RESEARCH MONOGRAPH

Galois Theory in Higher Categories

Descent, monodromy, and higher topoi

Manuscript for *Higher Categorical Semantics*

November 2026

Preface

The ordinary expression $L \otimes_K L$ in Galois theory serves two purposes. It describes the overlap of the extension with itself, and, when splitting holds, it encodes the symmetries by which copies of a descended object are identified. The higher analogue is the entire simplicial object of overlaps. Its realization recovers the base under effective descent. Its Segal operators describe multiplication in a groupoid, internal composition in a category, or a monad on a category of objects over a splitting base. These three conclusions require different hypotheses. Much of the work below is to identify them and prove the comparisons between their outputs.

The point of departure is Francis Borceux and George Janelidze's *Galois Theories* [1]. Their sequence of generalizations passes from subgroup correspondences to split algebras, profinite symmetry, Pierce spectra, relative reflections, coverings, and generalized descent. This monograph follows those questions through the passage from ordinary categories to $(\infty, 1)$-categories. The organizing results are homotopy-coherent descent in an ∞-topos and the classification of families by transport. They are joined by versions in stable and directed settings where the relevant additional structure survives.

Reading the book

Part I develops the correspondence between families and actions and proves the restriction of effective descent to chosen coefficients. Part II treats continuity of monodromy, the shape of a higher topos, and obstructions to lifting actions through truncations. Part III compares Pierce descent, connective spectral étale descent, and Galois extensions in stable homotopy theory. Part IV studies admissible reflections in slices, factorization, and centrality. Part V turns from invertible transport to right and left fibrations, directed descent, and exodromy. Part VI returns to toposes themselves and compares internal, geometric, and external presentation theorems.

The correspondence with the original book is thematic. Its Chapters 1–2 are treated in Part I, its Chapter 3 in Part II, and its Chapter 4 and appendix in Part III. Its Chapter 5 chiefly informs Part IV. Its Chapter 6 is distributed over Parts I, II, and IV, where the notions of covering, connectedness, and local triviality can be stated independently. Its Chapter 7 motivates Parts V and VI. The exposition assumes familiarity with adjunctions, slices, and elementary homotopy theory. It gives universe and variance conventions before using any classification theorem. Deep inputs receive references at their points of use.

Abstract

We develop higher analogues of the categorical Galois theories of Borceux and Janelidze. Effective descent in an ∞-topos yields actions of Čech groupoids and classifies torsors by delooping. Pro-spaces encode continuous monodromy under coefficient hypotheses, while Postnikov extensions give obstructions to lifting actions. We compare Pierce and spectral descent, admissible reflections and relative factorization, directed transport and exodromy, and groupoid presentations of topoi. The ordinary and higher statements are connected by specified truncations, duality, and base-change comparisons.

Manuscript disclosure

This manuscript is an independent mathematical treatment with the source book explicitly cited. Generative AI assisted with drafting, bibliography research, mathematical consistency checks, and typesetting. The journal submission metadata provides the responsible author or research group.

Contents

Preface **ii**

I Torsors and descent **1**

1 The language of coherent descent **2**
1.1 Conventions and the ambient category 2
1.2 Cartesianness and realizations 3
1.3 The overlap diagram 4
1.4 Restricting coefficients 5

2 Torsors and monodromy **6**
2.1 Pointed connected objects 6
2.2 Actions and classifying maps 7
2.3 Transport over a space 9
2.4 The subgroup correspondence and its higher replacement . . 10
2.5 Regular intermediates and delooped quotients 13

3 Effective descent and Čech actions **14**
3.1 Descent data and their realization 14
3.2 Mapping spaces and properties local on the base 15
3.3 The descent monad 16
3.4 Refinements and changes of groupoid presentation 17
3.5 Fixed split type and the torsor of frames 18
3.6 The boundary between Čech descent and hyperdescent 19

II Monodromy and shape **21**

4 Pro-spaces and continuous monodromy **22**
4.1 Pro-objects and the role of the test space 22
4.2 A representation category at stages 23
4.3 From one-stage actions to continuous discrete actions 25
4.4 Closed stabilizers 26
4.5 Continuous cochains and twisted coefficients 27

5 The shape of a topos and its local systems **29**
5.1 Shape as a functor of constant tests 29
5.2 From stage representations to objects 30
5.3 Two comparison theorems 32
5.4 The classifying topos 33
5.5 A failure of unrestricted stage monodromy 34
5.6 Higher transport 35
5.7 Hypercompletion and shape 35

6 Obstructions to higher monodromy **37**
6.1 Automorphisms of an Eilenberg–Mac Lane space 37
6.2 The Postnikov lifting problem 38
6.3 Continuous lifting and convergence 39

III Algebraic Galois theories **42**

7 Pierce spectra and derived étale descent **43**
7.1 Idempotents, localizations, and stalks 43
7.2 A relative Stone adjunction 44
7.3 The Galois groupoid of a normal Pierce splitting 47
7.4 Flat descent and the geometric étale site 49
7.5 What the two restrictions detect 50

8 Spectral Galois extensions and stable monodromy **51**
8.1 Three effectivity inputs 51
8.2 The Amitsur complex of a Rognes extension 52
8.3 Intermediate split algebras 54
8.4 The stable Galois category and its range 55
8.5 Periodic rings and higher local systems 56

IV Reflections and factorization **58**

9 Reflective Galois structures and factorization **59**
9.1 The adjunction seen from a varying base 59
9.2 The comparison over one descent map 61
9.3 Two factorization systems arising from truncation 63
9.4 Localizing the reflective pair 65

10 Central extensions and higher bands **67**
10.1 The ordinary abelianization comparison 67
10.2 Higher central extensions as pointed principal bundles 69
10.3 Higher cohomology and the role of trivialization 72

10.4 Abelianization after passing to group objects in spaces 73

V Directed Galois theories 75

11 Directed symmetries and generalized descent 76
11.1 Complete Segal objects and groupoid objects 76
11.2 Coherent presheaves and direction 77
11.3 Two factorization systems 78
11.4 Descent along an augmented diagram 80
11.5 Descent of directed representations 82

12 Exit paths and constructible monodromy 84
12.1 The topological classification 84
12.2 A conical calculation with higher stalks 86
12.3 A van Kampen principle with directed coefficients 87
12.4 Recovering the directed category from its coefficients 88
12.5 The profinite directed case 90

VI Higher topoi and presentation 91

13 Geometric morphisms and the classification of sections 92
13.1 The category of geometric morphisms 92
13.2 The universal property of a site 93
13.3 Étale maps and dependent substitution 95
13.4 Internal classifying objects and external classifiers 96

14 Groupoid presentations of higher topoi 97
14.1 The descent formula for a diagram of slices 97
14.2 From a groupoid to its quotient 98
14.3 The classical localic theorem 100
14.4 When a higher étale atlas has localic overlaps 100
14.5 Hyperdescent and the role of truncation 102

15 Reconstruction and comparison 104
15.1 Recovering the space of fibers 104
15.2 What truncation preserves 105
15.3 A comparison of the classification results 106

Bibliography 108

Subject index 109

Part I

Torsors and descent

CHAPTER 1

The language of coherent descent

In its simplest form a Galois correspondence compares subobjects of a given extension with subgroups of its symmetries. A stronger statement recovers the category of objects that split after a base change from trivial objects on that base and their gluing data. When gluing takes place in an ∞-category, a gluing map comes with compatibilities among homotopies and coherences among those compatibilities. A simplicial overlap diagram organizes this information without requiring a separate definition at each level.

1.1 Conventions and the ambient category

Fix a universe of small spaces and write $\mathcal{S}$ for their ∞-category. Functor categories, categories of large coefficients, and pro-categories are formed in larger universes when needed. A space means a homotopy type. An ∞-category means an $(\infty, 1)$-category, so the morphisms between two given objects form a space $\mathrm{Map}(-,-)$. For a topological space T, we distinguish its topology, its singular homotopy type, and the ∞-topos of sheaves on its open-set site. These inputs agree on certain classifications under the local hypotheses given below.

Limits, colimits, slices, and fibers are understood in the higher categorical sense. The space $\mathrm{Map}_{\mathrm{Cat}_\infty}(\mathcal{C},\mathcal{D}) = \mathrm{Fun}(\mathcal{C},\mathcal{D})^{\simeq}$ is the maximal subgroupoid of the functor category. The functor category itself uses all natural transformations, an important distinction when we study geometric morphisms in Chapter 13.

Definition 1.1. An ∞-topos is an accessible left exact localization of a presheaf ∞-category $\mathrm{Fun}(C^{\mathrm{op}}, \mathcal{S})$ for a small ∞-category C. A geometric morphism $f : \mathcal{Y} \to \mathcal{X}$ is specified by its inverse-image functor $f^* : \mathcal{X} \to \mathcal{Y}$, which preserves colimits and finite limits. Its right adjoint is denoted f_*.

A geometric morphism points in the opposite direction from its inverse image. For a map $p : U \to B$ inside $\mathcal{X}$, the same notation p^* means base change from $\mathcal{X}_{/B}$ to $\mathcal{X}_{/U}$. The left adjoint $p_!$ composes a map with p. The right adjoint p_* is dependent product. Its existence follows from local cartesian closure. If p is a map of spaces and $V \to U$ is a family, the fiber of $p_!V$ over b is the total space over the homotopy fiber U_b. The fiber of p_*V over b is the space of sections over U_b.

Lemma 1.2 (Mapping from a section)**.** *Given $q : E \to B$ and $a : A \to B$ in $\mathcal{X}$, one has*

$$\mathrm{Map}_{\mathcal{X}_{/B}}(A, E) \simeq \mathrm{fib}_a\big(\mathrm{Map}_{\mathcal{X}}(A, E) \xrightarrow{q\circ -} \mathrm{Map}_{\mathcal{X}}(A, B)\big).$$

This is also the space of sections of $a^*E \to A$.

Proof. A point of the fiber consists of a map to E and a homotopy identifying its composite with a. These are the data in the slice mapping space, with the same assertion for families parametrized by a test space. The pullback universal property rewrites that homotopy fiber as maps from the terminal object of $\mathcal{X}_{/A}$ to a^*E. □

If F is a fixed fiber, its space of automorphisms $\mathrm{Aut}(F)$ consists of the equivalence components of $\mathrm{Map}(F, F)$. Composition makes it a group object in spaces. Its group of components $\pi_0 \mathrm{Aut}(F)$ classifies automorphisms up to homotopy. A family with fiber F can depend on homotopies within a single component of $\mathrm{Aut}(F)$, even when the group of components of the automorphism object acts trivially.

1.2 Cartesianness and realizations

A cartesian family over a diagram $U : I \to \mathcal{X}$ is a natural transformation $V \to U$ whose square over every arrow is a pullback. Let $B = \mathrm{colim}_I U$. Pullback defines a comparison

$$\mathcal{X}_{/B} \longrightarrow \lim_{i \in I^{\mathrm{op}}} \mathcal{X}_{/U_i}. \tag{1.1}$$

The target consists of cartesian families. The identifications over arrows and their higher coherences are part of the data. A colimit for which this comparison is an equivalence is called van Kampen.

Proposition 1.3 (A test for descent of a diagram)**.** *Suppose the ambient category admits pullbacks and the required colimits, and these colimits are preserved by base change. The comparison* (1.1) *is fully faithful. It is an equivalence exactly when for each cartesian $V \to U$, the squares*

$$\begin{array}{ccc} V_i & \longrightarrow & \mathrm{colim}_j V_j \\ \downarrow & & \downarrow \\ U_i & \longrightarrow & B \end{array}$$

are pullbacks.

Proof. Realization takes V to $\mathrm{colim}_I V \to B$. For any $W \to B$, maps from the realization to W are the limit of maps from $V_i \to B$ to W, by the universal property of the colimit in the slice. Base change rewrites the latter mapping spaces as maps from $V_i \to U_i$ to $W \times_B U_i$. Their limit is the space of maps of cartesian families into the comparison image of W. Realization is therefore left adjoint to the comparison. The counit $\mathrm{colim}_I(W \times_B U_i) \to W$ is an equivalence

by the base-change assumption, so the right adjoint is fully faithful. For each V, the unit is an equivalence exactly when its components $V_i \to U_i \times_B \operatorname{colim}_j V_j$ are equivalences. □

A candidate over the base is determined by its restrictions as soon as the comparison is fully faithful. Existence of a candidate for every cartesian family requires the unit condition. Each covering theorem below uses a specific argument for existence.

Theorem 1.4 (Van Kampen property of higher topoi)**.** *Small colimits in an ∞-topos are van Kampen.*

Proof. In spaces, straightening over a space K gives $\mathcal{S}_{/K} \simeq \mathrm{Fun}(K, \mathcal{S})$, covariantly for a family viewed by transport. The inclusion of spaces into Cat_∞ preserves colimits because it is left adjoint to the maximal subgroupoid functor. The internal-hom universal property turns colimits of bases into limits of the functor categories. In a presheaf ∞-category, the two criteria in Proposition 1.3 then hold objectwise.

Write $L : \mathcal{P} \rightleftarrows \mathcal{X} : i$ for a left exact localization. A diagram of local objects and its cartesian families can be viewed in $\mathcal{P}$, since the inclusion preserves limits. Realizing there gives pullback squares as in Proposition 1.3. Applying L preserves those pullbacks and gives their realizations in $\mathcal{X}$. Universality of colimits in $\mathcal{X}$ follows by the same localization argument. If B_P is the colimit of the diagram in $\mathcal{P}$, pull back that colimit along $iW \times_{iB} B_P \to B_P$ and apply L. This identifies W as the colimit of the diagram base changed along $W \to B$. The two criteria now apply in $\mathcal{X}$. The straightening input is developed in [2, Ch. 2]. □

1.3 The overlap diagram

The Čech nerve of $p : U \to B$ has k-simplices

$$U_k = \underbrace{U \times_B \cdots \times_B U}_{k+1}.$$

Its operators omit or repeat factors. It is a groupoid object because the edges encode pairs of local lifts and the inverse operation reverses a pair. The map p is an effective epimorphism if the augmentation $U_\bullet \to B$ exhibits B as its realization. In spaces this amounts to surjectivity on components. In a topos of sheaves it expresses local existence, so a section over the full base is unnecessary.

Corollary 1.5 (Effective descent over one cover)**.** *For an effective epimorphism $p : U \to B$ in $\mathcal{X}$, the natural comparison is an equivalence*

$$\mathcal{X}_{/B} \simeq \mathrm{Tot}\big([k] \longmapsto \mathcal{X}_{/U_k}\big).$$

Moreover p^* is conservative.

Proof. Apply Theorem 1.4 to the Čech nerve. For a morphism in the limit whose degree-zero component is an equivalence, each degree-k component is an equivalence by restriction along a vertex map $[0] \to [k]$. Equivalences in the limit are detected componentwise. □

A stable category or a category of algebras can require a different proof. Theorem 1.4 concerns slices of an ∞-topos. A comparison for modules over a spectral ring computes a totalization of module categories using scalar extension. Its conservativity and continuity require hypotheses on that extension. Chapter 8 gives such hypotheses.

1.4 Restricting coefficients

Let $D^\bullet$ be a cosimplicial diagram of ∞-categories. Let E^k be replete full subcategories of D^k, preserved by all structure functors.

Lemma 1.6 (Restricting full coefficients)**.** *The functor $\mathrm{Tot}\, E^\bullet \to \mathrm{Tot}\, D^\bullet$ is fully faithful. An object of $\mathrm{Tot}\, D^\bullet$ lies in its essential image exactly when its degree-zero component lies in E^0.*

Proof. For any pair of compatible families, the mapping space is the limit of the corresponding levelwise mapping-space diagram. Fullness makes the two diagrams equivalent in every degree. The essential image consists of objects whose components lie in the replete full subcategories. Each structure equivalence and coherence of the given diagram belongs to the full subcategory as soon as its vertices do. Applying a vertex map $[0] \to [k]$ to the degree-zero component puts the degree-k component in E^k. □

In applications one must establish the unrestricted comparison and then check stability of the coefficients or construct a comparison on which stability holds. In algebraic Galois theory the second step can follow from a relative left exactness condition. In stratified geometry the comparison must respect which strata are connected by directed arrows. An equivalence of cores classifies objects and their equivalences. We use a comparison of full categories when morphisms between the coefficients also matter.

CHAPTER 2

Torsors, monodromy, and intermediate objects

The first two chapters of Borceux and Janelidze pass from intermediate fields to a category of split algebras. The Galois group acts on a set of embeddings. Inside the resulting action category, the quotients of the regular action recover the intermediate fields, with arrows reversed by the passage from algebras to actions [1, Chapters 1–2]. The argument suggests two distinct operations. One classifies objects trivialized by a cover. The other identifies which pointed intermediate objects are governed by subgroups. In an ∞-category these operations have different answers unless a truncation condition is imposed.

