

RESEARCH LIBRARY

Stratified Universality for Resolved Representations

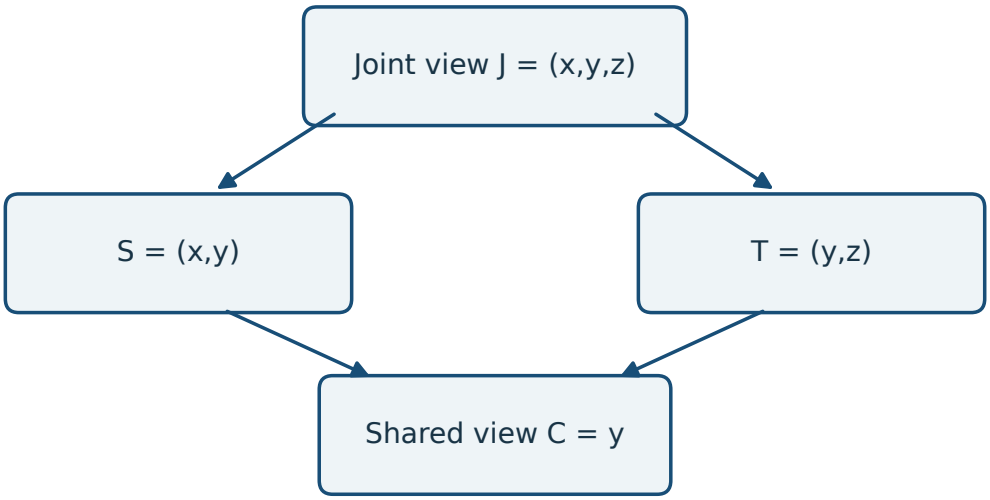
Profiled lane transport, semantic quotients, carrier change, and certified descent

Visual guide

Komorebi P.

Standardized publication edition / 8 October 2026

Reading scope



A joint representation preserves both declared views. Their shared quotient supplies the compatibility interface.

Contents

1 Reading this guide 3

2 A stratum has a passport 4

3 A theorem crosses a verified profile 5

4 Exact packets can be oblique 6

5 Transfer blocks retain the full map 7

6 Composition counts compatible paths 8

7 Coarser complete coordinates can be lossless 9

8 A selector names what it discards 10

9 Residuals propagate through composition 11

10 A fold has a collision kernel 12

11 Chart paths and the flat core 13

12 Twisted lanes still transport exactly 14

13 Compatible support and nonzero action differ 15

14 A semantic slice declares the task 16

15 Finite partition refinement discovers future distinctions 17

16 A linear invisible core must survive continuations 18

17 Task strengthening retains more information 19

18 Two views form a semantic diamond 20

19 A relative residual supplies the missing view 21

20 Carrier replacement is a factorization question 22

21 The same route graph can use different carrier laws 23

22 Observer change creates a relation object 24

23 An infinite discrete quotient can have finite linear rank 25

24 Value, witness, and history are different tasks 26

25 A behavioral bridge has two independent defects 27

26 Admissibility change is a different order 28

27 Convex and quadratic geometry have different bounds 29

28 Zero infimum need not be attained 30

29 Metric completion is a holonomy gate 31

30 Forgetting structure and rebuilding it differ 32

31 Finite atlas assembly needs effective gluing 33

32 Example ledger and reproducibility 34

1 Reading this guide

The paper asks when several descriptions of one mathematical object can be compared, transported, or assembled without losing a distinction that matters. Its answer is relative: first declare the operations, the permitted realizations, the continuation rules, and the observations. Then prove the relevant comparison.

This guide follows one thread. We begin with exact finite packet coordinates. We then compare complete coordinate changes with selectors and folds, study chart loops, and turn to task-dependent quotients. The last examples separate carrier change, observer change, and admissibility change. Each step introduces a different obligation.

Every worked example uses a stated finite model. Exact rational arithmetic checks its displayed identities. A finite computation supports the example, while the general theorem remains the written argument in the manuscript. The metric curve near the end is a floating-point illustration of an analytically proved limit.

Question	Mathematical object to inspect
Can the proof cross this bridge?	Preservation profile and target hypotheses
Is a coordinate record complete?	Analysis and synthesis inverses
What was discarded?	Selector residual, fold kernel, or quotient kernel
Do chart paths agree?	Closed-word holonomy and coherent comparisons
Which distinctions matter?	Future observations and behavioral kernel
Can one carrier replace another?	Semantic fiber containment
Can a missing completion be realized?	Admissible class and attainment

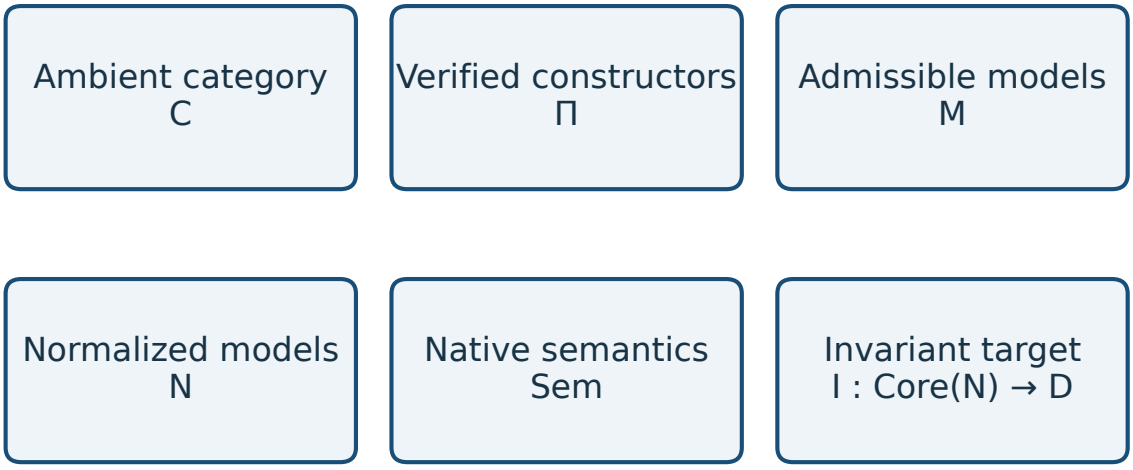
The guide is a companion, not a substitute for the 262-page standardized traditional paper. The labels printed on each page point to its exact definitions and results. A label identifies a manuscript result; it is not a Lean declaration or a new certification claim.

Manuscript anchor: `sec:resolved-foundations;sec:block19-admissibility-lifting`

2 A stratum has a passport

An object is represented inside an ambient category. Its allowed models form a smaller category, and its classification theorem may hold only on a normalized subcategory. A passport keeps these levels separate.

$$\mathbf{S}_\alpha = (\mathcal{C}_\alpha, \Pi_\alpha, \mathcal{M}_\alpha, \mathcal{N}_\alpha, \mathrm{Sem}_\alpha, \mathcal{I}_\alpha).$$



The six fields specify where a model lives, which operations are preserved, and exactly what is classified.

For example, a rational generating function and a finite recurrence can represent the same scalar sequence. A reduced generating function and a minimal reachable-observable state realization may carry a classification theorem. An arbitrary larger state presentation can contain unused coordinates. Equality of the scalar sequence therefore does not identify every such larger presentation.

A normalization witness includes a functor and a natural comparison of semantics. It also records its domain, auxiliary choices, and any restriction. “Normalize” is not permission to delete structure merely because it seems redundant.

An invariant is complete on its declared target when

$$\mathcal{I}(X) \cong \mathcal{I}(Y) \quad \Longleftrightarrow \quad X \cong Y \qquad (X, Y \in \mathcal{N}).$$

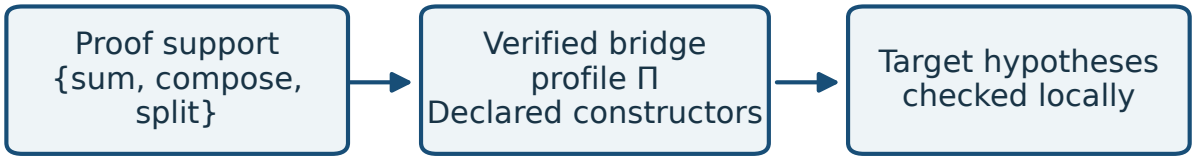
This quantifier is the boundary of the classification claim. To transport that claim across an overlap, the bridge must preserve the normalization and the relevant isomorphisms.

Manuscript anchor: `def:profiled-stratum-passport;def:normalization-witness;def:complete-invariant-target`

3 A theorem crosses a verified profile

A preservation profile records the constructors a bridge actually preserves. A proof-support certificate records the constructors its proof actually uses. The transport gate compares these two records.

$$\mathrm{Supp}(\mathrm{proof}) \subseteq \Pi_F.$$



Both the constructor inclusion and the target-specific hypotheses must pass.

Suppose an argument uses finite sums, composition, and split packet idempotents. An additive bridge with the required packet comparisons transports that equational argument. An argument that also uses a kernel, tensor product, limit, positive metric, or completed sum requires the corresponding additional preservation witness.

For a chain of bridges, a safe common profile is the intersection of their verified profiles. A construction preserved by only the first bridge does not silently survive the second.

