Finite-to-Total Reconstruction

Analytic Boundaries, Witnessed Totalization, and Relative Multiplication

J. Scott Little
Independent Researcher

Status. An operation may be analytically well behaved on every finite portion of a source and nevertheless fail to survive passage to its natural total realization. This paper separates source exposure from continuation demand and constructs the resulting strict core reconstruction bifunctor and analytic admissible region. For relative multiplication in the Jones basic construction, infinite E_A-orthogonal multiplicity gives bounded finite witness windows with exact norm $\|m_F\| = \sqrt{|F|}$ while the directed totalization is maximally nonclosable, with $\mathrm{mul}\,\overline{\Gamma(m_0)} = L^2(N, \tau)$. A graph-controlled closable witness with bounded decoder gives a positive sufficient certificate for closable totalization; contrapositively, the obstructed multiplication face admits no certificate of that Hilbert type on the canonical total witness source. The same algebraic multiplication formula nevertheless survives through the canonical dual operator-valued weight. The paper therefore separates finite validity from total analytic survival, and realization failure from extinction of the underlying operation.

Research note – not peer reviewed. Trust, but verify. Always.

Version 1.0.0 – August 9, 2026

Contents

1 Introduction

An operation may be well behaved on every finite portion of a source and nevertheless fail at the passage to its natural analytic totalization. The local formulas may remain compatible, every finite restriction may be bounded, and the underlying source inclusions may remain exact. What can fail is the existence of a single analytic realization carrying the same continuation after the finite pieces have been assembled.

The problem of this paper is therefore not whether every finite context has a valid realization. It is whether finite validity contains enough information to certify survival at the limit.

What finite information certifies analytic survival under totalization?

The motivating example comes from relative multiplication in the Jones basic construction. Let

$$A \subset (N, \tau)$$

be a finite von Neumann inclusion with trace-preserving conditional expectation E_A, and let $N_1 = \langle N, e_A \rangle$. On the algebraic basic-construction core one has

$$m_0(xe_A y) = xy.$$

If the inclusion admits an infinite E_A-orthonormal unitary family, there are mutually orthogonal unit vectors p_i in the source with $m_0(p_i) = 1$. Hence on the finite span of $|F|$ such vectors,

$$\|m_F\| = \sqrt{|F|}.$$

Yet the averages $\xi_n = n^{-1} \sum_{i=1}^{n} p_i$ satisfy $\xi_n \to 0$ while $m_0(\xi_n) = 1$. The total multiplication face is therefore not merely unbounded; it is nonclosable. The left N-module structure amplifies this single vertical defect to the entire target:

$$\operatorname{mul} \overline{\Gamma(m_0)} = L^2(N, \tau).$$

This example also shows that a one-dimensional notion of context is too coarse. The continuation being tested does not change: multiplication is fixed throughout. What changes is the amount of source structure on which that continuation must remain analytically meaningful. In other reconstruction problems the source may be fixed while the continuation burden grows. These are different operations, and they must be recorded separately.

We therefore index a formal context by a pair

$$(F, P),$$

where F records source exposure and P records the continuations required to survive. Before any analytic structure is imposed, the required continuation data determine a canonical minimum source quotient

$$Q_{F,P} = D_F \Big/ \bigcap_{\alpha \in P} \ker(r_\alpha J_{\alpha,F}).$$

The source variable is covariant, the continuation variable contravariant, and the two passages commute strictly. Thus

$$Q : \mathsf{Win}(D) \times \mathsf{Ctx}_{\mathrm{all}}(\mathbf{J})^{\mathrm{op}} \longrightarrow \mathsf{Vect}$$

is a strict bifunctor. This is the algebraic control case for the paper. Any later missing context must therefore be analytic rather than a defect of the core reconstruction maps.

Fix a class $\mathcal{K}$ of permitted analytic source carriers and a class $\mathfrak{M}$ of permitted continuation morphisms. A formal context (F, P) is analytically admissible when the source window admits a carrier in $\mathcal{K}$ through which every continuation in P descends as a morphism in $\mathfrak{M}$. Under natural source-restriction stability, the admissible contexts form a downward-closed region $\mathcal{A}_{\mathcal{K}}^{\mathfrak{M}}$ inside the formal context space.

A new phenomenon can then occur. The stronger continuation demand may be admissible on a smaller source window, and the larger source window may be admissible under the weaker continuation demand, while no allowed realization carries both simultaneously. We call this a rectangular admissibility obstruction.

The first principal result shows that relative multiplication realizes this missing-corner geometry naturally. Under infinite E_A-orthogonal multiplicity, every finite witness window supports the multiplication face boundedly, with exact norm $\|m_F\| = \sqrt{|F|}$, while the total Hilbert realization is maximally nonclosable. Thus finite strong contexts and the weak total source context exist, but the strong total context does not.

The exact finite norm profile does not itself identify the obstruction. Divergent finite norms can also occur when the limit is a legitimate closed unbounded operator. What fails here is the graph: source vectors can converge to zero while their continuation values retain a nonzero limit. Any positive criterion must therefore control the graph itself.

This forces a second object. A witness is additional source-side data from which the continuation can be recovered and whose own analytic behavior remains controlled by the continuation graph. We prove that if

$$T_0 = SW_0$$

with S bounded, W_0 closable, and

$$\|W_0 x\| \le C\left(\|x\|^2 + \|T_0 x\|^2\right)^{1/2},$$

then T_0 is closable. More precisely,

$$\Gamma(\overline{T}_0) \subseteq \Gamma(S\overline{W}_0),$$

and under reverse graph control the two closed domains coincide and $\overline{T}_0 = S\overline{W}_0$. The certificate is therefore not finite validity itself. It is additional structure whose own totalization remains controlled.

Applied contrapositively, the witness theorem shows that the infinite-multiplicity multiplication face cannot admit a witness certificate of this Hilbert type on the canonical total witness source.

Once maximal Hilbert failure and the corresponding witness exclusion are known, another Hilbert estimate is no longer the natural next question. One must instead ask whether multiplication itself has disappeared or only this realization of it. The canonical dual operator-valued weight answers that question: on its finite ideal, its linear realization satisfies

$$\widehat{E}_A(xe_A y) = xy,$$

the same algebraic formula as m_0. Hence the natural Hilbert realization may fail maximally while the underlying algebraic operation remains canonically represented in a different analytic framework.

Failure of an analytic realization is not yet extinction of the operation.

The individual ingredients used below belong to established theories: quotient reconstruction, the Jones basic construction, closed linear relations, operator-valued weights, inverse systems, and context-indexed descent. The contribution here lies in a specific finite-to-total interaction among them. We separate source exposure from continuation demand, identify the resulting analytic validity region, realize a missing directed context through relative multiplication at infinite orthogonal multiplicity, and give a graph-controlled witness criterion that restores closable totalization. The dual operator-valued weight then shows that failure at the Hilbert boundary need not extinguish the underlying algebraic operation.

The relative multiplication phenomenon arose in our preceding operator-algebraic analysis of the basic construction. The present paper isolates the finite-to-total mechanism itself and reproves the necessary multiplication obstruction in a self-contained form. The earlier work supplies the motivating operator-algebraic environment; the reconstruction bifunctor, admissible-region formulation, and witness totalization theorem developed here do not logically depend on that earlier proof.

Context-indexed presheaf and global-section methods are well established in sheaf-theoretic contextuality [1]. The principal obstruction considered here is different in emphasis: it may occur with a single continuation fixed while the exposed source grows. When the source window is held fixed, however, the continuation variable admits a derived descent theory. Under an additional contextual reduction calculus, effective descent decomposes into admissible amalgamation and recognition.