2.1 Pointed connected objects

Write $\mathcal{S}$ for the ∞-category of spaces and $\mathcal{X}$ for an ∞-topos. Pullbacks, colimits, and fibers are taken in the ∞-categorical sense. We write 1 for a terminal object and call a map $1 \to B$ a point of B internal to $\mathcal{X}$. This means a section of an object in the topos. A geometric morphism $\mathcal{S} \to \mathcal{X}$ is a different operation.

A simplicial object $D_\bullet$ is a groupoid object if its Segal maps are equivalences and the arrows are invertible in the internal sense. For example, the Čech nerve of a map $q : U \to B$ is the groupoid object with

$$\check{C}(q)_n = \underbrace{U \times_B \cdots \times_B U}_{n+1 \text{ factors}}.$$

Its face maps delete vertices and its degeneracies repeat vertices. The map q is an effective epimorphism if the augmentation $|\check{C}(q)_\bullet| \to B$ is an equivalence. Effectivity of groupoids in an ∞-topos says that for every groupoid object $D_\bullet$, the canonical map $D_0 \to |D_\bullet|$ is an effective epimorphism and its Čech nerve identifies with $D_\bullet$ [2, Section 6.1.3].

A group object is a groupoid object $G_\bullet$ with terminal object of objects. Its arrow object is denoted by $G = G_1$. In degree n the Segal maps identify G_n with G^n. The two ways of composing three arrows agree through the simplicial structure, which also supplies all higher coherences. The delooping is $BG = |G_\bullet|$, pointed by the degree-zero inclusion. A pointed object $b : 1 \to B$ is connected if b is an effective epimorphism. In this situation,

$$\Omega_b B = 1 \times_B 1$$

is the based loop group object, with simplicial description $\check{C}(b)_\bullet$. The definition of connectedness here is equivalent to $\tau_{\leq 0}B \simeq 1$. In fact a subobject of B is classified by a map to the subobject classifier, which is a 0-truncated object. If $\tau_{\leq 0}B \simeq 1$, any subobject containing b is therefore all of B, so the image of b is all of B. Conversely, if b is an effective epimorphism, a map from B to a 0-truncated object is determined by its restriction along b. Every map from 1 to that object extends along $B \to 1$. The resulting universal property identifies $\tau_{\leq 0}B$ with 1.

Proposition 2.1 (Loop and delooping)**.** *Taking loops and taking delooping define inverse equivalences between the ∞-category of group objects in $\mathcal{X}$ and the ∞-category of pointed connected objects of $\mathcal{X}$.*

Proof. A map of group objects induces a pointed map on geometric realizations. A pointed map induces a map on Čech nerves of its points. For a pointed connected object the augmentation of its Čech nerve realizes to the object. For a group object its simplicial groupoid is the Čech nerve of the point of its realization by effectivity. Both identifications are canonical maps of diagrams. They give unit and counit equivalences of the two functors, including on mapping spaces. □

These operations also work in every slice $\mathcal{X}_{/B}$. The terminal object in that slice is the identity of B. A point is a section, and the corresponding notion of connectedness says that the section is an effective epimorphism. Slicing is necessary when a group varies with parameters. A single external group of global sections need not recover an internal group object. The family of mapping spaces $\mathrm{Map}_{\mathcal{X}}(T, G)$, with its action by change of T, does recover it by the ∞-categorical Yoneda lemma.

2.2 Actions and classifying maps

A left G-object can be presented as a simplicial object $E_\bullet$ over $G_\bullet$ such that restriction to the initial vertex induces

$$E_n \xrightarrow{\sim} E_0 \times G^n.$$

The map to the final vertex at $n = 1$ supplies the action. With this orientation, successive arrows g, h act by $h \cdot (g \cdot x) = (hg) \cdot x$. Here and below formulas in elements specify maps of objects and their coherent simplicial identities. The category of actions is the full subcategory of the category of simplicial objects over $G_\bullet$ specified by the displayed pullback property. In particular its morphisms are coherent equivariant morphisms. The quotient $E_0//G$ means $|E_\bullet|$.

Theorem 2.2 (Actions as objects over a delooping)**.** *The functor sending $E_\bullet \to G_\bullet$ to $|E_\bullet| \to BG$ is an equivalence*

$$\mathrm{Act}_G(\mathcal{X}) \simeq \mathcal{X}_{/BG}.$$

Its inverse sends $Y \to BG$ to the action diagram $G_\bullet \times_{BG} Y$.

Proof. For a groupoid, cartesianness over the initial vertex implies cartesianness over every vertex. In degree one, transport to the other vertex has inverse induced by inversion. In higher degrees the Segal maps express transport along a string of arrows as successive transport in degree one. The resulting squares are cartesian for all simplex operators. The Segal pullbacks for $G_\bullet$, together with the cartesian vertex squares, then give the Segal pullbacks for $E_\bullet$. Invertibility is inherited from the base groupoid by transport. Thus $E_\bullet$ is a groupoid object. Put $Y = |E_\bullet|$. Effectivity gives $E_1 \simeq E_0 \times_Y E_0$ and $G_1 \simeq 1 \times_{BG} 1$. The square

$$\begin{array}{ccc} E_0 & \longrightarrow & 1 \\ \downarrow & & \downarrow \\ Y & \longrightarrow & BG \end{array}$$

is cartesian. To see this, test the comparison $E_0 \to Y \times_{BG} 1$ after pullback along the effective epimorphism $E_0 \to Y$. The resulting map is

$$E_0 \times_Y E_0 \longrightarrow E_0 \times_{BG} 1.$$

Under the identifications above it is exactly the degree-one vertex equivalence $E_1 \simeq E_0 \times G_1$, using the first projection to E_0. Equivalences can be tested after an effective epimorphism, as follows by reconstructing both source and target by the base changes of its Čech nerve and realizing. The vertex equivalence now gives $E_\bullet \simeq G_\bullet \times_{BG} Y$ in all degrees. In the other direction, universality of colimits implies $|G_\bullet \times_{BG} Y| \simeq Y$. All comparisons are natural in the diagrams, so they give inverse equivalences of ∞-categories. □

A G-torsor over T is an action in $\mathcal{X}_{/T}$ on an object $P \to T$ such that $P \to T$ is an effective epimorphism and the transport map

$$P \times G \longrightarrow P \times_T P, \qquad (x, g) \longmapsto (x, g \cdot x)$$

is an equivalence over T. The universal torsor has underlying object $1 \to BG$. Its action is specified by the identification of its Čech groupoid with the bar groupoid of G. Conventions using right actions give the same classifier via inversion.

Corollary 2.3 (Torsor classification)**.** *The space of G-torsors over T, with equivariant equivalences as arrows, is naturally equivalent to $\mathrm{Map}_{\mathcal{X}}(T, BG)$. Pullback in T corresponds to precomposition with the classifying map.*

Proof. The classifying equivalence of Theorem 2.2 in $\mathcal{X}_{/T}$ sends an action on P to an object $P/\!/G \to T \times BG$ over T. The comparison from its action groupoid to the Čech nerve of $P \to T$ is an equivalence in degree one by the torsor condition. The two Segal properties make it an equivalence in every degree. Since $P \to T$ is effective, this means $P/\!/G \simeq T$ over T. Conversely, if $P/\!/G \simeq T$, effectivity identifies the action groupoid with the Čech nerve of $P \to T$, giving both torsor conditions. The core of objects $Y \to T \times BG$ over T with $Y \to T$ an equivalence is $\mathrm{Map}_{\mathcal{X}}(T, BG)$, by replacing the source by T. This identification and the resulting pullback of the universal torsor commute with base change. □

This formulation retains automorphisms of a torsor and their higher homotopies. At a classifying map $c : T \to BG$, the loop space of the mapping space in Corollary 2.3 is the space of equivariant automorphisms of the classified torsor. Choosing a section of the torsor identifies it with the trivial torsor through the transport equivalence. Changing the section can change that identification. For an ordinary nonabelian group the induced change is conjugation. For an internal abelian group A with Eilenberg–Mac Lane deloopings, the same corollary applied to $G = K(A, m)$ identifies torsors with maps to $K(A, m+1)$. With the convention $H^i(\mathcal{X}, A) = \pi_0 \mathrm{Map}_{\mathcal{X}}(1, K(A, i))$, global isomorphism classes form $H^{m+1}(\mathcal{X}, A)$. The based homotopy groups of the torsor space at the trivial class recover lower degrees, because looping mapping spaces commutes with looping the target.

2.3 Transport over a space

For a space B there is a canonical equivalence

$$\mathcal{S}_{/B} \simeq \mathrm{Fun}(B, \mathcal{S}). \tag{2.1}$$

The functor associated to $Y \to B$ sends b to $Y_b = \{b\} \times_B Y$. Its maps are transport along paths, and homotopies of paths act through the higher coherences of transport. One construction uses straightening of left fibrations [2, Section 2.2.1]. Over a groupoidal base, the Grothendieck construction of a space-valued functor has only invertible arrows. Its total space is the homotopy sum $\mathrm{colim}_{b \in B} F(b)$ over $\mathrm{colim}_{b \in B} * \simeq B$. In the other direction an arbitrary map of spaces can be represented by a Kan fibration, which in particular is a left fibration. Thus the straightening correspondence restricts to the displayed equivalence.

For connected B the point b is an effective epimorphism in $\mathcal{S}$, and the space B is the delooping of $\Omega_b B$. Equation (2.1) and Theorem 2.2 identify evaluation at b with the forgetful functor from $\Omega_b B$-actions. The point $b : * \to B$ has fiber $\Omega_b B$. Under the equivalence it is the regular action, also characterized as the free action on the point. Indeed maps $* \to Y$ over B are points of Y_b, and the

Yoneda lemma in $\mathrm{Fun}(B, \mathcal{S})$ characterizes the free action by the representable functor $\mathrm{Map}_B(b, -)$.

Pullback along a map $f : C \to B$ corresponds to restriction of representations. Both adjoints to restriction exist and are computed by Kan extension. At b, left Kan extension of $F : C \to \mathcal{S}$ takes the homotopy sum of F over the groupoid of pairs $(c, f(c) \to b)$. Right Kan extension takes the limit over pairs $(c, b \to f(c))$. Inversion identifies the two indexing groupoids. The left adjoint is composition of total-space maps over the base, since giving a map from that composite is equivalent to giving a map to the pullback by the universal property of a fiber product. In the right adjoint the fiber over b is the space of sections of the restricted family over the fiber of f. In particular the two adjoints along $B \to *$ are homotopy orbits and homotopy fixed points.

If a connected space B has a fiber F at b, then after choosing an identification with a specified space K, the action can be expressed by a based map

$$B \longrightarrow B\,\mathrm{Aut}(K).$$

Here $\mathrm{Aut}(K)$ is the space of equivalences $K \to K$ with composition. To verify the classification, take the core $\mathcal{S}^{\simeq}$ in a universe containing the given fibers, choose the component containing K, and use Proposition 2.1 to identify the component with $B\,\mathrm{Aut}(K)$. Every functor from a space lands in the core, since all its arrows are invertible, and connectedness forces the selected component. Taking the homotopy fiber of evaluation at b gives the based mapping space and records the chosen identification. Forgetting the identification gives the unpointed mapping space. Taking π_0 of $\mathrm{Aut}(K)$ loses the higher transition data. For instance, let $K = S^1$. The identity component of $\mathrm{Aut}(S^1)$ is equivalent to S^1, as follows by evaluating a degree-one map at a point and noting that the degree-one component of the based mapping space is contractible. Consequently

$$\pi_2\big(B\,\mathrm{Aut}(S^1)\big) \cong \mathbb{Z}.$$

There are based families with fiber S^1 over S^2 whose classifying maps represent a nonzero integer, even though $\pi_1(S^2)$ is trivial.

2.4 The subgroup correspondence and its higher replacement

Fix a pointed connected space (B, b). A pointed intermediate object is a factorization

$$* \xrightarrow{y} Y \xrightarrow{q} B$$

of b, including an identification of the composite with b. Its morphisms are maps of such factorizations, with coherent triangles. By (2.1) the category of factorizations is the category of actions equipped with a map from the regular action, or equivalently the category of actions with a specified point of the

fiber over b. This statement is true whether or not Y is connected. It follows because taking the undercategory of a given object preserves equivalences of ∞-categories.

Proposition 2.4 (Connected intermediates)**.** *The full subcategory of pointed intermediate objects with connected Y is equivalent to the slice of group objects in spaces over $G = \Omega_b B$. Under the equivalence $Y \to B$ corresponds to $H = \Omega_y Y \to G$. Its fiber over b, with the distinguished point, is the homotopy quotient $G//H$ of the right H-action on G through $H \to G$.*

Proof. Slicing the equivalence of Proposition 2.1 gives the categorical claim. For the fiber, model the map as $BH \to BG$. Pull back the family $* \to BG$ to BH. Its fiber at the point of BH is G. Choose the path coordinate of a point in this family to run from the image of the moving point of BH to the fixed point of BG. If $\phi : H \to G$ is the homomorphism, transport in this coordinate along h sends g to $g\phi(h)^{-1}$. This is the covariant presentation of the right action $g \cdot h = g\phi(h)$. The total space is its homotopy quotient $G//H$, and the chosen point comes from the identity in G. Equivalently, this quotient is the realization obtained by pulling the augmented bar diagram of H back along $* \to BG$, so its description fixes the same map and basepoint without choosing path coordinates. □

The connectedness condition on the total space is the homotopy-theoretic transitivity condition. In a pointed factorization, the map from the regular action to the fiber is the orbit map induced by the distinguished point. It is an effective epimorphism exactly when Y is connected. In one direction the orbit map is the base change of $* \to Y$ along $Y_b \to Y$, and in the other direction effective epimorphisms descend along the effective epimorphism $Y_b \to Y$. The latter map is a base change of $* \to B$. To justify the descent assertion, factor a map into its effective epimorphism and monomorphism. The subobject given by the second factor is terminal after base change precisely if it was terminal before base change, since base change along an effective epimorphism detects equivalences. For spaces this criterion says that the regular orbit is surjective on components and places no discreteness condition on the fiber.

Theorem 2.5 (Subgroups at the covering level)**.** *The category of pointed factorizations over (B, b) with Y connected and $q : Y \to B$ 0-truncated is equivalent to the poset of subgroups of $\pi_1(B, b)$, ordered by inclusion.*

Proof. Under (2.1), the condition on q says every value of the functor is discrete. Such functors factor through the ordinary fundamental groupoid of B. This follows from the universal property of the homotopy category, since the category of sets is ordinary and a functor on a groupoid takes its arrows to bijections. Choosing the object b of the connected fundamental groupoid gives a $\pi_1(B, b)$-action on the fiber. The distinguished point x and connectedness make the

action transitive, by the orbit criterion above. The stabilizer K of x is a subgroup, and the pointed action is isomorphic to the left action on $\pi_1(B,b)/K$ via $gK \mapsto g \cdot x$. A pointed equivariant map exists from the quotient by K to the quotient by L precisely when $K \subseteq L$, and is then uniquely determined by the displayed formula. Mapping spaces between the pointed actions are discrete because maps between discrete fibers form discrete spaces and naturality is imposed by a limit of such spaces. The result is the stated equivalence of categories. □

This is the regular-quotient argument of [1, Proposition 2.4.1] at the covering level. The opposite category reverses its containment order, as for intermediate fields. Subgroups up to conjugacy classify connected covers up to equivalence when the point in each intermediate cover is forgotten. For example, the equivariant endomorphisms of the transitive set $\pi_1(B,b)/K$ correspond to the cosets fixed under K. The group of invertible endomorphisms is $N_{\pi_1(B,b)}(K)/K$, up to inversion in the convention for right translation.

The loop homomorphism for a connected covering is the inclusion of a union of components of G. Indeed its fiberwise description as the stabilizer is the pullback of the orbit map $G \to Y_b$ along the point x, and Y_b is discrete. For a general intermediate object the same pullback is a homotopy fiber with a chosen nullhomotopy of motion of x. Such a fiber need not be a union of components. Take $B = *$ and $Y = K(\mathbb{Z},2)$. The map of loop groups is $K(\mathbb{Z},1) \to *$, whose only homotopy fiber is neither empty nor contractible. Hence it fails to be a monomorphism in spaces. Proposition 2.4 accommodates this data by allowing arbitrary homomorphisms into G.

Corollary 2.6 (Relative truncation of the stabilizer)**.** *For $n \geq 0$, a map $BH \to BG$ of pointed connected spaces is n-truncated if and only if $H \to G$ is $(n-1)$-truncated.*

Proof. Set $F = BH \times_{BG} *$, pointed by the chosen points. The homotopy fiber of $H \to G$ over the identity is ΩF, since based loops preserve limits. Every fiber of $H \to G$ over a point in its image on components is equivalent to the identity fiber by translation, and all other fibers are empty. The degree-one orbit argument also shows that $G \to F$ meets every component of F, and the left G-action on F gives equivalences carrying the distinguished point into each component. Thus F is n-truncated if and only if its loop space at that point is $(n-1)$-truncated. The fibers of $BH \to BG$ are all equivalent to F since BG is connected. Combining these observations proves the claim. □

For $n = 0$, this identifies the covering condition with monomorphisms of loop groups. For $n = 1$, the stabilizer map can have nontrivial discrete homotopy fibers, as for $Y = B\mathbb{Z}$ over $B = *$. A classification by subgroup inclusions already fails at this level. The full slice of group objects records how the

stabilizer maps into the ambient group, including the coherent data through which it acts.