The distinction is practical. Calling a frequency band, a residue class, and a latent channel a “packet” gives a shared vocabulary. To apply the finite packet theorem, one must still supply idempotents, annihilation, reconstruction, and the chosen splittings. To claim an energy identity, one must also supply the Hilbert and orthogonality hypotheses.

Proof step	Additional gate
Finite block reconstruction	Additivity and complete split packets
Exact quotient transport	Required quotient preservation
Norm estimate	Declared norm and its inequalities
Global gluing	Coherent comparisons and effective descent

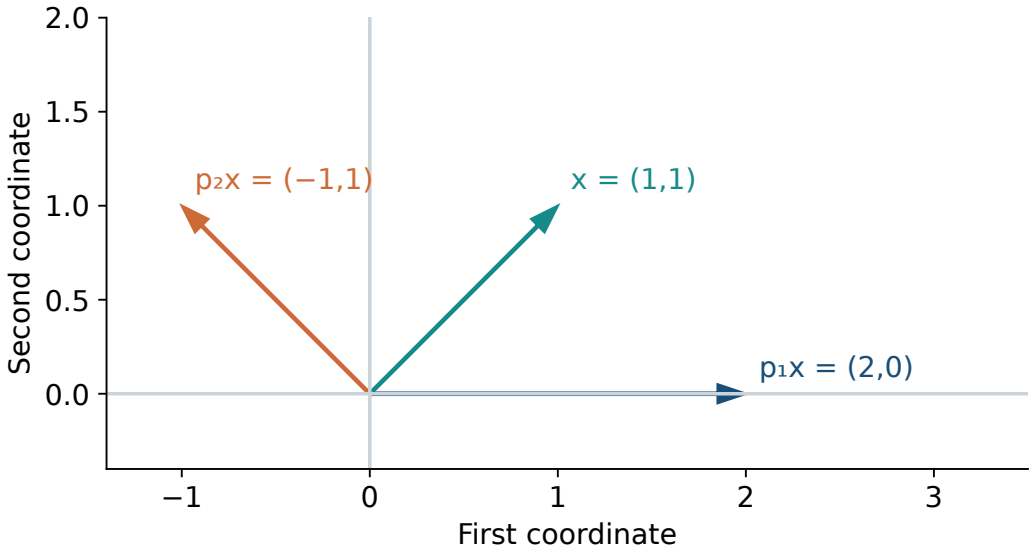
Manuscript anchor: `def:verified-preservation-profile-expanded;def:proof-support-certificate;prop:profile-safe-theorem-transport`

4 Exact packets can be oblique

Consider two endomorphisms of a two-dimensional vector space:

$$p_1 = \begin{pmatrix} 1 & 1 \\ 0 & 0 \end{pmatrix}, \qquad p_2 = \begin{pmatrix} 0 & -1 \\ 0 & 1 \end{pmatrix}.$$

They satisfy $p_i^2 = p_i$, $p_1 p_2 = p_2 p_1 = 0$, and $p_1 + p_2 = I$. For $x = (1, 1)$, the packets are $p_1 x = (2, 0)$ and $p_2 x = (-1, 1)$.



The packet vectors sum to the original vector. Their directions are not orthogonal in the standard metric.

The analysis and synthesis matrices are

$$A = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}, \qquad V = \begin{pmatrix} 1 & -1 \\ 0 & 1 \end{pmatrix}, \qquad AV = VA = I.$$

Analysis records the two packet coefficients; synthesis puts them back. The inverse identities use algebraic splitting, not a metric. This example prevents a common mistaken inference: a complete decomposition need not be an orthogonal decomposition.

The exact-check script verifies both inverse identities and all packet equations with rational arithmetic. Neither a numerical tolerance nor an approximate pseudoinverse enters this example.

Manuscript anchor: `ex:oblique-packets;thm:exact-packet-coordinate-expanded;bnd:algebraic-vs-metric`

5 Transfer blocks retain the full map

Let $f : X \to Y$ have complete finite source and target packets. The ambient corners are $q_a f p_i$. With chosen splittings, the smaller lane maps are

$$f_{ai} = \varpi_a f \iota_i : X_i \to Y_a.$$

These objects have different types. The small map lives between packet objects; its ambient corner is $\jmath_a f_{ai} \pi_i$. Reconstruction is

$$f = \sum_{a,i} q_a f p_i = \sum_{a,i} \jmath_a f_{ai} \pi_i.$$

For the preceding oblique coordinates, take

$$f = \begin{pmatrix} 2 & 3 \\ 5 & 7 \end{pmatrix}.$$

Its complete lane matrix is

$$h_f = AfV = \begin{pmatrix} 7 & 3 \\ 5 & 2 \end{pmatrix}.$$

The entries change because the basis changes. They still retain the entire map:

$$V h_f A = f.$$

The word “complete” matters. Retaining only a trace, a norm, a total sum, or an endpoint label generally gives a different map from this full lane analysis. Such a summary can have a nontrivial kernel.

The theorem is an isomorphism between the morphism group and the full finite matrix record. It does not assert that every scalar statistic of the record is a complete invariant.

Manuscript anchor: `def:mixed-transfer-blocks-expanded;thm:transfer-reconstruction-expanded;thm:block15-universal-lane-transport`

6 Composition counts compatible paths

For $X \xrightarrow{f} Y \xrightarrow{g} Z$, insert the complete intermediate packet identity. The composite block is

$$(gf)_{ci} = \sum_a g_{ca} f_{ai}.$$

For $g = \left(\begin{smallmatrix} 1 & 2 \\ 3 & 4 \end{smallmatrix}\right)$, the same oblique coordinate matrices give

$$h_g = AgV = \begin{pmatrix} 4 & 2 \\ 3 & 1 \end{pmatrix}, \qquad h_g h_f = \begin{pmatrix} 38 & 16 \\ 26 & 11 \end{pmatrix}.$$

Direct analysis of gf produces that same matrix. Each summand corresponds to a compatible intermediate lane; longer strings expand over all compatible lane paths.

Support propagation is an inclusion. If the blocks of f are supported on E and those of g on F, then the blocks of gf are supported on the relational composite $F \circ E$. Cancellation can make the actual support smaller.

For example,

$$f = \begin{pmatrix} 1 \\ 1 \end{pmatrix}, \qquad g = \begin{pmatrix} 1 & -1 \end{pmatrix}, \qquad gf = 0.$$

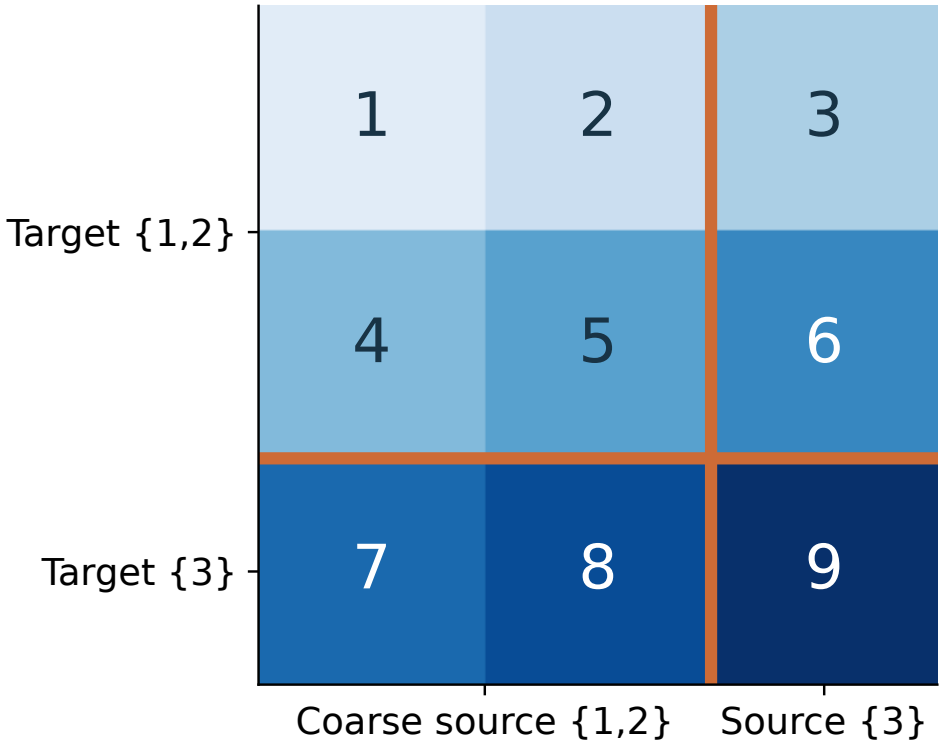
There are two nonzero compatible path contributions, 1 and -1. Their sum vanishes. A Boolean path ledger records where a contribution can occur; it does not replace the additive sum of contributions.

This is one reason carrier laws matter later in the guide. Signed addition, support union, and maximum combine the same path skeleton differently.

Manuscript anchor: `cor:packet-path-expansion;prop:support-propagation;thm:block15-matrix-category-equivalence`

7 Coarser complete coordinates can be lossless

Refinement groups fine packet addresses inside coarse ancestors. The coarse projectors are sums of the corresponding fine projectors. A coarse corner therefore retains a map between larger packet objects.



The four coarse corners have sizes 2×2, 2×1, 1×2, and 1×1. Together they retain all nine entries.