A second boundary phenomenon occurs even when no continuation becomes nonclosable. Every finite continuation context may admit an admissible carrier and the finite carriers may form a coherent inverse system, while no admissible total carrier exists. The coordinate model of section 9 separates this carrier-level failure from the operation-level obstruction furnished by relative multiplication.

The present work does not claim a general classification of finite-to-total reconstruction obstructions. The witness theorem is a sufficient criterion, not a characterization. The fixed-source descent theorem is conditional on the contextual reduction axioms of section 8. Nor do we claim that every failed Hilbert realization survives in an operator-valued framework. Finally, nothing in the present results forces higher categorical coherence: the core reconstruction bifunctor is strict, and the main failure lies in analytic realizability.

The main results may be summarized as follows.

(1) The minimal core carriers $Q_{F,P}$ form a strict bifunctor in source exposure and continuation demand.

(2) Under source-restriction stability, analytically realizable contexts form a downward-closed admissible region.

(3) Infinite E_A-orthogonal multiplicity gives a natural finite-to-total analytic boundary for relative multiplication, with $\|m_F\| = \sqrt{|F|}$ on finite witness windows and $\operatorname{mul}\overline{\Gamma(m_0)} = L^2(N,\tau)$ at totalization.

(4) A closable witness uniformly controlled by the continuation graph and admitting a bounded decoder forces closable totalization.

(5) The same core multiplication formula survives through the canonical dual operator-valued weight despite maximal failure of the natural Hilbert realization.

Section 2 constructs the core reconstruction bifunctor. Section 3 introduces analytic realizations and rectangular obstruction. Sections 4 and 5 develop the basic-construction example and prove the maximal finite-to-total boundary theorem. Section 6 proves the witness totalization criterion and its exclusion consequence for relative multiplication. Section 7 compares the failed Hilbert realization with the canonical dual operator-valued weight. Finally, sections 8 and 9 record the fixed-source descent shadow and the distinct failure of admissible total carrier representation.

2 Source Windows and Core Reconstruction

An analytic obstruction is meaningful only after the underlying algebraic reconstruction problem has been separated from it. We therefore begin before topology or completion enters. Two pieces of data will vary independently. The first records how much of the source has been exposed. The second records which continuations must remain distinguishable. The purpose of this section is to determine the minimal algebraic object carrying those requirements and to show that the two variations are exactly compatible.

Definition 2.1 (Source-window system)**.** A *source-window system* consists of a directed category $\mathsf{Win}(D)$, linear spaces D_F indexed by $F \in \mathsf{Win}(D)$, and injective linear maps

$$\iota_{F,G} : D_F \longrightarrow D_G, \qquad F \le G,$$

such that $\iota_{F,F} = 1_{D_F}$ and $\iota_{G,H}\iota_{F,G} = \iota_{F,H}$.

The windows need not exhaust all subspaces of an ambient source. They specify the enlargements admitted by the reconstruction problem. In the basic-construction application below they will be finite witness spans.

Let Λ index a family of continuations. For each $\alpha \in \Lambda$, fix linear spaces E_α and Y_α, a target realization $r_\alpha : E_\alpha \to Y_\alpha$, and maps $J_{\alpha,F} : D_F \to E_\alpha$. We assume compatibility with source exposure:

$$J_{\alpha,G}\iota_{F,G} = J_{\alpha,F}, \qquad F \leq G. \tag{1}$$

Condition (1) says that enlarging the source adds new source data; it does not alter the continuation already defined on the smaller window.

Definition 2.2 (Continuation context)**.** A *continuation context* is a subset $P \subseteq \Lambda$. We write $\mathsf{Ctx}_{\mathrm{all}}(\mathbf{J})$ for the poset of continuation contexts ordered by inclusion.

The algebraic theory does not require P to be finite. Finiteness becomes relevant only when analytic validity is tested under totalization.

Fix (F, P). Two source elements should become indistinguishable precisely when every continuation presently required by P fails to distinguish them. The missing datum is therefore not another carrier chosen externally, but the quotient determined by the joint continuation kernel.

Definition 2.3 (Context kernel and core reconstruction carrier)**.** For a formal context (F, P), set

$$K_{F,P} := \bigcap_{\alpha \in P} \ker(r_\alpha J_{\alpha,F}), \tag{2}$$

with $K_{F,\varnothing} = D_F$, and define

$$Q_{F,P} := D_F / K_{F,P}. \tag{3}$$

Let $\pi_{F,P} : D_F \to Q_{F,P}$ denote the quotient map.

The object $Q_{F,P}$ is the minimal algebraic carrier of the distinctions forced by P. It is not, in general, the analytic source carrier used later. In particular, $Q_{F,\varnothing} = 0$, although an analytic realization of the source window D_F may be nonzero.

Proposition 2.4 (Universal core reconstruction property)**.** *Let $q : D_F \to X$ be linear. The following are equivalent:*

(i) *every continuation in P descends through q;*

(ii) $\ker q \subseteq K_{F,P}$*;*

(iii) *there exists a unique linear map $\Phi_q : q(D_F) \to Q_{F,P}$ satisfying $\pi_{F,P} = \Phi_q q$.*

Consequently, $Q_{F,P}$ is the coarsest algebraic realization retaining every distinction required by P.

Proof. Suppose first that each $\alpha \in P$ descends through q. If $q(d) = 0$, then $r_\alpha J_{\alpha,F}(d) = 0$ for every $\alpha \in P$, hence $d \in K_{F,P}$. Conversely, if $\ker q \subseteq K_{F,P}$, define $T_\alpha(q(d)) := r_\alpha J_{\alpha,F}(d)$. The kernel inclusion makes T_α well defined. Under the same hypothesis, $\Phi_q(q(d)) := [d]_{F,P}$ is well defined and is plainly the unique map satisfying $\pi_{F,P} = \Phi_q q$. □

If $P \subseteq Q$, then $K_{F,Q} \subseteq K_{F,P}$. Hence there is a canonical quotient map

$$\rho^F_{Q,P} : Q_{F,Q} \longrightarrow Q_{F,P}, \qquad [d]_{F,Q} \longmapsto [d]_{F,P}.$$

Lemma 2.5. *For $P \subseteq Q \subseteq R$,*

$$\rho^F_{P,P} = 1, \qquad \rho^F_{Q,P}\rho^F_{R,Q} = \rho^F_{R,P}.$$

Proof. Both identities follow directly from representatives. □

Source enlargement requires one additional check: a distinction already visible on a smaller window must not disappear merely because more of the source has been exposed. This is exactly what (1) guarantees.

Proposition 2.6 (Source exposure)**.** *For $F \le G$ and every continuation context P, there is a unique injective map*

$$\iota^P_{F,G} : Q_{F,P} \longrightarrow Q_{G,P}$$

satisfying $\iota^P_{F,G}\pi_{F,P} = \pi_{G,P}\iota_{F,G}$. Explicitly,

$$\iota^P_{F,G}([d]_{F,P}) = [\iota_{F,G}(d)]_{G,P}.$$

Proof. If $d \in K_{F,P}$, then for every $\alpha \in P$,

$$r_\alpha J_{\alpha,G}\iota_{F,G}(d) = r_\alpha J_{\alpha,F}(d) = 0,$$

so the map is well defined. If $\iota^P_{F,G}([d]_{F,P}) = 0$, then $\iota_{F,G}(d) \in K_{G,P}$, and the same compatibility identity gives $d \in K_{F,P}$. Thus the induced map is injective. □

Theorem 2.7 (Core Reconstruction Bifunctor)**.** *Under* (1)*, the assignment*

$$(F,P) \longmapsto Q_{F,P} = D_F \Big/ \bigcap_{\alpha\in P} \ker(r_\alpha J_{\alpha,F})$$

defines a strict bifunctor

$$Q : \mathsf{Win}(D) \times \mathsf{Ctx}_{\mathrm{all}}(\mathbf{J})^{\mathrm{op}} \longrightarrow \mathsf{Vect}.$$

For every $F \le G$ and $P \subseteq Q$,

$$\rho^G_{Q,P}\,\iota^Q_{F,G} = \iota^P_{F,G}\,\rho^F_{Q,P}. \tag{4}$$

Moreover, every source-exposure map $\iota^P_{F,G}$ is injective.