2.5 Regular intermediates and delooped quotients

For a pointed connected cover classified by a subgroup $K \subseteq \pi_1(B, b)$, the deck transformations act transitively on the fiber precisely when K is normal. Indeed the orbit of the distinguished coset under deck transformations is indexed by $N_{\pi_1(B,b)}(K)/K$, so it is all of $\pi_1(B, b)/K$ precisely when the normalizer is all of $\pi_1(B, b)$. In that case the quotient is itself a group and the cover is the pullback of

$$* \longrightarrow B(\pi_1(B, b)/K)$$

along the map $B \to B(\pi_1(B, b)/K)$ induced by the quotient homomorphism. The path space construction gives left transport by the quotient group, and the right deck transformations commute with it.

At higher truncation levels the homotopy quotient $G//H$ is always defined. It can have higher homotopy groups even when G is discrete. The existence of a delooped quotient is a separate requirement and can be expressed in terms of the pointed action already obtained from the factorization.

Proposition 2.7 (Criterion for a delooped quotient)**.** *Let $i : H \to G$ be a map of group objects in spaces and let Q be another group object. For a homomorphism $\phi : G \to Q$, there is a pointed equivalence over BG*

$$BH \simeq BG \times_{BQ} *$$

if and only if $G//H$, as a pointed left G-space, is equivalent to the regular left Q-space restricted along ϕ. In the pointed actions the specified points are those induced by the identities.

Proof. The object $* \to BQ$ corresponds to the regular action by (2.1). Pullback along $B\phi$ restricts the action. Proposition 2.4 computes the fiber of $BH \to BG$ as $G//H$. Transport from the fixed base along loops of BG acts on this fiber by left multiplication. This action commutes coherently with the right H-action used to form the quotient. Specifying a point in the fiber is the same as giving a pointed map from the regular action of G. Hence (2.1), applied also to the undercategory of the point, gives the stated equivalence of conditions. □

For a solution ϕ, the action on components of Q is transitive because BH is connected. Thus $\pi_0 G \to \pi_0 Q$ is surjective. This condition is automatic for the ordinary quotient by a normal subgroup, but the full equivalence of pointed actions in Proposition 2.7 probes every homotopy group of the fiber and its transport.

CHAPTER 3

Effective descent and Čech actions

The split-algebra equivalence in [1, Chapter 2] classifies objects by transport among their trivializations. Later chapters of that book make the descent morphism explicit and replace a single group by an internal groupoid. In a higher topos this construction starts with the entire Čech nerve of an effective epimorphism. Each composable string of changes of trivialization comes with its own coherence, and these coherences are part of the gluing datum.

3.1 Descent data and their realization

For a map $q : U \to B$, put $U_n = U^{\times_B (n+1)}$. The cosimplicial diagram of slices

$$[n] \longmapsto \mathcal{X}_{/U_n}$$

acts on a simplex operator by pullback. Define $\mathrm{Desc}(q)$ to be its limit in ∞-categories. This limit can be described as the category of simplicial objects $V_\bullet \to U_\bullet$ such that for every map $[m] \to [n]$ in the simplex category the square

$$\begin{array}{ccc} V_n & \longrightarrow & V_m \\ \downarrow & & \downarrow \\ U_n & \longrightarrow & U_m \end{array}$$

is cartesian. Indeed the straightening of a diagram taking an object A to $\mathcal{X}_{/A}$ is the target projection of the arrow category with its cartesian morphisms the pullback squares. A limit of such a diagram classifies its cartesian sections. In this description the two vertex identifications in degree one give an equivalence of the two pullbacks of V_0 to U_1. The higher simplices specify coherent composition of those equivalences.

Theorem 3.1 (Effective descent)**.** *For an effective epimorphism $q : U \to B$ in an ∞-topos, the comparison*

$$\mathcal{X}_{/B} \longrightarrow \mathrm{Desc}(q), \qquad E \longmapsto (E \times_B U_n)_{n \geq 0},$$

is an equivalence. Its inverse is geometric realization over B.

Proof. The realization of the displayed diagram for E is E, by universality of colimits and $|U_\bullet| \simeq B$. For the other composite let $V_\bullet$ be a descent datum and put $Y = |V_\bullet|$. First, $V_\bullet$ is a groupoid object. The Segal squares can be checked

using a vertex at the beginning of a string together with cartesian transport to each next vertex. A composable string of arrows from V_0 then consists of its initial point and the underlying composable string in $U_\bullet$. Transport along the inverse of a base arrow gives the inverse arrow of $V_\bullet$, again by the cartesian vertex condition.

By effectivity $V_0 \to Y$ is an effective epimorphism and $V_1 \simeq V_0 \times_Y V_0$. There is a canonical comparison $V_0 \to Y \times_B U$ over Y. After pullback along $V_0 \to Y$, this comparison becomes

$$V_0 \times_Y V_0 \longrightarrow V_0 \times_B U.$$

The source is V_1, and the target is $V_0 \times_U U_1$ where $U_1 \to U$ uses the initial vertex. The comparison is therefore the cartesian vertex equivalence of the datum. To descend that equivalence, pull it and both its domain and codomain to the entire Čech nerve of $V_0 \to Y$. The comparison is an equivalence in every such degree. Each original object is recovered by realizing its pullback diagram, by universality of colimits. Hence $V_0 \simeq Y \times_B U$. The vertex comparison in each degree identifies $V_n \simeq Y \times_B U_n$, compatibly with all simplex operators. This gives the natural equivalences of the two composites. □

This gives the underlying mechanism of the action equivalence of Theorem 2.2. When $U = 1$ and $B = BG$, the Čech nerve is the bar groupoid of G. At the other extreme, there may be no distinguished global point of U or B. The internal groupoid then describes transitions between varying objects of objects. Replacing it by the loops at a chosen object would require extra hypotheses that are absent for most covers in a topos.

Corollary 3.2 (What a general map sees)**.** *For any $q : U \to B$, let $i : I \hookrightarrow B$ denote its image and $p : U \to I$ the associated effective epimorphism. There is a canonical equivalence $\mathrm{Desc}(q) \simeq \mathcal{X}_{/I}$, under which comparison from $\mathcal{X}_{/B}$ is restriction along i. Comparison is an equivalence if and only if q is an effective epimorphism.*

Proof. The
vCech diagrams for p and q coincide because i is a monomorphism. More explicitly $I \times_B I \simeq I$, so the comparison $U \times_I U \to U \times_B U$ is an equivalence, and the same holds for every iterated pullback. Theorem 3.1 applied to p gives the conclusion about restriction. If i^* is an equivalence, it reflects equivalences and carries $(I \to B) \to (B \to B)$ to an equivalence. Thus i is an equivalence. □

3.2 Mapping spaces and properties local on the base

Descent reconstructs morphisms as well as objects. For $E, F \to B$ and q effective, it identifies

$$\mathrm{Map}_{\mathcal{X}_{/B}}(E, F) \simeq \lim_{[n]\in\Delta} \mathrm{Map}_{\mathcal{X}_{/U_n}}(E \times_B U_n, F \times_B U_n). \tag{3.1}$$

To check this directly, realize the augmented diagram $E \times_B U_\bullet \to E$ in the slice over B. Mapping into F turns its colimit into a limit of mapping spaces. The adjunction between composition and pullback along $U_n \to B$ identifies each term of that limit with the corresponding term on the right of (3.1). An element of the right-hand side specifies a local map together with compatibility homotopies over intersections and higher coherent comparisons on all iterated intersections.

Let $\mathcal{D}(B) \subseteq \mathcal{X}_{/B}$ be a full subcategory closed under equivalences, for each B, and assume pullback takes objects of $\mathcal{D}$ to objects of $\mathcal{D}$. Call membership local on the base if $E \to B$ belongs to $\mathcal{D}(B)$ whenever its pullback belongs to $\mathcal{D}(U)$ along an effective epimorphism $U \to B$.

Proposition 3.3 (Descent for full subcategories)**.** *If membership in $\mathcal{D}$ is local on the base, then for every effective epimorphism $q : U \to B$,*

$$\mathcal{D}(B) \simeq \lim_{[n]\in\Delta} \mathcal{D}(U_n).$$

Conversely, if this comparison is an equivalence for all effective epimorphisms, membership is local on the base.

Proof. The proposed target is a full subcategory of $\mathrm{Desc}(q)$, since limits preserve full inclusions of ∞-categories. A cartesian section lies in it exactly when its degree-zero object is in $\mathcal{D}(U)$, since the other entries are equivalent to pullbacks along a vertex. Theorem 3.1 identifies such sections with the objects over B whose pullback along q belongs to $\mathcal{D}(U)$. This proves both directions. □

For example, the class of n-truncated maps is local on the base for each integer $n \geq -2$. The case $n = -2$ is the conservativity of pullback along an effective epimorphism. For the induction step, a map $E \to B$ is n-truncated precisely if its diagonal $E \to E \times_B E$ is $(n-1)$-truncated. Forming the diagonal commutes with base change. The cover on its codomain after pullback along $U \to B$ is effective, so the induction hypothesis applies. In particular this yields descent for monomorphisms and for covering objects with discrete fibers. Effective epimorphisms themselves are local, by applying the equivalence test to the image inclusion as in the orbit argument of Chapter 2. The locality proof uses finite-level properties and is valid in a nonhypercomplete topos.

3.3 The descent monad

Effective descent also has a monadic formulation. The functor $q^* : \mathcal{X}_{/B} \to \mathcal{X}_{/U}$ has left adjoint $q_!$ given by composition with q. It preserves colimits by universality and detects equivalences by the Čech argument. The ∞-categorical Barr–Beck theorem then gives

$$\mathcal{X}_{/B} \simeq \mathrm{Alg}_T(\mathcal{X}_{/U}), \qquad T = q^* q_!.$$

In the language of the Barr–Beck hypotheses, preserving all colimits implies preservation of realizations of those simplicial objects that become split after q^* [6, Theorem 4.7.3.5]. This proves monadicity without assuming that the underlying morphism q has a section.

The monad can be expressed in terms of the two projections $s, t : U_1 \rightrightarrows U$, with s the source and t the target. For $A \to U$,

$$T(A) \simeq t_! s^* A \simeq U \times_B A.$$

In the last expression the first coordinate is the target in U. The unit is induced by the diagonal of U over B. The multiplication takes (u_2, u_1, a) to (u_2, a) with the prescribed composition of paths in B. Its simplicial powers are indexed by strings in the Čech groupoid. For $A = E \times_B U$, the monad action identifies the point e over the preceding vertex with the same point over the new vertex. Thus the map $T(A) \to A$ is the transport correspondence from the canonical descent datum. The monadic comparison and the Čech comparison agree on every pullback from B, and both comparisons are equivalences. This relates their coherent algebra and coherent transport descriptions without an extra hypothesis on q.

As with the classical treatment in [1, Section 5.1], the monad lives on a slice over the splitting object. If a geometric problem begins with algebras and replaces them by opposite affine objects, geometric pullback usually becomes extension of scalars. Passing to opposites converts the preceding monadic expression into a comonadic one with the corresponding variance of the adjunction. In a general algebraic category, extension of scalars still must be proved to satisfy a descent theorem. The argument in this chapter uses groupoid effectivity and universal colimits in a specific ambient topos, which are the hypotheses making arbitrary effective epimorphisms work here.

3.4 Refinements and changes of groupoid presentation

Every effective internal groupoid presents a base. Let $D_\bullet$ be an internal groupoid, let $B = |D_\bullet|$, and apply Theorem 3.1 to $D_0 \to B$. This identifies cartesian objects over $D_\bullet$ with $\mathcal{X}_{/B}$. A map of groupoids induces a pullback functor between their cartesian action categories, and under this equivalence it is pullback along the map of realizations. In fact, pulling back $E \times_{|D_\bullet|} D_\bullet$ along the map of groupoids gives its pullback from E degree by degree. Refinement of an effective cover is a particular case, in which the map of realizations is the identity. Restricting a datum to a refinement changes the representatives for its local objects and their transition maps while preserving its realized object.

The following criterion gives a specific test for changing presentations. For a map $F_\bullet \to D_\bullet$ of internal groupoids, call it fully faithful if the comparison

$$F_1 \longrightarrow (F_0 \times F_0) \times_{D_0 \times D_0} D_1$$

is an equivalence. Call it essentially surjective if the map

$$F_0 \times_{D_0,s} D_1 \xrightarrow{t} D_0$$

is an effective epimorphism.

Proposition 3.4 (Change of presentation)**.** *If a map of internal groupoids is fully faithful and essentially surjective, it induces an equivalence of realizations and an equivalence of cartesian action categories. Conversely, a map inducing an equivalence of realizations is fully faithful and essentially surjective.*

Proof. Write $A = |F_\bullet|$ and $B = |D_\bullet|$. Effectivity gives

$$(F_0 \times F_0) \times_{D_0 \times D_0} D_1 \simeq F_0 \times_B F_0.$$

Full faithfulness thus identifies the degree-one terms of the Čech nerves of $F_0 \to A$ and $F_0 \to B$, compatibly with both vertices. By the Segal identities their entire nerves agree. Their realizations identify A with the image of $F_0 \to B$. On the other hand $F_0 \times_{D_0,s} D_1 \simeq F_0 \times_B D_0$, so essential surjectivity is equivalent to $F_0 \to B$ being effective after base change along $D_0 \to B$. Effectivity descends along that cover. The image is therefore B. If $A \to B$ is an equivalence, the same formulas give full faithfulness, and essential surjectivity follows by base change from the effective map $F_0 \to A$. The result for actions follows from the descent comparisons. □

An equivalence here allows the objects of objects to change. In later applications one can replace a groupoid associated to a cover by a more convenient presentation as long as the change of realization, or the two conditions in the proposition, can be verified.

3.5 Fixed split type and the torsor of frames

Fix $K \in \mathcal{X}$. An object $E \to B$ is a form of K if there exists an effective epimorphism $U \to B$ and an equivalence $E \times_B U \simeq U \times K$ over U. For a specified effective epimorphism q, the objects it splits with this type correspond in $\mathrm{Desc}(q)$ to those data with degree-zero entry equivalent to $U \times K$. The category of these objects, with all maps over B, is full by Theorem 3.1. To compute the space of forms we restrict to equivalences of total objects.

The group object $G = \underline{\mathrm{Aut}}(K)$ is the object of self-equivalences with its coherent composition. Concretely, a section of G over W is an equivalence $W \times K \to W \times K$ over W, and these spaces of sections are represented by the equivalence part of the internal hom. For a form E, its frame object over B represents equivalences $W \times K \to E \times_B W$ over W, for test objects $W \to B$. Local cartesian closedness gives the internal hom over B, and the equivalence locus can be represented by requiring a two-sided inverse. More precisely, the

object of maps with left and right inverses and witnesses for each gives this locus. At a map that is an equivalence the space of choices of each witness together with its inverse is contractible, by cancellation along an equivalence. At any other map at least one of the two spaces of choices is empty.

Proposition 3.5 (Classification of forms)**.** *The frame construction identifies the space of forms of K over B with the space of right G-torsors over B, hence with* $\mathrm{Map}_{\mathcal{X}}(B, BG)$.

Proof. Precomposition equips the frame object $\mathrm{Fr}(E)$ with a right G-action. A choice of local equivalence with $U \times K$ gives a section over U and identifies the action there with the regular right action. The torsor comparison

$$\mathrm{Fr}(E) \times G \longrightarrow \mathrm{Fr}(E) \times_B \mathrm{Fr}(E)$$

is an equivalence after pulling back along $U \to B$, and is therefore an equivalence globally. The frame map to B is effective since it has a section over U. Conversely, for a right G-torsor P form $P \times^G K$ by realizing the diagonal action groupoid, where g acts by $(p, x) \mapsto (pg^{-1}, gx)$. Over the effective epimorphism $P \to B$ the pulled-back torsor has its tautological section. The associated object then identifies with the product with K. The map from P to the frames of this associated object sends a local point p to $x \mapsto [p, x]$. It is equivariant by the relation $[pg, x] = [p, gx]$, and becomes an equivalence where P is trivialized, so it is an equivalence. The evaluation of a frame on K gives $\mathrm{Fr}(E) \times^G K \to E$, which becomes the product equivalence on the same trivializing cover and is thus an equivalence. Both maps are natural, giving inverse equivalences on cores. Corollary 2.3 gives the mapping-space description. □

If a finite-level truncation of K is imposed, its automorphism object also has bounded truncation degree, because mapping into an n-truncated object gives n-truncated mapping spaces. The classification then uses a correspondingly truncated target BG. For general K, it involves coherence at every degree. A section of the frame torsor is a chosen trivialization, and two sections may differ by a nontrivial automorphism. This distinguishes a property, being split, from the moduli of its witnesses.

