REALIZATION CALCULUS

VOLUME I

Komorebi Space

Coherent Realizations, Obstructions, and Descent

$$\mathfrak{K}(F) = \Gamma(F)$$

Declared views. Coherent families.
Obstruction, descent, and quantitative repair.

Komorebi P.

The two volumes

REALIZATION CALCULUS

These volumes study realization under an explicit contract: which data are available, which transformations are allowed, and which objects count as answers. A compatible description, an admissible realization, and a recoverable source are different mathematical achievements. Their relation has to be proved.

Volume I, Komorebi Space, constructs spaces of coherent families from declared systems of views. It develops extension obstructions, symmetry, descent, transport, and quantitative repair. Its analytic layers distinguish exact compatibility from membership in an intended normed realization, and a small residual from a small correction. Each conclusion retains the hypotheses that make it valid.

Volume II, Realization Boundaries, examines the information that observations and completion can preserve or discard. A concrete family of countable abelian groups has one common profinite completion and one common unmarked finite-quotient system, while its abstract group types and effective presentations retain an infinite parameter. Subgroup classification, homomorphism rigidity, root structure, computability, and grammar escape measure that difference precisely.

The point of contact is the semantic comparison in Volume I: defining a category of coherent data does not establish that a functor from intended objects is faithful, full, or essentially surjective. Volume II supplies explicit categorical and computational tests in one group theoretic setting. Its examples preserve the universal property of the ambient limit. They show that the limit alone need not recover the chosen source image.

The volumes can be read independently. Read them together to move from the construction of coherent answers to the classification of what those answers determine. A finite resolution level in Volume I need not have a finite value set. The finite quotients of Volume II are finite groups. Likewise, escape from a fixed countable grammar is a relative theorem about that grammar and its semantics, rather than an assertion that the groups escape every possible system of views.

Komorebi P.

A route through the pair

The following connections give a short route across the two books. Section and chapter numbers within each volume retain their original mathematical numbering.

Question	Volume I: Komorebi Space	Volume II: Realization Boundaries
What is coherent?	Views and coherent realizations. Compatible evaluations determine a family.	Chapters 2 and 9. The prime-coordinate product and the intended dense source.
What do observations retain?	Finite observations and realized transformations. The comparison data specify what survives.	Chapters 3 and 9. Membership observations retain a spectrum; unmarked finite quotients forget the selector.
Which ambient points are intended?	Coherent points and Hilbert realizations. Compatibility and range membership are separate.	Chapters 4 and 9. Pure intermediate groups share an ambient product without becoming the same source.
Which maps lift?	Internal realization and intended objects. Faithfulness, fullness, and essential surjectivity are separate questions.	Chapter 9. Faithful completion maps, explicit non-lifts, and qualified reflection results.
What can a procedure supply?	Certificates and implementation boundaries. An existence proof and an executable repair need separate evidence.	Chapters 8 and 10. Oracle presentation semantics and escape from a fixed countable grammar.

Formal companions

Volume I carries its accepted Komorebi Space development: 162 production modules and four entry-points, with formal counterparts for all sixteen principal theorem groups. This does not assign a separate Lean theorem to every illustration or printed proof route. Volume II carries its accepted Realization Boundaries development: 66 production modules and all forty-five registered contracts complete. Each volume's formal appendix explains its own scope.

This paired edition adds common presentation and reading connections. It preserves the mathematical bodies and the separately accepted proof sources. Its editorial checks establish source identity and publication quality. They do not constitute a new joint Lean replay. The source package records the accepted receipts, dependency trust boundaries, and the precise editorial changes.

Preface

A mathematical object can appear in several forms: as local data, as a sequence of finite observations, as a family of coordinates, or as the output of an approximate calculation. These descriptions may agree, fail to agree, or agree only after a correction. This monograph studies the conditions under which they describe one coherent realization.

The starting point is a declared system of views and maps between them. Its coherent families form a limit. Prescribed data then define an extension problem, whose obstruction records whether a completion exists and whose ambiguity records the freedom in choosing it. Resolution towers make these questions sensitive to scale. Groupoids retain the symmetries of a realization, while descent asks whether compatible local realizations assemble. In normed and Hilbert models, the same questions acquire quantitative forms: how much correction is needed, which information can be preserved, and whether small residuals certify proximity to an exact answer.

These constructions give a common language for realization, obstruction, and repair. Their connections are made through explicit maps and estimates. Existence, uniqueness, and stability remain separate assertions. A change of representation must identify both the transformed object and the data it preserves.

The underlying limit constructions, gauge groupoids, stack definitions, orthogonal projections, and bounded inverse principles are part of established mathematics. The contribution of this presentation is their integration with specified view systems, protected correction spaces, resolution differentials, graph transports, and worked examples. The comparison maps and hypotheses carry the substance of that integration. Formal verification establishes the recorded declarations, without making a claim of historical novelty for each ingredient.

> The word *space* does not assert one universal topological object containing all of mathematics. The strict construction is a set of coherent families. Symmetry is retained in a groupoid, and effective descent is proved for a specified site. Analytic estimates require a declared normed or Hilbert model. Each layer has its own hypotheses and its own notion of realization.

The book presents proofs together with examples that show why their hypotheses matter. Its companion visual guide emphasizes geometry and explicit calculations. Its Lean guide explains the formal definitions, interfaces, and evidence. Mathematical figures are reproducible vector diagrams or plots of explicit formulas.

Reading guide

Part I constructs coherent sections and follows their obstructions as more data are prescribed. Part II develops resolution towers, finite observations, split repair coordinates, and multiscale reconstruction. Part III retains symmetries, proves effective descent, and distinguishes coherent phases from removable ones. Part IV develops analytic stability and a geometry of transformation costs. Part V follows approximate answers to their limits, separates a coherent point from its intended Hilbert realization, and derives budget, regularization, and stopping criteria. Part VI works through finite loop models and returns to mathematical semantics.

A short technical route is Sections 1, 3, 13, 18, and 21. For packets and residuals, include Sections 8 and 10. The graph models in Section 17 show how a categorical realization problem acquires a numerical condition number. Sections 7 and 9 explain how the constructions behave under changes of representation.

For concrete motivations, begin with Sections 33, 34, and 26. The last connects finite realization, closure, and failure of protected infinite existence in one calculation. Sections 27 through 30 develop its positive converse and regularized solutions. Sections 31 and 32 explain how to certify an inexact calculation and select its tolerance.

Appendix A contains fifteen solved problems. Appendix B collects the assumptions behind the principal results. Appendix C supplies notation and a map into the formal source. Appendix D describes formal coverage, and Appendix E proves the analytic background used in the text.

Proofs and formal scope

The accompanying Lean development covers the foundational constructions and all 16 principal theorem groups on bounded realization, budgets, regularization, and certificates. This includes the actual infinite diagonal family and its tail bounds, the sharp asymptotic over all real budgets, the rounded noisy and inexact rates, and the concrete nonclosed-range example. Appendix D identifies the formal counterparts and their scope.

A formal counterpart can establish a theorem by a different route from its printed proof. In particular, the bounded-realization result is formalized through a compactness argument without assigning a separate formal counterpart to the printed weak-subsequence lemma. Explanatory prose, background proofs, and numerical illustrations are not all individually machine checked.

The complete source has passed a fresh integrated compilation, declaration audit, and diagnostic-control replay. The earlier records remain available as historical evidence. Lean's compiler and kernel, its stated standard axioms, and the pinned external dependency cache remain part of the trust boundary. The cache is recorded but was not rebuilt by this audit. Appendix D gives the exact counts and their interpretation.

Contents

I Realizations and their obstructions 1
1 Views and coherent realizations 1
2 Partial data, gluing, and correspondence fibers 3
3 The additive obstruction and the space of extensions 6
4 Protected correction as a general principle 9
5 Changing what is prescribed 11

II Resolution, repair, and retained detail 15
6 Resolution towers and exact transport 15
7 Finite observations and realized transformations 17
8 Banach repair and constructed split coordinates 23
9 Repair under a change of representation 28
10 Multiscale residuals and the limits of reconstruction 33

III Symmetry and effective descent 38
11 Nonabelian coherence and relative trivialization 38
12 Effective gluing of objects and their identifications 41
13 A genuine stack on the composition site 44
14 Equivalent presentations and identified boundaries 47
15 A boundary identification is data 50
16 Coherence need not mean removable phase 52

IV The geometry of exact realization 55
17 Graph views, holonomy, and reconstruction 55
18 Closed range is the stability boundary 58
19 Holonomy repair and transport amplification 61
20 Quantitative descent in a common Hilbert model 64
21 Controlled motion of exact spaces 67
22 Reading a stability certificate 72
23 A geometry of transformations needs its own cost 76

V From approximation to exact realization 79
24 Exactification without a fixed contraction rate 79
25 The residual floor and the closure of admissible changes 86
26 Exact finite repairs can escape the intended ambient space 89
27 A bounded repair budget can recover existence 93
28 The attainable residual at a fixed correction budget 97
29 Regularization selects a protected correction 102
30 A complete regularization diagnostic for the escaping model 108
31 Certifying an inexact protected correction 113
32 Coordinates, noise, and a justified stopping rule 118

VI Worked realization problems 123
33 A closed signal loop 123
34 A loop with a nontrivial fixed space 126
35 Mathematical meaning and the remaining boundaries 128
A A reader's workshop 134
B A hypothesis atlas 138
C Notation and formal correspondence 140
D Proof and formal verification 145
E Analytic facts used in the proofs 150
F Background and references 155

Part I

Realizations and their obstructions

A global object is not added to local data by terminology. It is a family satisfying every prescribed transition equation. The first task is to construct the space of such families and to keep its possible emptiness visible.

1 Views and coherent realizations

1.1 The data of a view system

Fix a small category $\mathcal{V}$, or work in fixed universes so that the products below exist in the chosen ambient universe. A *view* is an object of $\mathcal{V}$. It can encode a resolution level, a coordinate presentation, or a context in which part of an object is visible. These interpretations are mathematical choices, not additional axioms.

Definition 1.1 (View system)**.** A view system consists of sets F_v and maps

$$f^* : F_w \longrightarrow F_v \qquad (f : v \longrightarrow w)$$

satisfying $(1_v)^* = \mathrm{id}$ and $(g \circ f)^* = f^* g^*$. Equivalently, it is a functor $F : \mathcal{V}^{\mathrm{op}} \to \mathbf{Set}$.

The direction matters. An arrow $v \to w$ induces a restriction from the fiber at w to the fiber at v. A restriction need not be invertible. A category with several arrows between two views can retain different routes between their descriptions. Replacing it by a preorder can lose those routes and the constraints they impose.

Definition 1.2 (Strict Komorebi realization)**.** The ambient realization set is

$$\mathfrak{K}(F) = \Gamma(F) = \left\{ x \in \prod_{v \in \mathrm{Ob}\,\mathcal{V}} F_v : f^*(x_w) = x_v \text{ for every } f : v \to w \right\}. \tag{1.1}$$

A point of $\mathfrak{K}(F)$ is a whole coherent family. Its value in view v is the evaluation $\mathrm{ev}_v(x) = x_v$.

The raw product contains all families, including inconsistent ones. Equation (1.1) cuts out the exact families. No nonemptiness hypothesis is built into the definition.

Theorem 1.3 (Joint determination and the universal property)**.** *The evaluations jointly determine a realization. If a set T carries maps $c_v : T \to F_v$ satisfying $f^* c_w = c_v$, there is a unique map $c : T \to \Gamma(F)$ satisfying $\mathrm{ev}_v\, c = c_v$ for every view.*

Proof. We first verify that the evaluation maps themselves form a compatible family. If $x \in \Gamma(F)$ and $f : v \to w$, the defining equation for x gives

$$f^*(\mathrm{ev}_w\, x) = f^*(x_w) = x_v = \mathrm{ev}_v\, x.$$

If $x, y \in \Gamma(F)$ have equal evaluations, then $x_v = y_v$ for every v. Equality of functions into a dependent product gives equality of their underlying families. Membership in $\Gamma(F)$ is a property of a family, not a further choice of coordinate. Thus $x = y$.

Now fix the compatible maps $c_v : T \to F_v$. For each $t \in T$ define a family by $x(t)_v = c_v(t)$. For every arrow,

$$f^* x(t)_w = f^* c_w(t) = c_v(t) = x(t)_v.$$

This calculation verifies membership in $\Gamma(F)$, so the rule $c(t) = x(t)$ defines an actual function $c : T \to \Gamma(F)$. Evaluating it gives $\mathrm{ev}_v\, c(t) = c_v(t)$, as required.

To prove uniqueness, let $b : T \to \Gamma(F)$ be another such function. At every t and v,

$$(b(t))_v = \mathrm{ev}_v\, b(t) = c_v(t) = \mathrm{ev}_v\, c(t).$$

Joint determination first yields $b(t) = c(t)$, and function extensionality then yields $b = c$. Nothing in this argument assumes that T or $\Gamma(F)$ is inhabited. For example, when T is empty there is the unique empty function. When T is inhabited, existence of the compatible family (c_v) is already substantive data. □

Thus the strict construction is precisely the ordinary limit of the view diagram, with the usual limit interpretation in the sequential case [1]. This identifies its mathematical content and its universal property. It does not supply a universal view diagram for every possible mathematical question.

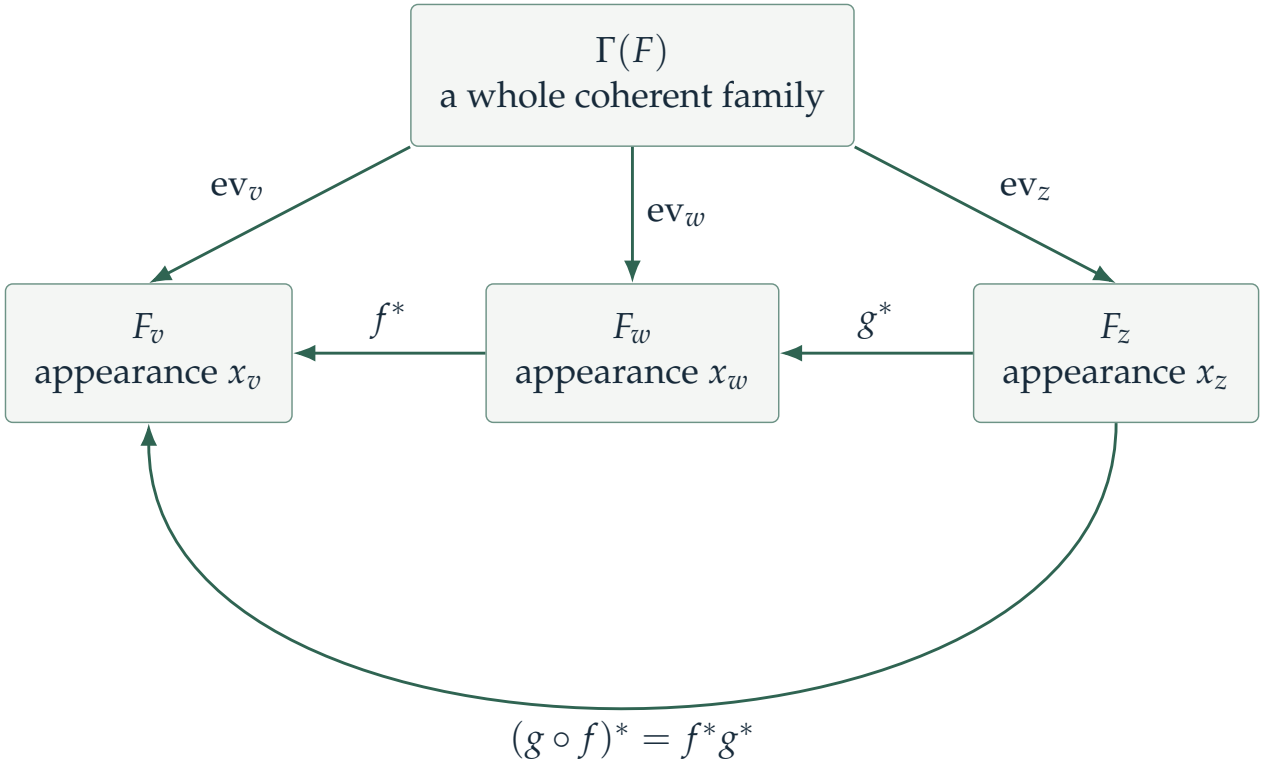


Figure 1: The same realization is evaluated in several views. Compatibility is imposed along the arrows, not inferred from the fact that all displayed fibers exist. The ambient family belongs to the simultaneous solution set of the transition equations.

1.2 Maps of realizations

Let $\alpha : F \to G$ consist of maps $\alpha_v : F_v \to G_v$ with

$$\alpha_v f_F^* = f_G^* \alpha_w.$$

The induced map on realizations is $\Gamma(\alpha)(x)_v = \alpha_v(x_v)$. The naturality square is exactly what is needed to show that the image family is coherent.

Proposition 1.4 (Functorial transport)**.** *These induced maps preserve identities and composition:*

$$\Gamma(\mathrm{id}_F) = \mathrm{id}_{\Gamma(F)}, \qquad \Gamma(\beta\alpha) = \Gamma(\beta)\Gamma(\alpha).$$

Proof. Evaluate each side at an arbitrary view and apply the corresponding identity for the coordinate maps. Theorem 1.3 then identifies the maps on the whole realization set. □

An arbitrary collection of fiber maps need not define a map on realizations. Its compatibility equations are part of the theorem, not a convention that can be omitted after a notation has been chosen.

1.3 A global view and a useful warning

Proposition 1.5 (Terminal evaluation)**.** *If t is terminal in $\mathcal{V}$, evaluation induces an equivalence $\Gamma(F) \simeq F_t$.*

Proof. For $a \in F_t$, take its restriction along the unique arrow $v \to t$ at each view. Uniqueness of the arrow to t, together with the functor laws, proves coherence. Evaluation at t recovers a. Conversely, any coherent family is the restriction of its coordinate at t. □

This proposition is both useful and diagnostic. If the supposedly reconstructed global object has already been placed in a terminal fiber, the limit recovers that fiber. The proposition does not prove that arbitrary partial data have a preimage there.

The category of elements $\int_{\mathcal{V}} F$ and $\Gamma(F)$ also have different roles. The former organizes individual appearances and their transitions. The latter selects coherent families across all views. The category of elements serves here as explanatory notation, while the Lean formalization directly represents families and their coherence.

2 Partial data, gluing, and correspondence fibers

2.1 Prescribed views

For $U \subseteq \mathrm{Ob}\,\mathcal{V}$, use the full subcategory on U. A partial section $s \in \Gamma_U(F)$ satisfies all transition equations whose endpoints lie in U. Define

$$\mathrm{Ext}_U(s) = \{x \in \Gamma(F) : x|_U = s\}. \tag{2.1}$$

This fiber can be empty, a singleton, or much larger. A partial section is already coherent on its prescribed domain. A raw assignment that fails an internal transition equation is not a partial section of this kind.

Let $(U_i)_{i\in I}$ be full patches. A matching family consists of $s_i \in \Gamma_{U_i}(F)$ with equal values on every overlap. The relevant covering condition is

$$\text{for every } f : v \to w, \quad \text{some } U_i \text{ contains both } v \text{ and } w. \tag{2.2}$$

Identities imply that this condition also covers every object.

Theorem 2.1 (Strict gluing)**.** *Under (2.2), restriction induces an equivalence*

$$\Gamma(F) \simeq \mathrm{Match}(F, (U_i)).$$

Every fixed matching family has a unique global glue.

Proof. Step 1: every coordinate can be selected. Apply the cover condition to the identity arrow of v. It gives an index containing v. Choose one such index $i(v)$ and put

$$x_v = (s_{i(v)})_v.$$

The choice is an existence construction. It is not an algorithm selecting patches in a preferred order. If $v \in U_j$, matching on the overlap $U_{i(v)} \cap U_j$ gives

$$x_v = (s_j)_v. \tag{*}$$

In particular the value is independent of the selected index.

Step 2: the selected coordinates satisfy every arrow equation. For $f : v \to w$, choose a single index j with $v, w \in U_j$. Formula (*) applies to both endpoints. Since the patch is full, it contains this actual arrow, even when other arrows have the same endpoints. Coherence of s_j gives

$$f^* x_w = f^*(s_j)_w = (s_j)_v = x_v.$$

Thus $x \in \Gamma(F)$. Object coverage alone would not supply the common patch used in this step.

Step 3: assembly and restriction are inverse. Formula (*) proves $x|_{U_j} = s_j$ for every j. Conversely, start with $y \in \Gamma(F)$ and assemble its restrictions. The resulting coordinate at v is y_v, regardless of which containing patch is chosen. Joint determination gives the original y. These two identities are exactly the inverse identities for the displayed equivalence.

For uniqueness with a fixed matching family, any other glue has coordinate $(s_{i(v)})_v = x_v$ at every v, hence equals x. If the view category is empty, its product has one empty family and the same conclusion follows without making any choices. □

The proof explains why object coverage alone is insufficient. Coherence is an equation attached to an arrow, so existence of a glue must see that equation somewhere in the local data.

Example 2.2 (An object cover with no glue)**.** Take two views $0 \to 1$, Boolean fibers, and the identity restriction. The singleton patches cover all objects. A local value false at 0 and true at 1 agrees on the empty overlap, yet no global section realizes both values. A global family would have to assign the same Boolean to both views.

Object coverage does prove uniqueness of a possible glue. It does not prove its existence. This distinction persists when sets are replaced by objects with symmetries.

2.2 Correspondences must also be coherent

Suppose natural maps form a span $F \xleftarrow{p} M \xrightarrow{q} G$, and fix $a \in \Gamma(F)$, $b \in \Gamma(G)$. At each view define the ordinary fiber

$$B_{a,b}(v) = \{m \in M_v : p_v(m) = a_v,\ q_v(m) = b_v\}.$$

Naturality and coherence provide its restrictions. A global bridge is an element of

$$\mathfrak{B}(a, b) = \{z \in \Gamma(M) : \Gamma(p)z = a,\ \Gamma(q)z = b\}.$$

Theorem 2.3 (Fiber and section compatibility)**.** *There is a coordinate preserving equivalence*

$$\Gamma(B_{a,b}) \simeq \mathfrak{B}(a, b).$$

Under a transition cover, matching local bridge sections glue to a unique global bridge with both prescribed projections unchanged.

Proof. We must first check that the displayed fibers really form a view system. Let $m \in B_{a,b}(w)$ and $f : v \to w$. Naturality of p and coherence of a give

$$p_v(f_M^* m) = f_F^* p_w(m) = f_F^* a_w = a_v.$$

The same calculation with q, b gives $q_v(f_M^* m) = b_v$. Therefore $f_M^* m$ belongs to $B_{a,b}(v)$. Identity and composition of these restricted maps are inherited from M, because equality in a subset is equality of its underlying elements.

Given $u \in \Gamma(B_{a,b})$, forget each fiber-membership condition and call the underlying family $z_v \in M_v$. Coherence of u is precisely coherence of z, so $z \in \Gamma(M)$. The equations $p_v z_v = a_v$ hold at every coordinate. Joint determination in $\Gamma(F)$ yields $\Gamma(p)z = a$, and the equations for q similarly yield $\Gamma(q)z = b$. Thus z is a global bridge.

In the opposite direction, let $z \in \mathfrak{B}(a, b)$. Evaluate its two global projection equations at v. They show that z_v belongs to $B_{a,b}(v)$, and its original coherence shows that these elements form a section of the fiber system. Forgetting and reattaching membership conditions preserves every coordinate, so the two constructions are inverse.

Finally apply strict gluing to this specific fiber system. The glued section retains its membership in every $B_{a,b}(v)$, which is why both prescribed projections survive gluing. Uniqueness is for the fixed matching bridge family. Different matching families for the same (a, b) can still yield different global bridges. □

These are ordinary fibers. They do not retain comparison isomorphisms or automorphisms. Section 14 constructs the appropriate groupoid version when this information matters.

2.3 Locally inhabited does not mean globally inhabited

Example 2.4 (The loop with no coherent choice)**.** Let the view category have one object and automorphism group C_2. Let its nonidentity arrow act on a set with two elements by swapping the elements. The fiber is nonempty and every transport is bijective. Nevertheless a global section would be a fixed point of the swap, and there is none.

The problem is not missing local elements or noninvertible transition maps. It is the absence of an element satisfying the loop equation. Inserting a point in each fiber does not certify a coherent selection of those points.

> The strict theory already distinguishes three levels: possible local appearances, compatible partial descriptions, and an actual global realization. A correspondence must pass through the same levels. Local nonemptiness of its fibers is not a globalization theorem.

3 The additive obstruction and the space of extensions

The strict definition gives a system of equations. In an additive setting, that system has an actual differential, and the failure to extend prescribed data has a quotient valued obstruction.

3.1 The compatibility differential

Assume each fiber A_v is an abelian group and each restriction is additive. Set

$$E = \prod_v A_v, \qquad R = \prod_{f:v\to w} A_v, \qquad (\delta x)_f = f^*(x_w) - x_v.$$

The group R retains an entry for every arrow. It is not silently identified with a cohomology group.

Proposition 3.1 (Differential identities)**.** *The map $\delta : E \to R$ is additive, $\Gamma(A) = \ker \delta$, and*

$$(\delta x)_{1_v} = 0, \qquad (\delta x)_{g\circ f} = f^*((\delta x)_g) + (\delta x)_f.$$

Proof. Additivity follows coordinate by coordinate. Vanishing is exactly the defining coherence equation. For composition, add and subtract f^*x_w:

$$f^*g^*x_z - x_v = f^*(g^*x_z - x_w) + (f^*x_w - x_v).$$

The identity equation uses $(1_v)^* = \mathrm{id}$. □

This exact formula is more informative than a scalar error score. It specifies how defects transform and combine. Norms can be added later, but they do not replace the typed residual.

3.2 The admissible corrections

Fix $U \subseteq \mathrm{Ob}\,\mathcal{V}$. Write $p_U : E \to \prod_{v\in U} A_v$ for raw restriction. A correction preserving the prescribed values belongs to $\ker p_U$. Its possible defects form the subgroup

$$B_U = \delta(\ker p_U) \leq R.$$

For a partial section s, choose any raw extension $\widetilde{s}$, for example zero off U, and define

$$Q_U = R/B_U, \qquad \mathfrak{o}_U(s) = [\delta\widetilde{s}] \in Q_U. \tag{3.1}$$

If another raw extension is chosen, their difference vanishes on U, so its differential belongs to B_U. Thus the class is independent of the lift.

Theorem 3.2 (Extension criterion)**.** *For every coherent partial section,*

$$\mathfrak{o}_U(s) = 0 \quad\Longleftrightarrow\quad \mathrm{Ext}_U(s) \neq \varnothing.$$

Proof. Let $\pi_U : R \to R/B_U$ be the quotient homomorphism. By the definition of a quotient by a subgroup,

$$\pi_U(r) = 0 \quad\Longleftrightarrow\quad r \in B_U.$$

Here B_U is the *image* of $\ker p_U$ under δ. Consequently membership of $\delta\widetilde{s}$ means the existence of a specific $h \in E$ satisfying both $p_U h = 0$ and $\delta h = \delta\widetilde{s}$. The subgroup need not be closed in any topology. No topology has been introduced.

If such h exists, set $x = \widetilde{s} - h$. Additivity gives the two equations separately:

$$\delta x = \delta\widetilde{s} - \delta h = 0, \qquad p_U x = p_U\widetilde{s} - p_U h = s.$$

The first equation puts x in $\Gamma(A)$, and the second puts that coherent family in $\mathrm{Ext}_U(s)$.

Conversely, take $x \in \mathrm{Ext}_U(s)$ and set $h = \widetilde{s} - x$. Since both raw families restrict to s, one has $p_U h = 0$. Since $\delta x = 0$, one has $\delta h = \delta\widetilde{s}$. This witnesses $\delta\widetilde{s} \in B_U$, hence vanishing of its quotient class.

The proof is independent of the selected raw lift. Indeed, if $\widetilde{s}'$ is another lift, $k = \widetilde{s}' - \widetilde{s}$ lies in $\ker p_U$, and

$$\delta\widetilde{s}' - \delta\widetilde{s} = \delta k \in B_U.$$

Thus the two quotient classes agree before either feasibility direction is invoked. A raw lift always exists in this additive model by using zero outside U. Constructing a raw lift does not solve the coherence equations. □

The quotient is constructed before the solvability test. The theorem therefore does not define obstruction circularly by saying that it vanishes when a solution exists. It also does not give a decision procedure for membership in an arbitrary subgroup of an infinite product.

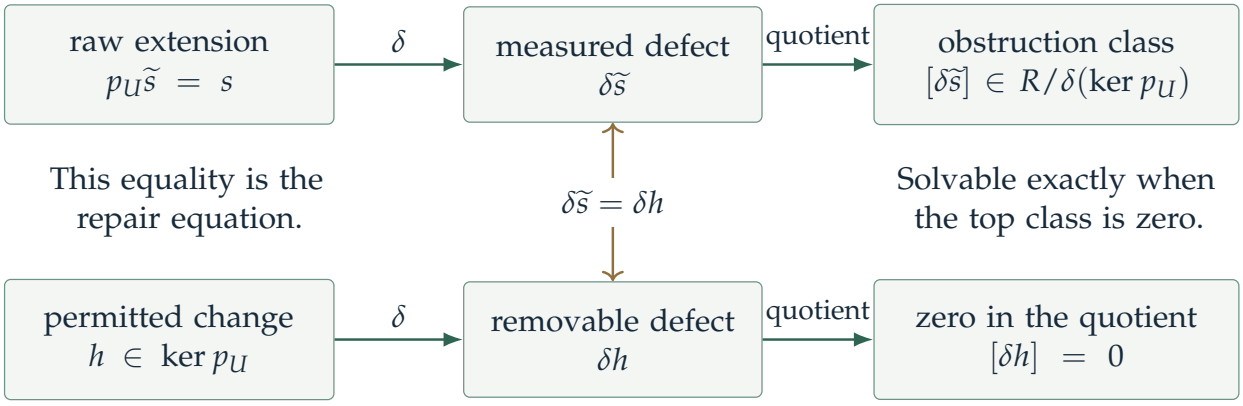


Figure 2: The raw family is repaired only along directions invisible on the prescribed views. The obstruction records whether its defect lies in the image of those admissible directions. Once an extension exists, a different kernel measures the remaining freedom.

3.3 Existence is not uniqueness

Define the ambiguity group

$$H_U = \ker\delta \cap \ker p_U.$$

It consists of coherent changes that are invisible on the prescribed domain.

Theorem 3.3 (Classification of a nonempty extension fiber)**.** *If $x_0 \in \mathrm{Ext}_U(s)$, translation gives an equivalence*

$$H_U \simeq \mathrm{Ext}_U(s), \qquad h \longmapsto x_0 + h.$$

Consequently the extension is unique exactly when $\mathfrak{o}_U(s) = 0$ and $H_U = 0$.

Proof. Define $\Theta : H_U \to \mathrm{Ext}_U(s)$ by $\Theta(h) = x_0 + h$. It lands in the extension set because

$$\delta(x_0 + h) = \delta x_0 + \delta h = 0, \qquad p_U(x_0 + h) = s + 0 = s.$$

Define $\Psi : \mathrm{Ext}_U(s) \to H_U$ by $\Psi(x) = x - x_0$. Its differential is zero, and its restriction is $s - s = 0$, so the claimed target is correct. The formulas give

$$\Psi\Theta(h) = x_0 + h - x_0 = h, \qquad \Theta\Psi(x) = x_0 + (x - x_0) = x.$$

This proves the equivalence rather than only a comparison of cardinalities.

If $H_U = \{0\}$, the parameter of every extension is zero, so every extension equals x_0. Conversely, if x_0 is the only extension, then for any $h \in H_U$ the equality $x_0 + h = x_0$ forces $h = 0$ by cancellation. Combining this observation with the existence criterion proves the stated necessary and sufficient condition for a unique extension.

The choice of x_0 only supplies an origin for the parameters. If $x_1 = x_0 + h_1$ is used instead, the parameter of x changes from $x - x_0$ to $x - x_1 = (x - x_0) - h_1$. The affine set of answers itself has not changed. □

An extension fiber is therefore an affine copy, or torsor, of the invisible coherent group. Selecting one solution gives coordinates on the fiber. Without a selected solution, the affine structure does not designate an origin.

Example 3.4 (Parity and an independent ambiguity)**.** For two views with $f^*(m) = 2m$ on $\mathbb{Z}$, a prescribed value k at the coarse view extends exactly when k is even. The extension is $(k, k/2)$ and is unique. In particular 1 is obstructed and 2 extends nontrivially. The parity statement is an instance of Theorem 3.2. If a separate, disconnected integer view is added and left unprescribed, every extendable seed has infinitely many completions. This last variant is a direct application of Theorem 3.3, not a claim that the parity obstruction has changed.

3.4 The exact sequence behind the criterion

The additive structures make restriction a homomorphism $r_U : \Gamma(A) \to \Gamma_U(A)$. Using the zero extension, or lift independence, $\mathfrak{o}_U$ is also a homomorphism.

Proposition 3.5 (Exactness at partial sections)**.** *The sequence*

$$\Gamma(A) \xrightarrow{r_U} \Gamma_U(A) \xrightarrow{\mathfrak{o}_U} Q_U$$

is exact at $\Gamma_U(A)$*. Thus* $\ker \mathfrak{o}_U = \mathrm{im}\, r_U$*, and restriction is onto precisely when* $\mathfrak{o}_U$ *is the zero map.*

Proof. First verify that $\mathfrak{o}_U$ is a homomorphism, since exactness is being asserted in an additive setting. Given $s, t \in \Gamma_U(A)$, choose raw lifts $\widetilde{s}, \widetilde{t}$. Their sum lifts $s + t$. Lift independence and additivity imply

$$\mathfrak{o}_U(s + t) = [\delta(\widetilde{s} + \widetilde{t})] = [\delta\widetilde{s}] + [\delta\widetilde{t}] = \mathfrak{o}_U(s) + \mathfrak{o}_U(t).$$

The zero lift gives $\mathfrak{o}_U(0) = 0$. Restriction r_U is a homomorphism because addition of global sections is coordinatewise.

If $s = r_U x$ for a global section x, use x itself as a raw lift. Its differential vanishes, so $\mathfrak{o}_U(s) = 0$. This proves $\mathrm{im}\, r_U \subseteq \ker \mathfrak{o}_U$. If $s \in \ker \mathfrak{o}_U$, the extension criterion constructs a global x restricting to s, proving the reverse inclusion.

Finally, r_U is surjective exactly when its image is all of $\Gamma_U(A)$. By the equality just proved, this is equivalent to every partial section having zero obstruction, which says that $\mathfrak{o}_U$ is the zero homomorphism. Nothing here claims that its image is all of Q_U. □

No surjectivity onto the entire residual quotient is asserted. The residual group contains many entries that need not arise as the differential of a partial lift.

3.5 Extension in stages

For $U \subseteq W \subseteq \mathrm{Ob}\,\mathcal{V}$, an extension first to W and then globally must retain the intermediate choice. There is an explicit equivalence

$$\mathrm{Ext}_U(s) \simeq \coprod_{\substack{t \in \Gamma_W(A) \\ t|_U = s}} \mathrm{Ext}_W(t). \tag{3.2}$$

It sends a global extension to its restriction to W, together with itself as an extension of that restriction. The inverse forgets the intermediate coordinate. These operations are inverse because the required intermediate coordinate of a global section is its actual restriction. A coherent intermediate choice can still be obstructed, even when a different intermediate choice completes. Equation (3.2) is why extension by stages is a dependent sum rather than a statement that every branch succeeds.

4 Protected correction as a general principle

The additive extension theorem is a particular instance of a useful general construction. Let E, R, S be abelian groups, with additive maps $d : E \to R$ and $p : E \to S$. The residual d measures compatibility, and p specifies the information that repair must preserve.

Theorem 4.1 (Semantics preserving correction)**.** *For $x \in E$, the following conditions are equivalent:*

$$\begin{aligned} &[dx] = 0 \text{ in } R/d(\ker p), \\ &\exists h \in E :\ p(h) = 0, \quad d(h) = d(x), \\ &\exists y \in E :\ d(y) = 0, \quad p(y) = p(x). \end{aligned}$$

The class of dx is unchanged when an admissible correction is subtracted from x.

Proof. The image $d(\ker p)$ is a subgroup of R. The defining property of the quotient says

$$[dx] = 0 \iff dx \in d(\ker p) \iff \exists h \in E :\ ph = 0,\ dh = dx.$$

This is the first equivalence with its witness made explicit.

For the second, a permitted h gives $y = x - h$. Then

$$dy = dx - dh = 0, \qquad py = px - ph = px.$$

Conversely an exact y with $py = px$ gives $h = x - y$, which satisfies $ph = 0$ and $dh = dx$. The two witness transformations are inverse: subtracting h from x and then subtracting that answer from x recovers h, and conversely.

To verify invariance, fix any h with $ph = 0$. Then $d(x - h) = dx - dh$, and dh lies in the subgroup divided out. Therefore $[d(x - h)] = [dx]$. This does not assert that the residual itself is unchanged. It asserts that permitted changes can alter the residual only within its already correctable coset. □

Here *semantics* means the declared mathematical observable p. It is not an assertion that one homomorphism captures every meaning of every object. In a tower, p might preserve a seed operator and every new output component. In a graph, it might preserve designated vertex values. Different choices give different problems.

Proposition 4.2 (A certificate of impossibility)**.** *Let $L : R \to D$ be additive. If L vanishes on $d(\ker p)$ but $L(dx) \neq 0$, no exact y can preserve $p(x)$.*

Proof. Suppose, for contradiction, that $dy = 0$ and $py = px$. Then $h = x - y$ satisfies $ph = 0$ and $dh = dx$. By the annihilation hypothesis,

$$L(dx) = L(dh) = 0,$$

contrary to the assumed nonzero value. This already proves impossibility directly.

Equivalently, if $r - r' \in d(\ker p)$, then $L(r - r') = 0$, so $L(r) = L(r')$. Hence $\overline{L}([r]) = L(r)$ is a well-defined homomorphism on $R/d(\ker p)$. Its value on $[dx]$ is nonzero. This supplies the quotient interpretation of the same contradiction, without assuming that every nonzero quotient class can be detected by a map into one fixed preferred coefficient group. □

This is often the most economical negative proof. One need not classify every possible correction. It is enough to find an invariant annihilating all admissible corrections but not the observed defect.

4.1 Why unrestricted cancellation is empty information

If every correction is permitted, taking $h = x$ leaves the exact candidate zero. Thus $\exists h : d(x - h) = 0$ is automatic. It says nothing about recovering an object of interest. A repair theorem becomes informative only when it preserves something, is quantitatively controlled, or satisfies a comparison property that excludes arbitrary replacement.

Example 4.3 (A rigid obstruction)**.** Take $E = R = S = \mathbb{Z}$, $d = p = \mathrm{id}$, and $x = 1$. The only admissible correction is zero, while the residual is nonzero. Hence there is no exact answer preserving the observed value. This is not a failure of the formalism. It is the correct certificate for an incompatible demand.

4.2 What an obstruction does not yet provide

An algebraic quotient does not automatically carry a useful norm. A zero infimum of distances to its boundary subgroup does not imply membership unless closure is controlled. Nor does vanishing supply a preferred correction when $\ker d \cap \ker p$ is nontrivial. These two issues lead to different parts of the book: ambiguity is treated algebraically and through groupoids, while closed range and condition numbers enter the analytic theory.

> A repair problem is not determined by its defect alone. It is determined by the defect *and* the admissible changes. Keeping the pair (d, p) explicit prevents a successful cancellation from being mistaken for a successful reconstruction.

5 Changing what is prescribed

The observable is part of a realization problem, not an annotation added after it has been solved. Prescribing more views can remove ambiguity, but it can also exclude every completion. This section follows both effects through the actual obstruction groups. It then works through a three-view model in which an initially harmless choice creates a genuine dead end.

5.1 Three maps that must not be confused

Keep the additive view system of Section 3. For a prescribed domain U, there are three distinct maps:

$$p_U : E \longrightarrow \prod_{u \in U} A_u, \qquad r_U : \Gamma(A) \longrightarrow \Gamma_U(A), \qquad \mathfrak{o}_U : \Gamma_U(A) \longrightarrow Q_U.$$

The first restricts a raw family, whether or not it is coherent. The second restricts a global coherent family. The third detects whether an already coherent partial family has a global completion. Their domains are different. In particular, surjectivity of the raw coordinate projection does not imply surjectivity of r_U.

The obstruction group is $Q_U = R/B_U$, where $B_U = \delta(\ker p_U)$. The ambiguity group is $H_U = \ker \delta \cap \ker p_U$. Both depend on what has been fixed. Enlarging the prescription reduces the permitted correction space and the invisible exact space at the same time.

5.2 Forgetting a prescription gives a quotient map

Let $U \subseteq W$. Then

$$\ker p_W \subseteq \ker p_U, \qquad B_W \subseteq B_U, \qquad H_W \subseteq H_U.$$

A correction invisible on all of W is certainly invisible on U. The direction of the quotient map is consequently

$$\pi_{W,U} : Q_W \longrightarrow Q_U, \qquad [r]_{B_W} \longmapsto [r]_{B_U}.$$

The larger prescribed domain appears in the source. More protection means a smaller subgroup is divided out, so more obstruction information can remain.

Theorem 5.1 (Functorial forgetting of protected views)**.** *The homomorphism $\pi_{W,U}$ is well defined and surjective. For $U \subseteq W \subseteq Z$, it satisfies*

$$\pi_{U,U} = I, \qquad \pi_{W,U}\pi_{Z,W} = \pi_{Z,U}.$$

For every $s \in \Gamma_W(A)$,

$$\pi_{W,U}\mathfrak{o}_W(s) = \mathfrak{o}_U(s|_U). \tag{5.1}$$

In particular an obstructed restriction remains obstructed when more compatible prescribed data are retained. A successful extension on the larger prescribed domain gives a successful extension of its restriction.

Proof. If $r - r' \in B_W$, then $r - r' \in B_U$, proving well-definedness. Every class modulo B_U has a representative r, and the class of that same representative modulo B_W maps to it. This proves surjectivity. Identity and composition follow by applying the formulas to a representative.

Choose a raw extension $\widetilde{s}$ of the partial section on W. It is also a raw extension of $s|_U$. Lift independence of the two extension obstructions therefore identifies both sides of (5.1) with the class of $\delta\widetilde{s}$ modulo B_U. A homomorphism sends zero to zero. Thus nonvanishing on the smaller domain rules out vanishing on the larger one, while vanishing on the larger domain implies vanishing after restriction. □

The converse implication does not hold. The quotient map can annihilate the obstruction created by a newly fixed value. Surjectivity of $\pi_{W,U}$ is also not a theorem that arbitrary larger partial sections extend. It is a statement about representatives in two obstruction groups, not about the image of restriction from global objects.

5.3 An initially free parameter can become an obstruction

Consider the chain of views $0 \to 1 \to 2$, with fiber $\mathbb{Z}^2$ at each view. Write the maps in the direction in which data are restricted:

$$\rho_{0,1}(a,b) = (a,0), \qquad \rho_{1,2}(a,b) = (a,2b), \qquad \rho_{0,2}(a,b) = (a,0).$$

The composite identity $\rho_{0,1}\rho_{1,2} = \rho_{0,2}$ is literal. A global family is exactly

$$x_0 = (k,0), \qquad x_1 = (k,2n), \qquad x_2 = (k,n), \qquad k,n \in \mathbb{Z}. \tag{5.2}$$

Indeed the first coordinate must agree at every view. The second coordinate at view zero is forced to vanish, and the second coordinate at view one must be twice the second coordinate at view two. These conditions also suffice.

First prescribe only $U = \{0\}$, with seed $(k,0)$. Every integer n in (5.2) gives a completion. Now prescribe $W = \{0,1\}$, with values $(k,0)$ and (k,m). This is a coherent partial section for every integer m, because $\rho_{0,1}$ forgets its second coordinate. Global extendability is a stronger requirement: m must be even.

Proposition 5.2 (A coherent branch can be a dead end)**.** *Every initial seed $(k,0)$ in this model has a completion set in bijection with $\mathbb{Z}$. A fixed intermediate value (k,m) extends precisely when $m \in 2\mathbb{Z}$, and then has a unique completion. In particular the nonzero initial seed $(7,0)$ extends, while the coherent intermediate choice $(7,1)$ does not.*

Proof. Equation (5.2) supplies a completion for each n, and the second coordinate of x_2 recovers the parameter. Fixing the intermediate value imposes $2n = m$. This equation has no integer solution for odd m, and exactly one for even m. Taking $k = 7$ and $m = 1$ proves the final assertion. □

The invisible coherent changes for the initial prescription are

$$((0,0),(0,2n),(0,n)), \qquad n \in \mathbb{Z}.$$

Thus $H_U \simeq \mathbb{Z}$. Requiring invisibility also at view one forces $2n = 0$, hence $n = 0$, so $H_W = 0$. More observation removes all ambiguity in the successful branches. It does not turn the unsuccessful branches into solutions.

5.4 Why the extension space is a dependent sum

Let s be the initial partial section. The stage theorem gives

$$\mathrm{Ext}_U(s) \simeq \coprod_{\substack{t \in \Gamma_W(A) \\ t|_U = s}} \mathrm{Ext}_W(t).$$

In the present example the index t is the integer m. The summand is empty for odd m and a singleton for even m. Taking their disjoint union gives one integer's worth of completions again. No contradiction appears between many initial completions and unique completion after a successful intermediate choice.

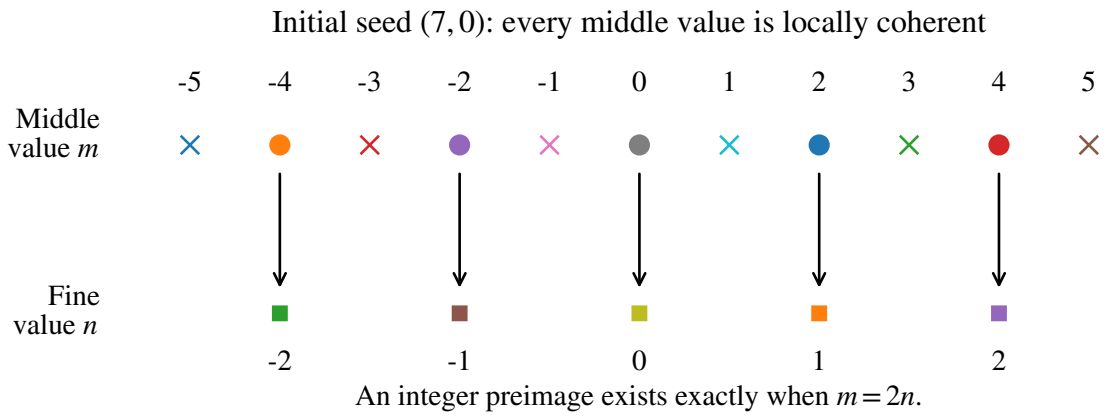


Figure 3: Every displayed middle value is internally coherent with the initial seed seven. Only even middle values lift to the last view. The initial problem has many global answers, but choosing an odd middle value destroys extendability. The arrows show actual integer preimages, not a probabilistic search policy.

This formula also explains a logical error in greedy reconstruction. The fact that some branch extends does not justify replacing the existential choice by an arbitrary coherent branch. A method that fixes an intermediate value must either retain the remaining extension obligation or supply a rule proved to preserve extendability. Coherence on the current patch alone provides no such guarantee.

For an increasing sequence of prescribed domains, the same warning applies repeatedly. The finite stage formula keeps each intermediate choice explicit. It does not by itself manufacture one compatible infinite sequence of choices. Proposition 35.2 supplies an existence criterion for towers of nonempty finite sets.

5.5 Changing the coefficients changes both questions

The equations behind the example can be read over a general abelian coefficient group T. With the initial first coordinate fixed, the possible completions are parametrized by $n \in T$. Prescribing the middle value m gives the equation $2n = m$. Direct substitution therefore gives

$$\text{feasible middle values} = 2T, \qquad \text{ambiguity after prescribing the middle} = \ker(2 : T \to T).$$

This calculation specializes as follows.

Coefficient group	Which middle values extend?	Completions of an extendable value
$\mathbb{Z}$	Even integers	One
$\mathbb{Q}$	Every rational	One
$\mathbb{F}_2$	Only zero	Two
$\mathbb{F}_3$	Every element	One

The integer instance is the actual formal example cited above. The other rows are worked coefficient calculations from the same two equations. In characteristic two, stronger prescription does not eliminate the hidden parameter because doubling annihilates it. Over the rationals it creates no parity obstruction because doubling is invertible.

A solver that changes the coefficient domain while leaving its user-facing description untouched changes the mathematical problem. The rational answer $1/2$ is not an integral completion of the odd middle value one. Conversely, a failure over $\mathbb{Z}$ is not evidence of a failure over $\mathbb{Q}$. The coefficient structure belongs beside the views, restrictions, and preserved values in the problem specification.

More information can reduce the set of answers without making the problem easier. The obstruction quotient records which new demands become impossible, while the ambiguity kernel records which freedom survives the demands that can be met.

Part II

Resolution, repair, and retained detail

Resolution gives a concrete family of views. Its compatibility differential connects exact maps, protected information, and residual coordinates. Recovering every finite coordinate is then separated from convergence in a chosen norm.

6 Resolution towers and exact transport

6.1 The sequential model

A tower in a category $\mathcal{C}$ is a sequence

$$\cdots \xrightarrow{\rho_2} X_2 \xrightarrow{\rho_1} X_1 \xrightarrow{\rho_0} X_0.$$

For sets, a point is a sequence $x_n \in X_n$ with $\rho_n(x_{n+1}) = x_n$. Composite restrictions turn the tower into a view system on $\mathbb{N}$, with an arrow $i \to j$ when $i \leq j$. Its section space is the same coherent point type, with the same coordinates.

A candidate transport from X to Y consists of maps $f_n : X_n \to Y_n$. It is exact when

$$\rho_n^Y f_{n+1} = f_n \rho_n^X \quad \text{for all } n. \tag{6.1}$$

Proposition 6.1 (The exact tower category)**.** *Identity maps and composites of exact transports are exact. An exact transport induces a map of coherent point spaces by $x \mapsto (f_n x_n)_n$, preserving identities and composition.*

Proof. For two exact transports, substitute $\rho_n^Z g_{n+1} = g_n \rho_n^Y$ and $\rho_n^Y f_{n+1} = f_n \rho_n^X$ into the composite square. For a coherent point,

$$\rho_n^Y f_{n+1} x_{n+1} = f_n \rho_n^X x_{n+1} = f_n x_n.$$

Identity, associativity, and composition on points are checked coordinatewise. □

This is a category on a fixed schedule. A full procategory allows additional reindexing and equivalence relations on maps. Those constructions must be supplied rather than identified with the fixed schedule by terminology.

6.2 The actual residual of a transport

Assume $\mathcal{C}$ is preadditive, so differences of maps make sense. The tower differential is

$$(\partial f)_n = \rho_n^Y f_{n+1} - f_n \rho_n^X.$$

Its target has components $\mathrm{Hom}(X_{n+1}, Y_n)$, not $\mathrm{Hom}(X_n, Y_n)$.

Proposition 6.2 (The transport Leibniz identity)**.** *For composable candidates,*

$$\partial(gf)_n = (\partial g)_n f_{n+1} + g_n (\partial f)_n.$$

In particular exactness is equivalent to vanishing of the differential.

Proof. Expand the right side. The terms $-g_n\rho_n^Y f_{n+1}$ and $+g_n\rho_n^Y f_{n+1}$ cancel, leaving $\rho_n^Z g_{n+1}f_{n+1} - g_n f_n \rho_n^X$. This is the differential of the composite. The exactness statement is (6.1). □

There is no additional mixed defect term in this identity. Such a term would require a different definition of the defect or a different expansion. Typed cancellation is a useful safeguard against importing an attractive but incorrect error formula.

6.3 Cofinal views preserve points, not every measurement

Let $\phi : \mathbb{N} \to \mathbb{N}$ be monotone and unbounded. The sampled tower uses levels $X_{\phi(n)}$ and their composite restrictions.

Theorem 6.3 (Cofinal reconstruction of points)**.** *Sampling induces an equivalence between the coherent points of X and of its sampled tower.*

Proof. Write $\rho_{i,j} : X_j \to X_i$ for the composite restriction when $i \le j$, with $\rho_{i,i} = I$. These maps satisfy $\rho_{i,j}\rho_{j,k} = \rho_{i,k}$. A coherent sampled point obeys

$$\rho_{\phi(n),\phi(m)} y_m = y_n \qquad (n \le m),$$

by iterating its adjacent coherence equations.

For each k, unboundedness gives an index n with $k \le \phi(n)$. Set $x_k = \rho_{k,\phi(n)} y_n$. If m is another permissible index, choose $\ell \ge n, m$. Monotonicity permits both composite restrictions, and sampled coherence gives

$$\rho_{k,\phi(n)} y_n = \rho_{k,\phi(n)}\rho_{\phi(n),\phi(\ell)} y_\ell = \rho_{k,\phi(\ell)} y_\ell = \rho_{k,\phi(m)} y_m.$$

Thus x_k is independent of the choice, including when the schedule repeats a level.

For the adjacent equation at k, choose one n with $k+1 \le \phi(n)$. Then

$$\rho_{k,k+1} x_{k+1} = \rho_{k,k+1}\rho_{k+1,\phi(n)} y_n = \rho_{k,\phi(n)} y_n = x_k.$$

The recovered family is therefore a point of the original tower.

Sampling a full point and recovering its kth coordinate gives $\rho_{k,\phi(n)} x_{\phi(n)} = x_k$ by its original coherence. Recovering a sampled point and evaluating at $k = \phi(n)$ allows the index n itself, giving $\rho_{\phi(n),\phi(n)} y_n = y_n$. These are both inverse identities. No surjectivity of the bonding maps or of the point evaluations was used. □

This theorem is about exact coherent points. Raw residual size is a different quantity. On identity scalar towers, $f_n = (-1)^n$ has nonzero fine differences, but its even subsequence is constant. The coarse residual vanishes because adjacent fine defects cancel. A cofinal equivalence of limits does not imply equality of residual norms or an absence of aliasing in uncorrected data.

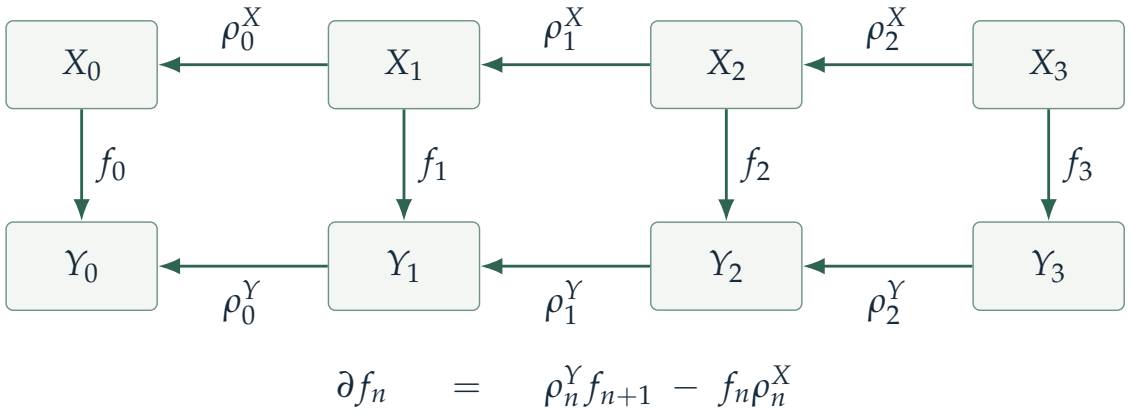


Figure 4: Exact transport commutes with every adjacent restriction. A cofinal subsequence retains coherent points, but a raw defect can cancel during coarsening. The retained coarse value alone need not record the omitted fine discrepancy.

7 Finite observations and realized transformations

The point-space construction has a converse question. Given a function between coherent realization spaces, can it be represented by compatible maps between finite views? The answer is not automatically yes. This section constructs the relevant category of reindexed maps and identifies its realized functions under an explicit observability hypothesis. Here a finite view means a view at a finite resolution level. Its set of possible values need not be finite.

7.1 An output level may require a finer input level

Let X and Y be sequential set-valued resolutions. Write $\rho^X_{i,j} : X_j \to X_i$ for composite restriction. A reindexed transport consists of a monotone schedule $\phi : \mathbb{N} \to \mathbb{N}$ and maps

$$f_n : X_{\phi(n)} \longrightarrow Y_n$$

satisfying, for $i \leq j$,

$$\rho^Y_{i,j} f_j = f_i \rho^X_{\phi(i),\phi(j)}. \tag{7.1}$$

The schedule records how much input is used for each output level. Monotonicity means that asking for more output never forces the representation to discard input already available. It does not require strict growth, a bounded look-ahead, or an efficiency estimate.

The realized map is

$$|f| : \mathrm{Point}(X) \to \mathrm{Point}(Y), \qquad |f|(x)_n = f_n(x_{\phi(n)}).$$

Equation (7.1) proves that this is a coherent output. A second transport $g_n : Y_{\psi(n)} \to Z_n$ composes with f using schedule $\phi \circ \psi$ and coordinate map $g_n f_{\psi(n)}$. The order of the schedules is fixed by the types, not by the order in which their letters are written in prose.

7.2 Equality is witnessed at finite stages

Two different schedules can represent the same transformation. To compare (ϕ, f) and (ψ, g), require that for every output level n there be a finite level $k \geq \phi(n), \psi(n)$ with

$$f_n \rho^X_{\phi(n),k} = g_n \rho^X_{\psi(n),k} \quad \text{as functions on all of } X_k. \tag{7.2}$$

Call this eventual equality. The witnessing k may depend on n. Equality is required on the whole finite-level fiber, not only on states that happen to extend globally.

Theorem 7.1 (The eventual transport category)**.** *Eventual equality is an equivalence relation and is preserved by composition. Its quotient therefore gives a category of resolutions and eventual transports. Realization descends to a functor into sets, preserving identities and composition.*

Proof. A finite-stage equality remains true after moving to a later stage. Indeed, if it holds on X_k and $\ell \geq k$, precompose both sides with $\rho^X_{k,\ell}$ and use the composite-restriction laws. Reflexivity uses $k = \phi(n)$, symmetry reverses the equality, and transitivity moves two witnesses to their maximum. The shared middle presentation then cancels as a literal function on that common fiber. This proves an equivalence relation without referring to global points.