Proof. Functoriality in each variable follows from theorems 2.5 and 2.6. For interchange, take $[d]_{F,Q} \in Q_{F,Q}$. Both composites in (4) send it to $[\iota_{F,G}(d)]_{G,P}$. Injectivity was proved in theorem 2.6. □

$$\begin{tikzcd}
Q_{F,Q} \arrow[r, "\iota^Q_{F,G}"] \arrow[d, "\rho^F_{Q,P}"'] & Q_{G,Q} \arrow[d, "\rho^G_{Q,P}"] \\
Q_{F,P} \arrow[r, "\iota^P_{F,G}"'] & Q_{G,P}
\end{tikzcd}$$

Thus the two contextual directions are already exactly compatible at the algebraic level. A later failure of analytic realization cannot be attributed to noncommutativity of these core passages: the formal upper-right corner exists, and the two routes through the square agree.

The quotient $Q_{F,P}$ answers the algebraic question: it records the minimum source information required by the continuation context. It does not answer whether that information admits a realization in a prescribed analytic class. The next section fixes such a class and asks which formal contexts survive completion, domain restrictions, and the chosen notion of analytic morphism.

3 Analytic Realization and the Admissible Region

The core reconstruction problem is exact. For every formal context (F, P) the quotient $Q_{F,P}$ exists, source exposure is covariant, continuation reduction is contravariant, and the resulting squares commute strictly. The obstruction begins only after an analytic realization class is specified. Completion introduces topology, domain conditions, and a notion of admissible morphism that are invisible to the core quotient. A formal context may therefore remain algebraically well defined while admitting no realization in the prescribed analytic class.

We fix these analytic requirements before testing any continuation. Otherwise admissibility would merely rename success.

Let $\mathcal{K}$ be a class of permitted analytic source carriers and let $\mathfrak{M}$ be a class of permitted analytic continuation morphisms.

Definition 3.1 (Analytic realization)**.** A $(\mathcal{K}, \mathfrak{M})$-*realization* of a formal context (F, P) consists of a linear map $q : D_F \to X$ with dense image and $X \in \mathcal{K}$, together with, for every $\alpha \in P$, an operator

$$T_\alpha : D(T_\alpha) \subseteq X \longrightarrow Y_\alpha$$

belonging to $\mathfrak{M}$ such that $q(D_F) \subseteq D(T_\alpha)$ and

$$T_\alpha q = r_\alpha J_{\alpha,F}. \tag{5}$$

We write $(F, P) \in \mathcal{A}^{\mathfrak{M}}_{\mathcal{K}}$ when such a realization exists.

The analytic carrier realizes the source window D_F itself. It need not coincide with the minimal algebraic carrier $Q_{F,P}$. No minimality condition is imposed on X; in particular, if a carrier realizes P, it realizes any smaller context by forgetting unused continuation certifications.

Continuation weakening is automatic at this level. Source restriction is not: after shrinking the source window, the closed subcarrier generated by the restricted source must remain inside the permitted analytic class.

Definition 3.2 (Source-restriction stability)**.** The pair $(\mathcal{K}, \mathfrak{M})$ is *source-restriction stable* if whenever $q : D_F \to X$ realizes (F, P) and $F' \leq F$, the closed subspace

$$X' := \overline{q(D_{F'})}^{X}$$

belongs to $\mathcal{K}$, and for every $\alpha \in P$ the restriction of T_α to the dense domain $q(D_{F'}) \subseteq X'$ belongs to $\mathfrak{M}$.

Proposition 3.3. *Bounded Hilbert-space operators are stable under restriction to closed source-generated subspaces.*

Proof. If $T : X \to Y$ is bounded and $X' \subseteq X$ is closed, then $T|_{X'}$ is bounded and $\|T|_{X'}\| \leq \|T\|$. □

Proposition 3.4. *Let $T : D(T) \subseteq X \to Y$ be closable. If $D' \subseteq D(T)$ and $X' = \overline{D'}^{X}$, then $T|_{D'} : D' \subseteq X' \to Y$ is closable.*

Proof. Suppose $x_n \in D'$ satisfies $x_n \to 0$ in X' and $Tx_n \to y$ in Y. The same convergence holds in X, so closability of T gives $y = 0$. □

Theorem 3.5 (Downward Closure of Analytic Admissibility)**.** *Assume $(\mathcal{K}, \mathfrak{M})$ is source-restriction stable. If*

$$(F, P) \in \mathcal{A}^{\mathfrak{M}}_{\mathcal{K}}, \qquad F' \leq F, \qquad P' \subseteq P,$$

then $(F', P') \in \mathcal{A}^{\mathfrak{M}}_{\mathcal{K}}$.

Proof. Let $q : D_F \to X$ be an admissible realization of (F, P) and set $X' := \overline{q(D_{F'})}^X$. By source-restriction stability, $X' \in \mathcal{K}$, and each continuation in P restricts to an allowed morphism on the dense domain $q(D_{F'}) \subseteq X'$. Hence (F', P) is admissible. Since $P' \subseteq P$, forgetting the continuations indexed by $P \setminus P'$ yields an admissible realization of (F', P'). □

Thus $\mathcal{A}_{\mathcal{K}}^{\mathfrak{M}}$ is a downward-closed subset of formal context space.

In many applications the analytic carrier is not arbitrary. A trace, module structure, or other intrinsic datum selects a distinguished completion. We write $\mathcal{K}_{\mathrm{can}}$ for such a canonical realization package when needed. Failure of one chosen carrier does not imply nonrealizability in every conceivable analytic model, but failure in a canonical package is already a genuine obstruction for that analytic problem.

Let $F \le G$ and $P \subseteq Q$. Suppose the stronger continuation context Q is realizable on the smaller source window while the larger source window is realizable under the weaker continuation demand. By theorem 3.5, (F, P) is then automatically admissible. Nothing, however, forces the remaining context (G, Q) to exist analytically.

Definition 3.6 (Rectangular obstruction)**.** A quadruple $(F, G; P, Q)$, with $F \le G$ and $P \subseteq Q$, is a *rectangular admissibility obstruction* relative to $(\mathcal{K}, \mathfrak{M})$ if

$$(F, Q), (G, P) \in \mathcal{A}_{\mathcal{K}}^{\mathfrak{M}}$$

but

$$(G, Q) \notin \mathcal{A}_{\mathcal{K}}^{\mathfrak{M}}.$$

Proposition 3.7 (Rectangular Obstruction Principle)**.** *Let $(F, G; P, Q)$ be a rectangular admissibility obstruction. Then the associated core reconstruction square exists and commutes strictly. Hence the obstruction is a failure of simultaneous analytic realizability, not a failure of algebraic coherence.*

Proof. The core square is provided by theorem 2.7, and its commutativity is (4). The missing analytic corner is precisely $(G, Q) \notin \mathcal{A}_{\mathcal{K}}^{\mathfrak{M}}$. □

Definition 3.8 (Directed source rectangular obstruction)**.** Let $F_1 \le F_2 \le \cdots$ be a directed source exhaustion with specified total source context F_∞, and let $P \subseteq Q$. A *directed source rectangular obstruction* occurs if

$$(F_n, Q) \in \mathcal{A}_{\mathcal{K}}^{\mathfrak{M}} \quad \text{for every } n,$$

and

$$(F_\infty, P) \in \mathcal{A}_{\mathcal{K}}^{\mathfrak{M}},$$

but

$$(F_\infty, Q) \notin \mathcal{A}_{\mathcal{K}}^{\mathfrak{M}}.$$

The notation F_∞ always denotes a specified total source object together with its prescribed analytic completion; it is not assumed to be an object of the original finite-window category.

