixed-scale lifting-stack nonemptiness, carrier regularity needed for incision, finite adapted atlases, positivity of the actual fixed-stage thickening margin, category recovery, thick realization identification, CP4, and objectwise coercive provenance are not Paper IV obligations.*

6 Application to the positive-thickness Yang–Mills front end

The integrated front end contains eleven major interfaces. Each must be read as a bundle of typed obligations rather than a slogan.

6.1 P1: regional time-zero extension

P1 asks whether objectwise finite physical packages extend to a regional prestack. The essential burden is not the existence of the finite packages but the construction of restriction maps preserving:

- physical Wilson quotients;
- complete-ground projections;
- common analytic cores;
- Gauss and Ward data;
- coefficient lines or correspondences;
- branch markings;
- selected scale compatibility.

P1 therefore belongs to $\mathrm{III}-\mathrm{A}$ unless an existing Paper I–II theorem already supplies exactly this regional structure. Objectwise coercivity does not by itself imply regional functoriality.

6.2 P2: selected Cauchy extraction and anchoring

P2 identifies spacetime candidate data with selected regional time-zero data through tuples

$$\mathfrak{c} = (N, S, \iota, B). \tag{15}$$

The burden includes existence, extraction, anchoring, finite comparison zigzags, and preservation of complete-ground and analytic data. Since later thick reconstruction uses these anchors, P2 is a Paper III obligation.

The restriction to selected admissible Cauchy slices is substantive. Proving the theorem for one selected comparison groupoid is not the same as proving unrestricted Cauchy independence.

6.3 P3 and P4: defect naturality and exactness

P3 constructs the common-core defect calculus and proves naturality. P4 must prove the converse direction: the declared vanishing conditions are sufficient to reconstruct an effective anchored realization.

A frequent proof-architecture error is to define the zero-defect subcategory to be the realizable subcategory and then cite the definition as a sufficiency theorem. A genuine P4 proof must construct the lift and prove that extraction and reconstruction are inverse up to the stated higher cells.

Both P3 and P4 are Paper III obligations because they define which thick local candidates are physically realizable.

6.4 P5: fixed-scale and cofinal lifting

P5 contains several scopes:

$$\begin{aligned}&\text{local objectwise lifting,}\\&\text{fixed-scale effective descent,}\\&\text{selected scale pseudocone compatibility,}\\&\text{full multiparameter cofinal uniformity.}\end{aligned}\tag{16}$$

The first three are needed to define the thick branch and therefore belong in Paper III. Uniformity required only to take the thin limit may be transferred. The distinction must be made clause by clause.

6.5 H1, H2, and H3: from a formal carrier to usable surgery geometry

The dynamically derived carrier

$$\Sigma_{\mathrm{obs}}^{\Gamma_{\mathrm{cof}}} = M \setminus M_{\mathrm{reg}}^{\Gamma_{\mathrm{cof}}}\tag{17}$$

is formally closed under suitable hypotheses. The later surgery, however, uses more: stratification, tubular neighborhoods, finite adapted atlases, collars, and positive clearance margins.

H1 asks for regular support, H2 for finite obstruction-adapted atlases, and H3 for positive numerical margins. These are not interchangeable.

Interface	**Qualitative role**	**Forbidden promotion**
H1	regularity of the derived carrier	closedness $\not\Rightarrow$ Whitney stratification or collarability
H2	finite adapted atlas formation	local charts $\not\Rightarrow$ a common finite atlas for every finite diagram
H3	positive margins for the selected thickening	existence of some small radius $\not\Rightarrow$ the actual selected ε_r fits

All fixed-stage geometry required to define the thick incised category belongs to Paper III. Only uniform control along the thinning tail may be transferred.

6.6 CIR: category recovery

Controlled incised category recovery must establish that decorated pre-incision certificates localize to the intended unadorned geometric category. It therefore requires recognition, functoriality, certificate refinement, contractible fibers, full faithfulness, and coverage.

A local equivalence on a convenient test subcategory is not automatically a recovery theorem for the basis used by the remainder of Paper III. Coverage is part of the theorem, not editorial bookkeeping.

6.7 RID and CP4: identifying the constructed object

RID compares the lifting gerbe with the retained Paper III realization target. It must preserve objects, morphisms, 2-cells, regional maps, selected Cauchy data, scale data, mixed cells, complete-ground projections, analytic cores, coefficient lines, roots, and branch markings.

CP4 concerns presentation independence of the ordinary regional observable algebra. It does not imply that the represented gerbe is neutral.

