From Initial Diagrams to Coherent Congruence

Solutions, rewrites, and monads beyond sets

Anonymous 8 November 2026

Abstract

Evan Patterson interprets diagrammatic equation proofs as changes of shape that preserve cones. We develop a calculus in which the allowed change depends on the homotopy information carried by solutions. For a shape map $u : A \to K$, the covariant weight $k \mapsto \tau_{\leq n}|A_{/k}|$ detects preservation of all limits with n-truncated mapping spaces and classifies diagram shapes up to the resulting zigzags. We give an explicit orthogonal factorization through a left fibration with n-truncated fibers and prove a cartesian form that can be tested using only product-preserving models. We establish relative versions in slices of an infinity-topos and construct an ordinary internal congruence quotient that commutes with geometric inverse image. Worked examples distinguish smooth representability, groupoid automorphisms, sheaf gluing, and higher descent. Ordinary term monads, a strict monoidal 2-monad, and analytic monads show which parts of the calculus survive substitution and which depend on the level of observation.

Keywords. E-graphs, initial functors, left fibrations, truncated weights, congruence, smooth sheaves, descent, analytic monads.

MSC 2020. 18A30, 18N60, 18C15, 18F20, 08A30, 68T15.

Contents

1 From equations to solution functors **3**
1.1 Conventions and the two uses of a diagram 4

2 Ordinary initiality and unary calculations **5**
2.1 A cone calculation with comma categories 5
2.2 Motifs, pushouts, and shared endpoints 6
2.3 Pasting two squares 7
2.4 Two cycles and their congruence 8

3 Truncated weights and coherent cones **9**
3.1 The covariant weight 9
3.2 Transport through an attachment 10
3.3 Sharpness of the levels 11

4 Comprehensive factorization at each level **12**
4.1 A right class and an explicit left factor 12
4.2 Cartesian Kan extension by a comma adjunction 13
4.3 Classification and the ordinary specialization 15
4.4 A groupoid obstruction in a product theory 16

5 Smooth equations and geometric solution objects **18**
5.1 A noncommutative equation in a smooth product context 18
5.2 Representable equalizers 19
5.3 A singular equalizer and its smooth sheaf 20
5.4 Relative smooth solutions and loss of submersivity 21

6 Groupoids and the geometry of witnesses **22**
6.1 Isomorphism witnesses in a pullback 22
6.2 A double-coset calculation 22
6.3 Actions and coherent fixed points 23

7 Sheaves and internal congruence **24**
7.1 A solution sheaf and its sections 24
7.2 Stages, stalks, and exactness 25
7.3 Constructing the internal congruence 25
7.4 A congruence in the topos of group actions 27

8 Relative weights and descent in an infinity-topos **28**
8.1 A test over a parameter object 28
8.2 Čech gluing and coherent local choices 29
8.3 A sheaf of witnesses with nontrivial transport 30
8.4 A locally solvable lift that does not glue 31

9 Monads, substitution, and higher arities **32**
9.1 Free syntax in a topos 32
9.2 Solution diagrams of algebras 34
9.3 A 2-monad and its pentagon 34
9.4 Analytic monads and isotropy of operations 35

10 Relations as chosen witnesses **36**
10.1 The universal algebra with a specified path 36
10.2 A free commutative monoid with one redundant equation 38

11 Additive and stable examples **39**
11.1 Eliminating a variable in an additive category 39
11.2 Derived fibers of the reduced equation 40

12 Changing the index and imposing finite-limit contexts **41**
12.1 A pullback theorem for cartesian fibrations 41
12.2 Composition of weights 42
12.3 Finite limits and the effect of truncation 42

13 Partial operations and geometric groupoids **44**
13.1 Definitional extension in a fixed domain 44
13.2 An internal category as a finite-limit theory 45
13.3 Free actions, effective relations, and stabilizers 45

14 Using the calculus as a certificate **46**
14.1 Local certificates and their scope 46
14.2 Worked comparison of three candidate rewrites 47
14.3 What each target tests 48

References **49**

1 From equations to solution functors

A diagram can record an expression, a collection of constraints, or the choice of a witness. The difference matters even for the simplest rewrite. If a diagram contains x, an operation f, and a new variable $y = f(x)$, the new variable is determined by x. If it contains x and an arbitrary loop at x, the loop may be an independent part of a solution. When solutions are points of a set, that distinction can disappear. For groupoids, it is the difference between a redundant variable and an automorphism. For sheaves of spaces, the same issue can recur at every stage of gluing.

The starting point is Patterson's first article on initial functors and e-graphs [16]. Its two calculations use unary operations in an arbitrary category. The later installments introduce e-diagrams and product contexts [17, 18]. The subsequent preprint proves ordinary and cartesian soundness and semantic completeness via the image copresheaf and comprehensive factorization [19, Sections 3–4 and Appendices A–B]. These results are the classical baseline. The use of connected comma categories, the ordinary saturation of diagrams by discrete opfibrations, and the passage from unary terms to finite-product contexts belong to that baseline.

Equality saturation uses an e-graph to retain many equivalent terms at once, with rebuilding or congruence closure ensuring that equivalent inputs give equivalent outputs [15, 25]. Other categorical models describe e-graphs using labeled hypergraphs and rewriting systems [3, 24]. Those presentations address how shared syntax is organized. Here the invariant of a diagram is instead defined by how the diagram acts on solutions. The relation with earlier higher treatments is direct. The journal's earlier note [1] already treats higher cartesian Kan extension and truncated soundness. The present work uses an explicit comma adjunction, proves the truncated orthogonality and cartesian detector classification, and develops internal congruence and monadic substitution with geometric

base change. The calculations in manifolds, groupoids, and sheaves specify what the detector sees in concrete models.

We work at each truncation level with the full covariant weight of a shape map. The formula for the weight is a pointwise left Kan extension. Its component set gives the ordinary image copresheaf. Its higher homotopy groups account for obstructions to replacing shapes when cones retain coherent choices. The main statements are the truncated detection theorem (Theorem 3.2), the factorization with an explicit formula for the slices of its unit (Theorem 4.1), and the cartesian semantic classification (Theorem 4.5). The classification includes transformations specified by a map of shapes and abstract equivalences of solution functors.

There are two kinds of change of context in the applications. One changes the indexing shape, with the model held fixed. The other changes the ambient category or a parameter object, with the underlying system held fixed. The former is tested by comma weights. The latter requires exactness of a functor, a representability argument, or a descent theorem. For instance, a restriction of a smooth solution sheaf is always defined, but a singular solution sheaf need not come from a smooth manifold. The smooth argument below proves this failure directly. In a Grothendieck topos, the ordinary congruence of a finitary algebra can instead be constructed from finite inference steps and then quotiented. Its stability under geometric inverse image follows from the operations used in its construction.

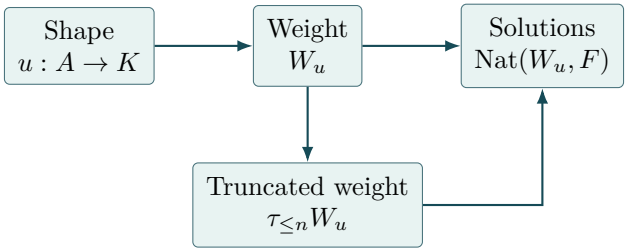


Figure 1.1. A shape determines a weight. An n-truncated model pairs with the truncated weight.

The paper uses a single distinction in several settings. An ordinary congruence is a relation that can be saturated and quotiented. A space of solutions represents points, paths, and higher paths. The relation of lying in the same connected component can be recovered from that space, and doing so is an additional operation. Neither construction presupposes a terminating equality-saturation procedure. The equivalences proved here are about diagram semantics and canonical comparison maps.

1.1 Conventions and the two uses of a diagram

Write $\mathcal{S}$ for spaces and Cat_∞ for small infinity-categories. Here an infinity-category means an $(\infty, 1)$-category. Limits and colimits in such a category are coherent. Ordinary categories and their strict constructions are explicitly indicated. We use a pair of universes so that a diagram shape is small and a category of models can live in the larger universe. A cartesian category has finite products, and a cartesian functor preserves them. We write composition gf for first f, then g. For an ordinary diagram, Hom

denotes a set. For a coherent diagram, Map and Nat denote spaces.

A diagram $D : A \to \mathcal{C}$ gives a functor of test objects

$$\mathrm{Sol}_D(Q) = \mathrm{Nat}(\Delta Q, D),$$

where ΔQ is constant at Q. A limit is a representing object for this functor. Alternatively, fix a signature category K and vary its interpretation $F : K \to \mathcal{C}$. A shape $u : A \to K$ then assigns to F the solution object of Fu, when it exists. We use both viewpoints. In Man the first is important because the functor of probes still exists when it is unrepresentable. In the semantic classification the second is important because the variation in F detects the weight.

In a product context $\Gamma = X_1 \times \cdots \times X_m$, variables are projections and a term of sort Y is an arrow $\Gamma \to Y$. For $m = 0$, Γ is terminal. An equation in context may either be a theorem about two arrows already equal in a theory or an additional constraint to be imposed in a model. The first gives a commutative triangle or a specified coherent filler in the theory. The second gives an equalizer or a homotopy pullback in the model. The certificate for changing a shape must respect this distinction.

2 Ordinary initiality and unary calculations

2.1 A cone calculation with comma categories

Let $u : A \to K$ be an ordinary functor. Its comma category $A_{/k}$ has objects $(a, h : u(a) \to k)$. An arrow to (a', h') is an arrow $r : a \to a'$ with $h = h'u(r)$. Connected means nonempty and connected by finite zigzags. We call u initial when every $A_{/k}$ is connected. This is the orientation for restriction of limits, as in [16, 20].

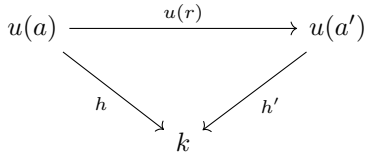


Figure 2.1. A morphism in $A_{/k}$. The arrows in the comma category point toward the fixed object k.

Proposition 2.1 (Ordinary cone invariance)**.** *The following are equivalent for u. Every $A_{/k}$ is connected. For every ordinary $F : K \to \mathsf{Set}$, restriction gives a bijection on sets of cones with vertex 1. For every ordinary $D : K \to \mathcal{C}$, restriction gives a bijection on cones with any vertex Q. In the last case representability on either side implies representability on the other.*

Proof. For a cone $(x_a : Q \to D(u(a)))_{a \in A}$, define its leg at k by $D(h)x_a$ using any $(a, h) \in A_{/k}$. A comma arrow makes the two resulting legs equal by the cone identity, and connectedness extends this to a zigzag. Postcomposition of h proves naturality at an arrow $k \to k'$. Taking $h = \mathrm{id}_{u(a)}$ proves that the construction inverts restriction, and the same formula proves uniqueness.

For the converse, form the set-valued functor

$$W_{u,0}(k) = \pi_0 A_{/k}.$$

A map of copresheaves $W_{u,0} \to F$ is determined by the elements of $F(u(a))$ assigned to $[a, \mathrm{id}]$. The only compatibility of these elements is the cone identity, since the value at $[a, h]$ must be $F(h)$ applied to the value at $[a, \mathrm{id}]$. Thus $\mathrm{Nat}(W_{u,0}, F) \cong \lim_A Fu$. Restriction along u corresponds to precomposition with $W_{u,0} \to \mathbf{1}$, where $\mathbf{1}$ is the terminal copresheaf. If this gives bijections for all F, the Yoneda lemma in $\mathrm{Fun}(K, \mathsf{Set})$ makes the augmentation an isomorphism. Representability follows from naturality in Q. □

The proof explains why choosing a vertex in a free term graph can certify equalities in a category without general limits. If A has a root r which is initial, its singleton inclusion is initial. Every cone is determined by its component at r. A sound change from this diagram retains the solution represented by $D(r)$, even if the new shape has many equations. When the cone leg at the retained root is the identity, any two parallel paths out of the root with a shared endpoint have equal images. This argument uses the identity as a test leg. Agreement on global elements of $D(r)$ would only prove arrow equality in a category in which those elements detect arrows.

2.2 Motifs, pushouts, and shared endpoints

We give three certificates for the motifs used in [16, 19]. Write $I = [1]$, rooted at 0, and $I^{(j)}$ for the free category on a root with j outgoing edges to distinct leaves. The unique root of $I^{(0)}$ is its only vertex. The fold $I^{(j)} \to I$ takes each edge to the generating arrow. It is initial, since the functors from the root into source and target are initial and restriction through their composite is a bijection. When a diagram is interpreted in $\mathcal{C}$, the fold is over $\mathcal{C}$ precisely when all arms have the same operation and the same interpretations of the endpoints. In a product context a common input tuple takes the place of the root object.

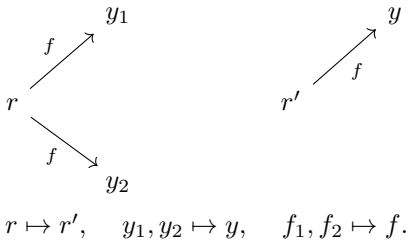


Figure 2.2. Operation introduction uses the root inclusion. Congruence reduction uses the fold from the two-arm shape to the one-arm shape.

For the path motif, let p and q be composable strings of lengths m and n with the same source and target. If their composites agree in $\mathcal{C}$, the interval over that composite maps to both chains, sending its generating arrow to the respective full path. Both maps preserve roots and are initial. A path of length zero is represented by $[0]$, with the interval arrow sent to an identity. The equality of composites is what allows the two maps to form a span over $\mathcal{C}$.

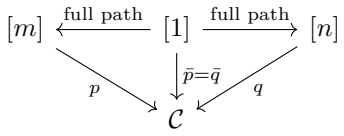


Figure 2.3. A path comparison through the single arrow over the common composite. The span also works when one side is an identity.

For contextual application, take a pushout $H = B \sqcup_A J$ of $u : A \to B$ along $A \to J$, with a compatible diagram on H. Cones on H with fixed vertex Q are pairs of cones on B, J with matching restrictions to A. If u is initial, each cone on J has exactly one compatible extension on B. This proves that $J \to H$ is initial, without assuming that the target category has all limits. The pushout is taken in the category of shapes. If two parallel paths are to be identified, the pushout may introduce relations between arrows of the indexing category. An equality of images in $\mathcal{C}$ by itself does not identify the paths inside the shape.

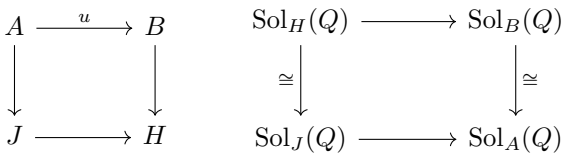


Figure 2.4. A pushout of shapes induces a pullback of cones. The right-hand vertical bijection propagates to the left.

An endpoint merge requires a certificate. In the free category on $p, q : x \rightrightarrows y$, the inclusion of x has two components in its comma category at y. In a set model with $D(x) = D(y) = \{0, 1\}$, $D(p) = \mathrm{id}$, and $D(q)$ the transposition, the cone set is empty. Before the merge, the fork with separate sinks has two cones. Even if both labels are rendered by the same symbol, the initiality claim concerns the indicated map of shapes. A quotient forcing $p = q$ gives a different indexing category with a root, which is valid over interpretations in which the path equation holds.

2.3 Pasting two squares

Use the presented category with the diagram in Figure 2.5 and relations $kf = f'h$ and $\ell g = g'k$. We want a certificate from the two equations to $\ell g f = g' f' h$. The outside paths are terms in one input $x : X$. In a cone over the single-root diagram their values are the two possible images of x.

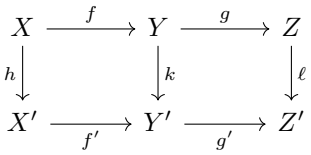


Figure 2.5. The inner squares give $kf = f'h$ and $\ell g = g'k$. The outside rectangle is the goal.

The path motif first replaces the fragment ℓg by $g'k$, keeping f as a prefix. It then replaces kf by $f'h$, keeping g' as a suffix. To see these as actual attachments, start with

an interval labelled by the composite of the first fragment and map it into the full first path. Apply the two contextual path attachments along that interval, read the resulting pair of paths, and use the common endpoints. Repeat with the second fragment. All maps from the two path replacements are initial by Figure 2.4. The resulting calculation in the signature category is

$$\ell g f = (g'k)f = g'(kf) = g'(f'h) = g'f'h.$$

The two middle rebracketings are composition in the shape, independent of any equations of the model. If a model is a category of smooth maps, the same proof establishes equality of smooth maps. If a model is coherent and the input squares carry homotopies, the corresponding pasting gives a path between the outside composites. The choice of that path may itself carry information.

2.4 Two cycles and their congruence

Consider the one-object theory generated by an endomorphism f with $f^3 = \mathrm{id} = f^5$. This example goes back to congruence closure [15, 16]. Start with the shape whose graph consists of a 3-cycle and a 5-cycle meeting at r. In the indexing category impose that each full cycle from r is the identity. This quotient shape is over the theory. The root r is initial because a path from r to any vertex consists of full turns at r followed by a unique partial turn, and each full turn can be removed by the imposed relations.