The formalism is now sufficient to state the obstruction, but it does not yet show that the missing corner occurs naturally. The next two sections return to the Jones basic construction. An infinite E_A-orthogonal unitary family will produce finite source windows on which relative multiplication is bounded. Their total Hilbert realization will then supply a directed rectangular obstruction in the strongest possible graph-theoretic form.

4 Relative Multiplication and Orthogonal Witness Windows

The rectangular obstruction of section 3 is only a formal possibility until one produces an analytic problem in which the strong finite contexts exist while the strong total context does not. The Jones basic construction provides such a problem. The mechanism is small: infinite relative orthogonality produces a family of orthogonal source vectors that the multiplication face identifies. Finite spans are boundedly harmless. Their directed average is not.

Let $A \subset (N, \tau)$ be a unital inclusion of finite von Neumann algebras, and let $E_A : N \to A$ be the trace-preserving conditional expectation. Write $N_1 = \langle N, e_A \rangle$ for the Jones basic construction and Tr_1 for its canonical semifinite trace. We use the standard basic-construction identities [4, 5]

$$e_A x e_A = E_A(x) e_A, \tag{6}$$

and

$$\mathrm{Tr}_1(x e_A y) = \tau(yx) \tag{7}$$

on the canonical finite ideal. Set $C_A := \operatorname{span} N e_A N$. In the canonical semifinite realization,

$$\overline{C_A}^{L^2(N_1,\mathrm{Tr}_1)} = L^2(N_1, \mathrm{Tr}_1).$$

On C_A define the relative multiplication face

$$m_0 : C_A \to N, \qquad m_0(x e_A y) = xy. \tag{8}$$

When analytic questions are considered, the target is $L^2(N, \tau)$.

Assumption 4.1 (Infinite E_A-orthogonal multiplicity)**.** There exists a sequence $\{u_i\}_{i \geq 1} \subset \mathcal{U}(N)$ such that

$$E_A(u_i^* u_j) = \delta_{ij} 1. \tag{9}$$

For each i, set $p_i := u_i e_A u_i^*$.

Lemma 4.2 (Orthogonal witness projections)**.** *Under theorem 4.1,*

$$p_i p_j = \delta_{ij} p_i.$$

Moreover,

$$\mathrm{Tr}_1(p_i) = 1, \qquad \|p_i\|_{2,\mathrm{Tr}_1} = 1.$$

Proof. Using (6),

$$\begin{aligned} p_i p_j &= u_i e_A u_i^* u_j e_A u_j^* \\ &= u_i E_A(u_i^* u_j) e_A u_j^* \\ &= \delta_{ij} p_i. \end{aligned}$$

Thus the p_i are mutually orthogonal projections. By (7), $\mathrm{Tr}_1(p_i) = \tau(u_i^* u_i) = 1$, hence $\|p_i\|_{2,\mathrm{Tr}_1} = 1$. □

Lemma 4.3 (Equal multiplication image)**.** *For every i, $m_0(p_i) = 1$.*

Proof. By (8), $m_0(u_i e_A u_i^*) = u_i u_i^* = 1$. □

Thus $\{p_i\}$ is an infinite orthonormal family in the source whose images under m_0 are all identical.

For $F \Subset \mathbb{N}$, define

$$D_F := \operatorname{span}\{p_i : i \in F\}.$$

If $F \subseteq G$, the source-exposure map is the inclusion $D_F \hookrightarrow D_G$. Set

$$D_{\mathrm{wit}} := \bigcup_{F \Subset \mathbb{N}} D_F = \operatorname{span}\{p_i : i \geq 1\},$$

and

$$X_{\mathrm{wit}} := \overline{D_{\mathrm{wit}}}^{L^2(N_1, \mathrm{Tr}_1)}.$$

Since the p_i are orthonormal, $X_{\mathrm{wit}} \cong \ell^2(\mathbb{N})$.

For finite F, write $m_F := m_0|_{D_F}$.

Proposition 4.4 (Exact finite-window norm). *For every nonempty finite $F \Subset \mathbb{N}$,*

$$\|m_F\| = \sqrt{|F|}. \tag{10}$$

Proof. Let $\xi = \sum_{i \in F} a_i p_i$. By orthonormality,

$$\|\xi\|_{2,\mathrm{Tr}_1}^2 = \sum_{i \in F} |a_i|^2.$$

By theorem 4.3,

$$m_F(\xi) = \left(\sum_{i \in F} a_i \right) 1.$$

Hence

$$\|m_F(\xi)\|_{2,\tau} \leq \sqrt{|F|} \left(\sum_{i \in F} |a_i|^2 \right)^{1/2}.$$

Thus $\|m_F\| \leq \sqrt{|F|}$, and equality is attained for $a_i = |F|^{-1/2}$. □

The norm growth rules out a uniformly bounded total extension, but it does not yet distinguish a legitimate closed unbounded limit from a nonclosable one. The obstruction must therefore be read from the graph.

Lemma 4.5 (Orthogonal equal-image mechanism). *Let $T_0 : D(T_0) \subset X \to Y$ be linear, where X and Y are Hilbert spaces. Suppose there is an orthonormal family $\{v_i\}_{i \geq 1} \subset D(T_0)$ and a nonzero $y \in Y$ such that $T_0 v_i = y$ for every i. For finite nonempty F, let $X_F := \operatorname{span}\{v_i : i \in F\}$ and $T_F := T_0|_{X_F}$. Then*

$$\|T_F\| = \sqrt{|F|}\, \|y\|.$$

Moreover, $x_F := |F|^{-1} \sum_{i \in F} v_i$ satisfies $\|x_F\| = |F|^{-1/2}$ and $T_0 x_F = y$. Consequently T_0 is not closable on the Hilbert closure of the algebraic span of the v_i.

Proof. For $x = \sum_{i \in F} a_i v_i$, orthonormality gives $\|x\|^2 = \sum |a_i|^2$, while $T_F x = (\sum a_i) y$. Cauchy–Schwarz gives the norm bound, with equality for equal coefficients. The averaging statement follows immediately, and the closability criterion then fails. □

The relative multiplication system is the specialization $v_i = p_i$, $y = 1$. The basic-construction structure will additionally amplify the single vertical direction 1 to the whole target.

We have now constructed the finite source windows and identified their exact analytic profile. Every finite multiplication face is bounded, but the same orthogonal equal-image geometry already contains a sequence converging to zero whose multiplication value does not. The next section turns that sequence into the total obstruction and uses the left N-module structure to show that the resulting graph defect is maximal.

5 The Finite-to-Total Analytic Boundary

The finite theory has no singularity: every m_F is bounded. The obstruction appears only after the witness windows are totalized. For the countable witness exhaustion used below, the relevant failure is detected by the vertical part of the graph closure.

Definition 5.1 (Directed graph defect)**.** Let $T_0 : \mathcal{D} \subset X \to Y$ be densely defined on the algebraic union of a countable increasing family of source windows. Define

$$\mathrm{Def}_\infty(T_\bullet) := \{y \in Y : \exists x_n \in \mathcal{D},\ x_n \to 0,\ T_0 x_n \to y\}.$$

Equivalently,

$$\mathrm{Def}_\infty(T_\bullet) = \mathrm{mul}\,\overline{\Gamma(T_0)}.$$

Hence T_0 is closable if and only if $\mathrm{Def}_\infty(T_\bullet) = \{0\}$.