3.6 The boundary between Čech descent and hyperdescent

A hypercover $V_\bullet \to B$ is an augmented simplicial object whose relative matching maps are effective epimorphisms, including $V_0 \to B$. The Čech nerve of an effective epimorphism is a hypercover, with the matching maps in positive degrees equivalences. A general hypercover may have additional local choices in higher degrees. It need not be a groupoid, so the effectivity argument for Theorem 3.1 does not prove descent along it.

For an object E of $\mathcal{X}$, hypercompleteness means that mapping into E inverts all infinitely connective maps. The hypercover criterion identifies this with the requirement that the augmentation induce an equivalence

$$\mathrm{Map}_{\mathcal{X}}(B, E) \longrightarrow \lim_{[n]\in\Delta} \mathrm{Map}_{\mathcal{X}}(V_n, E)$$

for all hypercovers $V_\bullet \to B$ [2, Theorem 6.5.3.12]. This uses the fact that $|V_\bullet| \to B$ is infinitely connective. In a hypercomplete topos it is an equivalence. A truncated target E also satisfies the displayed condition in an arbitrary ∞-topos, because maps into a target of fixed finite truncation cannot detect an infinitely connective discrepancy. The distinction matters when local trivializations or coefficients are given as unbounded towers. Categorical descent for one effective epimorphism here avoids any completeness hypothesis because the augmentation of that particular nerve is already an equivalence.

Part II

Monodromy and shape

CHAPTER 4

Pro-spaces and continuous monodromy

The passage from a single splitting field to a directed family of splitting fields in [1, Ch. 3] changes the meaning of an action. At each finite stage the Galois group acts discretely. An orbit in the limit need only be controlled by a quotient large enough to see that orbit. A continuous action on an entire space can require a more restrictive or a more flexible definition, depending on whether one demands a single trivializing stage for the object or permits a filtered assembly of objects with that property. This chapter distinguishes those constructions before passing to the shape of a topos.

4.1 Pro-objects and the role of the test space

Write $\mathcal{S}$ for the ∞-category of spaces. A pro-space P is represented by a small cofiltered diagram $i \mapsto P_i$. For a second pro-space $Q = \{Q_j\}_{j \in J}$, the mapping formula is

$$\mathrm{Map}_{\mathrm{Pro}(\mathcal{S})}(P, Q) \simeq \lim_{j \in J} \operatorname*{colim}_{i \in I^{\mathrm{op}}} \mathrm{Map}_{\mathcal{S}}(P_i, Q_j). \tag{4.1}$$

We place colim_i over the filtered direction throughout this chapter. Formula (4.1) comes from the embedding of $\mathrm{Pro}(\mathcal{S})^{\mathrm{op}}$ in $\mathrm{Fun}(\mathcal{S}, \mathcal{S})$ taking P to

$$F_P(T) = \operatorname*{colim}_i \mathrm{Map}(P_i, T).$$

Indeed the representable $\mathrm{Map}(Q_j, -)$ is a term in the filtered presentation of F_Q. A natural transformation from F_Q to F_P is a compatible family of transformations from those representables. The Yoneda lemma identifies the family with the right side of (4.1). In particular, an equivalence of pro-spaces can be checked against every constant test T. Here and below an ambient size enlargement may be used for the categories of test functors.

The realized inverse limit $\lim_i P_i$ loses this test property. Take $P_m = \{m, m+1, \dots\}$ with the inclusions as transition maps. As discrete spaces their inverse limit is empty. Yet $F_P(\varnothing)$ is empty because every P_m is nonempty. The constant empty pro-space has $F_{\varnothing}(\varnothing) \simeq *$. Thus even this tower of discrete spaces cannot be replaced by its realized inverse limit for a classification problem.

Definition 4.1. A space is *π-finite* if it is truncated, has finitely many components, and has finite positive homotopy groups at every basepoint. Let $\mathcal{S}_\pi$ be the

full subcategory on these spaces, and let $\mathcal{S}_{<\infty}$ be the full subcategory of all truncated spaces. The corresponding completions of P are written $P^{\wedge}_{\pi} \in \mathrm{Pro}(\mathcal{S}_{\pi})$ and $P_{<\infty} \in \mathrm{Pro}(\mathcal{S}_{<\infty})$.

These completions are characterized by retaining exactly the tests of the indicated kind. In the functor presentation, restrict F_P to the indicated full subcategory. Each subcategory admits finite limits. In the π-finite case, this follows from the homotopy exact sequence for a pullback and the finiteness of the component sets. Both restrictions preserve finite limits because filtered colimits of spaces do so. The pro-object characterization by accessible finite-limit-preserving functors then gives the completions. For truncated spaces the accessibility issue can be handled by the cofiltered diagram $\{\tau_{\le d} P_i\}_{i,d}$, since maps to any fixed truncated test stabilize in the Postnikov coordinate. A particular consequence that will be used repeatedly is

$$\mathrm{Map}(P^{\wedge}_{\pi}, T) \simeq \mathrm{Map}(P, T) \qquad (T \in \mathcal{S}_{\pi}). \tag{4.2}$$

There is an analogous formula for $P_{<\infty}$ with $T \in \mathcal{S}_{<\infty}$. Either construction is left adjoint to the inclusion of the corresponding pro-category. This follows for a general target by taking the limit of the equivalences for its constant stages. The same argument in degree 0 recovers the equivalence between pro-finite sets and Stone spaces used in [1, Ch. 3]. The higher completion tests Eilenberg–Mac Lane spaces, classifying spaces, and their finite Postnikov combinations.

4.2 A representation category at stages

For a coefficient ∞-category $\mathcal{D}$, set

$$\mathrm{Rep}_{\mathcal{D}}(P) = \operatorname*{colim}_{i} \mathrm{Fun}(P_i, \mathcal{D}), \tag{4.3}$$

where the colimit is in Cat_{∞}, and the arrows are given by restriction. A representative of an object is a family and all its transport coherence at a single stage. A representative of a morphism is a natural transformation after passing to a common stage. More precisely, if F and G are already represented at a common stage i, their mapping space is

$$\mathrm{Map}_{\mathrm{Rep}_{\mathcal{D}}(P)}(F, G) \simeq \operatorname*{colim}_{(k \to i)} \mathrm{Map}_{\mathrm{Fun}(P_k, \mathcal{D})}(F|_{P_k}, G|_{P_k}). \tag{4.4}$$

This follows by computing mapping spaces as fibers of the core of the arrow category over the core of pairs of objects. Filtered colimits of ∞-categories commute with forming these cores and arrow categories, and filtered colimits of spaces commute with the fibers. The reason for the first commutation can also be seen in simplicial models: every prescribed simplex and every witness that an edge is invertible involves a finite diagram, so it is already present at

a single refinement. Two stages that represent a prescribed datum become comparable after a further refinement. This proves invariance under cofiltered change of presentation and recovers the functorial construction of [3, §2].

The coefficient category may have infinitely many objects up to equivalence. A finite stage can still carry an arbitrary object of $\mathcal{D}$. The crucial finiteness for completion is instead the size and truncation of the spaces classifying possible transport and coherent maps.

Lemma 4.2. *If A and B are π-finite, then $\mathrm{Map}(A, B)$ is π-finite. The component of A in the core of spaces is $B\,\mathrm{Aut}(A)$ and is π-finite.*

Proof. A π-finite space has a simplicial model with finitely many simplices in every degree. For a connected 1-type use the nerve of its finite fundamental group. At a higher Postnikov step, the fiber $K(M, n)$ for a finite abelian group M has a simplicial abelian model with finitely many simplices in each degree. Represent the step by its twisted k-invariant into the $K(M, n+1)$-fibration over the nerve of the finite fundamental group. This target fibration can be modeled by the Borel construction for that group acting on $K(M, n+1)$. Its total space and the fiberwise path fibration are finite in each simplicial degree. Pulling the path fibration back along the twisted invariant gives a degreewise finite model for the extension. This inductive construction and a finite disjoint union give the asserted model.

If B is n-truncated with $n \geq 0$, the inclusion of the $(n+1)$-skeleton of this model into A induces an equivalence on maps to B, since the cells subsequently attached have dimension at least $n+2$. The skeleton is a finite cell complex. Mapping from a cell attachment into B forms a finite homotopy pullback of mapping spaces, starting with B and a point. Closure under finite homotopy limits gives the claim. A (-1)-truncated target is either empty or contractible and the conclusion holds directly. The equivalences $A \to A$ form a union of components of $\mathrm{Map}(A, A)$, hence a π-finite group-like space. Its delooping is connected and shifts its homotopy groups by one degree, proving the second assertion. □

To specify the completion comparison for all morphisms, let $\mathrm{Rep}^{\mathrm{fr}}_{\mathcal{S}_\pi}(P)$ consist of representations admitting a representative at a stage with a finite range of equivalence classes of fibers. Here fr refers to the range of isomorphism types, and the superscript imposes a condition on objects only. Thus all maps in (4.4) between those objects are retained.

Proposition 4.3. *Precomposition with $P \to P^\wedge_\pi$ gives an equivalence*

$$\mathrm{Rep}_{\mathcal{S}_\pi}(P^\wedge_\pi) \simeq \mathrm{Rep}^{\mathrm{fr}}_{\mathcal{S}_\pi}(P).$$

When P admits a presentation by connected spaces, the finite-range condition on the right is automatic.

Proof. A representation over a π-finite stage has finite range because the stage has finitely many components. Conversely, for a finite list D of equivalence classes of fibers, let C_D be the union of those components in the core of $\mathcal{S}_\pi$. It is π-finite by Lemma 4.2. The moduli space of stage objects with range in D is $\mathrm{Map}(P, C_D)$, because a functor from a space factors through the core of its target. Thus every object lifts through the completion by (4.2).

For the map space between two such objects choose finite lists D_0, D_1 for the source and target. Form the core of the arrow category and restrict its endpoints to $C_{D_0} \times C_{D_1}$. The resulting space E has a projection to that product with fiber $\mathrm{Map}(A, B)$ at (A, B). Each fiber and the base are π-finite. A uniform truncation bound for the fibers exists because the lists are finite. The homotopy exact sequence shows that the total space is also π-finite. The space of morphisms for stage representations with given endpoints is the fiber of

$$\mathrm{Map}(P, E) \longrightarrow \mathrm{Map}(P, C_{D_0} \times C_{D_1}).$$

This fiber description follows at each stage and passes through the filtered colimit. Both its source and target are preserved under (4.2), which proves full faithfulness. A connected stage maps to one component of the core, proving the final assertion. □

For example, take the constant pro-space $\mathbb{N}$, with $\mathbb{N}$ discrete. A family of finite discrete spaces with fiber cardinality $m + 1$ over $m \in \mathbb{N}$ belongs to $\mathrm{Rep}_{\mathcal{S}_\pi}(\mathbb{N})$. It cannot descend from its profinite completion by Proposition 4.3. Each individual fiber satisfies the requisite finiteness condition. The failure comes from the infinitely many equivalence classes occurring in a single family.

4.3 From one-stage actions to continuous discrete actions

Let $H = \lim_i H/U_i$ be a profinite group, where the U_i form the directed system of open normal subgroups, and set

$$B_{\mathrm{pro}}H = \{B(H/U_i)\}_i.$$

A left action on a discrete set M is continuous exactly when the stabilizer of each element is open. One direction follows by taking the preimage of the singleton under each orbit map. For the converse, the action map restricted to $H \times \{m\}$ is locally constant for each m, and these pieces are open in $H \times M$.

The stage category $\mathrm{Rep}_{\mathrm{Set}}(B_{\mathrm{pro}}H)$ realizes the full subcategory of discrete actions with a common open normal subgroup acting trivially. An action at one stage has that property. Conversely the common subgroup allows the action to be written over a quotient. If two actions come from H/U_i and H/U_j, then a function between them is H-equivariant exactly when it is equivariant for $H/(U_i \cap U_j)$, using surjectivity of the quotient map. Therefore passing to stages

identifies their mapping sets. Every finite continuous H-set belongs to this subcategory, as a finite intersection of open stabilizers is open.

Every continuous discrete H-set M is a filtered union of finite continuous H-sets. Every orbit is finite by the compactness of H and the openness of its stabilizer, and one can take the finite unions of orbits. A finite H-set is compact among continuous discrete H-sets. To see this, transport a map into a filtered colimit to some stage as a function on its finite underlying set. Choose an open normal subgroup fixing all elements of the source and their chosen representatives at that stage. Equivariance can now be checked against the finitely many elements of the quotient by that subgroup. The required equalities hold after some further refinement, by the description of filtered colimits of sets. The same argument with finitely many pointwise equalities proves uniqueness in the colimit of mapping sets. Thus the category of all continuous discrete actions is the ind-completion of its finite subcategory.

The difference between stage representations and this ind-completion already occurs for $H = \mathbb{Z}_p$. On the disjoint union $\coprod_{m\geq 1} \mathbb{Z}/p^m\mathbb{Z}$, the translation action is continuous and requires arbitrarily deep quotients. Its kernel is trivial, which is a closed subgroup that is not open. This object is built from stage actions by taking a filtered union, but no single stage represents it.

A different construction first discards the topology and takes $B(H^\delta)$, the delooping of the underlying discrete group. For $H = \prod_{m\geq 1} \mathbb{F}_2$, there is an abstract homomorphism $H \to \mathbb{F}_2$ that annihilates $\bigoplus_{m\geq 1} \mathbb{F}_2$ and takes the all-ones sequence to 1. Indeed, that sequence defines a nonzero vector in the quotient by the direct sum. Extend this vector to an algebraic basis and project to its coordinate. The resulting homomorphism is discontinuous because its kernel is dense and proper. Its action by translations on $\mathbb{F}_2$ is an action of H^δ with no continuous counterpart. Thus using the abstract classifying space instead of the pro-classifying space can change even the finite-coefficient theory.

4.4 Closed stabilizers and the infinite correspondence

The compact Hausdorff replacement of a finite discrete transitive action is a transitive action on a profinite space. For every closed subgroup $L \leq H$, there is a homeomorphism of H-spaces

$$H/L \longrightarrow \lim_i (H/U_i)/(LU_i/U_i). \tag{4.5}$$

Here H/L denotes the left H-space of cosets hL, with the quotient topology. To prove injectivity use $L = \bigcap_i LU_i$. Indeed, if $h \notin L$, some coset neighborhood hU_i is disjoint from L, by closedness and the neighborhood basis at the identity. A compatible point of the inverse limit specifies a family of closed cosets of the groups LU_i. Any finite subfamily intersects by refining the indices. Compactness gives a point in their intersection, proving surjectivity. A continuous bijection from a compact space to a Hausdorff space is a homeomorphism.

Conversely, given a transitive continuous action on a nonempty profinite space Y and a point $y \in Y$, its stabilizer L is closed as the inverse image of $\{y\}$ under the orbit map. The induced continuous bijection $H/L \to Y$ is again a homeomorphism. Between pointed homogeneous spaces $(H/L, L)$ and $(H/M, M)$ there is a unique pointed equivariant map precisely when $L \subseteq M$. In fact its value at hL is forced to be hM, and the subgroup condition makes this formula well defined and continuous. An unpointed equivariant map can send L to any coset hM with $L \subseteq hMh^{-1}$; the map is then determined by its value there. In the discrete continuous theory, transitive objects instead correspond to open stabilizers.

For a Galois field extension $K \subseteq L$, the ordinary subextension correspondence uses closed subgroups of $H = \mathrm{Gal}(L/K)$. The equality of a subgroup with the limit of its finite images in (4.5) shows why closedness is the condition visible to all finite subextensions. In greater algebraic generality the stages of the classical correspondence can likewise assemble into profinite H-spaces, as in [1, §3.5]. Continuous actions on profinite spaces and continuous actions on discrete sets therefore enter at different points of infinitary Galois theory. Replacing both by one category of H^{δ}-sets identifies constructions that have different stabilizers and different colimits.

4.5 Continuous cochains and twisted coefficients

For a discrete continuous H-module M, consider the cochain groups

$$C^q_{\mathrm{cts}}(H, M) = \mathrm{Cont}(H^q, M) \qquad (q \geq 0)$$

with the group-cohomology differential. For a fixed cochain the image is finite, since a compact space has finite image in a discrete space. Moreover continuity implies that the cochain factors as a function through some $(H/U_i)^q$. Pulling back its finitely many fibers gives a clopen partition of H^q. Each point has an open product of cosets contained in its partition member. A finite subcover by such products exists, and refining all the coset neighborhoods to one U_i proves the factorization. After shrinking U_i further, it fixes every element of the image. Thus the cochain corresponds to an element of $C^q(H/U_i, M^{U_i})$. In degree 0 the identical conclusion holds by openness of the stabilizer of the single value. The differentials respect refinement of the quotient and inclusion of the fixed submodules, giving an isomorphism of complexes

$$C^{\bullet}_{\mathrm{cts}}(H, M) \cong \operatorname*{colim}_i C^{\bullet}(H/U_i, M^{U_i}).$$

Filtered colimits are exact in abelian groups. Taking cohomology therefore gives

$$H^q_{\mathrm{cts}}(H, M) \cong \operatorname*{colim}_i H^q(H/U_i, M^{U_i}). \tag{4.6}$$

If the action on M factors through a fixed stage, the subgroups contained in its kernel form a cofinal subsystem and one can use M throughout on the right.