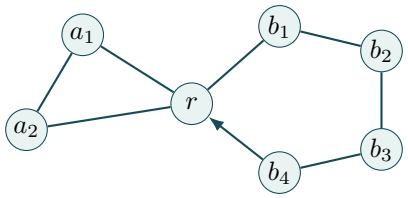


Figure 2.6. Every edge has label f. Both complete turns at r are identities in the indexing category.

Step	Source class	Targets merged by the fold
1	r	a_1, b_1
2	$[a_1] = [b_1]$	a_2, b_2
3	$[a_2] = [b_2]$	r, b_3
4	$[r] = [b_3]$	a_1, b_4
5	$[a_1] = [b_4]$	a_2, r
6	$[a_2] = [r]$	r, a_1

Each row is a pushout of the two-arm fold. In row 5, for example, $a_1 \to a_2$ and $b_4 \to r$ start in the same class by row 4. The quotient functor in each row sends the indicated pair of arrows to one arrow and respects all earlier relations. After row 6, all vertices have the class of r. The composite from the root singleton is still initial. For the cone with identity leg at this vertex, the retained self-loop forces $f\mathrm{id} = \mathrm{id}$. This is an explicit certificate in an arbitrary category with the indicated endomorphism, including Man and a category of sheaves.

More generally, for positive m, n, the monoid presentation $f^m = f^n = 1$ presents $\mathbb{Z}/d\mathbb{Z}$, where $d = \gcd(m, n)$. Indeed, f^{m-1} is a two-sided inverse for f. Bézout's identity therefore implies $f^d = 1$. The map to $\mathbb{Z}/d\mathbb{Z}$ sending f to 1 mod d satisfies both equations, so the period cannot be reduced further by the presentation. In a model, however, the order may divide d. For a smooth rotation of the circle by $2\pi/d$, the order is exactly d. The arithmetic calculation is a statement about equations. A choice of homotopies witnessing two periods for a space action could carry additional coherence, which the ordinary calculation does not classify.

3 Truncated weights and coherent cones

3.1 The covariant weight

For a small infinity-category A, its groupoid completion $|A| = \mathrm{colim}_A *$ is a space. For an ordinary category this is the homotopy type of its nerve. For $u : A \to K$ define

$$A_{/k} = A \times_K K_{/k}, \qquad W_u = \mathrm{Lan}_u \mathbf{1}_A, \qquad W_{u,n} = \tau_{\leq n} W_u. \tag{3.1}$$

The projection $K_{/k} \to K$ takes the domain of an arrow to k. Pointwise left Kan extension gives $W_u(k) \simeq |A_{/k}|$ [6, 14, 12]. Postcomposition makes it covariant in k. A map $v : A \to A'$ over K induces $W_u \to W_{u'}$, and hence $W_{u,n} \to W_{u',n}$. We take $n \in \{-1, 0, 1, \ldots, \infty\}$. The convention is $\tau_{\leq\infty} = \mathrm{id}$, and (-1)-truncated spaces are the empty and the contractible spaces up to equivalence.

For $F : K \to \mathcal{S}$, a point of $\mathrm{Nat}(W_u, F)$ is a coherent section of Fu. This follows from the Kan extension adjunction:

$$\mathrm{Nat}_K(W_u, F) \simeq \mathrm{Nat}_A(\mathbf{1}_A, Fu) \simeq \lim_A Fu. \tag{3.2}$$

Under this equivalence the map from $\lim_K F$ is precomposition with the canonical augmentation $W_u \to \mathbf{1}_K$. For example, a generator (a, h) in $W_u(k)$ must be sent to $F(h)$ applied to the solution at a. A morphism in the comma category supplies the path witnessing compatibility of this prescription. The adjunction packages the higher compatibilities.

Definition 3.1. We call u n-initial if $W_{u,n} \to \mathbf{1}_K$ is an equivalence. Write L_n for this class. For $n = \infty$ this is the homotopy-coherent initiality criterion, also called left cofinality. Call $\mathcal{C}$ locally n-truncated if all its mapping spaces are n-truncated.

An n-connected space is a space with contractible n-truncation. An n-connected map is a map whose homotopy fibers have that property. Thus a 0-connected space is nonempty and connected. A 1-connected space is nonempty and simply connected. The expression n-initial refers to this connectivity test, even though the models used to detect it have n-truncated mapping spaces. An ordinary u is initial in the sense of Section 2 exactly when its nerve lies in L_0.

Theorem 3.2 (Detection and transport of cones)**.** *For $u : A \to K$, the following conditions are equivalent.*

1. $u \in \mathsf{L}_n$.
2. *For every* $F : K \to \mathcal{S}_{\leq n}$, *the canonical map* $\lim_K F \to \lim_A Fu$ *is an equivalence.*
3. *For every locally* n*-truncated* $\mathcal{C}$, *diagram* $D : K \to \mathcal{C}$, *and test object* $Q \in \mathcal{C}$, *the restriction* $\mathrm{Sol}_D(Q) \to \mathrm{Sol}_{Du}(Q)$ *is an equivalence.*

Under these conditions a limit of one diagram exists if and only if a limit of the other exists. If they exist, the comparison of limits is an equivalence.

Proof. Pointwise truncation is left adjoint to inclusion on the functor category. When F is n-truncated valued, (3.2) therefore pairs F with $W_{u,n}$. The asserted comparison is induced by $W_{u,n} \to \mathbf{1}_K$. This gives $1 \Rightarrow 2$. Conversely, the Yoneda lemma in $\mathrm{Fun}(K, \mathcal{S}_{\leq n})$ detects that map by testing all F in that category. This gives $2 \Rightarrow 1$. For $2 \Rightarrow 3$, use $F(k) = \mathrm{Map}_{\mathcal{C}}(Q, D(k))$. The space of cones with vertex Q is its limit. The case $\mathcal{C} = \mathcal{S}_{\leq n}$ and $Q = *$ gives $3 \Rightarrow 2$. The statements about limits follow by representability and naturality in Q. □

$$\begin{array}{ccc} \lim_K F & \longrightarrow & \lim_A Fu \\ \simeq\downarrow & & \downarrow\simeq \\ \mathrm{Nat}(\mathbf{1}_K, F) & \xrightarrow[\text{precompose}]{} & \mathrm{Nat}(W_{u,n}, F) \end{array}$$

Figure 3.1. The lower map tests the augmentation of the weight. Its direction is opposite to $W_{u,n} \to \mathbf{1}_K$.

The theorem concerns restriction itself, as opposed to equivalences of solutions chosen separately for each model. The qualification on mapping spaces can be localized at a single probe. Even if $\mathcal{C}$ has objects with arbitrarily high automorphisms, the comparison at Q holds whenever the spaces $\mathrm{Map}_{\mathcal{C}}(Q, D(k))$ are n-truncated for the given Q.

3.2 Transport through an attachment

Proposition 3.3 (Contextual closure)**.** *The class* L_n *contains equivalences and functors with a right adjoint. It is closed under composition, and if* u *and* vu *are in* L_n, *then so is* v. *It is stable under pushout in* Cat_∞.

Proof. A functor with a right adjoint has a terminal object in each comma category of (3.1). Composition and the stated cancellation follow by applying Theorem 3.2 to the sequence of restrictions, with a diagram on the last shape. For the pushout $H = B \sqcup_A J$, the universal property gives

$$\mathrm{Fun}(H, \mathcal{C}) \simeq \mathrm{Fun}(B, \mathcal{C}) \times_{\mathrm{Fun}(A,\mathcal{C})} \mathrm{Fun}(J, \mathcal{C}).$$

Indeed Cat_∞ is cartesian closed, so testing with arbitrary parameter categories upgrades the mapping-space equivalence to this equivalence of functor categories. Taking the mapping space from the constant diagram at Q to a diagram on H gives the homotopy pullback of the three cone spaces. With $\mathcal{C} = \mathcal{S}_{\leq n}$ and $u : A \to B$ in L_n, the restriction to A on the B-factor is an equivalence. The other vertical comparison is therefore an equivalence. Vary the diagram and use Theorem 3.2. □

This argument justifies coherent versions of the rooted motifs. For example, if $p, q : X \to Y$ are equipped with a path in $\mathrm{Map}_{\mathcal{C}}(X, Y)$, two interval diagrams labelled by p and q are equivalent over $\mathcal{C}$. Indeed, the path gives a natural transformation whose endpoint components are identities, and it is therefore a natural equivalence. After introducing paths from their composites, the shape maps are in L_∞ because they preserve initial roots. If the available datum is that a path exists, a particular map over $\mathcal{C}$ still requires a choice of a path. If two path labels are to be identified, a higher comparison may be required to make the intended attachment commute.

There is a difference between an ordinary pushout and the pushout in Cat_∞. The latter retains the equivalence data relating the two restrictions in its defining square. A quotient imposing path equality in an ordinary indexing category is useful when the model already supplies the equality. In a coherent model it is the coherent universal property that supplies the map on cone spaces. Identifying the displayed endpoints without using that universal property can alter the answer.

3.3 Sharpness of the levels

The classes form a decreasing chain

$$\mathsf{L}_{-1} \supseteq \mathsf{L}_0 \supseteq \mathsf{L}_1 \supseteq \cdots \supseteq \mathsf{L}_\infty.$$

A test over the point already separates adjacent levels. For the shape $A \to *$ and constant coefficient Y, the comparison is

$$Y \longrightarrow \mathrm{Map}(|A|, Y), \tag{3.3}$$

which sends a value to the constant map. Two isolated vertices give a (-1)-initial map but a non-0-initial one, as seen by the diagonal on a set with two elements. The circle shape gives a 0-initial map that is not 1-initial. More generally, for $m \geq 1$, take the poset of nonempty proper faces of an $(m+1)$-simplex. Its nerve is the barycentric subdivision of the boundary sphere S^m. The collapse to a point is $(m-1)$-initial. It fails the m-test with $Y = K(\mathbb{Z}, m)$, since a based map of class $1 \in \pi_m Y$ cannot be homotopic to a constant map. Evaluation at a base point has fiber $\Omega^m Y$, which has components indexed by $\mathbb{Z}$. These counterexamples use finite ordinary posets as shapes.

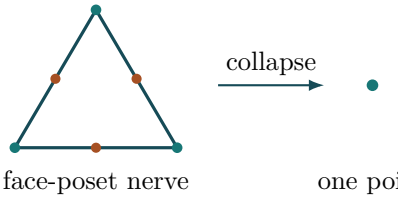


Figure 3.2. The face poset of the triangle boundary has six nonidentity incidences. Its realization is a circle, which disappears at level 0 and survives at level 1.

The comma test is relative. Contractibility of $|A|$ and $|K|$ does not suffice. The inclusion of the terminal vertex $\{1\} \to [1]$ induces an equivalence on groupoid completions, but its comma category at 0 is empty. Conversely, the inclusion of the initial vertex is in

L_∞. Its pullback along $\{1\} \to [1]$ is $\varnothing \to \{1\}$, which fails even the (-1)-test. These two diagrams distinguish invariance under pushout from stability under arbitrary pullback of shapes. Pullback of a parameter in a topos, studied below, is a different operation.

4 Comprehensive factorization at each level

4.1 A right class and an explicit left factor

A left fibration $p : E \to K$ is classified by a covariant functor $K \to \mathcal{S}$ under straightening. An object in its fiber over k is a point of the value at k. Transport over $k \to k'$ is the map of values of the classifier. We use the notation $\int_K F$ for its total category. A map over K from an arbitrary $u : A \to K$ into this total category is a section of u^*p. The space of such sections is $\lim_A Fu$. These are the straightening and section descriptions [6, 12, 14]. Define R_n to be the class of left fibrations whose fibers are n-truncated, up to equivalence over the base.

Theorem 4.1 (Truncated comprehensive factorization)**.** *The pair $(\mathsf{L}_n, \mathsf{R}_n)$ is an orthogonal factorization system on Cat_∞. For $u : A \to K$ its factorization is*

$$A \xrightarrow{\eta_{u,n}} E_{u,n} := \textstyle\int_K W_{u,n} \xrightarrow{p_{u,n}} K. \tag{4.1}$$

For $e = (k, w)$ in $E_{u,n}$ there is a canonical equivalence

$$|A_{/e}| \;\simeq\; \mathrm{hofib}_w(W_u(k) \longrightarrow W_{u,n}(k)). \tag{4.2}$$

Here $A_{/e}$ uses $\eta_{u,n}$. The factorization is natural up to coherent equivalence. For an ordinary functor and $n = 0$ it gives the initial–discrete-opfibration factorization of Street and Walters.

Proof. The section defining $\eta_{u,n}$ is the adjunction unit from $\mathbf{1}_A$ to $u^*W_{u,n}$. At a it uses the class of $(a, \mathrm{id}_{u(a)})$. The space of lifts of u into a right map classified by $F : K \to \mathcal{S}_{\le n}$ is

$$\mathrm{Map}_{/K}(A, \textstyle\int_K F) \simeq \mathrm{Nat}(W_u, F) \simeq \mathrm{Nat}(W_{u,n}, F) \simeq \mathrm{Map}_{/K}(E_{u,n}, \textstyle\int_K F).$$

The equivalences are natural in F and are induced by the indicated unit. Thus (4.1) is obtained by a reflection, and the class of reflected objects is closed under retracts in the slice over K.

For the comma formula, an arrow from $\eta_{u,n}(a)$ to e consists of $h : u(a) \to k$ and a path from the class of (a, h) to w in $W_{u,n}(k)$. This description is natural in the comma object (a, h). The resulting category over $A_{/k}$ is the category of elements of the pullback of the path family ending at w. The groupoid completion of the category of elements of a space-valued functor is its colimit. Indeed, maps from that completion to a space Z are precisely coherent maps from all the fibers to Z, by straightening and the universal property of a constant target. Here the family factors through $|A_{/k}|$. Its colimit is the total space of the corresponding family on $|A_{/k}|$. To check this reduction, a cocone into any space can be computed as a map from the family to the constant copresheaf on that space. Both copresheaves invert every arrow, so these maps factor, including their higher

compatibilities, through groupoid completion. The total space is the homotopy fiber in (4.2). The fibers of a space truncation map have contractible n-truncation. Hence $\eta_{u,n} \in \mathsf{L}_n$.

For orthogonality, pull an arbitrary right map in a proposed lifting square back along the bottom map. This reduces the lifting problem to extending a section of a left fibration with n-truncated fibers. By Theorem 3.2, the restriction on its spaces of sections is an equivalence. Each homotopy fiber of this restriction is contractible. Conversely, testing left orthogonality against all such left fibrations over the target, with bottom map its identity, recovers condition 2 of that theorem. A map right orthogonal to L_n is a retract of its right factor over the target by lifting against its own left factor. Closure under those retracts follows from the slice reflection above. This proves both characterizations. The universal property gives uniqueness of the factorization and its functoriality. At level 0 a classifier for an ordinary base is a covariant set-valued functor, whose ordinary category of elements is a discrete opfibration. The ordinary comma-component formula is precisely the classical construction [23, 19]. □

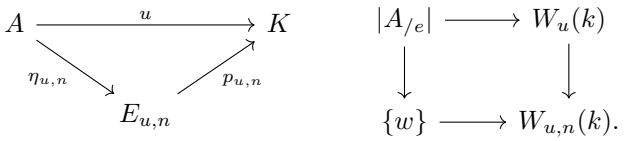


Figure 4.1. The left map inserts a canonical witness. The comma space at a right-hand object is the homotopy pullback on the right.

An empty comma weight at k yields no element of $E_{u,n}$ over k. There is then no condition on the left factor over that missing fiber. If $n = -1$, the right factor keeps only the support of the classifier. If $n = 0$, two points in a connected comma component define the same element on the right. If $n = 1$, the right factor can distinguish automorphisms in that component. The canonical maps $E_{u,n+1} \to E_{u,n}$ over K come from truncation and commute with the left maps in (4.1). These maps compare levels of saturation for the same input.

4.2 Cartesian Kan extension by a comma adjunction

Let A and K have finite products. We first prove a product-preservation statement for left Kan extension of an arbitrary cartesian diagram of spaces, using the comma categories directly. This supplies a coherent counterpart of the cartesian Kan extension used in the ordinary result [4, 18].

Theorem 4.2 (Cartesian extension)**.** *For any functor $u : A \to K$ between small cartesian infinity-categories and any finite-product-preserving $V : A \to \mathcal{S}$, the pointwise* $\mathrm{Lan}_u V$ *preserves finite products. The same holds with $\mathcal{S}$ replaced by an infinity-topos $\mathcal{X}$, or with $\mathcal{S}_{\leq n}$ in place of $\mathcal{S}$. In particular $W_{u,n}$ is cartesian.*

Proof. Put $P_x = A_{/x}$. The projections from $x \times y$ give $\delta : P_{x\times y} \to P_x \times P_y$. There is a right adjoint ρ. On $(a, h : ua \to x)$ and $(b, g : ub \to y)$ its value is

$$\rho((a,h),(b,g)) = \big(a \times b,\ \langle h\, u(\pi_a), g\, u(\pi_b)\rangle\big).$$

The product of A converts a map into $a \times b$ into a pair of maps. The two triangle constraints are exactly the projections of the triangle in K for a comma arrow. At the level of mapping spaces, the equivalences for the two products identify the corresponding homotopy fibers. This proves $\delta \dashv \rho$. The functor u only has to send projections to arrows. The formula allows a noninvertible comparison $u(a \times b) \to ua \times ub$.