For $F_n = \{1, \dots, n\}$, set

$$\xi_n := \frac{1}{n}\sum_{i=1}^{n} p_i. \tag{11}$$

Then $\|\xi_n\|_{2,\mathrm{Tr}_1} = n^{-1/2} \to 0$, while

$$m_0(\xi_n) = 1. \tag{12}$$

Proposition 5.2 (Witness-space nonclosability)**.** *The operator*

$$m_{\mathrm{wit},0} : D_{\mathrm{wit}} \subset X_{\mathrm{wit}} \longrightarrow L^2(N, \tau)$$

is not closable.

Proof. Equation (11) gives $\xi_n \to 0$, while (12) gives $m_{\mathrm{wit},0}\xi_n \to 1$. □

The witness-space defect initially supplies only the vertical direction 1. The basic construction carries more structure: m_0 is left N-linear on C_A.

Theorem 5.3 (Maximal graph defect)**.** *Under theorem 4.1,*

$$\mathrm{mul}\,\overline{\Gamma(m_0)} = L^2(N, \tau). \tag{13}$$

Proof. Let $z \in N$. Left multiplication by z is bounded on $L^2(N_1, \mathrm{Tr}_1)$, so $z\xi_n \to 0$. Since $z\xi_n \in C_A$ and m_0 is left N-linear,

$$m_0(z\xi_n) = z\, m_0(\xi_n) = z.$$

Hence $N \subseteq \mathrm{mul}\,\overline{\Gamma(m_0)}$. The multivalued part of a closed linear relation is closed in the target, while N is dense in $L^2(N, \tau)$. Therefore (13) holds. □

Corollary 5.4 (Full graph relation)**.** *Under the hypotheses of theorem 5.3,*

$$\overline{\Gamma(m_0)} = L^2(N_1, \mathrm{Tr}_1) \oplus L^2(N, \tau). \tag{14}$$

Proof. Set $\mathcal{G} := \overline{\Gamma(m_0)}$. By theorem 5.3, $\{0\} \oplus L^2(N, \tau) \subseteq \mathcal{G}$. Let $x \in L^2(N_1, \mathrm{Tr}_1)$. Choose $x_n \in C_A$ with $x_n \to x$. Since $(x_n, m_0 x_n) \in \mathcal{G}$ and $(0, -m_0 x_n) \in \mathcal{G}$, we have $(x_n, 0) \in \mathcal{G}$. Closedness gives $(x, 0) \in \mathcal{G}$. Adding an arbitrary vertical element $(0, y) \in \mathcal{G}$ yields $(x, y) \in \mathcal{G}$ for every $y \in L^2(N, \tau)$. □

We write F_∞ for the total source context determined by the algebraic union D_{wit} together with its canonical Hilbert completion X_{wit}. Let $P = \varnothing$ and $Q = \{\mu\}$, where μ denotes the requirement that relative multiplication survive as a closable Hilbert-space continuation.

Theorem 5.5 (Infinite Orthogonal-Multiplicity Analytic Boundary)**.** *Let $A \subset (N, \tau)$ satisfy theorem 4.1, and let the carrier package select the Hilbert realizations inherited from the basic construction. Then for every finite nonempty $F \Subset \mathbb{N}$, the strong finite context (F, Q) is admissible, with $\|m_F\| = \sqrt{|F|}$. The weak total source context (F_∞, P) is admissible, whereas the strong total context (F_∞, Q) is not. More strongly,*

$$\operatorname{mul}\overline{\Gamma(m_0)} = L^2(N, \tau).$$

Proof. Finite strong admissibility follows from theorem 4.4. The weak total source carrier X_{wit} exists by construction and imposes no multiplication requirement. The strong total context would require the induced multiplication face to be closable. This fails by theorem 5.2; the stronger maximal statement is theorem 5.3. □

$$\begin{tikzcd}
(F, P) \arrow[r] \arrow[d] & (F_\infty, P) \arrow[d] \\
(F, Q) \arrow[r] & \boxed{(F_\infty, Q)}
\end{tikzcd}$$

The first three contexts are admissible; the framed context is not. Thus theorem 5.5 realizes a directed rectangular admissibility obstruction in the sense of theorem 3.8.

Corollary 5.6 (Failure of source compactness)**.** *The closable-Hilbert admissible region for relative multiplication is not closed under directed source totalization: every finite strong context is admissible, while the strong total context is not.*

The finite norm profile is therefore not the obstruction by itself. What fails is the graph: a source sequence can disappear while its continuation value survives, and the module structure amplifies that single defect to the whole target. Any positive finite-to-total criterion must prevent precisely this vertical collapse. The next section asks what additional source-side datum can do so without requiring the continuation itself to remain uniformly bounded.

6 Uniform Witness Totalization

The obstruction in section 5 is graph-theoretic. Finite boundedness does not prevent a sequence from converging to zero in the source while retaining a nonzero continuation value at the limit. A positive finite-to-total criterion must therefore control more than the continuation norms. It must carry enough source-side information to prevent vertical graph collapse, and that additional information must itself survive totalization. We call such data a witness.

Let X, Y, Z be Hilbert spaces and let $\mathcal{D} \subset X$ be dense. Let $T_0 : \mathcal{D} \to Y$ be the continuation whose totalization is to be certified, and let $W_0 : \mathcal{D} \to Z$ be a second linear map. The witness must retain enough information to recover the continuation. We therefore assume there is a bounded map $S : Z \to Y$ such that

$$T_0 = SW_0. \tag{15}$$

Factorization alone is not sufficient. A sequence $W_0 x_n$ may wander in directions compressed by S while $SW_0 x_n$ converges. The witness must therefore be controlled by the graph of T_0.

Definition 6.1 (Graph-controlled witness)**.** The map W_0 is a T_0*-graph-controlled witness* if there exists $C < \infty$ such that

$$\|W_0 x\|_Z \le C\big(\|x\|_X^2 + \|T_0 x\|_Y^2\big)^{1/2} \qquad (x \in \mathcal{D}). \tag{16}$$

Equivalently, W_0 is bounded from $\mathcal{D}$ equipped with the graph norm $\|x\|_{T_0} := (\|x\|_X^2 + \|T_0 x\|_Y^2)^{1/2}$ into Z.

Definition 6.2 (Closably totalizable witness)**.** A witness $W_0 : \mathcal{D} \subset X \to Z$ is *closably totalizable* if W_0 is closable.

Theorem 6.3 (Uniform Witness Totalization)**.** *Let $T_0 : \mathcal{D} \subset X \to Y$ and $W_0 : \mathcal{D} \subset X \to Z$ be linear maps on a common dense domain. Assume:*

(i) *W_0 is closable;*

(ii) *W_0 is T_0-graph controlled;*

(iii) *there exists a bounded map $S : Z \to Y$ satisfying $T_0 = SW_0$.*

Then T_0 is closable. Moreover, if $x_n \in \mathcal{D}$, $x_n \to x$, and $T_0 x_n \to y$, then $x \in D(\overline{W}_0)$ and

$$y = S\overline{W}_0 x. \tag{17}$$

Consequently,

$$\Gamma(\overline{T}_0) \subseteq \Gamma(S\overline{W}_0). \tag{18}$$

Proof. Suppose first that $x_n \to 0$ and $T_0 x_n \to y$. Then (x_n) is Cauchy in the T_0-graph norm. By (16),

$$\|W_0 x_n - W_0 x_m\|_Z \leq C\|x_n - x_m\|_{T_0},$$

so $(W_0 x_n)$ is Cauchy in Z. Since Z is complete, $W_0 x_n \to z$ for some $z \in Z$. Closability of W_0 and $x_n \to 0$ imply $z = 0$. Hence $T_0 x_n = SW_0 x_n \to 0$, so $y = 0$. Thus T_0 is closable.