Let M be finite with continuous action, and take $q \geq 1$. The stage module actions induce families of pointed spaces $K(M, q)$ over $B(H/U_i)$ for all sufficiently deep i. The mapping space of sections at a stage is the homotopy fixed point space of $K(M, q)$, modeled by totalizing the usual bar cosimplicial object. Its set of components is $H^q(H/U_i, M)$. Taking components commutes with filtered colimits, so the stage mapping space has components $H^q_{\mathrm{cts}}(H, M)$ by (4.6). This identifies pro-monodromy with twisted cohomology directly from the stage dynamics. An untwisted test of the form $\mathrm{Map}(P, K(M, q))$ only computes the version with trivial action. In particular, truncation of an underlying homotopy type does not justify erasing the transport of its coefficient groups.

The three continuous constructions above use the same inverse system in different ways. Discrete ind-objects permit each orbit to descend separately. Profinite spaces permit closed stabilizers and keep the successive quotients of an orbit. Stagewise higher representations retain the full space of local coefficient transports at every stage, with refinement identifying equivalent presentations. The next chapter explains how a single ∞-topos supplies the pro-space that indexes such stages.

CHAPTER 5

The shape of a topos and its local systems

Covering theory in [1, Ch. 6] extracts a fundamental group from a category of coverings and a fiber functor. A higher topos contains families with fibers of all homotopy types. The role of the fundamental invariant is then played by a space or pro-space, and it can be constructed before deciding which coefficients the classification ought to allow.

5.1 Shape as a functor of constant tests

Let $\mathcal{X}$ be an ∞-topos with constant-object functor $\Delta\colon \mathcal{S} \to \mathcal{X}$ and global sections Γ. For $U \in \mathcal{X}$ define $\Pi(U)$ by

$$\mathrm{Map}_{\mathrm{Pro}(\mathcal{S})}(\Pi(U), T) \simeq \mathrm{Map}_{\mathcal{X}}(U, \Delta T). \tag{5.1}$$

The right side preserves finite limits in T. It is also accessible, since Δ preserves colimits and U is compact for some regular cardinal in the presentable category $\mathcal{X}$. The pro-representability criterion from Chapter 4 applies. Equivalently $\Pi(U)$ is the shape of the slice $\mathcal{X}_{/U}$, whose constant-object functor is $\Delta_U T = (U \times \Delta T \to U)$. The shape of $\mathcal{X}$ is $\Pi_\infty \mathcal{X} = \Pi(1)$ and corepresents the functor $T \mapsto \Gamma \Delta T$. These are the relative and absolute constructions of [2, §7.1.6] and [3, §2].

One presentation of $\Pi(U)$ uses the category of pairs (T, t) with $t\colon U \to \Delta T$. The maps are maps of test spaces together with specified coherent compatibilities with t. The category of pairs is cofiltered since the functor of tests preserves finite limits. Indeed, a finite diagram of such pairs determines a diagram of test spaces whose limit inherits a section by finite-limit preservation. The resulting pair supplies a cone. An essentially small cofiltered presentation exists by accessibility, and the projection to T pro-represents (5.1). Using all pairs is convenient for comparisons even when the class of tests is large.

There is a relation to actions internal to $\mathcal{X}$. For a point $a \in T$ put

$$P_a = U \times_{\Delta T} \Delta\{a\}.$$

The pullbacks inherit transport along T, producing a functor $a \mapsto P_a$ in $\mathcal{X}_{/U}$. The colimit of this functor is the terminal object of the slice. Indeed the family of all points of T, with all transport, has colimit the identity of T in $\mathcal{S}_{/T}$. Apply

Δ and then pull back along t, both of which preserve colimits. Conversely a functor $T \to \mathcal{X}_{/U}$ with colimit the terminal object is cartesian over the constant diagram of terminal objects. The correspondence between families over that constant diagram and objects over its colimit follows from universal descent in an ∞-topos. It gives a section $U \to U \times \Delta T$ in the slice. This is the sense in which the indexing pairs are principal T-actions, even when T has many components or has higher homotopy.

The functor $U \mapsto \Pi(U)$ sends colimits to colimits. For $U \simeq \operatorname{colim}_d U_d$, the right side of (5.1) is the limit of the corresponding right sides. Constant tests are jointly conservative, hence the comparison from the colimit of the $\Pi(U_d)$ is an equivalence. For a geometric morphism, pullback of constant objects along $f\colon \mathcal{Y} \to \mathcal{X}$ induces a map from $\Pi_\infty \mathcal{Y}$ to $\Pi_\infty \mathcal{X}$. It is an equivalence exactly when pullback gives equivalences on the spaces of all constant tests.

5.2 From stage representations to objects

For (T, t) as above and a family $E \to T$ of spaces, put

$$E_t = U \times_{\Delta T} \Delta E \longrightarrow U. \tag{5.2}$$

This operation respects both the reindexing of tests and natural transformations of families. Straightening over T and then taking the filtered colimit of stage representations gives a functor

$$R_U\colon \operatorname{Rep}_{\mathcal{S}}(\Pi(U)) \longrightarrow \mathcal{X}_{/U}.$$

Its values are locally constant. By definition this means that there is an effective epimorphism $\coprod_\alpha U_\alpha \to U$ such that the restriction to every U_α is equivalent in the slice to a constant object. In this case use the fibers $P_a \to U$ for one point a of each component of T. The sum of the selected points is an effective epimorphism onto T, and Δ and pullback preserve effective epimorphisms. Over P_a, the object E_t is constant with fiber $\operatorname{fib}_a(E)$. The same argument applies with $\mathcal{S}_{\le n}$ or $\mathcal{S}_\pi$ as the coefficient category.

A useful reduction in proving converse classification theorems is that stage representations satisfy descent on U. To prove this, let $\mathcal{D}$ be any chosen coefficient category, and write $W_m(\mathcal{D})$ for the core of $\operatorname{Fun}([m], \mathcal{D})$. The complete Segal nerve of $\operatorname{Rep}_{\mathcal{D}}(P)$ has m-th space

$$W_m(\operatorname{Rep}_{\mathcal{D}}(P)) \simeq \operatorname{Map}(P, W_m(\mathcal{D})). \tag{5.3}$$

Indeed, at a stage this follows by taking the core of $\operatorname{Fun}(P_i, \operatorname{Fun}([m], \mathcal{D}))$. Formation of W_m commutes with filtered colimits because $[m]$ is represented by a finite simplicial set and taking the core commutes with such colimits. Then use (4.1). Complete Segal nerves detect equivalences and preserve limits. Since

$\Pi(U)$ takes colimits to colimits, (5.3) proves that $U \mapsto \mathrm{Rep}_{\mathcal{D}}(\Pi(U))$ sends colimits in $\mathcal{X}$ to limits. All of the spaces $W_m(\mathcal{D})$ may be taken in a larger universe if required by the coefficient category.

Lemma 5.1. *Suppose $\mathcal{D}$ is a full subcategory of spaces and the comparison R_V induces an equivalence on mapping spaces between constant objects with values in $\mathcal{D}$ for every $V \in \mathcal{X}$. Then R_U with coefficients in $\mathcal{D}$ is fully faithful for all U, and its essential image is the subcategory of objects locally constant with values in $\mathcal{D}$.*

Proof. Given two stage representations, choose a common stage (T, t) and use its effective epimorphism $\coprod_a P_a \to U$ as above. In its Čech nerve, a component of an iterated fiber product of the cover is indexed by a tuple of points and maps into one of the P_a's. On that component both representations are constant. The mapping-space comparison is therefore an equivalence on each such component by the assumption. It is an equivalence on their coproduct because both representation and slice categories turn the coproduct into a product. Taking the limit over the Čech nerve gives the original mapping-space comparison. On the representation side this uses (5.3), and on the topos side it is effective descent.

Now take a locally constant object with the indicated coefficients. Choose an effective epimorphism giving trivializations. On every component in every degree of its Čech nerve the restrictions belong to the essential image, since constant objects have constant stage representations. Each realization functor there is fully faithful, so the coherent descent datum of this object is in the essential image of the limit of the functors. More explicitly, fullness of each image lifts all transition equivalences, and every higher simplex whose vertices are in the image lies in that full subcategory. Descent supplies a representation over U. □

The constant comparison in the lemma has an elementary form. For $A, B \in \mathcal{S}$ define $F_U(B) = \mathrm{Map}_{\mathcal{X}}(U, \Delta B)$. We have natural equivalences

$$\mathrm{Map}_{\mathrm{Rep}_{\mathcal{S}}(\Pi(U))}(A, B) \simeq F_U(\mathrm{Map}(A, B)),$$
$$\mathrm{Map}_{\mathcal{X}_{/U}}(\Delta_U A, \Delta_U B) \simeq \mathrm{Map}(A, F_U(B)).$$

The first is the formula for maps of constant diagrams at each stage. The second uses $\Delta_U A \simeq \mathrm{colim}_A U$ in the slice and then maps out of that colimit. If $\Pi(U) = \{P_i\}$, the comparison is consequently

$$\operatorname*{colim}_i \mathrm{Map}(A, \mathrm{Map}(P_i, B)) \longrightarrow \mathrm{Map}\big(A, \operatorname*{colim}_i \mathrm{Map}(P_i, B)\big). \tag{5.4}$$

This displays the relevant commutation requirement directly.

5.3 Two comparison theorems

Theorem 5.2 (Higher local monodromy)**.** *Fix $n \geq 0$, or let $n = \infty$. Suppose that the constant-object functor on n-truncated objects admits a genuine left adjoint*

$$L_n \colon \mathcal{X}_{\leq n} \longrightarrow \mathcal{S}_{\leq n}.$$

Then for each $U \in \mathcal{X}$, realization gives an equivalence

$$\mathrm{Rep}_{\mathcal{S}_{\leq n}}(\Pi(U)) \simeq \mathrm{LC}_{\leq n}(\mathcal{X}_{/U}).$$

Proof. For an n-truncated test T, ΔT is n-truncated. By adjunction,

$$F_U(T) \simeq \mathrm{Map}(L_n \tau_{\leq n} U, T).$$

If B is n-truncated, so is $\mathrm{Map}(A, B)$ for any space A. For constant coefficients in the theorem both sides of (5.4) therefore agree by the exponential law for spaces, with $L_n \tau_{\leq n} U$ as the representing space. Apply Lemma 5.1. For $n = \infty$, truncation in this argument means the identity. □

The hypothesis of Theorem 5.2 is local n-connectedness in the convention of [3, Def. 3.2]. For finite n, the genuine left adjoint represents the relative shape on n-truncated tests. Even then the monodromy category in general uses the *pro-$(n+1)$-truncation* of shape. Indeed the objects in $W_0(\mathcal{S}_{\leq n})$ have loops given by automorphism spaces, hence this classifier is $(n+1)$-truncated. The endpoint map of the arrow classifier has n-truncated fibers, so its total space is also $(n+1)$-truncated. By (5.3) and the mapping-space description using the arrow classifier, these two tests suffice to recognize equivalences of the representation categories. For $n = 0$ this recovers why a fundamental groupoid rather than just a space of connected components is needed to classify coverings.

Theorem 5.3 (Finite-coefficient higher Galois theory)**.** *For every ∞-topos and every U, realization induces an equivalence*

$$\mathrm{Rep}_{\mathcal{S}_\pi}(\Pi(U)) \simeq \mathrm{LC}_\pi(\mathcal{X}_{/U}).$$

Here local trivializations may have arbitrarily many different π-finite fiber types across the cover.

Proof. For $A \in \mathcal{S}_\pi$ and fixed finite n, $\mathrm{Map}(A, -)$ preserves filtered colimits of n-truncated spaces. Use the finite $(n+1)$-skeleton in the proof of Lemma 4.2. Mapping out of a finite cell complex is a finite limit of mapping spaces out of points, by repeated use of its cell attachments, and finite limits commute with filtered colimits in spaces. Filtered colimits of n-truncated spaces also stay n-truncated. Given $B \in \mathcal{S}_\pi$, choose n with B n-truncated. All spaces $\mathrm{Map}(P_i, B)$ in (5.4) are n-truncated. The comparison is therefore an equivalence. Lemma 5.1 applies without a connectivity hypothesis on $\mathcal{X}$. □

Theorem 5.3 is [3, Thm. 4.3]. Its coefficients are π-finite spaces, while its indexing object is the full shape. Proposition 4.3 permits using the profinite completion for the subcategory admitting a finite range of fiber types. The uncompleted shape can still be replaced by its *protruncation* for the full finite-coefficient theory when a uniform truncation bound is imposed on all fiber types. Without a uniform bound, a single classifier can have components of arbitrarily high degree. In particular, the phrase *finite-coefficient classification* must specify the category of coefficients and the tests of shape it claims to retain.

Consider the space

$$B = \coprod_{m \geq 2} K(\mathbb{Z}/2, m).$$

On each component, the based path fibration has π-finite fiber $K(\mathbb{Z}/2, m-1)$. If the family came from $B_{<\infty}$, it would descend to some Postnikov stage $\tau_{\leq d} B$. For $m > d$, the image of the m-th component in that stage is contractible, so the restricted family would have to be constant up to equivalence. The based path fibration has contractible total space, whereas the constant family on $K(\mathbb{Z}/2, m)$ with this fiber has nontrivial homotopy. The family therefore uses the full shape. This obstruction concerns a collection of higher degrees rather than its fundamental group.

5.4 The classifying topos and its limit

For a space T put $\beta(T) = \mathcal{S}_{/T}$. The classifying property needed for pro-spaces is

$$\mathrm{Map}_{\mathrm{Top}_\infty}(\mathcal{X}, \beta(T)) \simeq \mathrm{Map}_{\mathcal{X}}(1, \Delta T).$$

This is the étale slice universal property. The right side classifies a lift of $\mathcal{X} \to \mathcal{S}$ through $\mathcal{S}_{/T}$ by sending the universal point of its constant T to the chosen point. For a pro-space $P = \{P_i\}$, define $\beta(P) = \lim_i \beta(P_i)$ in Top_∞. The classifying property and (4.1) display an adjunction

$$\Pi_\infty \colon \mathrm{Top}_\infty \rightleftarrows \mathrm{Pro}(\mathcal{S}) : \beta.$$

Its unit is $\varphi \colon \mathcal{X} \to \beta(\Pi_\infty \mathcal{X})$. The realization functor at $U = 1$ factors as

$$\mathrm{Rep}_{\mathcal{S}}(\Pi_\infty \mathcal{X}) \longrightarrow \beta(\Pi_\infty \mathcal{X}) \xrightarrow{\varphi^*} \mathcal{X}.$$

The first arrow is induced by the projections from the topos limit, using their inverse images, while the category underlying that limit is computed by the direct images [3, §2]. Under the local n-connectedness hypothesis, Hoyois shows that both arrows, with n-truncated coefficients and objects, are fully faithful. Moreover the image of $\beta(\Pi_\infty \mathcal{X})_{\leq n}$ in $\mathcal{X}$ consists of filtered colimits of locally constant n-truncated objects [3, Thm. 3.13]. Theorem 5.2 above proves

the comparison on stage objects. The further claim about the whole classifying topos uses the fact that in this setting the shape admits a presentation by n-connected transition maps.

There is a direct description of the topos limit in degree zero for $P = B_{\mathrm{pro}}H$ from Chapter 4. An object of its 0-truncation is a coherent family of discrete H/U_i-sets M_i. If $U_j \subseteq U_i$, the direct image takes invariants under U_i/U_j, so the compatibility is $M_i \cong M_j^{U_i/U_j}$. These give compatible inclusions of inflated H-sets and their union is continuous. Conversely a continuous discrete H-set M yields the family of M^{U_i} and is its union. The family for the union of a compatible system returns that system since a fixed element in the union is represented by a fixed element at a common refinement. The construction on maps is by restriction and union. Thus the topos limit contains all continuous discrete actions, whereas the stage category only contains actions with a common open kernel.

The extension β to all pro-spaces is not fully faithful in general. Return to the tower of tails P_m in Chapter 4. The underlying category of $\beta(P_m)$ is a product of copies of spaces, with one coordinate for every element of the tail. The direct images extend a tail-indexed family by terminal objects at the missing coordinates. In a compatible family, each coordinate is eventually missing, so the corresponding objects in earlier stages are terminal. Hence $\beta(P)$ is the terminal category, the topos $\beta(\varnothing)$. But P and $\varnothing$ were distinguished by their constant empty test. This is the general pro-space obstruction recorded in [3, Rem. 2.8].