The right adjoint ρ is final for colimits. If V_x denotes V composed with the projection $P_x \to A$, then

$$(\mathrm{Lan}_u V)(x \times y) \simeq \operatorname*{colim}_{P_x \times P_y} (V_x \boxtimes V_y) \simeq (\operatorname*{colim}_{P_x} V_x) \times (\operatorname*{colim}_{P_y} V_y).$$

Here $V_x \boxtimes V_y$ has value $V(a) \times V(b)$, by the product preservation of V. The second equivalence uses preservation of colimits in each variable of the cartesian product. This is the product comparison: on each diagram value $V(a \times b)$, its two components are induced by $V(\pi_a)$ and $V(\pi_b)$. At a terminal object of K the comma category is equivalent to A, and A has a terminal object. The Kan extension takes that terminal value to a terminal value in the target. This proves the result for spaces and for an infinity-topos, whose product functors have right adjoints. For truncated spaces compute the extension in spaces and apply $\tau_{\leq n}$. Truncation preserves finite products, since products of n-connected maps are n-connected and products of n-truncated spaces are n-truncated. □

$$P_{x \times y} \underset{\rho}{\overset{\delta}{\rightleftarrows}} P_x \times P_y$$

$$(c, \langle f, g \rangle) \longmapsto ((c, f), (c, g))$$

$$(a \times b, \langle h\, u\pi_a, g\, u\pi_b \rangle) \longmapsto ((a, h), (b, g))$$

Figure 4.2. The comma adjunction separates the two outputs. The final functor ρ computes the product comparison by taking colimits over pairs.

The empty product is essential here. If the domain is empty and the codomain has a terminal object, the left Kan extension of the empty diagram is the constantly empty copresheaf in spaces. A product-preservation statement restricted to nonempty arities would therefore omit a hypothesis used in the cartesian factorization.

Lemma 4.3 (Products in a left fibration)**.** *Let K be cartesian and $p : E \to K$ be classified by $F : K \to \mathcal{S}$. If F preserves finite products, E is cartesian and p preserves and reflects finite product cones. Conversely, if E has finite products and p preserves them, then F preserves them.*

Proof. Given objects (k_i, z_i) of E, take $k = \prod_i k_i$. A point over k projecting to the tuple (z_i), with the comparison paths, exists in a contractible space. The mapping space from (d, t) to this point is the homotopy fiber of $\mathrm{Map}_K(d, k) \to F(k)$ at the chosen point. By the two product equivalences this fiber is the product of the mapping spaces to the factors. This establishes the product universal property, including the empty product. The same argument applies to any cone whose projection is a product.

For the converse, a product in E of the given tuple projects to a product in K, which supplies essential surjectivity for $F(\prod k_i) \to \prod F(k_i)$. To check full faithfulness, fix any target e over $k = \prod k_i$, form in E the product e' of its transported points over k_i, and compare $e \to e'$ over id_k. This comparison exists by the universal property, and is an equivalence since it is in the fiber. Thus e itself supplies a product of its transported points. Compare maps from any point of the fiber to e by that product. The condition of being over id_k is equivalent to lying over the projections on each factor. Maps over the projections are maps in the respective fibers from the transported point. This identifies the mapping spaces in the source and target of the classifier comparison. The empty tuple gives contractibility over the terminal object by the same mapping-space test. □

Corollary 4.4 (Cartesian factorization)**.** *The system of Theorem 4.1 restricts to* Cart_∞*, the infinity-category of small cartesian infinity-categories and product-preserving functors.*

Proof. Apply Theorem 4.2 to $\mathbf{1}_A$, and Lemma 4.3 to the resulting truncated classifier. The projection preserves and reflects finite products. Hence the left factor preserves them when the composite does. For a square of cartesian functors the ambient filler against a right map is cartesian, since its composite with the right map is cartesian and the right map reflects product cones. Thus the contractible space of ambient fillers consists of cartesian fillers. The two classes are closed under retracts in the cartesian category, since they are so in Cat_∞. The factorization argument for orthogonality then applies in this category as well. □

Cartesianity in this result means finite products. Finite limits involve equalizers and pullbacks of maps between product contexts. For example, truncating spaces to components preserves finite products but need not preserve the pullback of two points over a circle. The smooth and topos examples use separate hypotheses to treat equalizers.

4.3 Classification and the ordinary specialization

For $u : A \to K$ and $v : B \to K$, an n-initial zigzag over K is a finite chain of shape maps in L_n pointing in either direction, with the indicated coherent compatibilities with the maps to K. In the cartesian version every shape and map is cartesian. We call the space $\lim_A Fu$ the solution of u in F, as a functor of F.

Theorem 4.5 (Semantic classification)**.** *For two shape maps over K, the following are equivalent.*

1. *There is an n-initial zigzag between them over K.*
2. *$W_{u,n} \simeq W_{v,n}$ as copresheaves on K.*
3. *The two solution functors on $\mathrm{Fun}(K, \mathcal{S}_{\le n})$ are naturally equivalent.*

For cartesian shape maps these conditions are also equivalent to a cartesian n-initial zigzag and to natural equivalence of solutions on the full subcategory $\mathrm{Fun}^\times(K, \mathcal{S}_{\le n})$. Moreover, for any given map $j : A \to B$ over K, its induced comparison on solutions for all of the indicated models is an equivalence if and only if its induced map $W_{u,n} \to W_{v,n}$ is an equivalence.

Proof. An n-initial map of shapes over K gives a comparison of solutions in every model, and Theorem 3.2 makes this an equivalence. Conversely, (3.2) represents the solution functor by $W_{u,n}$ in the copresheaf category. The Yoneda lemma detects equivalence of the representing objects and, for a specified comparison, detects whether the representing map is an equivalence. Straightening transports an equivalence of weights to an equivalence of the total categories in (4.1), over K. The two left factors together with this equivalence form the required zigzag.

In the cartesian case, both weights are objects of $\mathrm{Fun}^{\times}(K, \mathcal{S}_{\leq n})$ by Theorem 4.2. This subcategory is full, so the same pairing is represented there by those weights. Yoneda applies with tests restricted to that subcategory. Corollary 4.4 gives cartesian left factors. The middle equivalence is cartesian because equivalences preserve products. □

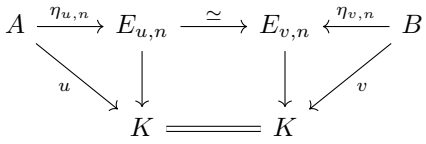


Figure 4.3. A natural equivalence of solution functors can be realized by an explicit zigzag through the two classified total categories.

The factorization also gives a terminal saturation. Fix u. Consider factorizations $A \xrightarrow{i} H \to K$ with $i \in \mathsf{L}_n$, and maps over K commuting with the indicated arrows from A. The object (4.1) is terminal in this infinity-category. Indeed, the space of maps from any such factorization is the space of fillers of a square with i on the left and $p_{u,n}$ on the right. Theorem 4.1 makes it contractible. This remains true in the cartesian category, by Corollary 4.4. The result is a universal property of an ideal saturation. Its total category may be infinite even for a finitely presented input.

For $n = 0$ and ordinary K, the weights are the image copresheaves of [19, Appendix A.3]. One obtains the ordinary statement directly: the unit sends a to $(u(a), [a, \mathrm{id}])$, the projection is the ordinary discrete opfibration, and Proposition 2.1 certifies the two ordinary left legs. In the cartesian case all objects in this construction are ordinary finite-product categories. This recovers the classification in terms of sets and refines it for groupoids and spaces.

The last assertion of Theorem 4.5 is relative to the fixed K. It can certify a map $A \to B$ over K even when the map itself fails to be n-initial as a functor with target B. For example, the inclusion of the terminal vertex into the arrow shape fails the comma test at the initial vertex, but over $K = *$ both shapes have the same weight. The two constant-coefficient solution functors agree. This is the difference between a local rewrite rule that can be used in arbitrary surroundings and a comparison valid in a chosen interpretation family.

4.4 A groupoid obstruction in a product theory

The sphere tests in Section 3 use bare shapes. One can exhibit the same separation of observational levels inside a cartesian theory. This requires some care: a cartesian shape

has a terminal object, so its groupoid completion over the point is contractible. The obstruction below lies in a weight over a nonterminal sort.

For a small infinity-category J, write $\mathrm{FP}(J)$ for its free finite-product completion. One construction takes the opposite of the full subcategory of presheaves on J^{op} spanned by finite coproducts of representables. Its objects can be written as finite tuples $(j_i)_{i\in I}$, including the empty tuple. A map to $(k_l)_{l\in L}$ is a choice of an input i for each output l, together with the corresponding maps in J. Thus

$$\mathrm{Map}_{\mathrm{FP}(J)}\big((j_i)_{i\in I},(k_l)_{l\in L}\big) \simeq \coprod_{\phi:L\to I}\prod_{l\in L}\mathrm{Map}_J(j_{\phi(l)},k_l).$$

The singleton inclusion is denoted by i_J. Concatenation is product. To extend a functor $F : J \to \mathcal{C}$ into a cartesian target, send a tuple to the product of its images, and send a map to the tuple of component maps after projection. The product universal property makes these prescriptions coherent and unique up to contractible choice, and does the same for natural transformations between extensions. Hence restriction gives $\mathrm{Fun}^{\times}(\mathrm{FP}(J),\mathcal{C}) \simeq \mathrm{Fun}(J,\mathcal{C})$.

Proposition 4.6 (The weight after free completion)**.** *Let $u : J \to K$ be any shape map. Restriction of the cartesian weight of $\mathrm{FP}(u) : \mathrm{FP}(J) \to \mathrm{FP}(K)$ along i_K gives $W_{u,n}$. The equivalence is natural in shape maps over K.*

Proof. For a product-preserving $D : \mathrm{FP}(J) \to \mathcal{S}_{\le n}$, the limit is $\mathrm{Nat}(\mathbf{1}, D)$. Both diagrams preserve products, so their mapping space is unchanged by considering it in the full subcategory of product-preserving diagrams. Restriction along i_J is an equivalence on that subcategory and takes the constant terminal diagram to the constant terminal diagram on J. Consequently $\lim_{\mathrm{FP}(J)} D \simeq \lim_J Di_J$, naturally in D. If D is the restriction of a cartesian model on $\mathrm{FP}(K)$, this identifies its solution functor for $\mathrm{FP}(u)$ with the solution functor for u on the singleton models of K. The two weight formulas represent those functors under the restriction equivalence. Yoneda gives the result, including its compatibility with maps of shapes. □

Take $K = [1]$, with arrow $b : 0 \to 1$, and let J_f be the fork $t_1 \leftarrow r \to t_2$ over b. Let J_p be the shape on two parallel arrows $r \rightrightarrows t$ over b. The map $J_f \to J_p$ identifies the outputs and keeps the two arrows. Both shapes have weight $*$ at 0. At 1 the fork has a contractible weight by its comma root. The parallel pair has weight S^1, the nerve of two parallel edges between the comma objects (r, b) and (t, id). Thus at $n = 0$ the weight map is an equivalence. At $n = 1$ the weight at 1 detects a noncontractible loop. Proposition 4.6 gives exactly the same computation on the singleton sorts in $\mathrm{FP}(K)$.

$$\begin{array}{ccc} J_f & \longrightarrow & J_p \\ {\scriptstyle i_{J_f}}\downarrow & & \downarrow{\scriptstyle i_{J_p}} \\ \mathrm{FP}(J_f) & \longrightarrow & \mathrm{FP}(J_p) \end{array} \qquad \begin{array}{ccc} * & \longrightarrow & S^1 \\ \| & & \downarrow{\scriptstyle \pi_0} \\ * & = & * \end{array}$$

Figure 4.4. Free product completion sends the output merge to a cartesian map. Its weight on the output sort still contains the loop before truncation.

In particular the cartesian solution functors become naturally equivalent on all set-valued models of the product theory. On groupoid-valued models, send the input sort to the terminal groupoid and the output sort to BC_2, with b the point inclusion. Extend to all contexts by finite products. In the original fork, each endpoint and its isomorphism to the image of the input is a contractible choice. In the merged-output diagram the two arrows share their endpoint. A first isomorphism determines its identification with the input image, leaving a second isomorphism as a loop at that image. The solutions are therefore the discrete two-element group C_2, rather than a contractible groupoid. The calculation is preserved by adjoining all product contexts, by the proposition. The obstruction is already on one sort and survives the theory's provision for arbitrary finite tuples.

Notice the role of the singleton inclusion. It compares solutions for product-preserving models. The inclusion itself need not be initial in the ambient category of shapes. For two unrelated sorts a, b in a discrete J, there is no map from either singleton to (a, b). Its comma category there is empty. Nevertheless, for a product-preserving model, a cone on the two singleton values determines its leg at the product. The full-subcategory hypothesis in Theorem 4.5 is what makes this selective testing work for the completed shapes.

5 Smooth equations and geometric solution objects

Let Man be the ordinary category of finite-dimensional Hausdorff second-countable smooth manifolds without boundary and smooth maps, with the empty manifold allowed. Products exist. An equation between terms in a context is equality of smooth maps. A constraint on a context asks for an equalizer or a pullback. The initiality test applies to cones in either case, including cases in which the solution presheaf has no representing manifold.

5.1 A noncommutative equation in a smooth product context

Let G be a Lie group with multiplication m, inversion ι, and unit $e : * \to G$. Set $d(x, y) = m(x, \iota(y))$. In the product context G^3, one can derive

$$d(d(x, y), d(z, y)) = d(x, z). \tag{5.1}$$

The input order on the left is important. With no commutativity assumption the left side expands to

$$(xy^{-1})(zy^{-1})^{-1} = (xy^{-1})(yz^{-1}) = x(y^{-1}y)z^{-1} = xz^{-1}.$$

The inverse-of-product equation used here follows from the two unit equations. Indeed $(zy^{-1})(yz^{-1}) = e$ and $(yz^{-1})(zy^{-1}) = e$. Uniqueness of inverse follows by inserting a unit between any proposed inverse and the given inverse. Each displayed equality is obtained by composing an equality in the signature with the appropriate projections and tupling maps.

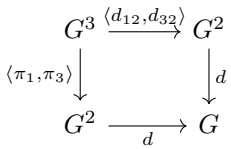


Figure 5.1. The division equation is a commutative square in Man. Here $d_{ij} = d\langle \pi_i, \pi_j \rangle$.

For a concrete model, use the affine Lie group of orientation-preserving maps of the real line, with elements (a, b) for $a > 0$, multiplication $(a, b)(c, d) = (ac, ad + b)$, and inverse $(a^{-1}, -b/a)$. The difference map is

$$d((a, b), (c, d)) = (a/c,\ b - ad/c).$$

Writing the third input as (e_0, f_0), the first coordinate of the outer difference is $(a/c)/(e_0/c) = a/e_0$. The second is

$$b - \frac{ad}{c} - \frac{a}{e_0}\left(f_0 - \frac{e_0 d}{c}\right) = b - \frac{af_0}{e_0}.$$

Thus the smooth computation agrees with the categorical one. This group is nonabelian: the dilation $(2, 0)$ composed with the translation $(1, 1)$ gives $(2, 2)$, whereas the opposite order gives $(2, 1)$. Consequently a proof using exchange of factors would not certify (5.1) in this model.

A mechanical encoding starts with the three projections from the context and introduces one node for each intermediate application of m or ι. The shape can reuse the node for y^{-1}. If that node was duplicated in the syntax, the cartesian fold identifies its two instances because their common input is the same projection. The identity $m(\iota(y), y) = e$ then replaces the middle application by the nullary unit, viewed over the original context. Evaluation at any test manifold M sends the whole construction to a diagram of functions on $\mathrm{Hom}(M, G)$. At $M = G^3$ the tuple of projections recovers the actual arrow identity, giving naturality in probes as well as pointwise evaluation.

5.2 Representable equalizers

Let $f, g : M \to N$ be smooth. If $(f, g) : M \to N \times N$ is transverse to the diagonal, its inverse image of the diagonal is an embedded smooth submanifold E [10]. The inclusion represents the equalizer. In fact, a smooth map $V \to M$ whose image lies in E factors uniquely as a smooth map into E, by the local submanifold charts. This condition gives the universal property on every smooth probe.

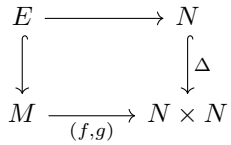


Figure 5.2. A transverse inverse image of the diagonal represents the equation $f = g$ in smooth probes.

For the fixed locus of inversion on S^1, take $f(z) = z$ and $g(z) = z^{-1}$. In angular coordinates, the derivative of $f - g$ in a local chart at a coincidence is 2, so the intersection

with the diagonal is transverse. The solution manifold is $\{1, -1\}$ with the discrete smooth structure. A probe $M \to S^1$ in the equalizer is exactly a locally constant choice of sign on M. If the domain of the probe is disconnected, there is one independent sign per component.

An alternative constraint uses the map $F : \mathbb{R}^2 \to \mathbb{R}^2$, $F(x, y) = (x^2 + y^2 - 1, x)$. Its zero set is $\{(0, 1), (0, -1)\}$. At either point the Jacobian has rows $(0, \pm 2)$, $(1, 0)$, so its rank is 2. Introducing a variable for the radius squared, or grouping the two scalar equations into a tuple, preserves the solution represented by those two points. In contrast, each individual scalar equation has a one-dimensional solution manifold. Taking their union would encode a disjunction and would change the constraint.