These are indispensable Paper III obligations. Without RID, the front-end construction and the later Paper III object may be two parallel architectures with similar notation rather than the same realization.

6.8 CPR: complete-ground coercive provenance

CPR must transport the actual inherited certificate

$$(\mathcal{C}^{\mathrm{prov}}, q^{\mathrm{prov}}, P_\Omega, I - P_\Omega, \mathbf{q}, \widehat{\kappa}) \tag{18}$$

through the regional, Cauchy, descent, stackification, and RID interfaces.

This is stronger than observing that the target notation also contains a symbol P_Ω or q. For an exact unitary comparison U, one must define

$$\mathcal{D}(q') = U\mathcal{D}(q), \qquad q'[U\psi] = q[\psi], \qquad P'_\Omega = UP_\Omega U^*. \tag{19}$$

For a coefficient line L, the scalar extension must be explicit:

$$q_L[\psi \otimes \ell] = q[\psi]\|\ell\|^2, \qquad P_{\Omega,L} = P_\Omega \otimes 1_L. \tag{20}$$

A controlled one-sided defect does not imply preservation of the same coercive constant.

Objectwise provenance belongs to Paper III. Uniformity over a full thinning tail may belong to Paper IV if it is not already a Paper II theorem on the selected branch.

7 Branchwise inhabitation and the hierarchy of witnesses

The word “witness” covers several logically different objects.

Definition 7.1 (Witness scope)**.** *A witness W has a scope tuple*

$$\mathrm{scope}(W) = (\textit{group}, \textit{branch}, \textit{scale range}, \textit{regions}, \textit{analytic level}, \textit{coherence level}). \tag{21}$$

Proposition 7.2 (Witness nonpromotion)**.** *Let W certify an obligation on a subdiagram $D' \subset D$. Then W certifies the corresponding obligation on D only if an extension theorem from D' to D is independently proved.*

Proof. The witness provides data only on D'. An assertion on D additionally requires data and compatibility on $D \setminus D'$ and across the interface. Those are precisely the contents of an extension theorem. □

Thus an $SU(2)$ fixed-stratum tensor-depth witness may certify finite package nonemptiness and depth-uniform control at one stratum. It does not by itself certify:

- the selected physical $SU(3)$ branch;
- thinning or cutoff removal;
- the physical linked-flux law;
- full cofinality;
- every P1–P5 or H1–H3 obligation.

Similarly, an $SU(3)$ linked-flux computation may certify a fixed-thickness coefficient law or sector detector without proving regional extension, Cauchy anchoring, graph-domain compatibility, or positive margin geometry.

7.1 The branch matrix

A trustworthy audit records obligations branchwise. A minimal matrix should include:

Branch	**Thick local**	**Cofinal**	**Thin**	**Spectral**
selected physical $SU(3)$	status required	status required	Paper IV	Paper V
strict $SU(2)$ depth witness	partial witness	not automatic	not automatic	not automatic
center-banded branch	branchwise	branchwise	Paper IV	Paper V
abstract conditional branch	conditional	conditional	conditional	conditional

The matrix prevents proofs from drifting between branches while retaining one common notation.

8 Exact, controlled, and limiting transport

Many obligation errors arise from collapsing three distinct comparison regimes.

Definition 8.1 (Exact transport). *A comparison is* exact *when the relevant object, domain, form, projection, or coherence datum is intertwined by an isomorphism, unitary, equivalence, or invertible higher cell of the declared type.*

Definition 8.2 (Controlled transport). *A comparison is* controlled *when a defect is bounded in a specified norm, graph norm, form norm, or categorical metric, but need not vanish.*

Definition 8.3 (Limiting transport). *A comparison is* limiting *when a family of exact or controlled comparisons converges in a declared topology and the limiting object is reconstructed from that convergence.*

These regimes support different conclusions.

Regime	**Typical valid conclusion**	**Typical invalid promotion**
Exact	identical scalar form value, transported projection, same certificate	no issue if domains are explicitly transported
Controlled	stability with a quantified degraded bound under two-sided hypotheses	same constant from a one-sided defect
Limiting	closed-form or algebra convergence after compactness/tightness	existence of the prelimit objects

The distinction also explains the paper boundary. Paper III primarily establishes exact and fixed-stage controlled transport. Paper IV studies limiting transport. Paper V studies spectral transfer after the limiting representation has been constructed.

9 Failure modes exposed by the obligation audit

The audit is valuable because it detects errors that ordinary line-by-line proofreading often misses.

9.1 Notation identity mistaken for object identity

Using the same symbol in two sections does not prove that the objects coincide. A source projection and a retained projection require an explicit comparison theorem. The same applies to forms, coefficient lines, branches, and scale maps.