We check compatibility with composition with all indices displayed. Write f, f' with schedules ϕ, ϕ' and g, g' with schedules ψ, ψ'. Fix an output level n. Choose $j \geq \psi(n), \psi'(n)$ with

$$g_n \rho^Y_{\psi(n),j} = g'_n \rho^Y_{\psi'(n),j}.$$

Next choose $k \geq \phi(j), \phi'(j)$ with

$$f_j \rho^X_{\phi(j),k} = f'_j \rho^X_{\phi'(j),k}.$$

Monotonicity implies that k is also at least both composite input levels. Applying the naturality equations and then the two witness equations gives

$$\begin{aligned} g_n f_{\psi(n)} \rho^X_{\phi(\psi(n)),k} &= g_n \rho^Y_{\psi(n),j} f_j \rho^X_{\phi(j),k} \\ &= g'_n \rho^Y_{\psi'(n),j} f_j \rho^X_{\phi(j),k} \\ &= g'_n \rho^Y_{\psi'(n),j} f'_j \rho^X_{\phi'(j),k} \\ &= g'_n f'_{\psi'(n)} \rho^X_{\phi'(\psi'(n)),k}. \end{aligned}$$

This is the required equality on every element of X_k. Thus composition descends to classes. The representative composition is associative because function composition and schedule composition are associative. The schedule $n \mapsto n$ with identity coordinate maps gives both identity laws. Hence the quotient has the claimed category structure.

Finally, evaluate a witness equation at the kth coordinate of a coherent point x. Its two restricted arguments become $x_{\phi(n)}$ and $x_{\psi(n)}$, proving equality of the realized outputs. Thus realization is well defined on the quotient. For a composite, direct evaluation gives

$$|g \circ f|(x)_n = g_n\big(f_{\psi(n)}(x_{\phi(\psi(n))})\big) = (|g|(|f|(x)))_n.$$

The identity law is equally coordinatewise. Joint determination of output points proves the functor laws. □

This is an actual quotient category in the formal source. No identification with the full procategory of an arbitrary base category is needed or claimed. The objects here are sequential set-valued resolutions, and the equivalence relation is the finite-stage relation just defined.

7.3 When global behavior determines the finite presentation

Say that X has surjective point projections when

$$\forall k \quad \mathrm{ev}_k : \mathrm{Point}(X) \to X_k \quad \text{is onto.}$$

This means that every finite input state really occurs in some coherent global input. It is stronger than nonemptiness of the global space.

Theorem 7.2 (Faithfulness on an observable source)**.** *If X has surjective point projections, two reindexed maps out of X are eventually equal exactly when they induce the same realized function. Surjective adjacent bonding maps are a sufficient condition for surjective point projections.*

Proof. Realized equality follows from eventual equality by the preceding theorem, so only its converse remains. Suppose $|f| = |g|$. Fix the output level n, and let $k = \max\{\phi(n), \psi(n)\}$. To prove equality of two maps defined on X_k, fix an arbitrary element $a \in X_k$. The assumed surjectivity of ev_k supplies a coherent point x with $x_k = a$. Equality of the two realized maps gives equality of their output coordinates:

$$f_n(x_{\phi(n)}) = g_n(x_{\psi(n)}).$$

Coherence of x rewrites this as

$$f_n\big(\rho^X_{\phi(n),k}(a)\big) = g_n\big(\rho^X_{\psi(n),k}(a)\big).$$

The index k was selected before a, so this is one common finite-stage equality for *all* elements of X_k. It is exactly the witness required by eventual equality at output n. Repeating for each n proves faithfulness.

For the sufficient condition on bonding maps, fix k and $a \in X_k$. Define coordinates below k by restriction of a. At each later stage use surjectivity of the adjacent bonding map to choose a preimage of the coordinate already chosen. Recursion with classical choice constructs the upper tail. Adjacent equations hold by construction, and induction on the interval length gives every composite equation. This produces a point evaluating to a. If a fiber is empty, the statement that each of its elements lifts is vacuous, as required by surjectivity onto that fiber. Without the projection-surjectivity assumption, equality on realizable inputs alone need not identify finite maps on values that never occur in a global point. □

The hypothesis is sufficient, not presented as a characterization of every source on which faithfulness happens to hold. Its role in the proof is precise: it converts a statement about globally realized inputs into a statement about all finite inputs.

Without it, global behavior can lose information. For example, use $X_n = \{m \in \mathbb{N} : m \geq n\}$ with inclusion bonds and a constant Boolean output tower. The source point space is empty. The two transports that always output false and always output true induce the same function from the empty set, but disagree at every finite stage. Every X_k is nonempty, so no eventual equality witness exists. This calculation uses the finite-prefix boundary of Example 35.1 and does not identify empty global behavior with equality of finite presentations.

7.4 The exact image of realization

For a function $F : \mathrm{Point}(X) \to \mathrm{Point}(Y)$, define finite observation by the quantifiers

$$\forall n\ \exists k\ \exists a : X_k \to Y_n\ \forall x : \quad F(x)_n = a(x_k). \tag{7.3}$$

The input level is uniform over all global inputs for each fixed output level. It may vary with the output level. This is not the weaker condition that each individual input can be inspected to some input-dependent finite depth.

Theorem 7.3 (Characterization of realized transformations)**.** *If X has surjective point projections, realization induces a bijection*

$$\mathrm{EventualHom}(X, Y) \simeq \{F : \mathrm{Point}(X) \to \mathrm{Point}(Y) : F \textit{ is finitely observed}\}.$$

Proof. For a reindexed map (ϕ, f), the formula $|f|(x)_n = f_n(x_{\phi(n)})$ immediately supplies a finite-observation witness for each n. Eventual equality does not affect the realized function by Theorem 7.1, so this defines a map from the quotient to the stated set of functions.

For surjectivity, let F be finitely observed. Choose one pair (k_n, a_n) at each output level, satisfying $F(x)_n = a_n(x_{k_n})$ for every coherent input. Define

$$\phi(n) = \max_{0 \le i \le n} k_i, \qquad f_n = a_n \rho^X_{k_n, \phi(n)}.$$

The maximum is of a nonempty finite set. It gives a monotone schedule with $k_n \le \phi(n)$. For every point x, coherence in X gives

$$f_n(x_{\phi(n)}) = a_n\big(\rho^X_{k_n,\phi(n)} x_{\phi(n)}\big) = a_n(x_{k_n}) = F(x)_n. \tag{*}$$

The remaining obligation is naturality on *all* finite states. Fix $i \le j$ and $u \in X_{\phi(j)}$. By surjectivity of that point projection, choose x with $x_{\phi(j)} = u$. Equation (*) and output coherence yield

$$\rho^Y_{i,j} f_j(u) = \rho^Y_{i,j} F(x)_j = F(x)_i,$$
$$f_i \rho^X_{\phi(i),\phi(j)}(u) = f_i(x_{\phi(i)}) = F(x)_i.$$

The state u was arbitrary. Hence the two finite-level functions are equal, which proves scheduled naturality. The resulting transport realizes F by (*).

Different witness choices can give different schedules and coordinate maps, but the realized function is still F. By Theorem 7.2 these transports are eventually equal, so their quotient class is independent of all choices. The same faithfulness result says that two classes with equal realized functions are identical. We have proved both surjectivity and injectivity, including the well-definedness of the inverse choice construction. No bound on the growth of ϕ follows from these quantifiers. □

The finite-level compatibility is derived in this proof. It is not inserted as an additional field in the hypothesis that F is finitely observed. This is one of the useful bridges between a global description of a transformation and an actual presentation in the resolution calculus.

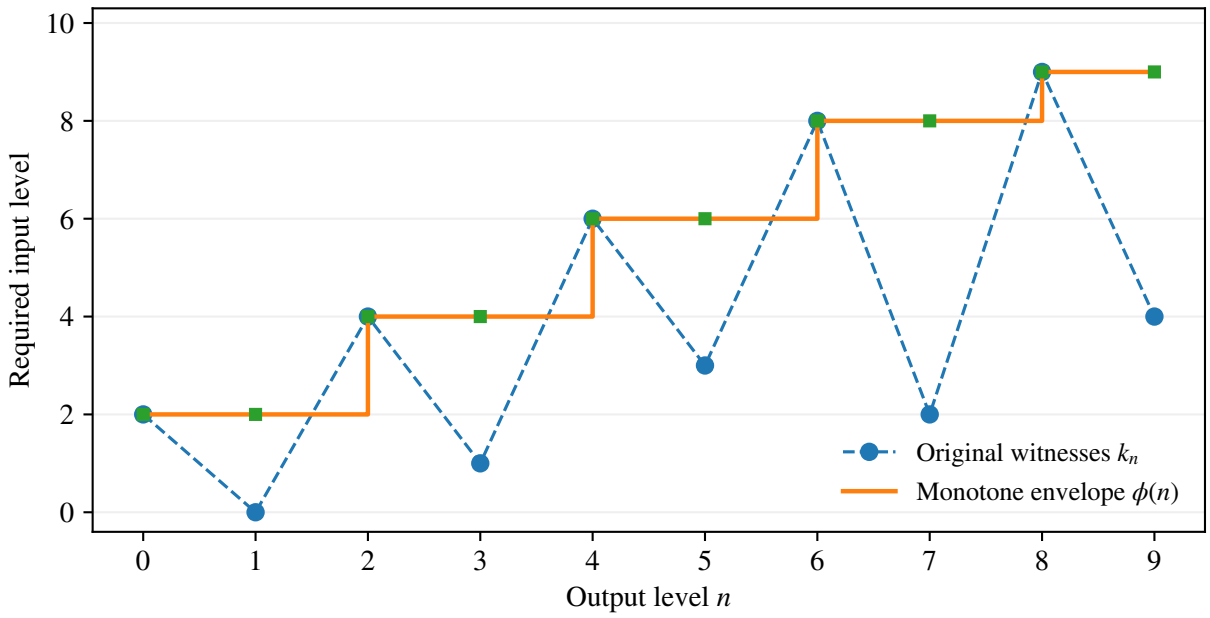


Figure 5: Independent finite-observation witnesses need not use a monotone input schedule. Taking the running maximum produces one without changing the realized function. The map at the enlarged level first restricts to the originally sufficient level. This construction gives no efficiency bound on how large those levels are.

7.5 Binary prefixes: a complete interpretation

Let X_n consist of binary words of length n, with prefix restriction. A coherent point is exactly an infinite binary sequence: its length-n view is its first n bits. Every prefix extends, for example by appending zeros. Hence the source satisfies the hypothesis of the characterization theorem.

Bitwise complementation is represented at level n by complementing the first n bits. The left shift $F(x)_i = x_{i+1}$ uses the schedule $\phi(n) = n + 1$. A nearest-neighbor difference over $\mathbb{F}_2$,

$$G(x)_i = x_i + x_{i+1},$$

also uses $n + 1$ input bits for n output bits. Applying G twice uses $n + 2$ bits, and the middle terms cancel:

$$G(G(x))_i = x_i + x_{i+2}.$$

The compatible finite maps make each statement meaningful on the entire coherent sequence, not merely on one chosen truncation. These are worked instances of the scheduled formulas. The finite regression suite checks their prefix compatibility independently.

Not every set function on these points is finitely observed. Consider the Boolean output that is one exactly when every input bit is zero, regarded as a point of the constant Boolean tower. Given any proposed finite level k, compare the all-zero sequence with a sequence whose first one occurs at position k. Their first k bits agree, but the outputs differ. Therefore no finite witness $a : X_k \to \{0, 1\}$ can compute even that one output coordinate. This rules out a reindexed presentation by the forward direction of the theorem.

The obstruction in this example is one of information, not a runtime benchmark. No amount of computation on the specified finite prefix distinguishes those two inputs. The theorem does not prohibit a different mathematical representation that supplies more information as part of each view. Changing the available view is changing the input model.

7.6 Cofinal changes also preserve the morphism calculus

A monotone unbounded schedule θ gives the sampled resolution X_θ. There is a transport from X to X_θ that reads the selected level. For the reverse direction choose, for each n, a level $j(n)$ with $n \le \theta(j(n))$, and restrict the available value down to X_n. Taking a least such index gives a monotone schedule.

Theorem 7.4 (Cofinal invariance of eventual transports)**.** *Sampling and recovery are inverse in the eventual transport category. Consequently, for monotone unbounded schedules θ and η, there is a bijection*

$$\mathrm{EventualHom}(X, Y) \simeq \mathrm{EventualHom}(X_\theta, Y_\eta).$$

On point spaces it agrees with the previously constructed cofinal realization equivalences.

Proof. Let $s : X \to X_\theta$ be sampling, with schedule θ and identity level maps. For recovery choose the least $j(n)$ with $n \le \theta(j(n))$. Unboundedness makes this set nonempty. Increasing n only strengthens the inequality, so j is monotone. The level map of $r : X_\theta \to X$ is $\rho^X_{n,\theta(j(n))}$. Composite restriction proves its naturality.

The composite $rs : X \to X$ reads level $\theta(j(n))$ and restricts to n. At that very stage its map agrees with the identity transport after restriction from $\theta(j(n))$ to n. Thus rs is eventually equal to the identity.

For $sr : X_\theta \to X_\theta$, the schedule at output n is $j(\theta(n))$, and the level map is $\rho^X_{\theta(n),\theta(j(\theta(n)))}$. Compare it with the identity at a sampled index

$$k \ge \max\{n, j(\theta(n))\}.$$

Restricting either route from level k of the sampled system gives the same original restriction $\rho^X_{\theta(n),\theta(k)}$. This proves eventual equality even when θ has repeated values. Hence the two classes are inverse isomorphisms, without requiring them to be literal inverse schedules.

Write s_X, r_X and s_Y, r_Y for the sampling and recovery isomorphisms of source and target. The asserted Hom map is

$$[f] \longmapsto s_Y[f]r_X,$$

with inverse $[g] \mapsto r_Y[g]s_X$. Associativity and the two inverse equations cancel the inserted pairs, proving both round trips. Applying the realization functor transports these equations to point spaces, where sampling and recovery are the cofinal point equivalences already constructed. No norm invariance or surjectivity of source point projections is needed for this categorical conjugation argument. □

This strengthens cofinal invariance from points to a specified category of transformations. It still does not make raw residual norms invariant, identify every global set function with a finite presentation, or preserve the numerical conditioning of a repair. Those are separate questions, addressed by the defect maps in the next sections.

> For a projection-surjective source, eventual finite presentations correspond exactly to uniformly finite observations at each output level. This is a precise machine-oriented representation theorem. It is not a claim that every represented map is computationally efficient, or that every arbitrary function of an infinite input has such a presentation.

8 Banach repair and constructed split coordinates

8.1 An actual complete candidate space

Let X_n, Y_n be Banach spaces over a complete nontrivially normed field, with continuous linear bonding maps and uniform bounds $\left\|\rho_n^X\right\| \le B_X$, $\left\|\rho_n^Y\right\| \le B_Y$. Use the spaces of bounded operator families

$$E = \ell_n^\infty \mathcal{L}(X_n, Y_n), \qquad R = \ell_n^\infty \mathcal{L}(X_{n+1}, Y_n).$$

The source formalization derives the required completeness from the operator spaces and their bounded product. The same differential satisfies

$$\|\partial f\|_\infty \le (B_X + B_Y)\, \|f\|_\infty .$$

Thus it is a continuous linear map, and evaluation at a fixed level is continuous. These facts connect algebraic compatibility to an analytic limit argument.

8.2 From a correction operator to a projection

The next statement is general. Let E be Banach, let $d : E \to R$, $p : E \to S$, and $A : R \to E$ be continuous linear maps. Suppose $pA = 0$ and, for $0 \le q < 1$,

$$\|df - dA(df)\| \le q\, \|df\| \qquad (f \in E). \tag{8.1}$$

The hypothesis concerns residuals of actual candidates. It need not hold for every independent vector in R.

Theorem 8.1 (Constructed geometric repair)**.** *Starting from $f^{(0)} = f$, define*

$$f^{(m+1)} = f^{(m)} - A(df^{(m)}).$$

There is a bounded linear projection $P_A : E \to \ker d$ with

$$f^{(m)} \longrightarrow P_A f, \qquad p(P_A f) = p(f),$$

and

$$\left\|f^{(m)} - P_A f\right\| \le \frac{\|A\|}{1-q} \left\|df^{(m)}\right\| .$$

It fixes every exact candidate and satisfies

$$\|P_A\| \le 1 + \frac{\|A\|\, \|d\|}{1-q}.$$

Proof. Put $r_m = df^{(m)}$. Additivity of d gives

$$r_{m+1} = r_m - dAr_m.$$

Each r_m is the residual of an actual candidate, so the stated hypothesis applies at every iteration. Induction yields, for $j \ge 0$,

$$\left\|r_{m+j}\right\| \le q^j \left\|r_m\right\| , \qquad \left\|f^{(m+j+1)} - f^{(m+j)}\right\| \le \|A\|\, q^j \left\|r_m\right\| .$$

For $n > m$, the finite triangle inequality therefore gives

$$\left\| f^{(n)} - f^{(m)} \right\| \le \|A\| \, \|r_m\| \sum_{j=0}^{n-m-1} q^j \le \frac{\|A\|}{1-q} \, \|r_m\| .$$

Since $\|r_m\| \le q^m \, \|df\|$ tends to zero, this is a Cauchy estimate. Completeness of E supplies a limit, denoted $P_A f$. Passing $n \to \infty$ in the same norm inequality proves the required tail estimate.

Continuity of d implies $d(P_A f) = \lim_m r_m = 0$. Protection holds before taking a limit: $pA = 0$ gives $pf^{(m+1)} = pf^{(m)}$, hence $pf^{(m)} = pf$ for every m. Continuity of p gives $p(P_A f) = pf$. This explains why continuity of the observable belongs to this analytic statement.

For linearity, write $L = I - Ad$, so $f^{(m)} = L^m f$. For fixed f, g and scalars λ, μ,

$$L^m(\lambda f + \mu g) = \lambda L^m f + \mu L^m g.$$

Continuity of addition and scalar multiplication allows passage to the norm limit and proves linearity of P_A. The estimate with $m = 0$ gives

$$\|P_A f\| \le \|f\| + \frac{\|A\|}{1-q} \, \|df\| \le \left(1 + \frac{\|A\| \, \|d\|}{1-q} \right) \|f\| ,$$

so the linear map is bounded. If $df = 0$, every iterate equals f, hence $P_A f = f$. Its image is contained in $\ker d$ and it fixes that kernel, proving surjectivity onto its stated image and $P_A^2 = P_A$. For $q = 0$ the residual vanishes after the first update and the same inequalities remain valid. □

This theorem constructs the dynamics from A. It does not construct a suitable A for every residual problem. That missing input can, however, be derived for a useful class of actual towers.

8.3 A right inverse from target splittings

Suppose each target bonding map has a continuous linear section

$$s_n : Y_n \to Y_{n+1}, \qquad \rho_n^Y s_n = I, \qquad \|s_n\| \le S.$$

Assume the propagation gap

$$SB_X < 1, \qquad K = \frac{S}{1 - SB_X}. \tag{8.2}$$

For $r \in R$, recursively set

$$(Ar)_0 = 0, \qquad (Ar)_{n+1} = s_n\big(r_n + (Ar)_n \rho_n^X\big). \tag{8.3}$$

Theorem 8.2 (Split tower correction)**.** *The recurrence defines a continuous linear map $A : R \to E$ with*

$$\partial A = I_R, \qquad \|A\| \le K.$$

The operator $P = I - A\partial$ is an exact projection and

$$(Pf)_0 = f_0, \qquad \|f - Pf\|_\infty \le K \, \|\partial f\|_\infty .$$

Proof. Let $\theta = SB_X < 1$. The number $K = S/(1-\theta)$ is nonnegative and satisfies

$$S + SB_X K = K.$$

Fix $r \in R$ and define its coordinates recursively as in (8.3). The initial coordinate has norm zero. If $\|(Ar)_n\| \le K\|r\|_\infty$, submultiplicativity of operator norms gives

$$\begin{aligned}\|(Ar)_{n+1}\| &\le \|s_n\|\left(\|r_n\| + \|(Ar)_n\|\left\|\rho_n^X\right\|\right)\\ &\le S(1+KB_X)\|r\|_\infty = K\|r\|_\infty .\end{aligned}$$

Induction proves the estimate at every coordinate, so the recursively defined family actually belongs to the bounded product E.

The coordinate recurrence preserves addition and scalar multiplication, starting from the zero initial coordinate. Induction therefore proves these laws for the family map A. Its uniform estimate proves $\|A\| \le K$, giving a continuous linear map with the claimed domain and codomain.

For every n, apply the target restriction to the recurrence. The equation $\rho_n^Y s_n = I$ gives

$$\rho_n^Y(Ar)_{n+1} = r_n + (Ar)_n\rho_n^X, \qquad (\partial Ar)_n = r_n.$$

Coordinate equality yields the operator identity $\partial A = I_R$. In particular this identity holds for every bounded residual, not just residuals already known to lie in the image.

Set $P = I - A\partial$. Then

$$\partial P = \partial - \partial A\partial = 0, \qquad P^2 = I - 2A\partial + A(\partial A)\partial = I - A\partial = P.$$

The expression $2A\partial$ abbreviates addition of the same operator twice and uses no division by two. Every exact f is fixed because $\partial f = 0$. Since $(Ar)_0 = 0$, the seed of Pf is f_0. Finally $f - Pf = A\partial f$, and the bound on A gives the asserted residual estimate. The denominators above are justified exactly by the strict propagation gap. □

The surjectivity is onto all bounded residual families. It is stronger than the assertion that some exact output can be assigned to each candidate. The gap and the chosen splittings are genuine inputs.

8.4 Protected information beyond the seed

When the target grows, its new kernel directions are independent information. Define

$$N_n(f) = (I - s_n\rho_n^Y)f_{n+1}, \qquad (p_Lf)_0 = f_0, \quad (p_Lf)_{n+1} = N_n(f).$$

Each $N_n(f)$ lands in $\ker\rho_n^Y$. The observable p_L retains both the seed and every new output component, relative to the chosen sections.

Theorem 8.3 (Split coordinate theorem)**.** *The observable is a bounded projection with $p_LA = 0$ and $p_LP = p_L$. It separates exact candidates. If*

$$\mathcal{S}_L = \{s \in E : p_Ls = s\},$$

then

$$\Phi : E \simeq_L \mathcal{S}_L \times R, \qquad \Phi(f) = (p_Lf, \partial f), \qquad \Phi^{-1}(s,r) = Ps + Ar.$$

Repair becomes $(s,r) \mapsto (s,0)$. It is the unique exact repair preserving p_Lf.

Proof. Put $B_n = I - s_n\rho_n^Y$. Since $\rho_n^Y s_n = I$, direct multiplication gives

$$B_n^2 = B_n, \qquad \rho_n^Y B_n = 0, \qquad B_n s_n = 0.$$

Moreover, $B_n y = y$ whenever $\rho_n^Y y = 0$. Thus B_n really projects onto the target kernel. Its norm is at most $1 + SB_Y$. Coordinate zero of p_L is the identity, and every later coordinate applies B_n, so $p_L^2 = p_L$ and $\|p_L\| \le 1 + SB_Y$. In particular the formula defines a map on the bounded product, not just on arbitrary coordinate families.

The recurrence for Ar gives zero seed and places $(Ar)_{n+1}$ in the image of s_n. The identities just proved make all its innovations zero. Therefore $p_L A = 0$, and

$$p_L P = p_L(I - A\partial) = p_L.$$

We next prove separation on the exact space. If f is exact, its restriction equation rewrites as

$$f_{n+1} = s_n \rho_n^Y f_{n+1} + B_n f_{n+1} = s_n f_n \rho_n^X + N_n(f).$$

Two exact families with the same seed and innovations agree first at level zero and then at every later level by this recurrence. More generally, if two arbitrary families have the same observable and differential, their difference is exact with zero observable, so separation on the exact space makes their difference zero.

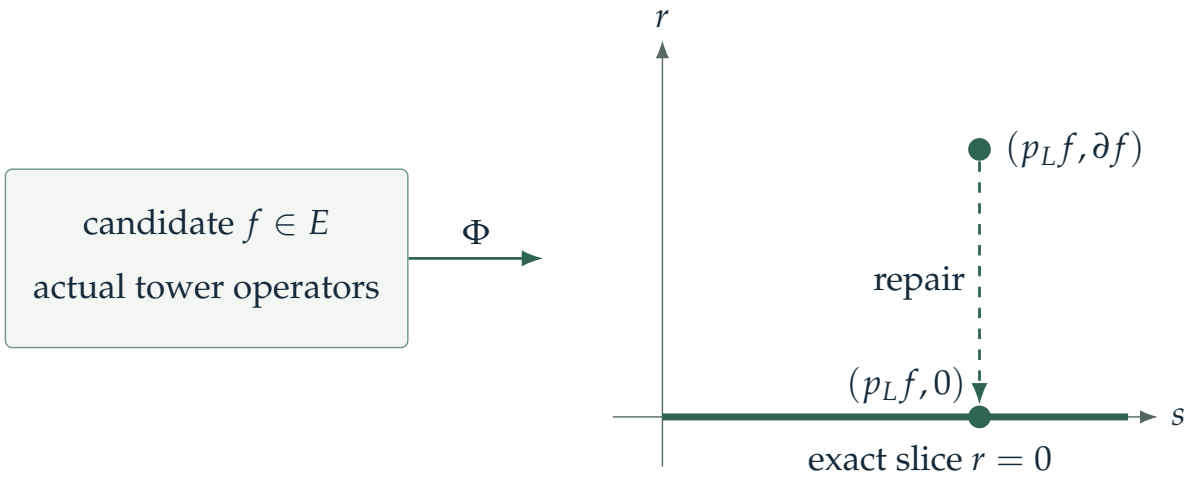
The proposed chart has the claimed target because $p_L(p_L f) = p_L f$. Given $s \in \mathcal{S}_L$ and $r \in R$, set $f = Ps + Ar$. The preceding identities and $\partial A = I$ give

$$p_L f = p_L s = s, \qquad \partial f = r.$$

This proves one round trip. For the other, f and $P(p_L f) + A\partial f$ have the same observable and residual, so the joint separation just proved identifies them. All maps are linear and bounded on their indicated subspaces and finite products, establishing a continuous linear equivalence.

Finally, Pf has coordinates $(p_L f, 0)$. Any exact answer with the same observable has those same two coordinates, hence equals Pf. This uniqueness is relative to the complete observable of seed and innovations. It is not obtained by assuming that the seed alone determines an exact growing-target family. □

This is an actual analytic coordinate equivalence, not generally an isometry. The protected coordinate depends on the declared splitting. It should not be called an intrinsic semantic coordinate without a further comparison theorem.



$$\Phi^{-1}(s,r) \;=\; Ps + Ar \qquad \text{The protected coordinate does not move.}$$

Figure 6: In the split coordinate chart, a candidate separates into retained information and compatibility residual. The exact subspace is the zero residual slice. The chart and its protected coordinate depend on the chosen target sections.

8.5 A growing example and the gap boundary

Use $X_n = Y_n = \mathbb{R}^{n+1}$ with maximum norms. Let π_n delete the last coordinate, take $\rho_n^X = \frac{1}{2}\pi_n$, $\rho_n^Y = \pi_n$, and let s_n append zero. Then $S = B_Y = 1$, $B_X = 1/2$, and $K = 2$.

Example 8.4 (Repair of the identity family)**.** The bounded candidate $f_n = I$ is not exact. Its repair is

$$(Pf)_n = \operatorname{diag}(2^{-n}, 2^{-(n-1)}, \dots, 2^{-1}, 1).$$

The new last output row is retained at every stage while inherited rows are propagated by the factor $1/2$.

Proof. At level zero the value is one. Exactness forces every inherited row at the next level to be one half of its previous value. Preservation of the innovation forces the new row to remain the corresponding identity row. Induction gives the displayed diagonal. □

For constant real scalar identity towers, by contrast, $(\partial f)_n = f_{n+1} - f_n$. The constant residual $r_n = 1$ has no bounded primitive, since any primitive satisfies $f_n = f_0 + n$. Here $SB_X = 1$. This counterexample excludes discarding the gap indiscriminately. It does not prove that the gap is necessary for every individual tower. Exact repair preserving the initial coordinate still exists in this constant example, by replacing a candidate with its initial constant family. Surjectivity onto all bounded residuals does not follow from that repair.

8.6 What is canonical

For fixed splitting and protected observable, the repair above is unique. Changing the splitting can change both the observable and the repaired answer. In the growing example, replace $s_n(y) = (y, 0)$ by $s_n'(y) = (y, y_0)$. At the first positive level, repair of the identity becomes

$$\begin{pmatrix} 1/2 & 0 \\ -1/2 & 1 \end{pmatrix} \quad \text{instead of} \quad \begin{pmatrix} 1/2 & 0 \\ 0 & 1 \end{pmatrix}.$$

Both are exact continuations for their respective innovation profiles and have the same seed. Thus a unique answer for a declared preservation problem is not an answer independent of all preservation choices.

9 Repair under a change of representation

An exact transformation should carry coherent inputs to coherent outputs. It need not carry the selected repair of an incoherent input to the selected repair of its transformed image. The distinction has an exact formula. In split coordinates, a compatible change of representation is generally block triangular rather than block diagonal. Its off-diagonal term records precisely the discrepancy between the two repair routes.

9.1 A section is a choice of residual coordinates

Let $d : E \to R$ be a continuous linear map between normed spaces over a nontrivially normed field. Suppose it has a continuous linear right inverse $A : R \to E$, so $dA = I_R$. Set

$$K = \ker d, \qquad P_A = I_E - Ad.$$

Then $dP_A = 0$, P_A fixes K, and $P_A^2 = P_A$. The right inverse chooses where each residual is represented inside the candidate space. Different right inverses can choose different representatives of the same residual.

Theorem 9.1 (The split defect chart)**.** *The maps*

$$\chi_A : E \simeq_L K \times R, \qquad \chi_A(x) = (P_A x, dx), \qquad \chi_A^{-1}(u, r) = u + Ar$$

are inverse continuous linear equivalences. In these coordinates, P_A sends (u, r) to $(u, 0)$. If B is another right inverse, then $A - B$ takes values in K, and

$$P_B x - P_A x = (A - B)dx.$$

The projections commute if and only if they are equal, and they are equal if and only if $A = B$.

Proof. The identity $dA = I$ gives $dP_A = d - dAd = 0$, so the proposed first chart coordinate actually belongs to K. If $u \in K$, then $du = 0$ and $P_A u = u$. Therefore, for $(u, r) \in K \times R$,

$$d(u + Ar) = r, \qquad P_A(u + Ar) = u + Ar - A(r) = u.$$

In the other direction, $P_A x + A(dx) = x - Adx + Adx = x$. These are both inverse equations, including the membership condition for the kernel coordinate. All component maps are bounded linear, so they define a continuous linear equivalence on the standard finite product. Applied to $P_A x$, the chart gives $(P_A x, 0)$, proving the repair formula.

For a second right inverse B, the equation $d(A - B) = I - I = 0$ places its difference from A in K. Direct subtraction gives

$$P_B - P_A = (I - Bd) - (I - Ad) = (A - B)d.$$

Both projections have range K and fix every point of K. Thus applying P_A after P_B gives P_B, while applying P_B after P_A gives P_A. These two compositions are equal if and only if the projections are equal.

If $P_A = P_B$, evaluate their difference at Ar. Since $dAr = r$, the difference formula yields $(A - B)r = 0$ for every r, hence $A = B$. Equality of sections conversely makes their projection formulas identical. The argument uses the right inverse to test every residual. It does not infer equality of sections merely from agreement on some restricted set of observed residuals. □

This is a linear splitting, not necessarily an orthogonal one. The exact coordinate u depends on the chosen section. The residual coordinate $r = dx$ does not. A specified Hilbert geometry can make one section distinguished through orthogonality, but the algebraic split alone does not do so.

The tower construction is stronger in a different direction. Its recurrence builds a right inverse from the target splittings, and its observable identifies what that repair preserves. The abstract chart here isolates the mechanism shared by those concrete constructions.

9.2 A change of representation has two components

Let $D : F \to T$ be a second continuous linear defect map. A compatible change consists of bounded linear maps

$$S : E \to F, \qquad C : R \to T, \qquad DS = Cd.$$

The equation is the commuting defect square. It says that transforming a candidate and then measuring its defect agrees with transforming its actual defect. No sections have yet been used.

In particular, S carries $\ker d$ into $\ker D$. It may still have a kernel, lose information, or fail to be onto. Calling it a change of representation does not assert an equivalence unless inverse maps have also been supplied.

Choose sections A of d and B of D. Define

$$M = SA - BC : R \longrightarrow F. \tag{9.1}$$

The commuting square and section identities give $DM = 0$. Thus this mismatch takes residual information into an exact direction of the target.

Theorem 9.2 (Triangular transport and repair mismatch)**.** *In split defect coordinates, the change of representation is*

$$\chi_B S \chi_A^{-1}(u, r) = (Su + Mr, Cr). \tag{9.2}$$

The two repair routes differ by

$$P_B Sx - SP_A x = M(dx). \tag{9.3}$$

Consequently they agree on exact inputs, and

$$\|P_B Sx - SP_A x\| \le \|M\| \, \|dx\|, \qquad \|M\| \le \|S\| \, \|A\| + \|B\| \, \|C\|.$$

They agree on all inputs if and only if $SA = BC$.

For two consecutive compatible changes (S_1, C_1), (S_2, C_2), with compatible choices of sections at the three ends, their mismatches satisfy

$$M_{21} = S_2 M_1 + M_2 C_1.$$

If $dx_j \to 0$ in norm for a fixed change and fixed sections, the repair discrepancy tends to zero.

Proof. First $DS = Cd$ sends $\ker d$ into $\ker D$, and

$$DM = DSA - DBC = CdA - C = C - C = 0.$$

Thus both terms in the proposed exact coordinate have the correct codomain. Given $u \in \ker d$, compute the residual and exact coordinates of $S(u + Ar)$ separately:

$$DS(u + Ar) = Cr, \qquad S(u + Ar) - BCr = Su + (SA - BC)r.$$

This proves the triangular formula as an equality in $\ker D \times T$.

For a general input, use the same commuting square to obtain

$$\begin{aligned} P_B S x - S P_A x &= Sx - BDSx - Sx + SAdx \\ &= (SA - BC)dx = M(dx). \end{aligned}$$

Its norm is at most $\|M\| \, \|dx\|$. The triangle inequality and submultiplicativity give $\|M\| \le \|S\| \, \|A\| + \|B\| \, \|C\|$. Exact inputs have $dx = 0$, so the two routes agree there regardless of the section choices. If they agree on *every* input, substitute $x = Ar$. Then $dAr = r$ makes $Mr = 0$ for each residual, hence $SA = BC$. That equality conversely makes the discrepancy vanish everywhere.

For composition, denote the three sections by A, B, A_2. The composite residual map is C_2C_1, so its mismatch is $S_2S_1A - A_2C_2C_1$. Insert and subtract S_2BC_1:

$$\begin{aligned} S_2S_1A - A_2C_2C_1 &= S_2(S_1A - BC_1) + (S_2B - A_2C_2)C_1 \\ &= S_2M_1 + M_2C_1. \end{aligned}$$

All terms map the first residual space into the last candidate space, in the written order. Finally, for fixed bounded M, the estimate $\|M(dx_j)\| \le \|M\| \, \|dx_j\|$ proves convergence of the discrepancy when $dx_j \to 0$. Allowing the section or representation norms to grow with j would require a separate uniform bound. □

The off-diagonal term in (9.2) has a specific meaning. A raw residual can become part of the target's chosen exact coordinate after a change of presentation. Deleting the residual before changing presentation removes that contribution. Deleting it afterwards does not. This is not a failure of exactness preservation. Both routes end in the target exact space, but they can end at different points in it.

9.3 A two-coordinate model in which the routes differ

Use $E = F = \mathbb{R}^2$, write candidates as (u, r), and take

$$d(u,r) = D(u,r) = r, \qquad A(r) = B(r) = (0,r).$$

The exact space is the horizontal axis and both repairs are $P(u,r) = (u,0)$. For a real parameter c, set

$$S_c(u,r) = (u + cr, r), \qquad C = I.$$

The defect square commutes because the second coordinate is unchanged. The mismatch is $M_c(r) = (cr, 0)$. Hence

$$PS_c(u,r) = (u + cr, 0), \qquad S_cP(u,r) = (u,0).$$

With $c = 2$ and input $(0,1)$, one route gives $(2,0)$ and the other gives $(0,0)$. Both have zero residual. A report containing only their final residuals would miss the entire difference.

Two shears compose by adding their parameters. In the mismatch law this appears as $M_{c_2+c_1} = S_{c_2}M_{c_1} + M_{c_2}$, and S_{c_2} fixes the exact horizontal axis. The identity is therefore a simple instance of the same composition rule, not an unrelated coordinate trick.

9.4 Actual tower resampling gives a defect map

Return to Banach towers with contractive source and target bonds. Let ϕ be a monotone schedule. Candidate sampling is

$$(S_\phi f)_n = f_{\phi(n)}.$$

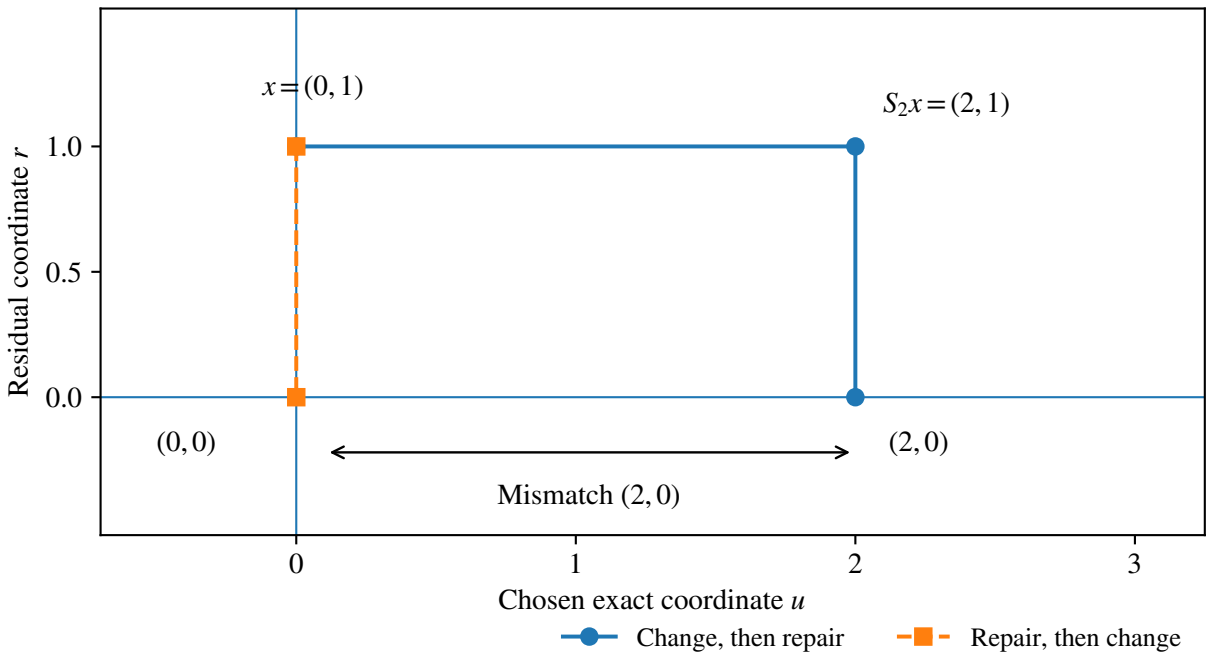


Figure 7: The shear preserves the residual coordinate but mixes it into the chosen exact coordinate. Repair after the shear reaches $(2,0)$. Repair before the shear reaches $(0,0)$. Their difference is the exact-valued mismatch applied to the original residual. The example uses the ordinary Euclidean drawing, while the algebraic identity does not depend on that drawing.

The sampled bonding map is the actual composite restriction between the selected levels. The residual map C_ϕ is not pointwise sampling of a scalar error score. It accumulates transported fine defects across each omitted block.

For a block from a to b, its value on a residual family r is

$$(C_{a,b}r) = \sum_{k=a}^{b-1} \rho^{Y}_{a,k}\, r_k\, \rho^{X}_{k+1,b}. \tag{9.4}$$

For the schedule ϕ, set $(C_\phi r)_n = C_{\phi(n),\phi(n+1)}r$. Every summand maps X_b into Y_a. These surrounding restrictions are necessary when the spaces change with the level. For an actual differential, the sum telescopes to

$$C_{a,b}(\partial f) = \rho^{Y}_{a,b} f_b - f_a \rho^{X}_{a,b}.$$

The dyadic coarsener in Section 10 is the case $b = a + 2$.

Theorem 9.3 (Bounded-gap resampling and selected repair)**.** *Suppose $\phi(n+1) - \phi(n) \le m$ for a fixed natural number m. Then the actual sampling maps are bounded and satisfy*

$$\|S_\phi\| \le 1, \qquad \|C_\phi\| \le m, \qquad \partial_\phi S_\phi = C_\phi \partial.$$

Exact candidates remain exact under every monotone common resampling, even without a gap bound. If A and B are bounded right inverses for the fine and sampled differentials, then

$$\left\| P_B S_\phi f - S_\phi P_A f \right\| \le (\|A\| + m\,\|B\|)\, \|\partial f\|\,.$$

Proof. For a bounded candidate family, $\sup_n \left\| f_{\phi(n)} \right\| \le \sup_k \|f_k\|$, proving the sampling bound. For a bounded residual family, each summand in (9.4) is a map from X_b to Y_a and has norm at most $\|r\|_\infty$, because every composite bond is contractive. A selected block has length at most m, so the triangle inequality bounds its sum by $m\,\|r\|_\infty$. Taking the supremum over blocks proves that C_ϕ maps into the bounded residual product and has the asserted operator norm. If a block has length zero its sum is zero, in agreement with the identity restriction.

To verify the commuting square, insert $r_k = \rho^Y_k f_{k+1} - f_k \rho^X_k$. The k-th term becomes

$$\rho^Y_{a,k+1} f_{k+1} \rho^X_{k+1,b} - \rho^Y_{a,k} f_k \rho^X_{k,b}.$$

Successive internal terms cancel in the finite sum, leaving $\rho^Y_{a,b} f_b - f_a \rho^X_{a,b}$. This is exactly the sampled differential. No infinite series interchange is involved.

If the fine differential vanishes, the finite telescope is zero on every selected interval. That conclusion remains true for unbounded gap lengths, because each individual interval is finite. What may fail without a gap bound is boundedness of the residual transformation on *arbitrary* bounded residual families, not preservation of exact inputs.

When the two bounded sections exist, Theorem 9.2 applies to the commuting square just established. Substitute $\left\|S_\phi\right\| \le 1$ and $\left\|C_\phi\right\| \le m$ in its estimate to get $(\|A\| + m\,\|B\|)\,\|\partial f\|$. Existence of those sections is an additional hypothesis, not a consequence of the sampling estimate. □

The existence of the two right inverses is not inferred merely from contractivity of the bonds. The split-tower theorem provides them under its additional section and propagation-gap conditions. Keeping these hypotheses separate prevents a resampling theorem from being mistaken for an unconditional reconstruction algorithm.

9.5 Two distinct losses under coarsening

On scalar identity towers, a fine residual block $(1, -1)$ has zero two-step sum. Thus $C_\phi r = 0$ need not imply $r = 0$, even for a bounded-gap schedule. This is cancellation. Retaining an independent detail coordinate remedies that information loss in a split coarsener, as the next section shows.

Unbounded gaps cause a different issue. With $\phi(n) = n^2$, the bounded residual family $r_k = 1$ has block sums $2n + 1$. These are unbounded. Therefore block accumulation cannot be a map from every bounded residual family to a bounded sampled residual family in this example. The constant residual is not claimed to be the differential of a bounded candidate on the identity tower. It is a witness about the proposed operator on the entire residual space.

Neither example contradicts cofinal invariance of coherent points or eventual transports. Those theorems concern exact compatibility. A norm on all raw defects, the retained fine detail, and the choice of a repair section are additional structures.

A compatible change of representation transports equations correctly. It transports the selected repair correctly only when the chosen residual sections also commute. The mismatch is an exact-valued operator, so it can remain invisible to a residual-only check.

10 Multiscale residuals and the limits of reconstruction

10.1 A split coarsener retains its missing detail

Let $C : R \to \overline{R}$ and $J : \overline{R} \to R$ be linear maps with $CJ = I$. Put $H = I - JC$. Then

$$r = JCr + Hr, \qquad CH = 0, \qquad H^2 = H.$$

The two coordinates (Cr, Hr) retain all information. The coarse coordinate alone generally does not.

The actual dyadic tower coarsener has the typed formula

$$(Cr)_n = \rho^{Y}_{2n} r_{2n+1} + r_{2n} \rho^{X}_{2n+1}. \tag{10.1}$$

It is obtained by telescoping two adjacent compatibility squares. With target sections, a lift is

$$(Jc)_{2n} = 0, \qquad (Jc)_{2n+1} = s_{2n} c_n.$$

The section identity gives $CJ = I$, so the missing detail is explicitly Hr, not an unspecified reconstruction error. The normed construction imposes the bounds needed to make these maps bounded on its operator family spaces.

10.2 A whole cascade

At scale j, let $C_j : R_j \to R_{j+1}$ and $J_j : R_{j+1} \to R_j$ be bounded linear maps with $C_j J_j = I$. For $x \in R_0$, define

$$x_0 = x, \quad x_{j+1} = C_j x_j, \quad L_0 = I, \quad L_N = J_0 \cdots J_{N-1},$$

$$h_j = x_j - J_j x_{j+1}, \qquad Q_N x = L_N x_N, \qquad S_N = I - Q_N.$$

The map Q_N restores the surviving coarse remainder to the original space. It is this restored remainder, not just x_N, that controls reconstruction.

Theorem 10.1 (Finite reconstruction and nested remainders)**.** *For every N,*

$$x = Q_N x + \sum_{j<N} L_j h_j, \qquad S_N x = \sum_{j<N} L_j h_j.$$

Each Q_N is a projection, and for $N \le M$,

$$Q_N Q_M = Q_M Q_N = Q_M.$$

The vectors with all details zero are exactly the persistent subspace

$$\mathcal{P}_\infty = \{x : Q_N x = x \text{ for every } N\}.$$

Proof. Let $D_0 = I$ and $D_N = C_{N-1} \cdots C_0 : R_0 \to R_N$. Thus $x_N = D_N x$, $Q_N = L_N D_N$, and repeated use of $C_j J_j = I$ gives

$$D_N L_N = I_{R_N}.$$

This cancellation takes place between adjacent maps with matching domains. No identification of different residual spaces is needed.

Since $L_{j+1} = L_j J_j$ and $x_{j+1} = C_j x_j$,

$$L_j h_j = L_j(x_j - J_j x_{j+1}) = Q_j x - Q_{j+1} x.$$

Summing from $j = 0$ to $N - 1$ leaves $x - Q_N x$ because $Q_0 = I$. This proves both finite reconstruction identities, including $N = 0$, when the detail sum is empty.

Idempotence follows from $D_N L_N = I$:

$$Q_N^2 = L_N(D_N L_N)D_N = Q_N.$$

For $M \geq N$, put $C_{N,M} = C_{M-1} \cdots C_N$ and $J_{N,M} = J_N \cdots J_{M-1}$, using identities for $M = N$. Then

$$D_M = C_{N,M} D_N, \quad L_M = L_N J_{N,M}, \quad C_{N,M} J_{N,M} = I.$$

The identities $D_N L_M = J_{N,M}$ and $D_M L_N = C_{N,M}$ now give

$$Q_N Q_M = L_N J_{N,M} D_M = Q_M, \qquad Q_M Q_N = L_M C_{N,M} D_N = Q_M.$$

Both orders have been checked, rather than inferred from one range inclusion.

If all h_j vanish, the finite telescope gives $Q_N x = x$ for every N. Conversely, if every $Q_N x = x$, then $L_j h_j = Q_j x - Q_{j+1} x = 0$. Apply the left inverse D_j of L_j to obtain $h_j = 0$. The injectivity of the restoration composite is the step that recovers the original detail from its transported value. □

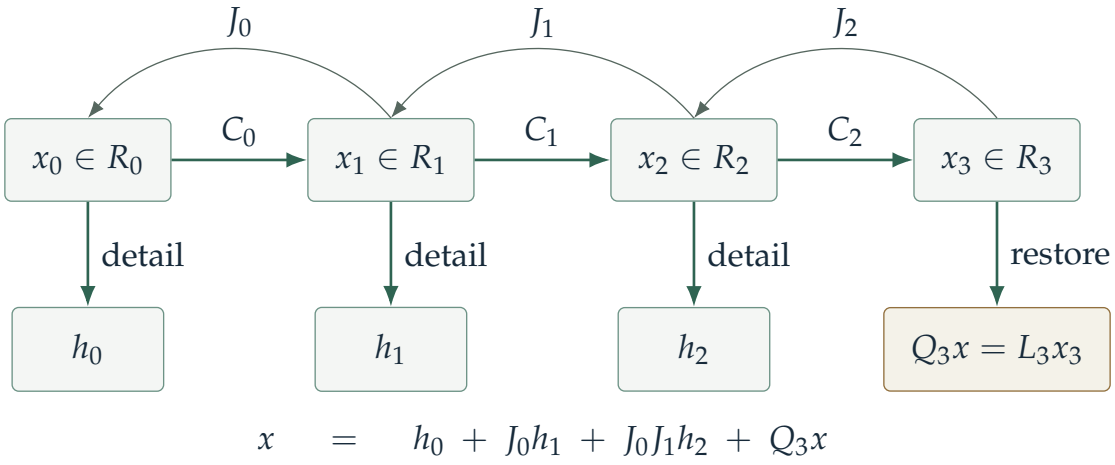


Figure 8: A split cascade retains every discarded component as a detail. The finite sum is always exact after the restored coarse remainder is included. An infinite detail expansion requires that remainder to vanish in the topology being used.

10.3 Convergence has its own theorem

Call x norm recoverable when $Q_N x \to 0$. This is equivalent to $S_N x \to x$, but it does not follow merely from vanishing persistent subspace. Let

$$\mathcal{F} = \bigcup_N \ker Q_N$$

be the finite core.

Theorem 10.2 (Density criterion for norm recovery)**.** *Suppose* $\sup_N \|Q_N\| \leq K < \infty$. *Then the norm recoverable vectors are exactly* $\overline{\mathcal{F}}$. *In particular all vectors are norm recoverable if and only if the finite core is dense.*

Proof. The nested identities first imply that the kernels increase: if $Q_M y = 0$ and $N \geq M$, then $Q_N y = Q_N Q_M y = 0$. Thus the finite core really is an increasing union of linear subspaces.

Suppose $Q_N x \to 0$. Put $y_N = S_N x = x - Q_N x$. Idempotence gives $Q_N y_N = Q_N x - Q_N^2 x = 0$, so $y_N \in \mathcal{F}$. Since $x - y_N = Q_N x \to 0$, these are finite-core approximants converging to x. Hence $x \in \overline{\mathcal{F}}$.

Conversely, fix $x \in \overline{\mathcal{F}}$ and $\epsilon > 0$. If $K = 0$, the common operator bound gives $\|Q_N x\| = 0$ for every N, so convergence is immediate. Suppose $K > 0$. Choose $y \in \mathcal{F}$ satisfying $\|x - y\| < \epsilon/(K+1)$. Some M has $Q_M y = 0$. For every $N \geq M$, nesting and the common operator bound give

$$\|Q_N x\| = \|Q_N(x-y)\| \leq K\,\|x-y\| < \epsilon \frac{K}{K+1} \leq \epsilon.$$

The approximation y is chosen before the depth N, giving precisely the quantifier order in norm convergence. This proves equality of the two sets. Applying it to every $x \in R_0$ gives the density criterion. □

There are other sufficient conditions. If the transported details are absolutely summable in the Banach space R_0, their sum exists and the remaining boundary lies in $\mathcal{P}_\infty$. If that persistent subspace is zero, the sum equals x. Alternatively,

$$\|x - S_N x\| \leq \left(\prod_{j<N} \|J_j\| \right) \|x_N\|$$

gives a direct sufficient decay test. A decreasing coarse norm by itself is not enough when restoration gains grow.

10.4 The actual dyadic result

For the dyadic residual hierarchy defined above, the persistent residual is zero. Zero detail at a scale reveals every even coordinate, while an odd coordinate is passed to a smaller index at the next scale. Strong induction on the coordinate index therefore proves that a residual with all details zero is zero.

Theorem 10.3 (Dyadic determination and finite coordinate recovery)**.** *In the actual dyadic system, all detail coordinates jointly determine the residual. Moreover,*

$$k+1 < 2^N \quad \Longrightarrow \quad (Q_N r)_k = 0, \qquad (S_N r)_k = r_k.$$

Thus every fixed residual coordinate is recovered exactly at a sufficiently large finite depth. The statement does not assert convergence in the supremum norm over all coordinates.

Proof. Write R_j for the residual space at scale j. Its lift has the actual coordinate formula

$$(J_j z)_{2m} = 0, \qquad (J_j z)_{2m+1} = s_{2m}^{(j)} z_m,$$

where $s_{2m}^{(j)}$ is the corresponding target section at that scale. The expression on the right is composition of the section with the residual operator. Its source and target are those of the odd coordinate at scale j.

Set $L_{j,0} = I_{R_j}$ and $L_{j,N+1} = J_j L_{j+1,N}$. We prove, simultaneously for all starting scales j, that

$$k+1 < 2^N \quad \Longrightarrow \quad (L_{j,N} z)_k = 0 \qquad (z \in R_{j+N}).$$

For $N = 0$, the inequality $k+1 < 1$ has no solution. Suppose the assertion holds at depth N. At depth $N+1$, an even coordinate vanishes by the displayed lift formula. For an odd coordinate $k = 2m+1$, the hypothesis gives

$$2(m+1) < 2^{N+1}, \qquad m+1 < 2^N.$$

The induction hypothesis at starting scale $j+1$ makes $(L_{j+1,N}z)_m$ zero. Composing that zero operator with $s_{2m}^{(j)}$ gives zero at coordinate k. This completes the induction without identifying residual spaces at different scales.

Now $Q_N r = L_{0,N} x_N$, so the support assertion applies to the restored remainder. Since $S_N r = r - Q_N r$, it gives the reconstruction equation. Finally, if two residuals have identical details at every scale, the finite telescope makes their S_N reconstructions equal for every N. At any fixed coordinate choose N with $k+1 < 2^N$. The reconstruction equation then makes their original coordinates equal. Hence the residuals are equal. □

The strict inequality matters: the next coordinate, $k = 2^N - 1$, can still have a nonzero remainder. There is also a sharper, nonconsecutive description of the support.

Proposition 10.4 (The dyadic support pattern)**.** *For the same lifts,*

$$2^N \nmid (k+1) \quad \Longrightarrow \quad (Q_N r)_k = 0.$$

Thus coordinate k is certainly recovered once $N > \nu_2(k+1)$, where $\nu_2(m)$ is the largest power-of-two exponent dividing the positive integer m. The only possible unrecovered coordinates at depth N are

$$2^N - 1, \quad 2 \cdot 2^N - 1, \quad 3 \cdot 2^N - 1, \quad \ldots.$$

Proof. Prove the assertion for $L_{j,N}$ by induction, uniformly in j. At depth zero every positive integer is divisible by one, so there is nothing to prove. At depth $N+1$, even coordinates vanish. If $k = 2m+1$, then

$$2^{N+1} \mid (k+1) \quad \Longleftrightarrow \quad 2^N \mid (m+1).$$

When divisibility fails, the inner coordinate in $J_j L_{j+1,N}$ vanishes by induction. The formula for the odd lift finishes the proof. Apply the result to $Q_N r = L_{0,N} x_N$. □

This refinement is formalized by a scale-uniform induction on the actual varying-fiber residual system in `DyadicSupport.lean`. It asserts where a remainder *may* survive, not that every allowed coordinate is nonzero. In the scalar identity-bonding model with $C(a,b) = a + b$, the residual constantly equal to one has value 2^N at every surviving position. That example proves sharpness of the uniform prefix cutoff.

10.5 Two boundary examples

On $\ell^\infty(\mathbb{N}, \mathbb{R})$, let C remove the first coordinate and let J prepend zero. Then $Q_N x$ is the tail of x, placed at its original coordinates with zeros before N. Every fixed coordinate is eventually zero in the remainder, and $\|Q_N\| \leq 1$. Nevertheless the constant sequence $\mathbf{1}$ satisfies $\|Q_N \mathbf{1}\| = 1$ for every N.

Proposition 10.5 (The recovery domain is c_0)**.** *For this shift cascade,*

$$S_N x \to x \text{ in } \ell^\infty \quad \Longleftrightarrow \quad x_n \to 0.$$

For $x_n = q^n$, $0 \leq q < 1$, the error is exactly q^N.

Proof. The error norm is $\sup_{n \geq N} |x_n|$. It tends to zero precisely when the sequence tends to zero. For a geometric sequence its largest tail value is the first one, q^N. □

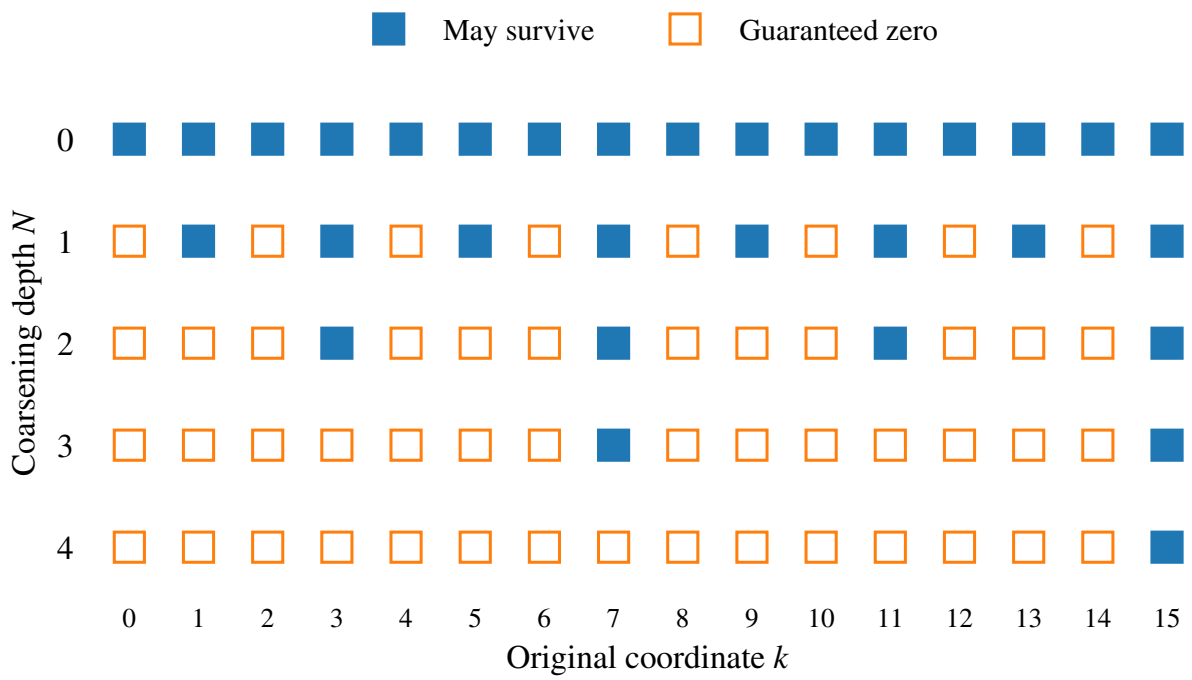


Figure 9: The possible support of the restored dyadic remainder is sparse and periodic. At depth N, only positions whose index plus one is divisible by 2^N can remain. Open cells are guaranteed recovered. Marked cells are allowed to survive, not asserted nonzero for every residual.