For the displayed 3×3 matrix, the coarse packets are $\{1, 2\}$ and $\{3\}$ in both source and target. The upper-left coarse block is the matrix $\left(\begin{smallmatrix} 1 & 2 \\ 4 & 5 \end{smallmatrix}\right)$, not its scalar sum.

The certified coordinate change is

$$\Gamma_{\mathrm{fine}}^{\mathrm{coarse}} = \mathfrak{a}_{\mathrm{coarse}} \mathfrak{a}_{\mathrm{fine}}^{-1}.$$

Both analyses are complete, so this map is invertible. Refinement chains compose strictly, and the coordinate changes intertwine matrix multiplication.

If we subsequently sum each coarse block into one number, information is discarded. That operation is a fold or statistic. Its kernel must be computed separately. A reduction in address count alone does not measure a reduction in retained linear dimension.

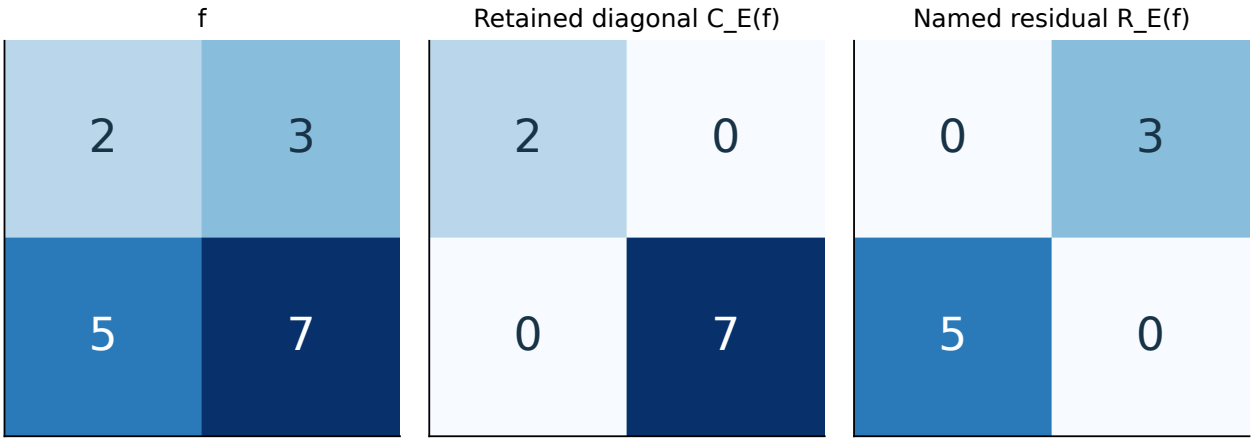
Manuscript anchor: `thm:exact-refinement-expanded;prop:block15-certified-coarsening;rem:block15-address-count-not-dimension`

8 A selector names what it discards

Fix source-target positions E. The selector and its residual are

$$C_E(f) = \sum_{(a,i)\in E} q_a f p_i, \qquad R_E(f) = f - C_E(f).$$

For standard coordinate packets and the diagonal support E, we obtain



The retained and residual matrices sum exactly to the original matrix.

The selector is an idempotent on the morphism space:

$$C_E^2 = C_E, \qquad C_E C_F = C_{E\cap F}, \qquad C_E R_E = R_E C_E = 0.$$

Here C_E acts on maps, rather than on source vectors. This type distinction matters when composing selectors with transfers or comparing supports.

The phrase “named residual” means the discarded map remains explicit. A later operation can inspect it, bound it, or combine it with other residuals. It is not automatically negligible.

If the packets are orthogonal projectors on Hilbert spaces and f is Hilbert–Schmidt, the selector is also an orthogonal projection in Hilbert–Schmidt operator space. In that extra profile, its residual satisfies an exact sum-of-squared-block-norm identity and the retained map is the best approximation on the selected support. The earlier oblique example does not automatically have those metric properties.

Manuscript anchor: `def:support-compression-expanded;prop:block-selector-algebra;thm:hs-optimality-expanded`

9 Residuals propagate through composition

Write $f = f_E + r_f$ and $g = g_F + r_g$. Expanding the product gives

$$gf - g_F f_E = r_g f_E + g_F r_f + r_g r_f.$$

All three terms are needed. The last term describes interaction between two discarded parts.

Using the standard-coordinate diagonal selectors from the preceding page,

$$f_E = \begin{pmatrix} 2 & 0 \\ 0 & 7 \end{pmatrix}, \quad r_f = \begin{pmatrix} 0 & 3 \\ 5 & 0 \end{pmatrix}, \quad g_F = \begin{pmatrix} 1 & 0 \\ 0 & 4 \end{pmatrix}, \quad r_g = \begin{pmatrix} 0 & 2 \\ 3 & 0 \end{pmatrix}.$$

The terms are

$$r_g f_E = \begin{pmatrix} 0 & 14 \\ 6 & 0 \end{pmatrix}, \qquad g_F r_f = \begin{pmatrix} 0 & 3 \\ 20 & 0 \end{pmatrix}, \qquad r_g r_f = \begin{pmatrix} 10 & 0 \\ 0 & 9 \end{pmatrix}.$$

Their sum is

$$gf - g_F f_E = \begin{pmatrix} 10 & 17 \\ 26 & 9 \end{pmatrix}.$$

In a declared submultiplicative norm, the identity implies

$$\|gf - g_F f_E\| \le \|r_g\|\|f_E\| + \|g_F\|\|r_f\| + \|r_g\|\|r_f\|.$$

This estimate is conditional on the normed-operator profile. The algebraic identity itself needs only the finite additive setting. A software approximation claim should report which residuals were charged and which norm supplied the estimate.

Manuscript anchor: `prop:exact-residual-product`

10 A fold has a collision kernel

Consider the additive fold

$$\Phi : \mathbb{Q}^2 \to \mathbb{Q}, \qquad \Phi(u, v) = u + v.$$

The inputs $(1, 0)$ and $(0, 1)$ share the same output. Their difference $(1, -1)$ lies in

$$\ker \Phi = \mathbb{Q}(1, -1).$$

The exact sequence

$$0 \to \mathbb{Q}(1, -1) \to \mathbb{Q}^2 \xrightarrow{\Phi} \mathbb{Q} \to 0$$

records precisely the distinctions that disappeared. A kernel is more informative than a vague statement that the record became “smaller.”

If a declared fold comes with a decoder D satisfying $D\Phi = I$ on its source, it is split lossless there. Such a decoder cannot exist for this full two-dimensional source: $D\Phi(1, -1) = 0$ cannot equal $(1, -1)$. It can exist after restricting to an explicitly chosen normalized core.

This gives a useful three-way distinction:

- A complete coarse corner record changes coordinates.
- A selector drops named positions and keeps a residual.
- A fold merges positions and has a collision kernel.

In an abelian target, image and cokernel add a second question: which target records are actually attainable? Kernel loss and essential-image failure are independent obligations.

Manuscript anchor: `thm:block15-universal-lane-transport;thm:block15-abelian-fold-exact-sequence;bnd:block15-coordinate-vs-statistic-coarsening`

11 Chart paths and the flat core

Take a connected finite graph of vector spaces with invertible edge transports, and let reverse edges carry inverse transports. Choose a root and a spanning tree. Tree transport identifies every fiber with the root fiber; each remaining edge contributes a closed-word holonomy H_e.

$$C_o = \bigcap_e \ker(H_e - I).$$

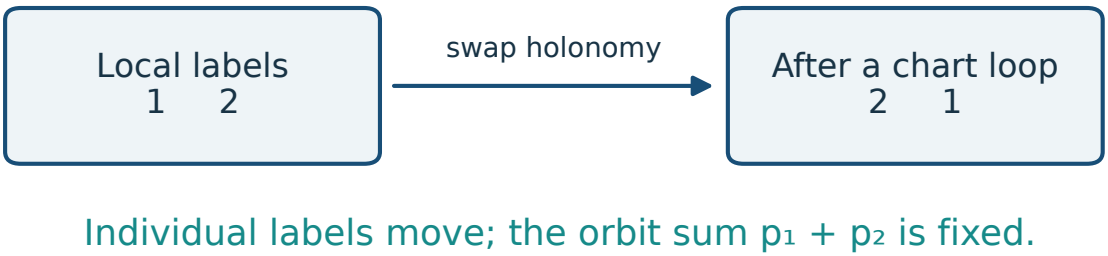
The root values of global parallel sections are exactly the vectors in C_o. The stacked defect map $\Delta x = ((H_e - I)x)_e$ yields

$$\dim C_o = \dim F_o - \operatorname{rank} \Delta.$$

For one swap holonomy

$$S = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \qquad C_o = \mathbb{Q}(1,1).$$

The defect has rank one. A compatible section survives, but the full two-dimensional fiber has no path-independent global frame.



A loop can move local labels while leaving their orbit sum fixed.

Changing the spanning tree or the fiber coordinates changes the displayed holonomies by the appropriate transport and conjugation. It does not change the intrinsic global-section space or its dimension. This is a finite graph theorem with declared invertible transports, not a theorem about arbitrary heterogeneous gluing.