9.2 Definition mistaken for sufficiency

Defining the realizable subcategory as the zero locus of a defect is legitimate only after a theorem proves that zero defect reconstructs the desired object. Otherwise the definition conceals the missing existence argument.

9.3 Connected enhancement fibers mistaken for contractible fibers

Forgetting auxiliary anchors or roots is harmless only when the enhancement fibers have the homotopy type required by the claimed equivalence. Connectedness alone does not remove isotropy.

9.4 A small auxiliary radius mistaken for the selected thickness

A local geometric lemma may prove that some $\delta > 0$ works. The reconstruction uses the already selected ε_r. H3 must prove $\varepsilon_r < \delta$, possibly after passing to an admissible tail.

9.5 Objectwise positivity mistaken for cofinal uniformity

The statements

$$\widehat{\kappa}_r > 0 \quad \text{for every } r \tag{22}$$

and

$$\inf_{r \in \Gamma_{\mathrm{cof}}} \widehat{\kappa}_r > 0 \tag{23}$$

are not equivalent. A family of positive constants may tend to zero.

9.6 Observable strictness mistaken for representation neutrality

Presentation independence of the ordinary algebra can coexist with nontrivial coefficient-line or correspondence twisting in its represented realizations. CP4 does not collapse the three-sector structure.

9.7 A later physical mechanism used to create an earlier habitat

Collar charges, boundary Ward characters, entropy surgery, post-incision state selectors, closed spacetime forms, invariant momentum, and spectral thresholds are later structures. Using them to prove the existence or regularity of the incision reverses the dependency chain.

10 A theorem-level workflow for obligation closure

The obligation audit should be performed as a proof construction, not as a checklist exercise.

10.1 Wave 0: source control

Identify the authoritative source, exact versions, labels, dependencies, and unavailable artifacts. No theorem should be repaired against an obsolete source without an explicit reason.

10.2 Wave 1: classify before proving

Assign every clause to III− A, III− B, IV, or NC. Classification stabilizes the proof boundary before local edits create new circularities.

10.3 Wave 2: prove in dependency order

A natural order is

$$P1, P2\ ;\quad P3, P4, P5\ ;\quad H1, H2, H3\ ;\quad \mathrm{CIR}\ ;\quad \mathrm{RID/CP4}\ ;\quad \mathrm{CPR}. \tag{24}$$

Each wave should produce manuscript-ready statements, source audits, integration instructions, coverage matrices, compile audits, and continuation state.

10.4 Wave 3: transfer only across a proved cut

After Paper III obligations are closed, construct the Paper III–Paper IV interface. The transferred datum should be a typed thick package, not a prose promise. Schematically,

$$\mathfrak{P}_{\mathrm{III}\to\mathrm{IV}} = \left(\mathcal{A}_\varepsilon^W, \mathfrak{G}_\varepsilon, P_{\Omega,\varepsilon}, q_\varepsilon^{\mathrm{prov}}, \mathcal{T}_{\varepsilon'\varepsilon}, \mathcal{O}_{\mathrm{IV}}\right), \tag{25}$$

where $\mathcal{O}_{\mathrm{IV}}$ is an explicit list of limiting obligations rather than an implicit assumption that all limits exist.

10.5 Wave 4: downgrade honestly when needed

If an indispensable Paper III obligation remains open, the theorem status must be downgraded. The correct outcome may be:

$$\text{conditional architecture only,} \tag{26}$$

or

$$\text{partially inhabited thick construction.} \tag{27}$$

This is not failure. It is the exact mathematical boundary of the current proof.

11 Why this is more than manuscript hygiene

The distinction between architecture and inhabitation reflects a deeper feature of modern mathematical physics.

11.1 Physical content lives at interfaces

The hardest part of a reconstruction is often not proving a powerful theorem inside one formalism. It is proving that objects constructed in different formalisms are the same objects:

- finite Hamiltonian packages and regional prestacks;
- time-zero data and spacetime Cauchy data;
- local lifting objects and retained W*-realizations;
- source quadratic forms and target preforms;
- chartwise representations and global observable algebras;
- fixed-thickness sectors and scale-coherent branches.

These are interface theorems. They carry the ontology of the theory.

11.2 The obstruction may lie in simultaneous coexistence

Each component of a proposed physical package may exist separately. The obstruction can arise because they cannot be chosen simultaneously with all required naturality and coherence. Thus model inhabitation is not reducible to checking a list of isolated properties. It is a gluing theorem for the complete packet.