As a second example, take scalar spaces with $C = \frac{1}{2}I$, $J = 2I$. The coarse value $x_N = 2^{-N}x$ tends to zero, but restoration gives $Q_N x = x$. Every detail is zero and the entire input is persistent. The example belongs to this scalar cascade, not to the actual dyadic hierarchy, whose persistent subspace was proved to vanish.

Retained packets make omission explicit. They do not guarantee convergence. Coordinate recovery, bounded reconstruction, summability, and attainment remain different mathematical properties.

Part III

Symmetry and effective descent

A set of solutions forgets the witnesses relating equivalent solutions. The next layer retains those witnesses, their composition, and their automorphisms. Gluing then has to respect the overlap identifications themselves.

11 Nonabelian coherence and relative trivialization

11.1 Coherent transition labels

Fix groups G_v and restriction homomorphisms $f^* : G_w \to G_v$, with the same identity and composition laws as before. The groups need not be abelian and the homomorphisms need not be bijective.

Definition 11.1 (Cocycle)**.** A cocycle is a family $c_f \in G_v$ for every $f : v \to w$, satisfying

$$c_{1_v} = 1, \qquad c_{g\circ f} = c_f f^*(c_g).$$

It determines a twisted view system with maps

$$T_c(f)(x) = c_f f^*(x).$$

A section of the twisted system satisfies $x_v = c_f f^*(x_w)$.

The cocycle law says that the transition rules compose coherently. It does not assert that the twisted system has a section. The coherent transition system and a global trivialization of that system are different objects.

For a family $a = (a_v) \in \prod_v G_v$, define

$$(a \cdot c)_f = a_v c_f [f^*(a_w)]^{-1}. \tag{11.1}$$

Proposition 11.2 (Gauge action and its groupoid)**.** *Equation* (11.1) *defines an action on cocycles. A gauge carrying c to d induces an equivalence of twisted section sets by $x_v \mapsto a_v x_v$. The cocycles and actual gauges form a groupoid: an arrow $a : c \to d$ followed by $b : d \to e$ has label ba.*

Proof. First, T_c really is a view system. For $f : v \to w$, $g : w \to z$, and $x \in G_z$, the homomorphism and cocycle laws give

$$T_c(f)T_c(g)(x) = c_f f^*(c_g g^* x) = c_f f^*(c_g) f^* g^* x = c_{g\circ f}(g \circ f)^* x.$$

The identity equation follows from $c_{1_v} = 1$. No rearrangement of factors is involved.

Put $d = a \cdot c$. Its identity label is $a_v 1 a_v^{-1} = 1$. For a composite, expand all factors in their original order:

$$\begin{aligned} d_f f^*(d_g) &= a_v c_f [f^*(a_w)]^{-1} f^*(a_w c_g [g^*(a_z)]^{-1}) \\ &= a_v c_f [f^*(a_w)]^{-1} f^*(a_w) f^*(c_g) [f^* g^*(a_z)]^{-1} \\ &= a_v c_{g\circ f} [(g \circ f)^*(a_z)]^{-1} = d_{g\circ f}. \end{aligned}$$

Thus the gauge formula preserves the cocycle condition. For two gauges a, b, substitution gives

$$(a \cdot (b \cdot c))_f = a_v b_v c_f [f^*(b_w)]^{-1} [f^*(a_w)]^{-1} = ((ab) \cdot c)_f.$$

The reversed order in the two inverse factors is exactly the formula for the inverse of a product. The identity gauge acts trivially.

If $d = a \cdot c$ and x is a section of T_c, put $y_v = a_v x_v$. Then

$$d_f f^*(y_w) = a_v c_f [f^*(a_w)]^{-1} f^*(a_w) f^*(x_w) = a_v x_v = y_v.$$

The inverse map is multiplication by a_v^{-1} at each view, proving an equivalence of the actual section sets.

For the groupoid, an arrow $a : c \to d$ means the equation $a \cdot c = d$. If $b : d \to e$, then $(ba) \cdot c = b \cdot (a \cdot c) = e$, giving composition. The identity label is 1 and the inverse label is a^{-1}. Associativity and the two inverse identities now follow from the corresponding multiplication laws in the product group. The label ba is therefore forced by the order of application, not an arbitrary notational convention. □

The action groupoid $Z^1(\mathcal{V}; G)//\prod_v G_v$ retains comparison witnesses. Its orbit set $Z^1(\mathcal{V}; G)/\prod_v G_v$ retains only isomorphism classes. There is no reason for the orbit set to be a group in this nonabelian setting.

11.2 Automorphisms and ambiguity are not the same question

The stabilizer of a cocycle is the group of families with

$$a_v = c_f f^*(a_w) c_f^{-1}.$$

If one gauge $a : c \to d$ is known, all gauges from c to d are uniquely of the form ah, with $h \in \mathrm{Stab}(c)$. Indeed $a^{-1}b$ stabilizes c for every other gauge b. This is the ordinary Hom torsor principle in a groupoid.

That symmetry group must not be confused with the set of global trivializations of one fixed cocycle. The latter will be parametrized by ordinary coherent group elements, after one trivialization has been selected.

11.3 The relative obstruction without commutativity

Prescribe a coherent partial section s of T_c on U. The based gauge group is

$$\mathcal{G}_U = \{a \in \prod_v G_v : a_u = 1 \text{ for } u \in U\}.$$

Choose a raw family q equal to s on U. Identity elements can be used elsewhere. Normalize the labels by

$$c_f^q = q_v^{-1} c_f f^*(q_w).$$

Its labels on arrows internal to U are one. Define the pointed orbit obstruction

$$\mathfrak{o}_U(c, s) = [c^q] \in Z^1(\mathcal{V}; G)/\mathcal{G}_U,$$

whose neutral point is the class of the trivial cocycle.

Theorem 11.3 (Nonabelian relative extension)**.** *The obstruction is independent of the raw lift. It is neutral exactly when s extends to a global section of T_c.*

Proof. The normalized label c^q is the gauge transform of c by q^{-1}, so Proposition 11.2 makes it a cocycle. If $f : v \to w$ is internal to U, the partial-section equation says $q_v = c_f f^*(q_w)$, and therefore $c^q_f = 1$ on that arrow.

If t is another raw lift, define $a_v = t_v^{-1} q_v$. On U both lifts equal s, so $a_v = 1$ and $a \in \mathcal{G}_U$. Its action is

$$\begin{aligned}(a \cdot c^q)_f &= t_v^{-1} q_v \, q_v^{-1} c_f f^*(q_w)[f^*(t_w^{-1} q_w)]^{-1} \\ &= t_v^{-1} c_f f^*(q_w)[f^*(q_w)]^{-1} f^*(t_w) = t_v^{-1} c_f f^*(t_w) = c^t_f.\end{aligned}$$

Thus changing the raw lift leaves the based orbit unchanged.

Suppose the orbit is neutral. Equality of orbit classes means there exists an actual based gauge b with $b \cdot c^q = 1$. At an arrow $f : v \to w$, this equation is

$$b_v q_v^{-1} c_f f^*(q_w)[f^*(b_w)]^{-1} = 1.$$

Multiply on the left by $q_v b_v^{-1}$. Homomorphic preservation of products then gives

$$c_f f^*(q_w b_w^{-1}) = q_v b_v^{-1}.$$

Therefore $x_v = q_v b_v^{-1}$ is a twisted coherent family. On U, $b_v = 1$ and $q_v = s_v$, so it extends the prescribed section.

Conversely, suppose x is a global twisted section extending s. Put $b_v = x_v^{-1} q_v$. This is based, and direct substitution gives

$$(b \cdot c^q)_f = x_v^{-1} c_f f^*(x_w) = 1,$$

where the last equality is precisely the section equation for x. Hence the obstruction is neutral. These arguments establish the result using group cancellation only. In particular they do not assume an additive quotient, a norm on labels, or a decision procedure for finding b. □

The orbit is defined from actual cocycles and an actual group action. It is not a Boolean test defined to be neutral precisely when an extension exists. Deciding orbit membership in arbitrary groups remains a separate problem.

Let

$$H_U = \{h : f^*(h_w) = h_v \text{ for every } f,\ h|_U = 1\}.$$

Theorem 11.4 (The nonabelian extension torsor)**.** *Given one extension x_0, the map $h \mapsto x_0 h$ is an equivalence from H_U to the extension set. The extension is unique if and only if the obstruction is neutral and H_U is trivial.*

Proof. Take $h \in H_U$. Its pointwise product with the chosen extension x_0 agrees with the prescribed data on U, because $h|_U = 1$. For an arrow $f : v \to w$, the coefficient homomorphism preserves multiplication and the two coherence equations give

$$c_f f^*((x_0)_w h_w) = c_f f^*((x_0)_w) f^*(h_w) = (x_0)_v h_v.$$

Thus $x_0 h$ is another twisted extension.

Conversely, if x is a twisted extension of the same prescribed data, define $h_v = (x_0)_v^{-1} x_v$. On U its value is one. For every arrow,

$$\begin{aligned}h_v &= [c_f f^*((x_0)_w)]^{-1}[c_f f^*(x_w)] \\ &= f^*((x_0)_w)^{-1} c_f^{-1} c_f f^*(x_w) \\ &= f^*((x_0)_w^{-1} x_w) = f^*(h_w).\end{aligned}$$

The cancellation uses inverse multiplication in its actual order. No commuting factors have been moved past one another. Hence $h \in H_U$. Pointwise group cancellation shows that $h \mapsto x_0 h$ and $x \mapsto x_0^{-1} x$ are inverse maps.

The extension set is nonempty precisely when the earlier orbit obstruction is neutral. Once it is nonempty, the displayed bijection makes it a singleton precisely when H_U has only its identity element. Both conditions are needed for unique existence. Changing the chosen extension changes the parametrization, not the extension set or the criterion. □

The additive theory is recovered by applying a multiplicative type tag to an additive view system and choosing the trivial cocycle. The normalized label then becomes $f^* q_w - q_v$. The source proves agreement of the actual extension sets and vanishing criteria. It does not identify the whole nonabelian orbit set with every element of the larger additive arrow residual quotient.

11.4 A small noncommutative model

Example 11.5 (A C_2 loop with S_3 coefficients)**.** Take one view with a loop σ satisfying $\sigma^2 = 1$, and constant coefficients in a group H. A cocycle is determined by an element $p \in H$ whose square is the identity. Gauge equivalence is conjugacy, and its automorphisms form the centralizer of p.

For $H = S_3$, these elements are the identity and the three transpositions. There are exactly two orbit classes. The identity cocycle has automorphism group S_3, which is noncommutative. A nonidentity involution gives a coherent cocycle whose regular twisted section equation $x = px$ has no solution.

Proof. The only nontrivial cocycle relation is $p^2 = 1$. At a single view the gauge action is $p \mapsto apa^{-1}$. All three transpositions are conjugate and none is conjugate to the identity. The stabilizer is therefore the centralizer. Finally $x = px$ implies $p = 1$ after multiplication by x^{-1}. □

The example makes an important boundary visible. A cocycle can exist and glue as an object even when it admits no global trivializing section. Stackhood will concern the former gluing problem, not a claim that every cocycle becomes trivial.

12 Effective gluing of objects and their identifications

12.1 Why pairwise agreement is not the right datum

For objects with symmetry, local descriptions need not be literally equal on overlaps. They may instead come with specified isomorphisms. Those isomorphisms must satisfy their own composition law. It is not enough to know that some isomorphism exists for each pair.

Let U_i be full patches in $\mathcal{V}$. Assume the composition cover condition

$$\forall (v \xrightarrow{f} w \xrightarrow{g} z), \quad \exists i : \ v, w, z \in U_i. \tag{12.1}$$

By inserting identity arrows, this implies coverage of objects and arrows.

A descent datum consists of local cocycles c_i and overlap gauges $a_{ij}(v)$ with

$$c_j(f) = a_{ij}(v) c_i(f) [f^* a_{ij}(w)]^{-1}, \tag{12.2}$$

$$a_{ik}(v) = a_{jk}(v) a_{ij}(v). \tag{12.3}$$

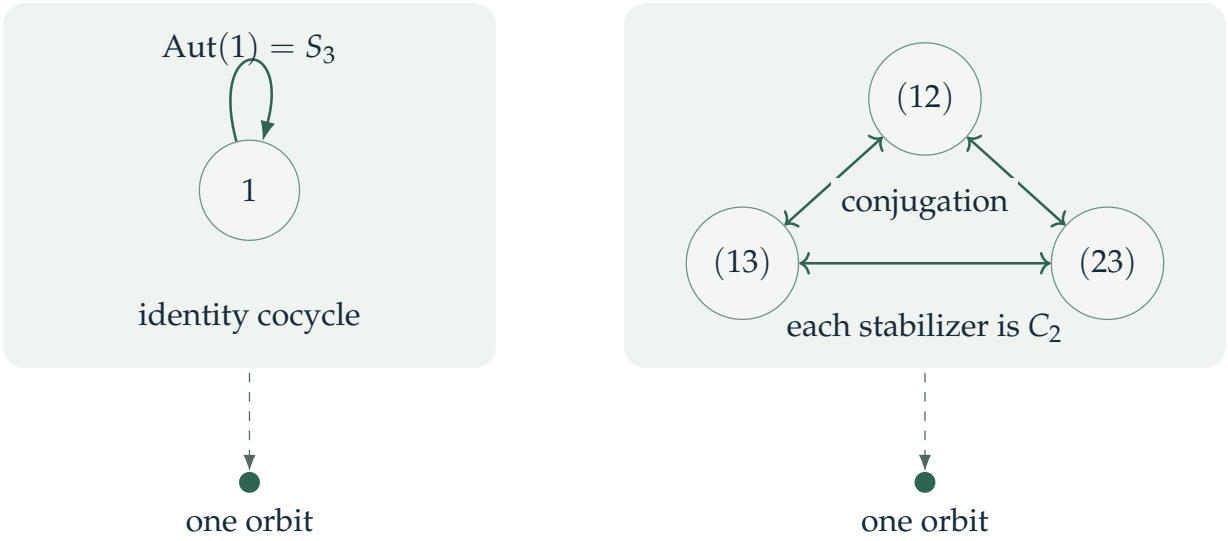


Figure 10: The orbit set sees two classes in the C_2-to-S_3 example. The groupoid also retains the arrows between conjugate transpositions and all automorphisms. Collapsing the arrows loses information that matters in descent.

The triple equation implies $a_{ii} = 1$ and $a_{ji} = a_{ij}^{-1}$. A morphism of descent data is a family of local gauges r_i satisfying

$$r_j(v)a_{ij}(v) = a'_{ij}(v)r_i(v).$$

These objects and arrows form a groupoid.

Theorem 12.1 (Effective cocycle descent)**.** *Under* (12.1)*, restriction is an equivalence between the global cocycle gauge groupoid and the groupoid of descent data. The overlap gauges supplied in the datum are realized, not discarded.*

Proof. Step 1: normalize the supplied identifications. The triple-overlap equation with all indices equal gives $a_{ii} = a_{ii}a_{ii}$, hence $a_{ii} = 1$. With indices i, j, i it gives $a_{ji}a_{ij} = 1$, hence $a_{ji} = a_{ij}^{-1}$. Composition coverage implies object coverage by inserting identity arrows. Choose a containing patch $\kappa(v)$ for each view and put

$$b_i(v) = a_{\kappa(v),i}(v) \quad (v \in U_i).$$

At a point of $U_i \cap U_j$, the triple law, now on patches $\kappa(v), i, j$, gives

$$b_j(v) = a_{ij}(v)b_i(v). \tag{*}$$

The two endpoints of an arrow may have different reference patches. No common reference patch is being assumed.

Step 2: obtain literal matching labels. Define

$$\widetilde{c}_i(f) = b_i(v)^{-1}c_i(f)f^*(b_i(w)) \quad (f : v \to w \text{ in } U_i).$$

This is a gauge transform and hence a cocycle on U_i. If the arrow lies also in U_j, substitute the original overlap equation and (*) at its two endpoints:

$$\begin{aligned}\widetilde{c}_j(f) &= b_i(v)^{-1}a_{ij}(v)^{-1}a_{ij}(v)c_i(f)[f^*a_{ij}(w)]^{-1}f^*(a_{ij}(w)b_i(w)) \\ &= b_i(v)^{-1}c_i(f)f^*(b_i(w)) = \widetilde{c}_i(f).\end{aligned}$$

Every cancellation is between neighboring inverse factors. Commutativity is not required.

Step 3: construct the global cocycle. Every arrow is contained in a patch, again by inserting an identity in a composable pair. Define c_f to be $\widetilde{c}_i(f)$ in any such patch. Step 2 proves independence. For an identity arrow, choose a patch through its view to obtain $c_{1_v} = 1$. For $f : v \to w$ and $g : w \to z$, composition coverage supplies a patch containing all three vertices. Fullness places f, g, and $g \circ f$ in that same patch. Its cocycle law therefore proves

$$c_{g\circ f} = \widetilde{c}_i(g \circ f) = \widetilde{c}_i(f) f^*(\widetilde{c}_i(g)) = c_f f^*(c_g).$$

This is the exact point where composition coverage, rather than arrow coverage alone, is used.

Undoing the normalization proves

$$b_i(v) c_f [f^*(b_i(w))]^{-1} = c_i(f).$$

Together with (*), these gauges constitute an isomorphism from the restriction of c to the supplied descent datum. They implement its original overlaps, proving essential surjectivity of restriction.

Step 4: retain every morphism. Let c, d be global cocycles. Their canonical restrictions have identity overlap gauges. A morphism between those restrictions is consequently a family r_i satisfying $r_j(v) = r_i(v)$ on each overlap. Assemble r_v from any containing patch. For an arrow choose a patch containing both endpoints. Its gauge equation proves

$$d_f = r_v c_f [f^*(r_w)]^{-1}.$$

Thus r is a global gauge. Its restriction is the original family. Conversely, assembling the restrictions of a global gauge preserves every coordinate and returns that gauge. Restriction is therefore bijective on each Hom set, not merely injective on isomorphism classes. This is full faithfulness.

Step 5: obtain an equivalence. For each descent object choose the global cocycle and isomorphism constructed above. For an arrow between descent objects, conjugate it by these two chosen isomorphisms and use the unique preimage under the Hom bijection of Step 4. This defines the arrow map of a quasi-inverse. Uniqueness of the preimage verifies its identity and composition laws, and the chosen isomorphisms provide the comparison to the identity functor. Hence the restriction functor is an equivalence of groupoids. Different patch choices can alter the representative but not this conclusion. □

12.2 The correct uniqueness statement

Suppose (c, b_i) and (c', b'_i) realize the same descent datum. The only possible compatible comparison is

$$g_v = b'_i(v)^{-1} b_i(v),$$

using any patch through v. The overlap identities make it independent of the patch, and $b'_i g = b_i$. It is therefore the unique global comparison respecting the fixed local identifications.

This is not a statement that the underlying cocycle has no automorphisms. Those automorphisms can reappear when the local identifications are forgotten. Nor does the construction choose one representative independently of every reference patch. It proves that the choices are related by the appropriate unique comparison.

12.3 A failed triple

Let $t \neq 1$ be an involution and ask for frames b_0, b_1, b_2 satisfying

$$b_1 = tb_0, \qquad b_2 = tb_1, \qquad b_2 = tb_0.$$

The first two give $b_2 = b_0$, while the third forces $t = 1$. Each displayed transition is individually invertible, but they cannot be the overlap identifications of a global frame. The defect lies in the triple relation.

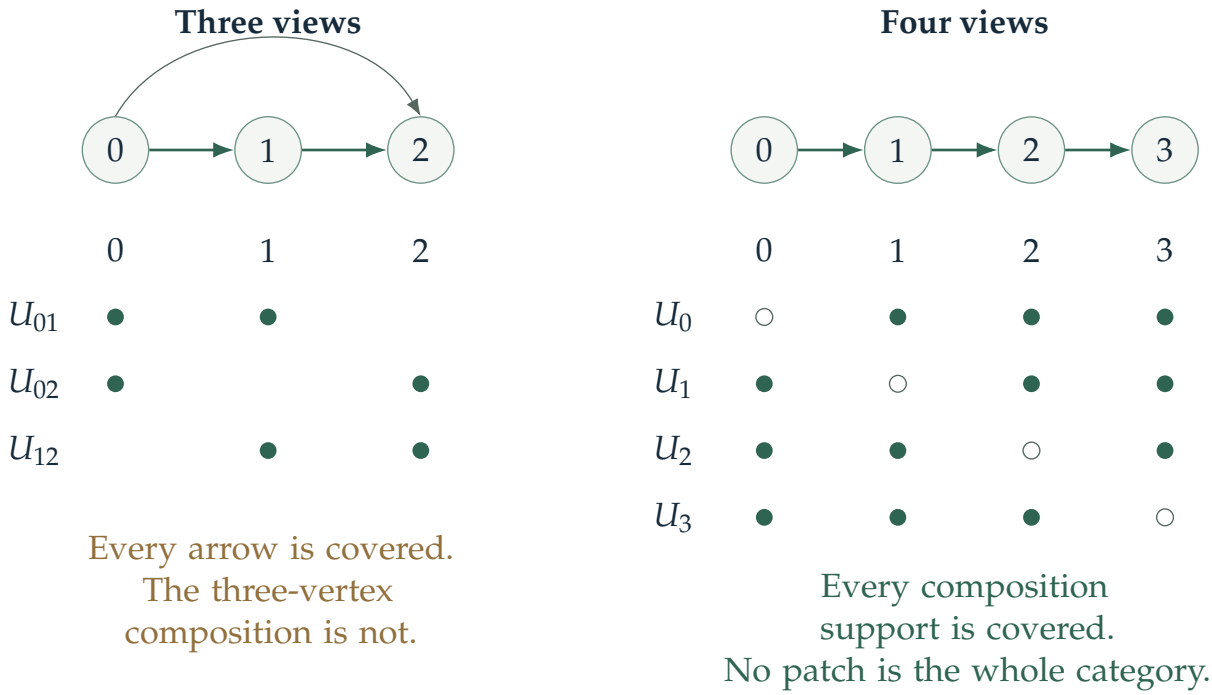


Figure 11: Object coverage, arrow coverage, and composition coverage are distinct. The proper two-view patches on a three-view chain see every arrow but no patch sees the relation $0 \to 1 \to 2$. The four proper three-view patches on a four-view chain do cover every composable pair.

13 A genuine stack on the composition site

The preceding theorem can be placed in a standard site, rather than called stack-like. We use the usual effective-descent meaning of a stack [2, 3]. The base of this site is the inclusion poset of full patches

$$\mathcal{P}(\mathcal{V}) = (\mathrm{Set}(\mathrm{Ob}\,\mathcal{V}), \subseteq).$$

This patch category is different from the original category of views.

13.1 The topology is defined without assuming effectivity

For a composable pair $q = (v \xrightarrow{f} w \xrightarrow{g} z)$, let its support be $|q| = \{v, w, z\}$. A sieve $\mathcal{R}$ on W is covering when

$$|q| \subseteq W \quad \Longrightarrow \quad (|q| \hookrightarrow W) \in \mathcal{R} \quad \text{for every composable } q. \tag{13.1}$$

This condition depends only on the view category, not on the coefficient groups or on a promised global realization.

Proposition 13.1 (The composition topology)**.** *Condition* (13.1) *defines a Grothendieck topology* J_{comp}*. The sieve generated by a family of inclusions* $U_i \subseteq W$ *is covering exactly when the family covers every composition support in* W*.*

Proof. In the inclusion poset, identify a sieve on W with a downward-closed collection of subsets of W. Pullback along $Z \subseteq W$ consists of those subsets of Z that were already members of the original sieve. We verify the three topology axioms using this description.

The maximal sieve contains every subset of W, hence every composition support. For stability under pullback, let $\mathcal{R}$ cover W and let q be a composable pair whose support lies in $Z \subseteq W$. The same pair lies in the full patch W, so $|q| \in \mathcal{R}$. Since $|q| \subseteq Z$, it belongs to the pulled-back sieve on Z. Thus that sieve covers Z.

For transitivity, suppose $\mathcal{R}$ covers W, and that for every $U \in \mathcal{R}$ the pullback of a sieve $\mathcal{T}$ to U covers U. Fix a composable pair q in W. Its support $S = |q|$ belongs to $\mathcal{R}$. Since patches are full, q is also a composable pair in S. The covering condition for the pullback of $\mathcal{T}$ to S therefore puts the identity inclusion $S \subseteq S$ in that pullback. By the definition of pullback, this says $S \in \mathcal{T}$. It holds for every q, so $\mathcal{T}$ covers W.

Finally, the sieve generated by $U_i \subseteq W$ consists exactly of the subpatches Z contained in at least one U_i. One inclusion follows because a sieve contains all restrictions of its generators. The reverse inclusion follows because this collection is itself downward closed and contains all generators, hence is the smallest such sieve. Consequently a support belongs to the generated sieve exactly when some U_i contains all its endpoints. This proves the family characterization. When W is empty, there are no composable pairs to test, so even the empty sieve is covering. □

Intersections of a covering family with a subpatch remain covering. Replacing each member by a covering family again gives a cover. Identities imply coverage of single views and arrows. The empty patch admits the empty cover.

13.2 The actual patch pseudofunctor

On a patch W, use

$$\mathfrak{K}_G(W) = Z^1(\mathcal{V}|_W; G|_W) /\!/ \prod_{v \in W} G_v.$$

Restriction keeps the existing cocycle values and gauge values on the smaller patch. It first gives a strict category valued functor on the opposite patch poset. Its standard pseudofunctor form has identity comparison isomorphisms, with the required coherence laws.

Theorem 13.2 (Stackhood on the composition site)**.** *The patch pseudofunctor* $\mathfrak{K}_G$ *is a stack for* J_{comp}*. It satisfies the standard condition of descent for morphisms and effective descent for objects on every covering sieve.*

Proof. We use the equivalence-of-descent-categories formulation of stackhood for the already constructed patch pseudofunctor. Its restriction maps act on actual cocycle labels and gauge values. Their identity and composition laws have been proved, so the descent condition is applied to a specified pseudofunctor, not assumed as part of its definition.

From standard descent data to overlap labels. For a covering family $(U_i \subseteq W)$, standard descent data provide local objects and comparison gauges on common subpatches, natural under further restriction. At a view $v \in U_i \cap U_j$, restrict the comparison to the singleton patch $\{v\}$ and call its value $a_{ij}(v)$. To verify the gauge equation along $f : v \to w$ in an overlap, use the comparison on $U_i \cap U_j$, which is a

full patch containing the arrow. Naturality identifies its two endpoint values with the singleton values. Its own gauge equation then becomes (12.2). Evaluating the standard triple-comparison identity at a singleton gives (12.3). Thus these labels form the concrete descent datum of Theorem 12.1.

From the concrete realization back to standard data. By Theorem 12.1 there is a global cocycle c and local gauges b_i implementing these overlaps. On any common subpatch $Z \subseteq U_i \cap U_j$, the required equation between gauge arrows is an equality of their values at every $v \in Z$. At such a view it is $b_j(v) = a_{ij}(v)b_i(v)$, already proved by the construction. Equality of coordinate functions proves equality of the arrows on Z. Thus every standard common-subpatch equation holds, not only the equations on singleton tests. This gives an isomorphism from the standard restriction of c to the original standard descent object.

Morphisms. A standard morphism between restrictions of global objects has local gauge components. Its comparison equations at singleton patches show that these components agree on overlaps. The Hom gluing map from Theorem 12.1 yields a unique global gauge. Conversely, a global gauge restricts to those components and satisfies all common-subpatch equations. Restriction and gluing are inverse on every Hom set. Standard restriction is therefore fully faithful as well as essentially surjective, hence an equivalence.

Every covering sieve. Let $\mathcal{R}$ be an arbitrary covering sieve on W. Index a family by all its member inclusions $U \subseteq W$. Each composition support belongs to this family by the covering definition, so it is a composition cover. Its generated sieve is exactly $\mathcal{R}$: one inclusion uses that $\mathcal{R}$ is downward closed, and the other uses each member as its own generator. The family argument therefore proves the descent equivalence on the original sieve. Applying it to every covering sieve proves both descent of morphisms and effectiveness of objects. These are the required standard stack conditions [2, 3]. □

The principal source declaration is `GroupViewSystem.patch_isStack`. Its target is mathlib's `Pseudofunctor.IsStack`, not a project predicate that repeats the conclusion as an assumption.

13.3 The cover is nontrivial, but specific

On a four-view chain, take the four patches obtained by deleting one vertex. All patches are proper, yet every composable pair has at most three endpoints and lies in one of them. The theorem genuinely glues pieces smaller than the whole object.

On the three-view chain, the analogous proper patches cover each object and each arrow but miss the support of $0 \to 1 \to 2$. They do not form a cover in this topology. Consequently the theorem is neither vacuous nor an unrestricted gluing principle.

Exact scope. This is an ordinary stack of strict group valued cocycle realizations on a specified combinatorial site. It is not an infinity-stack, a universal heterogeneous atlas, or an arithmetic existence theorem. The stack theorem glues coherent cocycles with their overlap isomorphisms. It does not assert that every resulting cocycle has a global trivialization.

14 Equivalent presentations and identified boundaries

14.1 Transporting the same question

Let $H : \mathcal{V} \to \mathcal{W}$ be part of a category equivalence, with inverse L and counit $\varepsilon : HL \Rightarrow \mathrm{id}_{\mathcal{W}}$. Pull a view system F on $\mathcal{W}$ back along H. Sampling a global section gives $(H^*x)_v = x_{Hv}$.

Proposition 14.1 (Equivalence of section spaces)**.** *Pullback induces $\Gamma(F) \simeq \Gamma(H^*F)$. If $U \subseteq \mathrm{Ob}\,\mathcal{W}$ is closed under isomorphism and $U' = H^{-1}U$, it also induces*

$$\mathrm{Ext}^{F}_{U}(s) \simeq \mathrm{Ext}^{H^*F}_{U'}(H^*s).$$

Existence and uniqueness of the prescribed extension are preserved and reflected.

Proof. Write the unit of the equivalence as $\eta : I_{\mathcal{V}} \Rightarrow LH$, and take the unit and counit with their triangle identities. Set $t_w = \varepsilon_w^{-1} : w \to HLw$. Given $y \in \Gamma(H^*F)$, define

$$(Ry)_w = F(t_w)y_{Lw}.$$

For an arrow $f : w \to w'$, naturality of the inverse counit states

$$HL(f)\, t_w = t_{w'} f.$$

Contravariance of F and coherence of y therefore give

$$\begin{aligned} F(f)(Ry)_{w'} &= F(f)F(t_{w'})y_{Lw'} \\ &= F(t_w)F(HL(f))y_{Lw'} = F(t_w)y_{Lw} = (Ry)_w. \end{aligned}$$

Thus recovery gives an actual coherent section.

If $x \in \Gamma(F)$, coherence along t_w gives

$$(RH^*x)_w = F(t_w)x_{HLw} = x_w.$$

For the other composite, the triangle identity says $t_{Hv} = H(\eta_v)$. Coherence of y along $\eta_v : v \to LHv$ yields

$$(H^*Ry)_v = F(t_{Hv})y_{LHv} = F(H\eta_v)y_{LHv} = y_v.$$

These coordinate equations prove both inverse identities.

For the relative claim, pullback clearly sends an extension of s to an extension of H^*s on $U' = H^{-1}U$. Conversely suppose y extends H^*s. If $w \in U$, the isomorphism $w \cong HLw$ and isomorphism closure imply $HLw \in U$. Hence $Lw \in U'$ and $y_{Lw} = s_{HLw}$. Both endpoints of t_w lie in U, so partial coherence gives

$$(Ry)_w = F(t_w)s_{HLw} = s_w.$$

Recovery therefore preserves the required boundary. The two inverse identities remain valid on the subsets of extensions. Bijections preserve nonemptiness and identify any two elements exactly when their images are equal, proving preservation and reflection of existence and uniqueness. □

The coefficients really are pulled back. This does not compare unrelated coefficient systems simply because their diagrams look alike. Isomorphism closure is a sufficient general hypothesis for the relative statement. The theorem does not assert that every unsaturated domain fails.

For group coefficients, the invisible coherent groups H_U are transported by a multiplicative equivalence, and the parameterizations $x_0 h$ commute with pullback. Thus feasibility and ambiguity are both retained. This is more than a comparison of the number of extensions. It is not an assertion that every possible based gauge orbit set is presentation independent.

14.2 Why a boundary must retain its identification

Let $R : \mathcal{C} \to \mathcal{D}$ be a functor between groupoids, and fix $b \in \mathcal{D}$. A realization over b should include a chosen isomorphism to b, rather than require literal equality of representatives.

Definition 14.2 (Marked realization groupoid)**.** The objects of $\mathcal{E}_b(R)$ are pairs (x, ϕ), with $x \in \mathcal{C}$ and $\phi : Rx \xrightarrow{\sim} b$. An arrow $(x, \phi) \to (y, \psi)$ is an arrow $a : x \to y$ with $\psi R(a) = \phi$.

Composition and inverses preserve this equation, so this is a groupoid. It is the explicit groupoid homotopy fiber construction, without invoking a higher categorical interface.

Proposition 14.3 (Existence and invisible automorphisms)**.** *The groupoid $\mathcal{E}_b(R)$ is nonempty exactly when b is in the essential image of R. An automorphism of (x, ϕ) is an automorphism a of x satisfying $R(a) = 1_{Rx}$.*

Proof. The first assertion is the object definition. For the second, the boundary equation is $\phi R(a) = \phi$. Cancel the isomorphism ϕ. Conversely $R(a) = 1$ makes the equation hold. □

This retains precisely the symmetries invisible at the identified boundary. It avoids freezing arbitrary representative coordinates and then mistaking the result for an invariant realization problem.

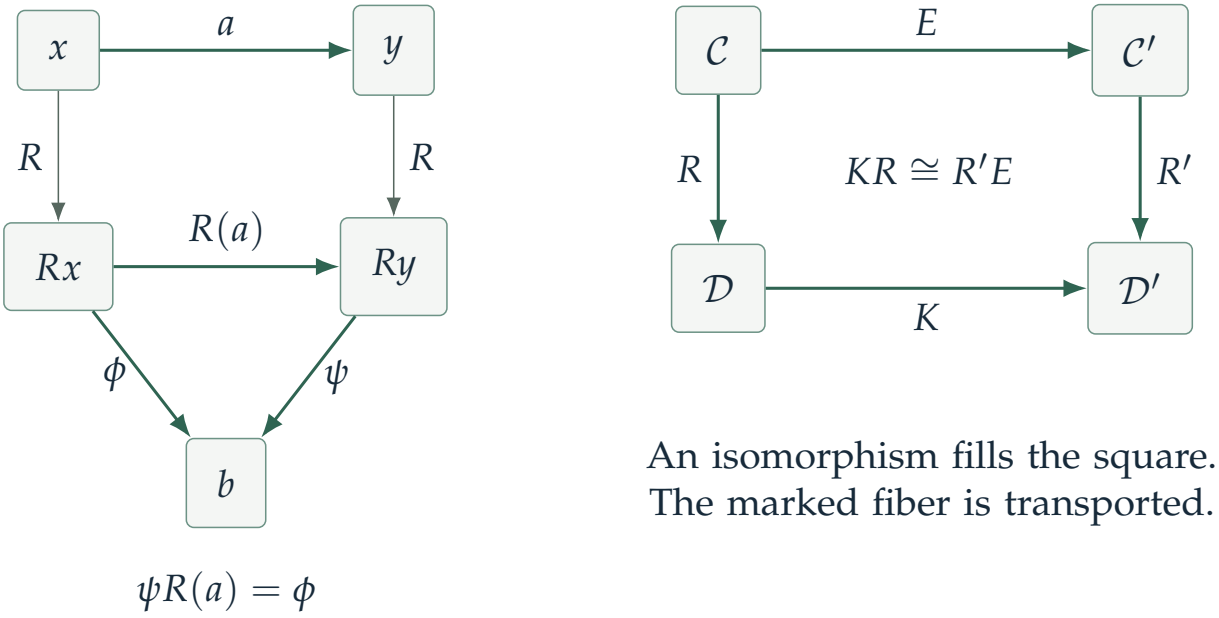


Figure 12: The chosen boundary identification is part of a realization. A comparison must commute with it. Under an equivalent presentation, the comparison square may commute through an isomorphism rather than a literal equality, and the marked realization groupoid is still transported correctly.

14.3 Weakly commuting squares are enough

Suppose $E : \mathcal{C} \to \mathcal{C}'$ and $K : \mathcal{D} \to \mathcal{D}'$ are equivalences, with functors R, R' and a natural isomorphism

$$\sigma : KR \xRightarrow{\sim} R'E.$$

The square is not required to commute literally.

Theorem 14.4 (Invariance of the marked realization problem)**.** *The map*

$$(x, \phi) \longmapsto (Ex, K(\phi)\sigma_x^{-1}), \qquad a \longmapsto E(a)$$

defines an equivalence $\mathcal{E}_b(R) \simeq \mathcal{E}_{Kb}(R')$.

Proof. For an object (x, ϕ) put $\beta_x = K(\phi)\sigma_x^{-1} : R'Ex \to Kb$. This is an isomorphism of the required type. If $a : (x, \phi) \to (y, \psi)$, naturality of σ says

$$R'E(a)\sigma_x = \sigma_y KR(a), \qquad \sigma_y^{-1} R'E(a) = KR(a)\sigma_x^{-1}.$$

Consequently

$$\beta_y R'E(a) = K(\psi)KR(a)\sigma_x^{-1} = K(\psi R(a))\sigma_x^{-1} = K(\phi)\sigma_x^{-1} = \beta_x.$$

So $E(a)$ is a marked arrow. Identity and composition are preserved because the underlying arrow map is that of the functor E.

To prove full faithfulness, start with a marked arrow $u : Ex \to Ey$ between the transported objects. Full faithfulness of E gives a unique $a : x \to y$ with $E(a) = u$. Its marked equation, rewritten by the naturality identity above, is

$$K(\psi R(a))\sigma_x^{-1} = K(\phi)\sigma_x^{-1}.$$

Cancel the isomorphism on the right and use faithfulness of K to obtain $\psi R(a) = \phi$. Thus the lifted arrow is marked. Uniqueness of the underlying lift also proves uniqueness as a marked arrow, giving a bijection on every Hom set.

For essential surjectivity, let (y, ψ) be a target realization. Choose x and an isomorphism $i : Ex \to y$, using essential surjectivity of E. Fullness of K lifts the well-typed composite

$$KRx \xrightarrow{\sigma_x} R'Ex \xrightarrow{R'(i)} R'y \xrightarrow{\psi} Kb$$

to an arrow $\phi : Rx \to b$. Since $\mathcal{D}$ is a groupoid, ϕ is automatically an isomorphism. Its image satisfies $K(\phi) = \psi R'(i)\sigma_x$, hence $\beta_x = \psi R'(i)$. This is precisely the condition for i to be an arrow from the transported realization to (y, ψ). It is an isomorphism because all categories involved are groupoids.

We have verified full faithfulness and essential surjectivity of the specified transport functor. Choosing these preimages and lifting arrows through the Hom bijections constructs a quasi-inverse, as in Theorem 12.1, proving the claimed equivalence. No literal equality $KR = R'E$ was used. □

Apply this to the actual restriction functors between cocycle patch groupoids and the equivalences induced by an equivalent view presentation. The full realization problem, including comparison witnesses and automorphisms, is transported.

Replacing b by an isomorphic boundary b' also gives an equivalence, by composing the identification with $b \to b'$. For $R = \mathrm{id}$, the marked realization groupoid has center $(b, 1_b)$, and between (x, ϕ) and (y, ψ) there is exactly one compatible arrow $\psi^{-1}\phi$. This is a useful example of contractible realization data, not a claim that the original groupoid has no symmetries.

15 A boundary identification is data

The groupoid formulation can seem more elaborate than a set of equivalence classes until a change of presentation is examined closely. A two-view example is enough to show what is lost by freezing representatives instead of retaining their comparison maps. It also explains why the marked realization construction uses an identification at the boundary.

15.1 Duplicate a view without adding information

Let the original view category have one object and only its identity arrow. Replace it by a category with two objects, 0 and 1, and exactly one arrow between each ordered pair. These arrows are inverse isomorphisms. The two categories are equivalent: selecting object zero and forgetting the duplicate give inverse equivalences up to the evident isomorphism.

Use a constant group H as coefficients, with identity pullback homomorphisms. A cocycle on the duplicated category is specified by its label $r \in H$ on the arrow $0 \to 1$. Its reverse label is r^{-1}, and identity labels are one. A gauge (a_0, a_1) carries r to

$$r' = a_0 r a_1^{-1}. \tag{15.1}$$

Every r is isomorphic to the trivial cocycle, for example with gauge $(1, r)$. The automorphisms of r are the pairs

$$(a_0, r^{-1} a_0 r), \qquad a_0 \in H.$$

Thus the unframed groupoid has one isomorphism class but retains a copy of H as the automorphism group of each object. This is compatible with the original one-view coefficient groupoid, which has one object and automorphism group H.

15.2 Freezing both representative gauges breaks that comparison

Now require a based gauge to be the identity at every view. On the duplicated category the only permitted gauge is $(1, 1)$. Equation (15.1) then carries r only to itself. Distinct labels cannot be compared.

Restrict to the selected view zero. Every cocycle becomes the same one-view identity cocycle. The distinction between r and 1 has disappeared in the target, even though the fully framed source has no comparison arrow between them.

Proposition 15.1 (Full representative framing is not invariant)**.** *For a nontrivial coefficient group, the restriction from the fully framed duplicated model to the fully framed single-view model need not be full, and hence need not be an equivalence. This failure occurs even though the underlying view categories are equivalent and the full prescribed domains are closed under isomorphism.*

Proof. Choose $r \neq 1$. In the fully framed duplicated model there is no arrow from r to 1, because the only allowed gauge leaves r unchanged. After restriction to view zero, both become the same cocycle. There is an identity arrow between their images. That target arrow has no preimage, so the functor is not full. Fullness is necessary for an equivalence. □

The failure is not repaired by declaring the two objects equivalent as sets of labels. The issue is exactly the missing arrow. Nor does it contradict the theorem for marked relative realizations. That theorem does not freeze every arbitrary representative coordinate. It transports an actual boundary object together with an identification to it.

15.3 The marked fiber keeps the comparison that matters

In the same example, let R restrict an unframed duplicated cocycle to view zero. Fix the unique boundary object. A marked realization has two pieces of data:

$$(r, \phi), \qquad r \in H, \quad \phi \in H.$$

The element ϕ is the chosen identification of the restricted object with the fixed boundary. It is not discarded just because the boundary has only one object: that object still has many automorphisms.

A comparison from (r, ϕ) to (r', ψ) is a gauge (a_0, a_1) satisfying both the cocycle equation and the boundary equation

$$r' = a_0 r a_1^{-1}, \qquad \psi a_0 = \phi.$$

These equations determine it uniquely:

$$a_0 = \psi^{-1}\phi, \qquad a_1 = (r')^{-1} a_0 r. \tag{15.2}$$

Substituting the formulas proves both equations without commuting any factors. Thus in this example there is exactly one identification-preserving arrow between any two marked realizations. The marked groupoid is nonempty and equivalent to the terminal groupoid, although its raw object set has many elements.

This is a worked two-view model of the marked transport mechanism. The generic groupoid construction and its equivalence theorem are Proposition 14.3 and Theorem 14.4. Formula (15.2) is checked in the independent finite suite for $H = S_3$, retaining the order of multiplication. This finite check is separate from the general formal theorem.

15.4 Four meanings of uniqueness

The example separates four statements that are often compressed into the word unique.

Statement	What it retains
One literal object	No different representative is present in the specified object set.
One isomorphism class	Different representatives may remain, as may their automorphisms.
A chosen realization	One object is selected, but no uniqueness follows from the choice.
One compatible comparison	The comparison respecting fixed identifications is unique, even when many representatives exist.

The unframed duplicated groupoid has one isomorphism class and nontrivial automorphisms. Its marked relative fiber has unique compatible comparisons. Its fully framed representative groupoid can instead have many mutually incomparable objects. These are three different constructions, not three descriptions of the same object with different terminology.

A descent theorem typically uses the last notion: two global realizations of the same specified local datum are related by a unique comparison respecting that datum. It does not force their underlying labels to be literally equal, or erase automorphisms after the local identifications have been forgotten.

> Boundary data are not only values. In a symmetry-sensitive problem they include the identification with those values. Retaining that identification is what allows a relative realization problem to

survive an equivalence of presentations.

16 Coherence need not mean removable phase

The preceding descent theory starts with transition data that already satisfy their composition laws. It is useful to see, in a completely finite calculation, how those laws can hold while a further simplification remains impossible. The following model separates *coherent composition* from *a choice of coordinates that removes every phase*.

The coefficient group and the degree of the data matter. The cocycles in Section 13 label individual arrows in a view category. Here a phase is attached to a *pair* of group elements. These are different constructions. The finite example is not obtained by renaming the earlier section obstruction as a general second cohomology group.

16.1 A two-bit phase law

Write $\mathbf{F}_2 = \mathbb{Z}/2\mathbb{Z}$ additively, and let

$$G = \mathbf{F}_2^2, \qquad x = (x_1, x_2), \qquad c(x, y) = x_2 y_1.$$

All entries and operations in this subsection lie in $\mathbf{F}_2$. In particular minus and plus agree, although retaining a minus sign in the coboundary formula helps explain its origin.

Proposition 16.1 (A normalized associative phase)**.** *The phase vanishes when either argument is zero and satisfies*

$$c(x, y) + c(x + y, z) = c(y, z) + c(x, y + z). \tag{16.1}$$

There is no function $a : G \to \mathbf{F}_2$ such that

$$c(x, y) = a(x) + a(y) - a(x + y) \quad \textit{for every } x, y \in G. \tag{16.2}$$

Proof. Normalization follows from multiplication by zero. Expanding the two sides of (16.1) gives, respectively,

$$x_2 y_1 + (x_2 + y_2) z_1, \qquad y_2 z_1 + x_2 (y_1 + z_1),$$

which are equal by distributivity.

Every expression $a(x) + a(y) - a(x + y)$ is symmetric in x, y, because G is abelian. With $e_1 = (1, 0)$ and $e_2 = (0, 1)$, however,

$$c(e_1, e_2) = 0, \qquad c(e_2, e_1) = 1.$$

This contradicts the symmetry required by (16.2). □

The proof rules out all sixteen possible scalar functions a, not just linear functions or selected trial gauges. The source `ProjectiveExample.lean` proves the normalization, cocycle, asymmetry, and no-scalar-flattening statements. The matrix interpretation and group construction below are worked derivations of the finite model, not separate formal declarations.

Phase table $c(x,y)=x_2y_1$ in $\mathbf{F}_2$

	0	e_1	e_2	e_1+e_2
0	0	0	0	0
e_1	0	0	0	0
e_2	0	1	0	1
e_1+e_2	0	1	0	1

$$c(e_1,e_2)=0 \text{ but } c(e_2,e_1)=1$$

Figure 13: The entry in row x and column y is $c(x,y)=x_2y_1$. The two entries singled out in the text, $c(e_1,e_2)=0$ and $c(e_2,e_1)=1$, already obstruct every scalar flattening. Associativity still holds for every triple.

16.2 The phase is visible in actual matrices

Over $\mathbb{R}$, define

$$X=\begin{pmatrix}0&1\\1&0\end{pmatrix}, \qquad Z=\begin{pmatrix}1&0\\0&-1\end{pmatrix}, \qquad U_{(p,q)}=X^pZ^q \quad (p,q\in\{0,1\}).$$

Direct multiplication gives $X^2=Z^2=I$ and $ZX=-XZ$. Moving Z^{x_2} past X^{y_1} therefore yields

$$U_xU_y=(-1)^{c(x,y)}U_{x+y}. \tag{16.3}$$

For example $U_{e_1}U_{e_2}=XZ$, while $U_{e_2}U_{e_1}=-XZ$. The two-bit labels add commutatively, but their chosen matrix lifts do not commute.

Associativity of three lifted matrices requires the same total sign on both parenthesizations. Equation (16.1) is exactly that requirement:

$$(U_xU_y)U_z=(-1)^{c(x,y)+c(x+y,z)}U_{x+y+z},$$

$$U_x(U_yU_z)=(-1)^{c(y,z)+c(x,y+z)}U_{x+y+z}.$$

Thus a nonzero phase is not automatically a failed associativity law. It can be the data that make the composition law explicit.

Changing lifts to $U'_x=(-1)^{a(x)}U_x$ changes the phase to

$$c'(x,y)=c(x,y)+a(x)+a(y)-a(x+y).$$

No such change makes c' vanish identically. The antisymmetric part is unchanged by every scalar change of lift. In this model it is

$$\kappa(x,y)=c(x,y)-c(y,x)=x_2y_1-y_2x_1.$$

This is a directly observable obstruction certificate for the model. Its nonzero value at (e_1, e_2) is enough, without a search over gauges.

The same symmetry argument rules out removing this sign multiplier by arbitrary nonzero complex scalar lifts as well. A multiplicative scalar coboundary $b(x)b(y)/b(x+y)$ on an abelian group is symmetric, whereas the multiplier in (16.3) is not. This last sentence is an algebraic consequence written here. The retained Lean example uses the stated $\mathbf{F}_2$-valued formulation.

16.3 Retaining the phase restores an ordinary multiplication law

Instead of trying to discard the phase, include it in the state. On the eight-element set $\widetilde{G} = \mathbf{F}_2 \times G$, put

$$(a, x) \cdot (b, y) = (a + b + c(x, y), x + y). \tag{16.4}$$

The first coordinate of the left-associated product of three elements is $a + b + d + c(x, y) + c(x + y, z)$. The right-associated coordinate is $a + b + d + c(y, z) + c(x, y + z)$. Equation (16.1) proves equality. Normalization gives identity $(0, 0)$, and

$$(a, x)^{-1} = (a + c(x, x), x).$$

Both inverse products are the identity because $x + x = 0$ and arithmetic is modulo two. Hence (16.4) defines a group.

Projection $(a, x) \mapsto x$ is a homomorphism onto G. Its kernel is $\{(0, 0), (1, 0)\}$, which commutes with every element. The map

$$(a, x) \longmapsto (-1)^a U_x$$

is a homomorphism by (16.3). Its eight matrices are distinct: they are $\pm I, \pm X, \pm Z, \pm XZ$. It is therefore faithful. The phase extension is not an assertion that G itself has acquired an ordinary representation with these same lifts. The represented group has changed.

A precise form of retaining the obstruction. The original four labels do not support a phase-free multiplication of the chosen matrices. Keeping one additional bit gives an eight-element multiplication that does. This repairs the description by retaining information, not by proving that the original phase was a coboundary.

16.4 How this example fits the realization calculus

There are three different questions in this small model. First, the phase satisfies the associativity relation. Second, scalar changes cannot trivialize it. Third, an enlarged state that retains the phase admits an ordinary multiplication law. Each question has a complete answer, but their answers are not interchangeable.

The example illustrates why an obstruction should be attached to the allowed changes and to a specified presentation problem. Enlarging the presentation can be mathematically productive, but it changes what is being realized and what information is carried. No general theorem that every obstruction can be removed by a finite enlargement is claimed. The finite checker enumerates all group triples, all scalar gauges, and all multiplication triples in $\widetilde{G}$, independently of the symmetry proof above.

Part IV

The geometry of exact realization

A categorical space of realizations does not come with a numerical distance by default. In a declared linear model, the compatibility equations become a bounded operator. Closed range, holonomy, and conditioning then determine what small residuals actually guarantee.

17 Graph views, holonomy, and reconstruction

17.1 Actual transports on a finite graph

Let a finite connected simple graph have vertex set V, root o, and a chosen route $p_i : o \to i$ for every vertex, with the route at o empty. Use a common real normed fiber H. An inverse consistent connection assigns continuous linear equivalences τ_{ij}, with $\tau_{ji} = \tau_{ij}^{-1}$. Our convention is that τ_{ij} carries the value at j into the coordinates at i.

The candidate and residual spaces are finite products with maximum norms:

$$E = \prod_{i \in V} H, \qquad R = \prod_{(i,j) \in \vec{E}} H, \qquad (dx)_{ij} = \tau_{ij} x_j - x_i.$$

Both orientations of an edge are retained. If M bounds the norms of all measured edge maps, then $\|d\| \le M + 1$. The empty edge case has the usual zero residual.

Distinct walks are retained as arrows of the view category. The product along a walk $q = (i_0, \dots, i_m)$ is

$$T_q = \tau_{i_0 i_1} \cdots \tau_{i_{m-1} i_m}.$$

Concatenation gives the restriction composition law. An exact vertex family defines a coherent walk-view section, and a coherent walk-view section satisfies every edge equation. Thus this graph construction returns to the original ambient realization type rather than creating a disconnected numerical model.

17.2 Transported residuals and path gains

Proposition 17.1 (The path telescope)**.** *For a walk $q = (i_0, \dots, i_m)$,*

$$T_q x_{i_m} - x_{i_0} = \sum_{k=0}^{m-1} (\tau_{i_0 i_1} \cdots \tau_{i_{k-1} i_k})(dx)_{i_k i_{k+1}}.$$

The empty prefix is the identity. If $\gamma(q)$ is the sum of the products of preceding edge norms, then

$$\left\| T_q x_{i_m} - x_{i_0} \right\| \le \gamma(q) \, \|dx\|_\infty .$$

Proof. Write $U_0 = I$ and $U_k = \tau_{i_0 i_1} \cdots \tau_{i_{k-1} i_k}$ for $1 \le k \le m$. The transport convention makes every U_k a map from the fiber at i_k to that at i_0. Since $U_{k+1} = U_k \tau_{i_k i_{k+1}}$, the k-th transported defect is

$$U_k (dx)_{i_k i_{k+1}} = U_{k+1} x_{i_{k+1}} - U_k x_{i_k}.$$

Summing over $k = 0, \dots, m-1$ cancels each interior value once with each sign. The remaining terms are $U_m x_{i_m} - U_0 x_{i_0} = T_q x_{i_m} - x_{i_0}$, proving the identity with its actual ordered products. For an empty walk both sides are zero.

The triangle inequality and the operator norm bound now give

$$\left\| T_q x_{i_m} - x_{i_0} \right\| \le \sum_{k=0}^{m-1} \|U_k\| \left\| (dx)_{i_k i_{k+1}} \right\| \le \left(\sum_{k=0}^{m-1} \prod_{j=0}^{k-1} \left\| \tau_{i_j i_{j+1}} \right\| \right) \|dx\|_\infty .$$

The bracket is $\gamma(q)$, with the empty product equal to one. Repeated vertices or edges do not affect the cancellation. Reordering the edge maps would, in general, invalidate it. □

For isometric maps, $\gamma(q) \le |q|$. For arbitrary bounded equivalences, route length cannot replace the amplification gain. Reversing a route can have a different gain, even though its transport is the inverse map.

17.3 Holonomy is the remaining exactness condition

Write $T_i = T_{p_i}$. Reconstruction from a proposed root value is

$$(R_p z)_i = T_i^{-1} z.$$

For each directed edge define the root loop operator

$$K_{ij} = T_i \tau_{ij} T_j^{-1}, \qquad K = \bigcap_{(i,j) \in \vec{E}} \ker(K_{ij} - I).$$

Each K_{ij} is the transport around an actual rooted loop. No flatness assumption has been made.

Theorem 17.2 (Holonomy classification)**.** *Evaluation at the root and reconstruction give an equivalence*

$$K \simeq_L \ker d.$$

A root value extends to an exact vertex family precisely when it lies in K, and that extension is unique. The space K is independent of the chosen routes. Reconstruction is route independent on K.

Proof. Every T_i is a continuous linear equivalence, being a finite ordered composition of such maps. Reconstruction is therefore well defined. For an edge (i, j), apply T_i to its defect:

$$T_i\left(\tau_{ij} T_j^{-1} z - T_i^{-1} z\right) = T_i \tau_{ij} T_j^{-1} z - z = K_{ij} z - z.$$

Because T_i is injective, that edge defect is zero if and only if the corresponding root loop fixes z. Thus $dR_p z = 0$ holds exactly for $z \in K$. The root route is empty, so $(R_p z)_o = z$.

Conversely, let $x \in \ker d$. The path telescope along p_i has all zero summands, hence $T_i x_i = x_o$ for every vertex. Applying T_i^{-1} gives $x_i = T_i^{-1} x_o$, so $x = R_p(x_o)$. The first part then places x_o in K. We have proved both inverse equations and uniqueness of an exact extension of a root value.

The maps are linear. On the finite vertex product,

$$\left\| R_p z \right\|_\infty \le \left(\max_i \left\| T_i^{-1} \right\| \right) \|z\| , \qquad \|x_o\| \le \|x\|_\infty .$$

These bounds also hold on the subspaces K and $\ker d$, proving continuity of both directions. No completeness is needed for these finite bounds.