Now let $x_n \to x$ and $T_0 x_n \to y$. The same graph-Cauchy argument gives $W_0 x_n \to z$ for some $z \in Z$. Since W_0 is closable, $x \in D(\overline{W}_0)$ and $z = \overline{W}_0 x$. Finally,

$$y = \lim_n T_0 x_n = \lim_n SW_0 x_n = S\overline{W}_0 x.$$

This proves (18). □

Remark 6.4. Theorem 6.3 gives a sufficient certificate for closability, not a characterization. In particular, a closable operator need not admit witness data satisfying the hypotheses above.

Theorem 6.5 (Graph-Equivalent Witness Realization)**.** *Assume the hypotheses of theorem 6.3. Suppose in addition there exists $C' < \infty$ such that*

$$\|T_0 x\|_Y \leq C'\left(\|x\|_X^2 + \|W_0 x\|_Z^2\right)^{1/2} \qquad (x \in \mathcal{D}). \tag{19}$$

Then the graph norms of T_0 and W_0 are equivalent on $\mathcal{D}$, and

$$D(\overline{T}_0) = D(\overline{W}_0). \tag{20}$$

Moreover,

$$\overline{T}_0 = S\overline{W}_0. \tag{21}$$

Proof. The estimates (16) and (19) show that the graph norms $\|\cdot\|_{T_0}$ and $\|\cdot\|_{W_0}$ are equivalent on $\mathcal{D}$. Both graph norms dominate the ambient X-norm. Hence a sequence that is Cauchy in either graph norm has the same X-limit under the canonical embedding into X. The corresponding graph completions therefore determine the same domain in X, which gives (20). If x belongs to this common domain and $x_n \in \mathcal{D}$ converges to x in the W_0-graph norm, then $W_0 x_n \to \overline{W}_0 x$. Graph-norm equivalence makes (x_n) Cauchy in the T_0-graph norm, and $T_0 x_n = SW_0 x_n \to S\overline{W}_0 x$. Therefore $\overline{T}_0 x = S\overline{W}_0 x$. □

Corollary 6.6 (Uniform contextual witness certificate)**.** *Let $X_1 \subseteq X_2 \subseteq \cdots$ be a countable increasing source system with dense algebraic union $\mathcal{D} \subset X$. Let $T_n : X_n \to Y$ and $W_n : X_n \to Z$ be compatible families inducing maps T_0 and W_0 on $\mathcal{D}$. Assume W_0 is closable, there exists C independent of n such that*

$$\|W_n x\|_Z \leq C\big(\|x\|_X^2 + \|T_n x\|_Y^2\big)^{1/2},$$

and there exists one bounded $S : Z \to Y$ with $T_n = SW_n$ for every n. Then the induced total continuation T_0 is closable.

Proof. The uniform estimate makes W_0 T_0-graph controlled on the algebraic union. Apply theorem 6.3. □

Corollary 6.7 (Certificate-induced source totalization)**.** *Let $P \Subset \Lambda$ be a finite continuation context and let $F_1 \leq F_2 \leq \cdots \to F_\infty$ be a specified source totalization. Suppose the finite contexts (F_n, P) admit compatible analytic realizations on a common total source carrier. If every continuation required by P admits a uniform contextual witness certificate as in theorem 6.6, and the resulting carriers remain in the prescribed analytic class, then the total context (F_∞, P) is admissible in the closable morphism class.*

Corollary 6.8 (No uniform Hilbert witness certificate)**.** *Under the hypotheses of theorem 5.5, the total witness-space multiplication operator*

$$m_{\mathrm{wit},0} : D_{\mathrm{wit}} \subset X_{\mathrm{wit}} \longrightarrow L^2(N, \tau)$$

admits no witness certificate satisfying the hypotheses of theorem 6.3.

Proof. Any such certificate would imply that $m_{\mathrm{wit},0}$ is closable, contradicting theorem 5.2. □

Consequently, any compatible global Hilbert witness construction whose restriction to D_{wit} satisfies the same hypotheses is likewise excluded.

$$\begin{tikzcd} \mathcal{D} \arrow[r, "W_0"] \arrow[rr, bend right, "T_0"'] & Z \arrow[r, "S"] & Y \end{tikzcd} \qquad T_0 = SW_0.$$

The witness theorem remains inside the original Hilbert analytic problem. It repairs totalization, when possible, by retaining additional source-side data whose own graph remains controlled. For relative multiplication this route is unavailable: theorem 6.8 rules out the corresponding certificate class. The remaining question is therefore no longer whether another Hilbert estimate repairs the same graph. It is whether the algebraic operation survives in a different analytic realization.

7 Analytic Realization Separation

The preceding sections leave no ambiguity about the Hilbert-space problem. Under infinite E_A-orthogonal multiplicity, the relative multiplication face is maximally nonclosable, and theorem 6.8 excludes the corresponding class of graph-controlled Hilbert witness repairs.

The remaining question is therefore not whether another Hilbert estimate has been missed. It is whether the algebraic multiplication operation itself has disappeared, or only its natural Hilbert realization.

Recall the algebraic basic-construction core

$$C_A = \operatorname{span} N e_A N$$

and the relative multiplication face

$$m_0 : C_A \to N, \qquad m_0(xe_A y) = xy.$$

Nothing in this algebraic formula failed in theorem 5.3; what failed was its realization as a closable operator

$$C_A \subset L^2(N_1, \mathrm{Tr}_1) \longrightarrow L^2(N, \tau).$$

Let $\widehat{E}_A$ denote the canonical dual operator-valued weight associated with the basic construction; see [2, 3, 5]. We use the standard finite-ideal formula, stated explicitly in the basic-construction literature (for example [6, §1.3.1]), that the linear extension of $\widehat{E}_A$ to its canonical finite ideal contains C_A and satisfies

$$\widehat{E}_A(xe_A y) = xy \qquad (x, y \in N). \tag{22}$$

Thus, on the common algebraic core,

$$\widehat{E}_A(xe_A y) = m_0(xe_A y).$$

Definition 7.1 (Relative analytic survival)**.** Let J_0 be an algebraic operation admitting two analytic realization problems that agree with J_0 on a common algebraic core. We say that J_0 exhibits *relative analytic survival* from the first realization framework to the second if it fails to admit the first realization but admits the second.

Theorem 7.2 (Analytic Realization Separation for Relative Multiplication)**.** *Let $A \subset (N, \tau)$ satisfy theorem 4.1. Then the relative multiplication face $m_0 : C_A \to N$ has the following two analytic statuses. First, as a Hilbert-space operator*

$$C_A \subset L^2(N_1, \mathrm{Tr}_1) \longrightarrow L^2(N, \tau),$$

it is maximally nonclosable:

$$\mathrm{mul}\,\overline{\Gamma(m_0)} = L^2(N, \tau).$$

Second, the linear finite-ideal realization of the canonical dual operator-valued weight satisfies

$$\widehat{E}_A(xe_A y) = xy = m_0(xe_A y)$$

on C_A. Consequently the relative multiplication operation exhibits relative analytic survival from the natural closable-Hilbert realization framework to the dual operator-valued-weight framework.

Proof. The maximal Hilbert-space failure is theorem 5.3. The second statement is the standard finite-ideal formula (22) for the canonical dual operator-valued weight. Both analytic realizations therefore agree with the same algebraic multiplication formula on C_A, while only the operator-valued-weight realization survives. □

Remark 7.3 (Scope of the comparison). Theorem 7.2 does not identify $\widehat{E}_A$ with a closable Hilbert-space operator, does not make it a witness in the sense of theorem 6.3, and does not assert an inclusion between Hilbert-space morphisms and operator-valued weights. The comparison is through the common algebraic multiplication formula on C_A.