11.3 A conditional theorem can still be profound

The demand for inhabitation does not diminish the value of a conditional architecture. A well-typed conditional theorem can reveal the correct site, the correct obstruction carrier, the correct higher coherence, and the exact analytic interfaces. It converts an undefined physical problem into a finite family of explicit proof obligations.

The mistake is not writing a conditional theorem. The mistake is presenting it as an inhabited physical theorem without the model certificate.

12 Consequences for the division among Papers III, IV, and V

The paper boundary should follow the logical type of the objects.

Principle 12.1 (Paper III: thick existence)**.** *Paper III must construct and identify the selected positive-thickness realization, including the dynamically derived incision, the controlled regional category, the ordinary observable net, the represented descent packet, and objectwise complete-ground provenance.*

Principle 12.2 (Paper IV: thin reconstruction)**.** *Paper IV must prove that the selected thick system admits an admissible thin limit with the required convergence, compactness, no-leakage, locality, time-slice, and cofinal-independence properties.*

Principle 12.3 (Paper V: spectral transfer)**.** *Paper V may transfer positive-thickness coercive information to the limiting representation only after the thin object, its dynamics, its complete-ground sector, and the no-leakage interfaces have been constructed.*

The division is therefore not merely expository. It is a factorization of the proof by logical type:

$$\boxed{\text{existence} \longrightarrow \text{limit} \longrightarrow \text{spectrum.}} \tag{28}$$

13 A compact obligation-transfer theorem

The preceding discussion can be summarized in one theorem.

Theorem 13.1 (Architecture–inhabitation–transfer theorem)**.** *Let $\mathcal{O}$ be the obligation category of a multi-stage reconstruction, and let*

$$\mathcal{O}_{\mathsf{Thick}}, \quad \mathcal{O}_{\mathsf{Thin}}, \quad \mathcal{O}_{\mathrm{spec}} \tag{29}$$

be the full subcategories assigned to thick existence, thin reconstruction, and spectral transfer. Suppose:

1. *every predecessor in $\mathcal{O}$ of a thick-existence obligation lies in $\mathcal{O}_{\mathsf{Thick}}$;*
2. *every obligation transferred to $\mathcal{O}_{\mathsf{Thin}}$ satisfies the dependency-cut criterion of Definition 5.2;*
3. *every spectral obligation receives a fully constructed thin object and a verified no-leakage interface;*
4. *branch labels, domains, scales, and comparison maps are preserved by every certificate assignment.*

Then the proof factorizes without circularity as

$$\mathsf{Cert}_{\mathsf{Thick}} \longrightarrow \mathsf{Cert}_{\mathsf{Thin}} \longrightarrow \mathsf{Cert}_{\mathrm{spec}}. \tag{30}$$

If, in addition, $\mathsf{Cert}_{\mathsf{Thick}}$ contains a certificate for a selected physical branch, the thick theorem is inhabited on that branch. Otherwise the thick theorem remains conditional regardless of the completeness of the later formal architecture.

Proof. Condition 1 prevents a thick object from depending on a later-stage certificate. Condition 2 ensures that transferred obligations concern only limiting properties of defined thick data. Condition 3 prevents spectral conclusions from being attached to an unidentified or leaking limiting sector. Condition 4 prevents certificates from changing type or branch under transport. The factorization follows by composition of the three certificate functors. The final statement is the definition of model inhabitation. □

14 Research directions

The obligation framework suggests several broader questions.

Research Question 14.1 (Canonical obligation categories)**.** *Can a proof assistant or manuscript compiler extract a canonical dependency category from labelled theorem hypotheses, references, and source-level domains?*

Research Question 14.2 (Homotopy type of certificate spaces)**.** *When auxiliary choices such as charts, roots, anchors, or presentations are forgotten, what homotopy conditions on the certificate space are sufficient to preserve the reconstructed object?*

Research Question 14.3 (Quantitative transfer defects). *Can one develop a general two-sided perturbation calculus that assigns a controlled degradation of coercive constants across nonexact comparison seams?*

Research Question 14.4 (Obligation descent for computer-assisted proofs). *How should machine certificates, finite witnesses, and symbolic computations be typed so that they glue to analytic and geometric proofs without scope promotion?*

Research Question 14.5 (Model-class stratification). *Can the admissible model class itself be stratified by which obligation subcategories are inhabited, thereby separating strict, scale-twisted, and fixed-thickness-gerbe-twisted phases before the thin limit?*

15 Conclusion

The distinction developed here is simple to state but structurally decisive:

$$\boxed{\text{A conditional reconstruction architecture is not yet an inhabited physical theory.}} \tag{31}$$

For positive-thickness Yang–Mills reconstruction, this means that Paper III must prove the existence and identification of the thick object it passes forward. Paper IV may prove that this object converges, survives, and becomes a thin Haag–Kastler net. Paper V may then analyze spectral transfer. The later papers cannot retroactively create the input objects required by the earlier one.