For route independence, the proved equivalence identifies K with the image of root evaluation from the intrinsically defined space $\ker d$. That image does not refer to any chosen routes, so every route system yields the same K. For $z \in K$, reconstruction using either system is an exact family with root z, and uniqueness makes the families equal. For $z \notin K$, this uniqueness argument does not assert equality of the two generally inconsistent reconstructed families. □

With isometric transports the equivalence is an isometry for the vertex maximum norm, because every reconstructed coordinate has the root norm. With general bounded equivalences, it need not be an isometry.

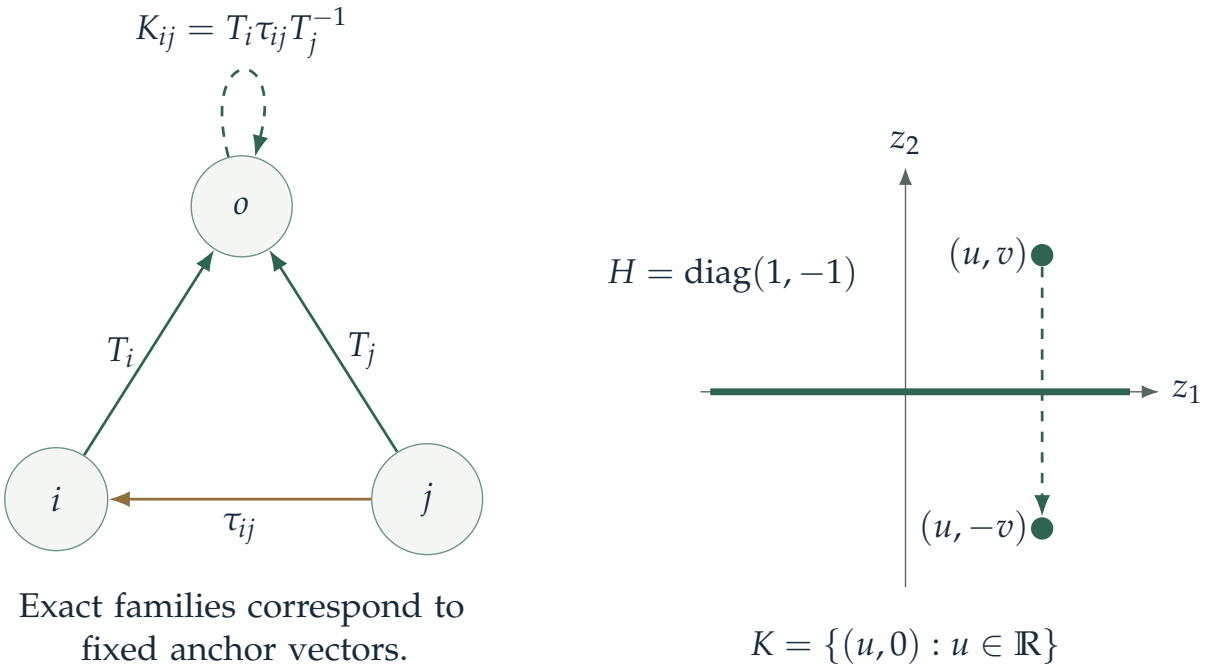


Figure 14: Tree routes reconstruct from an anchor. The remaining edges close loops and impose fixed-vector equations at that anchor. A nontrivial loop can obstruct an arbitrary seed without preventing every nonzero seed from extending. The fixed space, rather than the identity of the whole loop operator, is the relevant realization condition.

17.4 Prescribed observations remain visible

If a prescribed vertex set U contains o, a candidate x has an exact repair preserving $x|_U$ precisely when

$$x_o \in K \quad \text{and} \quad (R_p x_o)|_U = x|_U.$$

Necessity follows from the classification. Sufficiency is the reconstructed family itself. The source identifies this criterion with the earlier obstruction class in $R/d(\ker p_U)$. A graph seed test is therefore an instance of the same protected correction problem, not a replacement for it.

18 Closed range is the stability boundary

18.1 Projection exists before a residual estimate

Let E be a real Hilbert space and R a real Banach space. Let $D : E \to R$ be continuous linear. The kernel $K = \ker D$ is closed, so its orthogonal projection $q = P_K$ exists [4].

Proposition 18.1 (Exact orthogonal repair)**.** *For every $x \in E$,*

$$D(qx) = 0, \quad q^2 = q, \quad \|qx\| \le \|x\|,$$

and qx is the unique closest point to x in K.

Proof. Boundedness of D makes $K = \ker D$ closed. Lemma E.1 constructs $q = P_K$ and gives the orthogonal decomposition $x = qx + (x - qx)$. The first term belongs to K, so $Dqx = 0$. Projection fixes every point of K, giving $q(qx) = qx$.

Pythagoras applied to that decomposition gives $\|x\|^2 = \|qx\|^2 + \|x - qx\|^2$, hence $\|qx\| \le \|x\|$. For any $y \in K$, the vector $qx - y$ lies in K, whereas $x - qx$ lies in its orthogonal complement. Consequently

$$\|x - y\|^2 = \|x - qx\|^2 + \|qx - y\|^2 \ge \|x - qx\|^2.$$

Equality holds only when $\|qx - y\| = 0$, or $y = qx$. This proves nearest-point uniqueness without any closedness assumption on $\operatorname{ran} D$. Closedness of the kernel and of the range serve different purposes. □

No closed range assumption on D is needed for this projection. Such an assumption enters when the size of the residual is asked to control the distance to the exact space.

Theorem 18.2 (Closed range and quantitative exactness)**.** *The following are equivalent:*

1. $\operatorname{ran} D$ *is closed.*
2. *There is $C \ge 0$ with $\|x - qx\| \le C \|Dx\|$ for all x.*
3. *There is $C \ge 0$ such that every x has some exact y with $\|x - y\| \le C \|Dx\|$.*

When the range is closed, the least admissible constant is the operator norm of the inverse of $D : K^\perp \to \operatorname{ran} D$. In the zero operator case that inverse has zero dimensional domain and its norm is zero.

Proof. Let $K = \ker D$ and $K^\perp$ carry the induced Hilbert norm. Define $D_\perp u = Du$, but regard its target as $\operatorname{ran} D$ with the norm inherited from R. It is injective: $Du = 0$ puts u in both K and $K^\perp$, so $u = 0$. It is surjective onto that target: if $r = Dx$, then $u = x - qx \in K^\perp$ and $Du = Dx = r$. Thus it is a bounded linear bijection before any continuity of its inverse is claimed.

Suppose first that the range is closed. Because R is Banach, $\operatorname{ran} D$ is then Banach. The domain $K^\perp$ is also complete. The bounded inverse theorem, proved in Appendix E, gives a bounded inverse $A = D_\perp^{-1}$. Applying it to $Dx = D(x - qx)$ yields

$$x - qx = A(Dx), \qquad \|x - qx\| \le \|A\| \, \|Dx\|.$$

Hence condition 2 holds. If any $C \ge 0$ satisfies condition 2, then for every $r \in \operatorname{ran} D$ apply it to $u = Ar \in K^\perp$. Since $qu = 0$, it gives $\|Ar\| \le C \|r\|$. Taking the operator norm proves $\|A\| \le C$. This establishes the least constant, not just a sufficient one. If $D = 0$, both $K^\perp$ and the range are zero spaces and $\|A\| = 0$.

Now assume condition 2. To prove closedness, take an arbitrary sequence $r_n \in \operatorname{ran} D$ with $r_n \to r$ in R. Choose x_n with $Dx_n = r_n$ and put $u_n = x_n - qx_n$. Then $u_n \in K^\perp$ and $Du_n = r_n$. Applying the estimate to $u_n - u_m$ gives

$$\|u_n - u_m\| \le C \|r_n - r_m\| .$$

Thus (u_n) is Cauchy in the complete space $K^\perp$. Let its limit be u. Continuity gives $Du = \lim_n r_n = r$, so $r \in \operatorname{ran} D$. In a normed space, a set containing the limits of all its convergent sequences is closed. This proves condition 1 without presupposing that the initially unbounded inverse is continuous.

Condition 2 implies condition 3 by choosing $y = qx$. Conversely, if condition 3 supplies an exact y for a given x, the nearest-point property of qx gives

$$\|x - qx\| \le \|x - y\| \le C \|Dx\| .$$

The same C works for all x, proving condition 2. This last implication does not require the selected answers y to depend linearly or continuously on x. □

The last equivalence is important. A nonlinear choice of exact answers cannot bypass the failure of a uniform residual bound. Failure of closed range is not merely a defect of one selected projection algorithm.

18.2 A small residual far from exactness

On $\ell^2(\mathbb{N})$, index the standard basis from one and define $De_n = n^{-1}e_n$. The map is bounded and injective, so its exact space is zero. Yet

$$\|e_n\| = 1, \qquad \|De_n\| = 1/n \longrightarrow 0.$$

No finite C can bound $\operatorname{dist}(e_n, \ker D)$ by $C \|De_n\|$. The range is not closed. For example finite truncations of $(1/n)_{n\ge1}$ belong to the range and converge in ℓ^2, but a preimage of the limit would be the constant sequence, which does not belong to ℓ^2.

This diagonal model is formalized on the infinite Hilbert space used in Section 30. In the namespace `KomorebiSpace.DiagonalRegularization`, `T_not_closed_range` applies the general obstruction theorem to actual unit coordinate vectors. The declaration `w_mem_closure_not_range` proves that the displayed limiting datum belongs to the closure of the range but not to the range itself. Formal coordinate j represents the printed coordinate $n = j + 1$.

18.3 Norm zero versus algebraic zero

If an additive correction image is not closed, distance to that image can vanish without actual membership. Thus an infimum of residual error can be zero while the exact protected correction is unattained. The algebraic obstruction quotient and its possible seminorm answer different questions. One must identify a closed image, compactness condition, or other attainment principle before replacing one question by the other.

18.4 Stable repair with protected data

The preceding orthogonal projection does not generally preserve a separately declared observable. To join the analytic theory to Section 4, let $p : E \to S$ be continuous linear into a normed space. Define

$$M = \ker p, \qquad B = D|_M : M \to R, \qquad H = M \cap \ker D.$$

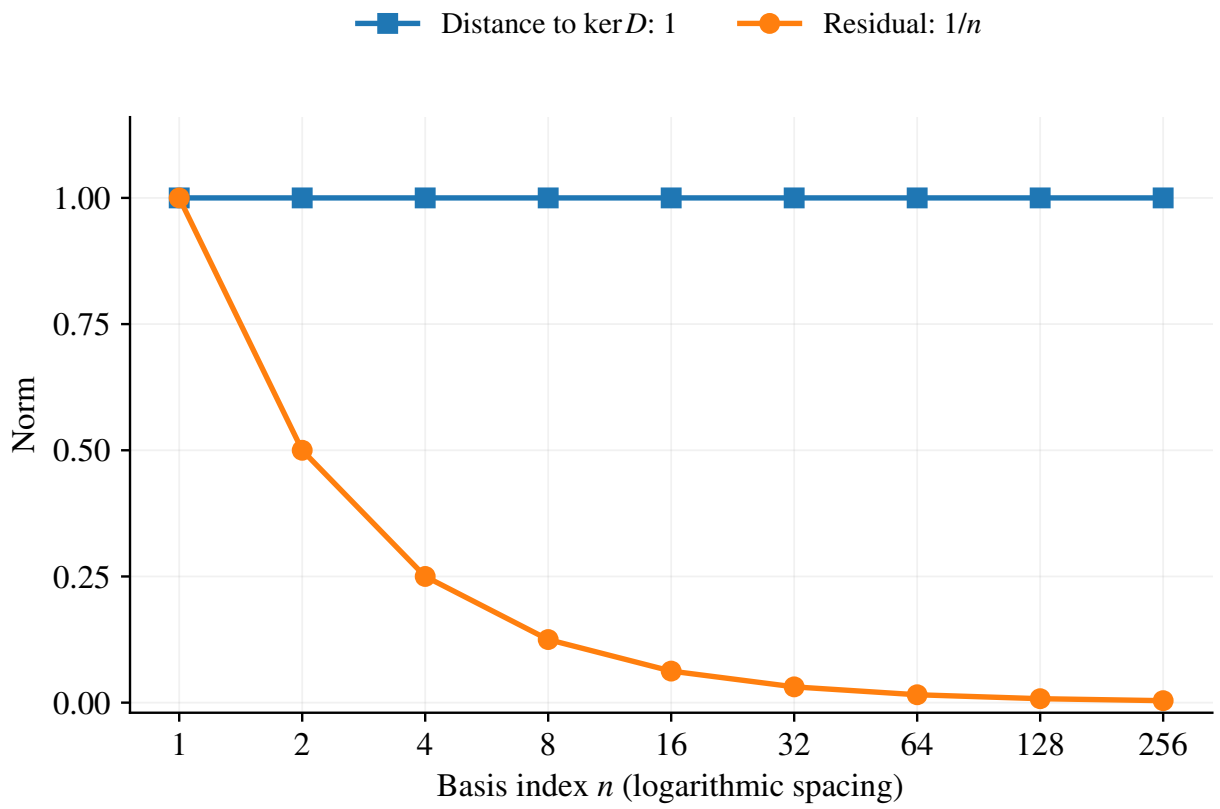


Figure 15: Residual size controls distance to the exact space only when transverse directions have a positive lower bound. The diagonal example sends unit vectors to residuals tending to zero, so no uniform inverse gain exists. Exact projection still exists.

The closed subspace M inherits its Hilbert structure from E. A candidate x is *feasible* for protected repair when an exact y with $p(y) = p(x)$ exists. The algebraic criterion is still $Dx \in \operatorname{ran} B$.

Theorem 18.3 (Nearest feasible protected repair)**.** *Every feasible x has a unique closest exact answer y_* preserving $p(x)$. Its correction $h_* = x - y_*$ is the unique vector in the orthogonal complement of* $\ker B$ *inside M satisfying $Bh_* = Dx$. For every other feasible answer y,*

$$\|x - y\|^2 = \|x - y_*\|^2 + \|y_* - y\|^2 .$$

Moreover, the following conditions are equivalent:

$$\operatorname{ran} B \text{ is closed} \quad \Longleftrightarrow \quad \exists C \geq 0 \; \forall \text{ feasible } x : \|x - y_*\| \leq C \, \|Dx\| .$$

When these conditions hold, the least nonnegative constant is the transverse inverse norm of B, including the zero-range case.

Proof. Feasibility gives some $h_0 = x - y_0 \in M$ with $Bh_0 = Dx$. Conversely, every $h \in M$ satisfying $Bh = Dx$ produces an exact protected answer $y = x - h$. Thus the answer problem and the correction problem are equivalent, with correction set

$$\{h \in M : Bh = Dx\} = h_0 + \ker B.$$

This is a nonempty closed affine subset of the Hilbert space M. Closedness of $\ker B$ follows from continuity of B, independently of whether its image is closed.

Let $k_0 = P^M_{\ker B} h_0$ and $h_* = h_0 - k_0$. Then $Bh_* = Bh_0 = Dx$ and $h_* \perp \ker B$ in M. If a different initial correction is chosen, it differs from h_0 by an element of $\ker B$. Its orthogonal-complement component is therefore again h_*. This proves that the construction is independent of the feasibility witness.

Every other correction can be written uniquely as $h_* + k$ with $k \in \ker B$. Orthogonality gives

$$\|h_* + k\|^2 = \|h_*\|^2 + \|k\|^2 .$$

Writing $y_* = x - h_*$ and $y = x - h_* - k$, we have $x - y = h_* + k$ and $y_* - y = k$. Substitution proves the stated Pythagorean identity. Its last term vanishes only when $y = y_*$, establishing the unique closest answer.

If $\operatorname{ran} B$ is closed, Theorem 18.2 applied on M gives a bounded inverse of B on $(\ker B)^{\perp_M}$. The vector h_* is exactly that inverse applied to Dx, so its norm obeys the required bound.

For the converse, suppose the bound holds for every feasible raw candidate. Every $x \in M$ is feasible, because $y = 0$ has $Dy = 0$ and $py = px = 0$. Its closest protected exact answer is $P^M_{\ker B} x$. Thus the assumed estimate reduces on M to

$$\left\| x - P^M_{\ker B} x \right\| \le C \, \|Bx\| \qquad (x \in M).$$

Theorem 18.2 forces $\operatorname{ran} B$ to be closed. The same restriction also proves that no smaller nonnegative constant than the transverse inverse norm can work. If the range is zero, feasible candidates are already exact under permitted correction and the least correction is zero, so the optimal gain is zero. □

The relevant stability boundary is therefore $D(\ker p)$, not automatically $D(E)$. Restricting permitted changes can turn a well-conditioned unconstrained problem into an obstructed or poorly conditioned protected one. The nearest protected answer is generally not $P_{\ker D} x$.

For $D(x_1, x_2, x_3) = x_1 - x_2$, $p(x) = (x_1, x_3)$, and $x = (2, 0, 5)$, the ordinary orthogonal answer is $(1, 1, 5)$. The protected answer is $(2, 2, 5)$. The respective squared errors are two and four. Both have zero residual, but only the latter preserves the declared data. The theorem is formalized in `ProtectedHilbertRepair.lean`, using the actual operator on $\ker p$. The displayed three-coordinate values are an illustrative calculation, not a separate universal algorithm.

19 Holonomy repair and transport amplification

19.1 An orthogonal anchor, not a global least squares answer

Return to the graph model with a real Hilbert fiber. The root residual is

$$\mathcal{H} z = (K_{ij} z - z)_{ij}, \qquad K = \ker \mathcal{H}.$$

Let $q = P_K$, and define

$$Px = R_p(q x_o).$$

This is always exact, fixes exact families, and is idempotent. It is the unique exact family preserving the feasible component $q x_o$ of the anchor. Its orthogonality is on the anchor fiber. The vertex and edge norms are still maximum norms.

For example on a two-vertex identity connection, $(0, 2)$ with root value zero is repaired to $(0, 0)$. The constant family $(1, 1)$ is closer in the maximum norm, but it does not preserve the anchor. The operator P is not being claimed to be a global nearest-point projection in that norm.

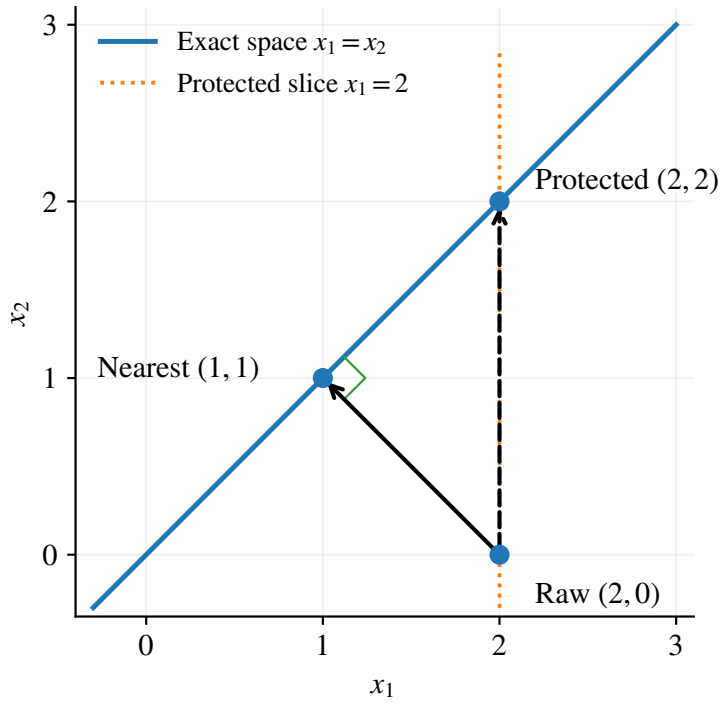


Figure 16: The same compatibility equation permits different repair problems. Orthogonal repair chooses the nearest point on the diagonal. Preserving the first coordinate forces a vertical correction to a different point. The third coordinate is fixed at five throughout.

19.2 Every factor in the bound has a meaning

Define the route gain

$$A_p = \max_i \gamma(\overline{p_i}),$$

using reverse routes so that the error is measured in each destination view. The path telescope gives

$$\left\| x - R_p(x_o) \right\|_\infty \le A_p \left\| dx \right\|_\infty .$$

Let N_p be the residual normalization $(N_p r)_{ij} = T_i r_{ij}$. The exact identity $\mathcal{H} = N_p d R_p$ implies

$$\left\| \mathcal{H}(x_o) \right\|_\infty \le \left\| N_p \right\| \left(\left\| d \right\| A_p + 1 \right) \left\| dx \right\|_\infty .$$

Theorem 19.1 (Amplified holonomy repair)**.** *If* ran $\mathcal{H}$ *is closed and* κ *is its transverse inverse gain, then*

$$\left\| x - Px \right\|_\infty \le \left[A_p + \left\| R_p \right\| \kappa \left\| N_p \right\| \left(\left\| d \right\| A_p + 1 \right) \right] \left\| dx \right\|_\infty . \tag{19.1}$$

The range condition is necessary for the existence of any uniform bound of distance to an exact family by $\left\| dx \right\|_\infty$*. It is automatic for a fixed finite dimensional fiber, but not uniformly quantitative across all connections.*

Proof. First quantify the residual seen at the root. The identity $\mathcal{H} = N_p d R_p$, applied to x_o, gives

$$\begin{aligned} \left\| \mathcal{H}(x_o) \right\|_\infty &= \left\| N_p \big(d(R_p x_o - x) + dx \big) \right\|_\infty \\ &\le \left\| N_p \right\| \left(\left\| d \right\| \left\| R_p x_o - x \right\|_\infty + \left\| dx \right\|_\infty \right) \\ &\le \left\| N_p \right\| \left(\left\| d \right\| A_p + 1 \right) \left\| dx \right\|_\infty . \end{aligned}$$

Closed range of $\mathcal{H}$ lets Theorem 18.2 control the anchor correction by $\|x_o - qx_o\| \le \kappa \|\mathcal{H}(x_o)\|_\infty$. Insert the uncorrected reconstruction between x and $Px = R_p q x_o$:

$$\|x - Px\|_\infty \le \|x - R_p x_o\|_\infty + \|R_p\| \, \|x_o - qx_o\| .$$

Substituting the two preceding bounds yields (19.1). The finite maximum-norm residual product is a Banach space because its fiber is Hilbert, so the target completeness needed in Theorem 18.2 is available.

For necessity, suppose a constant C satisfies $\operatorname{dist}(w, \ker d) \le C \|dw\|_\infty$ for every vertex family w. Fix $z \in K^\perp$ and take $w = R_p z$. Every exact family e has $e_o \in K$, and root evaluation has norm at most one. Therefore

$$\|z\| = \operatorname{dist}(z, K) \le \|z - e_o\| \le \|R_p z - e\|_\infty .$$

Taking the infimum over $e \in \ker d$ is legitimate without assuming that this global nearest point is attained. It gives

$$\|z\| \le C \|dR_p z\|_\infty \le C \left\|N_p^{-1}\right\| \|\mathcal{H} z\|_\infty .$$

Here N_p is a continuous linear equivalence on the finite edge product because every component T_i is one. For a general root vector, remove its projection onto K; this leaves the same $\mathcal{H}$-image and puts it in $K^\perp$. Thus the displayed estimate is a uniform residual bound of the kind in Theorem 18.2, which forces $\operatorname{ran} \mathcal{H}$ to be closed. A finite-dimensional fiber makes that image finite dimensional and hence closed. It does not give a uniform numerical value of the inverse gain across different connections. □

For isometric transports, one can use $\|R_p\| \le 1$, $\|N_p\| \le 1$, $\|d\| \le 2$, and $A_p \le r$, where r is the route radius. The bound reduces to

$$\|x - Px\|_\infty \le [r + \kappa(2r + 1)] \, \|dx\|_\infty .$$

The exact repair is route independent for a fixed root and inner product. A chosen upper bound involving route gains need not have the same numerical value for different routes.

19.3 Flatness does not remove amplification

Consider a scalar chain with $n + 1$ vertices, where $n \ge 1$, and coordinates $e_i(z) = 2^i z$. Set $\tau_{ij} = e_i e_j^{-1}$. This connection is flat, but reverse transport along one edge multiplies norms by two. Let

$$x_i = 8 \cdot 2^i - 1, \qquad y_i = 7 \cdot 2^i .$$

The family y is exact and preserves $x_o = 7$. On a forward edge the residual of x is $1/2$, and on its reverse it is -1. Thus

$$\|dx\|_\infty = 1, \qquad \|x - Px\|_\infty = 2^n - 1 .$$

For $n \ge 2$, a route length bound of n fails. This is amplification with trivial holonomy, distinct from poor conditioning of the fixed space of a nearly trivial loop.

19.4 Small holonomy can still be poorly conditioned

In $\mathbb{R}^2$, the rational rotation

$$U_t = \frac{1}{1 + t^2} \begin{pmatrix} 1 - t^2 & -2t \\ 2t & 1 - t^2 \end{pmatrix}$$

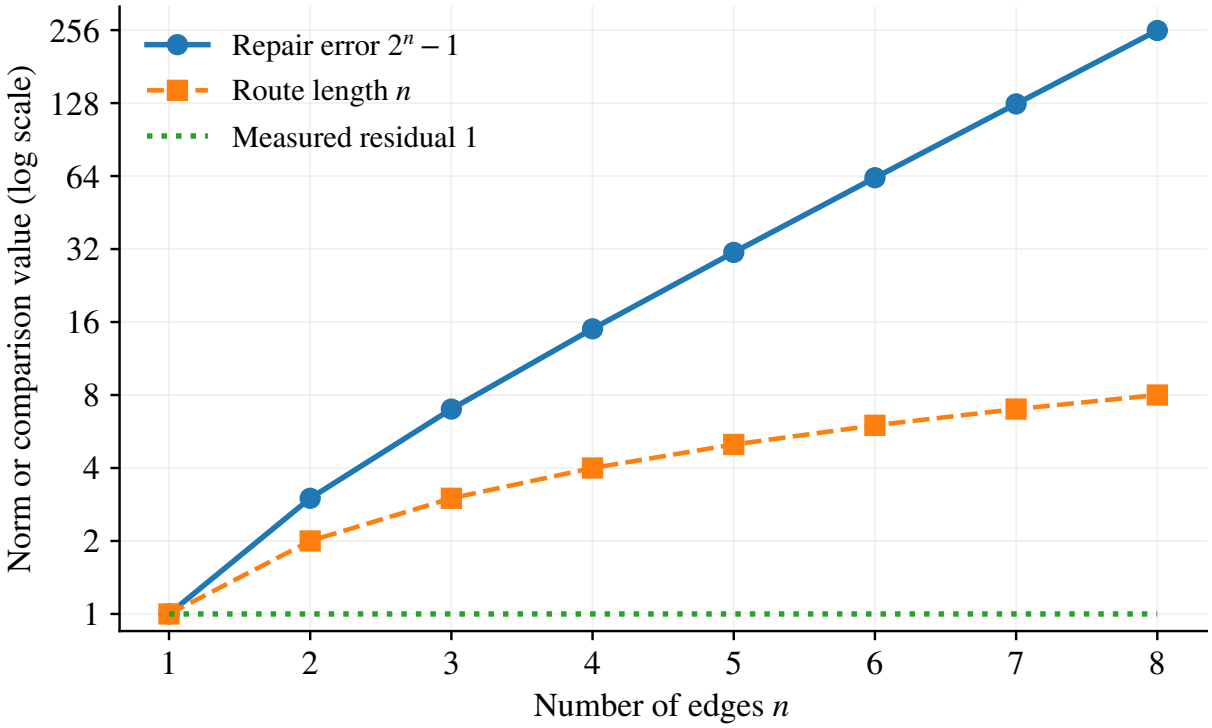


Figure 17: A flat connection can still amplify a fixed residual. Both orientations of every edge are measured, giving residual norm one. Root-preserving repair has error $2^n - 1$, which exceeds the route-length estimate once $n \geq 2$. The vertical scale is logarithmic.

satisfies

$$\|(U_t - I)z\|^2 = \frac{4t^2}{1+t^2}\,\|z\|^2 .$$

For $t \neq 0$ its fixed space is zero and its residual gain is $\sqrt{1+t^2}/(2|t|)$, which diverges as $t \to 0$. At $t = 0$, the entire fiber is fixed. Therefore small loop discrepancy does not supply a uniform stable realization estimate. The fixed directions and their conditioning must be analyzed together.

20 Quantitative descent in a common Hilbert model

20.1 Local exactness and matching are separate operators

Let E be a real Hilbert space of raw local candidates. Let $D : E \to R$ and $C : E \to Y$ be continuous linear maps into Banach spaces, measuring local equation defects and overlap mismatch, respectively. Put

$$L = \ker D, \qquad B = C|_L, \qquad M = \ker D \cap \ker C.$$

The exact matched families form M. First project to L, then project inside L to $\ker B$. Denote the resulting map by P.

Theorem 20.1 (The two-stage repair is the joint projection)**.** *The operator P is the orthogonal projection onto M, equivalently onto the kernel of $\Delta = (D, C)$. If $u = x - P_L x$ and $v = P_L x - Px$, then*

$$u \perp v, \qquad \|x - Px\|^2 = \|u\|^2 + \|v\|^2 .$$

Proof. The subspace $L = \ker D$ is closed, hence Hilbert with the inherited inner product. The restriction $B = C|_L$ is continuous, so its kernel is a closed subspace of L. Under the inclusion into E, this kernel is exactly $M = L \cap \ker C$. Thus both projections used to define P exist.

For an input x, write $\ell = P_L x$ and let Px be the projection of ℓ onto $\ker B$ inside L. Then $Px \in M$, $u = x - \ell \in L^\perp$, and $v = \ell - Px \in L$ is orthogonal to M. Since $M \subseteq L$, the vector u is also orthogonal to M. Their sum $x - Px = u + v$ is therefore orthogonal to M. A vector in M with that error is the unique orthogonal projection of x, so $P = P_M$.

The two errors themselves are orthogonal because $u \in L^\perp$ and $v \in L$. Expanding the squared norm of their sum gives $\|x - Px\|^2 = \|u\|^2 + \|v\|^2$. Finally $\ker(D, C) = \ker D \cap \ker C = M$, so this is also the projection onto the kernel of the joint defect. The proof uses a second projection *inside* L. Replacing it by the ambient projection onto $\ker C$ would invalidate the membership step and need not preserve local exactness. □

This identity does not require the local and overlap operators to commute. The second projection is taken in the already exact subspace, not in the whole raw space without restriction.

For a concrete failure of the unrestricted procedure, take $L = \{(a, 0)\}$ and $\ker C = \{(a, a)\}$ in $\mathbb{R}^2$. Their intersection is zero. Starting from $(1, 0)$, projection to L does nothing, while subsequent projection to $\ker C$ gives $(1/2, 1/2)$, which is no longer in L. The restricted second stage instead gives the correct joint answer zero.

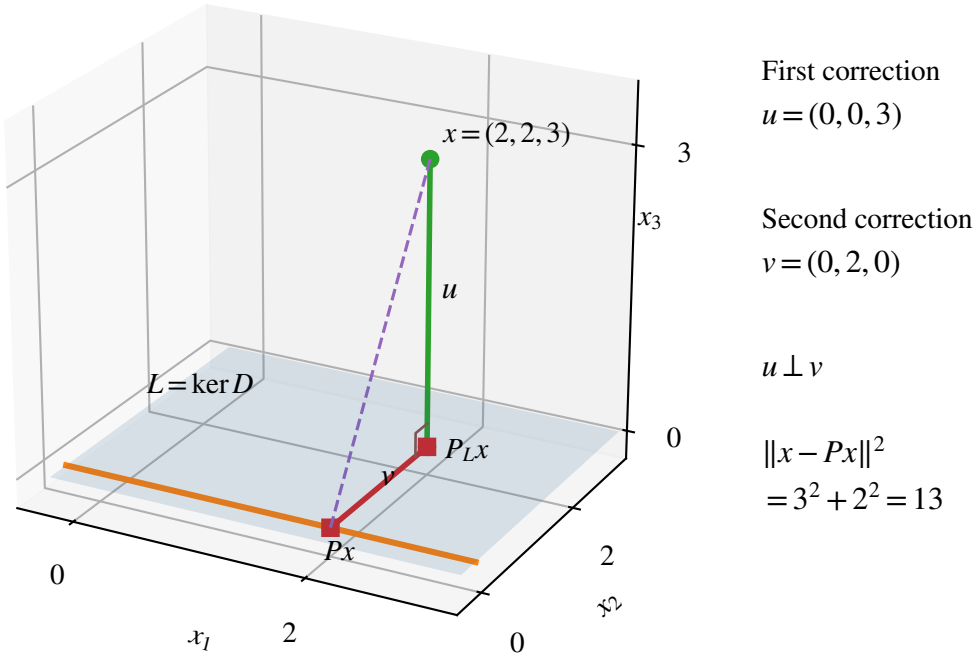


Figure 18: Local correction and overlap correction are orthogonal errors when the second problem is restricted to the locally exact subspace. Their composition is the projection onto the joint exact space. The overlap range inside that subspace is the additional stability obligation.

20.2 A stable repair requires a stable matching problem

Suppose ran D and ran B are closed, with residual gains a, b. Then

$$\|u\| \le a\,\|Dx\|, \qquad \|v\| \le b\,\|CP_L x\| \le b(\|Cx\| + \|C\|\ a\,\|Dx\|).$$

Therefore

$$\|x - Px\| \le a(1 + b\,\|C\|)\,\|Dx\| + b\,\|Cx\|. \tag{20.1}$$

The Pythagorean identity gives the sharper estimate

$$\|x - Px\|^2 \le a^2\,\|Dx\|^2 + b^2\big(\|Cx\| + a\,\|C\|\,\|Dx\|\big)^2.$$

For the maximum norm on $R \times Y$, two sufficient gains are therefore

$$G_2 = \sqrt{a^2 + b^2(1 + a\,\|C\|)^2} \le G_1 = a(1 + b\,\|C\|) + b.$$

The inequality follows from $\sqrt{u^2 + v^2} \le u + v$ for nonnegative u, v. The formal interface supplies the sum-type gain G_1, while its squared-error theorem also yields G_2. Neither is claimed to be the least joint gain. In either bound, $\|C\|$ may be replaced by any certified bound on C restricted to $(\ker D)^\perp$.

Theorem 20.2 (Closed range for the joint defect)**.** *Assume* ran D *is closed. Then*

$$\operatorname{ran}(D, C) \text{ is closed} \quad \Longleftrightarrow \quad \operatorname{ran}(C|_{\ker D}) \text{ is closed}.$$

Under either condition, a finite joint residual-to-repair bound exists.

Proof. Suppose the range of the joint map is closed. To test closedness of ran B, take a sequence $Bk_n \to y$ with $k_n \in \ker D$. The joint images are

$$(Dk_n, Ck_n) = (0, Bk_n) \longrightarrow (0, y).$$

Closedness supplies $x \in E$ with $(Dx, Cx) = (0, y)$. The first coordinate gives $x \in \ker D$, and the second gives $Bx = y$. Hence the matching range is closed. This direction did not require closedness of ran D.

For the reverse direction, assume both ran D and ran B are closed, and take joint images converging to (r, y). Thus $Dx_n \to r$ and $Cx_n \to y$. Decompose each input orthogonally as

$$x_n = u_n + k_n, \quad u_n \in (\ker D)^\perp, \quad k_n \in \ker D.$$

Because $r \in \operatorname{ran} D$ and the transverse inverse A of D is bounded by Theorem 18.2, the first components determine

$$u_n = A(Dx_n) \longrightarrow A(r) = u, \qquad Du = r.$$

Continuity of C now gives

$$Bk_n = Ck_n = Cx_n - Cu_n \longrightarrow y - Cu.$$

Closedness of ran B supplies $k \in \ker D$ with $Bk = y - Cu$. There is no need for the sequence k_n itself to converge. The single vector $u + k$ satisfies

$$D(u + k) = r, \qquad C(u + k) = Cu + Bk = y.$$

Thus the joint limit lies in the joint range, proving closedness. The target product is Banach under its maximum norm, so Theorem 18.2 applies to the joint map and gives a finite repair gain. The earlier two-stage estimate supplies an explicit, though not necessarily least, value. □

This Hilbert model is a quantitative counterpart of gluing, but it is not a norm on an arbitrary categorical stack. It starts from a fixed raw candidate space, fixed norms, and actual local and overlap operators. Changing those choices requires its own comparison.

21 Controlled motion of exact spaces

21.1 A two-sided condition number theorem

Let $D, D' : E \to R$ be continuous linear maps from a real Hilbert space to the same normed target. Write $q = P_{\ker D}$, $q' = P_{\ker D'}$. Suppose

$$\|(I-q)x\| \le C\,\|Dx\|, \qquad \|(I-q')x\| \le C'\,\|D'x\|, \qquad C, C' \ge 0,$$

and $\|D - D'\| \le \varepsilon$.

Theorem 21.1 (Motion of the exact projector)**.** *Without either kernel inclusion as a hypothesis,*

$$\|q - q'\| \le \max(C, C')\varepsilon. \tag{21.1}$$

In particular a uniform residual gain M makes $D \mapsto P_{\ker D}$ M-Lipschitz on the class of operators carrying that gain.

Proof. For $z \in \ker D'$, the old residual bound gives

$$\|(I-q)z\| \le C\,\|Dz\| = C\,\|(D-D')z\| \le C\varepsilon\,\|z\|.$$

Interchanging the roles gives

$$\|(I-q')w\| \le C'\varepsilon\,\|w\| \qquad (w \in \ker D). \tag{*}$$

We need a dual form of (*). For $v \in (\ker D')^\perp$ and any $y \in E$, self-adjointness of q and $v = (I-q')v$ give

$$|\langle qv, y\rangle| = |\langle v, qy\rangle| = |\langle v, (I-q')qy\rangle| \le C'\varepsilon\,\|v\|\,\|qy\| \le C'\varepsilon\,\|v\|\,\|y\|.$$

The Hilbert norm equals the supremum of these pairings on the unit ball, so $\|qv\| \le C'\varepsilon\,\|v\|$. This proves the precise restriction estimate needed on $(I-q')x$, not just a coarser bound in terms of $\|x\|$.

Now write

$$(q-q')x = q(I-q')x - (I-q)q'x.$$

The two displayed terms lie in $\ker D$ and $(\ker D)^\perp$, respectively, so they are orthogonal. The input components $(I-q')x$ and $q'x$ are also orthogonal. Applying the dual estimate to the first and the original cross estimate to the second gives

$$\begin{aligned}
\|(q-q')x\|^2 &= \|q(I-q')x\|^2 + \|(I-q)q'x\|^2 \\
&\le (C'\varepsilon)^2\,\|(I-q')x\|^2 + (C\varepsilon)^2\,\|q'x\|^2 \\
&\le \varepsilon^2 \max(C, C')^2\big(\|(I-q')x\|^2 + \|q'x\|^2\big) \\
&= \varepsilon^2 \max(C, C')^2\,\|x\|^2.
\end{aligned}$$

The bound $\|D - D'\| \le \varepsilon$ implies $\varepsilon \ge 0$, so taking square roots yields the operator estimate. If the common gain is M, substitute $C = C' = M$ and $\varepsilon = \|D - D'\|$. This is exactly the Lipschitz inequality on the class carrying that gain. The use of two orthogonal decompositions explains the maximum coefficient instead of a sum. □

There is no smallness threshold when both certificates are already supplied. The estimate can be uninformative when their gains are large. Knowing only the old certificate is a different situation.

For changing input as well as model, the same theorem gives

$$\|qx - q'y\| \le \|x - y\| + \max(C, C')\varepsilon\,\|y\|.$$

It also bounds the change in distance to the exact space by $\max(C, C')\varepsilon\,\|x\|$.

21.2 Constructing the new certificate

Assume the target R is Banach and $A : R \to E$ is a continuous linear right inverse of D, with $DA = I$, $\|A\| \le a$, and $a\varepsilon < 1$. Set

$$T = (D - D')A, \qquad Z = (I - T)^{-1}, \qquad A' = AZ.$$

The inverse is supplied by the convergent Neumann series because $\|T\| \le a\varepsilon < 1$.

Theorem 21.2 (Stable surjectivity and derived repair)**.** *The perturbed map is surjective, with*

$$D'A' = I, \qquad \|A'\| \le \frac{a}{1 - a\varepsilon}.$$

Its exact projection satisfies

$$\|x - q'x\| \le \frac{a}{1 - a\varepsilon} \|D'x\|, \qquad \|q - q'\| \le \frac{a\varepsilon}{1 - a\varepsilon}.$$

Proof. The map $T = (D - D')A$ acts on the Banach target R and satisfies $\|T\| \le a\varepsilon < 1$. The operator series

$$Z = \sum_{n=0}^{\infty} T^n$$

converges in operator norm and obeys $\|Z\| \le (1 - a\varepsilon)^{-1}$. Multiplying its finite partial sums on either side by $I - T$ gives $I - T^{N+1}$. Passing to the operator-norm limit proves $(I - T)Z = Z(I - T) = I$. Appendix E supplies the full series argument, including completeness of the operator space.

Since $DA = I$, one has $D'A = DA - (D - D')A = I - T$. Therefore $A' = AZ$ satisfies

$$D'A' = (I - T)Z = I, \qquad \|A'\| \le \|A\| \, \|Z\| \le \frac{a}{1 - a\varepsilon}.$$

The first equation explicitly supplies a preimage of every target vector, proving surjectivity.

For any x, put $y = x - A'D'x$. Then $D'y = D'x - D'A'D'x = 0$, so y is a valid exact competitor for orthogonal projection. Its error gives

$$\|x - q'x\| \le \|x - y\| = \|A'D'x\| \le \frac{a}{1 - a\varepsilon} \|D'x\|.$$

The same argument using A gives $\|x - qx\| \le a \|Dx\|$. Both coefficients are nonnegative and $a/(1 - a\varepsilon) \ge a$. Substitution into Theorem 21.1 yields

$$\|q - q'\| \le \frac{a\varepsilon}{1 - a\varepsilon}.$$

No assertion that $I - A'D'$ is orthogonal was needed. It supplies an exact comparison vector, while best approximation supplies the orthogonal bound. □

If a measured residual y satisfies $\|D'x - y\| \le \eta$, the distance to the new exact space is at most $a(\|y\| + \eta)/(1 - a\varepsilon)$. The uncertainty radius remains a quantity to justify in the application. The theorem does not estimate it statistically.

For a surjective Hilbert-source operator, the right inverse need not be a separate unexplained input. The inverse of $D : (\ker D)^{\perp} \to R$, followed by inclusion, constructs it. A known residual certificate bounds its norm. Surjectivity is still substantive, and cannot be assumed for a redundant residual map whose target contains unattainable directions.

21.3 An equivalence between the moving kernels

Theorem 21.3 (Explicit kernel transport)**.** *Under the hypotheses of Theorem 21.2,*

$$J : \ker D \simeq_L \ker D', \qquad Jx = x - A'D'x,$$

with inverse $y \mapsto y - ADy$. Moreover

$$\|Jx - x\| \le \frac{a\varepsilon}{1 - a\varepsilon} \|x\|.$$

Proof. Let $R_0 = I - AD$ and $R_1 = I - A'D'$. From $DA = D'A' = I$ we obtain $DR_0 = 0$ and $D'R_1 = 0$. Each R_i fixes its corresponding kernel. Thus R_1 restricts to a map from $\ker D$ to $\ker D'$, while R_0 restricts in the other direction.

The inverse identities require two additional cancellations. First,

$$R_0A' = (I - AD)AZ = (A - ADA)Z = 0.$$

Second, using $D'A = I - T$ and $Z(I - T) = I$,

$$R_1A = A - AZD'A = A - AZ(I - T) = 0.$$

For $x \in \ker D$, these identities give

$$R_0R_1x = R_0(x - A'D'x) = R_0x = x.$$

For $y \in \ker D'$, they give

$$R_1R_0y = R_1(y - ADy) = R_1y = y.$$

Both restricted maps are bounded and linear as restrictions of bounded linear operators. They are therefore inverse continuous linear maps, establishing the asserted equivalence.

Finally, if $Dx = 0$, then $D'x = (D' - D)x$. Hence

$$\|Jx - x\| = \|A'D'x\| \le \|A'\| \varepsilon \|x\| \le \frac{a\varepsilon}{1 - a\varepsilon} \|x\|.$$

This derivation uses the selected right inverse A throughout. It proves a controlled transport of exact vectors, not a claim that the selected map is the nearest-point projection onto the new kernel. □

The map J depends on the selected splitting. It is not generally the restriction of the orthogonal projection q'. Equivalence of exact spaces and nearest-point repair answer different questions.

21.4 An all-real worked model

Let $D_t(x, y) = x + ty$. Its kernel is the line $\{(-tb, b) : b \in \mathbb{R}\}$, and its orthogonal projector is

$$q_t = \frac{1}{1 + t^2} \begin{pmatrix} t^2 & -t \\ -t & 1 \end{pmatrix}. \tag{21.2}$$

The common right inverse $r \mapsto (r, 0)$ supplies gain at most one for every t, and $\|D_s - D_t\| = |s - t|$. The general theorem gives $\|q_s - q_t\| \le |s - t|$.

Proposition 21.4 (Exact motion of the line)**.** *For every real t,*

$$\|(q_t - q_0)x\|^2 = \frac{t^2}{1+t^2}\|x\|^2, \qquad \|q_t - q_0\| = \frac{|t|}{\sqrt{1+t^2}}.$$

In the small perturbation interface, kernel transport sends $(0, b)$ to $(-tb, b)$. Orthogonal projection instead sends it to $(-tb/(1+t^2), b/(1+t^2))$.

Proof. The vector $v_t = (-t, 1)$ spans $\ker D_t$ and has squared norm $1 + t^2 > 0$. Projection onto this line is

$$x \longmapsto \frac{\langle x, v_t\rangle}{1+t^2} v_t,$$

which gives the matrix in (21.2). In particular it is the orthogonal projector, not merely a map whose image is contained in the kernel.

Subtract $q_0 = \operatorname{diag}(0, 1)$, obtaining the symmetric matrix

$$B_t = q_t - q_0 = \frac{1}{1+t^2}\begin{pmatrix} t^2 & -t \\ -t & -t^2 \end{pmatrix}.$$

The two off-diagonal entries of its square are zero, since $-t^3 + t^3 = 0$. Both diagonal entries are $(t^4 + t^2)/(1+t^2)^2 = t^2/(1+t^2)$. Thus

$$B_t^{\mathsf{T}} B_t = B_t^2 = \frac{t^2}{1+t^2} I.$$

Pairing this identity with an arbitrary vector proves the all-vector norm formula. It bounds the operator norm above by $|t|/\sqrt{1+t^2}$, and evaluating any unit vector gives equality. The argument includes $t = 0$, where both sides vanish.

For the common section $A(r) = (r, 0)$, the perturbation inverse leaves A unchanged because $D_t A = I$. The constructed transport is therefore $Jx = x - AD_t x$ on the old kernel. At $x = (0, b)$, it gives $(-tb, b)$. Substitution into the projection formula gives instead $(-tb/(1+t^2), b/(1+t^2))$. The strict-radius theorem applies when $|t| < 1$, but these explicit line and projection calculations are valid for every real t. □

21.5 What fails without the hypotheses

Take $D_0 = \operatorname{diag}(1, 0)$ and $D_t = \operatorname{diag}(1, t)$. The operator difference has norm $|t|$. For every nonzero t, however, the new kernel is zero, while the old kernel is the vertical axis. Their projections differ in norm by one. Operator closeness alone does not imply kernel continuity.

More generally, suppose a unit vector x is exact for the original operator but orthogonal to the new kernel. Every new residual gain C' obeys

$$1 = \|x - q'x\| \le C' \|D'x\| \le C'\varepsilon.$$

Losing an exact direction under a small operator change forces the new gain to grow. At the strict radius boundary, $D = I$, $A = I$, and $D' = 0$ have $a = \varepsilon = 1$, yet D' is not surjective. Thus the strict condition in Theorem 21.2 cannot be replaced universally by a non-strict one.

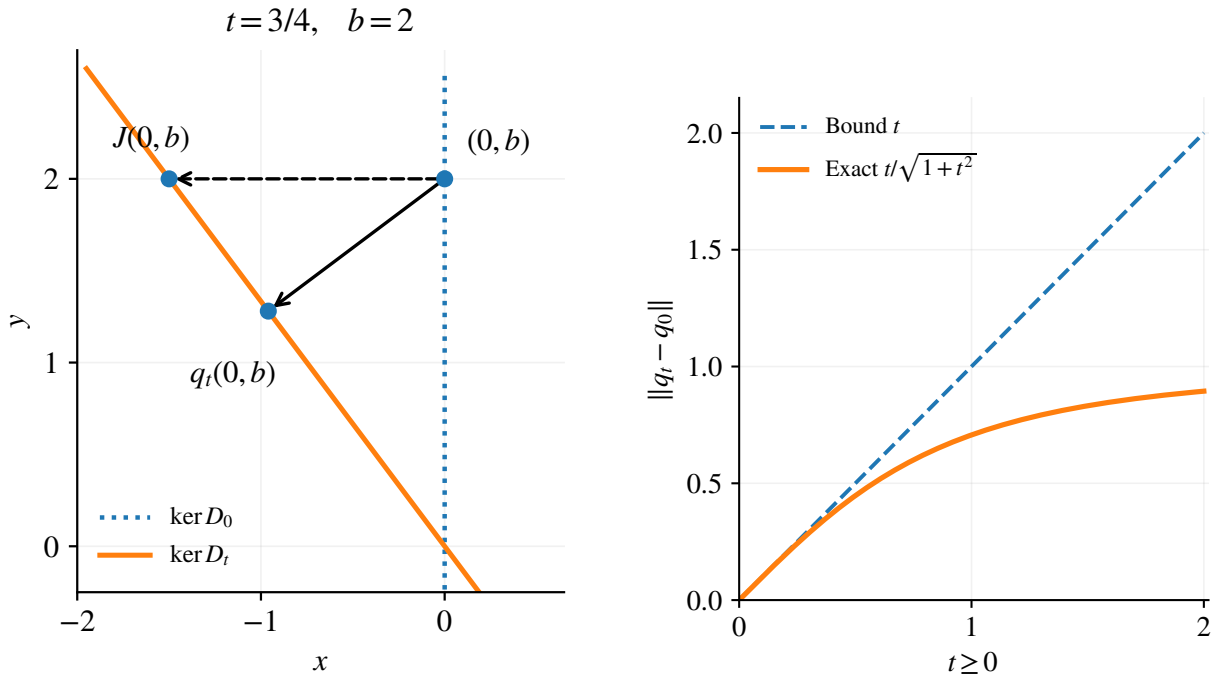


Figure 19: The exact space can move continuously without remaining equal to the original kernel. Transport along a chosen splitting preserves a coordinate, while orthogonal projection minimizes distance. In the line example these two maps have different formulas even though both land in the same exact space.

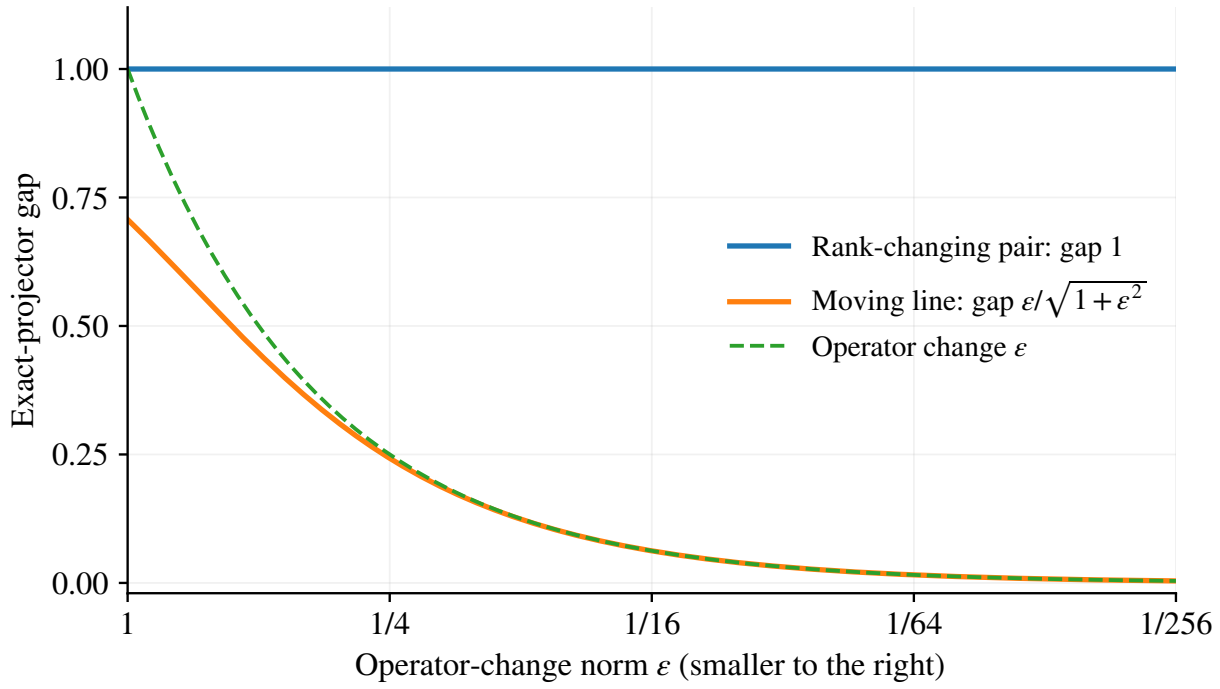


Figure 20: Two families can have the same small operator-change norm and very different exact-space motion. A moving line retains a uniform residual gain. The rank-changing diagonal family loses an exact direction for every nonzero parameter, so its projector gap remains one.

21.6 Motion of the actual descent repair

Compare pairs (D, C) and (D', C') on the same raw Hilbert space. Suppose each pair has closed local range and closed overlap range on its own locally exact subspace. Let G, G' be the resulting joint residual gains. If

$$\|D - D'\| \le \varepsilon_D, \qquad \|C - C'\| \le \varepsilon_C,$$

then the maximum norm on the target product gives

$$\|(D, C) - (D', C')\| \le \max(\varepsilon_D, \varepsilon_C).$$

Applying Theorem 21.1 to the joint operators, and using the actual two-stage identity of Theorem 20.1, gives

$$\|P_{D,C} - P_{D',C'}\| \le \max(G, G') \max(\varepsilon_D, \varepsilon_C). \tag{21.3}$$

This is a comparison of the original repair maps, not a replacement definition that happens to have a similar kernel. Both the local equations and overlap conditions can change. The raw space and its inner product remain fixed.

22 Reading a stability certificate

A numerical residual becomes useful only after it is tied to an operator, a norm, and a conclusion. This section separates uncertainty in a measured residual from uncertainty in the operator itself. It also distinguishes closeness to a reference exact space from closeness to a changed one. All estimates below refer to the actual canonical Hilbert projection used earlier.

22.1 The statement carried by a certificate

Let $D : E \to R$ be bounded linear from a real Hilbert space into a real normed space. A residual certificate with gain $C \ge 0$ is the universally quantified estimate

$$\|x - qx\| \le C \|Dx\| \quad \text{for every } x \in E, \qquad q = P_{\ker D}.$$

A reported number C is not a certificate unless this estimate has been justified in the stated model. A successful test at one input is likewise not the universally quantified estimate.

Suppose the measured vector is y, with a certified error radius η such that $\|Dx - y\| \le \eta$. The triangle inequality converts the measured ball into a bound for the true residual.

Proposition 22.1 (Measurement and operator uncertainty)**.** *With the reference certificate above,*

$$\|x - qx\| \le C(\|y\| + \eta) \quad \text{if } \|Dx - y\| \le \eta.$$

More generally, let $\widetilde{D} : E \to R$ be the bounded linear operator actually used in the measurement. If

$$\left\|D - \widetilde{D}\right\| \le \varepsilon, \qquad \left\|\widetilde{D}x - y\right\| \le \eta,$$

then

$$\|x - qx\| \le C(\|y\| + \eta + \varepsilon \|x\|). \tag{22.1}$$

The projection in this conclusion is still the projection for the reference operator D.

Proof. The first assertion follows from $\|Dx\| \le \|y\| + \eta$. For the second, decompose

$$Dx = (D - \widetilde{D})x + (\widetilde{D}x - y) + y.$$

Its norm is at most $\varepsilon \|x\| + \eta + \|y\|$. Apply the reference residual certificate. □

The input-size term has a mathematical purpose. A small error in an operator can produce a large error on a large input. It cannot generally be replaced by a number depending only on the measured residual. Even a measurement of zero may reflect a difference between the measured kernel and the reference kernel.

22.2 A vanishing measured residual can still miss the reference space

For a concrete family, take $E = \mathbb{R}^2$ with the Euclidean norm and $R = \mathbb{R}$, and define

$$D(a, b) = a, \qquad \widetilde{D}(a, b) = a + tb.$$

The reference projection is $q(a, b) = (0, b)$, its least residual gain is one, and $\left\|D - \widetilde{D}\right\| \le |t|$. On the input $x = (-tB, B)$, the measured operator returns exactly zero. Its distance to the reference kernel is nevertheless $|tB|$.

Setting $\eta = 0$ and $y = 0$ in (22.1) gives the valid upper bound $|t|\sqrt{1 + t^2}|B|$. Deleting the input-size term would give zero and would be false whenever $tB \neq 0$. The input is exact for the changed model and not exact for the reference model. This is precisely the distinction that the estimate is designed to retain.

There is no statistical assertion here. The error radii are mathematical inputs. Estimating them from noisy observations, unknown dynamics, or a fitted model is an additional task that requires its own assumptions and evidence.

22.3 A preserved kernel gives a stronger conclusion

A useful regime occurs when the changed operator annihilates every old exact vector:

$$\widetilde{D}q = 0.$$

This says that the old kernel is contained in the new kernel. Equality is not assumed. Let $\left\|D - \widetilde{D}\right\| \le \varepsilon$ and suppose the strict budget $C\varepsilon < 1$ holds.

Theorem 22.2 (Kernel preservation and a strict error budget)**.** *Under these hypotheses,*

$$\ker \widetilde{D} = \ker D, \qquad \|x - qx\| \le \frac{C}{1 - C\varepsilon} \left\|\widetilde{D}x\right\|.$$

For a measurement with $\left\|\widetilde{D}x - y\right\| \le \eta$*, the bound becomes*

$$\|x - qx\| \le \frac{C}{1 - C\varepsilon}(\|y\| + \eta).$$

Thus the exact space is unchanged and the residual gain degrades by an explicit factor.