Theorems 6.3 and 7.2 describe two different responses to analytic failure. The witness theorem remains inside the original Hilbert problem: additional source-side data certify that the continuation itself has a closable totalization. The present theorem does not repair that failed Hilbert graph. Instead, it shows that the same algebraic operation belongs canonically to a different analytic structure in which it remains meaningful.

internal repair	relative survival
retain controlled source data	change analytic realization
Hilbert continuation survives	Hilbert continuation still fails

Failure of an analytic realization is not yet extinction of the algebraic operation.

The main analytic arc of the paper is now complete. The remaining sections record two shadows of the same reconstruction problem. First, fixing the source produces a continuation-context descent problem under an additional contextual reduction calculus. Second, coherent finite contexts may fail to possess an admissible total carrier even when no continuation develops a graph defect.

8 Fixed-Source Descent

The principal obstruction developed above is source-directed: the continuation demand may remain fixed while the exposed source grows. If the source window is held fixed, however, the continuation variable carries its own local-to-global structure.

At the algebraic level, continuation reduction is already canonical: for $P' \subseteq P$,

$$Q_{F,P} \longrightarrow Q_{F,P'}.$$

Analytically, merely forgetting continuation labels need not remove information retained solely for the stronger context. A genuine fixed-source descent theory therefore requires an additional contextual reduction operation. The present section is conditional on the existence of such a reduction calculus.

Fix a source window F for the remainder of this section and write $D := D_F$. We suppress F from the notation for continuation contexts and their realizations.

Assumption 8.1 (Contextual reduction calculus)**.** For every admissible realization U and continuation context P, there is an admissible P-sufficient realization $\mathcal{R}_P U$ together with a passage $U \to \mathcal{R}_P U$ such that:

(R1) $\mathcal{R}_P U$ is admissible and P-sufficient;

(R2) $\mathcal{R}_P(\mathcal{R}_P U) \simeq \mathcal{R}_P U$;

(R3) if $P' \subseteq P$, then $\mathcal{R}_{P'}\mathcal{R}_P U \simeq \mathcal{R}_{P'} U$;

(R4) if $U \to V$ and $\mathcal{R}_P V \simeq V$, then $\mathcal{R}_P U \simeq V$.

Definition 8.2 (P-tight realization)**.** An admissible P-sufficient realization U is P*-tight* if

$$\mathcal{R}_P U \simeq U.$$

Thus tightness means that contextual reduction has stabilized: there is no further P-relevant information to remove within the chosen reduction calculus.

If U is P-tight and $P' \subseteq P$, define its restriction by

$$U|_{P'} := \mathcal{R}_{P'} U. \tag{23}$$

By theorem 8.1, for $P'' \subseteq P' \subseteq P$ one has $(U|_{P'})|_{P''} \simeq U|_{P''}$.

Let $\mathscr{T}(P)$ denote the set of equivalence classes of P-tight admissible realizations. For $P' \subseteq P$, define

$$\mathrm{res}_{P,P'}[U] := [\mathcal{R}_{P'} U].$$

Proposition 8.3 (Tight reconstruction presheaf)**.** *Under theorem 8.1, the assignment $P \mapsto \mathscr{T}(P)$ defines a contravariant set-valued presheaf on the continuation context poset.*

Proof. Idempotence gives identity restriction, and nested reduction gives composition on equivalence classes. □

Let $P = \bigcup_{i \in I} P_i$ be a cover of continuation contexts.

Definition 8.4 (Matching family)**.** A family $U_i \in \mathscr{T}(P_i)$ is *matching* if, for every i, j,

$$U_i|_{P_i \cap P_j} \simeq U_j|_{P_i \cap P_j}.$$

Definition 8.5 (Effective matching family)**.** A matching family $\{U_i\}$ is *effective* if there exists $U \in \mathscr{T}(P)$ such that $U|_{P_i} \simeq U_i$ for every i.

Definition 8.6 (Admissible amalgam)**.** Let $\{U_i \in \mathscr{T}(P_i)\}$ be a matching family. An *admissible amalgam* is an admissible realization A that is sufficient for $P = \bigcup_i P_i$ and admits passages $A \to U_i$ for every i.

Theorem 8.7 (Admissible Amalgamation)**.** *Assume theorem 8.1. Let $P = \bigcup_i P_i$ and let $\{U_i \in \mathscr{T}(P_i)\}$ be a matching family. If the family admits an admissible amalgam A, then $U := \mathcal{R}_P A$ is a global P-tight realization satisfying $U|_{P_i} \simeq U_i$ for every i.*

Proof. By (R1), $U = \mathcal{R}_P A$ is admissible and P-sufficient. By (R2), it is P-tight. For each i,

$$U|_{P_i} = \mathcal{R}_{P_i}\mathcal{R}_P A \simeq \mathcal{R}_{P_i} A$$

by (R3). Since $A \to U_i$ and U_i is P_i-tight, (R4) gives $\mathcal{R}_{P_i} A \simeq U_i$. □

Corollary 8.8 (Effectivity and amalgamation)**.** *Under theorem 8.1, a matching family is effective if and only if it admits an admissible amalgam.*

Proof. The forward implication is immediate: a global realization inducing the family is itself an admissible amalgam. The reverse implication is theorem 8.7. □

Definition 8.9 (Recognition)**.** A cover $\{P_i \to P\}$ has the *recognition property* if $U|_{P_i} \simeq V|_{P_i}$ for every i implies $U \simeq V$ for all $U, V \in \mathscr{T}(P)$.

For a cover $\mathcal{U} = \{P_i \to P\}$, let $\mathrm{Match}_{\mathcal{U}}(\mathscr{T})$ denote the matching families and define

$$\Delta_{\mathcal{U}} : \mathscr{T}(P) \longrightarrow \mathrm{Match}_{\mathcal{U}}(\mathscr{T}), \qquad U \longmapsto (U|_{P_i})_i.$$

Theorem 8.10 (Fixed-Source Descent Decomposition)**.** *Assume theorem 8.1. The tight reconstruction presheaf satisfies the sheaf condition on a cover $\mathcal{U} = \{P_i \to P\}$ if and only if:*

(i) *every matching family admits an admissible amalgam;*

(ii) *the cover has the recognition property.*

Equivalently,

$$\boxed{\textit{effective descent} = \textit{admissible amalgamation} + \textit{recognition}.}$$

Proof. The comparison map $\Delta_{\mathcal{U}}$ is surjective exactly when every matching family is effective. By theorem 8.8, this is equivalent to existence of an admissible amalgam for every matching family. The map $\Delta_{\mathcal{U}}$ is injective exactly when the cover recognizes global tight realizations. Hence bijectivity is equivalent to (i) and (ii). □

8.1 A capacity obstruction

Let

$$D = \mathbb{R}^2, \qquad J_1(x, y) = x, \qquad J_2(x, y) = y,$$

and require admissible linear carriers to have rank at most one. The singleton contexts $P_1 = \{1\}$ and $P_2 = \{2\}$ admit tight realizations $q_1(x, y) = x$ and $q_2(x, y) = y$. Their overlap is empty, so the pair is matching. Any realization sufficient for $P_1 \cup P_2 = \{1, 2\}$ must distinguish both coordinates and therefore has rank at least two. No admissible amalgam exists. By theorem 8.8, the matching family is not effective.

Thus local compatibility at fixed source need not imply global admissible realization.

Remark 8.11. The capacity example fails by nonexistence of an admissible amalgam. It does not exhibit recognition failure. In the unrestricted core theory, the quotient $Q_{F,P}$ is canonical, so genuine recognition failure requires additional admissible structure that is invisible after every local contextual reduction.

The descent theorem above is intentionally conditional and set-valued. Nothing in the present analysis forces a higher comparison object: the core bifunctor remains strict, and the contextual reduction calculus has been quotiented by the chosen equivalence relation. If a future analytic realization theory produces nontrivial comparison data that cannot be removed in this way, a higher-valued descent problem may become necessary. No such claim is made here.