Manuscript anchor: `thm:maximal-flat-core;cor:global-section-versus-frame`

12 Twisted lanes still transport exactly

Under finite address holonomy, there need not be one global label for every lane. The full matrix record instead lives in a matrix local system. A chart transition carries each position to its transported source-target address, together with the splitting comparisons.

$$(\gamma_* h)_{\lambda^Y(a),\lambda^X(i)} = \beta_a^Y F_\gamma(h_{ai})(\beta_i^X)^{-1}.$$

A global record is a compatible section: $\gamma_* h_U = h_V$. Exact analysis and synthesis identify descended morphisms with these sections. The compatible-section condition is part of the global type of the record.

For the swap representation of C_2, equivariant endomorphisms satisfy $Sh = hS$, hence

$$h = \begin{pmatrix} a & b \\ b & a \end{pmatrix}.$$

Composition remains exact:

$$\begin{pmatrix} c & d \\ d & c \end{pmatrix} \begin{pmatrix} a & b \\ b & a \end{pmatrix} = \begin{pmatrix} ca+db & cb+da \\ cb+da & ca+db \end{pmatrix}.$$

The orbit sum $p_1 + p_2 = I$ descends, while the individual projectors exchange under the loop. There is one object-address orbit but two matrix-position orbits. Thus failure of individual global labels does not erase the entire transport algebra.

The theorem keeps additive preservation, address cocycles, equivariant refinement, and splitting comparisons. Removing the global address trivialization does not remove these obligations.

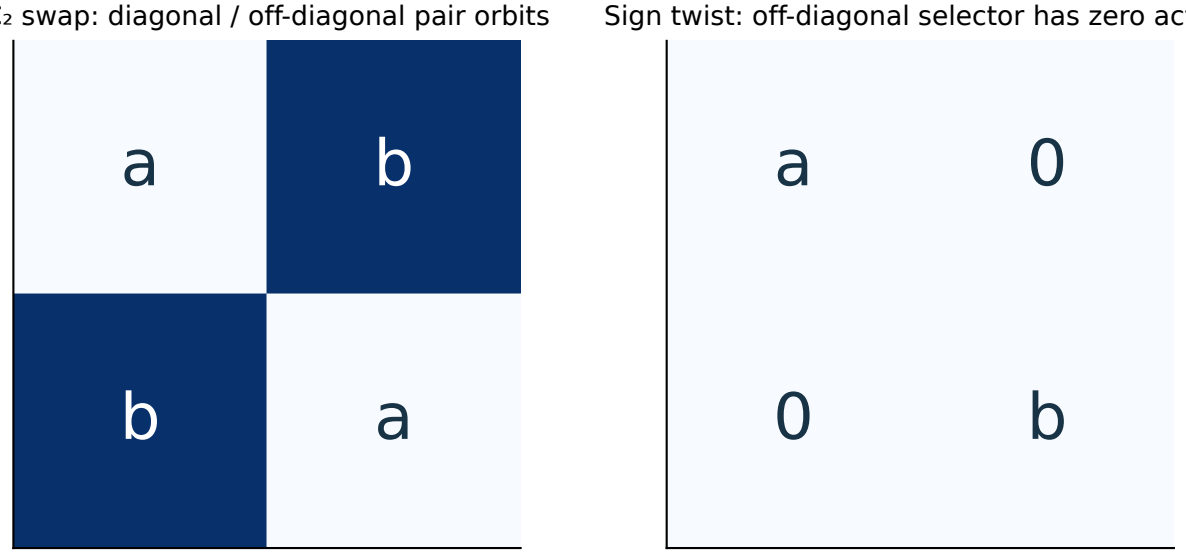
Manuscript anchor: `thm:block16-holonomy-twisted-lane-transport;ex:block16-c2-swap;rem:block16-removed-triviality-only`

13 Compatible support and nonzero action differ

In the swap example, the pair orbits are

$$\Delta = \{(1,1),(2,2)\}, \qquad \Omega = \{(1,2),(2,1)\}.$$

Their support selectors retain the diagonal and off-diagonal parts of h. A singleton such as $\{(1,1)\}$ is not compatible: it sends the invariant identity to a projector that moves under the swap.



Pair-orbit supports are minimal nonempty compatible supports. Their action also depends on the splitting twist.

Now keep both addresses fixed, but let the loop act on the two lane fibers by

$$D = \mathrm{diag}(1,-1)$$

over $\mathbb{Q}$. Equivariance forces $Dh = hD$, so $h = \mathrm{diag}(a,b)$. Every singleton position is a compatible support because the address action is trivial. Nevertheless, the nonempty off-diagonal support $\{(1,2)\}$ selects zero on every equivariant map.

This example separates an address-level statement from a statement about the section space. A minimal nonempty support need not define a nonzero endomorphism of that space. The permutation-module case has a stronger orbit-indicator basis theorem; an arbitrary twisted local system does not inherit it merely from its address orbits.

Manuscript anchor: `thm:block16-orbit-descent;thm:block16-permutation-orbital-calculus`

14 A semantic slice declares the task

The next question is not which coordinates reconstruct the entire object. It is which distinctions matter for a declared family of future observations.

$$\Sigma = (\mathcal{A}, \mathcal{G}, \mathcal{V}, \mathcal{O}).$$

The carrier doctrine $\mathcal{A}$ specifies how alternatives combine. The continuation grammar $\mathcal{G}$ specifies legal futures. The observation category $\mathcal{V}$ specifies the kind of observed values. The family $\mathcal{O}$ specifies the task.

For a continuation monoid M, action ρ, and observations o_j, behavior is

$$\mathrm{Beh}(x) = \big(o_j(\rho(m)x)\big)_{m,j}.$$

Two states are behaviorally equivalent when all these values agree. In a module category, the invisible submodule is

$$N = \bigcap_{m,j} \ker(o_j \rho(m)), \qquad Q = X/N.$$

The quotient is the coarsest exact representation for that task whenever the required quotient exists. Any surjective representation retaining all future observations admits a unique map to Q.

The semantic slice and the preservation profile serve different roles. The profile tells us what a bridge preserves. The slice tells us which distinctions the task asks us to retain. A bridge can preserve the carrier operations while failing to preserve the declared task distinctions.

Changing the continuation grammar can change the quotient even when the present observation is unchanged. The next two examples show why present invisibility alone is insufficient.

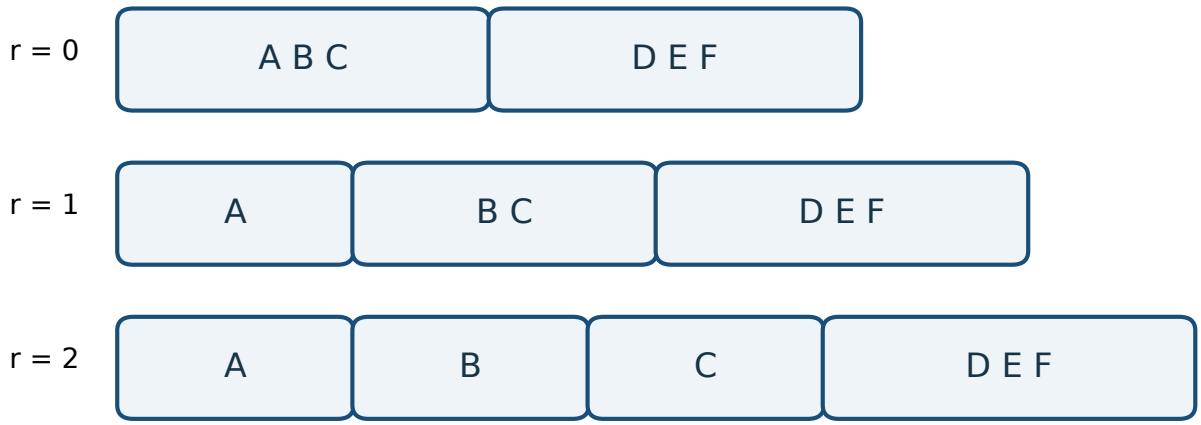
Manuscript anchor: `def:block17-semantic-slice;def:block17-observational-system;thm:block17-universal-behavioral-coimage`

15 Finite partition refinement discovers future distinctions

Let the state set be $\{A, B, C, D, E, F\}$. The observation is zero on A, B, C and one on D, E, F. There is one continuation generator with transitions

$$A \mapsto D,\ B \mapsto A,\ C \mapsto B,\ D \mapsto F,\ E \mapsto D,\ F \mapsto D.$$

States with the same present observation start in the same class. Split a class whenever its members transition into different current classes.



After two strict refinements, the partition stabilizes at $\{A\}, \{B\}, \{C\}, \{D, E, F\}$.

The first split distinguishes A from B, C: one step reaches observation one from A, but not from B, C. The next split distinguishes B from C. The three states D, E, F remain indistinguishable because every future stays inside the observation-one class.

At stabilization, every class has constant observation and transitions to a single class. Induction on continuation length shows all future observations agree within a class. Conversely, each split records a finite future that distinguishes the separated states.

This is a finite certificate for this system. The general behavioral-core theorem explains why the stabilized partition is the greatest continuation-stable invisible congruence.