Proof. Let $q = P_{\ker D}$, and write $u = x - qx$. The stated kernel inclusion makes both Dqx and $\widetilde{D}qx$ zero. Hence $Du = Dx$ and $\widetilde{D}u = \widetilde{D}x$. Applying the reference residual bound to u, whose projection onto $\ker D$ is zero, gives

$$\begin{aligned}\|u\| &\le C\,\|Du\| \le C\left\|\widetilde{D}u\right\| + C\left\|(D-\widetilde{D})u\right\| \\ &\le C\left\|\widetilde{D}x\right\| + C\varepsilon\,\|u\|\,.\end{aligned}$$

Subtract the final term. The strict inequality $C\varepsilon < 1$ makes $1 - C\varepsilon$ positive, so division preserves the inequality and gives the claimed residual estimate. This also covers $C = 0$ without division by C.

If $\widetilde{D}x = 0$, the same estimate forces $u = 0$, hence $x = qx \in \ker D$. This is the inclusion not initially assumed. Combining it with the supplied inclusion proves $\ker \widetilde{D} = \ker D$, and their orthogonal projections coincide in the fixed Hilbert geometry. Finally $\left\|\widetilde{D}x - y\right\| \le \eta$ gives $\left\|\widetilde{D}x\right\| \le \|y\| + \eta$. Substitution proves the measured certificate. The kernel-preservation hypothesis is what removed qx before applying the perturbation estimate. Without it, one obtains a term proportional to $\|x\|$, not the absorbable term proportional to $\|u\|$. □

This result does not require the reference operator to be surjective. It is different from the split-surjective perturbation theorem, which allows the kernel to move and constructs an actual transport between the old and new kernels. Both theorems have a strict quantitative radius, but they retain different structural data.

For $D(a,b) = a$ and $\widetilde{D}(a,b) = (1-t)a$, with $0 \le t < 1$, the kernel is preserved and the exact gain is $1/(1-t)$. The bound is attained on transverse inputs. At $t = 1$, the changed operator is zero and its kernel becomes all of $\mathbb{R}^2$. This illustrates why the strict endpoint cannot simply be included.

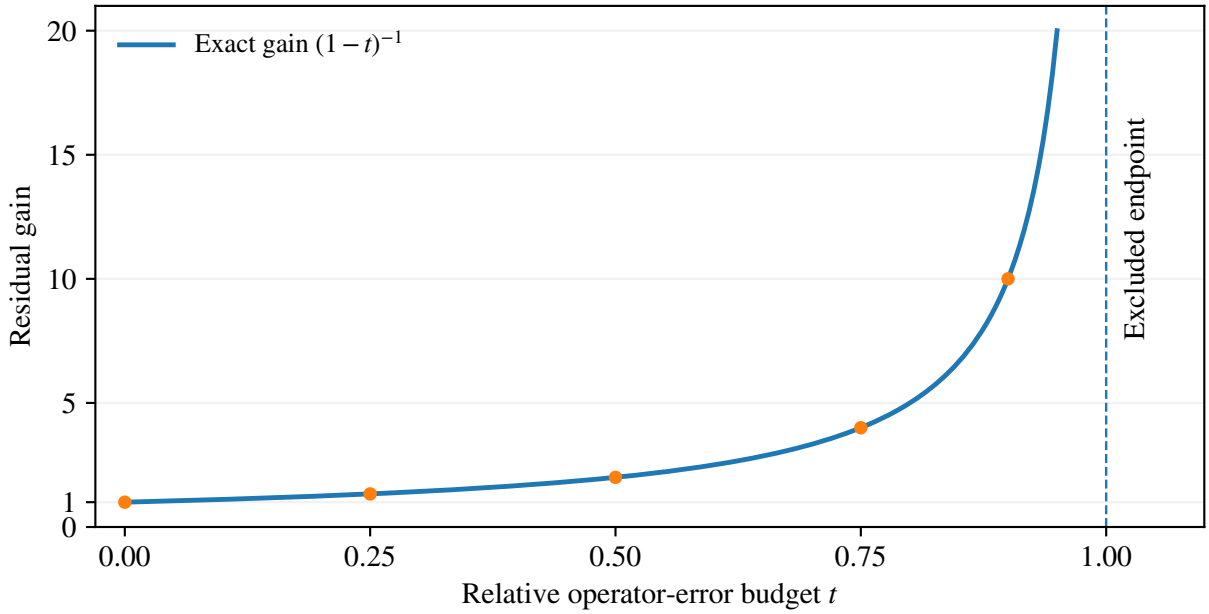


Figure 21: For the kernel-preserving family $\widetilde{D}_t(a,b) = (1-t)a$, the exact residual gain is $1/(1-t)$. It deteriorates as the strict budget is exhausted. At the excluded endpoint the kernel changes, so the finite-radius estimate does not extend by substituting $t = 1$.

22.4 Select the theorem by its conclusion

The following distinctions organize the estimates without turning them into competing formulas.

Available information	Structural hypothesis	Certified conclusion
Reference gain and measured error	None on a changed kernel	Distance to the reference kernel, with an input-size term
Reference gain and small operator change	Old kernel preserved	Same kernel and a degraded gain
Two residual gains	Same source Hilbert geometry	Distance between the two orthogonal projectors
A right inverse and a strict radius	Split surjectivity	A new right inverse and an equivalence of moving kernels

A claim about the new exact space cannot be obtained by silently replacing the reference projector in the first row. Conversely, a theorem that establishes equality of kernels under preservation need not be weakened to a moving-kernel estimate when the stronger hypothesis is actually available.

22.5 Norm conventions are part of the certificate

On a finite family of real Hilbert fibers, the maximum norm and the Euclidean product norm satisfy

$$\|x\|_\infty \le \|x\|_2 \le \sqrt{n}\,\|x\|_\infty ,$$

where n is the number of components. The first inequality follows because each squared component is one term in the sum. The second follows by bounding every component by the maximum. These inequalities transfer bounds with explicit dimension factors. They do not assert that the same numerical best constant remains optimal.

The graph chapters use the maximum norm on vertex and directed-edge products unless a different geometry is declared. The nearest-point statements use a Hilbert source norm. In the worked graph calculations that follow, selecting one orientation and using a sum of squared residuals defines a different numerical operator model. It has the same edge-exactness equations for an isometric connection, but a different residual norm.

A repair certificate should therefore identify its candidate space, residual map, preserved data, norms, gain, and error budgets. The exact equation alone does not specify these quantities. Explicitly retaining them is what makes a numerical conclusion transportable between examples.

23 A geometry of transformations needs its own cost

The preceding analytic chapters measure the size of a residual or the distance from a candidate to an exact subspace. Neither measurement automatically defines a distance between mathematical objects. In an additive category, the zero map is often perfectly compatible and perfectly uninformative. A geometry that assigns zero cost to every compatible map can therefore identify all objects at distance zero.

A cost of transformation must be declared separately. It may measure an admissible computational expense, a loss of retained information, or another specified quantity. The interpretation alone proves none of its mathematical laws. The construction below begins only after the required laws have been established.

23.1 Extended directed distance

Let $\mathcal{C}$ be a category and let W assign each arrow a value in the extended nonnegative reals $[0, \infty]$. Suppose

$$W(\mathrm{id}_A) = 0, \qquad W(gf) \leq W(f) + W(g).$$

The order of ordinary composition is first f, then g. Define

$$d_W(A, B) = \inf_{f:A\to B} W(f).$$

The infimum of an empty family is ∞. No assumption of connectedness, finite cost, symmetry, or attainment is made.

Proposition 23.1 (Directed cost geometry)**.** *For every arrow $f : A \to B$,*

$$d_W(A, B) \leq W(f), \qquad d_W(A, A) = 0,$$

and for all objects,

$$d_W(A, C) \leq d_W(A, B) + d_W(B, C).$$

The symmetric function

$$d_W^{\mathrm{sym}}(A, B) = \max\{d_W(A, B), d_W(B, A)\}$$

has zero diagonal, is symmetric, and satisfies the extended triangle inequality. The relation $d_W^{\mathrm{sym}}(A, B) = 0$ is an equivalence relation.

Proof. The infimum is bounded above by each of its entries. Apply this to the identity to obtain the zero diagonal. If either term on the right of the triangle inequality is infinite, there is nothing to prove. Otherwise, for every $\varepsilon > 0$, choose arrows with costs less than the corresponding infima plus ε. Their composite has cost less than $d_W(A, B) + d_W(B, C) + 2\varepsilon$. Letting ε tend to zero gives the result. This argument does not require an optimizing arrow.

Each of the two directed distances from A to C and from C to A is bounded by $d_W^{\mathrm{sym}}(A, B) + d_W^{\mathrm{sym}}(B, C)$. Taking their maximum proves the symmetric triangle law. Zero diagonal and symmetry give reflexivity and symmetry of the zero-distance relation. The triangle inequality and nonnegativity give transitivity. □

The formal source proves these laws directly in $\mathbb{R}^{\infty}_{\geq 0}$ and constructs the zero-distance setoid. It does not identify zero distance with categorical isomorphism. Passing to the quotient can separate zero-distance classes, but would still not prove that those classes are the intended semantic objects.

23.2 A concrete path category

Consider four vertices A, B, C, D, and the free path category generated by

$$A \xrightarrow{0} B, \qquad B \xrightarrow{0} A, \qquad A \xrightarrow{2} C, \qquad C \xrightarrow{5} A.$$

There are no arrows incident to D, apart from its identity. A path costs the sum of its edge labels. Concatenation is additive and empty paths cost zero, so all cost axioms hold. The directed distances are

d_W	A	B	C	D
A	0	0	2	∞
B	0	0	2	∞
C	5	5	0	∞
D	∞	∞	∞	0

For example, the path from B through A to C costs two. Every path from C to B must first pay the edge cost five and can then move from A to B for free. Nonnegative loops cannot improve either cost.

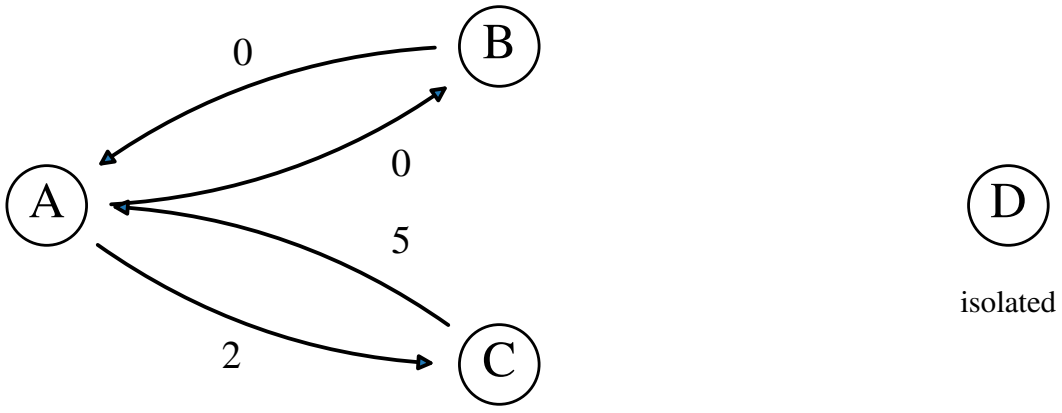


Figure 22: Directed path costs need not be symmetric or finite. Vertices A and B have zero cost in both directions, but they are not isomorphic in the free path category. The distinct zero-cost arrows have no imposed inverse relations.

The symmetric zero-distance classes are $\{A, B\}$, $\{C\}$, and $\{D\}$. The distance between the first two classes is five. The third class is infinitely far from each of them. Yet A and B are not isomorphic in this category. Any path from one to the other has positive length, so concatenating it with a return path cannot give the length-zero identity. This is a categorical distinction even though the numerical cost is zero.

One can instead impose relations declaring the two arrows inverses. That changes the category. The distance table alone does not record whether those relations were imposed.

23.3 Why compatibility alone can collapse the geometry

Proposition 23.2 (The zero-arrow collapse)**.** *Suppose $\mathcal{C}$ has zero morphisms and $W(0 : A \to B) = 0$ for every pair of objects. Then*

$$d_W(A, B) = 0 \quad \textit{for all } A, B.$$

Proof. The zero arrow is one member of the family over which the infimum is taken. Thus $0 \leq d_W(A, B) \leq W(0) = 0$. □

For the preadditive tower differential, the zero family of maps is exact. Consequently any proposed nonnegative arrow cost that vanishes on every exact map also vanishes on every zero map. If it satisfies the cost axioms at all, its object distance collapses by the proposition.

The qualification is important. A raw operator defect norm need not satisfy the unweighted composition inequality required of W. The transport Leibniz identity gives coefficients involving the surrounding maps. It is not legitimate to first call an arbitrary defect norm a subadditive cost and then use the metric construction.

23.4 Preservation changes the admissible arrows

One response to collapse is to restrict arrows to those satisfying a declared preservation contract. For objects equipped with maps $p_A : A \to S$, a morphism can be required to satisfy

$$p_B f = p_A.$$

Identities and composition preserve this equation. Unless $p_A = 0$, the zero arrow is not admissible. In other applications the arrows might be required to be isomorphisms or to retain a particular observable. These are different choices of category, not automatic consequences of small residuals.

Even after this restriction, separation remains a theorem to prove rather than a consequence of terminology. Distinct objects can still be connected by arbitrarily cheap preserving maps. Symmetry can fail, and the infimum can fail to be attained. A quantitative comparison therefore needs four pieces of information: the objects, the admissible arrows, the actual cost, and the meaning of zero distance.

The role of this construction is modest but useful. It prevents a compatibility measure from being silently promoted to a universal geometry. It also gives a valid directed geometry when the separate categorical and cost assumptions really hold.

Part V

From approximation to exact realization

A vanishing defect is a statement about measurements. A limiting realization also requires control of the candidate itself. This part follows the movement of approximate answers, the closure of admissible correction images, and the meaning of the ambient completion.

24 Exactification without a fixed contraction rate

The constructed linear repair in Section 8 is a useful special case of a more general principle. One can prove convergence without one fixed contraction factor, without a linear update rule, and without a globally Lipschitz iteration. What must not be lost is control of the total movement.

24.1 The two limits that must be established

Let (M, dist) be a complete metric space. Let $\omega : M \to [0, \infty)$ be a nonnegative residual function, and call a point exact when $\omega(x) = 0$. We regard ω as a real-valued function when discussing continuity. Assume it is lower semicontinuous, so that every sublevel set $\{x : \omega(x) \le a\}$ is closed.

For an approximation sequence f_n, the statements

$$\omega(f_n) \longrightarrow 0 \quad \text{and} \quad f_n \longrightarrow x$$

are different. The first concerns the observed discrepancy. The second concerns the mathematical candidate in the chosen metric. Neither implies the other without additional hypotheses.

Theorem 24.1 (Summable movement gives exactification)**.** *Suppose $b_n \ge 0$, $\sum_{n=0}^{\infty} b_n < \infty$, and*

$$\mathrm{dist}(f_n, f_{n+1}) \le b_n, \qquad \omega(f_n) \longrightarrow 0.$$

Then there is a limit $x \in M$ satisfying

$$f_n \longrightarrow x, \qquad \omega(x) = 0, \qquad \mathrm{dist}(f_n, x) \le \sum_{k=n}^{\infty} b_k.$$

Proof. For integers $m > n$, apply the triangle inequality successively along the finite chain $f_n, f_{n+1}, \dots, f_m$. It gives

$$\mathrm{dist}(f_n, f_m) \le \sum_{k=n}^{m-1} \mathrm{dist}(f_k, f_{k+1}) \le \sum_{k=n}^{m-1} b_k \le \sum_{k=n}^{\infty} b_k.$$

Given $\varepsilon > 0$, convergence of the nonnegative series supplies N for which the final tail is below ε whenever $n \ge N$. Interchanging m, n when necessary proves the Cauchy condition for all pairs of indices at least N. Completeness, not residual decay, is the hypothesis that now supplies a limit $x \in M$.

Keep n fixed. The reverse triangle inequality says

$$\big|\mathrm{dist}(f_n, f_m) - \mathrm{dist}(f_n, x)\big| \le \mathrm{dist}(f_m, x) \longrightarrow 0.$$

Passing to the limit in the uniform tail bound proves the claimed distance certificate at that n.

It remains to prove exactness of the limit. Fix $\varepsilon > 0$. Since $\omega(f_m) \to 0$, an entire tail of the sequence belongs to the sublevel set $C_\varepsilon = \{z : \omega(z) \le \varepsilon\}$. Lower semicontinuity makes this set closed, so the limit of that tail belongs to it. Hence $\omega(x) \le \varepsilon$ for every positive ε. If $\omega(x) > 0$, choosing $\varepsilon = \omega(x)/2$ contradicts this inequality. Since ω is nonnegative, $\omega(x) = 0$. We have not used, or asserted, continuity of ω or the stronger equality $\omega(f_m) \to \omega(x)$ for arbitrary convergent inputs. □

The formal theorem is slightly more general: it is stated in a complete pseudometric space. In a metric space the limit of this fixed sequence is unique. That does not mean different admissible approximation procedures select the same exact point.

24.2 Variable rates and a movement budget

Let $q_n, C_n \ge 0$, and define the prefix product

$$Q_0 = 1, \qquad Q_{n+1} = q_n Q_n.$$

Assume

$$\omega(f_{n+1}) \le q_n \omega(f_n), \qquad \operatorname{dist}(f_n, f_{n+1}) \le C_n \omega(f_n). \tag{24.1}$$

Theorem 24.2 (Variable-rate exactification)**.** *If $Q_n \to 0$ and $\sum_n C_n Q_n < \infty$, then (f_n) has an exact limit and*

$$\operatorname{dist}(f_n, x) \le \omega(f_0) \sum_{k=n}^{\infty} C_k Q_k. \tag{24.2}$$

For constant $q_n = q \in [0, 1)$ and constant $C_n = C$, the sharper restarted estimate is

$$\operatorname{dist}(f_n, x) \le \frac{C}{1-q}\, \omega(f_n).$$

Proof. The estimate $\omega(f_n) \le Q_n \omega(f_0)$ follows by induction. At $n = 0$ it is equality. If it holds at n, nonnegativity of q_n permits multiplication and gives

$$\omega(f_{n+1}) \le q_n \omega(f_n) \le q_n Q_n \omega(f_0) = Q_{n+1} \omega(f_0).$$

Because $Q_n \to 0$ and $\omega(f_0)$ is a fixed finite number, the residual tends to zero. The movement assumption then gives

$$\operatorname{dist}(f_n, f_{n+1}) \le C_n Q_n \omega(f_0).$$

This is a summable nonnegative budget by the hypothesis on $\sum C_n Q_n$. All assumptions of Theorem 24.1 now hold, and its tail certificate is exactly (24.2). If $\omega(f_0) = 0$, the same estimates make every step zero and every residual zero. No cancellation of that value was used.

For constant rates, the global estimate can be sharpened by restarting at the current index. Fix n. A new induction on j gives

$$\omega(f_{n+j}) \le q^j \omega(f_n), \qquad \operatorname{dist}(f_{n+j}, f_{n+j+1}) \le C q^j \omega(f_n).$$

Sum the finite chain from n to $n + m$, and then let $m \to \infty$ using the already established limit. The geometric series equals $1/(1-q)$, including $q = 0$, where its first term is one and all later terms vanish. This proves the restarted bound. The variable-rate proof itself requires only nonnegative q_n, decay of their cumulative products, and summability of the movement weights. It does not insert an unstated upper bound $q_n < 1$ at each step. □

The individual factors q_n need not all be below one. Temporary residual increases are permitted if the prefix products decay and the weighted movements remain summable. Conversely, a decaying product is not enough when its movement budget diverges.

24.3 A convergent path with no uniform geometric factor

In $\mathbb{R}^2$ with the maximum norm, fix $s \in \mathbb{R}$ and take

$$f_n = (s, (n+1)^{-1}), \qquad \omega(u, r) = |r|.$$

The protected coordinate is the first one. Direct calculation gives

$$q_n = \frac{n+1}{n+2}, \quad Q_n = \frac{1}{n+1}, \quad C_n = \frac{1}{n+2}, \quad \operatorname{dist}(f_n, f_{n+1}) = \frac{1}{(n+1)(n+2)}.$$

No constant $q < 1$ bounds all the ratios q_n, since they tend to one. Nevertheless

$$\sum_{k=n}^{\infty} C_k Q_k = \sum_{k=n}^{\infty} \left(\frac{1}{k+1} - \frac{1}{k+2} \right) = \frac{1}{n+1}.$$

The limit is $(s, 0)$, and the error bound is an equality. The model shows why the variable-rate theorem is not merely a rephrasing of a fixed geometric contraction.

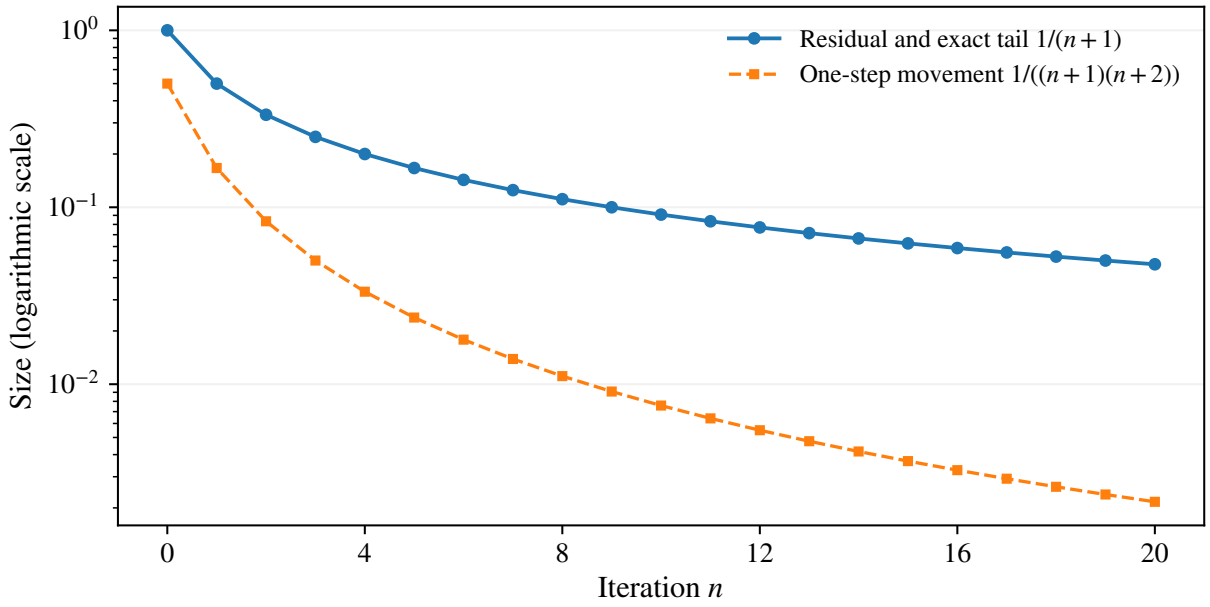


Figure 23: Polynomial residual decay can still give an exact limit with a sharp error certificate. The residual $1/(n+1)$ is not summable, but the actual movement $1/((n+1)(n+2))$ is. The theorem uses the movement series.

24.4 Residual bursts can be compatible with convergence

Set $q_{2k} = 2$, $q_{2k+1} = 1/8$. Then

$$Q_{2k} = 4^{-k}, \qquad Q_{2k+1} = 2 \cdot 4^{-k}.$$

For $f_n = (s, Q_n)$, the residual doubles at the start of every pair of steps. Choose $C_{2k} = 1$ and $C_{2k+1} = 7/8$. These give equality in the movement bound, and

$$\sum_{n=0}^{\infty} C_n Q_n = \sum_{k=0}^{\infty} \left(1 + \frac{7}{4} \right) 4^{-k} = \frac{11}{3}.$$

Thus $f_n \to (s, 0)$ despite infinitely many increases. Starting at an even index $2k$, the remaining movement budget is $\frac{11}{3}4^{-k}$. At an odd index $2k+1$, it is $\frac{8}{3}4^{-k}$. These bounds are larger than the actual distance because the path repeatedly moves away and returns.

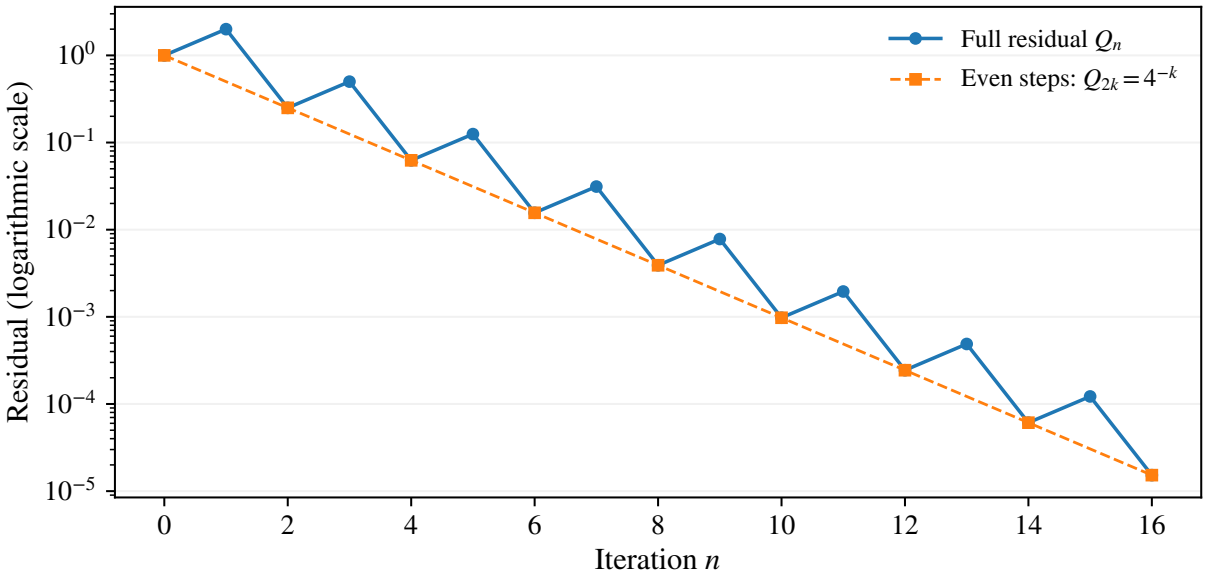


Figure 24: The residual increases at every even-to-odd step, but each two-step product is $1/4$. A convergent product and a summable movement budget control the whole path. Monotonicity of individual residual values is not assumed by the variable-rate theorem.

24.5 Protected meaning survives a controlled limit

Return to Banach and normed linear spaces, with continuous linear maps $d : E \to R$ and $p : E \to S$. Put $\omega(f) = \|df\|$. The residual is continuous, hence lower semicontinuous.

Proposition 24.3 (Controlled paths kill the actual obstruction)**.** *Suppose E is complete, $\|df_n\| \to 0$, $p(f_n) = p(f_0)$, and $\operatorname{dist}(f_n, f_{n+1}) \le b_n$ for a summable nonnegative budget. Then there is x with*

$$f_n \to x, \qquad dx = 0, \qquad p(x) = p(f_0),$$

and the constrained obstruction $[df_0] \in R/d(\ker p)$ vanishes.

Proof. The function $f \mapsto \|df\|$ is continuous because d is bounded linear and the norm is continuous. It is therefore an admissible lower semicontinuous residual for Theorem 24.1. That theorem supplies a limit $x \in E$, together with $\|dx\| = 0$, which means $dx = 0$.

Continuity of p gives $p(f_n) \to p(x)$. The left side is the constant sequence with value $p(f_0)$. Uniqueness of limits in the normed target therefore gives $p(x) = p(f_0)$. Define the actual vector $a = f_0 - x$. Linearity, rather than a closure argument, now yields

$$p(a) = p(f_0) - p(x) = 0, \qquad da = df_0 - dx = df_0.$$

Thus $a \in \ker p$ and $df_0 \in d(\ker p)$. Membership in this algebraic image is precisely what makes the class $[df_0]$ zero. Merely knowing membership in the closure of that image would not justify the same conclusion. □

This proposition converts a controlled approximation argument into a statement about the algebraic quotient defined earlier. Merely showing that $df_n \to 0$ would not produce the admissible vector $f_0 - x$, because there might be no candidate limit x.

24.6 When two repair paths select the same point

Proposition 24.4 (Path independence requires a comparison)**.** *Let $f_n \to x$ and $g_n \to y$ in a metric space. If*

$$\operatorname{dist}(f_n, g_n) \le \beta_n, \qquad \beta_n \to 0,$$

then $x = y$. In contrast, exactness and preservation of a nonzero observable alone do not force a unique answer.

Proof. For each n, the triangle inequality gives

$$\operatorname{dist}(x, y) \le \operatorname{dist}(x, f_n) + \operatorname{dist}(f_n, g_n) + \operatorname{dist}(g_n, y) \le \operatorname{dist}(x, f_n) + \beta_n + \operatorname{dist}(g_n, y).$$

Given any $\varepsilon > 0$, choose n large enough that the three terms on the right are each below $\varepsilon/3$. Thus $\operatorname{dist}(x, y) \le \varepsilon$ for every positive ε, so the distance is zero. The metric separation axiom gives $x = y$. In a pseudometric space only zero distance would follow without an additional separation hypothesis.

For the negative assertion let $E = \mathbb{R} \times (\mathbb{R} \times \mathbb{R})$, $d(s, u, r) = r$, and $p(s, u, r) = s$. Both $x = (1, 0, 0)$ and $y = (1, 1, 0)$ have zero defect and protected value one, but their middle coordinates differ. They are distinct points of the same feasible fiber. Constant sequences at those points meet exactness and preservation but not a vanishing comparison bound, so these conditions alone cannot imply path independence. □

A stronger worked example uses actual convergent paths from the *same* initial input:

$$f_n^{(a)} = (s, a(1 - 2^{-n}), 2^{-n}) \qquad (a \in \mathbb{R}).$$

All paths begin at $(s, 0, 1)$, keep $p = s$, and have residual 2^{-n}. Their maximum-norm movements are $\frac{1}{2}\max\{|a|, 1\}2^{-n}$, so every path satisfies geometric exactification. Their limits are $(s, a, 0)$. For different a, b,

$$\operatorname{dist}(f_n^{(a)}, f_n^{(b)}) = |a - b|(1 - 2^{-n}) \longrightarrow |a - b|,$$

not zero. The missing uniqueness hypothesis is visible in the exact direction along which the paths separate.

24.7 Changing the sampling schedule does not preserve a rate

A residual profile η_n indexed by resolution is not an iteration trajectory. Nevertheless a useful convergence property is invariant under a cofinal schedule when the profile is monotone.

Proposition 24.5 (Cofinal sampling of a decreasing profile)**.** *Suppose $\eta_n \ge 0$ is nonincreasing and $\phi : \mathbb{N} \to \mathbb{N}$ is monotone and unbounded. Then*

$$\eta_n \to 0 \quad \Longleftrightarrow \quad \eta_{\phi(n)} \to 0.$$

Without the monotonicity of η, the reverse implication can fail even for a cofinal schedule.

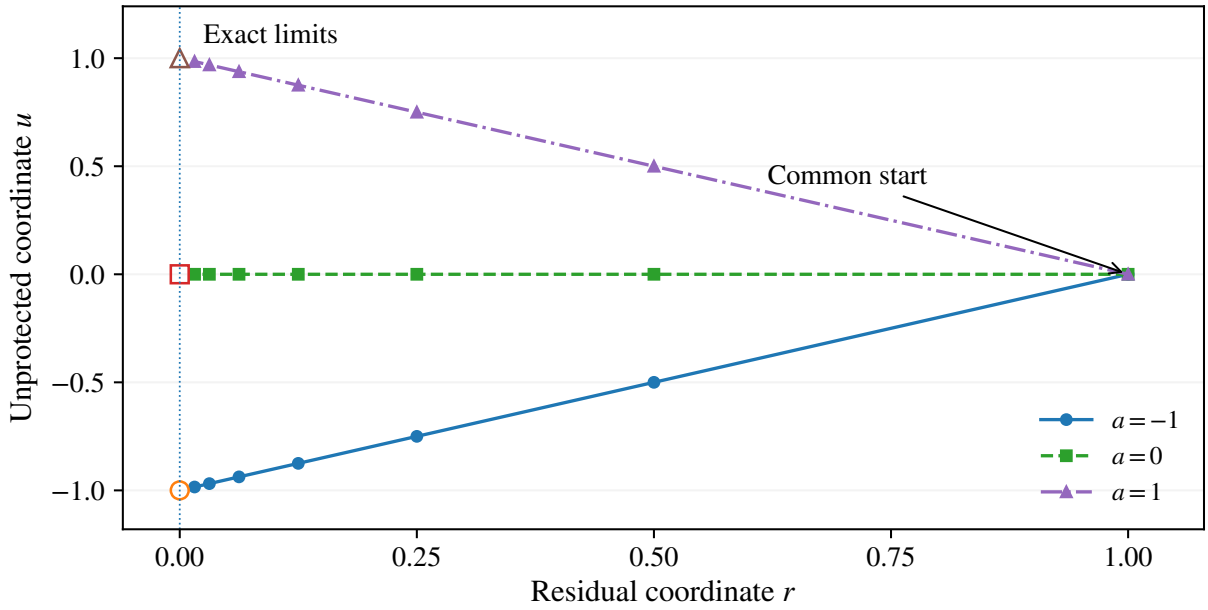


Figure 25: All paths start at the same candidate, have the same residual at each step, and preserve the same coordinate s. They converge to different exact points because the coordinate u is unprotected. A vanishing residual is not a path-comparison budget.

Proof. The schedule tends to infinity, so convergence of the full profile implies convergence of the sampled one. Conversely, choose a sampled value below any prescribed $\varepsilon > 0$. Every later index of the original profile has no larger value, by monotonicity. This proves convergence of the full profile.

For failure without monotonicity, take $\eta_{2n} = 0$, $\eta_{2n+1} = 1$, and $\phi(n) = 2n$. All sampled values vanish, but infinitely many original values equal one. □

Even in the decreasing case, the numerical rate changes. The profile $\eta_n = 1/(n+1)$, sampled at $\phi(n) = 2^n - 1$, becomes 2^{-n}. This is exponential decay in the *number of observations*, not a proof of a faster correction process at the original resolution scale.

24.8 Three hypotheses that cannot simply be dropped

Summable movement. In maximum-norm $\mathbb{R}^3$, let

$$f_n = (s, H_n, (n+1)^{-1}), \qquad H_n = \sum_{j=1}^{n} \frac{1}{j}.$$

The first coordinate is protected and the residual is the last coordinate. It tends to zero with the same ratios as the polynomial example. The movement equals $1/(n+1)$, however, and the unprotected coordinate diverges. Here $C_n = 1$, so $\sum C_n Q_n$ is the divergent harmonic series. Residual decay and semantic preservation hold, but exactification of this path fails.

Completeness. In the metric space $(0, \infty)$, take $f_n = 1/(n+1)$ and $\omega(x) = x$. The steps are summable and the residual tends to zero, but there is no limit in this space. Completing the space adds the missing zero. It is essential to say whether that completion belongs to the intended model.

Lower semicontinuity. In $\mathbb{R}$, define $\omega(0) = 1$ and $\omega(x) = 0$ for $x \neq 0$. The same positive sequence converges with summable movement and residual zero at every step, yet its limit is not exact. The zero

sublevel set is not closed. This example isolates why convergence of measurements along a sequence must interact correctly with the ambient topology.

25 The residual floor and the closure of admissible changes

An algebraic obstruction asks whether a residual is exactly the defect of an allowed correction. A numerical floor asks how closely allowed defects can approximate it. These questions agree under a closedness hypothesis. Without that hypothesis they can disagree even in a complete Hilbert setting.

25.1 Distance to the actual correction image

Let R be a normed additive group and $B \leq R$ an additive subgroup. In the protected linear model, $B = d(\ker p)$. For a residual $\Omega \in R$, define

$$\nu_B(\Omega) = \inf_{b \in B} \|\Omega - b\| .$$

The subgroup contains zero, so this is a finite nonnegative real number. It is an infimum, not an assertion that a nearest admissible residual exists.

Proposition 25.1 (The floor detects closure)**.** *For every $b \in B$,*

$$0 \leq \nu_B(\Omega) \leq \|\Omega - b\| .$$

Moreover,

$$\nu_B(\Omega) = 0 \quad \Longleftrightarrow \quad \Omega \in \overline{B}.$$

If B is closed, this is equivalent to $\Omega \in B$.

Proof. The set of candidate distances is nonempty because $0 \in B$, bounded below by zero, and contains the finite value $\|\Omega\|$. Its infimum is therefore a finite nonnegative real, and the defining lower-bound property gives $\nu_B(\Omega) \leq \|\Omega - b\|$ for every $b \in B$.

Suppose $\nu_B(\Omega) = 0$ and fix $\varepsilon > 0$. There must be $b \in B$ with $\|\Omega - b\| < \varepsilon$. Otherwise every candidate distance would be at least ε, making ε a lower bound and forcing the infimum to be at least ε, a contradiction. Thus every open ball about Ω meets B, which is the closure criterion.

Conversely, if $\Omega \in \overline{B}$, every such ball supplies b with $\|\Omega - b\| < \varepsilon$. The infimum then lies between zero and ε for every positive ε, so it is zero. If B is closed, membership in its closure is exactly membership in B. No step asserts that the infimum has a minimizing element when B is not closed, and closedness alone is being used here only for the zero-floor membership conclusion. □

Translation by a member $b_0 \in B$ does not change the floor: replacing b by $b - b_0$ gives the same set of distances. Thus the floor is well defined on the algebraic quotient R/B. When B is a linear subspace, this is the usual quotient seminorm. A nonzero algebraic class can have seminorm zero when B is not closed. The retained source implements the floor on residual representatives and proves the closure criterion. This quotient interpretation is its elementary translation-invariance calculation, not a new claim that every obstruction quotient carries a separated norm.

25.2 A precise three-way distinction in this fixed model

For this specified subgroup and residual there are exactly three possibilities:

Residual location	Algebraic class	Meaning of the floor
$\Omega \in B$	Zero	Exact cancellation exists, and the floor is zero.
$\Omega \in \overline{B} \setminus B$	Nonzero	Arbitrarily small corrected residuals exist, but no exact cancellation exists.
$\Omega \notin \overline{B}$	Nonzero	The floor is strictly positive and blocks every allowed correction.

This is a classification for the stated normed-group correction problem, not a universal trichotomy for every possible Komorebi model. Testing which case holds can itself be difficult. The definitions do not supply a decision algorithm.

In finite-dimensional real or complex linear models, every linear subspace is closed. The middle row is therefore absent. It becomes possible in infinite dimensions, or for additive subgroups that are not closed real-linear subspaces. For example the rational subgroup of $\mathbb{R}$ is dense, so the distance from any irrational real to that subgroup is zero even though its coset is nonzero.

Closedness also should not be confused with a numerical error bound back in the candidate space. The statement $\Omega \in B$ supplies an admissible correction, but not a bound on its size. The closed-range estimates of Sections 18 and 21, or their protected counterparts, are the additional analytic tools that control that size in their declared settings.

25.3 A persistent floor blocks every correction sequence

The correction image can vary with the resolution or the admissibility rule. Let $B_n \leq R$ and let $\Omega_n \in R$ be given.

Proposition 25.2 (Uniformly positive floors obstruct asymptotic repair)**.** *Suppose $c > 0$ and eventually*

$$c \leq \nu_{B_n}(\Omega_n).$$

For any choices $b_n \in B_n$, it is impossible that

$$\|\Omega_n - b_n\| \longrightarrow 0.$$

Proof. Every allowed correction satisfies $\|\Omega_n - b_n\| \geq \nu_{B_n}(\Omega_n)$. The corrected norms are therefore eventually at least c. They cannot converge to zero. □

The quantifier covers all admissible sequences, not just the output of one repair routine. On the other hand, the hypothesis is an eventually *uniform* positive lower bound. Having a nonzero obstruction class at each stage is not the same statement. The classes might have zero floors, or their positive floors might tend to zero.

A finite-dimensional illustration uses $R = \mathbb{R}^2$, $B = \mathbb{R} \times \{0\}$, and $\Omega_n = (n, 1)$. Every correction leaves the second coordinate equal to one, so the floor is one independently of n. Replacing the residuals by $(n, 1/(n+1))$ gives nonzero algebraic classes at every stage, but their floors tend to zero. The persistent-floor theorem applies to the first sequence and not the second.

25.4 The bridge back to exactification

Proposition 24.3 and Proposition 25.1 fit together closely. If corrections produce a norm-convergent, protected path with vanishing residual, the actual quotient class vanishes. If only the corrected residuals tend to zero, all that follows is membership in $\overline{d(\ker p)}$.

In particular, when the class is nonzero but its floor is zero, *no* protected path with vanishing residual can satisfy all the exactification hypotheses. Something must fail: the path may escape, its movements may not be summable, the candidate space may not be complete, or its residual and semantic maps may not have the required continuity. The next model isolates escape while keeping the spaces Hilbert and both maps continuous.

26 Exact finite repairs can escape the intended ambient space

This model links the algebraic obstruction, the numerical floor, the movement of approximations, and the semantic image of an inverse limit. It uses a diagonal operator whose entries are completely explicit. Every finite problem is solvable. The protected infinite problem is not.

26.1 The complete spaces and the preservation rule

Index sequences by the positive integers. Set

$$E = \ell^2(\mathbb{N}_{>0};\mathbb{R}) \oplus \mathbb{R}, \qquad R = \ell^2(\mathbb{N}_{>0};\mathbb{R}), \qquad \|(u,t)\|_E^2 = \sum_{n\geq 1} |u_n|^2 + |t|^2.$$

Let $w = (1/n)_{n\geq 1}$. It belongs to ℓ^2, because

$$\sum_{n\geq 1} \frac{1}{n^2} \leq 1 + \sum_{n\geq 2} \frac{1}{n(n-1)} = 2.$$

Define the defect and protected observable by

$$d(u,t)_n = \frac{u_n + t}{n}, \qquad p(u,t) = t.$$

The diagonal map $T : u \mapsto (u_n/n)_n$ has operator norm one. Since $d(u,t) = Tu + tw$,

$$\|d(u,t)\| \leq \|u\| + |t|\,\|w\| \leq \sqrt{1 + \|w\|^2}\;\|(u,t)\|_E\,.$$

Thus both d and p are continuous linear maps between the declared complete spaces. No failure of completeness or continuity is built into the model.

Take the raw candidate $f_0 = (0,1)$. Its residual is w, and the coordinate that must be preserved is $t = 1$. An exact protected answer would require

$$u_n = -1 \quad \text{for every } n.$$

That sequence does not lie in ℓ^2. Consequently no answer in E is both exact and protected.

In fact $\ker d = \{0\}$. If $d(u,t) = 0$, then $u_n = -t$ for every n. A constant square-summable sequence must be zero, so $t = 0$ and $u = 0$. There is an exact answer after abandoning the preservation rule, namely zero, but it changes $p(f_0) = 1$ to zero. Even a trivial ambiguity space does not imply feasibility.

26.2 The obstruction class is nonzero while its floor is zero

The admissible correction space is $\ker p = \ell^2 \oplus \{0\}$, with correction image

$$B = d(\ker p) = \operatorname{ran} T = \{r \in \ell^2 : (nr_n)_{n\geq 1} \in \ell^2\}.$$

It contains every finitely supported sequence, so it is dense in R. But $w \notin B$, because $(nw_n)_n = (1,1,\ldots)$ is not square summable. Therefore

$$[w] \neq 0 \text{ in } R/B, \qquad \nu_B(w) = 0.$$

This is the middle row of the preceding classification. The example is a worked Hilbert-space derivation without a separate Lean instance. Its finite-coordinate calculations are independently tested. The general floor and exactification results in Sections 24 and 25 have formal counterparts.

26.3 A complete calculation of the finite repairs

Let $\mathbf{1}_N$ be one in coordinates $1, \dots, N$ and zero elsewhere. Put

$$a_N = (\mathbf{1}_N, 0), \qquad f_N = f_0 - a_N = (-\mathbf{1}_N, 1).$$

Then $p(f_N) = 1$, and the first N residual coordinates vanish exactly. The remaining residual and movement satisfy

$$d(f_N)_n = \begin{cases} 0 & n \le N, \\ 1/n & n > N, \end{cases} \qquad \|a_N\|^2 = N, \qquad \|f_N\|^2 = N + 1.$$

For every $N \ge 1$, telescoping comparisons give

$$\frac{1}{N+1} \le \|d(f_N)\|^2 = \sum_{n>N} \frac{1}{n^2} \le \frac{1}{N}. \tag{26.1}$$

Indeed $1/[n(n+1)] \le 1/n^2 \le 1/[n(n-1)]$ for $n \ge 2$. Summing the two telescoping bounds proves (26.1) without numerically evaluating an infinite series.

The residual tends to zero, but the corrections grow like $\sqrt{N}$. Successive candidates differ by a unit coordinate, so

$$\|f_{N+1} - f_N\| = 1.$$

The movement series diverges and the sequence is not Cauchy. This is exactly the hypothesis missing from the summable exactification theorem.

Repaired coordinates N	$\|a_N\|$	Residual square interval	Finite inverse gain
1	1	$[1/2, 1]$	1
4	2	$[1/5, 1/4]$	4
16	4	$[1/17, 1/16]$	16
64	8	$[1/65, 1/64]$	64

The intervals bound the *infinite* residual remaining after the displayed repair. The inverse-gain column concerns a different finite operator, defined below. Keeping those meanings separate prevents a finite success from being mistaken for a uniform infinite certificate.

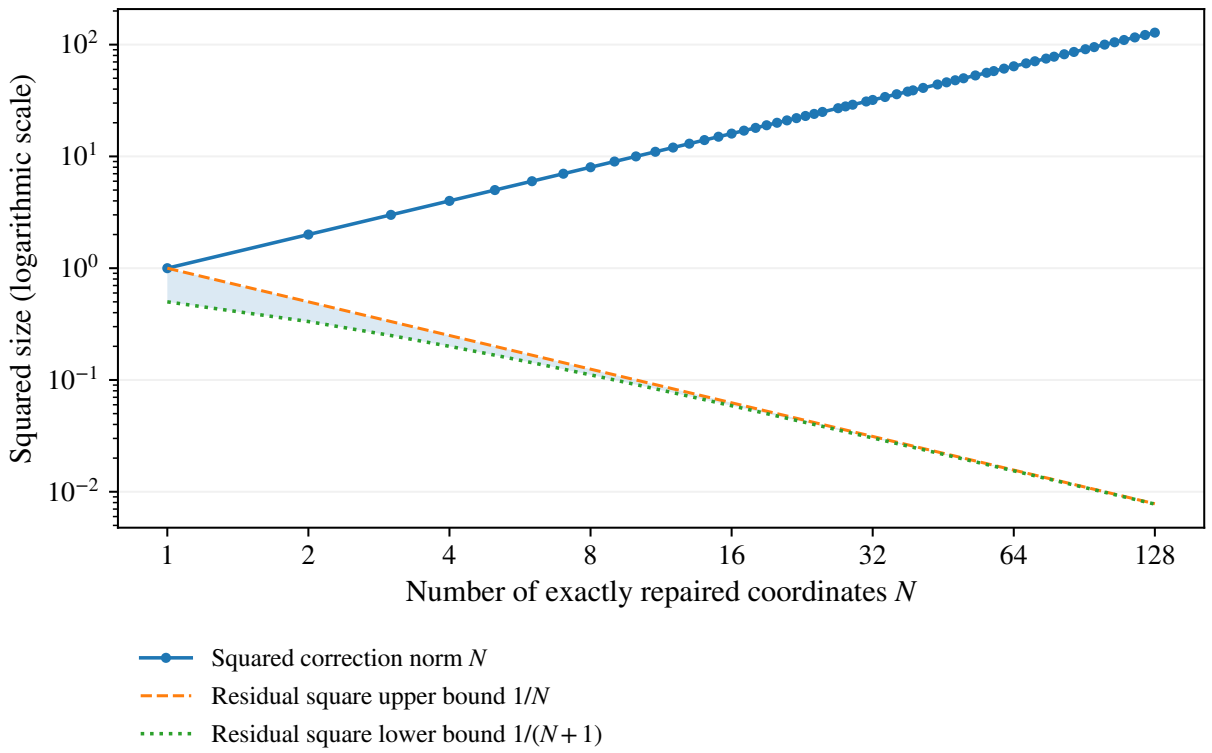


Figure 26: The squared correction size is N, while the infinite residual square lies between $1/(N+1)$ and $1/N$. The plotted band consists of proved bounds, not a numerical truncation of the tail. Smaller residuals require larger changes and do not produce a convergent candidate.

26.4 Escape is not an accident of this particular algorithm

Suppose an arbitrary protected candidate $(u,1)$ satisfies $\|d(u,1)\| \le \varepsilon$. On its first N coordinates,

$$\sum_{n=1}^{N} |u_n+1|^2 \le N^2 \sum_{n=1}^{N} \frac{|u_n+1|^2}{n^2} \le N^2\varepsilon^2.$$

The reverse triangle inequality in $\mathbb{R}^N$ gives

$$\|u\| \ge \sqrt{N} - N\varepsilon \qquad (N \ge 1). \tag{26.2}$$

If a sequence of protected candidates has residual tending to zero, then for each fixed N, its eventual correction norms are at least, for example, $\sqrt{N}/2$. Taking N arbitrarily large shows that these norms must tend to infinity. There is no uniformly bounded family of protected corrections with arbitrarily small residual.

Thus changing the iteration rule cannot fix the obstruction while preserving all the declared requirements. One must change the input, the preservation contract, or the intended ambient class. The inequality is a quantitative obstruction to a bounded repair procedure, rather than a failure report from one chosen recurrence.

26.5 Every finite operator is invertible, but the gains diverge

On $\mathbb{R}^N$, let

$$T_N = \operatorname{diag}(1, 1/2, \dots, 1/N).$$

It is invertible and $\left\|T_N^{-1}\right\| = N$ in the Euclidean norm. The upper bound follows by multiplying each coordinate by an integer at most N, and equality is attained at the last coordinate vector. Full finite rank therefore supplies no uniform gain as $N \to \infty$.

In the infinite space, $\|e_n\| = 1$ but $\|Te_n\| = 1/n$, while $\ker T = \{0\}$. A uniform estimate $\operatorname{dist}(u, \ker T) \leq C\,\|Tu\|$ would force $1 \leq C/n$ for every n, which is impossible. This agrees with the nonclosed range already established directly. The low residual directions do not reflect new exact vectors. They reflect increasingly ill-conditioned approximate directions.

26.6 A coherent inverse-limit point still exists in a larger space

At each finite level, the protected exact equations have the unique solution

$$(-1, \ldots, -1; 1) \in \mathbb{R}^N \oplus \mathbb{R}.$$

Truncation maps the solution at level $N+1$ to the one at level N. These finite solution sets are singletons, so their inverse limit is also a singleton. Its coordinate description is

$$(-1, -1, \ldots; 1) \in \mathbb{R}^{\mathbb{N}_{>0}} \times \mathbb{R}.$$

The coherent point is not missing. It lies outside the intended ℓ^2-based semantic space.

This example does not contradict finite-fiber existence. It distinguishes the ordinary inverse limit of finite views from the image of a separate comparison map that sends square-summable semantic objects to their finite prefixes. The prescribed coherent point is not in that image. Replacing ℓ^2 by all sequences makes the equation solvable, but changes the topology and the class of allowed objects.

Four statements, four answers. Every finite prefix is exactly solvable. The ordinary coherent inverse-limit point exists. Its protected Hilbert realization does not exist. Approximate protected Hilbert realizations have residuals tending to zero only at the cost of unbounded correction size. All four facts hold simultaneously in one explicit model.

27 A bounded repair budget can recover existence

The preceding example does not merely lack a convergent iteration. It has no bounded family of increasingly accurate protected corrections. This suggests a converse question. If approximate corrections stay in one bounded region, does an exact answer exist?

For bounded linear equations in a Hilbert space, the answer is yes. This is an existence statement, not a claim that the given sequence converges in norm. A stronger conclusion becomes available when the candidates are the least norm solutions of nested finite constraint systems.

The two bounded-realization theorems in this chapter have formal counterparts in the accompanying Lean project. The printed proof of bounded approximate feasibility uses the weak-subsequence lemma below, whose argument remains written background material. A formal counterpart does not by itself verify that printed proof route. The later budget, regularization-limit, noise, and solver-certificate results also have formal counterparts under their stated hypotheses, as do the explicit infinite diagonal estimates, asymptotic budget constant, and rounded noisy rate. Exact finite models have separate executable checks. The Hilbert and regularization background is classical [6]. No novelty is claimed for weak compactness or quadratic regularization.

27.1 The correction space is the protected kernel

Let E, Y, Z be real Hilbert spaces, let $d : E \to Y$ and $p : E \to Z$ be bounded linear maps, and fix a raw candidate $f \in E$. Write

$$H = \ker p, \qquad T = d|_H : H \to Y, \qquad r = df.$$

The space H is a closed Hilbert subspace with the norm inherited from E. An allowed correction is $a \in H$, and its repaired candidate is $f - a$. Thus

$$p(f - a) = pf, \qquad d(f - a) = r - Ta.$$

The protected exact problem is precisely $Ta = r$. Restricting the domain before optimizing is important. A penalty on $\|pa\|$ would only discourage a change in the protected data, not prohibit it.

The adjoint below always means $T^* : Y \to H$, taken with the inherited Hilbert metric on H. In ambient coordinates this is the projection of d^* onto H. A coordinate basis of $\ker p$ need not be orthonormal. Its Gram matrix must then be included in any calculation of correction energy.

27.2 The compactness needed for the existence argument

Lemma 27.1 (Bounded sequences have weakly convergent subsequences)**.** *Every bounded sequence (a_n) in a real Hilbert space has a subsequence converging weakly to some a. Moreover,*

$$\|a\| \le \liminf_k \|a_{n_k}\| .$$

Proof. Choose $M \ge 0$ with $\|a_n\| \le M$. If $M = 0$, every term is zero and the conclusion is immediate. Otherwise let H_0 be the closed linear span of the countably many vectors a_n. Finite rational linear combinations of those vectors form a countable dense set, so H_0 is separable. Apply Gram–Schmidt to a countable dense sequence, omitting zero remainders. This gives either a finite orthonormal basis or a countable orthonormal sequence (e_j) whose closed span is H_0. The ambient Hilbert space itself need not be separable.

Each scalar sequence $\langle a_n, e_j\rangle$ is bounded by M. For the first coordinate choose a convergent subsequence, for the second choose a subsequence of that subsequence, and continue. Taking the kth term from the kth nested subsequence, with increasing original indices, yields a_{n_k} for which every fixed coordinate converges. Call the coordinate limits c_j. In the finite-dimensional case only finitely many successive selections are required.

For each finite J, Bessel's inequality and passage to the limit give

$$\sum_{j=1}^{J} c_j^2 = \lim_k \sum_{j=1}^{J} \langle a_{n_k}, e_j\rangle^2 \le M^2.$$

The nondecreasing partial sums are therefore bounded. Orthogonality makes the partial vector sums $\sum_{j\le J} c_j e_j$ Cauchy, so completeness defines $a \in H_0$ with $\|a\| \le M$.

To verify weak convergence against $h \in H_0$, choose a finite basis combination h_J close to h. Then

$$|\langle a_{n_k} - a, h\rangle| \le |\langle a_{n_k} - a, h_J\rangle| + 2M \|h - h_J\|.$$

For a prescribed tolerance, choose h_J first to control the second term and then k to control the first by finitely many coordinate limits. For a general h in the ambient space, write $h = P_{H_0}h + h_\perp$. The orthogonal component pairs to zero with every a_{n_k} and with a, reducing to the preceding case. This proves weak convergence against every test vector.

If $a \neq 0$, testing against $a/\|a\|$ gives

$$\|a\| = \lim_k \langle a_{n_k}, a/\|a\|\rangle \le \liminf_k \|a_{n_k}\|.$$

The zero limit case follows from nonnegativity. Thus the lower-semicontinuity conclusion concerns the actual norms on the selected subsequence, not only the original bound M. □

Theorem 27.2 (Bounded approximate feasibility)**.** *For a bounded linear map $T : H \to Y$ between real Hilbert spaces and $r \in Y$, the following are equivalent:*

$$r \in \operatorname{ran} T,$$
$$\text{there is a bounded sequence } (a_n) \subset H \text{ with } \|Ta_n - r\| \longrightarrow 0.$$

If all $\|a_n\| \le M$, an exact solution exists with norm at most M.

Proof. If $Ta = r$, the constant sequence $a_n = a$ is bounded and has zero residual. For the converse, let $\|a_n\| \le M$ and $Ta_n \to r$ in norm. By the preceding lemma there is a subsequence a_{n_k} converging weakly to $a \in H$ with $\|a\| \le M$.

For every fixed $y \in Y$, the adjoint identity and weak convergence give

$$\langle Ta, y\rangle = \langle a, T^*y\rangle = \lim_k \langle a_{n_k}, T^*y\rangle = \lim_k \langle Ta_{n_k}, y\rangle.$$

The last limit equals $\langle r, y\rangle$ because norm convergence implies convergence of each continuous pairing. Thus $\langle Ta - r, y\rangle = 0$ for every y. Taking $y = Ta - r$ gives $\|Ta - r\|^2 = 0$, hence $Ta = r$.

This proof does not claim that Ta_{n_k} converges in norm merely because a_{n_k} converges weakly. Its norm convergence to r was a separate hypothesis. The boundedness of the candidate corrections is used to obtain a weakly convergent subsequence, and their accurate residuals are used to identify its image exactly. The bound $\|a\| \le M$ comes from weak lower semicontinuity, so the exact answer stays inside the same closed budget ball. □

No closed range assumption appears. The theorem concerns a single datum and a bounded family of its corrections. Closed range is the stronger, uniform statement that every attainable datum admits a correction controlled by one residual gain.

An immediate contrapositive is useful. If $r \in \overline{\operatorname{ran} T} \setminus \operatorname{ran} T$, then every sequence with $Ta_n \to r$ satisfies $\|a_n\| \to \infty$. Otherwise infinitely many terms would lie in one bounded ball, contradicting the theorem along that subsequence. This recovers the escape mechanism of Section 26 without using the diagonal formula.

27.3 Existence does not make an arbitrary sequence converge strongly

Take $H = Y = \ell^2$ and $Te_n = e_n/n$. The vectors $a_n = e_n$ have norm one and satisfy $Ta_n \to 0$. An exact answer exists, namely zero. Yet $\|e_n - e_m\| = \sqrt{2}$ for $n \neq m$, so no subsequence converges in norm.