The obligation audit is therefore not a bureaucratic exercise. It identifies the exact location of physical existence inside a large formal argument. It separates abstract possibility from realized structure, local certificates from compatible global proof, and finite witnesses from the branches they genuinely inhabit.

The architecture tells us what a world of solutions would have to look like. The inhabitation theorem proves that the selected Yang–Mills branch actually lives there.

A Obligation classification template

For each obligation Q, record the following fields.

Field	Required entry
Identifier	Exact theorem, hypothesis, or subclause label.
Object type	Region, algebra, Hilbert packet, form, projection, category, 2-cell, margin, or limit datum.
Branch	Strict, center-banded, fixed-thickness-gerbe-twisted, correspondence-valued, or abstract.
Scale scope	Objectwise, fixed scale, finite demand, selected tail, or full cofinal family.

Field	Required entry
Dependencies	Immediate predecessor obligations.
Status	III− A, III− B, IV, or NC.
Certificate	Exact theorem, construction, computation, or witness.
Interface proof	How the certificate is identified with the object used downstream.
Uniformity	Exact quantified scope of constants or margins.
Nonclaims	Stronger conclusions explicitly excluded.
Manuscript action	Prove, cite, transfer, weaken, or delete.

B A first-pass Paper III versus Paper IV boundary

Obligation	Default home	Reason
Regional time-zero restriction maps	Paper III	define the thick regional input
Selected Cauchy extraction and anchor	Paper III	define the thick spacetime/time-zero interface
Common-core defect naturality	Paper III	define physical comparison of local candidates
Zero-defect reconstruction/effectivity	Paper III	establish local realizability
Fixed-scale lifting-stack nonemptiness	Paper III	construct thick objects
Selected scale pseudocone	Paper III	construct the thick diagram
Carrier regularity for incision	Paper III	define the surgery habitat
Finite adapted atlas	Paper III	define local surgery charts
Positive margin for actual ε_r	Paper III	validate selected thickening
Uniform margin on thinning tail	Paper IV	limiting uniformity
Controlled incised category recovery	Paper III	identify thick indexing category
RID/CP4	Paper III	identify constructed and retained thick objects
Objectwise coercive provenance	Paper III	identify inherited thick certificate
Cofinal uniform coercivity	Paper II or IV	depends on exact imported profile status
Expectation-algebra convergence	Paper IV	limiting assertion
Mosco/strong-resolvent convergence	Paper IV	limiting form/operator assertion
No leakage of complete-ground sector	Paper IV	limiting sector stability
Thin Haag–Kastler net	Paper IV	thin reconstruction output
Invariant-mass spectral transfer	Paper V	spectral consequence

Obligation	Default home	Reason
Confinement and scattering	Separate work	not implied by the reconstruction chain alone

References

[1] R. Haag, *Local Quantum Physics: Fields, Particles, Algebras*, 2nd ed., Springer, Berlin, 1996.

[2] T. Kato, *Perturbation Theory for Linear Operators*, 2nd ed., Springer, Berlin, 1976.

[3] S. Mac Lane, *Categories for the Working Mathematician*, 2nd ed., Springer, New York, 1998.

[4] J. Giraud, *Cohomologie non abelienne*, Springer, Berlin, 1971.

[5] R. Brunetti, K. Fredenhagen, and R. Verch, "The generally covariant locality principle: a new paradigm for local quantum field theory," *Communications in Mathematical Physics* **237** (2003), 31–68.

[6] J. S. Little, "Coarse transport, central twist, and finite-scale coercivity for Yang–Mills Wilson-tube Hamiltonians," research manuscript, 2026.

[7] J. S. Little, "Scale-bordism criteria for uniform coercivity and center-banded vacuum descent in Yang–Mills Wilson-tube systems," research manuscript, 2026.

[8] J. S. Little, "Positive-thickness reconstruction for Yang–Mills Wilson-tube systems: local-normal nets and controlled center-banded descent," research manuscript, 2026.

[9] J. S. Little, "Entropy-surgered thick-to-thin reconstruction of the Yang–Mills observable system," research manuscript, 2026.

[10] J. S. Little, "Spectral transfer and no-leakage for thin Yang–Mills observable representations," research manuscript, 2026.