5.3 A singular equalizer and its smooth sheaf

Consider $f, g : \mathbb{R}^2 \to \mathbb{R}$ given by $f(x, y) = xy$, $g(x, y) = 0$. Its set of point solutions is the union of the two axes. In this case the functor of smooth solutions is not representable in Man.

Proposition 5.1. *The parallel pair f, g has no equalizer in* Man.

Proof. Suppose $j : E \to \mathbb{R}^2$ is an equalizer. Testing with the one-point manifold identifies the points of E bijectively with the cross. The axis parametrizations $a(t) = (t, 0)$ and $b(t) = (0, t)$ lift to smooth maps into E. Their lifts pass through the same point e over the origin. The derivative $d_e j$ has both $(1, 0)$ and $(0, 1)$ in its image by differentiating the two lifted maps. Thus j is a submersion at e, and its image contains an open neighborhood of the origin by the submersion theorem. Its image is contained in the cross, a contradiction. □

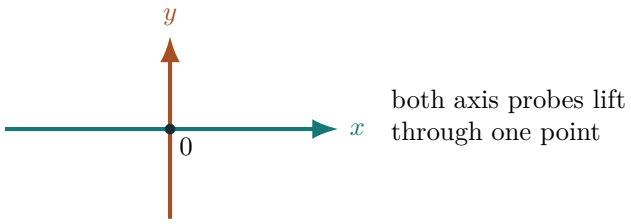


Figure 5.3. Each axis probe contributes an independent tangent direction at the putative lift of the origin.

There is a canonical solution in sheaves. Use the site of open subsets of Euclidean spaces and smooth maps, with open covers. The smooth functor of points $y(M)$ is a sheaf on this site. The restricted Yoneda functor is fully faithful: evaluate a natural transformation on the chart inclusions of M, glue the resulting maps on chart overlaps, and then check on all probes by locality. On any existing limit in Man, the functor y preserves that limit, since each probe computes it by the universal property. The equalizer of $y(f), y(g)$ is the sheaf

$$\mathcal{Q}(V) = \{(a, b) \in C^\infty(V, \mathbb{R})^2 : ab = 0\}. \tag{5.2}$$

The proof of Proposition 5.1 also shows that this sheaf is not representable by a manifold. Its restrictions give more than a topological picture of the point set. For a smooth one-dimensional probe take $a(t) = e^{-1/t^2}$ for $t > 0$ and 0 otherwise, and $b(t) = e^{-1/t^2}$ for $t < 0$ and 0 otherwise. Both are smooth, $ab = 0$, and no neighborhood of 0 maps entirely into one axis. Thus the two axis submanifolds alone do not account for all local plots.

$$\begin{array}{ccc} \mathcal{Q} & \hookrightarrow & y(\mathbb{R}^2) \underset{y(g)}{\overset{y(f)}{\rightrightarrows}} y(\mathbb{R}) \\ \text{probe} & & \\ V & & (a,b) \longmapsto ab \end{array}$$

Figure 5.4. The cross equation has a sheaf equalizer whose sections are smooth probes satisfying the equation.

The same functor on plots is realized by the cross with its subspace diffeology. In that category a plot into the cross is exactly a plot into $\mathbb{R}^2$ whose values lie in the cross. More generally, limits of diffeological spaces equip the set-theoretic limit with the initial diffeology for its projections [2]. The concrete diffeological model and the sheaf model agree on these plot sections. The initiality proofs operate on the probe functor and therefore apply in both settings, even though they differ in their available colimits and ambient structure.

5.4 Relative smooth solutions and loss of submersivity

A constraint may have a smooth total space but variable fiber behavior. In $\mathsf{Man}/\mathbb{R}$, let the parameter be t, let the unknown be x, and consider $x^2 = t$. The solution object is the graph $P = \{(t, x) : t = x^2\}$ with its map to the t-line. The graph is a manifold because the derivative of $x^2 - t$ in the t direction is -1. In the coordinate x on P, the derivative of its projection to the base is $2x$. Hence the projection fails to be a submersion at $x = 0$. Its fibers consist of two points for $t > 0$, one point for $t = 0$, and no points for $t < 0$.

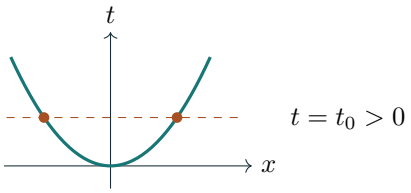


Figure 5.5. The graph is smooth. Its projection to the parameter line has a critical point and a changing number of solutions.

There are smooth branches over $(0, \infty)$, and no smooth section on any open neighborhood of 0. In particular, a point of the fiber at 0 does not extend to a local family. A cone-preserving rewrite over the parameter preserves sections for every probe with its parameter map. If it is followed by applying a smooth pullback, the latter step has its own existence question. Passing to the sheaf solution separates that existence issue from the invariance of the equations.

6 Groupoids and the geometry of witnesses

The ordinary category of small groupoids and functors has strict limits. We use the same groupoids to model 1-truncated spaces, with natural isomorphisms as higher arrows and equivalence of groupoids as equivalence. In this coherent interpretation a cone over a cospan includes an isomorphism relating the two images. It can be meaningful even when the corresponding strictly equal objects do not exist in the given presentation of the groupoids.

6.1 Isomorphism witnesses in a pullback

For functors $F : \mathcal{A} \to \mathcal{B}$ and $H : \mathcal{C}_0 \to \mathcal{B}$ between groupoids, the coherent pullback is modeled by the iso-comma groupoid. Its objects are $(a, c, \gamma : F(a) \xrightarrow{\cong} H(c))$. Its arrows to (a', c', γ') are pairs $(r : a \to a', s : c \to c')$ with $H(s)\gamma = \gamma' F(r)$. For every test groupoid Q, giving a functor into this groupoid is giving two functors and a natural isomorphism between their images, with the analogous description for natural isomorphisms of tests. Thus this model has the required mapping-space pullback property.

$$\begin{array}{ccc} F(a) & \xrightarrow{\gamma} & H(c) \\ {\scriptstyle F(r)}\downarrow & & \downarrow{\scriptstyle H(s)} \\ F(a') & \xrightarrow[\gamma']{} & H(c') \end{array}$$

Figure 6.1. The square defining a morphism of witnesses in the iso-comma groupoid.

If F is an isofibration, the inclusion from the strict pullback is an equivalence of groupoids. For a witness γ, lift it along F from a, producing an isomorphism $r : a \to a'$ with $F(a') = H(c)$ and $F(r) = \gamma$. Then (r, id_c) identifies the original object with the identity-witness object. The inclusion is fully faithful because the square between identity witnesses is exactly equality of the two arrows in $\mathcal{B}$. This gives a sufficient hypothesis for computing the coherent object by a strict one. For the span $* \to BG \leftarrow *$ selecting the unique object of a nontrivial one-object groupoid BG, the strict pullback has one object and the coherent pullback is the discrete set G. The missing datum in the strict computation is the choice of the witness $\gamma \in G$.

6.2 A double-coset calculation

Let $\alpha : H_0 \to G$ and $\beta : L \to G$ be group homomorphisms. For the cospan $BH_0 \to BG \leftarrow BL$, the iso-comma objects are $g \in G$. The arrows are precisely the action of $L \times H_0$ on G by

$$(\ell, h) \cdot g = \beta(\ell) g \alpha(h)^{-1}.$$

Its components are the double cosets $\beta(L)\backslash G / \alpha(H_0)$. The automorphism group at g is the stabilizer of g in this action. If α, β are inclusions, projection to H_0 identifies the stabilizer with $H_0 \cap g^{-1} L g$. For noninjective homomorphisms the pair of lifts must be kept, and that intersection formula can lose kernel automorphisms.

Example 6.1 (Two types of intersection in S_3)**.** Take $G = S_3$ and $L = H_0 = \langle(12)\rangle$, with the inclusion maps. At the identity the double coset has size 2 and the stabilizer is C_2. At $g = (23)$ the conjugate subgroup is $\langle(13)\rangle$, so the intersection is trivial and the orbit has size 4. The two sizes sum to 6. Therefore the coherent pullback is equivalent to $BC_2 \amalg *$. It has two isomorphism classes but only one class has nonidentity automorphisms.

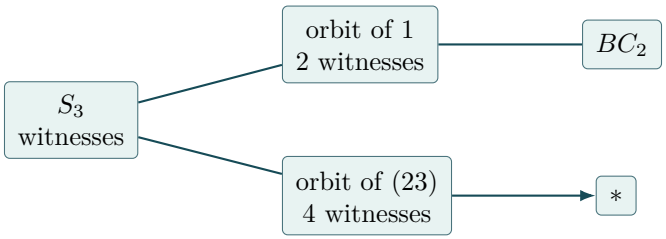


Figure 6.2. Double cosets record components. Each component is equivalent to the classifying groupoid of its stabilizer.

Factoring one leg of the cospan into two functors does not change this answer when the composite is held fixed and the intermediate value is introduced by the rooted path motif. The old arrow shape maps into the new path by the composite. It is initial, and the cospan is a context into which this rule can be pushed out. Concretely, a cone over the factored leg contains an object at the intermediate groupoid and an isomorphism from the image of the original object to that intermediate object. Such pairs form a contractible space of choices once the original object is fixed. Removing them keeps the final witness at BG, including its stabilizer.

6.3 Actions and coherent fixed points

A group G acting on a groupoid $\mathcal{A}$ gives a diagram $BG \to \mathcal{S}_{\leq 1}$. Its limit is the groupoid of equivariant objects with coherent isomorphisms, often called the homotopy fixed points. For a strict action this can be written by choosing $x \in \mathcal{A}$, arrows $a_g : x \to g(x)$, with $a_1 = \mathrm{id}$ and $a_{gh} = g(a_h)a_g$. A morphism $b : x \to x'$ satisfies $g(b)a_g = a'_g b$. These equations are the ordinary groupoid form of a section of the corresponding left fibration over BG.

For a constant action on BH_0 one gets $\mathrm{Map}(BG, BH_0)$. A functor $BG \to BH_0$ is a homomorphism $G \to H_0$, and a natural isomorphism is conjugation by an element of H_0. Taking $G = H_0 = C_2$ gives $BC_2 \amalg BC_2$, one component per homomorphism. Restriction along the collapse $BC_2 \to *$ selects the component of the trivial homomorphism. It is an equivalence on limits for all set-valued models on the point, since BC_2 is connected, and it fails for this 1-truncated model.

Example 6.2 (Inversion action on BC_4)**.** Let the nonidentity element σ of C_2 act on BC_4 by $a \mapsto -a$. Write C_4 additively. An equivariant object is determined by $a = a_\sigma$, and the identity condition for σ^2 is $a - a = 0$. An arrow $b : a \to a'$ has $-b + a = a' + b$, so $a' = a - 2b$. The two orbits are the parity classes. The stabilizer in either is $\{0, 2\}$. Thus the coherent fixed-point groupoid is $BC_2 \amalg BC_2$. The strict fixed subgroupoid has the single object and automorphisms $\{0, 2\}$. The additional coherent component consists of the odd witnesses, which cannot be made strict by a change of trivialization.

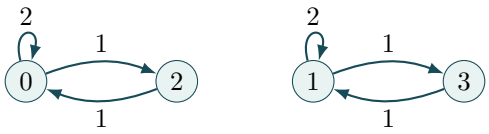


Figure 6.3. For inversion on BC_4, adding a change of trivialization by 1 moves within a parity class. The loop labelled 2 generates its automorphisms.

The calculation distinguishes two operations on witnesses. A cone-preserving move transports the complete groupoid with its morphisms. Passing to components afterwards identifies isomorphic choices. Performing that second operation on the coefficients before forming the limit can give a different result: in this example the coefficient BC_4 has one component and the limit of the component diagram is a singleton, whereas the groupoid limit has two components. The same failure recurs for higher truncation whenever an automorphism is part of the solution data.

7 Sheaves and internal congruence

7.1 A solution sheaf and its sections

Let X be a topological space. A diagram D in the ordinary topos $\mathrm{Sh}(X)$ has a limit computed on every open $V \subseteq X$, since limits of presheaves of sets preserve the sheaf condition [22]. Its solution sheaf takes V to the set of tuples satisfying the diagram on V. If a shape map is ordinary initial, the comparison is a bijection on every open and commutes with restrictions. The inverse extends a local tuple by the comma formula from Proposition 2.1, using restrictions of the same diagram arrows. This explains why a rewrite preserving solution sheaves simultaneously preserves all their global and local sections.

Example 7.1 (Square roots on the circle)**.** Let $X = S^1$ and let $\mathcal{M}(V) = C^\infty(V, S^1)$. The map $s \mapsto s^2$ defines an endomorphism of this sheaf. Let z denote the global section given by the identity of S^1, and set $\mathcal{L}(V) = \{s : s^2 = z|_V\}$. On an arc with a continuous angular coordinate θ, the branches are $\pm e^{i\theta/2}$. On overlaps any two choices differ by a locally constant sign. Thus $\mathcal{L}$ is a torsor for the sheaf $\underline{C}_2$. A global section would be a map $s : S^1 \to S^1$ with $2\deg(s) = 1$, which is impossible. On a full turn the displayed local branch changes sign.

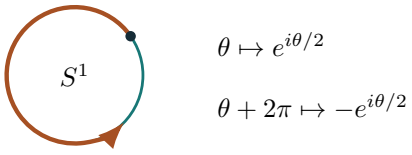


Figure 7.1. The local square-root branch transports to its negative around the circle. The torsor is locally inhabited and has no global section.

There is a useful rewriting test in this example. A variable r for s^2 can be introduced together with its defining operation, yielding a comparison of solution sheaves for $s^2 = z$ and the pair $r = s^2$, $r = z$. Over every open, the inverse to forgetting r sends s to

(s, s^2). It preserves the monodromy of the solution sheaf. Replacing the equation by the assertion that some branch exists on a cover instead takes the image of $\mathcal{L} \to 1$. That image is terminal, and the result has a global section. These are different operations on the same system.

The phenomenon persists for additive equations. On $X = S^1$, let $\omega = x\,dy - y\,dx$ restricted from the plane. The sheaf of primitives $\{h : dh = \omega\}$ is a torsor under locally constant real functions. Local angles provide primitives. A global h would have $\int_{S^1} dh = 0$, whereas $\int_{S^1} \omega = 2\pi$. Thus even when all operations are linear over $\mathbb{R}$ and every local solve uses an affine solution space, there may be a global obstruction. The subobject of all primitives and its support have different semantics.

7.2 Stages, stalks, and exactness

A stalk functor on sheaves of sets is left exact. It is a filtered colimit over neighborhoods, and filtered colimits of sets commute with finite limits. A product of infinitely many constraints need not commute with a stalk. For a concrete example, on $\mathbb{R}$ let $\mathcal{A}$ be the sheaf of smooth real functions. Choose a smooth bump b_i supported in $(3/(4i), 5/(4i))$ with $b_i(1/i) = 1$. The germ of every b_i at 0 is zero. The germ of the tuple $(b_i)_{i\ge1}$ in $(\prod_i \mathcal{A})_0$ is nonzero, since no neighborhood of 0 makes the entire tuple zero. The canonical map

$$(\prod_i \mathcal{A})_0 \longrightarrow \prod_i \mathcal{A}_0$$

is therefore noninjective. The failure arises from the quantifier order: a common neighborhood for all coordinates is stronger than a separate neighborhood for each coordinate.

Global sections preserve all existing limits because they are represented by the terminal sheaf. They need not detect isomorphisms or preserve effective epimorphisms. In Example 7.1, the torsor map $\mathcal{L} \to 1$ is an effective epimorphism while its global section map is $\varnothing \to *$. On a general Grothendieck site, testing a morphism on all stages via representable sheaves after sheafification detects isomorphisms, and the test does not require enough points. More generally, an inverse image of a geometric morphism preserves finite limits and all colimits. That exactness applies to particular constructions below. It cannot be replaced by a test only on global points.

7.3 Constructing the internal congruence

Let $\mathcal{E}$ be a Grothendieck topos and fix a small set of sorts and a finitary many-sorted signature. An algebra A consists of objects A_s and operation maps $m_\sigma : \prod_i A_{s_i} \to A_t$, with any specified equational laws. A seed relation is a family $S_s \hookrightarrow A_s \times A_s$. For instance, if two terms are evaluated on the same parameter object P, their pairing $P \to A_s \times A_s$ gives such a seed by taking its image. This presentation allows a family of equations that has local instances without a list of global elements.