The conclusion of Theorem 27.2 is the existence of a weak limit along a subsequence that solves the equation. It is not summable movement, a strong convergence rate, or a selection rule for the original candidates. The minimum correction for the zero datum is zero, which this particular sequence never approaches in norm.

27.4 Minimum energy finite repairs do converge strongly

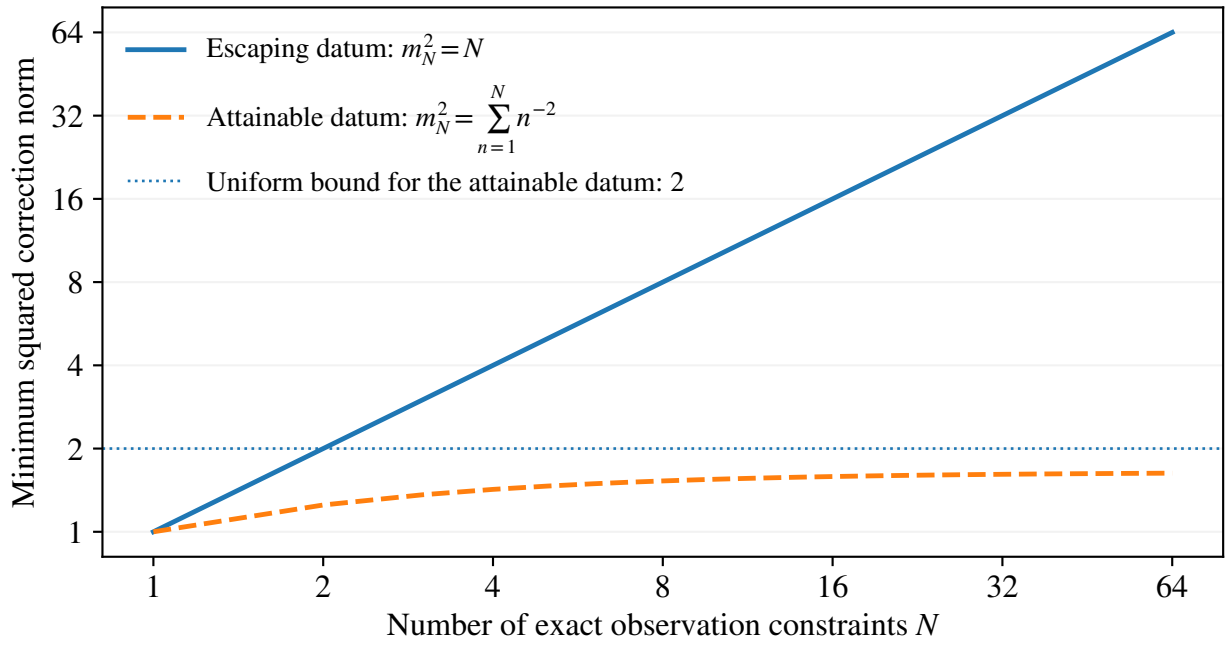


Figure 27: Minimum energy finite corrections behave differently for two data under the same diagonal operator. Bounded partial energy produces a strong Hilbert realization. Escaping partial energy accompanies a coherent product-space point outside the Hilbert space. The attainable energy lies below the explicit bound two.

Suppose now that $A_1 \supseteq A_2 \supseteq \cdots$ are nonempty closed affine subsets of a real Hilbert correction space H. A typical A_n consists of corrections satisfying the first n bounded linear observation equations. The sets are affine because the equations are linear, and closed because the observations are continuous.

Let a_n be the unique point of A_n nearest to zero, and put $m_n = \|a_n\|$. The Hilbert projection theorem supplies these points. It can also be proved directly: a sequence approaching the minimum norm is

Cauchy by the parallelogram identity, its limit belongs to the closed convex set, and the same identity gives uniqueness.

Theorem 27.3 (A uniform finite budget yields a strong realization)**.** *For the nested affine sets above, m_n is nondecreasing, and for $m \geq n$,*

$$\|a_m - a_n\|^2 = m_m^2 - m_n^2. \tag{27.1}$$

Their intersection is nonempty if and only if $\sup_n m_n < \infty$. In that case a_n converges in norm to the unique least norm point $a_* \in \bigcap_n A_n$, with

$$\|a_* - a_n\|^2 = \|a_*\|^2 - \|a_n\|^2 .$$

If the intersection is empty, then $m_n \to \infty$.

Proof. The nearest-point theorem for nonempty closed convex subsets of a Hilbert space, proved in Appendix E, supplies a_n. Write $A_n = a_n + L_n$, where L_n is its closed direction subspace. For $\ell \in L_n$ and every real t, the vector $a_n + t\ell$ belongs to A_n. Minimality implies

$$0 \leq \|a_n + t\ell\|^2 - \|a_n\|^2 = 2t\langle a_n, \ell\rangle + t^2 \|\ell\|^2 .$$

Allowing both signs of arbitrarily small t forces $\langle a_n, \ell\rangle = 0$.

If $m \geq n$, nesting gives $a_m \in A_n$, so $a_m - a_n \in L_n$. Expanding its norm therefore gives

$$\|a_m\|^2 = \|a_n + (a_m - a_n)\|^2 = \|a_n\|^2 + \|a_m - a_n\|^2 .$$

This is the exact energy identity. It also makes m_n^2 nondecreasing.

Suppose these energies are bounded, and let $L = \lim_n m_n^2$. Given $\epsilon > 0$, choose N with $L - m_N^2 < \epsilon^2$. For $m \geq n \geq N$,

$$\|a_m - a_n\|^2 = m_m^2 - m_n^2 \leq L - m_N^2 < \epsilon^2.$$

Thus (a_n) is Cauchy and has a norm limit a_*. For each fixed n, all terms from that index onward belong to the closed set A_n, so $a_* \in A_n$. This places a_* in their entire intersection.

If a lies in the intersection, then $\|a\| \geq \|a_n\|$ for every n. Norm continuity gives $\|a\| \geq \|a_*\|$. Uniqueness follows from the nearest-point theorem applied to the intersection, which is closed and convex and now known to be nonempty. Alternatively, the midpoint of two different minimizers would have strictly smaller norm by the parallelogram identity. Taking $m \to \infty$ in the finite energy identity yields

$$\|a_* - a_n\|^2 = \|a_*\|^2 - \|a_n\|^2 .$$

Conversely, any element of the intersection bounds every minimum norm from above. Therefore an empty intersection forces the nondecreasing sequence m_n to be unbounded. A nondecreasing unbounded nonnegative sequence tends to infinity, giving the last conclusion. □

The finite answers need not extend one another coordinate by coordinate. Their shared feature is minimum energy in one common ambient Hilbert space. This is the extra geometry that forces their differences to be orthogonal to the earlier minimizers.

For the attainable diagonal datum $r_n = 1/n^2$ under $Tu = (u_n/n)$, with $n \geq 1$, the first N equations force correction coordinate n to equal $1/n$ for $n \leq N$. Their minimum energy completion is zero beyond N, and

$$m_N^2 = \sum_{n=1}^{N} \frac{1}{n^2} \leq 2.$$

These corrections converge strongly to $(1/n)$. For the escaping datum $r_n = 1/n$, the corresponding energy is $m_N^2 = N$. Every finite equation is again solved, but the budget test now rules out an ambient solution.

Finite consistency is not enough by itself. For nested closed affine problems in one Hilbert space, a uniform bound on the minimum correction norms is exactly what upgrades finite feasibility to ambient feasibility. This is a datum-specific test, not a uniform inverse bound for the operator.

28 The attainable residual at a fixed correction budget

The unrestricted floor $\operatorname{dist}(r, \operatorname{ran} T)$ forgets how large a correction must be. To retain that information, fix a nonnegative budget M and define

$$\beta_{T,r}(M) = \inf_{\|a\| \le M} \|r - Ta\| .$$

The norm is the actual correction norm in $H = \ker p$. It is not the norm of arbitrary coordinates used to represent that kernel. The function β describes the smallest residual compatible with a declared amount of allowed change.

28.1 At every finite budget the residual minimum exists

Theorem 28.1 (Attainment and shape of the budget profile)**.** *For a bounded linear map between real Hilbert spaces, the infimum defining $\beta(M)$ is attained for every $M \ge 0$. The profile is nonincreasing and convex. For $0 \le M_1 \le M_2$,*

$$0 \le \beta(M_1) - \beta(M_2) \le \|T\| \, (M_2 - M_1).$$

Writing $C = \overline{\operatorname{ran} T}$, $b = P_C r$, and $q = \|r - b\|$, one has

$$\lim_{M \to \infty} \beta(M) = q.$$

The equality $\beta(M) = q$ at a finite budget is equivalent to the existence of a with $Ta = b$ and $\|a\| \le M$. In particular,

$$\beta(M) = 0 \quad \Longleftrightarrow \quad \exists a \in H : \; Ta = r, \; \|a\| \le M.$$

Proof. The feasible ball is nonempty because it contains zero, and the infimum lies between zero and $\|r\|$. Choose a_n in that ball with $\|r - Ta_n\| \to \beta(M)$. A weakly convergent subsequence has limit a with $\|a\| \le M$. Bounded linearity makes $r - Ta_{n_k}$ converge weakly to $r - Ta$. Applying weak lower semicontinuity in the target gives

$$\|r - Ta\| \le \liminf_k \|r - Ta_{n_k}\| = \beta(M).$$

The reverse inequality holds by the definition of the infimum, so a attains it. When $M = 0$ the only feasible correction is zero, and the same conclusion reads $\beta(0) = \|r\|$.

For convexity, take minimizers a_1, a_2 at budgets M_1, M_2 and $t \in [0, 1]$. The candidate $a_t = ta_1 + (1 - t)a_2$ satisfies

$$\|a_t\| \le tM_1 + (1 - t)M_2,$$

and its residual is $t(r - Ta_1) + (1-t)(r - Ta_2)$. The triangle inequality therefore gives

$$\beta(tM_1 + (1-t)M_2) \le t\beta(M_1) + (1-t)\beta(M_2).$$

Monotonicity follows from inclusion of balls. If $0 \le M_1 \le M_2$ and $M_2 > 0$, choose a minimizer a_2 at M_2 and put $a_1' = (M_1/M_2)a_2$. Then $\|a_1'\| \le M_1$ and $\|a_2 - a_1'\| \le M_2 - M_1$, so

$$\beta(M_1) \le \left\|r - Ta_1'\right\| \le \beta(M_2) + \|T\|\,(M_2 - M_1).$$

Together with monotonicity this proves the Lipschitz estimate. The omitted case $M_2 = 0$ has $M_1 = 0$ and is equality.

Since $r - b$ is orthogonal to C and $b - Ta \in C$, every correction satisfies

$$\|r - Ta\|^2 = q^2 + \|b - Ta\|^2. \tag{28.1}$$

Thus $\beta(M) \ge q$. Given $\epsilon > 0$, range density in C supplies one correction a with $\|b - Ta\| < \epsilon$. For every $M \ge \|a\|$ it follows that

$$q \le \beta(M) \le \sqrt{q^2 + \epsilon^2}.$$

First letting the budget increase and then letting $\epsilon \downarrow 0$ proves the limit.

Finally, if $\beta(M) = q$, choose an attaining correction in that ball. The orthogonal identity forces $\|b - Ta\|^2 = 0$, hence $Ta = b$. Conversely such a correction attains residual q. Setting $q = 0$ proves the exact-feasibility criterion. Attainment at fixed budget is essential to the forward implication: a zero unrestricted infimum would not itself supply a preimage. □

Although ran T may not be closed, the image $T(\{a : \|a\| \le M\})$ of each closed budget ball is closed and convex. To see closedness, take a convergent sequence in that image, lift it to the bounded ball, and apply the weak limit argument. The full range is the increasing union of these closed images. Passing to an unlimited budget can therefore lose closedness even though every fixed-budget optimization is well posed.

The nearest reachable image $y_M = P_{T(\{a:\|a\|\le M\})}r$ is unique, so the best residual vector $r - y_M$ is unique as well. Corrections mapping to y_M can differ by vectors in $\ker T$. Among those corrections there is a unique least norm correction. These two uniqueness statements should not be confused with injectivity of T.

28.2 A dual witness can rule out every correction in the budget

Theorem 28.2 (A sharp budget certificate)**.** *For every $M \ge 0$,*

$$\beta_{T,r}(M) = \max_{\|z\|\le 1} \big(\langle r, z\rangle - M\,\|T^*z\|\big). \tag{28.2}$$

Thus any single z with $\|z\| \le 1$ and

$$\langle r, z\rangle - M\,\|T^*z\| > \varepsilon$$

certifies that no correction of norm at most M has residual at most ε.

Proof. For every $\|a\| \le M$ and $\|z\| \le 1$, the adjoint identity and Cauchy–Schwarz give

$$\begin{aligned}\langle r, z\rangle - M\,\|T^*z\| &= \langle r - Ta, z\rangle + \langle a, T^*z\rangle - M\,\|T^*z\| \\ &\le \langle r - Ta, z\rangle \le \|r - Ta\|.\end{aligned}$$

Taking the infimum over corrections proves that every proposed dual value is at most $\beta(M)$. This part alone already proves the impossibility certificate.

To prove sharpness, let a_M be an attaining correction from Theorem 28.1 and put $v = r - Ta_M$. If $v = 0$, then $\beta(M) = 0$ and $z = 0$ attains the dual bound. Assume henceforth $v \neq 0$ and set $z = v/\|v\|$.

Fix any feasible a and put $h = a - a_M$. The segment $a_M + th$ is feasible for $0 \leq t \leq 1$. Minimality gives, for $t > 0$,

$$0 \leq \|v - tTh\|^2 - \|v\|^2 = -2t\langle v, Th\rangle + t^2 \|Th\|^2 .$$

Divide by t and let $t \downarrow 0$. It follows that $\langle v, T(a - a_M)\rangle \leq 0$, hence

$$\langle T^*z, a\rangle \leq \langle T^*z, a_M\rangle \quad (\|a\| \leq M).$$

The support value of a Hilbert ball can be computed directly. If $T^*z \neq 0$ and $M > 0$, it is attained at $a = MT^*z/\|T^*z\|$ and equals $M\|T^*z\|$. If either factor is zero, every feasible pairing is zero and the same formula holds. Thus

$$\langle T^*z, a_M\rangle = M\|T^*z\| .$$

Substituting this equality gives

$$\langle r, z\rangle - M\|T^*z\| = \langle r - Ta_M, z\rangle = \|v\| = \beta(M).$$

We have constructed an attaining witness and thereby justified the word “maximum.” The proof uses one-sided variations in a convex ball, not an unproved exchange of an infimum and a supremum. □

In finite dimensions this is a checkable inequality, not merely an optimization prescription. A program may propose the correction or the witness. The displayed inequality can be verified independently of the proposal procedure. In a protected ambient calculation, use the restricted adjoint, not d^* without projection onto $\ker p$.

28.3 A two-dimensional calculation of the supporting witness

Take

$$T = \begin{pmatrix} 1 & 0 \\ 0 & 1/2 \end{pmatrix}, \qquad r = (1,1), \qquad a = (1/2, 2/5), \qquad v = r - Ta = (1/2, 4/5).$$

Then $T^*v = a$. At the budget $M = \|a\| = \sqrt{41}/10$, the normalized vector $z = v/\|v\|$ gives

$$\langle r, v\rangle - \|a\|\,\|T^*v\| = \frac{13}{10} - \frac{41}{100} = \frac{89}{100} = \|v\|^2 .$$

Consequently a is a globally optimal budgeted correction and $\beta(M) = \sqrt{89}/10$. It is not necessary to compare it with a dense numerical grid of competing vectors. The supporting witness proves optimality over the entire ball.

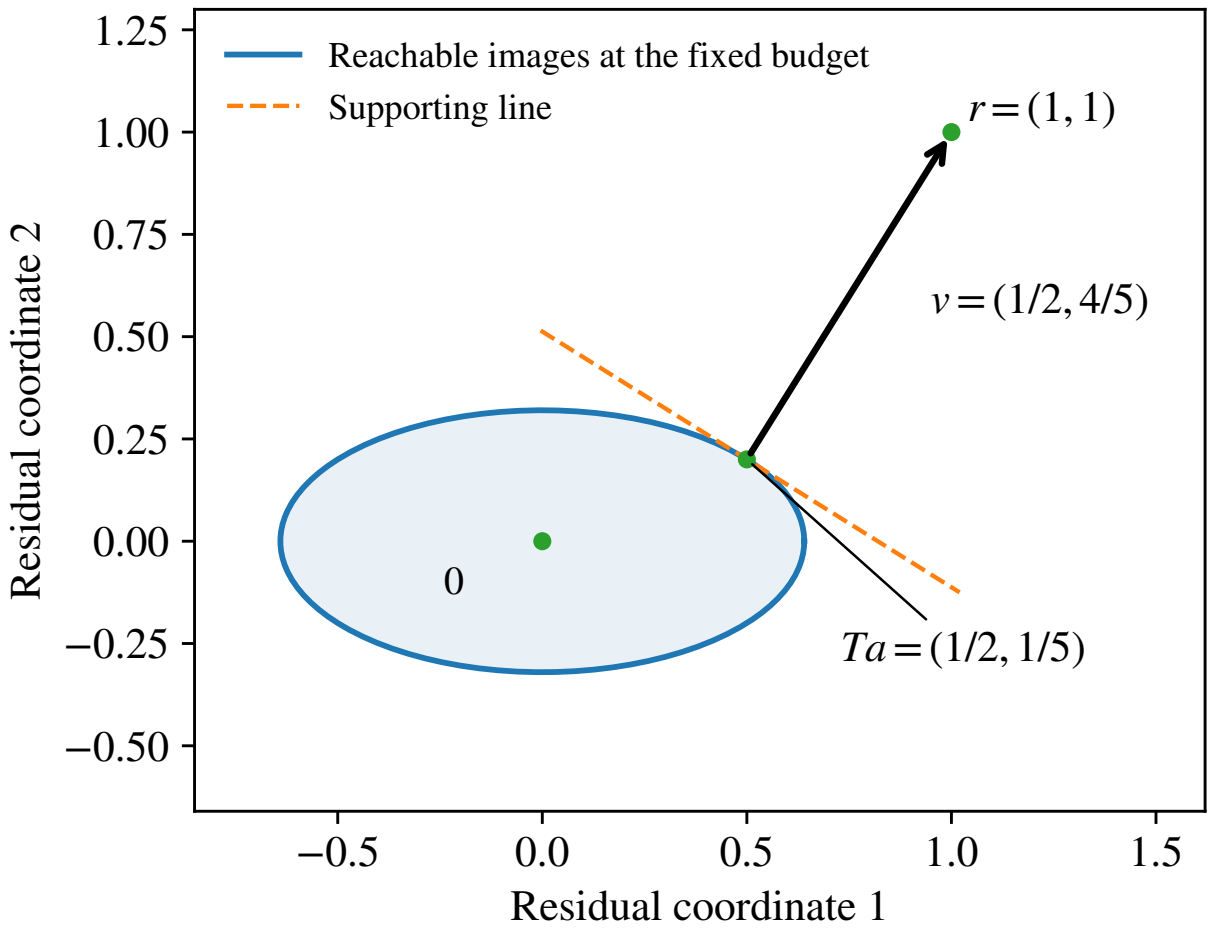


Figure 28: The image of a Euclidean correction ball is an ellipse. The nearest reachable point has a residual normal to a supporting line. That same residual, normalized, is the dual witness in the exact budget formula. The plotted point and budget are the rational example in the text.

28.4 How measurement and model uncertainty affect the budget

Theorem 28.3 (A robust budget profile). *Suppose the bounded linear map $\widetilde{T} : H \to Y$ and the datum $\widetilde{r} \in Y$ satisfy*

$$\left\|T - \widetilde{T}\right\| \le \eta, \qquad \|r - \widetilde{r}\| \le \delta, \qquad \eta, \delta \ge 0.$$

For every $M \ge 0$,

$$\left|\beta_{T,r}(M) - \beta_{\widetilde{T},\widetilde{r}}(M)\right| \le \delta + \eta M.$$

In particular, for $\|z\| \le 1$,

$$\beta_{T,r}(M) \ge \langle \widetilde{r}, z\rangle - M\left\|\widetilde{T}^* z\right\| - \delta - \eta M.$$

Proof. Fix a feasible correction a. The two residual vectors have difference

$$(r - Ta) - (\widetilde{r} - \widetilde{T}a) = (r - \widetilde{r}) - (T - \widetilde{T})a,$$

whose norm is at most $c = \delta + \eta M$. The reverse triangle inequality therefore gives both inequalities

$$\|r - Ta\| \le \left\|\widetilde{r} - \widetilde{T}a\right\| + c, \qquad \left\|\widetilde{r} - \widetilde{T}a\right\| \le \|r - Ta\| + c.$$

The domain of feasible corrections is identical in the two problems. Taking the infimum of each inequality over that common domain yields

$$\beta_{T,r}(M) \le \beta_{\widetilde{T},\widetilde{r}}(M) + c, \qquad \beta_{\widetilde{T},\widetilde{r}}(M) \le \beta_{T,r}(M) + c.$$

These are precisely the absolute-value bound. One could equivalently evaluate at the two attained minimizers, but no relation between those minimizers is needed.

For any admissible z, Theorem 28.2 for the nominal problem gives

$$\beta_{\widetilde{T},\widetilde{r}}(M) \ge \langle \widetilde{r}, z\rangle - M \left\| \widetilde{T}^* z \right\|.$$

Combine this with the lower half of the preceding comparison. The result is a certificate against every correction in the original fixed ball, reduced by the explicit uncertainty allowance. If the correction space itself changed, the common-domain step would fail and this proof would no longer apply. □

This estimate identifies a practical boundary of increasingly large corrections. A model error η is amplified by the correction budget even if the measured residual is tiny. The candidate-level inequality is

$$\|r - Ta\| \le \left\| \widetilde{r} - \widetilde{T}a \right\| + \delta + \eta \, \|a\|.$$

The theorem compares maps on one fixed correction space. It does not cover an uncertain preservation map whose kernel itself changes.

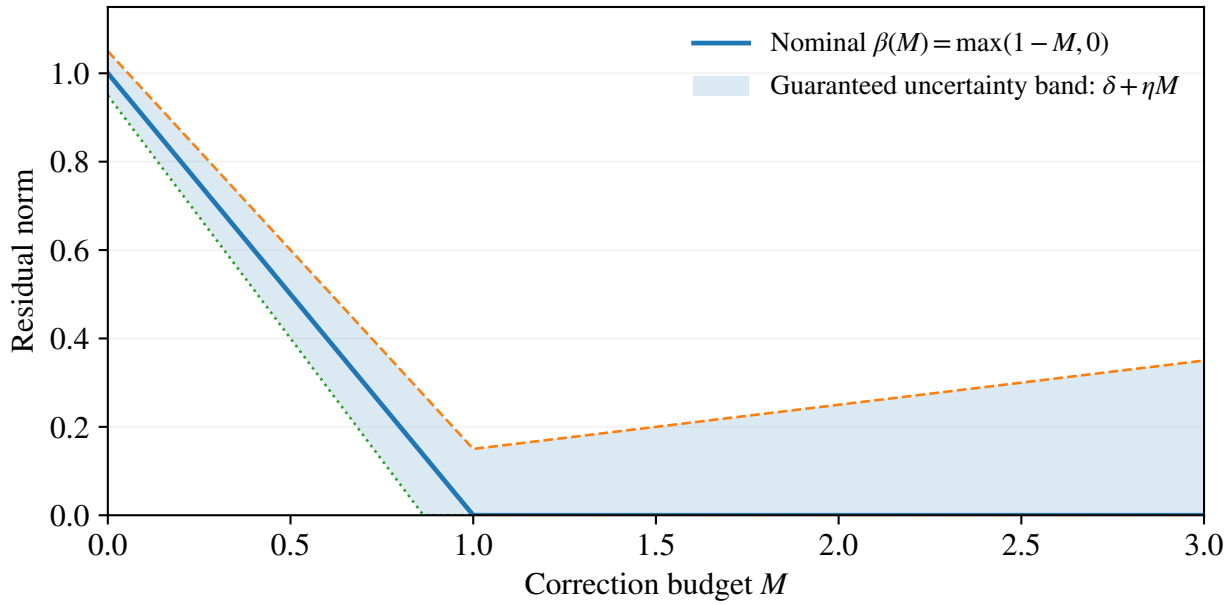


Figure 29: The nominal budget profile is surrounded by the rigorous uncertainty band $\delta + \eta M$. A large budget can improve the nominal fit while enlarging the uncertainty in that conclusion. The example uses $T = 1$, $r = 1$, $\delta = 1/20$, and $\eta = 1/10$.

29 Regularization selects a protected correction

A fixed budget asks how small the residual can become without exceeding an allowed correction size. Quadratic regularization reverses the emphasis. It asks for the best compromise between the residual and the size of the change. The resulting problem has a unique answer even when the exact protected equation has no solution.

That last sentence is both its usefulness and its limitation. The existence of a regularized answer is not a proof of exact realizability. The limiting behavior, rather than the mere existence of the minimizer, supplies the diagnosis.

29.1 The normal equation is derived in the protected space

For $\alpha > 0$, define on $H = \ker p$

$$J_\alpha(a) = \|Ta - r\|^2 + \alpha \|a\|^2 .$$

Both norms are the declared Hilbert norms. All minimizers below are corrections in H, so preservation is exact for every parameter, not just in the limit. This is ordinary quadratic Tikhonov regularization applied to the actual protected correction operator. The standard regularization background is discussed in [6]. The proof here does not require T to be compact.

Theorem 29.1 (A unique regularized correction)**.** *For every $\alpha > 0$, there is a unique minimizer*

$$a_\alpha = (T^*T + \alpha I)^{-1} T^* r.$$

It satisfies $T^*(r - Ta_\alpha) = \alpha a_\alpha$. For every $a \in H$,

$$J_\alpha(a) - J_\alpha(a_\alpha) = \|T(a - a_\alpha)\|^2 + \alpha \|a - a_\alpha\|^2 . \tag{29.1}$$

Proof. Step 1: the normal operator has an actual bounded inverse. Let $B = T^*T + \alpha I : H \to H$. It is bounded and self-adjoint. For every $h \in H$,

$$\langle Bh, h\rangle = \|Th\|^2 + \alpha \|h\|^2 \geq \alpha \|h\|^2 .$$

If $h \neq 0$, Cauchy–Schwarz and division by $\|h\|$ give $\|Bh\| \geq \alpha \|h\|$. The same inequality holds at zero. It implies injectivity of B.

To prove closedness of its range, suppose $Bh_n \to y$. Applying the lower bound to $h_n - h_m$ gives

$$\|h_n - h_m\| \leq \alpha^{-1} \|Bh_n - Bh_m\| .$$

Thus h_n is Cauchy. Completeness of H supplies a limit h, and continuity gives $Bh = y$. If y is orthogonal to the range of B, then for every h,

$$0 = \langle Bh, y\rangle = \langle h, By\rangle.$$

Taking $h = By$ gives $By = 0$, and injectivity gives $y = 0$. A closed subspace with zero orthogonal complement is the whole Hilbert space, so $\operatorname{ran} B = H$.

The inverse therefore exists. Its linearity follows from uniqueness of preimages, and its bound is $\left\|B^{-1}y\right\| \leq \alpha^{-1} \|y\|$. This construction does not assume closedness of $\operatorname{ran} T$. It uses the positive term αI in the normal operator.

Step 2: define the optimizer and prove its normal equation. Set $a_\alpha = B^{-1}T^*r$. Then $Ba_\alpha = T^*r$, or equivalently

$$T^*(Ta_\alpha - r) + \alpha a_\alpha = 0.$$

Because the construction takes place in $H = \ker p$, it already satisfies the exact protection requirement.

Step 3: prove global minimality, rather than just stationarity. Write any proposed correction as $a = a_\alpha + h$ and set $s = Ta_\alpha - r$. Expanding both squared norms gives

$$\begin{aligned} J_\alpha(a_\alpha + h) - J_\alpha(a_\alpha) &= \|Th\|^2 + \alpha \|h\|^2 + 2\langle s, Th\rangle + 2\alpha\langle a_\alpha, h\rangle \\ &= \|Th\|^2 + \alpha \|h\|^2 + 2\langle T^* s + \alpha a_\alpha, h\rangle \\ &= \|Th\|^2 + \alpha \|h\|^2 . \end{aligned}$$

This is the asserted energy identity. It is nonnegative and, since $\alpha > 0$, vanishes only for $h = 0$. Therefore a_α is the unique global minimizer. In the zero correction space, the same construction and identities reduce to its single element, zero. □

A declared regularization parameter is part of the problem. A calculation that changes α is solving a different optimization problem, even if the protected data and forward operator stay fixed. On the other hand, the preservation condition itself does not change with α.

29.2 Coordinate bases must retain the correction metric

In finite dimensions, choose orthonormal coordinates in the ambient candidate and residual spaces. Let B have linearly independent columns spanning $\ker p$. The columns of B need not be orthonormal. Write $a = Bc$. The coefficient vector c is a representation of the correction, not necessarily an isometric one. Set

$$S = dB, \qquad G = B^{\mathsf{T}} B.$$

The correct coordinate equation is

$$(S^{\mathsf{T}} S + \alpha G)c = S^{\mathsf{T}} r. \tag{29.2}$$

The matrix G is positive definite because the columns are independent. Replacing G by the identity changes the meaning of correction size unless the chosen basis is orthonormal.

For example, let $d(a_1, a_2) = a_1$, $p(a_1, a_2) = a_1 - a_2$, $r = 1$, and $\alpha = 1$. A basis of $\ker p$ is the column $B = (1, 1)^{\mathsf{T}}$. Here $S = 1$ and $G = 2$, so the correction is $(1/3, 1/3)$. Penalizing the coordinate with $G = 1$ instead gives $(1/2, 1/2)$. The latter is not the optimizer for the ambient Euclidean norm. The difference is a metric change, not a discrepancy between equivalent calculations.

The executable example in this edition constructs the nullspace basis over the rationals, retains its Gram matrix, solves (29.2), and returns both the correction and an exact normal-equation residual. It also handles a zero-dimensional protected kernel without pretending that a nonzero correction exists.

29.3 Parameter monotonicity and the fixed-budget problem

Proposition 29.2 (Regularized answers lie on the budget frontier)**.** *If $0 < \alpha < \gamma$, then*

$$\|a_\alpha\| \geq \|a_\gamma\|, \qquad \|r - Ta_\alpha\| \leq \|r - Ta_\gamma\| .$$

At $M_\alpha = \|a_\alpha\|$,

$$\beta(M_\alpha) = \|r - Ta_\alpha\| .$$

When $v_\alpha = r - Ta_\alpha \neq 0$, its normalization is an optimizing witness in Theorem 28.2 at that budget.

Proof. Let $A_\alpha = \|a_\alpha\|^2$ and $R_\alpha = \|r - Ta_\alpha\|^2$. Comparing the two admissible minimizers in each other's objectives gives

$$R_\alpha + \alpha A_\alpha \le R_\gamma + \alpha A_\gamma, \qquad R_\gamma + \gamma A_\gamma \le R_\alpha + \gamma A_\alpha.$$

Add these inequalities and cancel the residual terms. The result is

$$(\gamma - \alpha)(A_\alpha - A_\gamma) \ge 0.$$

Since $\gamma > \alpha$, it follows that $A_\alpha \ge A_\gamma$. Return to the first comparison to obtain

$$R_\alpha - R_\gamma \le \alpha(A_\gamma - A_\alpha) \le 0.$$

Taking square roots, which preserves order for nonnegative numbers, proves both monotonicity statements.

At $M_\alpha = \|a_\alpha\|$, the vector a_α is budget-feasible. If another correction a in that ball had strictly smaller residual norm, then

$$J_\alpha(a) = \|r - Ta\|^2 + \alpha \|a\|^2 < \|r - Ta_\alpha\|^2 + \alpha \|a_\alpha\|^2,$$

contradicting its defining minimality. Thus it attains the budget frontier.

For the witness statement, let $v = r - Ta_\alpha \ne 0$. The normal equation gives $T^*v = \alpha a_\alpha$. Also $r = v + Ta_\alpha$, hence

$$\langle r, v\rangle = \|v\|^2 + \langle a_\alpha, T^*v\rangle = \|v\|^2 + \alpha \|a_\alpha\|^2.$$

For $z = v/\|v\|$, the support penalty is

$$M_\alpha \|T^*z\| = \frac{\alpha \|a_\alpha\|^2}{\|v\|}.$$

Subtracting it from $\langle r, z\rangle$ yields $\|v\| = \beta(M_\alpha)$, so this is an optimizing witness. This calculation also covers a zero correction a_α with nonzero residual. If $v = 0$, normalization is unavailable, but the zero witness attains the zero budget residual. □

This links a penalized calculation to a hard-budget certificate. It does not say that every larger budget must produce a different answer. Once the least-squares floor is attained, excess budget can be irrelevant. If $T^*r = 0$, every regularized correction is zero and the entire profile is constant.

29.4 The limit separates projection from actual realizability

Keep $C = \overline{\operatorname{ran} T}$, $b = P_C r$, and $q = \|r - b\|$. Since $r - b \in \ker T^*$, the regularized corrections depend on b, not on the component that is orthogonal to the closure of all attainable residuals.

Theorem 29.3 (The regularization limit alternative)**.** *As $\alpha \downarrow 0$,*

$$Ta_\alpha \longrightarrow b, \qquad \|r - Ta_\alpha\| \longrightarrow q, \qquad \alpha \|a_\alpha\|^2 \longrightarrow 0.$$

If $b \in \operatorname{ran} T$, let $a^\dagger$ be its unique minimum norm preimage. Then

$$a_\alpha \longrightarrow a^\dagger \quad \text{in norm.}$$

If $b \notin \operatorname{ran} T$, then $\|a_\alpha\| \to \infty$.

Proof. Step 1: identify the unavoidable component. Since $r - b \perp \overline{\operatorname{ran} T}$, the adjoint identity gives $T^*(r-b) = 0$. Thus $a_\alpha(r) = a_\alpha(b)$. For every correction the orthogonal decomposition is

$$\|r - Ta\|^2 = q^2 + \|b - Ta\|^2 .$$

In particular both the residual square and the minimum objective are at least q^2.

Step 2: take the small-parameter limit in the correct order. Let $\epsilon > 0$. Since b lies in the closure of the range, choose a single $a \in H$ with $\|r - Ta\|^2 < q^2 + \epsilon$. This a is chosen before α. Minimality gives

$$q^2 \le \|r - Ta_\alpha\|^2 \le J_\alpha(a_\alpha) \le q^2 + \epsilon + \alpha \|a\|^2 .$$

For sufficiently small α, the last term is below $q^2 + 2\epsilon$ (and if $a = 0$ it is already below $q^2 + \epsilon$). Hence both intervening quantities tend to q^2. Their difference is $\alpha \|a_\alpha\|^2$, which tends to zero. The orthogonal decomposition now implies $\|Ta_\alpha - b\|^2 \to 0$, and taking square roots gives the first two limits.

Step 3: when the projected datum is attained, obtain a uniform bound. Suppose $b \in \operatorname{ran} T$. The preimage set is a nonempty closed affine subset, so projection of zero supplies its unique minimum-norm element $a^\dagger$. Comparing with it and subtracting q^2 gives

$$\|Ta_\alpha - b\|^2 + \alpha \|a_\alpha\|^2 \le \alpha \left\|a^\dagger\right\|^2 .$$

Thus $\|a_\alpha\| \le \left\|a^\dagger\right\|$ for every $\alpha > 0$. This bound does not follow merely from the previously proved vanishing weighted penalty.

Step 4: strengthen weak convergence to norm convergence. Take an arbitrary sequence $\alpha_n \downarrow 0$, or more generally positive $\alpha_n \to 0$. The bound supplies a weakly convergent subsequence $a_{\alpha_{n_k}} \rightharpoonup a$. Since its images converge in norm to b, the adjoint pairing argument of Theorem 27.2 yields $Ta = b$. Weak lower semicontinuity gives

$$\|a\| \le \liminf_k \left\|a_{\alpha_{n_k}}\right\| \le \left\|a^\dagger\right\| .$$

Minimality and uniqueness of $a^\dagger$ force $a = a^\dagger$. Moreover the same lower bound, together with the uniform upper bound, gives convergence of the norms to $\left\|a^\dagger\right\|$. In a real Hilbert space,

$$\left\|a_{\alpha_{n_k}} - a^\dagger\right\|^2 = \left\|a_{\alpha_{n_k}}\right\|^2 + \left\|a^\dagger\right\|^2 - 2\langle a_{\alpha_{n_k}}, a^\dagger\rangle \longrightarrow 0.$$

Thus this subsequence converges strongly. If the full positive-parameter family failed to converge strongly, there would be an $\epsilon > 0$ and parameters $\alpha_n \to 0$ whose corrections all stayed at distance at least ϵ from $a^\dagger$. The argument just given would produce a strongly convergent subsequence of them, a contradiction. This proves the limit for the entire parameter family.

Step 5: in the unattained case the norms tend to infinity. If $b \notin \operatorname{ran} T$ and the norms failed to tend to infinity, there would exist $M < \infty$ and positive $\alpha_n \to 0$ with $\|a_{\alpha_n}\| \le M$. Their images still tend to b by Step 2. Theorem 27.2 would then put b in the actual range, a contradiction. This proves divergence of the norms, not just unboundedness somewhere along the path. □

The vanishing penalty $\alpha \|a_\alpha\|^2$ is not a boundedness certificate for a_α. Its coefficient is tending to zero. In the escaping case, the penalty and residual can both tend to zero while the actual corrections diverge.

Residual floor	Is b attained?	Limit of the protected corrections
$q = 0$	Yes	Strong convergence to the nearest exact protected correction.
$q = 0$	No	Residual vanishes, correction norms diverge, no exact protected answer.
$q > 0$	Yes	Strong convergence to the least-squares correction of minimum norm, not an exact answer.
$q > 0$	No	A positive residual floor and diverging correction norms occur together.

The last case is not hypothetical. Let $Tu = (u_n/n)$ on ℓ^2, define $Su = (Tu, 0)$ into $\ell^2 \oplus \mathbb{R}$, and take $r = ((1/n), 1)$. Then $b = ((1/n), 0)$ is in the closure but not the range, while $q = 1$. The same escaping corrections as before coexist with an irreducible orthogonal residual. Positive floor and unbounded correction cost are independent features.

29.5 Noise is controlled at a fixed parameter

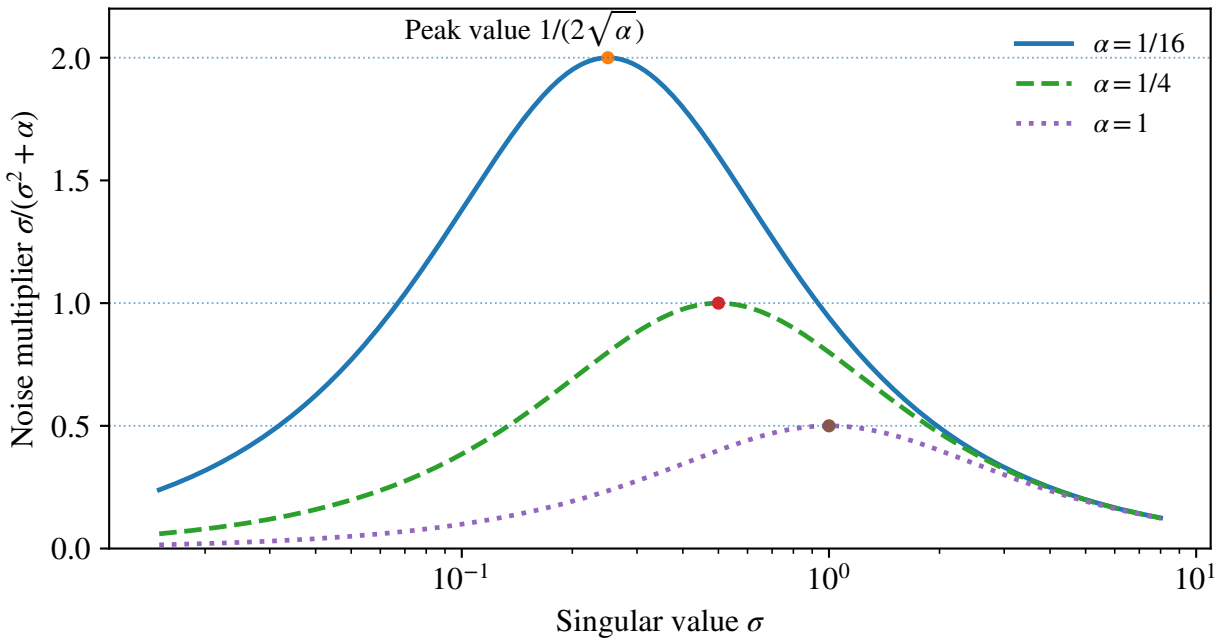


Figure 30: For one singular value σ, the response to residual-data noise is $\sigma/(\sigma^2 + \alpha)$. Its maximum occurs at $\sigma = \sqrt{\alpha}$ and equals $1/(2\sqrt{\alpha})$. The plot evaluates this formula and the universal horizontal bounds at three explicit positive parameters.

Theorem 29.4 (Data sensitivity and a convergent parameter choice)**.** *For two residual data vectors $r, \tilde{r}$ with $\|r - \tilde{r}\| \le \delta$ and the same T, α,*

$$\|a_\alpha(r) - a_\alpha(\tilde{r})\| \le \frac{\delta}{2\sqrt{\alpha}}, \qquad \|T(a_\alpha(r) - a_\alpha(\tilde{r}))\| \le \delta.$$

The coefficient $1/(2\sqrt{\alpha})$ *is sharp over bounded linear Hilbert operators. If* $b = P_C r \in \operatorname{ran} T$, $\alpha(\delta) \to 0$, *and* $\delta/\sqrt{\alpha(\delta)} \to 0$, *then the noisy corrections converge strongly to* $a^\dagger$ *as* $\delta \to 0$.

Proof. Both corrections are constructed at the same positive parameter and in the same Hilbert correction space. Let $h = a_\alpha(\widetilde{r}) - a_\alpha(r)$ and $e = \widetilde{r} - r$. Subtract the two normal equations to obtain

$$(T^*T + \alpha I)h = T^*e.$$

Pair with h and use the adjoint identity:

$$\|Th\|^2 + \alpha \|h\|^2 = \langle e, Th\rangle \le \delta \|Th\| .$$

Put $t = \|Th\| \ge 0$. If $t > 0$, dropping the nonnegative penalty and dividing by t gives $t \le \delta$. If $t = 0$, that inequality is already true. Completing the scalar square gives

$$\alpha \|h\|^2 \le \delta t - t^2 = \frac{\delta^2}{4} - (t - \delta/2)^2 \le \frac{\delta^2}{4}.$$

Because $\alpha > 0$, taking square roots proves the correction bound. This also explains why the coefficient is $1/(2\sqrt{\alpha})$ rather than the coarser $\|T\| /\alpha$ obtained by bounding the inverse alone.

For sharpness at a fixed parameter, take $H = Y = \mathbb{R}$ and $T = \sqrt{\alpha}\, I$. The solution map is multiplication by

$$\frac{\sqrt{\alpha}}{\alpha + \alpha} = \frac{1}{2\sqrt{\alpha}}.$$

Any nonzero data perturbation attains the first bound. This is worst-case sharpness across operators at that parameter, not a claim of equality for every fixed operator and every parameter.

For the limiting statement, insert the exact-data correction at the selected parameter:

$$\left\| a_{\alpha(\delta)}(\widetilde{r}) - a^\dagger \right\| \le \frac{\delta}{2\sqrt{\alpha(\delta)}} + \left\| a_{\alpha(\delta)}(r) - a^\dagger \right\| .$$

The first term tends to zero by the noise/parameter hypothesis. The second tends to zero by the attained-projection case of Theorem 29.3, because the selected parameter tends to zero. Both statements require the same fixed operator and Hilbert metrics. The conclusion is a least-squares correction when the unavoidable residual component is nonzero. □

The theorem does not manufacture an exact solution when r is outside the range. For the positive-floor attainable case it converges to a least-squares correction. For the unattainable projected datum it cannot produce a bounded exact-data limit. The range hypothesis must therefore precede the parameter-choice conclusion.

Regularization gives a unique, exactly protected answer at every positive parameter. Whether those answers converge to an exact realization, converge only to a best approximation, or escape is determined by the projected datum's actual range membership. A smooth regularized path does not erase that distinction.

30 A complete regularization diagnostic for the escaping model

The abstract alternatives become explicit for the diagonal model. They also answer a natural objection to the hard prefix corrections: perhaps a smoother correction could fit increasingly well without escaping. It can fit more efficiently, but it cannot avoid the obstruction.

30.1 The entire regularized family is explicit

Use $Tu = (u_n/n)$ and $r = w = (1/n)$ as in Section 26. Choose $\alpha = m^{-2}$ for an integer $m \geq 1$. The normal equation in each coordinate gives

$$a_n^{(m)} = \frac{m^2}{m^2+n^2}, \qquad v_n^{(m)} = r_n - (Ta^{(m)})_n = \frac{n}{m^2+n^2}.$$

The repaired ambient candidate is $(-a^{(m)}, 1)$. Its protected coordinate is one for every m.

Proposition 30.1 (Energy, residual, and certified tails)**.** *The displayed $a^{(m)}$ belongs to ℓ^2 and is the unique regularized correction at $\alpha = m^{-2}$. Set*

$$A_m = \left\|a^{(m)}\right\|^2 = \sum_{n\geq 1} \frac{m^4}{(m^2+n^2)^2}, \qquad B_m = \left\|v^{(m)}\right\|^2 = \sum_{n\geq 1} \frac{n^2}{(m^2+n^2)^2}.$$

Then

$$\frac{m}{4} \leq A_m \leq \frac{4m}{3}, \qquad \frac{1}{4(m+1)} \leq B_m \leq \frac{2}{m}. \tag{30.1}$$

For the sums $A_{m,K}, B_{m,K}$ through any integer $K \geq m$,

$$0 \leq A_m - A_{m,K} \leq \frac{m^4}{3K^3}, \qquad 0 \leq B_m - B_{m,K} \leq \frac{1}{K}. \tag{30.2}$$

Proof. For $n > m$, one has $(a_n^{(m)})^2 \leq m^4/n^4$, so the correction sequence is square-summable. The diagonal operator is bounded with norm one, and $w = (1/n) \in \ell^2$, hence the residual sequence is also in ℓ^2. The operator is self-adjoint, and the coordinate equality

$$\frac{1}{n}\frac{n}{m^2+n^2} = m^{-2}\frac{m^2}{m^2+n^2}$$

proves $T^*v^{(m)} = \alpha a^{(m)}$ as an equality of Hilbert vectors. Theorem 29.1 therefore identifies this actual infinite vector as the unique minimizer, rather than just a coordinatewise candidate.

For $1 \leq n \leq m$, $m^2/(m^2+n^2) \geq 1/2$. There are m such coordinates, so $A_m \geq m/4$. For the upper bound, the first m squares are at most one, and the decreasing function x^{-4} gives

$$\sum_{n>m} \frac{1}{n^4} \leq \int_m^\infty x^{-4}\,dx = \frac{1}{3m^3}.$$

Thus $A_m \leq m + m^4/(3m^3) = 4m/3$.

For $n \leq m$, the residual square is at most n^2/m^4. The sum-of-squares estimate

$$\sum_{n=1}^{m} n^2 = \frac{m(m+1)(2m+1)}{6} \leq m^3$$

holds for $m \geq 1$, so this part contributes at most $1/m$. For $n > m$, each residual square is at most $1/n^2$, whose tail is at most $1/m$. This proves $B_m \leq 2/m$.

For the lower bound, $n > m$ implies $m^2 + n^2 \leq 2n^2$, hence

$$\frac{n^2}{(m^2+n^2)^2} \geq \frac{1}{4n^2}.$$

The other direction of the decreasing-function integral comparison gives

$$\sum_{n>m} \frac{1}{n^2} \geq \int_{m+1}^{\infty} x^{-2}\,dx = \frac{1}{m+1}.$$

This proves $B_m \geq 1/(4(m+1))$.

Finally, for $K \geq m$, both omitted tails consist of nonnegative terms. Bound the first by $m^4 \sum_{n>K} n^{-4}$ and the second by $\sum_{n>K} n^{-2}$. The same upper integral comparisons give the two certified remainders in (30.2). No finite truncation is substituted for an equality of infinite energies. □

Thus the correction norm grows on the order of $\sqrt{m}$, the residual norm decays on the order of $m^{-1/2}$, and the penalty $\alpha A_m = A_m/m^2$ tends to zero. The smooth correction avoids the hard discontinuity at a prefix boundary, but it still approaches the forbidden constant sequence coordinate by coordinate.

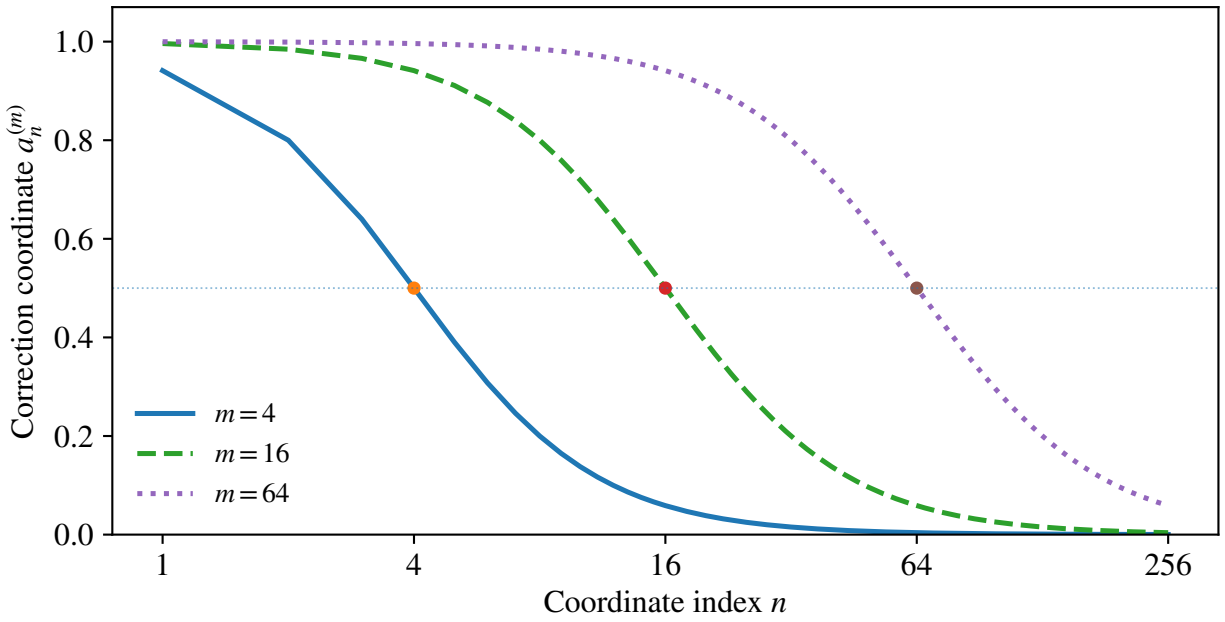


Figure 31: At each fixed coordinate, the soft correction approaches one as m tends to infinity. Its transition moves outward as m increases. The chart compares three exact coordinate formulas $m^2/(m^2+n^2)$, not finite approximations to an unspecified optimization procedure. Every displayed sequence has finite energy, but the family's energies are unbounded.

30.2 The sharp asymptotic correction cost

The preceding bounds establish divergence and decay. The leading constant can also be computed, which permits a meaningful comparison with hard prefix repair.

Theorem 30.2 (A sharp budget asymptotic for the diagonal datum)**.** *As $m \to \infty$,*

$$\frac{A_m}{m} \longrightarrow \frac{\pi}{4}, \qquad mB_m \longrightarrow \frac{\pi}{4}.$$

For the exact budget profile of this infinite problem,

$$\lim_{M\to\infty} M\beta_{T,w}(M) = \frac{\pi}{4}. \tag{30.3}$$

The hard prefix correction at budget $M = \sqrt{N}$ instead has

$$\sqrt{N}\ \|w - T\mathbf{1}_N\| \longrightarrow 1.$$

Proof. Step 1: control the infinite Riemann sums uniformly. Define

$$f(x) = \frac{1}{(1+x^2)^2}, \qquad g(x) = \frac{x^2}{(1+x^2)^2}.$$

Then $A_m/m = m^{-1}\sum_{n\ge1} f(n/m)$ and $mB_m = m^{-1}\sum_{n\ge1} g(n/m)$. Fix $L \ge 2$ and put $N = \lfloor mL \rfloor$. On $[0, L]$, both functions are uniformly continuous. If $\omega_f(h)$ is a modulus of continuity there, comparison on the N intervals of length $1/m$ gives

$$\left| \frac{1}{m}\sum_{n=1}^{N} f(n/m) - \int_0^{N/m} f(x)\,dx \right| \le L\omega_f(1/m).$$

The missing interval $[N/m, L]$ has length below $1/m$ and a bounded integrand, so it contributes a term tending to zero. The same argument applies to g. This proves convergence of each truncated sum to its integral on $[0, L]$.

The tails require a bound independent of sufficiently fine discretization. Since $N \ge mL/2$, the inequalities $f(x) \le x^{-4}$ and $g(x) \le x^{-2}$ give

$$\frac{1}{m}\sum_{n>N} f(n/m) \le m^3 \sum_{n>N} n^{-4} \le \frac{m^3}{3N^3} \le \frac{8}{3L^3},$$

$$\frac{1}{m}\sum_{n>N} g(n/m) \le m \sum_{n>N} n^{-2} \le \frac{m}{N} \le \frac{2}{L}.$$

The integral tails are at most $1/(3L^3)$ and $1/L$. Given an error tolerance, choose L first to control both discrete and integral tails, then choose m to control the compact-interval sums. This justifies passing from finite-interval Riemann convergence to the full half-line sums. No uncontrolled interchange of two infinite limits is involved.

Step 2: evaluate the integrals. The substitution $x = \tan\theta$ gives $dx = \sec^2\theta\,d\theta$ and $1 + x^2 = \sec^2\theta$, so

$$\int_0^\infty f(x)\,dx = \int_0^{\pi/2} \cos^2\theta\,d\theta = \frac{\pi}{4}, \qquad \int_0^\infty g(x)\,dx = \int_0^{\pi/2} \sin^2\theta\,d\theta = \frac{\pi}{4}.$$

Step 1 now proves the two stated asymptotics.

Step 3: move from regularization parameters to arbitrary budgets. Set $M_m = \sqrt{A_m}$. Every positive coordinate $m^2/(m^2+n^2)$ increases strictly with m, so A_m, and hence M_m, increases strictly. Since $A_m/m \to c = \pi/4 > 0$, these budgets tend to infinity and

$$\left(\frac{M_{m+1}}{M_m}\right)^2 = \frac{A_{m+1}/(m+1)}{A_m/m}\,\frac{m+1}{m} \longrightarrow 1.$$

Proposition 29.2 identifies $\beta(M_m) = \sqrt{B_m}$. Thus

$$M_m\beta(M_m) = \sqrt{(A_m/m)(mB_m)} \longrightarrow c.$$

For every sufficiently large M, choose m with $M_m \le M \le M_{m+1}$. Monotonicity of β gives

$$M_m\beta(M_{m+1}) \le M\beta(M) \le M_{m+1}\beta(M_m).$$

The lower bound is $(M_m/M_{m+1})[M_{m+1}\beta(M_{m+1})]$ and the upper is $(M_{m+1}/M_m)[M_m\beta(M_m)]$. Both tend to c, proving the full budget limit, not merely convergence on a selected subsequence of budgets.

Step 4: compare the hard-prefix method. That correction has norm $\sqrt{N}$ and residual square $\sum_{n>N} n^{-2}$. Integral comparison yields

$$\frac{N}{N+1} \le N\sum_{n>N}\frac{1}{n^2} \le 1.$$

Taking the limit and then a square root gives its leading product one. The two constants are consequently compared at large actual correction budgets, rather than at arbitrarily equated parameter labels. □

At large matched correction budgets, the optimal soft repair therefore has about $\pi/4$ times the hard-prefix residual. This improves the efficiency of approximation, not the existence result. Both residuals tend to zero only as the budget tends to infinity. The limiting constant is a calculation for this declared model, not a universal regularization efficiency factor.

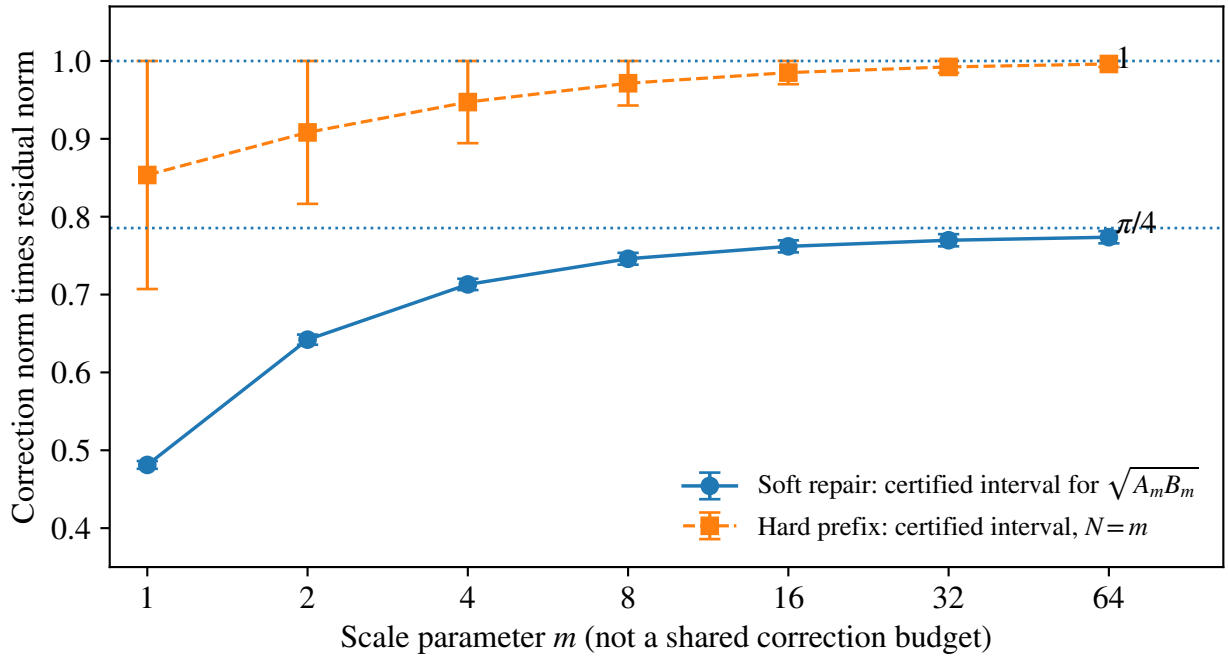


Figure 32: The plotted intervals enclose the actual infinite quantities $\sqrt{A_mB_m}$, using finite sums plus the proved tails. The corresponding hard-prefix intervals enclose $\sqrt{N\sum_{n>N} n^{-2}}$ with $N = m$. These have different natural budget sequences. Their limits give the matched-budget asymptotic constants $\pi/4$ and one.