Manuscript anchor: `thm:block17-behavioral-core;cor:block17-finite-certification`

16 A linear invisible core must survive continuations

Take $X = \mathbb{Q}^3$, one continuation matrix T, and one observation row o:

$$T = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \qquad o = \begin{pmatrix} 1 & 0 & 0 \end{pmatrix}.$$

The current observation ignores e_2 and e_3, so $N_0 = \operatorname{span}(e_2, e_3)$. But $Te_2 = e_1$, which becomes visible. In contrast, $Te_3 = e_3$ remains invisible at every future step.

$$oT = \begin{pmatrix} 0 & 1 & 0 \end{pmatrix}, \qquad oT^2 = 0, \qquad N = \mathbb{Q}e_3.$$

The exact future-observation stack has rank two, so the quotient has dimension two. This conclusion is certified by exact row reduction on the displayed rational matrices.

For a generating family A, define

$$\Phi(N) = N_0 \cap \bigcap_{a \in A} \rho(a)^{-1}(N).$$

Starting at N_0, the descending iteration gives $N^{(1)} = \mathbb{Q}e_3$, then $N^{(2)} = N^{(1)}$. The greatest invisible stable core is a fixed point of Φ, and every fixed point is invisible and stable. Smaller stable invisible submodules need not themselves be fixed points.

In finite dimension, each strict inclusion lowers dimension. The process therefore stabilizes after at most $\dim X$ strict dimension drops. Finite stabilization is not an effectivity theorem for every infinite-dimensional or unrestricted semantic slice.

Manuscript anchor: `thm:block17-coinductive-characterization;cor:block17-finite-certification`

17 Task strengthening retains more information

Adding observations or continuations can only shrink the invisible kernel. If $N_s \subseteq N_w$ are the strong and weak task kernels, then

$$Q_s = X/N_s \twoheadrightarrow X/N_w = Q_w.$$

The direction of this forgetting map is easy to reverse accidentally: the strong task retains more, so it can answer the weak task after forgetting.

For $X = \mathbb{Q}^3$ with identity continuations, let the weak task observe only x, and the strong task observe (x, y). Then

$$N_w = \operatorname{span}(e_2, e_3), \qquad N_s = \operatorname{span}(e_3).$$

The forgetting map is $(x, y) \mapsto x$. Its exact task-loss sequence is

$$0 \to \mathbb{Q}e_2 \to \mathbb{Q}^2 \to \mathbb{Q} \to 0.$$

The lost coordinate is specified by a quotient object N_w/N_s, not by the number of words used to describe the task.

For three nested task kernels, the total loss fits into

$$0 \to N_2/N_1 \to N_3/N_1 \to N_3/N_2 \to 0.$$

This is extension data. Over a field, a splitting can be chosen, but is generally not canonical. Over other modules, the extension may fail to split. For example, the quotient map $\mathbb{Z} \to \mathbb{Z}/2\mathbb{Z}$ has no additive section: an element of order two cannot map to a nonzero element of the torsion-free group $\mathbb{Z}$.

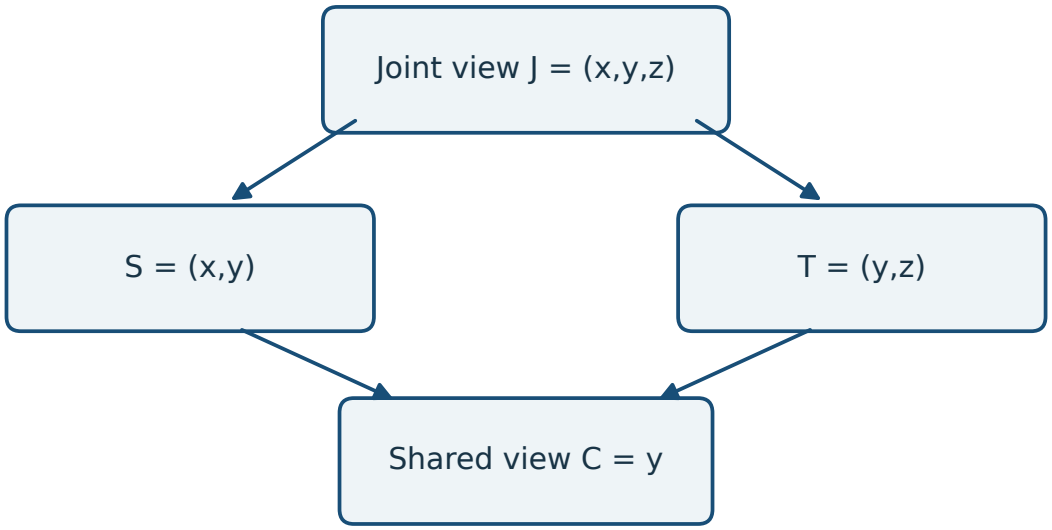
Manuscript anchor: `cor:block17-task-strengthening;thm:block17-task-change;cor:block17-task-loss-composition;bnd:block17-residual-splitting`

18 Two views form a semantic diamond

Let $X = \mathbb{Q}^3$, with two semantic maps

$$s(x,y,z) = (x,y), \qquad t(x,y,z) = (y,z).$$

Their kernels are $N = \mathbb{Q}e_3$ and $M = \mathbb{Q}e_1$. The joint quotient retains all three coordinates, while the shared quotient retains only y.



The shared coordinate is the exact compatibility condition for assembling the two views.

Thus

$$J = X/(N \cap M) = X, \qquad C = X/(N + M) \cong \mathbb{Q}.$$

A pair $((a,b),(c,d)) \in S \oplus T$ comes from a joint state precisely when $b = c$. The exact sequence is

$$0 \to \mathbb{Q}^3 \xrightarrow{\alpha} \mathbb{Q}^4 \xrightarrow{\beta} \mathbb{Q} \to 0,$$

with $\alpha(x,y,z) = (x,y,y,z)$ and $\beta(a,b,c,d) = b - c$. The guide's checker verifies ranks three and one, $\beta\alpha = 0$, and compatible and incompatible sample pairs.

Manuscript anchor: `thm:block17-semantic-diamond`

19 A relative residual supplies the missing view

In the diamond, the residual needed to augment S with T is

$$R_{T|S} = \ker(J \to S) \cong N/(N \cap M).$$

For $s(x,y,z) = (x,y)$ and $t(x,y,z) = (y,z)$, this residual is $\mathbb{Q}e_3$: the source view already knows y, but must add z. In the reverse direction, the residual is $\mathbb{Q}e_1$. Relative residuals need not be symmetric.

The factorization gate is

$$t \text{ factors through } s \quad \Longleftrightarrow \quad N \subseteq M \quad \Longleftrightarrow \quad R_{T|S} = 0.$$

The first example fails the gate. States $(0,0,0)$ and $(0,0,1)$ have the same S-view but different T-views. This pair is an exact witness that no map from S can recover the complete T-view.

If we replace t by $t'(x,y,z) = y$, the gate passes: t' is the second-coordinate projection of s. There is no missing relative residual for this weaker target task.

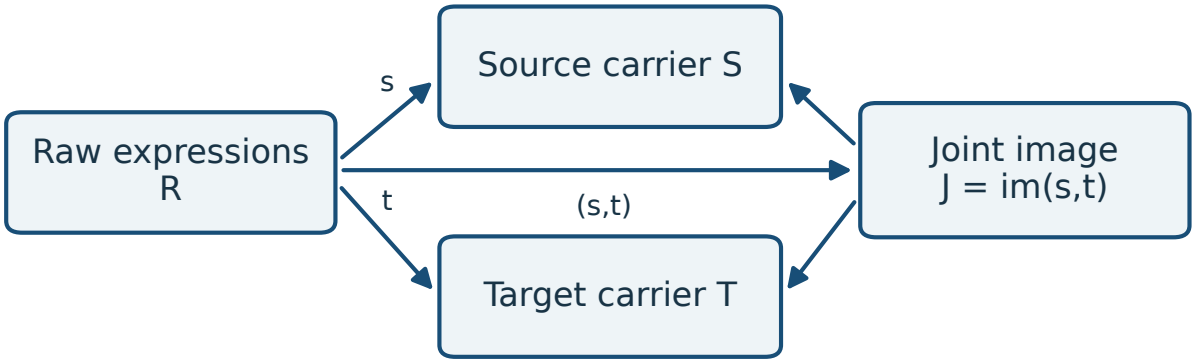
The exact sequence $0 \to R_{T|S} \to J \to S \to 0$ does not in general supply an independent residual coordinate. A constructor that stores S and the residual as a direct product must also provide the relevant splitting.

Manuscript anchor: `def:block17-relative-residual;cor:block17-factorization-criterion;bnd:block17-residual-splitting`

20 Carrier replacement is a factorization question

Let raw expressions $\mathcal{R}$ have two semantic interpretations, $s : \mathcal{R} \to S$ and $t : \mathcal{R} \to T$, preserving a fixed operation signature. An exact replacement map exists on $\operatorname{im} s$ precisely when

$$\operatorname{KerPair}(s) \subseteq \operatorname{KerPair}(t).$$



The joint image retains both semantics when the target cannot be reconstructed from the source.

The inclusion says that every distinction erased by s is also erased by t. If $s(x) = s(y)$ but $t(x) \neq t(y)$, then no replacement map can assign a unique target value to that source value.

For additive monoids, take raw pairs $\mathbb{N}^2$, $s(a, b) = a + b$, and $t(a, b) = (a, b)$. The pairs $(1, 0)$ and $(0, 1)$ obstruct replacement. Their sum is the same, while their provenance differs.