For relations on the same object, use R^{op} for interchange of the coordinates. The composite $R \circ R$ is the image of the pullback of two copies of R over the middle coordinate

under projection to the first and last coordinates. For a family $R = (R_s)_s$, write

$$C_\sigma(R) = \operatorname{im}\left(\prod_i R_{s_i} \longrightarrow A_t \times A_t\right),$$

where the arrow applies m_σ on each side. For nullary operations this is the diagonal pair of the constant. Define a sequence of subobjects by

$$R_s^0 = \Delta_{A_s} \vee S_s, \tag{7.1}$$

$$R_s^{j+1} = R_s^j \vee (R_s^j)^{\mathrm{op}} \vee (R_s^j \circ R_s^j) \vee \bigvee_{\sigma:\mathrm{out}(\sigma)=s} C_\sigma(R^j), \tag{7.2}$$

$$R_s^\omega = \bigvee_{j\geq 0} R_s^j. \tag{7.3}$$

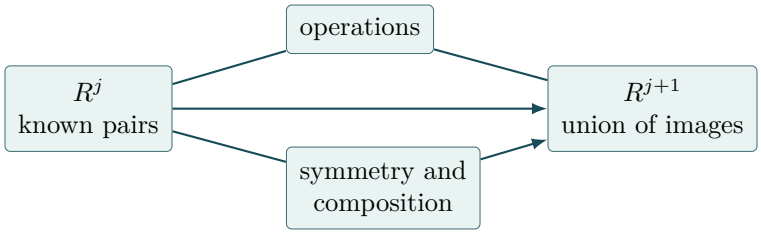


Figure 7.2. One closure stage uses finite products, a pullback for the middle witness, and images of relation-generating maps.

Theorem 7.2 (Internal saturation and geometric base change)**.** *The family R^ω is the least internal congruence containing S. Its effective quotient $q : A \to A/R^\omega$ in $\mathcal{E}$, with operations descended along the quotient maps, is universal among algebra maps that identify all seed pairs. A geometric inverse image carries* (7.1)–(7.3) *and this quotient to the corresponding construction in its target topos.*

Proof. Reflexivity follows from the first stage. Symmetry follows by taking the union of the symmetry steps. For transitivity, the pullback and the image defining composition distribute over joins in each input. Indeed pullback preserves coproducts and images in a topos, and a union is the image of a coproduct of subobjects. A pair of stages i, j can be replaced by their maximum, since the chain is increasing. Thus every summand of $R_s^\omega \circ R_s^\omega$ lies in R_s^ω. The operation test is analogous. Its arity is finite, so a tuple of stages is bounded by their maximum. This proves compatibility. Any congruence containing the seeds contains every stage by induction, proving minimality.

A Grothendieck topos is exact, so an internal equivalence relation is the kernel pair of its effective quotient. The product of the quotient maps of the inputs of an operation is a regular epimorphism with kernel pair the product of the corresponding relations. Compatibility makes the operation on representatives coequalize that kernel pair after composing with the output quotient. Hence it descends uniquely to the quotient. To check an equation in a finite context, precompose both sides with the regular epimorphism from the original context. The two composites agree by the law for A and the fact that the descended maps respect the operations. The universal property among algebra maps

now follows sort by sort from the quotient, with the algebra-map identities checked after precomposing with products of the quotient maps.

A geometric inverse image preserves the products and pullbacks in each stage. It preserves images because it preserves the regular epimorphism and monomorphism of an image factorization. It preserves all joins by the coproduct description of a union and preserves the coequalizers of the relations. Induction on the stages followed by the union proves the asserted compatibility. The descended operations agree by the uniqueness in the quotient construction. □

The construction admits finite derivations as witnesses. At stage 0 a witness is either a diagonal or a seed witness. At stage $j + 1$, a witness includes one of the previous witnesses, its reversal, two witnesses with a common middle value, or a tuple of witnesses for an operation. Products and pullbacks are needed before taking the images so that the shared middle value is the same. A section of the resulting image can be witnessed on a cover. Its internal truth does not require a global choice of witness. The effect of taking the image is to forget distinctions between different derivations of the same pair.

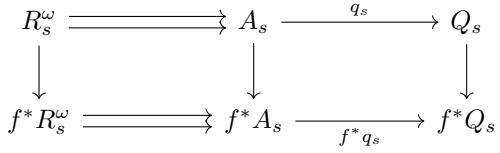


Figure 7.3. The lower row denotes the construction after applying geometric inverse image to the seed data. Its equivalence relation and quotient agree with the images of the upper ones.

The proof uses finitary arities at one exact point. An infinite tuple of witnesses can require unboundedly many stages, so the operation step need not preserve this increasing countable union. It also uses the exactness of the ambient category at the quotient step. For example, the differential-topological argument for the cross in Section 5 already shows that one cannot use the proof for general constraint diagrams in Man.

7.4 A congruence in the topos of group actions

Let $\mathcal{E}$ be the topos of C_2-sets and let A be the free monoid on a, b with the generator of C_2 exchanging the letters. Use the seed subobject consisting of (a, b) and (b, a). The congruence relates two words precisely when they have the same length. Length equality is a congruence containing the seeds. Conversely, the seeds can be inserted at any letter position by the monoid operations, so every word is equivalent to the all-a word of its length. The quotient is $\mathbb{N}$ under addition with trivial action.

The only globally fixed word of A is the empty word: each positive-length word changes in its first letter. Every length class in the quotient is globally fixed. Thus the fixed-point functor does not carry this quotient to a quotient of fixed words. The forgetful functor to Set, which is a geometric inverse image here, does carry it to the ordinary quotient by length. This makes the distinction between geometric inverse image and global sections concrete. It also shows why it is useful to allow seeds as subobjects. The relation is equivariant even though neither of the two displayed nonidentity seed pairs is a fixed pair.

8 Relative weights and descent in an infinity-topos

8.1 A test over a parameter object

Let $\mathcal{X}$ be an infinity-topos and let $B \in \mathcal{X}$. In this section shapes are small external infinity-categories, while the values and the solutions lie in the slice $\mathcal{X}/B$. Write 1_B for its terminal object, $\tau^B_{\leq n}$ for its truncation, and Δ_B for the constant-object functor from spaces to the slice. For $u : A \to K$ put

$$W^B_u = \mathrm{Lan}_u 1_B, \qquad W^B_{u,n} = \tau^B_{\leq n} W^B_u, \qquad W^B_u(k) \simeq \Delta_B |A_{/k}|. \tag{8.1}$$

The constant-object formula follows because Δ_B preserves colimits and carries a point to 1_B. A value of $D : K \to \mathcal{X}/B$ is n-truncated in this slice when its structure map to B is n-truncated. The parameter can be arbitrary, even if the fibers of each diagram value are truncated.

Theorem 8.1 (Relative cone criterion)**.** *The augmentation $W^B_{u,n} \to 1_B$ is an equivalence if and only if*

$$\lim_K D \longrightarrow \lim_A Du$$

is an equivalence for every diagram $D : K \to (\mathcal{X}/B)_{\leq n}$. More generally, equivalence of relative truncated weights for two shapes over K is equivalent to natural equivalence of their solution objects on all such diagrams. If the two shapes are cartesian, product-preserving relative models suffice. The weights and their augmentations commute with change of parameter by pullback and with geometric inverse image.

Proof. The constant diagram at a probe $Q \to B$ satisfies

$$\mathrm{Lan}_u \Delta Q \simeq Q \times_B W^B_u,$$

because product with Q in the slice preserves colimits. Using exponentiation in $\mathcal{X}/B$, the mapping space from Q into the solution of u is

$$\mathrm{Nat}_K(Q \times_B W^B_u, D) \simeq \mathrm{Nat}_K(W^B_u, D^Q) \simeq \mathrm{Nat}_K(W^B_{u,n}, D^Q).$$

The last equivalence holds since exponentiation by Q preserves n-truncated objects. This follows from its adjunction by testing mapping spaces into the exponential with an arbitrary object. These equivalences are natural in Q and D, and compatible with the augmentation. Thus equivalence of the weights implies equivalence of solution objects by Yoneda in the slice. Conversely, take $Q = 1_B$ and vary D, obtaining a natural equivalence of representables on $\mathrm{Fun}(K, (\mathcal{X}/B)_{\leq n})$. Yoneda detects the weights or the augmentation. For cartesian shapes the weights preserve products by the same comma adjunction as in Theorem 4.2, since products in the slice preserve colimits separately and truncation preserves finite products. When D is cartesian, so is D^Q, since exponentiation is a right adjoint. Thus the argument also works on the full subcategory of cartesian diagrams.

A parameter pullback $g^* : \mathcal{X}/B \to \mathcal{X}/B'$ is left exact and preserves colimits, by universality of colimits in an infinity-topos. Geometric inverse image induces a functor of the corresponding slices with the same exactness. Both commute with the pointwise

Kan extension and relative truncation in (8.1). For the truncation claim, n-truncated maps are characterized recursively by their diagonals, and n-connected maps by effective epimorphisms and their iterated diagonals, in the convention of Section 3 [12, Section 6.5.1]. These classes are preserved by a left exact functor preserving colimits. Such a functor therefore carries a connected–truncated factorization to one, which proves the claim. □

$$\begin{array}{ccccc} W_u^B & \longrightarrow & W_{u,n}^B & \longrightarrow & 1_B \\ {\scriptstyle g^*}\downarrow & & \downarrow{\scriptstyle g^*} & & \downarrow{\scriptstyle g^*} \\ W_u^{B'} & \longrightarrow & W_{u,n}^{B'} & \longrightarrow & 1_{B'} \end{array}$$

Figure 8.1. Parameter pullback sends the three displayed objects to their counterparts, up to the canonical equivalences in the slice over B'.

In particular an external n-initial map stays valid for all such sliced models. For a cover of B, it is enough that the augmentation be an equivalence after pullback to the cover, by descent. A general geometric inverse image need only preserve finite limits. The base change of weights, which are built by colimits and truncation, does not imply preservation of arbitrary solution limits. A change of parameter inside $\mathcal{X}$ does preserve all these limits because pullback is right adjoint to postcomposition along g.

The slice also clarifies the relation between support and coherent solutions. In $\mathcal{X}/B$ the support of $P \to B$ is $\tau^B_{\leq -1}P$, an object that records local existence. At level 0 one obtains the relative sheaf of components. Neither one in general determines the coherent solution object. For example in spaces, $* \times_{BC_2} * \simeq C_2$ is discrete with two points, whereas truncating the cospan at level 0 first gives the terminal pullback. This computation gives a concrete failure of left exactness of absolute 0-truncation, consistent with its commuting with base change of *relative* truncations.

8.2 Čech gluing and coherent local choices

Let $U \to B$ be an effective epimorphism. Its Čech nerve $U_\bullet$ has $U_j = U^{\times_B (j+1)}$, with faces that forget one entry and degeneracies that repeat it. The geometric realization of $U_\bullet$ is B. For any $P \to B$ this gives

$$\operatorname{Map}_{\mathcal{X}/B}(B, P) \simeq \operatorname{Tot}_{[j]\in\Delta} \operatorname{Map}_{\mathcal{X}/B}(U_j, P). \tag{8.2}$$

Indeed the Čech colimit is also a colimit in the slice, and mapping into P turns it into the indicated limit of spaces. This form of descent follows directly from effectiveness. The stronger equivalence $\mathcal{X}/B \simeq \lim_{[j]\in\Delta} \mathcal{X}/U_j$ follows from descent for slices of an infinity-topos [12, Section 6.1]. In particular a comparison in the slice can be tested after pullback to the cover. For relative diagram solutions, apply (8.2) with P the solution object. Equivalently, form the local solution objects first and use that change of parameter preserves limits.

$$U \times_B U \times_B U \;\rightrightarrows\; U \times_B U \;\rightrightarrows\; U \longrightarrow B$$

Figure 8.2. The initial stages of a Čech nerve. A solution on U has a comparison on $U \times_B U$ and a compatibility on the triple pullback.

For groupoid-valued local solutions on an open cover, choose objects x_i and isomorphisms $g_{ij} : x_j|_{U_{ij}} \to x_i|_{U_{ij}}$. On triple overlaps the condition is $g_{ij}g_{jk} = g_{ik}$. Isomorphisms between such data consist of families of local isomorphisms respecting the overlap maps. For space-valued solutions, the equation on triple overlaps is replaced by a compatible 2-simplex, with compatibilities on higher overlaps. A local rewrite natural over the Čech nerve induces an equivalence on the totalizations. Unrelated equivalences selected at each open give no such map until they have been made compatible.

8.3 A sheaf of witnesses with nontrivial transport

Let $B = S^1$ and $H = S_3$. In sheaves of spaces on B, let $\mathcal{B}H$ be the sheaf whose value on an open is the groupoid of torsors for the locally constant group H on that open. Gluing torsors and their isomorphisms gives the sheaf condition in groupoids. Choose a global torsor P with monodromy $g = (12)$. Concretely, form the quotient of $\mathbb{R} \times H$ by the action generated by $(t, h) \mapsto (t + 1, gh)$, with its right H-action. The torsor map to $\mathbb{R}/\mathbb{Z}$ has local sections. Over any connected trivializing arc its automorphisms are identified with H by a local trivialization.

A loop at P in the sheaf $\mathcal{B}H$ is precisely an automorphism of P in that sheaf. Write $\mathcal{A}_P = * \times_{\mathcal{B}H} *$, with both maps selecting P. Then $\mathcal{A}_P$ is the locally constant sheaf of groups with fiber H and monodromy $a \mapsto gag^{-1}$. Indeed an H-equivariant fiber map of a trivial right torsor is left multiplication by some a. Transporting the trivialization by the left multiplication g conjugates the fiber map by g. Thus the stalk has six elements, and the global section group is the fixed subgroup

$$\Gamma(B, \mathcal{A}_P) = C_{S_3}((12)) = \{1, (12)\}.$$

The last equality holds because an element commuting with (12) must preserve its two-element support, and there are only two such permutations in S_3.

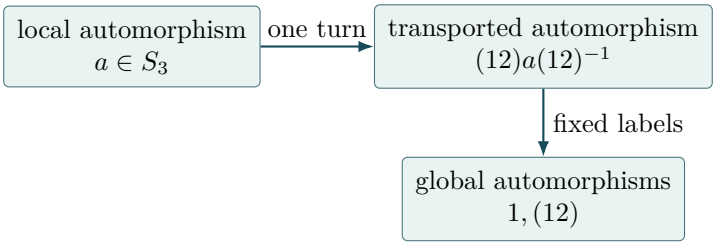


Figure 8.3. A local witness can vary under transport around the parameter circle. Global witnesses are the centralizer of the chosen monodromy.

Now interpret the arrow theory [1] by the selected point $* \to \mathcal{B}H$. In the fork with separate targets there are two independent endpoint-and-path choices. The solution sheaf is terminal by the rooted rule, including on every restriction of P. For the parallel pair with a single endpoint, eliminate the first path by that same rule. The remaining

path is a loop at P, so the solution sheaf is $\mathcal{A}_P$. This is the groupoid obstruction of Section 4.4 with a parameter and a nontrivial transition. Treating each local solution independently would yield a six-element answer on a trivializing arc. Descent retains the conjugation comparison, and the global answer has two elements. Applying 0-truncation to the *coefficient* $\mathcal{B}H$ before either limit would instead give the terminal sheaf, since all torsors are locally isomorphic. The example separates these three operations by their exact solution counts.

8.4 A locally solvable lift that does not glue

Take the topos of sheaves of spaces on $\mathbb{RP}^2$, and the central extension

$$0 \longrightarrow C_2 \longrightarrow C_4 \xrightarrow{\text{mod } 2} C_2 \longrightarrow 0.$$

Let $P = S^2 \to \mathbb{RP}^2$ be the nontrivial C_2-torsor. Consider the solution groupoid of C_4-torsors together with an identification of their pushforward C_2-torsor with P. The isomorphism in this definition is part of the constraint. This is the fiber over P of the map from the groupoid of C_4-torsors to that of C_2-torsors. Over contractible opens there are solutions, because P is trivial there.

A global solution would have monodromy a homomorphism $\pi_1(\mathbb{RP}^2) = C_2 \to C_4$ whose reduction is the identity. The image of the generator must have order dividing 2, so it is either 0 or 2, both of which reduce to 0. This proves there is no global lift. The same obstruction can be read from transition functions. Let $c_{ij} \in C_2$ be a cocycle for P on a good cover and let $\tilde{c}_{ij} \in C_4$ be chosen lifts. Then

$$o_{ijk} = \tilde{c}_{ij} + \tilde{c}_{jk} - \tilde{c}_{ik}$$

takes values in the kernel. It is a Čech 2-cocycle because it is a coboundary in the C_4-valued complex. Changing the lifts by a kernel-valued 1-cochain changes o by its coboundary. Vanishing of this class would allow the lifts to satisfy the cocycle equation and hence define a torsor. The monodromy computation proves the class is nonzero. This example locates the obstruction at the overlap compatibility, rather than in a failure to solve the equation locally.

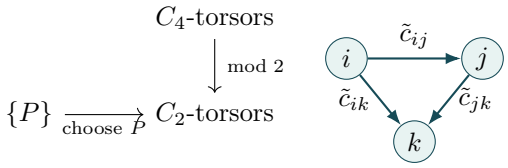


Figure 8.4. The global lift is a homotopy fiber. The discrepancy of the two paths on the triple overlap belongs to the central kernel.

This obstruction and the square-root torsor exhibit distinct levels of gluing. The first-order sheaf problem can fail because two local choices disagree on an overlap. For the lift problem one can match images on overlaps, yet a proposed family of lifts can fail the compatibility on triple overlaps. The relative weight criterion controls changes of diagram within either problem while retaining its designated level of solution data.