Fixed-source descent concerns compatibility across continuation contexts. It does not address the separate question of whether a coherent family of finite realizations possesses any admissible total carrier. The final section isolates that second finite-to-total failure.

9 Finite Coherence Without Admissible Total Representation

The obstruction of section 5 occurs after the total source carrier has already been constructed: relative multiplication fails to survive on that carrier. A different finite-to-total failure can occur one level earlier. Every finite continuation context may possess an admissible carrier, and those carriers may form a perfectly coherent inverse system, while no admissible carrier realizes the total continuation context.

The distinction is therefore

failure of an admitted operation

versus

failure of an admissible total carrier

Let

$$D = c_{00}(\mathbb{N}),$$

the vector space of finitely supported scalar sequences. For each $n \in \mathbb{N}$, define

$$J_n : D \to \mathbb{K}, \qquad J_n(x) = x_n.$$

Fix the source and let the continuation context vary through finite subsets $P \Subset \mathbb{N}$. The joint continuation map is

$$\Omega_P : D \to \mathbb{K}^P, \qquad \Omega_P(x) = (x_n)_{n \in P}.$$

Its kernel is

$$K_P = \{x \in c_{00}(\mathbb{N}) : x_n = 0 \text{ for every } n \in P\}.$$

Hence the core reconstruction carrier is

$$Q_P = D/K_P \cong \mathbb{K}^P. \tag{24}$$

Let the admissible carrier class consist of finite-dimensional vector spaces. For finite P, define

$$q_P : D \to \mathbb{K}^P, \qquad q_P(x) = (x_n)_{n \in P}.$$

Each required continuation is recovered by a coordinate projection, so every finite continuation context is admissibly realizable.

Proposition 9.1 (Finite realizability without an admissible total carrier)**.** *If admissible carriers are required to be finite dimensional, then every finite continuation context $P \Subset \mathbb{N}$ is realizable, while the total continuation context $P_\infty = \mathbb{N}$ is not.*

Proof. Finite realizability is provided by $q_P : D \to \mathbb{K}^P$. Suppose a finite-dimensional carrier X realized every coordinate continuation through some map $q : D \to X$. Then

$$\ker q \subseteq \bigcap_{n \geq 1} \ker J_n.$$

The coordinate family separates $c_{00}(\mathbb{N})$, so the intersection is $\{0\}$. Thus q would be injective, impossible because $c_{00}(\mathbb{N})$ is infinite dimensional while X is finite dimensional. □

The failure of total admissibility is not caused by inconsistency among the finite carriers. For $P \subseteq Q$, let $\rho_{Q,P} : \mathbb{K}^Q \to \mathbb{K}^P$ be coordinate restriction. These maps form a coherent inverse system

$$\left(\mathbb{K}^P, \rho_{Q,P}\right)_{P \Subset \mathbb{N}}.$$

In the category of vector spaces, its inverse limit exists:

$$\varprojlim_{P \Subset \mathbb{N}} \mathbb{K}^P \cong \mathbb{K}^{\mathbb{N}}. \tag{25}$$

Indeed, a coherent family of finite coordinate vectors is precisely an arbitrary scalar sequence. The finite system therefore possesses an ambient totalization. What fails is membership of that totalization in the prescribed finite-dimensional carrier class.

Definition 9.2 (Pro-realization)**.** A *pro-realization* of a continuation system is a coherent inverse system of admissible finite-context carriers realizing every finite continuation context.

The coordinate system above is therefore pro-realizable without being admissibly representable by a single total carrier.

The finite realizations induce a canonical map

$$q_\infty : c_{00}(\mathbb{N}) \longrightarrow \mathbb{K}^{\mathbb{N}},$$

which is simply the inclusion

$$c_{00}(\mathbb{N}) \hookrightarrow \mathbb{K}^{\mathbb{N}}. \tag{26}$$

This map is not surjective. Hence the ambient inverse limit contains coherent total points that do not arise from any state of the original source.

Proposition 9.3 (Carrier totalization and state effectivity are distinct)**.** *In the coordinate model:*

(i) *every finite continuation context admits an admissible carrier;*

(ii) *the coherent finite carrier system has ambient inverse limit $\mathbb{K}^{\mathbb{N}}$;*

(iii) *no admissible finite-dimensional total carrier exists;*

(iv) *the canonical source map $c_{00}(\mathbb{N}) \to \mathbb{K}^{\mathbb{N}}$ is not surjective.*

Thus admissible carrier totalization and effectivity of coherent limit points are logically distinct requirements.

Proof. Statements (i)–(iii) follow from theorem 9.1 and (25). Statement (iv) is (26), since most scalar sequences are not finitely supported. □

The coordinate example and the relative multiplication example therefore fail for different reasons. For relative multiplication, the total source carrier exists but an admitted continuation becomes nonclosable. For the coordinate system, every finite continuation is harmless and the finite carriers are mutually coherent, but the total carrier leaves the prescribed structural class. In the first case the obstruction is carried by the graph of an operation. In the second it is carried by the admissibility condition on the carrier itself.

Operation failure is not carrier failure.

The proper inclusion $c_{00}(\mathbb{N}) \subsetneq \mathbb{K}^{\mathbb{N}}$ exhibits a third logically distinct issue: even after an ambient limit carrier has been constructed, a coherent limit point need not come from the original source. The examples therefore distinguish at least three finite-to-total questions:

(a) does an admissible total carrier exist?

(b) do the required operations survive on it?

(c) are its coherent points represented by original source states?

We do not claim that these alternatives form an exhaustive classification of reconstruction obstructions.

9.1 The joint finite-to-total problem

The two principal examples vary one contextual coordinate at a time. Relative multiplication fixes the continuation demand and enlarges the source. The coordinate model fixes the source and enlarges the continuation context. The genuinely mixed problem allows both to vary:

$$F_1 \leq F_2 \leq \cdots \to F_\infty, \qquad P_1 \subseteq P_2 \subseteq \cdots \to P_\infty.$$

Open Problem 9.4 (Joint reconstruction compactness). Find structural hypotheses on the source-window system, analytic carrier class, continuation family, and morphism class under which admissibility of a directed family (F_n, P_n) forces admissibility of the joint total context (F_∞, P_∞).

Theorem 6.3 supplies one sufficient source-direction certificate for individual continuations. No general two-axis compactness theorem is asserted here.

Finite coherence is therefore weaker than total admissible representation. The distinction persists at several levels: carriers may leave the structural class, operations may fail after completion, and coherent limit points may fail to come from the original source. The common lesson is not that finite reconstruction is unreliable. It is that finite validity certifies only the finite problem. Passage to a total realization requires control of the class in which the carrier, its operations, and its represented states are required to remain.

The object being transported and the certificate authorizing its transport must survive the limiting passage together.

References

References

[1] S. Abramsky and A. Brandenburger, *The sheaf-theoretic structure of non-locality and contextuality*, New J. Phys. **13** (2011), 113036.

[2] U. Haagerup, *Operator valued weights in von Neumann algebras, I*, J. Funct. Anal. **32** (1979), 175–206.

[3] U. Haagerup, *Operator valued weights in von Neumann algebras, II*, J. Funct. Anal. **33** (1979), 339–361.

[4] V. F. R. Jones, *Index for subfactors*, Invent. Math. **72** (1983), 1–25.

[5] H. Kosaki, *Extension of Jones' theory on index to arbitrary factors*, J. Funct. Anal. **66** (1986), 123–140.

[6] S. Popa, *Strong rigidity of* II_1 *factors arising from malleable actions of w-rigid groups, I*, Invent. Math. **165** (2006), 369–408.