The joint completion $J_{s,t} = \operatorname{im}(s, t)$ is the coarsest quotient preserving both views. Its projection to the source view is invertible exactly when the factorization gate passes. This completion retains distinctions; it does not claim that every carrier embeds canonically into every other carrier.

Manuscript anchor: `thm:block18-carrier-factorization;thm:block18-universal-joint-completion`

21 The same route graph can use different carrier laws

An observer packet aggregates histories by a finite state set. Let B be a finite branch alphabet, $\delta(s,b)$ a transition, and $(A,\oplus,0)$ a commutative monoid. Each branch action L_b must preserve the monoid law.

$$a'_t = \bigoplus_{\delta(s,b)=t} L_b(a_s).$$

For finite raw-history families, advancing each history and then aggregating agrees with aggregating first and using this formula. Associativity, commutativity, and the branch homomorphism law justify the regrouping. Infinite-history sums require additional convergence data.

For the Haar example at $x=(1,0)$, reconstruction of the first coordinate has contributions $1/2, 1/2$, while the second has contributions $1/2, -1/2$. Signed addition returns $(1,0)$. Maximum returns $(1/2,1/2)$. The graph of routes did not change; its aggregation law did.

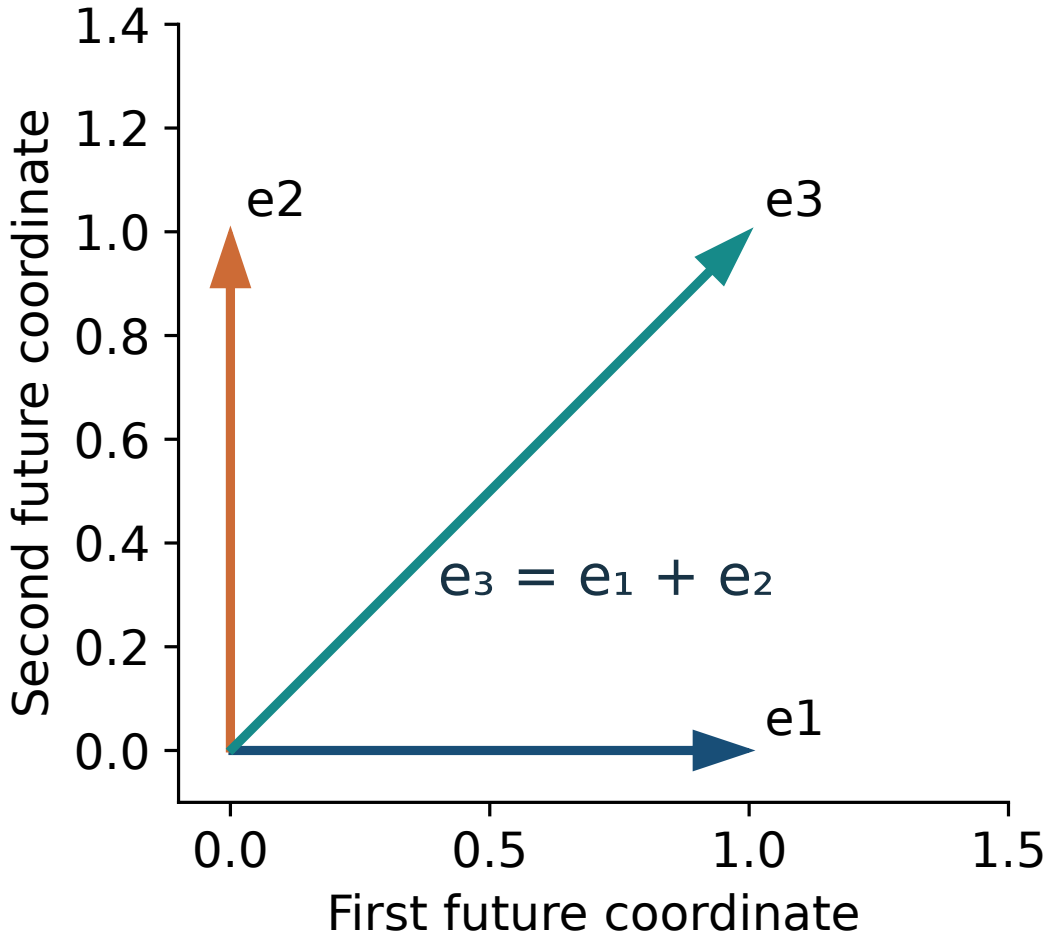
Two algebraic barriers make this precise. A nonzero ring cannot have a unital semiring homomorphism into a nontrivial zerosumfree semiring: the image of $1+(-1)=0$ would contradict zerosumfreeness. Conversely, an idempotent-additive nontrivial semiring cannot map unitally into a nontrivial additively cancellative semiring: the image of $1+1=1$ would force $1=0$.

These barriers apply to the specified homomorphism type. They do not forbid every task-restricted encoding or every enrichment preserving more information.

Manuscript anchor: `thm:block18-carrier-indexed-observer;ex:block18-haar-cancellation;thm:block18-cancellation-barrier;thm:block18-idempotence-barrier`

22 Observer change creates a relation object

Distinct future functions need not be linearly independent. A discrete observer retains equality classes of functions. A linear observer retains their span.



Three distinct vectors can generate a two-dimensional linear behavior space.

For $E = \{e_1, e_2, e_3\} \subseteq \mathbb{Q}^2$, where $e_1 = (1,0)$, $e_2 = (0,1)$, and $e_3 = (1,1)$, free linearization has three basis symbols. Evaluation is

$$\pi : \mathbb{Q}^{(E)} \to \operatorname{span} E, \qquad \pi = \begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 1 \end{pmatrix}.$$

Its relation space is $\mathbb{Q}(1,1,-1)$, so

$$0 \to \mathbb{Q}(1,1,-1) \to \mathbb{Q}^3 \xrightarrow{\pi} \mathbb{Q}^2 \to 0$$

is exact. The free symbols are distinct, but their evaluated functions satisfy a linear relation.

This is observer base change, not carrier replacement. The retained span can be smaller than the free vector space on the discrete residual states even when the carrier law is unchanged. The two minimizations commute exactly when the distinct residual functions are linearly independent.

Manuscript anchor: `thm:block18-observer-base-change;bnd:block17-no-absolute-minimum`

23 An infinite discrete quotient can have finite linear rank

Fix $N \geq 1$, $z \neq 0$, and the cyclic shift S_z on K^N, with wrap weight z. Put

$$M_N = I + S_z, \qquad f_N(n) = e_0^{\mathsf{T}} M_N^n e_0.$$

The $N \times N$ Hankel block $H_{ij} = f_N(i+j)$ has

$$\det H = \pm z^{N-1}.$$

Thus the linear residual dimension is exactly N. The lower bound comes from the nonzero Hankel minor; the upper bound comes from the displayed N-dimensional realization.

For $N = 2, z = 1$,

$$f_2(n) = 1, 1, 2, 4, 8, 16, \ldots, \qquad H = \begin{pmatrix} 1 & 1 \\ 1 & 2 \end{pmatrix}, \qquad \det H = 1.$$

The residual futures $k \mapsto f_2(n+k)$ are all distinct. Prefixes zero and one differ at suffix one, and prefixes from one onward are distinguished by their later exponentially growing values. The discrete exact future quotient is infinite even though the linear continuation dimension is two.

In general, for positive real z,

$$f_N(n) = \sum_{r=0}^{\lfloor n/N \rfloor} \binom{n}{Nr} z^r.$$

The source proves the required distinctions for every prefix. The guide verifies determinant identities on 24 exact finite trials, $1 \leq N \leq 6$ and $z \in \{1/2, 1, 2, -1\}$. Positivity is used for the discrete distinctness claim, not for the nonzero determinant claim.

Manuscript anchor: `thm:block18-polyphase-hankel-separation`

24 Value, witness, and history are different tasks

For convex value semantics, an observer state stores a convex body C_s. A finite branch update is

$$C'_t = \operatorname{conv} \bigcup_{\delta(s,b)=t} A_b^{\mathsf{T}} C_s.$$

With support function $h_C(x) = \sup_{c \in C} c^{\mathsf{T}} x$,

$$h_{C'_t}(x) = \max_{\delta(s,b)=t} h_{C_s}(A_b x).$$

Here the branch alphabet and stored atom families are finite. We use $\operatorname{conv}(\varnothing) = \varnothing$ and $h_\varnothing = -\infty$; the empty maximum is $-\infty$. An empty packet represents no available history.

Consider the square with vertices $(\pm 1, \pm 1)$ and an interior atom at the origin carrying a separate provenance label. The interior atom never maximizes a nonzero linear functional now. Under every invertible continuation, it remains interior in the transformed full-dimensional body and still cannot maximize.

If the zero continuation is admissible, every atom maps to the origin. The interior atom's provenance now ties with all others. Deleting it was harmless for value-only queries, but changes the all-witness task after that continuation.

A convex hull preserves support values. It does not by itself preserve which branch won, how many histories tied, or the full history. A pruning rule must therefore state both the continuation class and the query task. The singular example is a precise boundary on provenance deletion, not a failure of the convex value transport theorem.