9 Monads, substitution, and higher arities

9.1 Free syntax in a topos

A monad distinguishes the operations on terms from the equations imposed on them. Let Σ be the finitary signature of Section 7. For a family $X = (X_s)_s$ of objects of a Grothendieck topos $\mathcal{E}$, consider the polynomial

$$(P_X Y)_s = X_s \amalg \coprod_{\sigma:(s_1,\dots,s_r)\to s} \prod_{i=1}^{r} Y_{s_i}.$$

The initial algebra is the free Σ-term algebra $T_\Sigma X$, with variables in X. To construct it, start from the initial family, apply P_X repeatedly along the initial chain, and take its colimit. Filtered colimits commute with finite products in a Grothendieck topos, so P_X preserves this colimit. The tail of the chain gives the algebra structure on the colimit. Given a P_X-algebra, the compatible maps from the finite stages are defined recursively by its structure map. This yields existence and uniqueness of a map of algebras from the colimit. Equivalently, evaluating a term with a variable assignment is induction on its finite tree.

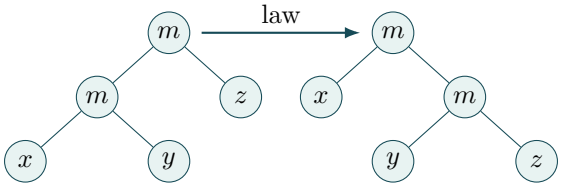


Figure 9.1. A term is a rooted operation tree. The associativity seed is evaluated for every triple of terms, not only for named variables.

Fix a small family E of equations in finite contexts. Evaluate both sides of each equation in $T_\Sigma X$ on the product of the term objects in its context. Take the image of this pair and form the union of such images in the output sort. Use that family as seeds in Theorem 7.2. Denote the quotient by $T_{\Sigma,E}X$.

Theorem 9.1 (Presented term monads)**.** *The quotient $T_{\Sigma,E}X$ is free on X among Σ-algebras satisfying E. The endofunctor $T_{\Sigma,E}$ is a monad, with unit insertion of variables and multiplication simultaneous substitution. For geometric inverse image $f^* : \mathcal{E} \to \mathcal{E}'$, there is a natural isomorphism*

$$f^* T_{\Sigma,E} \cong T'_{\Sigma,E} f^*$$

commuting with both monad structures. The statement includes a finite-product theory specified by chosen finitary generators and equations.

Proof. The quotient operation maps descend by Theorem 7.2. For each equation, the two maps from the corresponding quotient context become equal after precomposition with the product of the quotient maps of the input sorts. These product maps are regular epimorphisms, so the equations hold in the quotient. Given a variable assignment to

an algebra satisfying E, the evaluation from $T_\Sigma X$ identifies all seed pairs. Theorem 7.2 makes it factor uniquely through the quotient. This gives the free-forgetful adjunction.

The unit is the universal assignment. The multiplication is the unique extension of the identity assignment from $T_{\Sigma,E}X$ into itself. That extension evaluates a term whose variables are terms. The unit and associativity laws follow because the two maps in each law agree on the free generators of their common source. Geometric inverse image carries the initial chain and its colimit to those of the target because the chain uses finite products and small coproducts. It carries the equations' seed images and congruence quotients to the new ones by Theorem 7.2. The resulting isomorphism is natural and respects insertion and substitution by their descriptions on the generators. □

In particular, the algebraic theory can be presented as the finite-product category generated by the sorts and operations modulo the displayed equations [9]. Its models are the finite-product-preserving functors into the ambient category. The finitary monad here is determined by the same presentation in the specified topos. A different category of arities may give a different relationship between theories and monads. For a general treatment using a dense category of arities, see [5]. The construction above uses finite contexts to make the exactness of each stage explicit.

For a one-sorted example, the free monoid monad has

$$T_{\mathrm{list}}X = \coprod_{r\geq 0} X^r.$$

The evaluation of a binary term is its word of leaves, with the empty word for the unit. Two terms with the same word have a common representative in the quotient, obtained by removing units and reassociating to a fixed bracketing. Conversely, the word is invariant under those two laws. Equivalently, the word object represents maps from X to an internal monoid by concatenation and evaluation. Adding commutativity gives

$$T_{\mathrm{mult}}X = \coprod_{r\geq 0} X^r/\Sigma_r, \tag{9.1}$$

where the quotient denotes orbits in the ordinary topos. A commutative product of r factors is invariant under every adjacent transposition, which generates the symmetric group. Thus evaluating words in any commutative monoid descends to (9.1). Concatenation followed by the quotient gives the required multiplication on this free monoid. This argument is internal and uses no sorting order on X.

$$\begin{array}{ccc} T_\Sigma T_\Sigma X & \xrightarrow{\mu^\Sigma_X} & T_\Sigma X \\ \downarrow & & \downarrow{\scriptstyle q_X} \\ T_{\Sigma,E}T_{\Sigma,E}X & \xrightarrow[\mu^E_X]{} & T_{\Sigma,E}X \end{array}$$

Figure 9.2. The vertical map on the left applies the term quotient at both levels. The square expresses that simultaneous substitution respects the congruence.

9.2 Solution diagrams of algebras

Let T be an ordinary monad on $\mathcal{C}$. The forgetful functor $\mathrm{Alg}_T(\mathcal{C}) \to \mathcal{C}$ creates limits. For a diagram of algebras whose underlying limit is L, each projection $p_i : L \to A_i$ gives $TL \xrightarrow{Tp_i} TA_i \to A_i$. These maps are a cone, so they define $TL \to L$. The unit and multiplication identities are checked after all the projections, which jointly detect maps into L. The same argument proves that the limiting cone has the universal property among algebras. No preservation of the limit by T is used.

For a coherent monad, understood as an associative algebra in the monoidal infinity-category of endofunctors under composition, the Eilenberg–Moore forgetful functor likewise creates limits that exist in the underlying category and reflects equivalences [21, 13]. Here the action of the monad and its coherence are all part of the algebra structure. One way to see the limiting construction when the underlying category admits the limits of the indicated shape is to apply the limit functor to the diagram of algebras with the pointwise monad. The limit has the canonical comparison $T \lim D \to \lim TD$, obtained by applying T to the limiting cone. Those comparisons are compatible with the unit and multiplication by the same universal property, giving the induced algebra on the limit. This is the coherent form of lifting the cones on $T^j L$ for all j. The monadic limit theorem provides the universal mapping-space statement.

It follows that a rewrite in L_n induces an equivalence of the algebra limits if the underlying diagrams admit those limits and the relevant mapping spaces in the underlying category are n-truncated. For ordinary algebras the latter condition is automatic with $n = 0$. For coherent algebras the conclusion also holds at $n = \infty$ with no truncation hypothesis. To obtain a finer level one can test the underlying cone spaces against all probes and verify the local condition in Theorem 3.2. A class of algebras need not contain all detector models used by that theorem. A comparison that is invertible for one fixed theory may therefore impose fewer restrictions than universal n-initiality.

9.3 A 2-monad and its pentagon

There is a 2-monad M on the 2-category of small categories, functors, and natural transformations:

$$M\mathcal{C} = \coprod_{r \geq 0} \mathcal{C}^r.$$

Its unit takes a one-letter word, and its multiplication flattens a word of words. These operations are 2-natural. A strict algebra gives all arity products with equality under flattening and insertion. Choosing the empty and binary operations gives a strict monoidal category. A pseudoalgebra instead supplies coherent natural isomorphisms for the two laws. In the unbiased description it has an operation for each arity, and these comparison isomorphisms relate grouping in stages to the full arity operation. It determines a monoidal category with the usual associator and unitors, and conversely a monoidal category supplies such a pseudoalgebra after choosing bracketings [11].

For four objects the associator provides the two routes in the pentagon in Figure 9.3. The pseudoalgebra compatibility requires the equality of those composites. In a setting with invertible 3-cells one can instead supply a comparison between those composites

as part of the model, with further compatibility conditions. The pentagon illustrates why a collection of pairwise rewrite isomorphisms does not itself specify their composite comparisons.

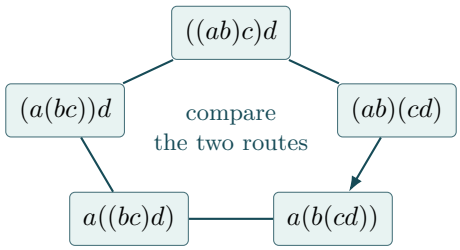


Figure 9.3. The pentagon compares two sequences of reassociation from the fully left-associated word to the fully right-associated word.

There is a test with small shapes for this distinction. The face poset of the boundary of a pentagon has groupoid completion S^1. Adding the 2-face, with all its incidences, gives a face poset with terminal object and contractible completion. For constant diagrams these two shapes are indistinguishable by set models, but the map from the boundary shape to the filled shape fails the 1-test with the coefficient BC_2. A functor from the circle to BC_2 with nontrivial monodromy does not extend over the filled cell. This topological test isolates the role of the coherence cell. In an actual monoidal pseudoalgebra the pentagon is a law, and its associators satisfy that law.

9.4 Analytic monads and isotropy of operations

In spaces, a polynomial functor on the slice over a space of colors I can be described by

$$I \xleftarrow{s} E \xrightarrow{p} B \xrightarrow{t} I, \qquad T = \Sigma_t \Pi_p s^*.$$

Here B parametrizes operations, E parametrizes slots of those operations, and s, t give input and output colors. The operations may have automorphisms. The analytic case has finite discrete slot fibers. An analytic monad has an analytic underlying functor and cartesian unit and multiplication, where a cartesian natural transformation has pullback naturality squares. In this terminology cartesian refers to pullbacks, as distinct from finite-product preservation of shapes. The theory of analytic monads on spaces identifies them, with varying color spaces, with infinity-operads modeled by dendroidal Segal spaces [8, 7]. The polynomial describes the arities and their symmetries. The monad structure supplies substitution and all its coherent laws.

$$\begin{array}{ccccccc} I & \xleftarrow{\ s\ } & E & \xrightarrow{\ p\ } & B & \xrightarrow{\ t\ } & I \\ \text{input colors} & & \text{slots} & & \text{operations} & & \text{output colors} \end{array}$$

Figure 9.4. The three maps determine pullback to the slots, dependent product over the slots of each operation, and dependent sum over the operations of each output color.

For an explicit comparison use the free associative and free commutative monads on spaces. Their underlying functors have formulas

$$T_{\mathrm{ass}}(X) = \coprod_{r \geq 0} X^r, \qquad T_{\mathrm{comm}}(X) = \coprod_{r \geq 0} (X^r)_{h\Sigma_r}. \tag{9.2}$$

For the associative monad the ordered positions of a word are distinguished. For the commutative monad the indexing groupoid is the groupoid of finite sets and bijections, with the functor sending a set S to the space of S-indexed tuples. Its colimit is the right side of (9.2). Disjoint union gives substitution of tuples, including the coherent permutation data. These are the free-algebra formulas for the associative and commutative infinity-operads [13, 8]. In arity 0 the unique tuple is the monad's empty word or empty multiset.

At $X = *$, the r-ary term of the associative monad is a point, while the r-ary part of the commutative monad is $B\Sigma_r$. In arity 2 this is BC_2. If $X = \{a, b\}$ is discrete, the arity 3 summand on the right has four components, represented by aaa, aab, abb, bbb. Their automorphism groups are $\Sigma_3, \Sigma_2, \Sigma_2, \Sigma_3$, respectively. Indeed the homotopy quotient is the action groupoid, and a permutation fixes a tuple exactly when it preserves the fiber of each label. The formula for any discrete label set X is a disjoint union over finite multisets, with stabilizer $\prod_{x \in X} \Sigma_{m_x}$ for multiplicities m_x. Only finitely many factors are nontrivial.

$$\begin{array}{ccc} (X^r)_{h\Sigma_r} & \longrightarrow & B\Sigma_r \\ {\scriptstyle \tau_{\leq 0}}\downarrow & & \downarrow{\scriptstyle \tau_{\leq 0}} \\ (\pi_0 X)^r/\Sigma_r & \longrightarrow & * \end{array}$$

Figure 9.5. The upper map forgets labels but keeps slot symmetries. The lower map is the component-level orbit computation.

Taking π_0 in (9.2) gives (9.1) on $\pi_0 X$, since π_0 commutes with colimits and finite products. The comparison commutes with multiplication: disjoint union of labelled finite sets induces addition of the multiplicities, which is the substitution on multisets. Thus the ordinary commutative term monad appears as the component shadow of the analytic example. The level-0 equation identifies permuted tuples, while the higher object still admits the stabilizers just calculated. This is a concrete example of the information retained when the weights and the models of an equation calculus are allowed to have positive truncation level.

10 Relations as chosen witnesses

The internal congruence of Theorem 7.2 asks for two elements to have the same image under a quotient. In a category of homotopy-coherent algebras, a relation can ask instead for a chosen path between the images. The distinction is detectable by the mapping property of the quotient. This section makes that property explicit and compares it to strict quotients and to linear equation solving.

10.1 The universal algebra with a specified path

Let T be a coherent monad on an infinity-category $\mathcal{C}$ with binary coproducts. Suppose its free and forgetful adjunction is $F : \mathcal{C} \rightleftarrows \mathrm{Alg}_T(\mathcal{C}) : U$. Let $r, s : P \to UA$ be maps into the underlying object of an algebra A. They correspond by adjunction to $\widehat{r}, \widehat{s} : FP \to A$.

If the displayed pushout exists in $\mathrm{Alg}_T(\mathcal{C})$, define $A\langle r \simeq s\rangle$ by

$$\begin{array}{ccc} F(P \amalg P) & \xrightarrow{\widehat{[r,s]}} & A \\ {\scriptstyle F(\nabla)}\downarrow & & \downarrow \\ FP & \longrightarrow & A\langle r \simeq s\rangle. \end{array} \tag{10.1}$$

Here ∇ is the codiagonal, and $F(P \amalg P)$ is a coproduct of two copies of FP by the free adjunction. The hypotheses hold, for example, for free associative or commutative monoids in spaces, whose categories of algebras are presentable [13, 8].

Proposition 10.1 (Mapping property of a coherent relation)**.** *Mapping $A\langle r \simeq s\rangle$ into an algebra B classifies an algebra map $f : A \to B$ equipped with a path*

$$Uf \circ r \simeq Uf \circ s \quad \textit{in} \quad \mathrm{Map}_{\mathcal{C}}(P, UB).$$

More precisely its mapping space is naturally equivalent to the homotopy pullback of the endpoint comparison $(- \circ r, - \circ s)$ against the diagonal on $\mathrm{Map}_{\mathcal{C}}(P, UB)$.

Proof. Apply $\mathrm{Map}_{\mathrm{Alg}_T(\mathcal{C})}(-, B)$ to (10.1). It takes the pushout to a homotopy pullback. By adjunction the two vertices coming from FP and $F(P \amalg P)$ are $M = \mathrm{Map}_{\mathcal{C}}(P, UB)$ and $M \times M$, respectively. The induced map $M \to M \times M$ is the diagonal. Replace it by the equivalent path-space map $M^{[0,1]} \to M \times M$ with its two endpoints, using that $[0, 1]$ is contractible. The pullback then has exactly the stated data. □

$$\begin{array}{ccc} \mathrm{Map}(A\langle r \simeq s\rangle, B) & \longrightarrow & \mathrm{Map}(A, B) \\ \downarrow & & \downarrow{\scriptstyle \text{two composites}} \\ \mathrm{Map}(P, UB)^{[0,1]} & \xrightarrow[\text{endpoints}]{} & \mathrm{Map}(P, UB) \times \mathrm{Map}(P, UB) \end{array}$$

Figure 10.1. Mapping the relation attachment into B produces the space of witnesses for that relation. The left vertical projection takes a model to its chosen path.

When all relevant mapping spaces are discrete, there is a unique path if the two endpoints agree. The mapping property becomes the ordinary coequalizer property. In a Grothendieck topos, for a finitary algebra and seed given by r, s, the ordinary version of (10.1) therefore agrees with the quotient of Theorem 7.2, by the universal property on algebra maps. The map $P \to UA \times UA$ can be replaced by its image in this ordinary comparison. In the coherent comparison, the space of choices of paths can depend on P, and the image in components generally has a different universal property.

As the smallest case, take the identity monad on spaces and $A = P = *$ with $r = s = \mathrm{id}$. In the ordinary version there is no change to A. In spaces, the homotopy pushout is S^1. For example, a homotopy pushout model has two vertices and two edges connecting them, one edge for each point of $P \amalg P$. Its mapping space into B consists of a point of B and a loop at it, in agreement with Proposition 10.1. Adding the same relation twice independently gives two chosen loops and is represented by $S^1 \vee S^1$. A comparison between the two choices amounts to another cell. Attaching a disk to a chosen loop in

the one-relation example makes that loop nullhomotopic and yields a contractible space. These three results have the same set of components. They differ in what a model must supply.

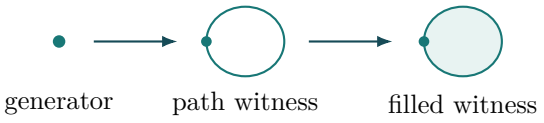


Figure 10.2. For a generator equated to itself, choosing a witness attaches a loop. Choosing a contraction of that witness fills the loop.

This construction changes the theory by adjoining a witness. It is separate from Theorem 3.2, which compares indexing shapes while interpreting them in a fixed theory. One can use both constructions in sequence. A relation attachment defines what a model is required to supply, and a cone comparison decides whether a change in the shape of a subsequent query preserves those models' solution spaces.