5.5 A failure of unrestricted stage monodromy

Let $C = \{0,1\}^{\mathbb{N}}$ with its product topology, and $\mathcal{X} = \mathrm{Shv}(C;\mathcal{S})$. Its shape is the system of finite clopen partitions of C, regarded as discrete spaces and ordered by refinement. Here is a direct computation. A basis is given by the clopen subsets, each of which is compact. A covering of a clopen subset by basis members has a finite disjoint clopen refinement, by choosing a finite subcover and successively subtracting earlier sets. Thus the topology on the basis is generated by finite disjoint clopen covering families. The constant sheaf with value T takes a clopen V to

$$\operatorname*{colim}_{\mathcal{P}} \prod_{D\in\mathcal{P}} T,$$

over finite clopen partitions of V. This prescription satisfies the disjoint-cover sheaf condition because partitions refining a fixed disjoint cover are cofinal and filtered colimits commute with finite products. Its map from the constant presheaf exhibits sheafification, since any natural map to a sheaf is forced, on a partition, by the maps from its parts and the disjoint-cover sheaf property. Taking $V = C$ computes every constant test of shape.

Consider the two constant sheaves with values $\mathbb{N}$ and $\{0,1\}$. A map between them can be specified by a sequence of locally constant functions on C with values in $\{0,1\}$. Indeed the source is a coproduct of countably many terminal objects. Choose the functions giving the successive coordinates of C. The sequence is a morphism of the constant sheaves, but no single finite partition simultaneously makes every coordinate function constant on each part. Such a partition would make the identity map of the underlying set of C factor through a finite set. The sequence therefore lies outside the image of the stage morphism space (5.4). This gives a failure of full faithfulness for unrestricted set-valued stage monodromy on a topos with profinite 0-dimensional shape. The finite-coefficient theorem still applies to this topos.

5.6 Covering spaces and higher transport

For a space B, the classifying topos $\mathcal{S}_{/B} \simeq \mathrm{Fun}(B, \mathcal{S})$ has shape B itself. Its global sections of a constant object T are $\mathrm{Map}(B, T)$. All its objects are locally constant, as pullback to the family of points of B gives their fibers. When B is connected and pointed, set-valued local systems are $\pi_1(B)$-sets, since their classifying core and arrow core are both 1-truncated. The fiber functor at b recovers $\pi_1(B, b)$ as its group of natural automorphisms. For left G-sets the forgetful functor is $\mathrm{Hom}_G(G, -)$. Equivariant self-maps of the left regular G-set are right translations, so its automorphism group is G^{op}. The contravariance of the Yoneda embedding in the representing object then gives G as the automorphism group of the fiber functor. This realizes the fundamental-group construction of [1, §6.7] in the representable case. For ordinary locally path-connected semilocally simply connected spaces, locally constant sheaves of sets are the sheaves of sections of covering spaces, and path lifting gives the same monodromy description on the ordinary fundamental groupoid. These local topological hypotheses ensure that every monodromy action reconstructs a cover.

Positive homotopy groups can enter even for connected base with trivial π_1. Take $B = K(\mathbb{Z}, 2)$, pointed at its unit. The based path fibration $P_b B \to B$ has fiber $K(\mathbb{Z}, 1)$ and contractible total space. The constant family with this fiber has total space $B \times K(\mathbb{Z}, 1)$, which is not contractible. Thus these are inequivalent families over B. Their set-valued monodromy cannot distinguish them, while the degree-two shape allows transport of the homotopies in the fiber.

5.7 Hypercompletion and shape

Let $h \colon \mathcal{X}^{\wedge} \to \mathcal{X}$ be the hypercompletion. Every truncated object is hypercomplete. Indeed an ∞-connected morphism is left orthogonal to truncated objects over the terminal object, so localizing at those morphisms fixes every truncated

object. For each truncated test T, ΔT is already truncated in $\mathcal{X}$, and

$$\mathrm{Map}_{\mathcal{X}}(1, \Delta T) \simeq \mathrm{Map}_{\mathcal{X}^\wedge}(1, h^*\Delta T).$$

Consequently hypercompletion induces equivalences on protruncated shape and on its profinite completion, as also observed in [3, Rem. 5.3]. That argument only tests truncated targets. Unbounded constant tests and filtered colimits of such tests require separate hypotheses. Likewise an attempted computation of full shape from Postnikov sections cannot exchange the uniformly bounded tests with their limit without establishing the required commutation. The distinction is quantified already by (4.1), which places a filtered colimit inside the limit over target stages.

For a concrete pro-space with trivial protruncation, use the tower

$$Q_i = \prod_{m \geq i} K(\mathbb{Z}, m),$$

with transitions that insert basepoints in the coordinates new to the target stage. For fixed d the stages $i > d$ are d-connected, so every d-truncated test treats Q as terminal. Yet the compatible inclusions $Q_i \to Q_1$ define a point of $\mathrm{Map}(Q, Q_1)$ that is not in the component of a constant map. At any refinement, the coordinate inclusion followed by projection to that coordinate induces the identity of $K(\mathbb{Z}, m)$. A constant homotopy class would yield the zero map on its positive homotopy group. This proves that unbounded tests can see data invisible in every fixed truncation.

CHAPTER 6

Obstructions to higher monodromy

A classical action on a discrete set is determined by its permutation representation. For a space F, actions on F are classified by maps into $B\,\mathrm{Aut}(F)$, whose higher homotopy groups control extensions of the permutation-level information. This chapter describes the lifting problem without assuming that the base has a discrete fundamental group as its entire homotopy type. The argument applies also to maps from a pro-space into a bounded target.

6.1 Automorphisms of an Eilenberg–Mac Lane space

For an abelian group A and $n \geq 1$, choose a functorial Eilenberg–Mac Lane space $K(A,n)$ represented by a simplicial abelian group. Transport through another model leaves the following equivalences invariant. Let Aut_* denote the space of based self-equivalences, with the identification at the basepoint included as data.

Lemma 6.1. *The map $\mathrm{Aut}_*(K(A,n)) \to \mathrm{Aut}(A)$ is an equivalence onto a discrete group. The full group object of self-equivalences is equivalent to*

$$\mathrm{Aut}(K(A,n)) \simeq K(A,n) \rtimes \mathrm{Aut}(A),$$

with the translation action on the first factor and the usual action of coefficient automorphisms.

Proof. Pointed maps $K(A,n) \to K(A,n)$ have components $\widetilde{H}^n(K(A,n);A) \cong \mathrm{Hom}(A,A)$. For $n \geq 2$, this follows from connectivity, Hurewicz, and the universal coefficient theorem. For $n = 1$, the identification uses $H_1(K(A,1);\mathbb{Z}) = A$ and the freeness of H_0. In both cases the higher homotopy groups of the pointed mapping space are $\widetilde{H}^{n-i}(K(A,n);A)$ in degree $i > 0$, by suspension and the representing property of Eilenberg–Mac Lane spaces. They vanish by connectivity and reducedness. A pointed map is an equivalence precisely when the induced endomorphism of A is invertible. The based maps induced by automorphisms of A therefore give an equivalence from the discrete group.

Evaluation of an unbased self-map at zero has this pointed mapping space as its fiber. Translation by the evaluation value provides a splitting on underlying spaces. In the chosen additive model the assignment $(a,\alpha) \mapsto (x \mapsto a + \alpha(x))$ is a map of group objects, using the semidirect product law. Its effect on the evaluation fibers and bases is an equivalence, which proves the assertion. □

Delooping gives a fibration with a distinguished section

$$K(A, n+1) \longrightarrow B\,\mathrm{Aut}(K(A, n)) \longrightarrow B\,\mathrm{Aut}(A). \tag{6.1}$$

The section is induced by the inclusion of coefficient automorphisms in the semidirect product. In the description by families over $B\,\mathrm{Aut}(A)$, the fibration is the local system of spaces $K(A, n+1)$ with its natural transport. Indeed looping the quotient retrieves the stated conjugation action on $K(A, n)$, and the section identifies the neutral point in each fiber.

A local system of abelian groups on a space Y is a functor $M : Y \to \mathrm{Ab}$, where Ab is the ordinary category of abelian groups and all arrows in the image of Y are isomorphisms. Write $K_Y(M, r) \to Y$ for the family associated to the functor $y \mapsto K(M_y, r)$. Define $H^r(Y; M)$ to be the component group of the section space $\Gamma_Y K_Y(M, r)$. This agrees with cohomology with local coefficients. In particular the zero section is part of the structure, and translation identifies each section component with the component of zero. For $0 \le i \le r$, looping the fiberwise group and using that sections preserve limits yields

$$\pi_i\big(\Gamma_Y K_Y(M, r), 0\big) = H^{r-i}(Y; M).$$

For $i > r$ the homotopy group vanishes. This constructive description needs no choice of cocycle model for the base.

6.2 The Postnikov lifting problem

Let Z be a connected space, $n \ge 2$, and $Y = \tau_{\le n-1} Z$. The map $\tau_{\le n} Z \to Y$ has fiber $K(A, n)$ for $A = \pi_n Z$ at a chosen basepoint, by the homotopy exact sequence. Transport gives the local coefficient system M associated to the usual $\pi_1(Z)$-action on A. The relative assertion below uses the full section space.

Theorem 6.2 (A Postnikov step as a space of nullhomotopies)**.** *The Postnikov extension $\tau_{\le n} Z \to Y$ determines a section k of $K_Y(M, n+1)$ defined up to equivalence with its classifying data. The extension is equivalent over Y to the homotopy pullback of k and the zero section. Given a map $f : T \to Y$, the space of lifts of f to $\tau_{\le n} Z$ is equivalent to the space of paths from 0 to f^*k in $\Gamma_T K_T(f^*M, n+1)$. It is nonempty exactly when $[f^*k] = 0$ in $H^{n+1}(T; f^*M)$. If a lift is fixed, the full lift space is equivalent to $\Gamma_T K_T(f^*M, n)$.*

Proof. A family over Y whose fibers are equivalent to $K(A, n)$ is classified by a map $Y \to B\,\mathrm{Aut}(K(A, n))$, by straightening and the component of $K(A, n)$ in the core of spaces. Composing with (6.1) gives the classifying map of M. Over that map, a lift to $B\,\mathrm{Aut}(K(A, n))$ is exactly a section k of $K_Y(M, n+1)$. One can identify the classified family itself by pulling back the bundle of paths that start at zero in the fiber $K(A, n+1)$. Such paths have endpoint $k(y)$ and their path space is a principal $K(A, n)$-space. Under trivialized coefficients

this is the universal path fibration, whose transport consists of translations. Coefficient automorphisms respect zero and act functorially on these paths. Across a change of coefficients the transport thus combines translations with automorphisms by the semidirect product of Lemma 6.1. This is the universal family for the classifying map from which k was obtained.

Pulling back along f and taking sections preserves the pullback because sections are represented by the terminal object in the slice. The stated path space therefore gives all lifts and their homotopies. If one path is fixed, composing with its inverse identifies this path space with the loop space at zero. Looping the fiberwise Eilenberg–Mac Lane space and then taking sections proves the final claim. □

In particular, when the obstruction is zero the set of equivalence classes of lifts is a torsor for $H^n(T; f^*M)$; if a lift is chosen, its higher automorphisms in the lift space are given by the lower cohomology groups. Applying this statement to the Postnikov tower of $B\,\mathrm{Aut}(F)$ gives successive lifting problems for monodromy with prescribed fiber F. The first stage records the discrete group $\pi_0\,\mathrm{Aut}(F)$, and later stages can change both existence and the space of choices. The usual Postnikov construction of spaces is reviewed in [18], and the corresponding viewpoint on cohomology in slices is developed in [4].

Example 6.3 (A nonzero obstruction over a torus)**.** Take $Y = B\mathbb{Z}^3$, the homotopy type of the three-torus, $M = \mathbb{Z}$ with trivial transport, and $n = 2$. A generator k of $H^3(Y;\mathbb{Z})$ defines a family $Z \to Y$ as the fiber of $k : Y \to K(\mathbb{Z},3)$. Its fibers are $K(\mathbb{Z},2)$ and the homotopy sequence shows that Z is a connected 2-type with $\tau_{\le 1}Z \simeq Y$. The identity of Y cannot lift. The map $\mathbb{T}^3 \to \mathbb{T}^3$ that sends the third coordinate to its identity and leaves the first two coordinates fixed does lift, because it factors through the two-torus. The lift components in the second case form a torsor for $H^2(\mathbb{T}^3;\mathbb{Z}) = \mathbb{Z}^3$. At a chosen lift the degree-one and degree-two homotopy groups of the lift space are $\mathbb{Z}^3$ and $\mathbb{Z}$, respectively. All these claims follow from the exterior-algebra cohomology of the torus and Theorem 6.2. They distinguish the vanishing test, the moduli of solutions, and the group of automorphisms of a solution.

6.3 Continuous lifting and convergence

Let $P = (P_i)$ be a pro-space as in Chapter 4. Suppose that a partial monodromy map to a space Y is represented by $f_i : P_i \to Y$. For j over i, write $f_j : P_j \to Y$ for its pullback. The lifting problem for a target map $Z \to Y$ has a useful stage description.

Proposition 6.4. *There is a natural equivalence*

$$\begin{aligned}\mathrm{fib}_{[f_i]}\big(\mathrm{Map}_{\mathrm{Pro}(\mathcal{S})}(P,Z) \longrightarrow \mathrm{Map}_{\mathrm{Pro}(\mathcal{S})}(P,Y)\big)\\ \simeq \operatorname*{colim}_{j\to i} \mathrm{fib}_{f_j}\big(\mathrm{Map}(P_j,Z) \longrightarrow \mathrm{Map}(P_j,Y)\big).\end{aligned}$$

For a Postnikov step as in Theorem 6.2, a lift exists exactly when the pulled-back obstruction vanishes at some refining stage.

Proof. Restricting to refinements of the chosen stage is cofinal in the filtered direction. The map in the proposition then compares the filtered colimits of the indicated diagrams of mapping spaces. The compatible basepoint diagram is constant at a point and induces $[f_i]$. Filtered colimits of spaces commute with the pullback defining its fiber. This proves the equivalence. A filtered colimit of spaces has a component exactly when some stage does, since π_0 preserves colimits. The result now follows from the vanishing criterion above. □

A space Z is the limit of its Postnikov tower. There is therefore a comparison of spaces

$$\operatorname*{colim}_i \lim_n \operatorname{Map}(P_i, \tau_{\le n} Z) \longrightarrow \lim_n \operatorname*{colim}_i \operatorname{Map}(P_i, \tau_{\le n} Z). \tag{6.2}$$

The left side classifies maps from P to the constant pro-space Z. The right side classifies maps from P to the pro-object consisting of the truncations of Z. A lift for each separate obstruction need not use one fixed stage i. A bounded target, or a constant source, gives an equivalence in (6.2). The following computation shows how the two sides differ without either condition.

Example 6.5 (A compatible tower with no common stage)**.** Fix an abelian group $A \neq 0$, for instance $\mathbb{Z}$. Let $P_i = \prod_{k \ge i} K(A, k)$ for $i \ge 2$, with maps $P_{i+1} \to P_i$ given by inserting the basepoint in the ith coordinate. Let $Z = P_2$. Every P_i is $(i-1)$-connected. For fixed n, the map $\operatorname{Map}(*, \tau_{\le n} Z) \to \operatorname{Map}(P_i, \tau_{\le n} Z)$ is an equivalence for $i > n$, since the target is n-truncated. The right side of (6.2) is consequently Z and is connected.

On the left side, inclusions $P_i \to Z$ represent one compatible element of $\operatorname{colim}_i \pi_0 \operatorname{Map}(P_i, Z)$. This element differs from the component of the constant map. Indeed, for every refining stage P_j, restrict the inclusion to $K(A, j)$ using the jth factor and project Z to that factor. The composite is the identity and induces id_A on π_j, whereas the constant map induces zero. Thus every bounded observation becomes trivial after refinement, while the unbounded monodromy test detects a nontrivial component.

There is a second issue even when the source is an ordinary space. The inverse limit of the spaces of lifts encodes compatibility among choices. Passing to their component sets before forming the limit can lose information. For a countable tower of pointed spaces, the homotopy limit is the homotopy equalizer of the identity and shift on the product. An explicit point of the homotopy equalizer consists of a point in each space and a path from that point to the image of the next point. This model yields a direct example.