30.3 The same operator has an attainable datum

Change only the residual to $b_n = 1/n^2$. Its minimum norm preimage is $a_n^\dagger = 1/n \in \ell^2$. At $\alpha = m^{-2}$, the regularized correction is

$$c_n^{(m)} = \frac{m^2}{n(m^2+n^2)}, \qquad a_n^\dagger - c_n^{(m)} = \frac{n}{m^2+n^2}.$$

The bias square is exactly the previously computed B_m. This links the two cases without conflating their existence hypotheses: the escaping residual of one problem is the regularization bias of an attainable problem.

Corollary 30.3 (An explicit noisy attainable repair)**.** *For the attainable datum $b = (1/n^2)$, suppose $\left\|\widetilde{b} - b\right\| \le \delta$ with $0 < \delta \le 1$. Choose*

$$m = \lceil \delta^{-2/3} \rceil, \qquad \alpha = m^{-2}.$$

Then the noisy regularized correction satisfies

$$\left\| a_\alpha(\widetilde{b}) - a^\dagger \right\| \le \sqrt{\frac{2}{m}} + \frac{\delta m}{2} \le (\sqrt{2}+1)\delta^{1/3}.$$

Proof. The exact-data bias has coordinates $n/(m^2+n^2)$, so its squared norm is precisely B_m. By Proposition 30.1,

$$\left\| a_{m^{-2}}(b) - a^\dagger \right\| \le \sqrt{2/m}.$$

Theorem 29.4 at the same parameter gives

$$\left\| a_{m^{-2}}(\widetilde{b}) - a_{m^{-2}}(b) \right\| \le \frac{\delta}{2\sqrt{m^{-2}}} = \frac{\delta m}{2}.$$

The triangle inequality proves the first bound.

Because $0 < \delta \le 1$, the number $u = \delta^{-2/3}$ is at least one. The ceiling satisfies $u \le m < u + 1 \le 2u$. Therefore

$$\sqrt{2/m} \le \sqrt{2}\,\delta^{1/3}, \qquad \frac{\delta m}{2} \le \delta u = \delta^{1/3}.$$

Adding these estimates proves the displayed rate with every ceiling factor accounted for. The positive-noise condition is needed to define this particular finite integer choice. The statement does not prescribe an infinite m at $\delta = 0$, nor does it use a probabilistic assumption on the perturbation. □

This is a deterministic guarantee for the specified datum and Hilbert noise bound. It is not a statistical confidence statement. The identical parameter rule in the escaping problem cannot produce convergence to an exact protected Hilbert solution, because no such solution exists.

30.4 What the calculation permits a machine to certify

The finite program can check a normal equation, exact preservation, and a quadratic optimality identity. It can return a dual witness ruling out a smaller residual at a given budget. In the infinite diagonal example it can return partial sums together with rational tail enclosures. Those are distinct certificates with distinct conclusions.

None of them alone establishes unrestricted semantic realizability. The bounded-budget theorem needs a genuinely uniform bound. The exactification theorem needs control of candidate movement.

The regularization limit needs actual attainment of the projected datum. The repeated appearance of a small residual does not replace any of those hypotheses.

The useful next diagnostic is therefore not just "Did the residual decrease?" It is "What happened to the correction norm, what quantity was preserved, and what existence theorem applies to this datum?"

31 Certifying an inexact protected correction

An existence theorem for a regularized optimizer does not guarantee that a numerical procedure has reached it. The procedure may stop early, use a poorly scaled basis, or report the residual of the wrong equation. These are mathematical distinctions, not implementation details that can be deferred until after a certificate has been interpreted.

The relevant question is now more specific. Given a proposed correction, what quantities can be checked without first computing the exact optimizer? The answer is a normal-equation residual and, independently, a primal and dual objective gap. Both give quantitative statements about the declared regularized problem. Neither promises that the original unregularized realization problem is solvable.

> The accompanying formalization constructs the positive-parameter optimizer and proves the energy identity, global uniqueness, and the estimates in this chapter, with a separate bridge to the actual protected kernel. Regularization limits and the combined inexact stopping theorem also have formal counterparts. The explicit infinite diagonal results also have formal counterparts. The printed weak-subsequence argument remains written background material. The exact Python certificate verifier is tested, not universally verified as a refinement of Lean.

31.1 There are three different residuals

Keep $H = \ker p$, $T = d|_H$, and a residual datum r. Fix $\alpha > 0$. Write a_α for the unique minimizer of Theorem 29.1 and let $a \in H$ be any proposed correction. Three quantities now have different meanings:

$$v = r - Ta, \qquad g = T^*(Ta - r) + \alpha a, \qquad e = a - a_\alpha.$$

The vector v is the fit residual. The vector g is the normal-equation residual, equal to one half of the derivative of the quadratic objective after identifying the dual with H. The vector e is the unknown correction error. The exact optimizer generally has $v \neq 0$ but always has $g = 0$.

The adjoint in g is the restricted adjoint on H, not the unrestricted ambient adjoint. Equivalently,

$$T^*v = P_H d^* v.$$

Thus the correction certificate must measure the part of the normal equation that belongs to the permitted correction space. A numerical program that forgets this projection can incorrectly reject an exact protected optimizer.

31.2 A posteriori correction and fit bounds

Theorem 31.1 (Normal residual as a solver certificate)**.** *For every proposed $a \in H$,*

$$\|a - a_\alpha\| \le \frac{\|g\|}{\alpha}, \tag{31.1}$$

$$\|T(a - a_\alpha)\| \le \frac{\|g\|}{2\sqrt{\alpha}}, \tag{31.2}$$

$$0 \le J_\alpha(a) - J_\alpha(a_\alpha) \le \frac{\|g\|^2}{\alpha}. \tag{31.3}$$

Proof. Let $e = a - a_\alpha$ and $B = T^*T + \alpha I$. The normal equation of the constructed optimizer gives

$$Be = (T^*T + \alpha I)a - T^*r = g.$$

Pairing with e yields the exact identity

$$\|Te\|^2 + \alpha \|e\|^2 = \langle g, e\rangle \le \|g\| \|e\|. \tag{31.4}$$

If $e = 0$, all three conclusions hold with zero error. Otherwise drop the first nonnegative term and divide by $\alpha \|e\| > 0$ to obtain $\|e\| \le \|g\| / \alpha$.

To obtain the sharper fitted-data bound, retain the quadratic dependence on $\|e\|$ instead of substituting the first estimate. Multiplying (31.4) by α and rearranging gives

$$\begin{aligned}\alpha \|Te\|^2 &\le \alpha \|g\| \|e\| - \alpha^2 \|e\|^2 \\ &= \frac{\|g\|^2}{4} - \left(\alpha \|e\| - \frac{\|g\|}{2}\right)^2 \le \frac{\|g\|^2}{4}.\end{aligned}$$

Division by α and square roots give the factor $1/(2\sqrt{\alpha})$.

By Theorem 29.1 the objective gap is exactly $\|Te\|^2 + \alpha \|e\|^2$, hence by (31.4) it is also $\langle g, e\rangle$. It is nonnegative and satisfies

$$J_\alpha(a) - J_\alpha(a_\alpha) \le \|g\| \|e\| \le \frac{\|g\|^2}{\alpha}.$$

Every certificate quantity on the right is computed from the proposal a and the declared problem. The unknown optimizer was used to prove the guarantee, but it is not required to evaluate the certificate. □

Sharp scalar examples. For sharpness of the correction-error bound, take $T = 0$ on $\mathbb{R}$ and a nonzero error. Then $g = \alpha e$. For sharpness of the fitted-data bound, take $T = \sqrt{\alpha}\, I$ on $\mathbb{R}$, so $g = 2\alpha e$ and $\|Te\| = \|g\| / (2\sqrt{\alpha})$. These examples establish worst-case sharpness, not simultaneous equality for every operator.

For a fixed raw vector f, the corresponding repaired candidates are $f - a$ and $f - a_\alpha$. Their distance is exactly $\|e\|$. Both preserve pf because both corrections belong to H. The theorem therefore controls a genuine protected candidate error, not just an intermediate coordinate vector.

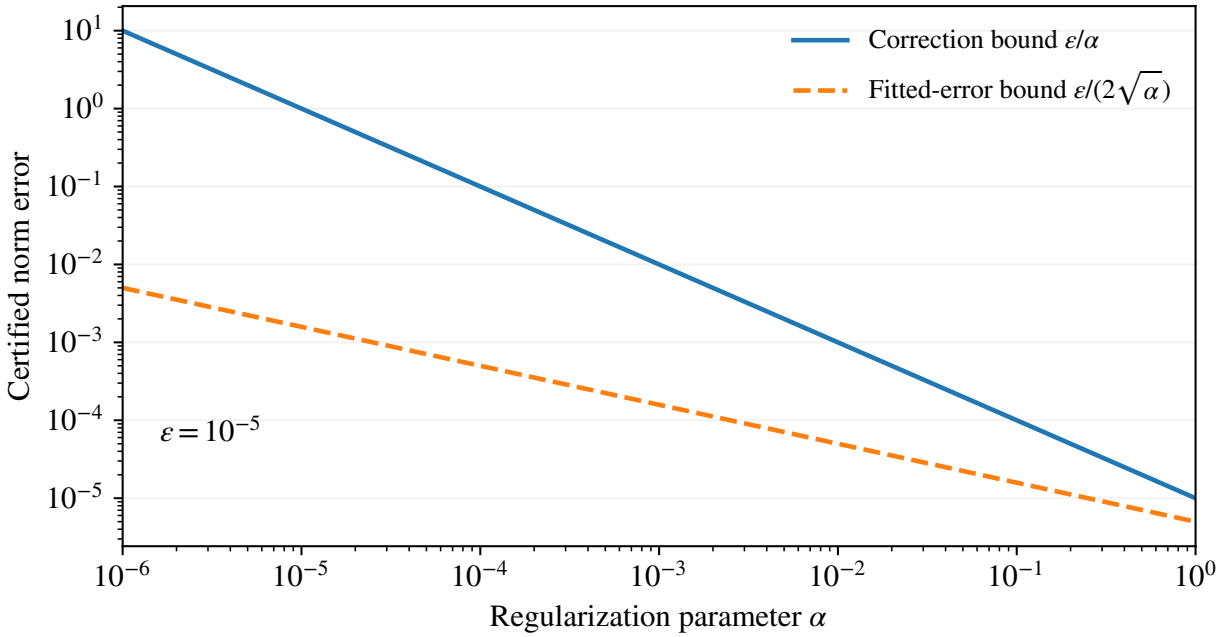


Figure 33: A fixed normal-equation tolerance has different consequences at different regularization parameters. The curves are universal certified upper bounds for a declared normal-residual norm, not observed errors of a particular experiment. The correction estimate grows as α^{-1} while the fitted-data estimate grows as $\alpha^{-1/2}$.

31.3 Why a vanishing normal residual is not enough

Let $T = 0$ on $\mathbb{R}$, let $r = 0$, and take $\alpha_n = n^{-2}$ for $n \geq 1$. The exact optimizer is zero for every such n. The fixed proposal $a_n = 1$ instead has

$$v_n = 0, \qquad g_n = \alpha_n \longrightarrow 0, \qquad J_{\alpha_n}(a_n) = \alpha_n \longrightarrow 0, \qquad |a_n - a_{\alpha_n}| = 1.$$

Thus the fit residual, the normal residual, and even the absolute objective value can all tend to zero without convergence to the selected optimizer. The certificate correctly reports $|g_n| / \alpha_n = 1$.

This is not an infinite-dimensional obstruction. It occurs in one dimension because the regularizing penalty itself disappears. The question is not whether the reported residual is small in isolation, but whether it is small relative to the parameter that converts it into an error bound.

31.4 A dual certificate does not require solving the normal equation

For an arbitrary $z \in Y$, define

$$\mathcal{D}_\alpha(z) = 2\langle r, z\rangle - \|z\|^2 - \frac{\|T^*z\|^2}{\alpha}.$$

Unlike the fixed-budget witness in Theorem 28.2, this dual witness has no unit-norm constraint. Its quadratic penalties supply the control.

Theorem 31.2 (Exact regularized primal and dual gap)**.** *For every $a \in H$ and $z \in Y$,*

$$J_\alpha(a) - \mathcal{D}_\alpha(z) = \|Ta - r + z\|^2 + \frac{\|T^*z - \alpha a\|^2}{\alpha}. \tag{31.5}$$

Consequently $\mathcal{D}_\alpha(z) \le J_\alpha(a)$ for every pair. At $a = a_\alpha$ and $z = z_\alpha = r - Ta_\alpha$, equality holds, so the dual maximum is attained and equals the primal minimum. Moreover,

$$\alpha \|a - a_\alpha\|^2 \le J_\alpha(a) - \mathcal{D}_\alpha(z).$$

For the easily computed choice $z = r - Ta$,

$$J_\alpha(a) - \mathcal{D}_\alpha(r - Ta) = \frac{\|g\|^2}{\alpha}. \tag{31.6}$$

Proof. For arbitrary a, z, expand the first square as

$$\|Ta - r + z\|^2 = \|Ta - r\|^2 + \|z\|^2 + 2\langle Ta, z\rangle - 2\langle r, z\rangle,$$

and the scaled second square as

$$\frac{\|T^*z - \alpha a\|^2}{\alpha} = \frac{\|T^*z\|^2}{\alpha} + \alpha \|a\|^2 - 2\langle T^*z, a\rangle.$$

The adjoint identity cancels $2\langle Ta, z\rangle$ against the last term. The remaining expression is

$$\|Ta - r\|^2 + \alpha \|a\|^2 - \left(2\langle r, z\rangle - \|z\|^2 - \alpha^{-1} \|T^*z\|^2\right),$$

which is the asserted gap. Its nonnegativity gives weak duality for every pair, without any constraint on the witness norm.

Set $a = a_\alpha$ and $z_\alpha = r - Ta_\alpha$. The first square is zero by this definition. The normal equation gives $T^*z_\alpha = \alpha a_\alpha$, so the second is zero as well. Therefore

$$\mathcal{D}_\alpha(z_\alpha) = J_\alpha(a_\alpha).$$

For every other z, weak duality with a_α gives $\mathcal{D}_\alpha(z) \le J_\alpha(a_\alpha)$. This proves attainment of the dual maximum and equality with the primal minimum directly.

For a general proposed pair (a, z), decompose its gap as

$$J_\alpha(a) - \mathcal{D}_\alpha(z) = [J_\alpha(a) - J_\alpha(a_\alpha)] + [J_\alpha(a_\alpha) - \mathcal{D}_\alpha(z)].$$

The second bracket is nonnegative by weak duality, while Theorem 29.1 bounds the first below by $\alpha \|a - a_\alpha\|^2$. This proves the correction-error certificate.

Finally take $z = r - Ta$. The first square in the gap is zero, while $T^*z - \alpha a = -g$. The second becomes $\|g\|^2 / \alpha$, proving the residual-witness formula. A small duality gap controls optimizer error only after division by the specified positive parameter. □

A proposed correction and a proposed dual witness can come from different procedures. The verifier only needs the mathematical problem, the two vectors, and the displayed identities. A poor witness remains valid as a lower bound, but can give a loose gap. The residual witness avoids a second optimization routine and recovers the bound in Theorem 31.1 exactly.

31.5 A protected example with every certificate calculated

Take

$$d(x,y) = x, \qquad p(x,y) = x - y, \qquad f = (1,0), \qquad r = df = 1, \qquad \alpha = 1.$$

Every allowed correction has the form (t,t). The objective is

$$J(t,t) = (t-1)^2 + 2t^2,$$

so $a_\alpha = (1/3, 1/3)$ and $J(a_\alpha) = 2/3$. Suppose a procedure returns only $a = (1/4, 1/4)$. Then

$$v = \frac{3}{4}, \qquad T^*v = (3/8, 3/8), \qquad g = (-1/8, -1/8), \qquad \|g\|^2 = \frac{1}{32}.$$

The actual correction-error square is $1/72$, while the certified upper bound is $1/32$. The actual fitted-error square is $1/144$, while the sharper bound in Theorem 31.1 is $1/128$.

The proposed objective is $11/16$. Its actual excess over the optimum is $1/48$. The residual witness gives

$$\mathcal{D}_1(3/4) = \frac{21}{32}, \qquad J(a) - \mathcal{D}_1(3/4) = \frac{1}{32}.$$

The duality gap is an upper bound for suboptimality, not necessarily the exact suboptimality. Both repaired candidates preserve $p(f) = 1$.

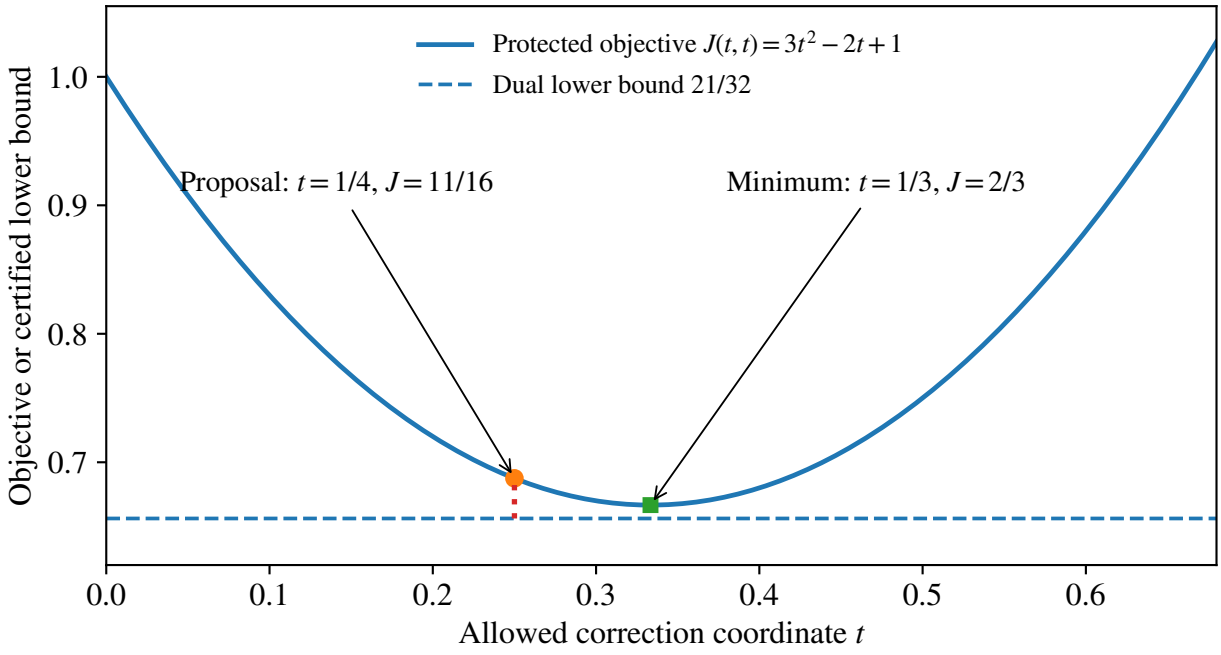


Figure 34: The objective on the actual protected line is $3t^2 - 2t + 1$. The proposed correction at $t = 1/4$, exact minimizer at $t = 1/3$, and residual-witness lower bound are distinct. The vertical gap to the dual level is the checkable certificate. The smaller gap to the minimum is the true suboptimality.

The supplied program `examples/inexact_repair_certificate.py` verifies this calculation without calling the optimizer. It also accepts a user-supplied dual witness. Every reported scalar is rational, and all norm error bounds are stored as squared quantities so that no rounding of a square root is required.

32 Coordinates, noise, and a justified stopping rule

A certificate can be exact and still be interpreted incorrectly. Two common errors are to measure a coordinate covector with the wrong norm, or to keep a solver tolerance fixed while sending the regularization parameter to zero. Both change what a small reported number means.

32.1 The residual of the coordinate equation is a covector

In finite dimensions, use orthonormal coordinates in the ambient candidate and residual spaces. Let the columns of B form an arbitrary basis of $H = \ker p$, not necessarily an orthonormal one. Write $a = Bc$, $S = dB$, and $G = B^{\mathsf{T}}B$. As before, G is positive definite. For an inexact coefficient vector c, the residual of the coordinate normal equation is

$$q = (S^{\mathsf{T}}S + \alpha G)c - S^{\mathsf{T}}r.$$

The vector q stores pairings against basis vectors. It is not the coordinate vector of the Hilbert normal residual unless the basis is orthonormal.

Proposition 32.1 (The invariant norm of a coordinate certificate)**.** *The actual normal residual is*

$$g = BG^{-1}q, \qquad \|g\|^2 = q^{\mathsf{T}}G^{-1}q. \tag{32.1}$$

For any invertible change of basis $B' = BC$, the same proposed correction has

$$c' = C^{-1}c, \quad G' = C^{\mathsf{T}}GC, \quad q' = C^{\mathsf{T}}q,$$

and

$$(q')^{\mathsf{T}}(G')^{-1}q' = q^{\mathsf{T}}G^{-1}q.$$

Proof. Because B has independent columns, for every nonzero coefficient vector c,

$$c^{\mathsf{T}}Gc = c^{\mathsf{T}}B^{\mathsf{T}}Bc = \|Bc\|^2 > 0.$$

Thus G is positive definite and invertible. The matrix $P = BG^{-1}B^{\mathsf{T}}$ is symmetric and obeys

$$P^2 = BG^{-1}(B^{\mathsf{T}}B)G^{-1}B^{\mathsf{T}} = P, \qquad PB = B.$$

Its range is precisely the column space H, so it is the orthogonal projection onto the entire permitted correction space.

The restricted Hilbert normal residual is

$$g = Pd^{\mathsf{T}}(dBc - r) + \alpha Bc.$$

Substitute $S = dB$, $G = B^{\mathsf{T}}B$ and the expression for P:

$$g = BG^{-1}(S^{\mathsf{T}}(Sc - r) + \alpha Gc) = BG^{-1}q.$$

Equivalently $q = B^{\mathsf{T}}g$ records the pairings of g with the basis columns. Its squared Hilbert norm is

$$g^{\mathsf{T}}g = q^{\mathsf{T}}G^{-1}B^{\mathsf{T}}BG^{-1}q = q^{\mathsf{T}}G^{-1}q,$$

using symmetry of G^{-1}. This proves both coordinate formulas.

For the basis change $B' = BC$, the same ambient vector is represented by $c' = C^{-1}c$. Also $S' = SC$ and $G' = C^{\mathsf{T}}GC$. Substitute these formulas into the new coordinate normal equation:

$$\begin{aligned} q' &= ((S')^{\mathsf{T}}S' + \alpha G')c' - (S')^{\mathsf{T}}r \\ &= C^{\mathsf{T}}(S^{\mathsf{T}}S + \alpha G)CC^{-1}c - C^{\mathsf{T}}S^{\mathsf{T}}r = C^{\mathsf{T}}q. \end{aligned}$$

The inverse metric is $(G')^{-1} = C^{-1}G^{-1}(C^{\mathsf{T}})^{-1}$. Hence

$$(q')^{\mathsf{T}}(G')^{-1}q' = q^{\mathsf{T}}CC^{-1}G^{-1}(C^{\mathsf{T}})^{-1}C^{\mathsf{T}}q = q^{\mathsf{T}}G^{-1}q.$$

For a zero-dimensional kernel these are empty matrices and zero vectors, and the resulting scalar norm is zero. The spanning assumption is indispensable: a basis of a proper subspace would compute a projection to a different correction problem. □

In the previous example, $B = (1,1)^{\mathsf{T}}$, $G = 2$, and the proposal $c = 1/4$ gives $q = -1/4$. Thus $q^{\mathsf{T}}G^{-1}q = 1/32$, as required. Replacing B by sB changes q to sq and G to s^2G. The Euclidean norm of the coordinate covector can be made arbitrarily small or large by changing the nonzero scalar s. The actual correction, normal residual, and certified error do not change.

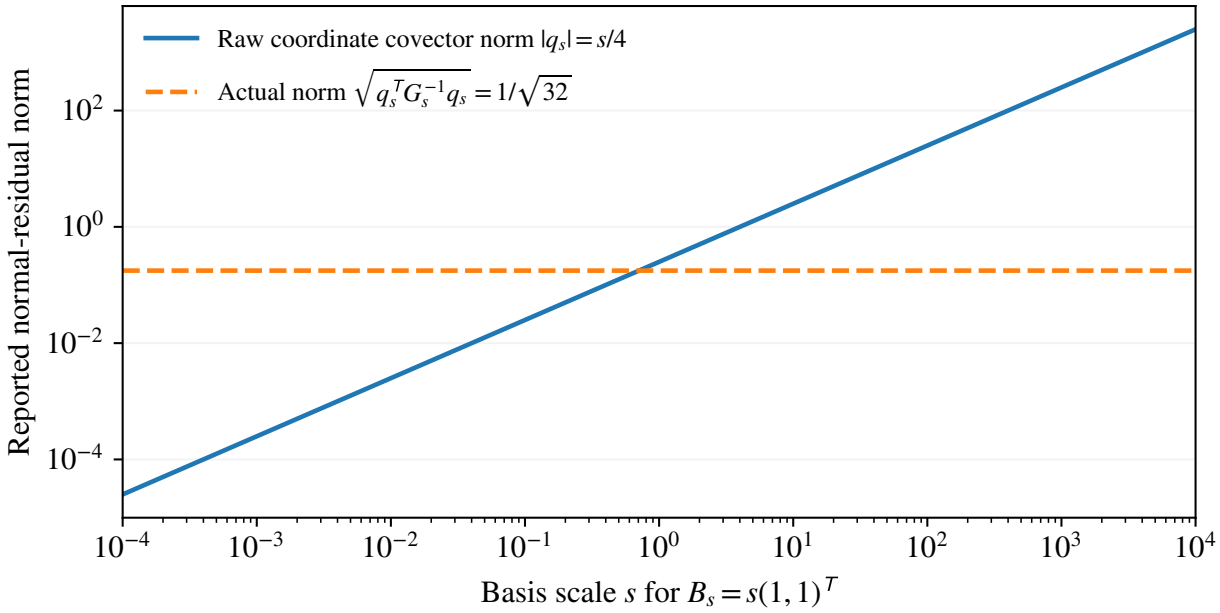


Figure 35: Rescaling the basis of the same protected line rescales the raw coordinate covector. Its Euclidean norm is therefore not a representation-invariant stopping criterion. The Gram-corrected norm stays equal to $1/\sqrt{32}$ for every displayed scale. All systems represent the same proposed correction.

The finite verifier checks that a supplied basis spans the *whole* protected kernel. A basis of a smaller subspace might give a correct certificate for that smaller optimization problem while missing an allowed direction that decreases the objective. The rank, kernel-membership, and independence checks therefore belong to the certificate contract.

32.2 Separate data error, numerical error, and regularization bias

Let r be the intended datum and $\widetilde{r}$ the measured datum. Suppose

$$\|\widetilde{r} - r\| \le \delta, \qquad \|T^*(Ta - \widetilde{r}) + \alpha a\| \le \varepsilon.$$

The maps T, p and their Hilbert metrics remain fixed. By Theorem 31.1 and the fixed-parameter part of Theorem 29.4,

$$\|a - a_\alpha(r)\| \le \frac{\varepsilon}{\alpha} + \frac{\delta}{2\sqrt{\alpha}}. \tag{32.2}$$

The first term is the solver's contribution. The second is the worst-case propagation of data error. They should not be collapsed into one unqualified residual tolerance.

Theorem 32.2 (A consistent inexact stopping rule)**.** *Assume the projected intended datum belongs to the actual range of T, with minimum-norm preimage $a^\dagger$. Along a sequence of positive parameters and measured problems, suppose*

$$\alpha_n \to 0, \qquad \frac{\delta_n}{\sqrt{\alpha_n}} \to 0, \qquad \frac{\varepsilon_n}{\alpha_n} \to 0.$$

If the proposed corrections $a_n \in H$ satisfy the displayed data and normal-residual bounds, then $a_n \to a^\dagger$ in norm. At every index,

$$\left\|a_n - a^\dagger\right\| \le \frac{\varepsilon_n}{\alpha_n} + \frac{\delta_n}{2\sqrt{\alpha_n}} + \left\|a_{\alpha_n}(r) - a^\dagger\right\|.$$

Proof. At index n, distinguish two exact minimizers: let $\widetilde{a}_n = a_{\alpha_n}(\widetilde{r}_n)$ solve the measured problem and let $a_n^0 = a_{\alpha_n}(r)$ solve the intended-data problem. All three corrections, including the proposal a_n, belong to the same space H.

The supplied normal-residual bound is for the measured datum. Therefore Theorem 31.1 applies first to that problem and gives

$$\|a_n - \widetilde{a}_n\| \le \frac{\varepsilon_n}{\alpha_n}.$$

The pointwise data estimate in Theorem 29.4, at the same parameter, then gives

$$\left\|\widetilde{a}_n - a_n^0\right\| \le \frac{\delta_n}{2\sqrt{\alpha_n}}.$$

Insert these two intermediate vectors between a_n and $a^\dagger$ in the triangle inequality to obtain the displayed three-term estimate.

The first term tends to zero by the solver-tolerance hypothesis and the second by the data/parameter hypothesis. Since the projected intended datum is attained and $\alpha_n \to 0$, the last term tends to zero by Theorem 29.3. For a direct epsilon argument, make each term smaller than one third of a prescribed tolerance beyond its own threshold and take the maximum of the three thresholds. This gives norm convergence of the proposals.

Protection is not merely asymptotic: $a_n \in \ker p$ implies $p(f - a_n) = pf$ at every index. The proof compares no changing operators or changing kernels. Both the pointwise certificate and the convergence conclusion have formal counterparts under these hypotheses. □

The accompanying Lean formalization includes the pointwise estimate (32.2), the regularization limit, and the resulting convergence theorem. Their statements retain the same fixed operator, Hilbert geometry, and attainment requirements.

32.3 Retaining the explicit diagonal rate with an inexact solver

For the attainable datum of Corollary 30.3 and $0 < \delta \le 1$, take $m = \lceil \delta^{-2/3} \rceil$ and $\alpha = m^{-2}$ as before. If an inexact solver satisfies

$$\varepsilon \le \kappa \alpha \delta^{1/3} \quad (\kappa \ge 0),$$

then the same argument gives

$$\left\| a - a^{\dagger} \right\| \le (\sqrt{2} + 1 + \kappa)\delta^{1/3}.$$

Thus the inner solve can retain the proven outer accuracy if its normal residual decreases on the scale $\alpha\delta^{1/3}$. Since α is on the scale $\delta^{4/3}$, the corresponding normal-residual scale is $\delta^{5/3}$, with the stated ceiling factors retained in the exact rule. This is a sufficient rule for this problem, not a universal optimal stopping law.

The declaration `DiagonalRegularization.diagonal_inexact_rounded_rate` formalizes this specialization for the actual infinite diagonal operator. Its hypotheses bound both the data error and the normal residual evaluated against the measured datum. It retains the literal ceiling in m and derives the exact-data bias estimate within the model.

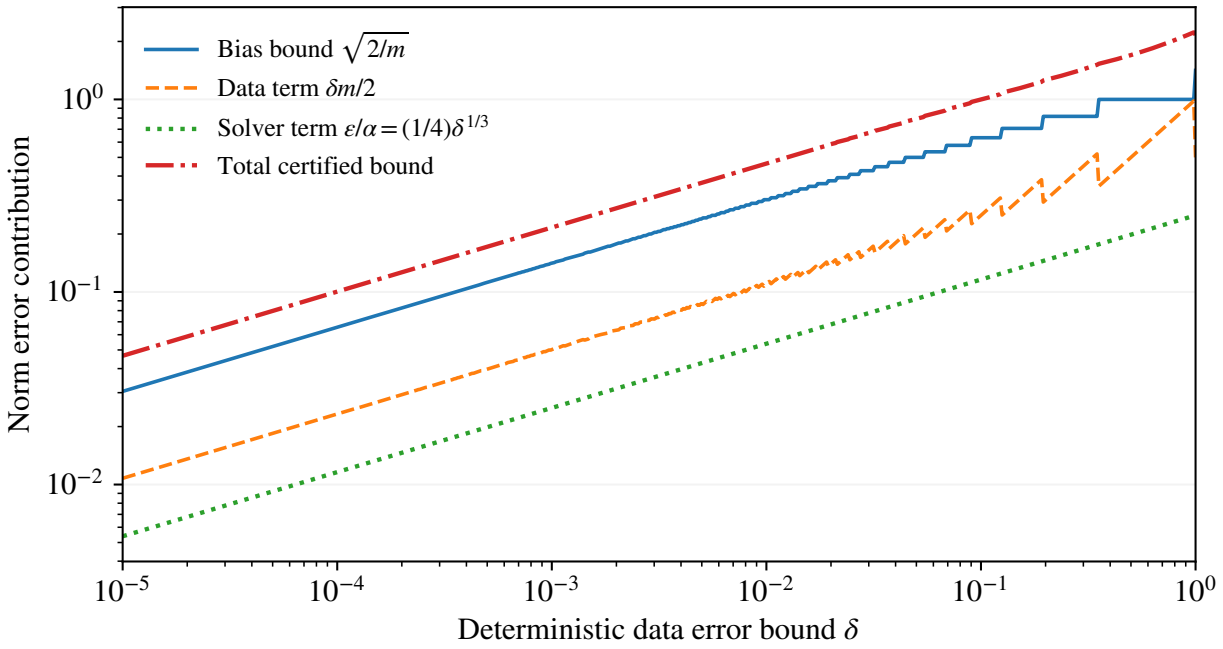


Figure 36: For the attainable diagonal example, the three displayed terms are the proved bias bound, propagated data-noise bound, and a solver-error budget chosen to preserve the same $\delta^{1/3}$ rate. The ceiling in $m = \lceil \delta^{-2/3} \rceil$ is evaluated explicitly. The plot shows bounds, not measurements from a trained system.

For the escaping datum, the same rule does not create convergence to a nonexistent protected Hilbert answer. It only makes the numerical proposals track the regularized path accurately. Tracking an escaping path and converging to a realization are different conclusions.

32.4 The executable verification contract

The program takes d, p, f, α and a proposed ambient correction, with optional measured residuals, a full kernel basis, and an independent dual witness. It rejects a correction that changes the protected data

and checks the basis's dimensions, rank, and kernel membership. It does not silently project a rejected proposal into another problem.

The output records $f - a$, protection before and after, measured fit, actual model residual, the Gram-corrected normal certificate, objective values, and squared error bounds. Verification recomputes the record without calling an optimizer. Independent determinant-based tests compare the bounds with actual errors in small ambient constrained systems. A certificate establishes consistency with its included problem, not authentication that this is the externally intended task.

A certificate needs a declared target. Fit measures agreement with data. The normal residual bounds optimizer error at a specified parameter. The duality gap bounds objective error. A realization theorem is still needed to conclude convergence to an intended exact object.

Part VI

Worked realization problems

The constructions become easier to distinguish when a complete example is carried through from its view equations to its obstruction, its ambiguity, and its numerical repair. The following calculations instantiate the preceding principles. They do not introduce a second, silently enlarged formal theory.

33 A closed signal loop

33.1 Prescribed differences as affine view transitions

Consider four scalar views arranged around a cycle. Their values are x_0, x_1, x_2, x_3, and the prescribed differences are r_0, r_1, r_2, r_3. Indices are taken modulo four. The compatibility equations are

$$x_{i+1} - x_i = r_i. \tag{33.1}$$

Work over $\mathbb{Q}$ or $\mathbb{R}$. These equations can be presented as a strict set-valued view system: an oriented edge has pullback $z \mapsto z - r_i$, and a path has the composite affine pullback. Use the category of paths, retaining the full loop as a path rather than imposing that it is an identity arrow in advance.

Every edge map is bijective. Around the cycle their composite is translation by $-\sum_i r_i$. A global realization must be fixed by that translation. Thus the example makes local invertibility and global compatibility visibly different without needing a noncommutative group.

33.2 The complete solvability calculation

Summing (33.1) gives the necessary condition

$$r_0 + r_1 + r_2 + r_3 = 0. \tag{33.2}$$

It is sufficient as well. Given an anchor a, set

$$x_0 = a, \quad x_1 = a + r_0, \quad x_2 = a + r_0 + r_1, \quad x_3 = a + r_0 + r_1 + r_2.$$

The first three equations hold by construction, and the last is exactly (33.2). Every solution has this form because the first three equations recursively determine it from x_0. Therefore a compatible residual has a one-dimensional family of primitives, with constant shifts as its ambiguity. Prescribing the anchor removes that ambiguity.

In matrix form, the difference operator is

$$D = \begin{pmatrix} -1 & 1 & 0 & 0 \\ 0 & -1 & 1 & 0 \\ 0 & 0 & -1 & 1 \\ 1 & 0 & 0 & -1 \end{pmatrix}, \qquad \ker D = \operatorname{span}\{(1,1,1,1)\}.$$

The explicit construction proves

$$\operatorname{ran} D = \ker \ell, \qquad \ell(r) = \sum_i r_i.$$

Thus ℓ is a complete feasibility test for this prescribed-difference problem. It annihilates every actual difference family, but takes a nonzero value on an incompatible one. This is the finite obstruction-certificate pattern of Proposition 4.2 and Proposition 35.3 in a fully visible model.

The same proof works on a cycle of any positive length: partial sums give a primitive, and the final wraparound equation is the zero-sum condition. No Fourier transform or numerical approximation is needed for this solvability argument.

33.3 Repairing an incompatible residual is a new problem

Take

$$r = (1, 0, 0, 0).$$

Each of the first three edge equations can be satisfied, but the full sum is one. There is no coherent closed-loop signal with exactly these four differences. Computing a long list of partial values does not remove the missing wraparound equation.

To repair the residual, now regard the vector of edge data as a candidate in its own right. Its defect is $\ell(r)$. The exact subspace is the zero-sum hyperplane. With the ordinary Euclidean norm and no protected edge coordinates, the nearest compatible residual is

$$\widehat{r} = r - \frac{\ell(r)}{4}\mathbf{1} = \left(\frac{3}{4}, -\frac{1}{4}, -\frac{1}{4}, -\frac{1}{4}\right).$$

The correction is perpendicular to the hyperplane, so the projection theorem proves it is the unique nearest vector there. Its squared size is $1/4$. This is a repair of the edge data, not a solution of the original incompatible equations with those data unchanged.

Anchor the repaired signal at zero. Its primitive is

$$\widehat{x} = \left(0, \frac{3}{4}, \frac{1}{2}, \frac{1}{4}\right). \tag{33.3}$$

The next partial sum is zero, so it closes. If instead the desired representative has mean zero, subtract the mean $3/8$:

$$\widehat{x}_{\mathrm{mean}} = \left(-\frac{3}{8}, \frac{3}{8}, \frac{1}{8}, -\frac{1}{8}\right).$$

Both primitives realize exactly $\widehat{r}$. They differ by an element of $\ker D$. The anchor condition and the mean condition are two different selections from the same affine solution fiber.

33.4 Protecting three edges forces a different repair

Suppose the first three values $(r_0, r_1, r_2) = (1, 0, 0)$ are required to remain fixed. The admissible corrections then affect only the final edge. Compatibility forces

$$\widetilde{r} = (1, 0, 0, -1), \qquad \widetilde{x} = (0, 1, 1, 1)$$

for anchor zero. This is the unique repair preserving those three edges. Its squared correction size is one rather than $1/4$. The larger error is not a failure of nearest repair: the feasible set has changed.

Here the general protected construction is completely explicit. With candidate space $E = \mathbb{R}^4$, defect $d = \ell$, and protected observable $p(r) = (r_0, r_1, r_2)$, the allowed direction is the last coordinate axis. The restricted operator $d|_{\ker p}$ has gain one. Without protection, the least gain for distance to $\ker \ell$ is $1/2$, since $\|\ell\| = 2$. The different constants measure different correction spaces.

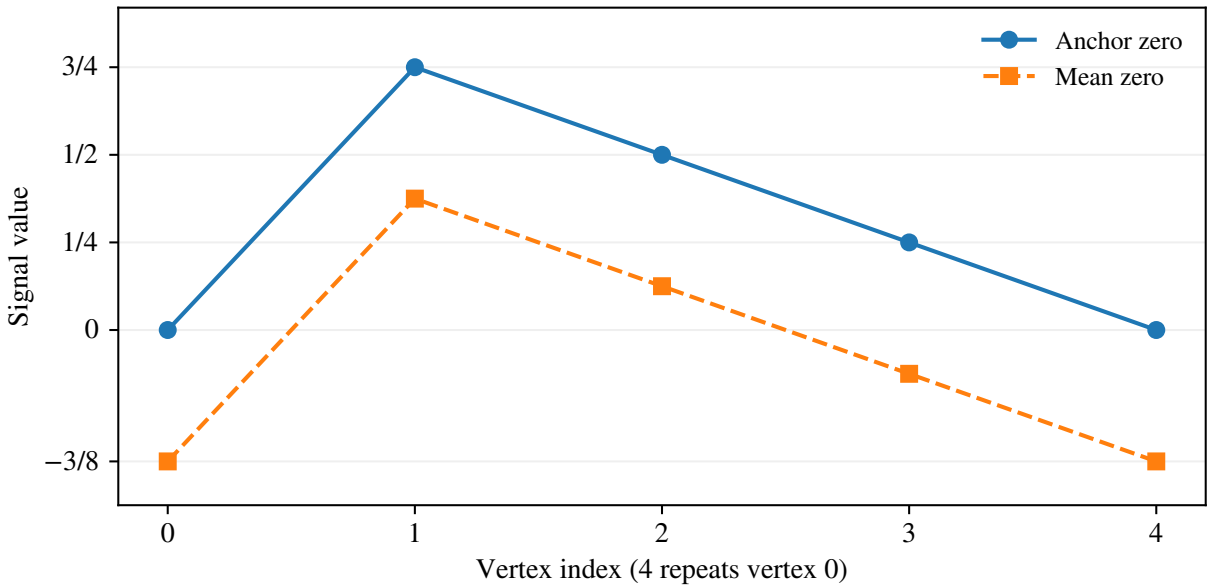


Figure 37: The repaired edge data integrate to a closed signal. The value at position four repeats the anchor at position zero. The mean-zero and anchor-zero primitives differ by a constant and therefore have the same four differences. The horizontal axis labels vertices, not successive optimization iterations.

33.5 Where the packet and transfer viewpoints enter

The raw edge family contains information about four transitions. Its zero-sum constraint links all four. Summing a block of differences transfers information to a coarser transition, while retaining the differences within the block preserves the detail that the sum omits. For example,

$$r_0 + r_1 = x_2 - x_0$$

contains the two-step transition but does not determine r_0 and r_1 separately. Replacing (r_0, r_1) by $(r_0 + t, r_1 - t)$ leaves that coarse transition unchanged.

A packet of coarse value and independent detail retains both directions. A residual-only summary can detect failure of the wraparound equation but cannot tell which protected edge may be changed. An anchor or a mean fixes the remaining signal ambiguity after feasible edge data have been chosen. These are three different roles: transition compression, feasibility repair, and representative selection.

This example is a complete finite calculation built from the strict realization, additive obstruction, and protected Hilbert principles. The companion program checks the displayed rational values, the zero-sum criterion, the primitive reconstruction, and the orthogonality equations. The calculation does not claim formal verification of a Fourier or signal-processing library.

34 A loop with a nontrivial fixed space

The previous cycle has scalar translations and a scalar total obstruction. A linear connection can behave differently: its loop may be nontrivial while retaining a whole subspace of feasible anchors. The following three-vertex model distinguishes full anchor preservation, feasible-anchor preservation, and global nearest-point repair in one calculation.

34.1 The connection and its exact families

Let the fiber at each of three vertices be $\mathbb{R}^2$. Use identity transport along the route $0 \to 1 \to 2$, and let the closing edge carry the reflection

$$J = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}.$$

For one chosen orientation of each edge, the residuals are

$$r_{01} = x_1 - x_0, \qquad r_{12} = x_2 - x_1, \qquad r_{20} = Jx_0 - x_2.$$

The reverse transports are inverses, so this defines an inverse-consistent isometric connection. Exactness first forces $x_1 = x_0$ and $x_2 = x_0$, and the remaining edge then forces $Jx_0 = x_0$. Hence every exact family is

$$((a,0),(a,0),(a,0)), \qquad a \in \mathbb{R}. \tag{34.1}$$

The loop is not the identity on the whole fiber. Its fixed space is the first coordinate axis. Every anchor on that axis extends uniquely, and every anchor with nonzero second coordinate is obstructed if its full value is to be preserved.

The three-edge residual also gives a direct telescope:

$$r_{01} + r_{12} + r_{20} = (J - I)x_0.$$

Its second component is $-2(x_0)_2$. In this chosen route normalization, the loop residual determines the transverse anchor component. It does not constrain the first component, which is the exact freedom seen in (34.1).

34.2 Three different questions for the same raw data

Take

$$x_0 = (2,1), \qquad x_1 = (4,-1), \qquad x_2 = (3,2).$$

The measured oriented residuals are

$$r_{01} = (2,-2), \qquad r_{12} = (-1,3), \qquad r_{20} = (-1,-3).$$

Preserving the entire original anchor is impossible, because $(2,1)$ is not fixed by J. The earlier obstruction criterion therefore gives a negative answer to that particular preservation problem.

Preserving only the feasible projection of the anchor gives a different answer. Orthogonally project $(2,1)$ to $(2,0)$, then reconstruct along the tree. The result is

$$P_{\text{anchor}}x = ((2,0),(2,0),(2,0)).$$

It is the unique exact family preserving the first component of the original anchor. It does not preserve its second component, and makes no claim to do so.

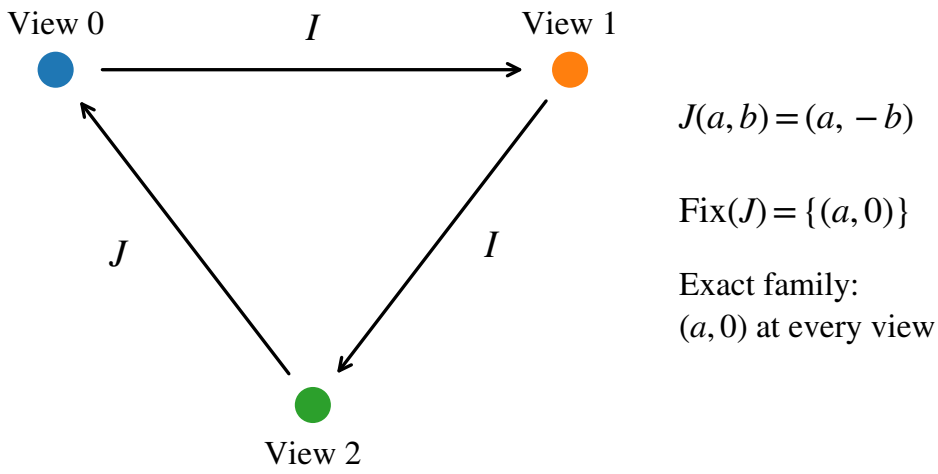


Figure 38: The tree edges identify the three fiber values. The reflection on the closing edge then removes only the transverse anchor component. The loop is nontrivial, yet all nonzero vectors on its fixed line define nonzero coherent families. The displayed connection uses a two-dimensional fiber at every vertex.

For global nearest-point repair, declare the Euclidean product norm on $(\mathbb{R}^2)^3$. The closest member of (34.1) uses the mean of the three first coordinates, namely three:

$$P_{\text{nearest}}x = ((3,0),(3,0),(3,0)).$$

To verify this without appealing only to a general theorem, write the squared error at parameter a:

$$(2-a)^2+(4-a)^2+(3-a)^2+1^2+(-1)^2+2^2 = 3(a-3)^2+8.$$

It is minimized uniquely at $a = 3$. The anchor-based answer has squared error eleven. The globally nearest answer has squared error eight. Their difference has squared norm three, giving the expected Pythagorean identity $11 = 8+3$.

No two of these conclusions conflict. Full anchor preservation is infeasible. Feasible-anchor preservation has a unique answer. Unprotected global nearest repair selects a different answer. A zero-residual output without its preservation contract would not distinguish them.

34.3 An exact conditioning calculation

For this paragraph use one orientation per edge and the Euclidean norm on the six residual coordinates. Group candidate coordinates by fiber direction. The two scalar blocks of the defect operator are

$$A = \begin{pmatrix} -1 & 1 & 0 \\ 0 & -1 & 1 \\ 1 & 0 & -1 \end{pmatrix}, \qquad B = \begin{pmatrix} -1 & 1 & 0 \\ 0 & -1 & 1 \\ -1 & 0 & -1 \end{pmatrix}.$$

The ordinary cycle block satisfies

$$A^{\mathsf{T}}A = 3I - \mathbf{1}\mathbf{1}^{\mathsf{T}}.$$

It is zero on the constant line and equals $3I$ on its orthogonal complement. The reflected block satisfies

$$B^{\mathsf{T}}B = I + vv^{\mathsf{T}}, \qquad v = (1, -1, 1)^{\mathsf{T}}.$$

Since $\|v\|^2 = 3$, this matrix has eigenvalue four in the direction of v and eigenvalue one on its orthogonal complement. In particular it is positive definite.

These two identities show that the least nonnegative gain for the full defect-to-distance estimate in this Euclidean model is one. The lower bound is attained, for example, on a candidate with zero first components and second-component vector $(1, 1, 0)$, which lies in $v^{\perp}$. The squared input distance and squared residual are both two.

For the displayed raw candidate, the squared residual is

$$(2^2 + (-2)^2) + ((-1)^2 + 3^2) + ((-1)^2 + (-3)^2) = 28.$$

Its squared nearest-point error is eight, consistent with the estimate. The certificate is universal for this fixed matrix because the Gram identities hold for all vectors, not because that particular example passed.

If both orientations are retained with a Euclidean sum-of-squares residual norm, each isometric edge contributes its squared residual norm twice. The residual norm is multiplied by $\sqrt{2}$, and the least gain becomes $1/\sqrt{2}$. With the maximum residual norm used in the general graph theorem, that duplication does not change the maximum. This is a concrete illustration of why the norm convention must travel with a certificate.

34.4 What this example connects

The strict section equations identify the exact families. The holonomy fixed space determines which anchors are feasible. The protected observable selects which exact answer is admissible. The Hilbert geometry selects a nearest representative when that is the question being asked. The Gram matrices determine the quantitative gain in one declared coordinate model.

The example does not require the loop to be trivial on the whole fiber. Nor does it require the anchor repair to be globally orthogonal. Both of those stronger conclusions would be false. The realization calculus is useful here because it keeps the correct weaker structures available rather than forcing the model into a stronger but incorrect description.

35 Mathematical meaning and the remaining boundaries

35.1 A relative theory can be complete without being absolute

The strict space $\Gamma(F)$ is canonical relative to its diagram: its universal property determines it up to the appropriate unique isomorphism. The additive obstruction is canonical relative to its admissible corrections. A marked realization problem is invariant under the specified comparison equivalences. A repair can be canonical relative to a protected observable, or orthogonal relative to a fixed inner product.

These are meaningful and different kinds of canonicity. They need not culminate in one object independent of every choice of views, coefficients, topology, semantics, and geometry. Some changes of view intentionally forget information. Others strengthen the question. Treating every such change as an equivalence would remove the very distinctions the calculus is designed to preserve.

The appropriate central statement is therefore not that every possible mathematical presentation lives in one fixed universal space. It is that a declared system of representations can be organized into coherent realizations, with explicit extension, comparison, and repair problems. A proposed bridge between two such systems has to show what structure it preserves and what information it loses.

35.2 Internal realization and intended objects

Suppose a separate category $\mathcal{S}$ contains the intended semantic objects of an application, and a realization construction gives a comparison functor

$$\mathcal{R} : \mathcal{S} \longrightarrow \mathcal{K},$$

where $\mathcal{K}$ is the relevant category or groupoid of coherent data. Defining $\mathcal{K}$ does not establish that $\mathcal{R}$ is faithful, full, or essentially surjective. These are new mathematical questions about the actual semantic comparison.

Faithfulness asks whether distinct semantic arrows with the same source and target can become indistinguishable. Fullness asks whether every arrow between the images of two semantic objects comes from an arrow between those objects. Essential surjectivity asks whether every object of $\mathcal{K}$ is isomorphic to the image of an actual semantic object. Depending on the application, one may seek only a restricted image or a classified obstruction to belonging to it.

This distinction prevents a formal globalization of encoded data from being mistaken for an unrestricted arithmetic or physical existence theorem. The abstract construction does not manufacture missing semantic objects.

35.3 Every finite test can pass while the infinite problem fails

Example 35.1 (Finite prefixes without a global point)**.** Let $X_n = \{m \in \mathbb{N} : m \geq n\}$, with bonding maps given by inclusion. Every level is nonempty, and every finite prefix has a coherent point. The inverse limit is nevertheless empty.

Proof. For a fixed terminal level N, the integer N belongs to every fiber X_n with $n \leq N$. Assign that same integer at each such level. Every restriction is an inclusion, so it preserves the assigned integer and all finite coherence equations hold. This proves finite solvability with the quantifiers in the order "for every N, there exists a prefix."

Suppose instead that $x = (x_n)$ is a global point. Equality along each inclusion says that the underlying natural numbers of x_{n+1} and x_n are equal. Induction therefore makes every underlying coordinate equal to the number m underlying x_0. Membership of x_{m+1} in X_{m+1} now gives $m \geq m + 1$, a contradiction. Thus no global point exists.

There is no conflict with a gluing theorem for prefixes. The prefix chosen above for length $N + 1$ restricts to the value $N + 1$, not to our chosen value N on the shorter prefix. Nor could another choice of one solution for every length give a matching family: its value at level zero would force the same forbidden global integer. Separate existence at every length is strictly weaker than existence of a compatible family of prefix solutions. □

This counterexample is formalized in `FinitePrefixBoundary.lean` using the original resolution and partial-section types. It explains why inverse-limit existence needs more than separate finite solvability. Finiteness, compactness, surjectivity, or another coherence principle must be proved in the relevant application. The example does not contradict the strict gluing theorem: a compatible matching family of all prefixes was never supplied.

35.4 A positive infinite-existence criterion

The failed example isolates a missing compactness or lifting hypothesis. Finite fibers supply a concrete positive criterion without requiring every bonding map to be surjective.

Proposition 35.2 (Finite fibers and persistent images)**.** *Let (X_n, ρ_n) be a sequential inverse system of nonempty finite sets. Its inverse limit is nonempty. More precisely, if $\rho_{n,m} : X_m \to X_n$ is composite restriction, then*

$$\operatorname{im}(\mathrm{ev}_n) = Y_n, \qquad Y_n = \bigcap_{m \geq n} \operatorname{im} \rho_{n,m}.$$

Thus a prescribed value at level n extends precisely when it survives every later image.

Proof. Write $I_{n,m} = \rho_{n,m}(X_m)$ for $m \geq n$. Each $I_{n,m}$ is nonempty. The composition law gives

$$I_{n,m+1} = \rho_{n,m}\big(\rho_m(X_{m+1})\big) \subseteq I_{n,m}.$$

These are subsets of the finite set X_n. Every strict decrease removes an element permanently, so only finitely many strict decreases are possible. There is an index $s_n \geq n$ such that $I_{n,m} = I_{n,s_n}$ for all $m \geq s_n$. Consequently $Y_n = I_{n,s_n}$ is nonempty.

We next prove that restriction induces a *surjection* $Y_{n+1} \to Y_n$, even though $X_{n+1} \to X_n$ need not be surjective. For membership, let $z \in Y_{n+1}$. Given $m \geq n+1$, write $z = \rho_{n+1,m}(w)$ for some $w \in X_m$. Then $\rho_n z = \rho_{n,m}(w)$. Together with the automatic condition at $m = n$, this shows $\rho_n z \in Y_n$.

For surjectivity, fix $y \in Y_n$ and choose

$$m \geq \max\{s_n, s_{n+1}, n+1\}.$$

Since $y \in I_{n,m}$, there is $w \in X_m$ with $\rho_{n,m}(w) = y$. Put $z = \rho_{n+1,m}(w)$. Stabilization at level $n+1$ gives $z \in Y_{n+1}$, and composition gives $\rho_n z = y$.

Choose $x_0 \in Y_0$ and recursively choose $x_{n+1} \in Y_{n+1}$ above x_n. The preceding surjectivity justifies each step. Adjacent coherence implies every composite coherence equation, so this gives a point. This is an existence construction using countably many choices, not an asserted effective procedure for locating the stabilization indices.

Finally, fix a level n. Every global point satisfies $x_n = \rho_{n,m}(x_m)$ for each $m \geq n$, hence $x_n \in Y_n$. Conversely, start with any $y \in Y_n$, lift recursively above level n through the surjections just proved, and define the lower coordinates by $x_i = \rho_{i,n}(y)$. The composite restriction laws verify coherence across and below level n. The resulting point evaluates to y. These two inclusions prove the asserted equality of evaluation images, not merely nonemptiness of the inverse limit. □

This is the standard finite-set inverse-limit principle, here with its sequential persistent-image proof [1]. Its formalization in `FiniteRealizationExistence.lean` constructs an explicit equivalence with the pinned library's section type, then applies the finite cofiltered-system and eventual-range theorems. The printed sequential proof and this formal route establish the same stated criterion. Both use ordinary classical choice. Surjective bonding maps on arbitrary nonempty sets give a second sufficient condition by direct recursive lifting. Neither criterion says that arbitrary data in a non-surjective finite tower extend.

35.5 A finite computation that exposes the preservation contract

Consider a raw vector $x = (2, 0, 5) \in \mathbb{R}^3$ with defect $d(x) = x_1 - x_2$. Unconstrained orthogonal repair gives $(1, 1, 5)$, with squared error two. If the first and third coordinates must be preserved, the unique exact repair is $(2, 2, 5)$, with squared error four.