10.2 A free commutative monoid with one redundant equation

Let F_{comm} be the free commutative monoid functor in spaces from (9.2). Take $A = F_{\mathrm{comm}}(*)$, $P = *$, and let both r, s select its generator. Each adjoint map $F_{\mathrm{comm}}(P) \to A$ is the identity. Since F_{comm} preserves coproducts and pushouts,

$$A\langle r \simeq s\rangle \simeq F_{\mathrm{comm}}(S^1).$$

The collapse $S^1 \to *$ induces the algebra map that sends the chosen path to the constant path. On components both algebras have the ordinary commutative monoid $\mathbb{N}$, and the induced map is its identity. On the length-one summand the map is $S^1 \to *$, which is not an equivalence. The map preserves the grading by length, so the difference cannot be accounted for by other summands. At length two the source is $((S^1)^2)_{hC_2}$ and the target is BC_2, with the map that forgets the two labels. The homotopy fiber of that map at the base point is $(S^1)^2$, by the associated family for the action. Thus a witness on the generator changes more than the generator's own component.

The set-level comparison also follows directly from adjunction. A commutative monoid in spaces can be defined as a special functor from finite pointed sets, whose value at $\langle r\rangle$ is the r-fold product of the value at $\langle 1\rangle$. Applying π_0 levelwise preserves the required product equivalences. For any ordinary commutative monoid viewed as such a functor with discrete values, the adjunction for pointwise truncation computes maps into it using the ordinary monoid of components. Thus taking components is left adjoint to this inclusion and preserves the pushout used to impose the relation. In this example the ordinary equation on the generator is already satisfied, so the ordinary monoid is unchanged. The nontrivial length-one space in the higher version records the extra witness supplied by the presentation.

A further base-change statement requires specifying what acts on algebras. For the identity monad in an infinity-topos, geometric inverse image preserves (10.1) because the construction is a colimit. For a general monad, if the inverse image lifts to a functor of algebra categories that carries the free adjunction on objects to the corresponding

free adjunction and preserves this pushout, it carries the relation attachment to the new attachment. These are sufficient hypotheses by the displayed diagram. Commutation with a particular monad on underlying objects alone leaves the algebra-colimit step to be justified.

11 Additive and stable examples

11.1 Eliminating a variable in an additive category

An additive category provides finite biproducts, and hence both tuples of inputs and sums of parallel maps. This structure permits a different family of equations from those of a general product category. Let

$$M = \begin{pmatrix} a & b \\ c & d \end{pmatrix} : X \oplus Y \longrightarrow P \oplus Q$$

be a block map, where $d : Y \to Q$ is an isomorphism. Set $s = a - bd^{-1}c : X \to P$. A solution on a probe V is a pair $x : V \to X,\ y : V \to Y$ for which $ax + by = 0$ and $cx + dy = 0$. The second equation determines $y = -d^{-1}cx$. Substituting into the first leaves $sx = 0$, and the substitution can be inverted by the projection to x. This gives an isomorphism of solution functors and hence of kernels whenever one of them exists.

There is also a diagram proof valid in every additive category. Define the two triangular isomorphisms

$$R = \begin{pmatrix} 1 & 0 \\ -d^{-1}c & 1 \end{pmatrix}, \qquad L = \begin{pmatrix} 1 & -bd^{-1} \\ 0 & 1 \end{pmatrix}.$$

Inverting either amounts to reversing the sign of the off-diagonal entry. Multiplication gives $LMR = s \oplus d$. A kernel is a pullback against zero, so pre- and postcomposing with these isomorphisms preserves the solution functor. The kernel of $s \oplus d$ is the kernel of s paired with the zero solution of d. The projection to X gives the same comparison as the probe calculation.

$$\begin{array}{ccc} X \oplus Y & \xrightarrow[\cong]{R} & X \oplus Y \\ {\scriptstyle s\oplus d}\downarrow & & \downarrow{\scriptstyle M} \\ P \oplus Q & \xrightarrow[L^{-1}]{\cong} & P \oplus Q \end{array}$$

Figure 11.1. Block elimination is a change of variables and equations by isomorphisms. The reduced constraint is $sx = 0$.

For a sheaf example, let $X = Y = P = Q$ be the sheaf of smooth functions on the real line. Choose $a = c = 1$, $b = t$, $d = 1$ acting by multiplication, with t the coordinate. The two equations $x + ty = 0,\ x + y = 0$ reduce to $(1 - t)x = 0$. The solution sheaf is zero: a smooth function annihilated by $1 - t$ vanishes off 1, and continuity makes it vanish at 1 as well. Solving the pointwise system only at $t = 1$ would instead give a one-dimensional space. Such a pointwise specialization is tensoring with the residue field, which need not preserve kernels. It is different from taking a stalk, whose solutions are germs on

a common neighborhood. The comparison is a property of the functor used to change coefficients, rather than of the matrix identity.

When d is no longer invertible, a local formula can still hold on the open set where d is invertible. For example, $ty = x$ gives unique elimination of y away from the origin. At the origin the image of the projected map on germs is the ideal of germs divisible by t. The equality of that ideal with germs vanishing at 0 follows by writing $x(t) = t \int_0^1 x'(st)\, ds$ for a representative vanishing at 0. Consequently the domain of this partial solve is proper. Its domain and its smooth extension behavior have to be included in a global certificate.

11.2 Derived fibers of the reduced equation

In a stable infinity-category the biproducts and mapping spaces support coherent analogues of the block maps [13, Chapter 1]. An invertible block again gives equivalent fibers. To see this, the biproduct descriptions turn the formula for LMR into a calculation of its four component maps up to coherent homotopy. Both triangular maps are equivalences because their displayed inverses compose to the identity in the homotopy category, which detects equivalences. The two fibers agree under these equivalences, and $\mathrm{fib}(s \oplus d) \simeq \mathrm{fib}(s) \oplus \mathrm{fib}(d) \simeq \mathrm{fib}(s)$. This retains the homotopy in the equation with zero. For any V, $\mathrm{Map}(V, \mathrm{fib}(s))$ is the homotopy fiber of $\mathrm{Map}(V, X) \to \mathrm{Map}(V, P)$ over the zero map.

For example, take the derived category of abelian groups as a stable infinity-category, with cochain complexes. The ordinary kernel of $2 : \mathbb{Z} \to \mathbb{Z}$ is zero. The fiber of $2 : \mathbb{Z}[0] \to \mathbb{Z}[0]$ in the derived category is represented by the complex $\mathbb{Z} \xrightarrow{2} \mathbb{Z}$ in cohomological degrees 0 and 1. It has zero cohomology in degree 0 and cohomology $\mathbb{Z}/2$ in degree 1. This also follows from the exact sequence of cohomology for the fiber triangle. Testing only maps from the unshifted $\mathbb{Z}[0]$ misses this group in degree 1, while testing maps from $\mathbb{Z}[-1]$ detects it. The representable cone functor of the stable target therefore distinguishes the fiber from zero.

$$\begin{array}{ccc} \mathrm{fib}(2) & \longrightarrow & \mathbb{Z}[0] \\ \downarrow & & \downarrow{\scriptstyle 2} \\ 0 & \longrightarrow & \mathbb{Z}[0] \end{array} \qquad 0 \longrightarrow \mathbb{Z} \xrightarrow{2} \mathbb{Z} \longrightarrow \mathbb{Z}/2 \longrightarrow 0.$$

Figure 11.2. The stable pullback gives the derived fiber. The exact sequence identifies its surviving cohomology in degree 1.

The additive calculation shows that one can use more than product congruence when a target has linear structure. It also illustrates why the level of the detector has to match the equations. Over ordinary modules a kernel records exact solutions. Over complexes a fiber records a map together with its nullhomotopy in the chosen model. A linear elimination by equivalences is sound in both settings, while a comparison from ordinary kernels to derived fibers requires additional exactness.

12 Changing the index and imposing finite-limit contexts

12.1 A pullback theorem for cartesian fibrations

The failure of arbitrary pullback of n-initial shape maps has a positive counterpart. Let $p : L \to K$ be a cartesian fibration. A cartesian lift of $h : k \to p(\ell)$ with target ℓ is an arrow $\ell_h \to \ell$ over h. Its defining property is that arrows into ℓ whose projection factors through h lift uniquely up to a contractible choice through ℓ_h, with the given factorization on the base. Right fibrations and projections of the form $K \times V \to K$ are examples. We use cartesian in this paragraph in the sense of cartesian fibrations, independently of the existence of products.

Theorem 12.1 (Indexed comparison of weights)**.** *Suppose $p : L \to K$ is a cartesian fibration, $u : A \to K$ is any functor, and $u' : A \times_K L \to L$ is the pullback. There is a canonical equivalence*

$$W_{u'} \simeq p^* W_u$$

of covariant weights on L. It induces the same equivalence for n-truncated weights. Consequently $u' \in \mathsf{L}_n$ if $u \in \mathsf{L}_n$. The factorization of u' from Theorem 4.1 is the pullback along p of the factorization of u.

Proof. Fix $\ell \in L$, with $k = p(\ell)$. The comparison of comma categories is the pullback along $A_{/k} \to K_{/k}$ of $L_{/\ell} \to K_{/k}$. Choosing cartesian lifts into ℓ gives a fully faithful right adjoint to $L_{/\ell} \to K_{/k}$, over $K_{/k}$. Indeed an arrow from any $g : m \to \ell$ into the cartesian lift of h is a lift through that cartesian arrow of a factorization of $p(g)$ through h. For each fixed factorization on the base its space of choices is contractible. This is precisely the mapping-space property of the adjunction. The same proof after pulling back to $A_{/k}$ gives a right adjoint to the comparison of comma categories, with invertible counit. The unit and counit then show that the comparison induces an equivalence on groupoid completions. The comparison itself is canonical and natural in ℓ by postcomposition. Thus the pointwise equivalences give the claimed equivalence of weights. Truncating gives its level-n version.

For the factorization, left fibrations pull back by precomposition of the classifier, including their truncated versions. The section defining the unit in (4.1) is compatible with the comma comparison: the identity of an object in the pullback projects to the corresponding identity in A. Therefore the pulled-back factorization identifies with the one constructed from u'. The preservation of n-initiality follows by comparing augmentations. □

$$\begin{array}{ccc} A \times_K L & \xrightarrow{u'} & L \\ \downarrow & & \downarrow{\scriptstyle p} \\ A & \xrightarrow[u]{} & K \end{array} \qquad \begin{array}{ccc} (A \times_K L)_{/\ell} & \longrightarrow & L_{/\ell} \\ \downarrow & & \downarrow \\ A_{/k} & \longrightarrow & K_{/k} \end{array}$$

Figure 12.1. On the right, a cartesian lift supplies the comparison object over each comma arrow with target k. The comma pullback has the same completion as its lower-left corner.

As a small example, regard [1] as a parameter shape with a single transition. Let a presheaf of sets on [1] be specified by a map $Y_1 \to Y_0$. Its category of elements is a right fibration. If the rewrite starts with the initial-vertex inclusion $\{0\} \to [1]$, the pullback indexes the elements of Y_0. For an element $y \in Y_1$, there is a unique lift of the transition with target y, beginning at its image in Y_0. This is the distinguished object of the relevant comma category, independently of the fibers of $Y_1 \to Y_0$. In the covariant direction the analogous statement fails: the functor $[1] \to \mathsf{Set}$ with values $\varnothing \to *$ gives a left fibration supported only at 1. Pulling the same rewrite back along it gives the empty shape over one vertex. Thus the type of transport in the index category matters.

12.2 Composition of weights

For $u : A \to K$ and $v : K \to L$, let $\mathrm{Lan}_v^{(n)}$ denote left Kan extension with values in $\mathcal{S}_{\leq n}$. The weights satisfy

$$W_{vu,n} \simeq \mathrm{Lan}_v^{(n)} W_{u,n}. \tag{12.1}$$

In fact, the left adjoints to restriction compose. Before truncation this gives $\mathrm{Lan}_{vu} \mathbf{1}_A \simeq \mathrm{Lan}_v \mathrm{Lan}_u \mathbf{1}_A$. For an n-truncated target copresheaf, mapping out of $\mathrm{Lan}_v W_u$ first restricts the target to K and then factors through $W_{u,n}$. This proves (12.1) and the compatibility of its augmentations. In pointwise form at $\ell \in L$, the right side is the colimit in $\mathcal{S}_{\leq n}$ of $W_{u,n}(k)$ over the arrows $v(k) \to \ell$.

This identity supplies a way to use intermediate classifications. Suppose a shape is first interpreted in a theory K, and then in a larger theory L by v. Equivalence of level-n weights over K implies equivalence of the resulting weights over L, even when v is not an equivalence. In particular the proof of an equality in an abstract signature can be transported along a functor interpreting the signature into a concrete model. Equality after that transport gives no converse without a conservativity hypothesis. For example, a functor from the category of groups to sets taking only the trivial group as a constant value identifies all parallel maps. Such a change of models loses much more than the level-0 quotient of a comma space.

12.3 Finite limits and the effect of truncation

Product contexts suffice for total operations. An essentially algebraic context can require a finite limit, such as the subobject of composable pairs of arrows. At the fully coherent level there is a finite-limit-preserving version of comprehensive factorization. The following proof also locates exactly where the extension to finite truncation fails.

Proposition 12.2 (Finite-limit preservation at the full level)**.** *Let A, K be small finitely complete infinity-categories. For any $u : A \to K$ the full weight $W_u : K \to \mathcal{S}$ preserves finite limits. If u preserves finite limits, the factorization $A \to E_{u,\infty} \to K$ has finitely complete middle category and both maps preserve finite limits.*

Proof. The covariant representables give the formula

$$W_u \simeq \operatorname*{colim}_{a \in A^{\mathrm{op}}} \mathrm{Map}_K(u(a), -).$$

Indeed, mapping out of the right side into F gives $\lim_A Fu$ by Yoneda, and the same is true of W_u by (3.2). The opposite of A has finite colimits, so it is filtered. Filtered colimits of spaces commute with finite limits [12, Section 5.3.3], and each displayed representable preserves all limits that exist in K. This proves the first claim.

For a finite diagram in $E_{u,\infty}$, take the limit k of its projection to K. Its given points upstairs form a coherent section of W_u along that diagram. The limit comparison for W_u identifies that section with a point of $W_u(k)$, with its comparison paths. To check the limiting cone, map into it from (d, x) and use the homotopy fiber of $\mathrm{Map}_K(d, k) \to W_u(k)$. The two limit comparisons identify this fiber with the limit of the corresponding mapping-space fibers. This proves the limit property. It proves reflection by the same argument for any cone whose projection is limiting. Therefore the left map of the factorization preserves these limits whenever its composite with the right map does. □

The same argument for ordinary finitely complete categories, ordinary filtered colimits, and set-valued representables gives an ordinary finite-limit result. The coherent argument cannot be followed by absolute n-truncation without further hypotheses, even for $n = 0$. The issue is the use of finite *homotopy* limits in the base category.

Proposition 12.3 (Failure at finite levels)**.** *For each finite $n \geq 0$ there is a finite-limit-preserving $u : A \to K$ for which the projection $E_{u,n} \to K$ cannot create finite limits. In fact a specific cospan upstairs has no pullback projecting to the pullback of its image.*

Proof. Let $Z = K(C_2, n+1)$ be a pointed Eilenberg–Mac Lane space. Take K to be a small full subcategory of spaces closed under finite limits and containing $*, Z$. Such a subcategory can be chosen by starting with these two objects and iteratively adding representatives of limits of all finite diagrams in the full subcategory at the previous stage. Let $u : * \to K$ pick the terminal object. It is finite-limit-preserving. By the covariant representable formula, $W_u(k) \simeq k$ and $W_{u,n}(k) \simeq \tau_{\leq n} k$.

Let e_0 be the object over $*$, e_Z the object over Z with its unique point in the truncation up to equivalence, and take the two identical arrows $e_0 \to e_Z$ over the chosen base point. The projected pullback is $\Omega Z \simeq K(C_2, n)$. At any object e over ΩZ, the mapping space $\mathrm{Map}(e_0, e)$ is a homotopy fiber of $\Omega Z \to \tau_{\leq n}\Omega Z$. It is contractible. On the other hand, mapping from e_0 into the cospan produces $* \to Z \leftarrow *$, because $\tau_{\leq n} Z \simeq *$. Its pullback is ΩZ, which is noncontractible. Thus no cone over the projected pullback can be a pullback upstairs. □

$$\begin{array}{ccc} \Omega Z & \longrightarrow & * \\ \downarrow & & \downarrow \\ * & \longrightarrow & Z \end{array} \qquad \begin{array}{ccc} \tau_{\leq n}\Omega Z & \longrightarrow & * \\ \downarrow & & \downarrow \\ * & \longrightarrow & \tau_{\leq n} Z \simeq * \end{array}$$

Figure 12.2. The left square is a homotopy pullback. The right loses the loop information needed for that pullback. Proposition 12.3 expresses the failure inside a left fibration of witnesses.

The contrast gives a choice of hypotheses for essentially algebraic examples. If a full weight is known to be n-truncated valued, truncating it does nothing and the finite-limit proof still applies. If only the observed solutions are n-truncated, Theorem 3.2 still gives

cone invariance, but it does not assert preservation of finite limits by the saturation projection. To handle a specified partial operation one can instead keep its domain explicitly as a limit context.