Example 6.6 (The degree-p tower of circles)**.** For a prime p, take the tower $B\mathbb{Z}$ with transition maps induced by multiplication by p. It is a tower of abelian

group objects in spaces. The difference of the identity and shift on the product is induced on π_1 by

$$D : \prod_{i\geq 0} \mathbb{Z} \longrightarrow \prod_{i\geq 0} \mathbb{Z}, \qquad (a_i) \longmapsto (a_i - p a_{i+1}).$$

The homotopy equalizer is its fiber as a map of group objects, so the long exact sequence identifies its component group with $\operatorname{coker} D$. The map $(b_i) \mapsto \sum_{i\geq 0} p^i b_i$ converges in $\mathbb{Z}_p$, is surjective by the p-adic expansion, and takes the image of D into the ordinary integers by telescoping. Conversely, if the sum is an integer s, put

$$a_i = p^{-i}\Big(s - \sum_{k<i} p^k b_k\Big) \in \mathbb{Z}.$$

Divisibility follows from the p-adic equality, and $a_i - p a_{i+1} = b_i$. Thus the component group is $\mathbb{Z}_p/\mathbb{Z}$, although the component set at each stage is a singleton. This is the elementary tower calculation behind the derived inverse limit term.

The two examples locate different coherence requirements. Maps into a pro-target can be compatible even when they have no common source stage. Choices in a tower of lift spaces can have additional components carried by the transition paths. The classification space keeps both effects visible.

Part III

Algebraic Galois theories

CHAPTER 7

Pierce spectra and derived étale descent

The ring-theoretic construction in [1, Chapter 4 and Appendix A] changes two ingredients of field-theoretic Galois theory. It uses the entire space of connected components as the space of objects of the Galois groupoid, and it interprets split algebras by sections of a sheaf of rings over a Stone space. Both constructions extend to commutative ring spectra. This extension has a particularly concrete feature. Idempotents determine the topology at every level of the construction, while the algebras of sections retain all the homotopy groups of the ring spectrum. The derived affine category and its Pierce reflection therefore carry different amounts of information.

Throughout this chapter, A denotes an E_∞-ring, possibly nonconnective, CAlg_A its ∞-category of commutative algebras, and Y a Stone space, that is, a compact Hausdorff space with a basis of clopen subsets. The empty space and the zero ring are allowed. Write $P(A)$ for the Stone spectrum of the Boolean algebra $\mathrm{Idem}(\pi_0 A)$. Points are ultrafilters. For an idempotent e, the clopen subset $D(e)$ consists of ultrafilters containing e, and every clopen of $P(A)$ arises uniquely in this way. A ring map $A \to C$ gives a continuous map $P(C) \to P(A)$ by inverse image of ultrafilters. All tensor products between ring spectra below are derived.

7.1 Idempotents, localizations, and stalks

Multiplication, complement, and join in $\mathrm{Idem}(\pi_0 A)$ are given by ef, $1-e$, and $e+f-ef$. These formulas only use degree zero. The localization $A[e^{-1}]$, by contrast, is an algebra over the entire spectrum. Its universal property says that the mapping space from $A[e^{-1}]$ into an A-algebra C is contractible when e maps to a unit of $\pi_0 C$, and empty otherwise. An idempotent is a unit exactly when it is the element 1. Localization satisfies

$$\pi_t(A[e^{-1}]) \cong (\pi_t A)[e^{-1}] \cong e\pi_t A \qquad (t \in \mathbf{Z}).$$

Here the last module is viewed over $e\pi_0 A$, whose unit is e. These facts are instances of spectral localization [6, Example 7.5.0.7]. They imply

$$A \xrightarrow{\sim} A[e^{-1}] \times A[(1-e)^{-1}].$$

Indeed the map induces $m \mapsto (em, (1-e)m)$ on every homotopy group. The proof also applies to any finite partition of unity by pairwise orthogonal idempotents. Equivalences of ring spectra are detected on the underlying spectra.

Proposition 7.1. *The rule $D(e) \mapsto A[e^{-1}]$, with the localization maps as restrictions, defines a sheaf on the clopen basis of $P(A)$. It extends uniquely, up to equivalence, to a sheaf of E_∞-rings on all open subsets. Its global sections are A, and its stalk at an ultrafilter x is*

$$A_x \simeq \operatorname*{colim}_{e \in x} A[e^{-1}],$$

where the transition maps restrict along inclusions of smaller clopens. The ring $\pi_0 A_x$ has precisely the two idempotents 0 and 1.

Proof. Every cover of a clopen by clopens has a finite subcover and a finite pairwise disjoint clopen refinement formed by intersections and complements. For a disjoint cover the sheaf comparison is exactly the product decomposition above, and intersections of distinct members have the zero ring as sections. The finite refinements yield the sheaf condition on the basis. Explicitly, for a finite cover, decompose every member and every overlap by the same finite Boolean partition. On each nonempty part, the indices of cover members containing it form a nonempty set, with a contractible Čech indexing groupoid. The limit of the restrictions over that part is its section ring. Taking the product over parts recovers the original section ring. The same argument applies to an arbitrary clopen cover by restricting over each part of a finite subordinate disjoint partition. On each part a chosen member contains the entire part, so its restricted cover has a section and the Čech limit is split. The extension and the stalk formula follow by restriction to the basis.

Filtered colimits commute with π_0, and the polynomial identity defining an idempotent descends to a stage. Every idempotent in $e\pi_0 A$ is an idempotent f of $\pi_0 A$ with $f \leq e$. An ultrafilter containing e contains either f or $1-f$. Restrict along the former to make the idempotent 1 and along the latter to make it 0. Each ring at a stage has distinct unit and zero because its idempotent belongs to an ultrafilter. They remain distinct in the filtered colimit. □

An ultrafilter stalk is indecomposable in the sense of products. The Zariski spectrum instead organizes tests at prime ideals and uses local rings. For degree-zero A, Proposition 7.1 recovers the Pierce sheaf in [1, Section 4.2]. Its spectral analog shows exactly where the extra information resides. The stalks and restrictions retain the groups $e\pi_t A$ before passage to the colimit.

7.2 A relative Stone adjunction

Let $q : Y \to P(A)$ be continuous. To build its algebra of sections, take pairs $(\mathcal{U}, W)$ where $\mathcal{U}$ is a finite partition of Y into nonempty clopens and $W =$

$(W_U)_{U\in\mathcal{U}}$ is a family of clopens in $P(A)$ satisfying $q(U)\subseteq W_U$. For the empty partition use the empty product, the zero algebra. A refinement partitions into smaller clopens and shrinks each chosen neighborhood inside the neighborhood of its parent. Denote by e_{W_U} the associated idempotent and put

$$A\langle Y\rangle = \operatorname*{colim}_{(\mathcal{U},W)} \prod_{U\in\mathcal{U}} A[e_{W_U}^{-1}]. \tag{7.1}$$

The transition map to a refinement repeats the old factor for each of its children and then restricts it along the smaller neighborhood. Intersections of partitions and neighborhoods give common refinements, so this is a filtered colimit. No point of $q(U)$ is required to have an open image, which is the reason for using neighborhoods of the images.

Lemma 7.2. *For an A-algebra C, the space of maps*

$$\operatorname{Map}_{\mathrm{CAlg}_A}\left(\prod_{j=1}^{r} A[e_j^{-1}], C\right)$$

is discrete. Its points are ordered clopen partitions $P(C)=\coprod_j V_j$, where empty parts are allowed and each V_j maps into $D(e_j)$ in $P(A)$.

Proof. The factors in the source define orthogonal idempotents summing to 1, whose images in $\pi_0 C$ give the partition (V_j). After this partition is fixed, the target decomposes as $\prod_j C[v_j^{-1}]$, for the associated idempotents v_j. Mapping to the j-th factor forces the source idempotent of the j-th factor to map to 1, so the map factors through the j-th source factor. The universal property of localization makes this factorwise mapping space contractible exactly when $e_j v_j = v_j$. To see that no higher ambiguity enters through the partition, one can use [6, Corollary 7.5.4.6]. The source is a spectral étale A-algebra because its degree-zero algebra is a finite product of idempotent localizations and its higher groups obey the flat base-change formula. Thus mappings to an arbitrary target are determined discretely by degree zero. The Boolean condition $e_j v_j = v_j$ is equivalent to the asserted containment of clopens. □

Theorem 7.3 (Relative Pierce adjunction)**.** *There is a natural equivalence*

$$\operatorname{Map}_{\mathrm{CAlg}_A}(A\langle Y\rangle, C) \simeq \operatorname{Hom}_{\mathsf{Stone}/P(A)}(P(C), Y).$$

The canonical continuous map $Y\to P(A\langle Y\rangle)$ is an isomorphism over $P(A)$. Consequently $Y\mapsto A\langle Y\rangle$ embeds $(\mathsf{Stone}/P(A))^{\mathrm{op}}$ fully faithfully into CAlg_A. Its image will be called the Pierce-split *A-algebras. On opposite categories the left adjoint is the functor $C\mapsto P(C)$ over $P(A)$. The corresponding reflection is therefore represented on algebras by the map $A\langle P(C)\rangle\to C$, and its fixed objects are exactly the Pierce-split algebras.*

Proof. For a pair $d = (\mathcal{U}, W)$, set $Y_d = \coprod_{U \in \mathcal{U}} W_U$ over $P(A)$. The basis containment condition in Lemma 7.2 gives the mapping set from $P(C)$ to Y_d over $P(A)$. The finite-stage ring has Pierce spectrum Y_d, since the Boolean algebra of its idempotents is the finite product of those of its factors. A refinement gives $Y_{d'} \to Y_d$ by inclusion into parent neighborhoods. There is an isomorphism in Stone spaces over $P(A)$

$$Y \xrightarrow{\sim} \lim_d Y_d.$$

For nonempty Y, a point of the limit has a single base coordinate $x \in P(A)$, as is seen by considering the coarsest partition with neighborhood $P(A)$. The coordinates at the unshrunk partitions choose a compatible member of each finite clopen partition. They determine a unique $y \in Y$ by compactness and separation by clopens. For any clopen neighborhood of $q(y)$, continuity gives a clopen piece about y mapping into that neighborhood, and the corresponding coordinate forces x to lie in it. Hence $x = q(y)$. Conversely y supplies all coordinates. The canonical continuous bijection is a homeomorphism by compactness and Hausdorffness. The statement for the empty space follows from the empty partition.

Mapping out of the colimit in (7.1) now gives the displayed equivalence since limits of discrete spaces remain discrete. Independently, filtered colimits of E_∞-rings induce filtered colimits of Boolean algebras of idempotents. An element and the equation making it idempotent occur at a finite stage, and equalities hold after a further refinement. Stone duality identifies the spectrum of that colimit with $\lim_d Y_d$. This proves the second claim and identifies the two constructions naturally. The last statement follows by taking $C = A\langle Z \rangle$. □

For an ordinary ring R, formula (7.1) gives the ring of sections of the inverse-image Pierce sheaf over Y. A section of the inverse image is locally represented by a section on a clopen neighborhood of its image. Compactness gives a finite clopen partition subordinate to these representatives. If two collections of representatives define the same section, local equality of their germs and a second finite clopen refinement make their representatives equal on sufficiently small clopen neighborhoods. Thus they become equal at a common stage. This identifies the construction with the relative section functor of [1, Section 4.3], including its variance. For a ring spectrum the same formula computes each homotopy group by taking the colimit of the corresponding finite product of modules $e_{W_U} \pi_t A$. Each finite-stage ring is finite étale over A, since its ordinary degree-zero algebra is a finite product of direct summands of the unit and its homotopy groups satisfy flat base change. Thus the Pierce-split algebras are filtered colimits of these particular finite étale algebras.

Example 7.4 (A compact parameter space)**.** Suppose $\pi_0 A$ is nonzero with no nontrivial idempotents and $Y = \{0,1\}^{\mathbf{N}}$ is the Cantor space. The relative algebra is the filtered colimit of products of copies of A indexed by clopen

partitions. Its degree-t group is the group of locally constant functions $Y \to \pi_t A$ with the target discrete. Every such function has finite image by compactness. For degree zero, even the characteristic function of a single point lies outside this group, since that point is closed and fails to be open. The product $\prod_{y \in Y} A$ has all functions as homotopy groups and has the Boolean algebra of all subsets of Y as degree-zero idempotents. The two constructions distinguish continuous parameters from arbitrary parameters.

Corollary 7.5 (Base change)**.** *For a map $A \to C$, there are natural equivalences*

$$C \otimes_A A\langle Y \rangle \simeq C\langle Y \times_{P(A)} P(C) \rangle, \qquad A\langle Y \rangle \otimes_A A\langle Z \rangle \simeq A\langle Y \times_{P(A)} Z \rangle.$$

In particular, P takes an affine pullback to the corresponding Stone pullback if one leg is the opposite of a Pierce-split algebra.

Proof. Apply the universal property of pushout to maps into an arbitrary C-algebra D. The right side of Theorem 7.3 then describes lifts of $P(D) \to P(A)$ to Y, which are precisely maps over $P(C)$ to the indicated pullback. The first equivalence follows from Yoneda. Taking $C = A\langle Z \rangle$ gives the second. □

7.3 The Galois groupoid of a normal Pierce splitting

We now impose an effectivity hypothesis separately from the classification of split algebras. For $A \to B$, write $B_n = B^{\otimes_A (n+1)}$, with Amitsur cofaces given by insertion of units and codegeneracies by multiplication. Say the map has *affine algebra descent* when extension of scalars induces

$$\mathrm{CAlg}_A \xrightarrow{\sim} \lim_{[n] \in \Delta} \mathrm{CAlg}_{B_n}. \tag{7.2}$$

In this limit, the transition functors are extensions of scalars. Say that the map is *normal for Pierce splitting* when $B \otimes_A B$, viewed as a B-algebra using the first factor, is Pierce-split. Symmetry gives the condition for the second factor. This condition is adapted to the inverse-limit construction above. It refers to one specific reflection, whereas affine algebra descent refers to the ambient affine category.

Theorem 7.6 (Pierce Galois correspondence)**.** *If $A \to B$ has affine algebra descent and is normal for Pierce splitting, the simplicial Stone space*

$$G_n = P(B_n)$$

is an internal groupoid. The opposite of the full subcategory of CAlg_A on those D for which $B \otimes_A D$ is Pierce-split is equivalent to the category of internal $G_\bullet$-actions in Stone spaces. The underlying space of the action associated to D is $P(B \otimes_A D)$ over $G_0 = P(B)$.

Proof. The tensor-power algebra in degree n, with its map from the factor at any vertex, is the n-fold tensor product over B of $B \otimes_A B$ with a common selected factor. Normality and Corollary 7.5 make it Pierce-split over that vertex. The usual Segal decomposition of the affine Čech nerve expresses its degree n as the iterated pullback of n consecutive arrows over $\mathrm{Spec}(B)$, where Spec here denotes passage to opposite algebras. Each pullback can use a Pierce-split leg at its vertex. Corollary 7.5 therefore identifies G_n with the corresponding Stone pullback of composable arrows. Applying the same corollary to the first and last vertices of a triangle, or to either pair of edges with a common tail or head, gives the outer horn isomorphisms for inverses. Concretely, the exchange of the two factors in B_1 supplies inversion, and its identities with multiplication follow from the corresponding identities in the affine Čech groupoid. Thus $G_\bullet$ is a groupoid object.

Under (7.2), requiring the object at degree zero to be Pierce-split is equivalent to requiring it at all degrees. Indeed each comparison with a vertex identifies the object at degree n with an extension of scalars of its degree-zero value, and Pierce splitting is preserved by extension of scalars. The equivalence of Theorem 7.3 is natural for precisely these transition functors by Corollary 7.5. Taking opposites therefore identifies the desired category with

$$\lim_{[n]\in\Delta} \mathsf{Stone}_{/G_n}.$$

For an internal groupoid this limit is the category of internal actions. Given an object $V \to G_0$, an identification of its two pullbacks to G_1 is an arrowwise transport map. The degenerate arrows impose its unit law and the pullbacks over G_2 impose composition. The Segal isomorphisms identify all further compatibilities with iteration of these two laws. Conversely a unitary associative action gives that identification and all its extensions by pullback. Morphisms are maps over G_0 commuting with the transport. This proves the equivalence and the claim about the underlying space. □

The condition on $B \otimes_A B$ also gives the closure of the trivial objects under the special base change that transports a B-algebra along $A \to B$. For $D = B\langle Y\rangle$, first apply Corollary 7.5 with the second B-structure on $B \otimes_A B$. The resulting algebra is Pierce-split over $B \otimes_A B$. If the latter, with its first B-structure, corresponds to a Stone space T over $P(B)$, the composite is Pierce-split over the first B. To verify this last assertion, express each algebra over $B \otimes_A B$ using (7.1). A finite product of clopen factors of $B\langle T\rangle$ corresponds over B to a finite union of clopen subspaces of T, by the relative mapping formula. Filtered colimits of these algebras correspond to inverse limits of their Stone spaces by that same formula. Conversely, taking $D = B$ in the closure condition recovers normality. In the ordinary ring setting, these are the split-object and groupoid mechanisms in [1, Sections 4.5–4.7].