Both outputs are exact. Only one satisfies the second preservation problem. Neither theorem contradicts the other because the admissible changes differ. An implementation that reports only zero residual would not distinguish the two tasks.

For a finite linear system over a field, the same question admits independently checkable positive and negative certificates. Let $A \in \mathbb{K}^{m\times n}$ encode compatibility, $P \in \mathbb{K}^{p\times n}$ the protected values, and $x \in \mathbb{K}^n$ the raw candidate. Let the columns of N be a basis of $\ker P$. The correction equation is

$$ANz = Ax, \qquad y = x - Nz.$$

It is equivalent to the stacked system $Ay = 0$, $Py = Px$.

Proposition 35.3 (A finite certificate alternative)**.** *Exactly one of the following occurs:*

1. *There is y with $Ay = 0$ and $Py = Px$. The full answer set is $y + H$, where $H = \ker A \cap \ker P$.*
2. *There are row vectors $\lambda \in \mathbb{K}^{1\times m}$ and $\mu \in \mathbb{K}^{1\times p}$ with*

$$\lambda A = \mu P, \qquad \lambda Ax = 1.$$

The negative witness annihilates every admissible correction defect, since $\lambda AN = \mu PN = 0$. The ambiguity dimension in the positive case is

$$\dim H = n - \operatorname{rank}\begin{pmatrix} A \\ P \end{pmatrix}.$$

Proof. Set $T = \begin{pmatrix} A \\ P \end{pmatrix}$ and $b = \begin{pmatrix} 0 \\ Px \end{pmatrix}$. A vector y is an exact protected answer if and only if $Ty = b$. Perform Gaussian row operations on the augmented matrix, keeping their product U. Each elementary row operation is invertible over the given field. Thus $Ty = b$ and $UTy = Ub$ have exactly the same solutions.

If no row of the reduced augmented system has zero left side and nonzero right side, assign arbitrary values to its free variables and solve for its pivot variables. This supplies a solution y_0. For any vector y,

$$Ty = b \quad \Longleftrightarrow \quad T(y - y_0) = 0.$$

Translation by y_0 is therefore a bijection from $\ker T$ to the answer set, with inverse $y \mapsto y - y_0$. Rank–nullity applied to the actual linear map $T : \mathbb{K}^n \to \mathbb{K}^{m+p}$ gives the ambiguity dimension. Dependent prescription rows do not alter this argument.

If the reduced system is inconsistent, choose a row with zero left side and right side $\alpha \neq 0$. Let w be the corresponding row of U, and split it according to the two blocks as $w = (u, v)$. The retained row operations give the literal identities

$$uA + vP = 0, \qquad vPx = \alpha.$$

Division is permitted because α is a nonzero field element. Define $\lambda = -\alpha^{-1}u$ and $\mu = \alpha^{-1}v$. Multiplying the displayed equations by α^{-1} gives

$$\lambda A = \mu P, \qquad \lambda Ax = \mu Px = 1.$$

For every admissible correction $a \in \ker P$, these same row equations imply $\lambda Aa = \mu Pa = 0$. The witness therefore annihilates the entire correction image while taking value one on the required defect.

The alternatives cannot occur together. If y were an exact protected answer in the presence of such a witness, then

$$1 = \lambda Ax = \mu Px = \mu Py = \lambda Ay = 0.$$

A field has $1 \neq 0$, so this is impossible. Consistency or inconsistency of the reduced system proves exhaustiveness, and the contradiction proves exclusivity. Empty row or column blocks cause no extra division: normalization is used only in the inconsistent case, where the nonzero scalar has already been obtained. The proof does not transfer to integral coefficients by silently allowing the same division. □

The formal proof gives these literal row-vector equations over an arbitrary field. It uses algebraic dual separation of the stacked image rather than verifying the row-reduction program. The affine answer equivalence and rank-nullity formula are proved separately. Thus a normalized obstruction is an actual checkable witness, while universal correctness of the supplied Python elimination implementation remains a different question.

35.6 One equation, three preservation problems

For the following three cases work over $\mathbb{Q}$, with $A = (1, -1, 0)$ and $x = (2, 0, 5)$. With $P(x) = (x_1, x_3)$, the unique answer is $(2, 2, 5)$. With $P(x) = x_3$, the entire answer set is $\{(t, t, 5) : t \in \mathbb{Q}\}$. With $P(x) = (x_1, x_2)$, no answer exists. In this last case

$$\lambda = (1/2), \qquad \mu = (1/2, -1/2)$$

is a normalized negative certificate over $\mathbb{Q}$ or $\mathbb{R}$: $\lambda A = \mu P$ and $\lambda A x = 1$.

The package implements this finite interface over $\mathbb{Q}$ in `examples/protected_repair.py`. It returns a particular exact answer and an ambiguity basis, or a normalized dual witness. Its verifier checks the certificate equations directly. A Gaussian particular answer is *not* advertised as the nearest answer of Theorem 18.3. That is a separate geometric selection.

Exact field elimination does not solve the integer parity problem by changing its coefficient domain. Over $\mathbb{Q}$, $2m = 1$ has the solution $1/2$. Over $\mathbb{Z}$, it has none. The rational interface rejects floating-point inputs rather than silently treating an approximate decimal as exact arithmetic. Its finite tests establish the tested implementation behavior, not an arbitrary-input Lean refinement theorem.

35.7 Why this is relevant to machine reasoning

The framework gives a machine explicit mathematical objects to manipulate: views, restrictions, partial realizations, typed defects, admissible corrections, and comparison witnesses. A successful step can carry both an output and a certificate of what remained fixed. A failed step can return an obstruction in a specified quotient rather than an unexplained failure flag.

The split coordinate theorem is especially direct. In its stated class, the same candidate is represented by protected information and residual data, and repair zeros only the residual coordinate. The descent theorem explains which local certificates can be assembled. The analytic results explain when a small measured defect has a quantitatively reliable meaning.

These observations motivate a machine oriented mathematical language. They do not establish an autonomous reasoning system, an empirical intelligence gain, or a general efficient algorithm. Classical choice, infinite products, quotient membership, and bounded inverse constructions can be mathematically valid without yielding executable finite procedures. Implementation claims need their own specifications and tests.

35.8 What this monograph closes, and what it leaves open

The book supplies a coherent account of strict realization spaces, additive and nonabelian extension problems, split tower repair, retained multiscale detail, a genuine composition-site stack, presentation

transport, graph holonomy, closed-range stability, and moving exact spaces. The convergence and obstruction theory separates small measurements from small movements, coherent coordinate families from their semantic image, and zero residual from zero transformation cost. The bounded-budget theory supplies a positive existence criterion and keeps range membership separate from regularized solvability. Its fixed-budget dual witness can certify a limit on every admissible correction, while its diagonal example computes the actual leading cost of approximation. Each layer is attached to the data that make it meaningful.

The following broader targets are not conclusions of this edition: arbitrary higher categorical coefficients and infinity-stack descent, unrestricted changes of view doctrine, norm reconstruction for every bounded infinite residual hierarchy, nonlinear or unknown-connection repair, perturbation across changing Hilbert geometries without new comparison data, and semantic essential surjectivity for unrestricted arithmetic or physical applications.

The accompanying Lean project contains 162 production modules and 4 entry points, with compiler records, declaration axiom reports, and rejection controls. The bounded-budget, regularization-limit, noise, inexact-certificate, and coordinate-invariance results have formal counterparts under their stated hypotheses. The explicit infinite diagonal estimates, leading budget constant over all real budgets, and rounded noisy rate also have formal counterparts. The budget results do not require finite-dimensionality, separability, injectivity, or closed range. Their robustness comparison uses one fixed correction space. Further extensions require their own hypotheses and evidence.

The organizing principle. Coherent realization, obstruction to extension, freedom of completion, and stability of repair belong to one mathematical conversation. Their unification is useful precisely when it does not identify them with one another.

A A reader's workshop

These short problems isolate choices that recur in the main development. Each has a complete solution. They are reading exercises and direct applications of the cited theory, not additional entries in the formal theorem inventory.

1. A stable limit does not make every seed persistent

At every level take $X_n = \{0, 1\}$, with each bonding map sending both elements to zero. Find the inverse limit and the image of evaluation at every level. Does nonemptiness of every finite fiber make the seed one extendable?

Solution. For a coherent family, $x_n = \rho_n(x_{n+1}) = 0$ for every n. The all-zero family is coherent and is therefore the only point. Every evaluation image is $\{0\}$. Equivalently, the image of restriction from any strictly later level is already $\{0\}$, so the persistent image is $\{0\}$. The seed one never extends. This is compatible with Proposition 35.2, which proves nonemptiness and characterizes the persistent image, not surjectivity of every evaluation.

2. A residual chart can change without changing the defect

Let $d : \mathbb{R}^2 \to \mathbb{R}$ be $d(u, r) = r$. For $a \in \mathbb{R}$, let $A_a(r) = (ar, r)$. Compute the exact repair P_a, compare it with P_b, and decide when they commute.

Solution. Since $dA_a = I$,

$$P_a(u, r) = (u - ar, 0), \qquad P_b(u, r) - P_a(u, r) = ((a - b)r, 0).$$

Both fix the horizontal exact axis. Thus $P_a P_b = P_b$ and $P_b P_a = P_a$. They commute exactly when they are equal, which occurs exactly when $a = b$, as can be checked at input $(0, 1)$. Every repair has zero residual. The different answers reflect the different right inverses, as in Theorem 9.1.

3. Pairwise invertibility does not supply a triple law

Let t be a nonidentity transposition in S_3. Suppose three proposed overlap comparisons require frames b_0, b_1, b_2 satisfying

$$b_1 = tb_0, \qquad b_2 = tb_1, \qquad b_2 = tb_0.$$

Every comparison is individually invertible. Can the three frames exist?

Solution. The first two equations give $b_2 = t^2 b_0 = b_0$. The third then gives $b_0 = tb_0$, and right cancellation gives $t = 1$, contrary to the assumption. Invertibility of every proposed overlap map does not prove the composition law needed for descent. Theorem 12.1 applies to data satisfying that law, not to arbitrary lists of pairwise isomorphisms.

4. Zero noisy output is not a certificate of reference exactness

Take $D(a, b) = a$, $\widetilde{D}(a, b) = a + \frac{1}{10}b$, and input $x = (-10, 100)$. The measured operator returns zero exactly. Find the distance to the reference kernel and evaluate the reference uncertainty bound.

Solution. The reference kernel is the vertical axis, so the distance is ten. Its residual gain is one, and the operator error norm is $1/10$. With $y = 0$ and $\eta = 0$, Proposition 22.1 gives

$$\operatorname{dist}(x, \ker D) \le \frac{1}{10}\sqrt{10^2 + 100^2} = \sqrt{101}.$$

The true distance ten is below this bound. A formula omitting the operator-error term would instead claim a zero distance, which is false. The input is exact for the measured model and not for the reference one.

5. A prefix computation can hide a wraparound defect

On a four-cycle, take proposed differences $(2, -1, 3, -3)$ and anchor zero. Integrate the first three edges. Which equation fails, and what changes when only the final edge is allowed to move?

Solution. The first three edges give values $(0, 2, 1, 4)$. Their required last difference is $x_0 - x_3 = -4$, not -3. The total prescribed sum is one, so the original loop has no realization. Keeping the first three edges forces the unique correction $(2, -1, 3, -4)$, whose sum is zero. This is a protected repair of the edge data. It is not a realization of the original data without modification.

6. Residual decay can hide unprotected drift

In maximum-norm $\mathbb{R}^3$, use $p(s, u, r) = s$ and $d(s, u, r) = r$. Compare the paths $(1, 0, (n+1)^{-1})$ and $(1, H_n, (n+1)^{-1})$, where $H_n = \sum_{j=1}^{n} 1/j$. Which hypotheses distinguish them?

Solution. Both preserve $p = 1$ and have the same residual ratios $(n+1)/(n+2)$. The first moves by $1/((n+1)(n+2))$ and converges to $(1, 0, 0)$. The second moves by $1/(n+1)$ and has an unbounded second coordinate. In the variable-rate theorem, the first has $C_n = 1/(n+2)$ and a summable weighted budget. The second has $C_n = 1$, making that budget harmonic and divergent. The difference is movement, not residual decay or semantic preservation.

7. How a finite observation changes the apparent rate

Let $\eta_n = 1/(n+1)$. Find its profile after sampling at $\phi(k) = 2^k - 1$. Does the new profile prove exponential convergence in the original index?

Solution. The sampled value is $\eta_{\phi(k)} = 2^{-k}$. Recovering the original index gives $k = \log_2(n+1)$ on the sampled stages, so the value is still $1/(n+1)$. The schedule preserves vanishing but changes the meaning of the horizontal axis. No faster original process has been proved. Proposition 24.5 is a convergence equivalence under a monotone cofinal schedule, not a rate-invariance theorem.

8. What does a zero floor fail to certify?

For the diagonal model of Section 26, can an exact protected answer exist in $\ell^2 \oplus \mathbb{R}$? Does a coherent family of exact finite prefixes exist? Explain why the answers differ without contradicting finite-fiber existence.

Solution. Exactness with protected value one forces every sequence coordinate to equal minus one. This is not square summable, so there is no protected Hilbert answer. Each finite prefix nevertheless has the unique solution $(-1, \ldots, -1; 1)$, and these solutions are coherent under truncation. Their inverse limit contains the all-minus-one sequence in the full product. The finite-fiber theorem supplies that coherent point. It does not prove that the point is in the image of the square-summable semantic space. The correction image is dense but does not contain the raw residual, so its floor is zero and its algebraic class is nonzero.

9. A phase witness without enumerating every gauge

For $c(x,y) = x_2y_1$ on $\mathbf{F}_2^2$, show that no scalar coboundary equals c. Then compute the square of the lifted label $(0, e_1 + e_2)$ in the eight-element extension.

Solution. Scalar coboundaries $a(x) + a(y) - a(x+y)$ are symmetric, but $c(e_1, e_2) = 0$ and $c(e_2, e_1) = 1$. Thus no scalar coboundary equals c. For $x = e_1 + e_2$, $c(x,x) = 1$, so $(0,x)^2 = (1,0)$ and its fourth power is the identity. The corresponding matrix is XZ, whose square is $-I$. The additional phase bit records precisely the sign lost by projecting back to the four-label group.

10. The finite energy identity

In $\mathbb{R}^2$, first prescribe $a_1 + a_2 = 1$, then also prescribe $a_1 = 3/4$. Find the two minimum norm corrections and check the exact nested-affine identity.

Solution. The first correction is $(1/2, 1/2)$, with energy $1/2$. The second is $(3/4, 1/4)$, with energy $5/8$. Their difference has squared norm $1/8 = 5/8 - 1/2$. The identity uses the affine nature of both constraint sets. For the nested closed convex sets $[1,\infty)$ and $[2,\infty)$ in $\mathbb{R}$, the minimum points instead have squared difference one and energy difference three.

11. A basis is not automatically an isometry

In Euclidean $\mathbb{R}^2$, let $p(a) = a_1 - 2a_2$, $d(a) = a_1$, $r = 1$, and $\alpha = 1$. Compute the protected regularized correction using the basis $(2,1)$ of the kernel. What changes if the Gram matrix is omitted?

Solution. The coordinate c gives $a = (2c, c)$, with $S = 2$ and $G = 5$. The normal equation is $9c = 2$, so $a = (4/9, 2/9)$. Its objective is $5/9$. Replacing G by one gives $c = 2/5$ and $a = (4/5, 2/5)$, whose actual ambient objective is $21/25 > 5/9$. That replacement solves a different metric problem.

12. A robust negative certificate

A nominal scalar system has $T = 1$, $r = 1$. The operator error is at most $1/10$ and the data error at most $1/20$. Can a correction of norm at most $1/2$ achieve true residual at most $3/10$? What does the same test say at budget two?

Solution. At budget $1/2$, the nominal optimum is $1/2$. Subtracting $1/20 + (1/10)(1/2)$ leaves $2/5 > 3/10$, so the proposed accuracy is excluded for every admissible true model. At budget two, the lower estimate is negative and clips to zero. It excludes nothing. The disappearance of this certificate is not proof of exact feasibility.

13. A smaller coordinate residual without a better correction

For the protected example in Section 31, keep the same correction $(1/4, 1/4)$ but replace the kernel basis $B = (1,1)^{\mathsf{T}}$ by $B' = B/100$. Compute the coordinate normal residual and its actual Hilbert norm.

Solution. In the original basis, $c = 1/4$, $G = 2$ and $q = 3c - 1 = -1/4$. In the new basis, $c' = 25$, $G' = 1/5000$ and $q' = -1/400$. The coordinate residual is one hundred times smaller, but

$$(q')^{\mathsf{T}}(G')^{-1}q' = \frac{5000}{160000} = \frac{1}{32} = q^{\mathsf{T}}G^{-1}q.$$

The physical correction and the certificate have not improved. The changed number records a changed basis.

14. How small must the solver residual be?

For $n \geq 1$, suppose $\alpha_n = n^{-2}$, the measured-data error is bounded by n^{-2}, and the exact-data regularized corrections converge to the desired minimum-norm answer. Does a solver normal-residual tolerance n^{-2} guarantee convergence of the proposals? Give a sufficient replacement.

Solution. The noise term is at most $1/(2n)$, but the solver term is at most one. Thus the displayed tolerance does not provide the required guarantee. The scalar counterexample $T = 0$, $r = 0$, $a_n = 1$ meets it and does not approach the exact optimizer. Replacing it by n^{-3} makes the solver term at most $1/n$, so the triangle inequality and the assumed vanishing bias prove convergence. This is a sufficient worst-case rule, not a necessary condition for every particular operator.

15. Exact optimization of measured data is not exact realization

Let d be the identity on $\mathbb{R}$, protect no nonzero observable, take raw input $f = 2$, measured residual $r = 1$, and $\alpha = 1$. Find the exact measured-data optimizer and compare the fit residual with the actual repaired residual.

Solution. The objective $(a-1)^2 + a^2$ is minimized at $a = 1/2$. Its normal residual is zero, so the solver certificate is exact. The measured fit residual is $r - da = 1/2$. The actual repaired candidate is $f - a = 3/2$, whose model residual is $3/2$. The difference is $df - r = 1$. A certificate about solving the measured optimization problem does not remove measurement error or the regularization bias.

A checklist for applying the theory

Before selecting a theorem, state the raw and exact objects, the coefficient structure, the maps that relate views, and the data that must remain fixed. Then ask which conclusion is needed: existence, a description of all answers, an identification-preserving comparison, a stable distance estimate, or an executable finite certificate. Finally record the hypotheses that make that conclusion available.

The examples above fail only when one of those distinctions is suppressed. The missing information is not supplied by giving the same construction a more general name.

B A hypothesis atlas

A longer monograph is useful only if the reader can locate the theorem that answers the actual question. This atlas groups the main obligations by the desired conclusion. The references lead to the statements and proofs, not just to examples that resemble the intended application.

B.1 Existence and identification

Desired conclusion	Data and hypotheses to establish	Main entry
A coherent family	A view system with its actual restrictions and all compatibility equations. Nonemptiness is separate.	1.3–1.5
A prescribed extension	The declared partial section and an actual vanishing relative obstruction, or matching data on a transition cover.	2.1, 3.2, 4.1
A unique extension	Existence and vanishing of the invisible coherent ambiguity, not existence alone.	3.3, 11.4
An infinite point from finite views	A sequential inverse system with nonempty finite fibers, or another proved lifting/compactness hypothesis. A prescribed value must lie in the persistent image.	35.1, 35.2
Descent of objects and witnesses	The actual groupoid data, coherent overlap isomorphisms, and a composition-covering family on the specified site.	12.1–13.2
Invariance of a boundary problem	An identified boundary and the specified comparison equivalences. Literal freezing of every representative is a different construction.	14.3–14.4, 15.1
The same limit from two paths	Convergence of both paths in a metric space and a pairwise distance budget tending to zero.	24.4

B.2 Approximation and numerical meaning

Desired conclusion	Data and hypotheses to establish	Main entry
An exact limit of approximations	Completeness, a nonnegative lower semicontinuous residual, vanishing residuals, and summable candidate movement.	24.1–24.2
A protected exact limit	In addition, continuity of the protected map and preservation along the candidate path.	24.3
Exact cancellation from zero floor	Closedness of the actual admissible residual image. Without it only closure membership follows.	25.1–25.2
A stable residual-to-distance bound	The required closed range or bounded inverse in the declared normed model. In a protected problem, use the restriction to $\ker p$.	18.2, 18.3

Desired conclusion	Data and hypotheses to establish	Main entry
A controlled moving exact space	Two residual certificates, or a proved split-surjective regime with the strict smallness condition. Norm closeness alone does not preserve rank.	Section 21
An independent object distance	An arrow category and an extended nonnegative cost with zero identity cost and composition subadditivity. Separation is another question.	23.1–23.2

B.3 Bounded budgets, regularization, and solver certificates

Every entry has a printed proof. Formal coverage includes the bounded-budget, regularization-limit, noise, solver-certificate, coordinate-invariance, and stopping results under their stated hypotheses. The explicit infinite diagonal estimates, asymptotic budget constant over all real budgets, and rounded noisy rate also have formal counterparts. The formal-source appendix gives the precise correspondence.

Desired conclusion	Data and hypotheses to establish	Main entry
An exact answer from bounded approximations	A fixed bounded linear equation between real Hilbert spaces, residual convergence, and a uniform correction bound.	27.2
Strong convergence of finite repairs	Nested nonempty closed affine constraints, minimum norm answers, and a uniform energy budget.	27.3
A sharp finite-budget exclusion	The correction ball, the restricted adjoint, and a dual witness in the unit ball.	28.1–28.2
Robustness to model uncertainty	A fixed correction space and certified operator and data error bounds.	28.3
A unique protected regularized answer	Positive regularization parameter and the actual Hilbert metric on the protected kernel.	29.1–29.2
A convergent regularization limit	Actual attainment of the projected datum, not merely its membership in the closure.	29.3
A noisy convergent answer	The same attained datum, vanishing parameter, and noise divided by the square root of the parameter tending to zero.	29.4, 30.3
A leading approximation-cost constant	The explicit diagonal model and a proof controlling its infinite Riemann-sum tails.	30.1–30.2
A certified inexact optimizer	Positive parameter, the actual normal residual in the protected Hilbert norm, or an exact primal and dual gap.	31.1–31.2
A basis-independent solver residual	A full independent kernel basis and its Gram inverse. The normal-equation coordinate residual is a covector.	32.1

Desired conclusion	Data and hypotheses to establish	Main entry
A convergent stopping schedule	Attained projected datum and separate vanishing bias, rescaled data noise, and rescaled normal residual.	32.2

B.4 Three questions before using a finite check

First determine whether the tested object is the actual object in the theorem. A scalar simulation can illustrate an operator-valued statement without checking every dependent fiber. Next determine whether the test checks a universal finite domain or a selected finite range. Finally identify any limit that remains outside the computation.

The phase extension is exhaustively checked on its eight elements. The diagonal example is different: rational finite sums verify its displayed enclosures, while the proof of nonexistence uses the infinite square-summability condition. Calling both results simply “tested” would conceal the most important distinction.

The companion Lean guide records these scopes and identifies the checked examples. General source declarations are checked in the accompanying formal project. Printed examples and conceptual connections have only the formal coverage explicitly documented for them.

C Notation and formal correspondence

C.1 Notation

Symbol	Meaning
$\mathcal{V}, F_v, f^*$	A declared view category, its fibers, and contravariant restrictions.
$\Gamma(F)$	The strict set of coherent global families.
$\Gamma_U(F), \mathrm{Ext}_U(s)$	Partial sections on a full patch and their global extension fiber.
δ, ∂, d, D	Respectively the view differential, tower differential, a model residual, and a Hilbert defect operator. Each target is declared where used.
p, p_U, p_L	A protected observable, restriction to prescribed views, and an innovation observable depending on the splitting.
$\mathfrak{o}_U, H_U$	The existence obstruction and the group of invisible coherent changes. Additive and nonabelian conventions are stated separately.
C_j, J_j, Q_N, S_N	Coarsener, chosen lift, restored coarse remainder, and retained finite reconstruction.
c_f, a_v	Coherent nonabelian transition labels and gauge values.
$\mathfrak{K}_G, J_{\mathrm{comp}}$	The cocycle patch groupoid and its composition-support topology.
$\mathcal{E}_b(R)$	Realizations equipped with an actual identification of their boundary with b.
T_i, K_{ij}, R_p	Transport to the root, the corresponding root loop operator, and reconstruction from the root.

Symbol	Meaning
$\mathcal{H}, q, P$	Root residual, an orthogonal kernel projector, and the explicitly defined model repair. A graph repair need not be orthogonal in the global maximum norm.
$\pi_{W,U}, H_U$	The map forgetting part of the prescription and the corresponding invisible coherent ambiguity.
ϕ, EventualHom	A monotone source schedule and transport classes modulo eventual equality.
$\chi_A, M = SA - BC$	Split exact/residual coordinates and the mismatch between chosen repair lifts.
ω, Q_n, C_n, b_n	A residual, its prefix-product rate, movement coefficient, and summable step budget. Here Q_n is not the dyadic remainder operator Q_N.
$\nu_B, d_W, d_W^{\mathrm{sym}}$	Residual floor, extended directed transport distance, and its symmetric maximum.
$c(x,y), \kappa(x,y)$	The two-bit pair phase and its antisymmetric scalar-gauge invariant.
H, T, r	In the budget chapters, $H = \ker p$ with its inherited Hilbert metric, $T = d\|_H$, and $r = df$ for the true raw residual. Measured data are declared separately.
$\beta_{T,r}(M), a_\alpha$	Minimum residual at correction budget M, and the quadratic regularized correction.
C, b, q	Closure of the correction image, its projection of the residual, and the irreducible residual floor.
$v, g, e, \mathcal{D}_\alpha$	Fit residual, Hilbert normal residual, optimizer error, and regularized dual objective in the solver chapters.
B, G, q	Kernel basis, its Gram matrix, and the coordinate normal covector. This q is not the scalar residual floor used earlier.
$\simeq, \simeq_L$	A specified equivalence, and a continuous linear equivalence. Neither notation alone asserts an isometry.

C.2 From the text to the formal source

The table below locates the foundational results in the accompanying Lean development. Numbered locations refer to results in the main text. Page references locate discussions without a separate numbered statement. Module names are relative to the `KomorebiSpace` directory in the `lean/` source tree. A row may collect several declarations that together express the displayed result.

The table is a navigation aid. A declaration's full type specifies its quantifiers, hypotheses, and conclusion. Those types are the authority for formal coverage. The correspondence does not assign a separate machine-checked counterpart to every explanatory sentence or illustrative calculation. Appendix D gives the complementary map for bounded realization, budgets, regularization, and certificates.

Location	Topic	Lean modules
1.3	Limit universal property	`ViewSections`

Location	Topic	Lean modules
1.4	Functorial transport	`ViewSections`
1.5	Terminal evaluation	`ViewSections`
2.1	Strict section gluing	`SectionGluing`
2.2	Object-only cover failure	`ViewBoundaryExamples`
2.3	Correspondence fibers	`ViewCorrespondence`
2.4	Nonempty fibers without a section	`ViewBoundaryExamples`
3.1	Additive view differential	`SectionObstruction`
3.2	Additive extension criterion	`SectionObstruction`
3.3	Extension ambiguity torsor	`SectionExtensionFibers`
p. 8	Parity seed example	`SectionExtensionExample`
3.5	Exactness at partial sections	`SectionExactSequence`
p. 9	Extension by stages	`SectionExtensionStages`
4.1	Protected correction	`Correction`
4.2	Obstruction witness	`Correction`
6.1	Exact tower maps and realized points	`Towers` `Realization`
6.2	Tower differential composition	`Towers` `TowerCorrection`
6.3	Cofinal point equivalence	`Cofinal`
8.1	Geometric operator repair	`RepairDynamics` `RepairProjection`
8.2	Constructed split tower correction	`SplitTowerRepair`
8.3	Split semantic and residual coordinates	`SplitTowerCoordinates` `SplitTowerSemantics`
8.4	Growing identity repair	`GrowingTower`
p. 27	A bounded residual without bounded primitive	`IdentityTowerObstruction`
p. 28	Dependence on target splitting	`SplittingDependence`
p. 33	Lossless dyadic residual split	`DyadicResidualSplit`
10.1	Cascade telescope and persistent part	`CascadeProjections` `CoarseningCascade`
10.2	Recovery by density	`CascadeDensity`
10.3	Dyadic finite-coordinate recovery	`DyadicInfiniteReconstruction`
10.5	Shift cascade and c0	`SequenceCascade`
11.2	Gauge groupoid and section transport	`GroupViewSystems` `CocycleGaugeGroupoid`
11.3	Nonabelian extension obstruction	`NonabelianExtension`
11.4	Nonabelian extension torsor	`NonabelianExtensionFibers`
11.5	S3 loop model	`CocycleExamples`

Location	Topic	Lean modules
12.1	Effective groupoid descent	`DescentEquivalence` `DescentComparison`
13.1	Composition topology	`CompositionSite`
13.2	Standard stackhood	`CompositionStack`
14.1	Equivalent view presentations	`SectionViewEquivalence` `RelativeSections` `RelativeAmbiguity`
14.3	Identified-boundary realization groupoid	`MarkedRealization`
14.4	Marked transport across a square	`MarkedTransport`
17.1	Transported path telescope	`AmplifiedTransport`
17.2	Holonomy classification	`AmplifiedHolonomy`
18.1	Orthogonal exact repair	`HilbertKernelProjection`
18.2	Closed range equivalence	`HilbertKernelProjection`
19.1	Amplified holonomy gain	`AmplifiedRepair` `HilbertHolonomyRepair`
p. 63	Flat transport amplification	`AmplifiedExamples`
20.1	Two-stage joint projection	`QuantitativeDescent`
20.2	Joint closed range	`DescentClosedRange`
21.1	Two-sided projector motion	`MovingKernelBounds`
21.2	Derived perturbed certificate	`SplitKernelPerturbation`
21.3	Transport between moving kernels	`SplitKernelPerturbation`
21.4	Exact moving-line example	`MovingKernelExamples`
p. 70	Rank loss and radius endpoint	`MovingKernelBounds` `MovingKernelExamples`
p. 72	Motion of quantitative descent	`MovingDescent`
35.1	Finite prefixes without a point	`FinitePrefixBoundary`
10.4	Sharper dyadic support criterion	`DyadicSupport`
18.3	Nearest protected Hilbert repair	`ProtectedHilbertRepair`
35.3	Finite positive and negative certificates	`ProtectedLinearCertificates`
35.2	Finite-fiber inverse-limit existence	`FiniteRealizationExistence`
5.1	Obstructions under stronger prescription	`ObstructionRestriction`
5.2	Prescribed choices and obstructed intermediate steps	`ExtensionChoiceExample`
7.1	Scheduled transports and eventual equality	`ReindexedTransports` `EventualCategory`
7.2	Faithfulness from observable finite stages	`RealizationFaithfulness`
7.3	Finite observation characterizes realized transport	`FiniteObservation`

Location	Topic	Lean modules
7.4	Cofinal transport equivalence	`CofinalTransports`
9.1	Split repair coordinates and choice dependence	`SplitDefectCharts`
9.2	Triangular representation change and repair mismatch	`ResidualNaturality`
9.3	Residual sampling and repair across scales	`ResidualSampling` `SamplingRepair`
15.1	Why literal full framing is presentation dependent	`FramingBoundaryExamples`
22.1	Measured and uncertain residual certificates	`ResidualPerturbation`
22.2	Stable repair when the old exact space is preserved	`ResidualPerturbation`
24.1	Exactification from summable movement	`Exactification`
24.2	Variable-rate exactification and posterior bounds	`Exactification`
24.3	Controlled approximation annihilates protected obstruction	`LinearExactification`
24.4	Path comparison and nonunique protected answers	`Exactification` `RepairExamples`
24.5	What cofinal residual-profile sampling preserves	`Profiles`
25.1	Obstruction floors detect closure, not exact membership	`ObstructionFloor`
25.2	Persistent floors exclude all admissible repair sequences	`ObstructionFloor`
23.1	Independent transport costs and directed geometry	`WeightedGeometry`
23.2	Zero arrows can collapse object distance	`WeightedGeometry`
16.1	An associative phase with no scalar flattening	`ProjectiveExample`

C.3 Using the correspondence

For a proof involving a restriction or a change of representation, compare the source and target spaces before identifying two displayed formulas. For an analytic estimate, check the norm, the permitted correction space, and the operator whose adjoint is used. The declarations preserve these distinctions explicitly. A continuous linear equivalence alone is not an isometry, and a coordinate equation alone does not fix the metric of its residual.

The source contains further definitions, lemmas, and examples beyond this table. The Lean guide explains its organization and the formal treatment of the principal analytic results. Verification records retain the exact source inventories and declaration types from the accepted compilation.

D Proof and formal verification

The printed arguments and the accompanying Lean development have complementary roles. The text explains constructions and their mathematical context. A Lean declaration fixes a precise statement whose proof is checked by the kernel. A correspondence between them records the actual hypotheses and conclusion, while allowing the formal proof to use a different route.

All 16 principal theorem groups on bounded realization, budgets, regularization, and certificates have complete formal counterparts. The table identifies these statements by their numbers in the main text. This coverage concerns the stated results, rather than every proof route or illustrative calculation.

Result	Topic	Formal counterpart
27.2	Bounded approximate feasibility	Verified
27.3	Nested affine minimum corrections	Verified
28.1	Attainment and shape of the budget profile	Verified
28.2	Sharp budget certificate	Verified
28.3	Robust budget comparison	Verified
29.1	Unique regularized correction	Verified
29.2	Parameter and budget comparison	Verified
29.3	Regularization limit alternative	Verified
29.4	Data sensitivity and convergent parameter choice	Verified
30.1	Infinite diagonal family and certified tails	Verified
30.2	Sharp diagonal budget asymptotic	Verified
30.3	Rounded noisy diagonal rate	Verified
31.1	Normal-residual solver certificate	Verified
31.2	Regularized primal and dual gap	Verified
32.1	Coordinate certificate and arbitrary basis change	Verified
32.2	Consistent inexact stopping rule	Verified

D.1 Bounded realization and budget geometry

The formal bounded-realization theorem concerns bounded linear maps between real Hilbert spaces. It gives an exact preimage within the same norm bound as the approximating corrections. Its proof uses compactness of a dual ball and the Riesz isometric equivalence. The printed proof uses the weak-subsequence lemma, Lemma 27.1, whose particular proof route is not separately formalized here. Neither route requires separability or closed range. Neither asserts strong convergence of the original approximating sequence.

The nested-constraint theorem constructs the minimum-norm points of nonempty closed affine sets. It proves the exact energy identity, norm monotonicity, and the equivalence between a nonempty intersection and bounded minimum norms. Under that bound, the corrections converge strongly to

the unique minimum-norm point of the intersection. If the intersection is empty, their norms tend to infinity. The affine assumption is retained: for general convex sets the corresponding energy statement is an inequality.

Namespace	Representative declarations

`BoundedApproximateRealization`	`exists_solution_of_bounded_residual_tendsto_zero`
	`mem_range_iff_bounded_approximation`
	`norm_tendsto_atTop_of_not_mem_range`

`NestedMinimumCorrections`	`nested_affine_energy`
	`nested_affine_intersection_iff_bounded`
	`nested_affine_strong_realization`
	`nested_affine_escape`

`ProtectedBoundedRealization`	`exists_exact_repair_of_bounded_approximation`

All namespaces in this appendix have the prefix `KomorebiSpace`. The protected application uses the actual subtype ker p with its inherited Hilbert norm. It constructs a repaired candidate preserving the declared observable and the original correction budget.

The budget-profile theorems concern bounded operators between real complete Hilbert spaces. The budget is nonnegative and may be zero. Attainment of a residual minimum over a bounded correction ball does not require the entire operator range to be closed. The sharp support formula supplies an actual maximizing witness, expressed formally by `IsGreatest`. It does not infer attainment from the mere existence of a supremum.

Namespace	Representative declarations

`BudgetProfile`	`exists_budget_minimizer`
	`beta_eq_iInf`
	`beta_tendsto_floor`
	`beta_eq_floor_iff`
	`beta_eq_zero_iff`

`BudgetProfileGeometry`	`beta_antitone`
	`beta_zero`
	`beta_convex_combination`
	`beta_profile_bounds`

`BudgetDuality`	`dualValue_le_beta`
	`exists_attaining_dual_witness`
	`sharp_budget_duality`
	`strict_exclusion_certificate`

`BudgetRobustness`	`candidate_residual_le`
	`beta_abs_sub_le`
	`nominal_dual_lower_certificate`
	`excludes_budget_of_nominal_dual_certificate`

The robustness comparison uses the same correction space for both problems. Its bound $\delta + \eta M$ controls data and operator uncertainty at budget M. A changing protected kernel is a different problem and is not covered by that estimate.

D.2 Regularization, noise, and stopping

For positive parameter α, the regularized normal operator is constructed as a continuous linear equivalence. Its inverse defines the correction. The formal development proves uniqueness, the energy identity, parameter monotonicity, budget optimality, and the two data-sensitivity estimates. The fixed-space protected application uses the restricted operator and its actual adjoint.

The limit theory distinguishes projection onto the closure of the range from membership in the range itself. Fitted values converge to the projected datum and residual norms converge to the irreducible floor. If that projected datum is attained, the corrections converge in norm to its unique minimum-norm preimage. Otherwise their norms tend to infinity. The data and stopping theorems use this limit theorem together with their explicit parameter, noise, and solver-tolerance assumptions.

Namespace	Representative declarations

`RegularizedHilbertRepair`	`exists_unique_minimizer`
	`energy_identity`
	`residual_best_at_budget`
	`parameter_norm_antitone`
	`parameter_residual_monotone`
	`data_noise_bound`
	`data_fitted_bound`

`RegularizationLimits`	`fitted_tendsto_projectedDatum`
	`residual_norm_tendsto_floor`
	`weighted_penalty_tendsto_zero`
	`norm_tendsto_atTop_of_not_mem_range`

`RegularizationStrongLimit`	`exists_unique_minimum_and_limit`

`RegularizationNoiseLimits`	`data_noise_coefficient_attained`
	`noisy_parameter_choice`
	`inexact_parameter_choice`
	`stopping_rule_sequence`

The stopping theorem compares proposals, exact regularized corrections, and the intended minimum correction in one fixed Hilbert model. It includes the pointwise three-term bound and the convergence conclusion. Protection is exact at every index. The explicit diagonal schedule at the end of Section 32 combines this solver bound with the formalized noisy rate in Corollary 30.3. The declaration `Diagonal Regularization.diagonal_inexact_rounded_rate` proves this specialization with the literal ceiling, the data-noise hypothesis, and the actual measured-data normal-residual bound. The exact-data bias estimate is derived rather than assumed.

D.3 Certificates and their coordinate metric

The solver certificate controls correction error, fitted error, and objective excess through the actual Hilbert normal residual. The primal and dual identity supplies objective bracketing and a correction-distance

bound for an arbitrary dual witness. These are universal statements over the declared Hilbert spaces, with $\alpha > 0$ wherever coercivity is used.

Namespace	Representative declarations

`RegularizedHilbertRepair`	`correction_error_bound`
	`fitted_error_bound`
	`objective_gap_bound`

`RegularizationCertificates`	`primal_dual_gap`
	`dual_maximizer`
	`gap_controls_correction`
	`residual_witness_gap`

`CoordinateCovector`	`gram_isUnit`
	`protected_coordinate_certificate`
	`computed_normalResidual_changedBasis`
	`computed_certificate_changedBasis`
	`changedBasis_coordinates`

The coordinate certificate uses a basis of the entire permitted correction space, with its inherited Hilbert metric. The proof derives invertibility of the Gram matrix from that basis, identifies the coordinate residual with the pairings of the Hilbert normal residual, and proves

$$g = BG^{-1}q, \qquad \|g\|^2 = q^{\mathsf{T}}G^{-1}q.$$

Every invertible basis change is allowed. The changed coordinates describe the same correction, the residual transforms as a covector, and the displayed squared norm is invariant. The proof includes the zero-dimensional correction space. It does not replace the Hilbert norm by a coordinate supremum norm or assume that the basis is orthonormal. These metric identities hold for every real α. The optimizer estimates still require positive α.

D.4 The infinite diagonal example

Proposition 30.1, Theorem 30.2, and Corollary 30.3 have formal counterparts in the namespace `Komorebi.DiagonalRegularization`. The formal space is the real Hilbert space `lp` at exponent two. Its natural index starts at zero, so formal coordinate j denotes printed coordinate $n = j + 1$. The operator, data, corrections, residuals, and hard prefixes are elements or maps of that actual infinite space.

Result	Declarations in `DiagonalRegularization`
30.1	`proposition_30_1` `a_unique_minimizer`
30.2	`sharp_diagonal_budget_asymptotic`
30.3	`corollary_30_3`
Inexact rate	`diagonal_inexact_rounded_rate`

Result	Declarations in `DiagonalRegularization`
Range	`T_not_closed_range` `w_mem_closure_not_range`

The first pair identifies the universal regularized correction, proves it is the unique global minimizer, equates the infinite squared norms with A_m and B_m, and proves all four global bounds and both nonnegative omitted-tail bounds. The asymptotic theorem proves both scalar limits, the budget-profile limit $\pi/4$ over all real M tending to infinity, and the hard-prefix comparison with constant one. Its budget conclusion is not restricted to the regularization mesh.

The noisy theorem uses the actual attainable datum and its minimum-norm preimage. It retains the literal ceiling in $m = \lceil \delta^{-2/3} \rceil$ and proves both displayed inequalities for every Hilbert perturbation within the stated noise bound. The bias estimate is derived within the model. The inexact rate adds the stated solver tolerance to this same model and parameter choice.

The range obstruction uses actual unit coordinate vectors with residual norm $1/(j+1)$ and applies the general unit-witness theorem. The hard-prefix images converge in Hilbert norm to w, placing it in the closure of the range. An exact preimage would have every coordinate equal to one, contradicting square summability. These declarations formalize the concrete example following Theorem 18.2. Exact finite enclosures remain separate illustrations of the infinite results.

D.5 Verification record and trust boundary

The source contains 166 compilation targets (162 production modules and 4 entry points), 1,859 explicit theorem or lemma declarations, and 5,339 exported declarations. The latter includes definitions and generated helpers; neither count measures independent mathematical discoveries.

Fresh compilation treated warnings as errors and disabled automatic implicit parameters. The complete declaration audit checked defining modules and recursive axioms, rejecting unsafe declarations and project-specific axioms. Source, compiler, cache, and generated-object identities bind this replay.

All 56 diagnostic controls had their specified outcomes: 18 successes, 35 intended elaboration rejections, and three axiom-policy rejections. The 52 inherited controls were rerun alongside four new controls. Earlier records remain historical.

The environment is Lean 4.34.0-rc2 with mathlib commit `aeb72f1e9981b6d35b426b16dca7759be3e7ba06`. Permitted axioms are `propext`, `Classical.choice`, and `Quot.sound`. Lean's compiler and kernel and the pinned external dependency cache remain trusted; the cache was recorded but not rebuilt. Finite executable examples provide separate computational evidence; they are not verified extractions of the Lean proofs.

E Analytic facts used in the proofs

This appendix supplies the analytic ingredients used at several points in the book. They are classical facts, not additional novelty claims. Writing out their proofs makes the dependence on completeness, convexity, the inner product, and the declared operator norm visible. The statements are restricted to the real Hilbert and real Banach settings used in the main arguments. They provide background support for the main results.

The projection and inverse constructions have standard formal counterparts in the mathematical library cited in [4, 5]. The regularization context is also discussed in [6]. The arguments below are ordinary written proofs. The accompanying Lean project does not by itself constitute a formal verification of these printed proofs.

E.1 Projection onto a closed convex set

Lemma E.1 (Existence, uniqueness, and the variational inequality)**.** *Let H be a real Hilbert space, let $C \subseteq H$ be nonempty, closed, and convex, and fix $x \in H$. There is a unique $c \in C$ minimizing $\|x - c\|$. It is characterized by*

$$\langle x - c, z - c\rangle \le 0 \qquad (z \in C). \tag{E.1}$$

If $C = c_0 + L$ is a closed affine subspace with direction space L, the condition is $x - c \perp L$.

Proof. **Produce a Cauchy minimizing sequence.** Put $d = \inf_{z\in C} \|x - z\|$. Nonemptiness makes this a finite nonnegative real number. Choose $c_n \in C$ so that $\|x - c_n\| \to d$. Convexity places $(c_n + c_m)/2$ in C. The parallelogram identity gives

$$\|c_n - c_m\|^2 = 2\|x - c_n\|^2 + 2\|x - c_m\|^2 - 4\left\|x - \frac{c_n + c_m}{2}\right\|^2 \le 2\|x - c_n\|^2 + 2\|x - c_m\|^2 - 4d^2.$$

The right side tends to zero as both indices tend to infinity. Thus (c_n) is Cauchy. Completeness supplies a limit $c \in H$, closedness puts it in C, and continuity of the norm gives $\|x - c\| = d$. A minimizing sequence in an arbitrary closed set need not have been Cauchy. The midpoint was the essential use of convexity.

Prove uniqueness. If $c, c' \in C$ both attain d, the same midpoint inequality gives

$$\|c - c'\|^2 \le 2d^2 + 2d^2 - 4d^2 = 0.$$

Therefore $c = c'$.

Derive the variational inequality. For $z \in C$ and $0 < t \le 1$, convexity makes $c + t(z - c)$ feasible. Minimality and the inner-product expansion yield

$$0 \le \|x - c - t(z - c)\|^2 - \|x - c\|^2 = -2t\langle x - c, z - c\rangle + t^2\|z - c\|^2.$$

Divide by the positive number t and let $t \downarrow 0$. This gives (E.1). Conversely, if that inequality holds, then for every $z \in C$,

$$\|x - z\|^2 = \|x - c\|^2 - 2\langle x - c, z - c\rangle + \|z - c\|^2 \ge \|x - c\|^2.$$

It therefore characterizes the unique minimizer.

For an affine subspace, $c + h$ and $c - h$ are feasible for every $h \in L$. Applying the inequality with both signs gives $\langle x - c, h\rangle = 0$. Orthogonality conversely implies the same variational inequality. In this case the preceding expansion becomes the Pythagorean equality, not merely an inequality. □

For a closed linear subspace L, denote its projection by P_L. The lemma gives a unique decomposition

$$x = P_L x + (x - P_L x), \qquad P_L x \in L, \quad x - P_L x \in L^\perp.$$

Adding two such decompositions, or multiplying one by a scalar, proves linearity by uniqueness. The Pythagorean identity gives $\|P_L x\| \le \|x\|$, so P_L is bounded. It fixes L, annihilates $L^\perp$, and satisfies $P_L^2 = P_L$. For $x, y \in H$, orthogonality gives

$$\langle P_L x, y\rangle = \langle P_L x, P_L y\rangle = \langle x, P_L y\rangle.$$

Thus P_L is self-adjoint. These properties are conclusions of the construction, not extra data assumed for a proposed projection.

E.2 Riesz representation and the adjoint

Lemma E.2 (Representing a bounded functional)**.** *For every bounded real linear functional $\ell : H \to \mathbb{R}$ on a real Hilbert space, there is a unique $u \in H$ with*

$$\ell(h) = \langle h, u\rangle \quad (h \in H), \qquad \|\ell\| = \|u\|.$$

Proof. If $\ell = 0$, take $u = 0$. Otherwise choose x with $\ell(x) \neq 0$. Boundedness makes $L = \ker \ell$ closed. Use Lemma E.1 to form $v = x - P_L x$. Then $v \perp L$ and $\ell(v) = \ell(x) \neq 0$, in particular $v \neq 0$.

For any $h \in H$, the vector $h - \ell(h)v/\ell(v)$ belongs to L. Pairing it with v therefore gives

$$\langle h, v\rangle = \frac{\ell(h)}{\ell(v)} \|v\|^2.$$

Consequently $u = \ell(v)v/\|v\|^2$ represents ℓ. If u' also represents it, pairing the difference of the two formulas with $h = u - u'$ gives $\|u - u'\|^2 = 0$, hence uniqueness.

Cauchy–Schwarz gives $|\ell(h)| \le \|u\| \|h\|$, so $\|\ell\| \le \|u\|$. If $u \neq 0$, the unit vector $h = u/\|u\|$ gives the reverse inequality. The zero case was already included. □

Let $T : H \to R$ be a bounded linear map between real Hilbert spaces. For fixed $z \in R$, the functional $h \mapsto \langle Th, z\rangle$ is bounded by $\|T\| \|z\| \|h\|$. The lemma therefore defines a unique vector $T^*z \in H$ with

$$\langle Th, z\rangle = \langle h, T^*z\rangle.$$

Uniqueness of representatives proves that T^* is linear. The same bound proves $\|T^*z\| \le \|T\| \|z\|$. Conversely,

$$\|Th\| = \sup_{\|z\|\le 1} |\langle Th, z\rangle| \le \|T^*\| \|h\|,$$

where the first equality follows by testing $z = Th/\|Th\|$ when $Th \neq 0$. Hence $\|T^*\| = \|T\|$.

The defining pairing also gives the identity

$$(\operatorname{ran} T)^\perp = \ker T^*. \tag{E.2}$$

Indeed, z is orthogonal to every Th exactly when $\langle h, T^*z\rangle = 0$ for every h, which is equivalent to $T^*z = 0$. Orthogonal complements do not change on taking closure. The projection decomposition then gives

$$\overline{\operatorname{ran} T} = (\ker T^*)^\perp.$$

This is why a self-adjoint injective operator has dense range. It does not yet say its range is closed. In the regularized normal-operator proof, the coercive lower bound supplies that separate fact.

E.3 Why a bounded bijection has a bounded inverse

The closed-range arguments use the inverse theorem only after proving that the relevant source and target subspaces are complete and that the restricted map is bijective. The following derivation records this analytic step explicitly.

Lemma E.3 (Baire's complete-space principle)**.** *A nonempty complete metric space cannot be the union of countably many closed sets with empty interior.*

Proof. Suppose $F_1, F_2, \dots$ are closed with empty interior. Choose a nonempty open ball whose closed ball B_1, of positive radius at most $1/2$, avoids F_1. Such a choice is possible because the complement of F_1 is open and meets every nonempty open set. Recursively choose a closed ball B_{n+1} of positive radius at most $2^{-(n+1)}$, contained in the open ball with the same center and radius as B_n, and disjoint from F_{n+1}. The intersection of that open ball with the complement of F_{n+1} is nonempty and open, so it contains a sufficiently small closed ball. This also covers the case of isolated points.

The centers form a Cauchy sequence: all centers after stage n lie in B_n, whose diameter is at most 2^{1-n}. Completeness gives a limit. For each fixed n, a tail of centers lies in the closed set B_n, so the limit lies there. It therefore belongs to no F_n, contradicting that these sets cover the space. □

Lemma E.4 (Continuous inverse for Banach spaces)**.** *Let $S : X \to Y$ be a bounded linear surjection between real Banach spaces. There is a constant $C < \infty$ such that every $y \in Y$ has a preimage x with $\|x\| \le C\,\|y\|$. If S is also injective, its inverse is bounded linear.*

Proof. If $Y = \{0\}$, take the zero preimage. Otherwise write B_X for the closed unit ball. Surjectivity gives

$$Y = \bigcup_{n \ge 1} \overline{S(nB_X)}.$$

By Lemma E.3, at least one of the closed sets on the right has nonempty interior. Thus there are $n \ge 1$, $y_0 \in Y$, and $\rho > 0$ such that

$$B(y_0, \rho) \subseteq \overline{S(nB_X)}.$$

Here $B(y_0, \rho)$ is an open ball. If $\|h\| < \rho$, both $y_0 + h$ and y_0 belong to this closure. Choose approximating image sequences for both. Subtracting their preimages shows $h \in \overline{S(2nB_X)}$. Scaling by $2n$ gives a constant $c = \rho/(2n) > 0$ with

$$B(0, c) \subseteq \overline{S(B_X)}. \tag{E.3}$$

This only gives approximate preimages. We next turn them into an actual preimage with a summable budget.

For a nonzero residual $r \in Y$, set $\lambda = 2\,\|r\|\,/c$. The vector r/λ has norm $c/2$, so (E.3) supplies $z \in B_X$ with

$$\|Sz - r/\lambda\| \le c/4.$$

Then $u = \lambda z$ satisfies

$$\|u\| \le 2\,\|r\|\,/c, \qquad \|r - Su\| \le \|r\|\,/2.$$

If the residual is zero, choose $u = 0$. Start from $r_0 = y$, repeat this construction, and define $r_{j+1} = r_j - Su_j$. Induction yields

$$\|r_j\| \le 2^{-j}\,\|y\|\,, \qquad \sum_{j \ge 0} \|u_j\| \le \frac{4}{c}\,\|y\|\,.$$

The series of preimages converges in X by completeness, to a vector x with the displayed norm bound. Continuity of S, together with the finite telescoping equation

$$S\left(\sum_{j=0}^{N-1} u_j\right) = y - r_N,$$

proves $Sx = y$. Thus we may take $C = 4/c$.

If S is bijective, the preimage is unique. Uniqueness makes the inverse additive and homogeneous, and the bound just obtained makes it continuous. In the merely surjective case the construction does not assert a linear choice of preimage. In particular it does not assert that every Banach-space surjection has a bounded linear right inverse. □

For example, in Theorem 18.2 the map to be inverted has source $(\ker D)^{\perp}$ and target $\operatorname{ran} D$. Both are closed subspaces of the stated Banach spaces. The inverse lemma applies only after the proof has established those closedness and bijectivity facts. It would not apply with an unclosed range equipped with its inherited norm.

E.4 Neumann inversion with both inverse identities

Lemma E.5 (A strict perturbation radius)**.** *Let W be a real Banach space and $T : W \to W$ bounded linear, with $\|T\| \le \theta < 1$ and $\theta \ge 0$. Then*

$$Z = \sum_{j=0}^{\infty} T^j$$

exists in operator norm and satisfies

$$(I - T)Z = Z(I - T) = I, \qquad \|Z\| \le \frac{1}{1-\theta}.$$

Proof. First recall why the space of bounded operators on W is complete. An operator-norm Cauchy sequence A_n makes $A_n x$ Cauchy for each x, and hence defines $Ax = \lim_n A_n x$. Pointwise limits preserve linearity. The sequence of operator norms is bounded, so the pointwise limit is a bounded operator. For a fixed sufficiently large n, pass $m \to \infty$ in $\|(A_n - A_m)x\| \le \varepsilon \|x\|$ to obtain $\|(A_n - A)x\| \le \varepsilon \|x\|$. Taking the supremum over the unit ball proves operator-norm convergence. Completeness of W is the step that produced each pointwise limit.

Submultiplicativity gives $\|T^j\| \le \theta^j$, with the convention $T^0 = I$. The finite partial sums are therefore operator-norm Cauchy, since for $M > N$,

$$\left\| \sum_{j=N+1}^{M} T^j \right\| \le \sum_{j=N+1}^{M} \theta^j \le \frac{\theta^{N+1}}{1-\theta}.$$

Their limit Z has norm at most $\sum_{j\ge 0} \theta^j = 1/(1-\theta)$. The operator identities for finite sums are

$$(I - T)\sum_{j=0}^{N} T^j = I - T^{N+1}, \qquad \left(\sum_{j=0}^{N} T^j\right)(I - T) = I - T^{N+1}.$$

Multiplication by a fixed bounded operator is continuous in operator norm, and $\|T^{N+1}\| \le \theta^{N+1} \to 0$. Passing to the limit proves both inverse identities. When W is the zero space the same statements hold, with its unique identity map. When $\theta = 1$, the summable tail and decay used here disappear, so the proof supplies no endpoint conclusion. □

E.5 The dependency order

The analytic arguments can now be read without a circular invocation of their conclusions. Projection gives the orthogonal decomposition and Riesz representation. The latter gives adjoints and the relation between range and adjoint kernel. Baire's principle independently gives the Banach inverse theorem. Those ingredients yield closed-range estimates and the positive regularized normal inverse. Neumann's series gives the separate small-perturbation inverse and its explicit radius.

The weak-subsequence proof in Section 27 uses projection onto the closed separable span generated by the sequence, followed by a coordinate diagonal argument. It does not assume that the whole ambient Hilbert space is separable. The budget-attainment and regularization-limit proofs use that sequential result. Strong convergence is proved later from the minimum-norm condition and the norm identities, not inferred from weak convergence alone.

F Background and references

The references identify standard sources for limits, descent, Hilbert projection, bounded inverse theorems, and regularization. Appendix E gives the analytic arguments used repeatedly in the text. The accompanying Lean source and its pinned mathematical-library revision determine the formal interfaces used in this edition. Online library documentation is a navigation aid and may describe a different revision. The consultation dates below are retained from the recorded source checks.

References

[1] The Stacks Project Authors. *Limits and colimits over preordered sets*. Section 4.21, Tag 002Z, including the finite-set inverse-limit theorem in Lemma 4.21.7. https://stacks.math.columbia.edu/tag/002Z. Accessed 10 October 2026.

[2] The Stacks Project Authors. *Stacks*. Section 8.4, Tag 0268. In particular the effective descent and equivalence-of-descent-categories formulations. https://stacks.math.columbia.edu/tag/0268. Accessed 10 October 2026.

[3] The mathlib Community. *CategoryTheory.Sites.Descent.IsStack*. Standard pseudofunctor stack interface. https://leanprover-community.github.io/mathlib4_docs/Mathlib/CategoryTheory/Sites/Descent/IsStack.html.

[4] The mathlib Community. *Analysis.InnerProductSpace.Projection.Basic*. Orthogonal projection onto complete subspaces. https://leanprover-community.github.io/mathlib4_docs/Mathlib/Analysis/InnerProductSpace/Projection/Basic.html.

[5] The mathlib Community. *Analysis.Normed.Operator.Banach*. Banach open mapping and continuous inverse constructions. https://leanprover-community.github.io/mathlib4_docs/Mathlib/Analysis/Normed/Operator/Banach.html.

[6] Christian Clason. *Regularization of Inverse Problems*. Graduate lecture notes, arXiv:2001.00617v2, 2021. Sections 2, 4, and 6 provide standard background on weak convergence, parameter choice, and Tikhonov regularization. https://arxiv.org/abs/2001.00617v2. Consulted 12 October 2026.