13 Partial operations and geometric groupoids

13.1 Definitional extension in a fixed domain

Let P be a context object in a category with the indicated pullbacks. A partial operation on X with values in Y consists here of a domain monomorphism $j : U \hookrightarrow X$ and a map $m : U \to Y$. For an input $a : P \to X$, the domain of definition in that context is $P_U = P \times_X U$. The operation defines a map $P_U \to Y$ and a graph in $P_U \times Y$. If a second equation uses this output, its variable must range over that graph, not all of $P \times Y$.

Proposition 13.1 (Introducing a dependent value)**.** *The context with the dependent variable $y = m(u)$, over the pullback P_U, represents the same solution functor as P_U itself. This isomorphism is natural under any pullback of the input context for which the relevant pullbacks exist. The corresponding statement for a map m in an infinity-category uses a path witness for $y = m(u)$ and gives an equivalence.*

Proof. For any probe Q, a point of the extended context consists of $p : Q \to P$, a lift $u : Q \to U$ with $ju = ap$, and a value $y : Q \to Y$ with $y = mu$. The last datum is uniquely determined by the first two, and the first two are a map $Q \to P_U$. This is natural in Q and under context pullbacks. In the coherent version use the homotopy pullback to account for the first comparison. For a fixed u, the space of pairs consisting of y and a path to mu in $\mathrm{Map}(Q, Y)$ is contractible. Thus adjoining that value and witness is again an equivalence, naturally in the probes. □

$$\begin{array}{ccccc} P_U & \longrightarrow & U & \xrightarrow{m} & Y \\ \downarrow & \lrcorner & \downarrow{\scriptstyle j} & & \\ P & \xrightarrow[a]{} & X & & \end{array}$$

Figure 13.1. The operation becomes total after passing to the pullback context. The new variable is the graph value over that context.

To compare two applications of a partial operation, the input expressions must be compared in the domain U. An equality in X on a probe can be used together with membership in the domain. If the domain is represented by a monomorphism, two lifts over the same arrow into X are already equal. For a general map $U \to X$ its fibers may contain several choices of input structure. Then equality after projection to X alone cannot justify the fold. For example the two points of a double cover can give opposite output signs. This is an instance where an equation about the projected input discards a premise of the operation. In an infinity-category with a (-1)-truncated domain inclusion, lifts again form an empty or contractible space. For higher fibers they can carry genuine automorphisms.

13.2 An internal category as a finite-limit theory

For an internal category with objects C_0 and arrows C_1, composition is defined on $C_2 = C_1 \times_{C_0} C_1$, using the target of the first arrow and the source of the second. The identities are a map $C_0 \to C_1$. Associativity is equality of two arrows $C_3 \to C_1$, where C_3 is the iterated pullback of three composable arrows. Product-context equations on C_1^3 would have a larger domain. In the partial syntax, introduction of the composite of the first two arrows is a definitional extension on C_3 as in Proposition 13.1. The source and target laws place the result in the domain of the second composition. Thus both paths through the diagram in Figure 13.2 are well typed before the associativity equation is used.

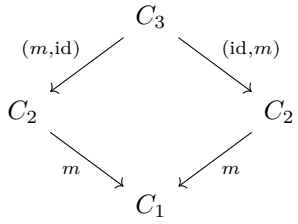


Figure 13.2. The associativity diagram lives over the context of three composable arrows. Each edge to C_2 uses the source and target identities for composition.

In Man a concrete family of examples comes from a surjective submersion $p : M \to B$. Define $C_0 = M$, $C_1 = M \times_B M$, with (x, y) an arrow from x to y. The source and target projections are submersions by base change. The finite fiber products defining C_2, C_3 exist and are the triple and quadruple products over B. Composition is

$$(x, y), (y, z) \longmapsto (x, z).$$

Both routes in Figure 13.2 send (x, y, z, w) to (x, w). The unit is (x, x) and inverse exchanges the entries. These formulas are smooth, since they are restrictions of product projections. On an arbitrary smooth probe the relevant fiber products have exactly the required equality of parameter maps. This is a Lie groupoid with a finite-limit context whose existence follows from submersivity.

13.3 Free actions, effective relations, and stabilizers

For $p : S^1 \to S^1$, $p(z) = z^2$, the groupoid just constructed has arrow manifold $S^1 \times C_2$. The pair (z, z') is in the same fiber precisely when $z' = \epsilon z$ for $\epsilon = \pm 1$, and ϵ is determined by the pair. The groupoid is the translation groupoid of the free action of C_2 on S^1 by sign. Its map to the target circle is a local submersion with local sections, so the smooth sheaf map $y(p) : y(S^1) \to y(S^1)$ is an effective epimorphism. Its Čech groupoid in smooth sheaves is the image of the preceding internal groupoid. The quotient as a sheaf of spaces is $y(S^1)$, by the Čech realization formula. Automorphisms in this groupoid are trivial, since z and $-z$ are distinct.

$$S^1 \times C_2 \underset{(z,\epsilon)\mapsto \epsilon z}{\overset{(z,\epsilon)\mapsto z}{\rightrightarrows}} S^1 \xrightarrow{\ p\ } S^1$$

Figure 13.3. The action groupoid of the free sign action is the Čech groupoid of the smooth double cover.

Contrast this with the sign action on $\mathbb{R}$ by $x \mapsto -x$. The smooth action groupoid has the same two possible group elements for a change of representative. Away from 0 the stabilizer is trivial. At 0 it is C_2. In the associated quotient stack, an object over a probe is a C_2-torsor over that probe with an equivariant map to $\mathbb{R}$, and morphisms are equivariant isomorphisms respecting the map. Over a one-point probe every torsor is trivial up to isomorphism, so the solution groupoid is the action groupoid of C_2 on the points of $\mathbb{R}$. The zero component is equivalent to BC_2, whereas each strictly positive orbit is a point up to equivalence. This direct description verifies that the quotient stack is not a representable smooth manifold: the functor of points of a manifold gives a discrete groupoid on the point, which cannot have the automorphism at zero. The ordinary orbit set, which is a ray as a topological space, only records the components of this test.

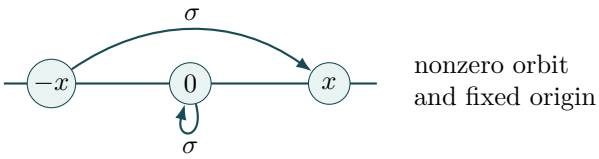


Figure 13.4. A nontrivial automorphism survives at the fixed point of the sign action on $\mathbb{R}$. It is visible in coherent solutions even after all orbit identifications.

The two actions demonstrate why internal quotient data can have two meanings. The Čech groupoid of a cover between discrete objects is the effective relation given by its kernel pair. It gives descent to its target. An action groupoid with stabilizers is a groupoid object whose quotient retains those stabilizers as loop spaces. For the first action the solution from the square-root problem is the fiber sheaf of the cover, which is nontrivial over the circle. For the second action a local test at the fixed orbit already detects a positive homotopy level. In both examples the domains and witnesses of composition must be respected by a diagram rewrite.

14 Using the calculus as a certificate

14.1 Local certificates and their scope

The preceding constructions identify different data to record in a proof. A shape certificate consists of source and target diagrams over a common theory, an oriented map or zigzag between their shapes, and a level n. To apply a local rule inside a larger diagram, the attaching map must identify the variables and operation labels of the rule with their actual occurrences. The induced shape map is a pushout. Its test is that the indicated comma spaces are n-connected, or that it is built from maps already satisfying that test. Checking only the two endpoint labels leaves out this attaching map.

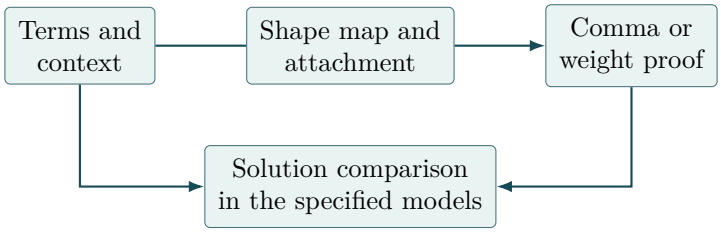


Figure 14.1. A rule is checked at the shape level, then interpreted in models of the declared observational level.

For the ordinary unary motifs one can give a connectedness certificate by exhibiting an initial root, a common path, or a finite chain of folds. The cycle proof in Section 2 is an example in which the quotient shape remains initial after each fold. Theorem 4.5 is completeness for semantic equivalence with respect to the indicated collection of all tests. Its converse constructs a zigzag via the comprehensive factorization. The two middle weights can involve infinitely many objects or higher paths. Thus this proof gives no bound on the size of a certificate consisting solely of finitely many chosen motifs.

There are also distinct operations on the witnesses themselves. In an ordinary topos the image of the object of finite derivations is the congruence relation. In a coherent theory a path can identify two terms, and a higher path can compare two identifications. An invariant computed only from the relation of existence of a path cannot decide whether the space of such identifications is contractible. For an explicit calculation, the two homotopy classes of loops in BC_2 differ by the nontrivial group element. Both have the same source and target point. The homotopy-level comparison must therefore be tested at the level of the cone or of the weight, even if the equation of components is settled.

14.2 Worked comparison of three candidate rewrites

Take $K = [1]$ with objects $0, 1$ and arrow $b : 0 \to 1$. Let $F : K \to \mathcal{S}$ send b to $f : X \to Y$. Consider the following three ways of recording an output of f.

1. The arrow shape K with its identity map to the theory. A cone has an input x and an output y with a path from $f(x)$ to y. The space is equivalent to X.
2. The fork shape with two separate targets, all arrows over b. It has an initial root, so its solution space is again X. The fold to the arrow shape is in L_∞. A point of its solution over x is a pair of independent choices of an output and its path from $f(x)$, each in a contractible space.
3. The free parallel-pair shape with both target vertices identified, still with both arrows over b. A coherent solution now has x, one y, and two paths from $f(x)$ to y. Over fixed x, the space of these data is equivalent to $\Omega_{f(x)}Y$. Indeed the first path and its endpoint can be eliminated by the contractible choice argument, leaving the second path as a loop.

The quotient from the fork to the parallel-pair shape is not initial at the target vertex, since its comma category there has two components. But over the theory K, the comparison of the two solution functors is an equivalence for every set-valued F. In

that setting there is exactly one witness of equality whenever the two images agree, and here they already agree because both arrows have label b. At level 1, taking $X = *$ and $Y = BC_2$ gives one solution component for the first two shapes and two for the third.

The weight computation makes this example relative and explicit. At 0, all three weights have a single point. For the arrow and fork, the weight at 1 is contractible. The arrow case has a terminal comma object $(1, \mathrm{id})$. The fork case has a comma category with initial object (r, b). For the parallel-pair shape, the comma category at 1 has the root over b, the common endpoint over id, and two parallel arrows between them. Its realization is a circle. The induced map from the fork's comma space to this circle is an equivalence after 0-truncation and fails after 1-truncation. This is the calculation underlying the different comparison results, without a need to choose a model in advance.

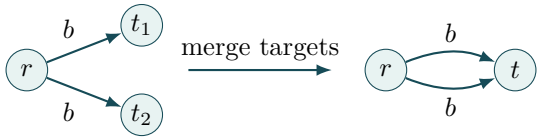


Figure 14.2. Identifying the output vertices while retaining two arrows creates a loop in the relative comma weight. Folding the arrows as well is a different shape map.

14.3 What each target tests

The table summarizes the operations justified by the preceding arguments. The row for algebras requires that the underlying limit and the designated detector hypotheses hold. The entries for sheaves distinguish preservation by shape changes from reindexing along a functor of ambient categories.

Target	Data in a solution	Hypothesis used in the examples
Smooth manifolds	Smooth maps from all probes, with strict equations	Ordinary initiality for cones; transversality or another representability proof for a smooth solution object
Groupoids up to equivalence	Objects, isomorphism witnesses, automorphisms	Level 1 for uniform changes of shape; iso-comma for pullbacks
Ordinary sheaves and toposes	Compatible local sections and internal relations	Ordinary initiality for cones; exactness for congruence quotients and geometric inverse image
Slices of an infinity-topos	Families of choices and coherences over a parameter	Relative truncated weights and effective-epimorphism descent
Algebras for a monad	Cones respecting action and substitution	Creation of limits by the underlying functor, plus the cone comparison in the underlying category

The example of the two parallel paths shows that the set level can certify a comparison which the groupoid level rejects for the same fixed theory. The crossing axes show that invariance of smooth cone functors is independent of representability by a smooth object. The two sign actions show that component-level quotienting and stack quotienting can agree for one action and differ for another, depending on stabilizers. In each case the

criterion is a condition on the objects or witnesses of the actual problem. It is not a property of the visual drawing alone.

Authorship and evidence

This manuscript is prepared under an anonymous byline for editorial submission. The submitting author is responsible for the claims, proofs, references, and the choice of submission. An AI assistant assisted with literature retrieval, proof development, preparation of the manuscript, and diagram typesetting. The evidence consists of the stated mathematical proofs, worked calculations, and cited foundational results.

References

[1] Anonymous. *Higher E-Graphs: Cartesian Factorization, Proof Spaces, and Topos Semantics*. Higher Categorical Semantics, October 22, 2026. Editorially reviewed. Journal record.

[2] John C. Baez and Alexander E. Hoffnung. *Convenient categories of smooth spaces*. Transactions of the American Mathematical Society 363 (2011), 5789–5825. arXiv:0807.1704.

[3] Roberto Biondo, Davide Castelnovo, and Fabio Gadducci. *EGGs are adhesive!* 2025. arXiv:2503.13678.

[4] Francis Borceux and Brian J. Day. *On product-preserving Kan extensions*. Bulletin of the Australian Mathematical Society 17 (1977), 247–255. DOI:10.1017/S0004972700010455.

[5] John Bourke and Richard Garner. *Monads and theories*. Advances in Mathematics 351 (2019), 1024–1071. arXiv:1805.04346.

[6] Denis-Charles Cisinski. *Higher Categories and Homotopical Algebra*. Cambridge University Press, 2019. Chapters 4–6. DOI:10.1017/9781108588737.

[7] Nicola Gambino and Joachim Kock. *Polynomial functors and polynomial monads*. Mathematical Proceedings of the Cambridge Philosophical Society 154 (2013), 153–192. arXiv:0906.4931.

[8] David Gepner, Rune Haugseng, and Joachim Kock. *∞-operads as analytic monads*. International Mathematics Research Notices 2022 (16), 12516–12624. arXiv:1712.06469.

[9] F. William Lawvere. *Functorial semantics of algebraic theories*. 1963. Reprinted with additional material in Reprints in Theory and Applications of Categories 5 (2004), 1–121. TAC reprint.

[10] John M. Lee. *Introduction to Smooth Manifolds*. Second edition, Springer, 2013. Chapters 4–6. DOI:10.1007/978-1-4419-9982-5.

[11] Tom Leinster. *Higher Operads, Higher Categories*. Cambridge University Press, 2004. arXiv:math/0305049.

[12] Jacob Lurie. *Higher Topos Theory*. Princeton University Press, 2009. arXiv:math/0608040.

[13] Jacob Lurie. *Higher Algebra*. September 2017. Section 3.1 and Chapter 4. Author's text.

[14] Jacob Lurie. *Kerodon*. Subsection 7.2.1, Cofinal Morphisms of Simplicial Sets. Tag 02N0.

[15] Greg Nelson and Derek C. Oppen. *Fast decision procedures based on congruence closure*. Journal of the ACM 27 (1980), 356–364. DOI:10.1145/322186.322198.

[16] Evan Patterson. *How to prove equations using diagrams, part 1. Initial functors*. May 27, 2025. Author's article.

[17] Evan Patterson. *How to prove equations using diagrams, part 2. Weak equivalences and e-diagrams*. June 10, 2025. Author's article.

[18] Evan Patterson. *How to prove equations using diagrams, part 3. The setting of algebraic theories*. October 14, 2026. Author's article.

[19] Evan Patterson. *How to prove equations using diagrams. A categorical and operational semantics for e-graphs*. November 4, 2026, version 1. arXiv:2611.05924v1.

[20] Emily Riehl. *Categorical Homotopy Theory*. Cambridge University Press, 2014. Lemma 8.3.4. Author's text.

[21] Emily Riehl and Dominic Verity. *Homotopy coherent adjunctions and the formal theory of monads*. Advances in Mathematics 286 (2016), 802–888. arXiv:1310.8279.

[22] The Stacks Project Authors. *The Stacks Project*, Section 6.29, Limits and colimits of sheaves. Tag 009E.

[23] Ross Street and Robert F. C. Walters. *The comprehensive factorization of a functor*. Bulletin of the American Mathematical Society 79 (1973), 936–941. DOI:10.1090/S0002-9904-1973-13268-9.

[24] Aleksei Tiurin, Chris Barrett, Dan R. Ghica, and Nick Hu. *Equivalence hypergraphs. DPO rewriting for monoidal e-graphs*. arXiv:2406.15882.

[25] Max Willsey, Chandrakana Nandi, Yisu Remy Wang, Oliver Flatt, Zachary Tatlock, and Pavel Panchekha. *Egg. Fast and extensible equality saturation*. Proceedings of the ACM on Programming Languages 5 (POPL), article 23, 2021. arXiv:2004.03082.