Manuscript anchor: `thm:block18-convex-packet-transport;prop:block18-singular-resurrection;prop:block18-invertible-hiding;bnd:block18-provenance-boundary`

25 A behavioral bridge has two independent defects

Let $f : X \to Y$ be a linear bridge with behavioral kernels N_X, N_Y. It induces a map of task quotients only if

$$f(N_X) \subseteq N_Y.$$

This compatibility condition says an invisible source difference cannot become a visible target difference. Once it holds, the exact sequence is

$$0 \to \frac{f^{-1}(N_Y)}{N_X} \to \frac{X}{N_X} \xrightarrow{\bar{f}} \frac{Y}{N_Y} \to \frac{Y}{N_Y + f(X)} \to 0.$$

The left defect measures collapsed source distinctions. The right defect measures unattainable target semantic states. Either one can vanish without the other.

With fully visible tasks, $N_X = N_Y = 0$, the projection $\mathbb{Q}^2 \to \mathbb{Q}$, $(x, y) \mapsto x$, is surjective but loses the line $\mathbb{Q}e_2$. The inclusion $\mathbb{Q} \to \mathbb{Q}^2$, $x \mapsto (x, 0)$, is injective but misses the target quotient represented by the second coordinate. An isomorphism requires both defects to vanish.

Under composition, loss fits into

$$0 \to \ker \bar{f} \to \ker(\bar{g}\bar{f}) \to \ker \bar{g} \cap \operatorname{im} \bar{f} \to 0.$$

The intersection matters. A target distinction killed by $\bar{g}$ contributes to composite loss only if the first bridge can reach it. One cannot add arbitrary stagewise kernel dimensions without inspecting this image condition.

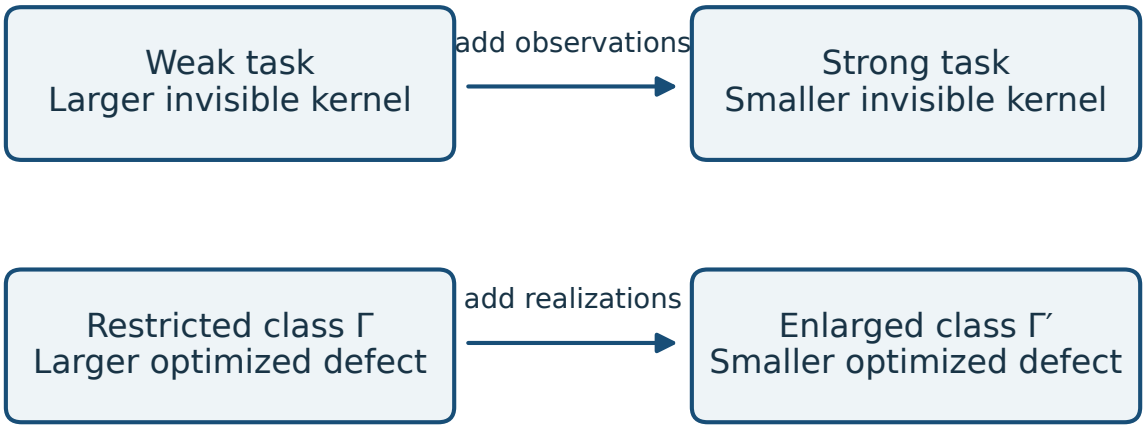
Manuscript anchor: `def:block18-task-compatible-bridge;thm:block18-bridge-loss-exact;thm:block18-loss-composition`

26 Admissibility change is a different order

For an admissible class Γ and nonnegative defect $d(X, g)$, define

$$\delta_\Gamma(X) = \inf_{g \in \Gamma} d(X, g).$$

Enlarging the admissible class can only decrease this optimized defect. It does not ask the system to preserve more observations.



Task strengthening and admissibility enlargement act on different mathematical data.

The class-relative defect profile includes attainment. There are three cases:

- Completion: zero defect, attained by an admissible realization.
- Gap: strictly positive optimized defect.
- Escape: zero infimum, with no attaining realization.

The distinction between gap and escape is essential. An approximation sequence can make the defect arbitrarily small while its conditioning diverges. That is not an exact completion in the declared class.

The next pages show two complementary examples. Changing from quadratic to general convex geometry removes a positive optimal-bound gap. Changing a quadratic metric along a degenerating sequence approaches a zero defect without reaching an invariant positive-definite metric.

Manuscript anchor: `thm:block19-admissibility-monotonicity;def:block19-trichotomy`

27 Convex and quadratic geometry have different bounds

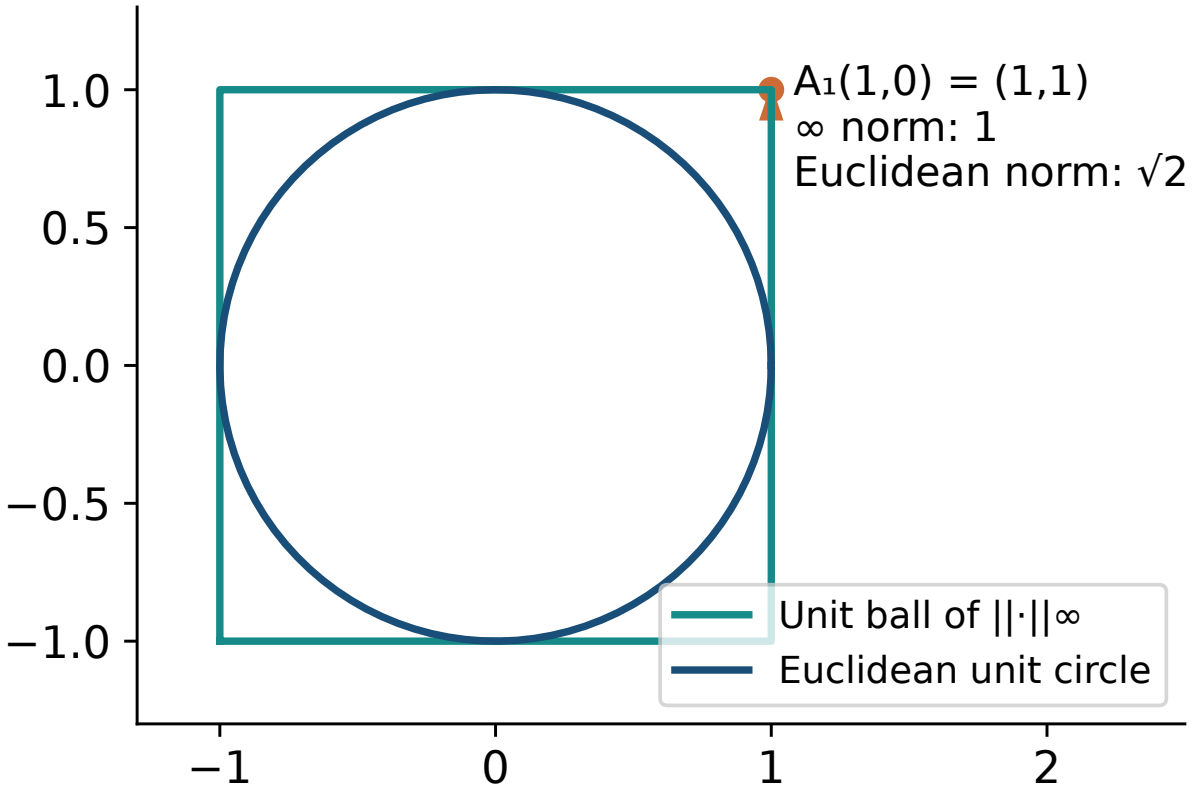
Consider

$$A_1 = \begin{pmatrix} 1 & 0 \\ 1 & 0 \end{pmatrix}, \qquad A_2 = \begin{pmatrix} 0 & 1 \\ 0 & -1 \end{pmatrix}.$$

They satisfy

$$\max(\|A_1 x\|_\infty, \|A_2 x\|_\infty) = \|x\|_\infty.$$

The exact best convex/Finsler bound is one. Every positive-definite quadratic metric instead has a common operator bound at least $\sqrt{2}$.



The image of $(1,0)$ under A_1 lies on the square boundary but outside the Euclidean unit circle.

For $Q = \left(\begin{smallmatrix} a & c \\ c & b \end{smallmatrix}\right)$, $d = ab - c^2 > 0$, the squared induced norms are

$$u = \frac{b(a+b+2c)}{d}, \qquad v = \frac{a(a+b-2c)}{d}.$$

The exact identity

$$ab((a+b)^2 - 4c^2) - 4d^2 = ab(a-b)^2 + 4c^2 d \geq 0$$

gives $uv \geq 4$, hence $\max(\sqrt{u}, \sqrt{v}) \geq \sqrt{2}$. Equality holds for scalar Euclidean metrics. This is a geometry-class separation, not a claim that no useful quadratic approximation exists.

Manuscript anchor: `thm:block19-convex-quadratic-separation`

28 Zero infimum need not be attained

For the unipotent matrix

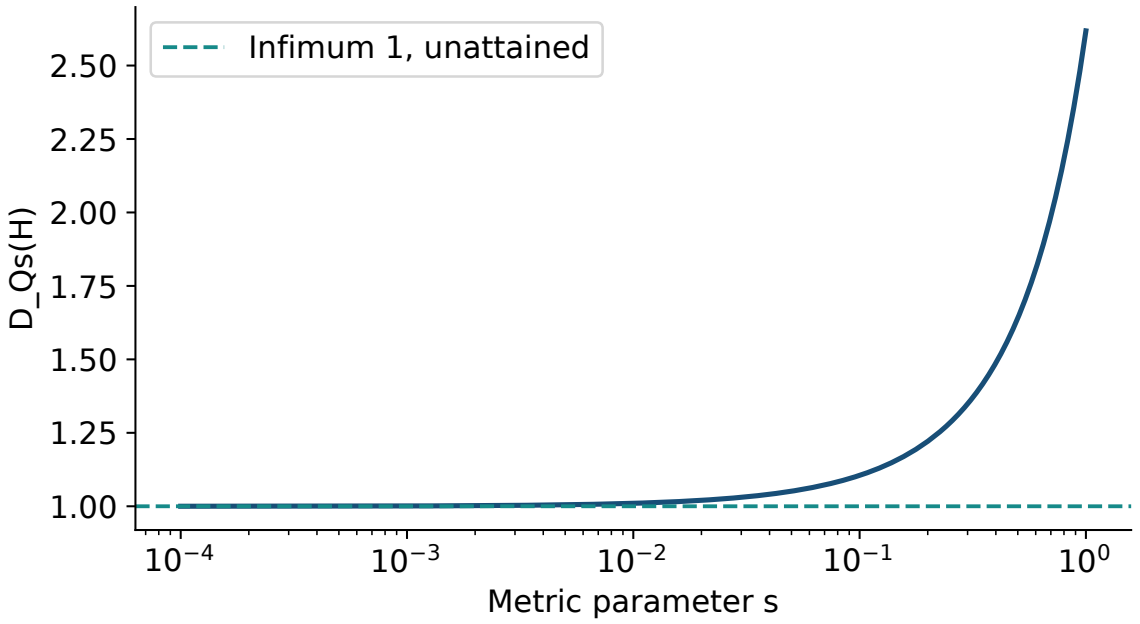
$$H = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix},$$

there is no positive-definite Q with $H^{\mathsf{T}} Q H = Q$. Indeed, the equation forces the upper-left entry of Q to vanish.

For $0 < s \leq 1$, however, choose $Q_s = \operatorname{diag}(s^2, 1)$. The metric-normalized map is

$$\begin{pmatrix} s & 0 \\ 0 & 1 \end{pmatrix} H \begin{pmatrix} s^{-1} & 0 \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} 1 & s \\ 0 & 1 \end{pmatrix}.$$

Its singular-value condition number tends to one, while $\kappa_2(Q_s) = s^{-2}$ diverges. With defect $d = \log D_Q(H)$, the infimum is zero but unattained.



Floating-point illustration of the proved limit. The matrices and the nonattainment argument are exact.

There is a clean compactness certificate. If a uniformly conditioned sublevel is compact and d is continuous, then it attains a minimum. If no zero exists, that minimum is positive. A sequence with $d \to 0$ must therefore eventually leave every such conditioned sublevel. This argument uses the attained minimum on the compact set; it does not assume every compact Hausdorff space is sequentially compact.

Manuscript anchor: `prop:block19-unipotent-escape;prop:block19-escape-compactness`

29 Metric completion is a holonomy gate

Let a finite connected graph carry invertible transports between positive-dimensional real or complex fibers of common dimension n. Reverse traversal uses inverse transport. A compatible conformal metric packet satisfies

$$T_e^* Q_{t(e)} T_e = \lambda_e^2 Q_{s(e)}, \qquad \lambda_e > 0.$$

Normalize a root holonomy by

$$\widehat{H} = |\det H|^{-1/n} H.$$

The positive-dimensional assumption is required for this expression. Compatible vertex metrics exist precisely when the normalized holonomy representation preserves a positive-definite root metric. In finite dimension, this is equivalent to unitarizability, or relative compactness of its image.

The swap holonomy from the earlier pages already preserves the Euclidean metric. The unipotent example has determinant one, so normalization leaves it unchanged. Its powers are unbounded and there is no invariant positive-definite metric. The degenerating Q_s sequence does not contradict the criterion because no individual member realizes exact invariance.

For an invertible map, intrinsic metric distortion is the singular-value ratio of $R_t T R_s^{-1}$, where $Q_s = R_s^* R_s$ and $Q_t = R_t^* R_t$. Simultaneously changing map and metric coordinates leaves it invariant. Along a composite, distortions multiply submultiplicatively; their logarithms add subadditively.

This gate concerns metric enrichment. The existence of an underlying algebraic transport or a nonzero global section alone does not supply the compatible positive-definite metric.

Manuscript anchor: `prop:block19-distortion-cocycle;thm:block19-metric-completion`

30 Forgetting structure and rebuilding it differ

Suppose richer categories have forgetful functors

$$U : \widetilde{\mathcal{C}} \to \mathcal{C}, \qquad V : \widetilde{\mathcal{D}} \to \mathcal{D}.$$

A rich bridge $\widetilde{B}$ automatically has an underlying bridge $V\widetilde{B}$. Given a bare bridge B, constructing $\widetilde{B}$ and a comparison

$$V\widetilde{B} \cong BU$$

is an additional lifting problem.

Underlying data	Additional requested structure
Invertible graph transports	Compatible positive-definite metric packet
Additive realization	Exact strong monoidal fiber realization
Projective semilinear data	Linear semilinear descent
Value-only quotient	Witness or full-history semantics
Bare quotient coordinates	Compatible topology, filtration, or metric

The corresponding gates can involve holonomy, tensor comparisons, a relative Brauer class, relation objects, or conditioning. The theory records the common shape of the lifting problem; it does not identify all these obstructions as one universal scalar defect.

Existence, uniqueness, and residual automorphisms belong in the lift passport. A construction may exist without being canonical, or may exist only on a restricted subcategory. A source proof about an underlying object does not establish those extra fields by itself.

This is the reason the final synthesis keeps a family of doctrine-indexed task fibers. Exact quotient calculus works within a fiber; crossing to another doctrine needs a verified bridge and, when richer structure is requested, an enhancement certificate.

Manuscript anchor: `prop:block19-forgetful-lifting;def:block19-task-fiber;bnd:block19-second-order-scope`

31 Finite atlas assembly needs effective gluing

Local chart classifications do not automatically classify global objects. A finite profiled atlas supplies chart categories, typed overlaps, coherent reversible comparisons, and semantic realizations. Its descent category consists of compatible families of local objects and compatible local morphisms.

The finite pseudolimit construction gives componentwise identities and composition. Invertibility is componentwise, so the maximal subgroupoid of descent agrees with descent of the local cores. A pseudonatural directional bridge induces a global functor; it becomes an equivalence only with the corresponding equivalence hypotheses.

For invariant completeness, three obligations must be kept together:

1. Each invariant is complete on its normalized local category.
2. The invariants and chart comparisons are coherent on overlaps.
3. The local invariant-compatible isomorphisms have effective descent.

The last condition is the gluing gate. Equal local invariant values can supply local isomorphisms without ensuring that those isomorphisms assemble globally. Holonomy, cokernel, Brauer, or higher-coherence defects may obstruct that assembly in their respective settings.

The manuscript verifies a selected finite certified atlas and separates its reversible connected components from directional bridges. It does not infer an unrestricted heterogeneous stack or one global coordinate system from the finite example.

The useful publication claim is exact: finite profiled assembly is proved under its stated coherence, preservation, and effectivity hypotheses, and the declared finite atlas supplies a concrete realization. Larger atlas classes remain separate obligations.

Manuscript anchor: `thm:atlas-assembly;thm:block15-lane-descent;thm:block19-relative-semantic-synthesis`

32 Example ledger and reproducibility

The source package contains `generate_examples_and_figures.py`, 14 vector figures with PNG companions, and `validation/visual_example_checks.json`. The example checker records 137 passing exact rational assertions. The metric curve alone uses floating point for illustration.

Example	Exact computation recorded
Oblique packets	Idempotence, annihilation, completeness, inverse coordinates
Transfer and residuals	Lane conjugation, contraction, selector laws, all three residual terms
Refinement and fold	Nine retained entries; exact collision-kernel witness
Holonomy	Swap equivariance, fixed-core rank, sign-twisted zero selector
Behavioral quotient	Complete partition trace; linear future-observation rank
Semantic diamond	Rank, compatibility map, zero composite and witnesses
Observer change	Relation vector and evaluation rank
Polyphase Hankel family	24 rational determinant trials, $1 \leq N \leq 6$
Geometry separation	75 exact polynomial-identity trials

The ledger's individual records carry the manuscript label that justifies each example. A finite-grid check is not a substitute for the theorem's universal proof. The guide does not claim Lean certification for its newly written explanatory examples.

To reproduce the figures, use a Python environment with NumPy and matplotlib. The exact algebra layer uses the standard-library `Fraction` class and custom rational row reduction. No SymPy dependency is required. Then compile the supplied TeX three times with pdfLaTeX so that the two-page contents and page references stabilize. The full Markdown guide preserves the same text, equations, and illustrations for readers and machines.

The mathematical boundaries throughout the guide are intentional: finite profiles, explicit quotient existence, finite branching, carrier homomorphism laws, task compatibility, and effective descent are part of the claims being explained.

Manuscript anchor: `sec:block15-universal-lane-transport;sec:block16-holonomy-twisted-lane-transport;sec:block17-semantic-quotients;sec:block18-carrier-observer;sec:block19-admissibility-lifting`