7.4 Flat descent and the geometric étale site

Call $A \to B$ faithfully flat when $\pi_0 A \to \pi_0 B$ is faithfully flat and

$$\pi_t A \otimes_{\pi_0 A} \pi_0 B \xrightarrow{\sim} \pi_t B \qquad (t \in \mathbf{Z}).$$

For any A-module M, flat base change gives $\pi_t(M \otimes_A B) = \pi_t M \otimes_{\pi_0 A} \pi_0 B$. One way to establish this formula is to compare the two homological functors in M. The candidate assembled from the homotopy groups and a flat module is exact and preserves coproducts, and the comparison is an isomorphism on all suspended copies of A. The ordinary ring $\pi_0 A$ acts on the homotopy groups, so the comparison is balanced. The class of modules for which it is an isomorphism in every degree is stable and closed under coproducts. The suspended unit generates Mod_A as a localizing subcategory. In particular extension of scalars detects equivalences and concentration of homotopy groups.

Faithfully flat spectral descent identifies Mod_A with the Amitsur limit of Mod_{B_n} by extension of scalars [8, Theorem 3.1]. The functors are symmetric monoidal, so the equivalence is symmetric monoidal with the pointwise structures on the limit. Passing to commutative algebra objects yields (7.2). Indeed an algebra is given by a functor out of the commutative operad satisfying the inert-arrow conditions, and functor formation and these conditions are checked in every coordinate of a limit. For ordinary rings, concentrate all homotopy groups in degree zero. A differential degree cannot survive base change and then become zero by faithful flatness. This recovers the affine statement of classical fpqc descent [9]. The symmetric monoidal comparison carries the multiplication and unit operations through module descent.

There is a different restriction of the same descent formula. An A-algebra C is *spectrally étale* in the geometric sense when $\pi_0 C$ is étale over $\pi_0 A$ in the ordinary sense and

$$\pi_t C \cong \pi_t A \otimes_{\pi_0 A} \pi_0 C \qquad (t \in \mathbf{Z}).$$

It is finite étale when this ordinary algebra is also finite as a module. Lurie's étale lifting theorem gives an equivalence

$$\mathrm{CAlg}_A^{\mathrm{et}} \simeq N(\mathrm{Alg}_{\pi_0 A}^{\mathrm{et}}) \tag{7.3}$$

and gives discrete mapping spaces from étale sources into any A-algebra, with mapping sets computed on π_0 [6, Theorem 7.5.0.6 and Corollary 7.5.4.6]. Formula (7.3) applies also to the finite étale subcategory. It is compatible with base change because pushouts of étale maps remain étale and the degree-zero comparison is ordinary base change. For connective bases, the corresponding spectral affine small étale sites agree with the classical sites for $\pi_0 A$. Their generating covers are tested by joint surjectivity on the same underlying spectra of ordinary rings [5]. The equivalence of sites yields an equivalence of their

∞-topoi of space-valued sheaves, which still have higher sheaf cohomology even though the affine étale probes and their maps are discrete.

Proposition 7.7. *If $A \to B$ is faithfully flat, the ∞-categories of étale and finite étale A-algebras are the respective Amitsur limits of the same subcategories over B_n.*

Proof. Both properties are stable under extension of scalars. Conversely let C be an A-algebra whose base change to B has either property. Ordinary faithfully flat descent of the étale and finite properties, together with $\pi_0(B \otimes_A C) \cong \pi_0 B \otimes_{\pi_0 A} \pi_0 C$, yields the required condition on degree zero [9]. For each t, tensor the homotopy comparison map for C over $\pi_0 A$ with $\pi_0 B$. The resulting map is the homotopy comparison for $B \otimes_A C$ as a B-algebra, by flat base change for modules. It is an isomorphism, so the original map is an isomorphism by faithful flatness. Hence the required subcategories are the restrictions of (7.2) selected in degree zero, which proves the assertion. □

7.5 What the two restrictions detect

The Pierce functor is built from components, while a small étale site is built from flat unramified maps and contains the Zariski opens. If A has connected π_0, the former sends its base to one point regardless of the prime spectrum. A geometrically étale map may have connected source and target with a genuinely different field of functions. Such a map is sent to the identity of a one-point Stone space. It can nevertheless become Pierce-split after a sufficiently large faithfully flat extension, and Theorem 7.6 then records its monodromy through the arrow space of that extension. Component data alone therefore cannot recover the geometric étale site on the base. The descent construction uses component data simultaneously on the cover and on its iterated overlaps.

Here is the field calculation at an infinite level. Let L/K be an infinite algebraic Galois extension, the union of finite Galois subextensions L_i/K, with Krull group $G = \lim_i \mathrm{Gal}(L_i/K)$. As L-algebras,

$$L \otimes_K L \cong \operatorname*{colim}_i \prod_{\mathrm{Gal}(L_i/K)} L \cong \{\text{continuous functions } G \to L_{\mathrm{disc}}\}.$$

At a finite stage the first isomorphism is the splitting of a separable Galois tensor product. For the second, a continuous function to the discrete field has finitely many clopen fibers by compactness, and a common finite-stage refinement detects all of them. Thus HL is normal over HK. The Pierce image of the Amitsur nerve is the groupoid with one object and arrow group G, and Theorem 7.6 classifies split algebras by continuous G-actions on Stone spaces, as in [1, Chapter 3]. All tensor powers here are flat over K, so passing to Eilenberg–Mac Lane spectra preserves their discreteness.

CHAPTER 8

Spectral Galois extensions and stable monodromy

A categorical Galois theory depends on which objects count as covers. For a commutative ring spectrum, ordinary étaleness of degree zero together with the flat homotopy base-change formula gives one answer. A different answer asks for commutative algebra objects that become products of components of the unit after an effective extension of scalars. The second answer extends the split-algebra construction of [1] into stable homotopy theory. It permits torsors such as the complexification map in periodic topological K-theory, which fail the homotopy flatness test. The two theories are related by descent, with different input conditions.

8.1 Three effectivity inputs

A *stable homotopy theory* here means a presentable symmetric monoidal stable ∞-category $\mathcal{C}$ in which the tensor product preserves colimits separately. Write $\mathbf{1}$ for its unit. An algebra $B \in \mathrm{CAlg}(\mathcal{C})$ is *faithful* if $B \otimes -$ is conservative. In a stable category this is equivalent to requiring $B \otimes M \simeq 0$ to imply $M \simeq 0$, by applying the test to cofibers. It is *descendable* if the thick tensor ideal it generates contains $\mathbf{1}$. A thick tensor ideal is closed under finite cofiber sequences, retracts, and tensor products with any object of $\mathcal{C}$.

The following criterion gives the difference between a conservative test and a finite proof of descent. Put $I = \mathrm{fib}(\mathbf{1} \to B)$, with its map $u : I \to \mathbf{1}$. Write $u_n : I^{\otimes n} \to \mathbf{1}$ for the tensor power of u.

Proposition 8.1 (Nilpotence criterion)**.** *The algebra B is descendable if and only if u_n is nullhomotopic for some $n \geq 1$.*

Proof. Say an object M belongs to $\mathcal{N}$ if $u_n \otimes M$ is nullhomotopic for some n. It is closed under tensoring with an arbitrary object and under retracts. Suppose $M_1 \to M \to M_2$ is a cofiber sequence and powers n and m annihilate the terms M_1 and M_2, respectively. The map $I^{\otimes m} \otimes M \to M$ becomes null after composition to M_2, by naturality. It lifts along $M_1 \to M$. Tensor that lift with $I^{\otimes n}$ and apply u_n on its target. The resulting composite factors through $u_n \otimes M_1$, hence vanishes. By naturality this equals $u_{n+m} \otimes M$. This shows extension closure, and suspension, desuspension and rotation of a cofiber sequence give the other two-out-of-three cases. Thus $\mathcal{N}$ is thick. It contains

B since multiplication splits $B \to B \otimes B$ and hence $B \otimes u$ is null. If B is descendable, $\mathbf{1} \in \mathcal{N}$.

Conversely, the cofiber of u_n is built from the cofibers in the chain

$$I^{\otimes n} \longrightarrow I^{\otimes (n-1)} \longrightarrow \cdots \longrightarrow I \longrightarrow \mathbf{1}.$$

These cofibers are $I^{\otimes j} \otimes B$ for $0 \leq j < n$, since the arrows are tensorings of u. Repeated use of the cofiber sequence for a composite puts the cofiber of u_n in the thick tensor ideal generated by B. If u_n is null, $\mathbf{1}$ is a retract of that cofiber. □

A descendable algebra is faithful. For a fixed M, the objects X satisfying $X \otimes M \simeq 0$ form a thick tensor ideal, so if that ideal contains B, it contains the unit. Descendability moreover implies effective Amitsur descent for $\mathcal{C}$ as a symmetric monoidal category [8, Proposition 3.22]. The decisive step in the comonadicity proof concerns a cosimplicial diagram that becomes split after tensoring with B. After this tensoring its partial-totalization tower is pro-constant. The objects X such that the tower for the diagram tensored by X is pro-constant form a thick tensor ideal. Finite limits commute with tensoring, and constant pro-objects form a stable full subcategory closed under retracts. Tensoring a pro-constant tower produces a pro-constant tower because tensoring defines a levelwise functor on pro-objects. Descendability puts the unit in the ideal. For a pro-constant tower, passage to the limit commutes with tensoring by B. The dual form of Barr–Beck then gives comonadicity, and the iterated base-change diagram gives the Amitsur limit.

Two other effectivity inputs serve different purposes. Faithfulness plus dualizability suffices for comonadicity without an assumption about compactness of the unit, because tensoring with a dualizable object preserves all limits. When the unit is compact these two conditions also imply descendability [8, Theorem 3.38]. Flat faithful descent of ring spectra, used in Chapter 7, is still another criterion. Its hypotheses are expressed on all homotopy groups and it establishes the Amitsur equivalence directly. In particular, a proof of equivalence of module categories via conservative, limit-preserving tensoring is available without checking homotopy flatness.

8.2 The Amitsur complex of a Rognes extension

Let G be a finite group and let $A \to B$ be a map of E_∞-rings with a coherent action of G on B by A-algebra maps. Its characteristic comparisons are

$$A \longrightarrow B^{hG}, \qquad \gamma : B \otimes_A B \longrightarrow \prod_{g \in G} B, \qquad \gamma_g = \mu \circ (1 \otimes g). \tag{8.1}$$

Here B^{hG} is the homotopy limit of the action diagram over BG, computed in ring spectra, and μ is multiplication. If both maps are equivalences, the

extension is G-Galois in the sense of Rognes. It is *faithful* if B is faithful as an A-module [7, Sections 4 and 6]. The faithfulness condition matters to the following categorical comparison. The fixed-point equation only describes the algebra of invariants, while descent must detect every module.

Lemma 8.2 (Bar description of the tensor powers)**.** *If γ is an equivalence, there is an equivalence of cosimplicial algebras*

$$B^{\otimes_A(\bullet+1)} \simeq \left([n] \longmapsto \prod_{(g_1,\dots,g_n)\in G^n} B \right),$$

where the right side is the cochain diagram of the G-action. The degree-n component at $(g_1,\dots,g_n)$ is the map that multiplies the factors after twisting factor i by $g_1\cdots g_i$, with the first factor untwisted.

Proof. For fixed n, use the first factor as a common B-base and write the tensor power as the n-fold tensor product over B of $B\otimes_A B$. Apply γ to each factor. Finite products of B-algebras are taken to finite products by base change, as can be checked on underlying modules, where finite products are also coproducts. This gives an equivalence indexed by independent elements $(h_1,\dots,h_n)$, with factor i twisted by h_i. Reindex by $h_i = g_1\cdots g_i$, a bijection of indexing sets.

For the first coface, the inserted unit in factor zero makes every later factor share the twist g_1. This is the coface applying g_1 after deleting it from the coordinate tuple. An interior insertion replaces adjacent group entries by their product. Insertion at the final position drops the last entry. A codegeneracy multiplies two adjacent factors and is identified with inserting the group identity at that place. These formulas are equalities for strict actions and are identities of composites of the multiplication and the action structure maps for a coherent action. Equivalently, the comparison is induced from the map of simplicial diagrams from the action groupoid of G on the affine represented by B to its affine Čech nerve, where an arrow and its first object determine the second. This construction gives the whole family with its coherences. □

Theorem 8.3 (Stable Galois descent)**.** *For a faithful G-Galois extension of E_∞-rings, extension of scalars induces symmetric monoidal equivalences*

$$\mathrm{Mod}_A \simeq (\mathrm{Mod}_B)^{hG}, \qquad \mathrm{CAlg}_A \simeq (\mathrm{CAlg}_B)^{hG}.$$

The group acts on the categories on the right by extension of scalars along its automorphisms of B. A module or algebra there has semilinear transport with respect to the action on its scalars.

Proof. Rognes's dualizability theorem makes B a dualizable A-module [7, Proposition 6.2.1]. Tensoring with it is therefore both conservative and limit-preserving. To check the latter, regard tensoring as an endofunctor of Mod_A

and use its left adjoint, tensoring with the dual $B^\vee$. Limits of B-modules are computed on their underlying A-modules, so extension of scalars also preserves them. By comonadicity, the Amitsur module comparison is an equivalence. The Čech comparison is symmetric monoidal, since every extension of scalars is, and the monoidal operations in the limit are pointwise.

Apply Lemma 8.2. Modules over $\prod_J B$ with J finite are products of modules over the factors. This follows by projecting to the orthogonal central idempotents, giving an inverse to taking the direct sum with its componentwise action. The first coface on the resulting category-valued bar construction is the twist by g_1, the middle cofaces compose group elements, and the last drops the last element. Thus its limit is the category of coherent semilinear descent data $(\mathrm{Mod}_B)^{hG}$. Passing to commutative algebra objects in the symmetric monoidal equivalence proves the second assertion. □

A fixed-point description follows at the level of individual objects. If D is a semilinear B-algebra in the right-hand category, its inverse image is the homotopy fixed-point algebra of the underlying action, under A. Indeed the right adjoint to extension of scalars with canonical descent data is formed by the limit of the forgetful diagram, including the transport. This is seen by taking mapping spaces out of any A-algebra and applying extension–restriction adjunction in every degree. The equivalence in Theorem 8.3 identifies this right adjoint with the inverse. The module statement holds by the same argument. Thus reconstructing the underlying algebra or module uses the full homotopy limit even when a classifying category of special algebras has discrete maps.

8.3 Intermediate split algebras

Suppose $\pi_0 B$ is nonzero and has only trivial idempotents. A B-algebra of the form $B^T = \prod_{t\in T} B$, with T a finite set, has a discrete space of maps to another such algebra B^U. By Lemma 7.2, that space is the set of maps $U \to T$. Extension of scalars along each $g : B \to B$ canonically carries B^T to a copy of B^T, since it preserves these products and takes the unit to the unit. Every coherent semilinear structure on B^T then amounts to a genuine action on T, by contravariance and discreteness. In formulas one may take the action on functions to be $(g\varphi)(t) = g(\varphi(g^{-1}t))$.

Corollary 8.4. *Under the hypotheses of Theorem 8.3 and connectedness of $\pi_0 B$ as above, the opposite of the category of A-algebras becoming B^T for finite T is equivalent to the category of finite G-sets. Connected objects equipped with a map of algebras to B correspond to pointed transitive G-sets, hence to subgroups $H \leq G$. For $T = G/H$ with point H, the associated intermediate algebra is B^{hH}.*

Proof. The full subcategory in question is preserved by the equivalence of Theorem 8.3 and is characterized after extension by the discrete products just

described. The object B corresponds to the regular G-set by (8.1). Hence a map of algebras to B gives a G-equivariant map from the regular G-set, which is specified by its value at the identity. A nonempty split object is connected precisely when the action is transitive, since a nontrivial product decomposition is exactly a disjoint union of invariant subsets. This also agrees with indecomposability of the underlying algebra. An idempotent would give a product decomposition, and each factor would remain split after base change because the idempotents of B^T are the characteristic functions of subsets of T. For transitive T, the stabilizer of the point yields the unique pointed isomorphism with G/H. The fixed-point construction for $B^{G/H}$ is a limit over the action groupoid $G \ltimes (G/H)$ with value B at each object and transport induced by the group action. This groupoid is equivalent to BH by inclusion of the stabilizer at H. Its limit is B^{hH}. Under the same identification, the map to B is evaluation at the selected object. □

The correspondence is formulated in the category of the split algebras and their maps. In spectral algebra there is generally no useful criterion for an intermediate object that simply asks for injectivity of a map on every homotopy group. Connected split objects have a categorical role inside the category defined by base change and decomposition into components. The stabilizer is read from that category.

Corollary 8.5 (Normal subgroups)**.** *In the situation of Corollary 8.4, if H is normal in G then $A \to B^{hH}$ is a faithful (G/H)-Galois extension.*

Proof. Set $Q = G/H$ and $D = B^{hH}$. The right translations of Q on G/H commute with the left G-action. Contravariance transports them to an action by A-algebra maps on D. After tensoring with B, the action is the regular permutation action on $\prod_Q B$. The torsor tensor-square comparison becomes the canonical equivalence obtained from the bijection between pairs of elements of Q and a point of Q together with a translation. The operation $B \otimes_A -$ detects equivalences and commutes with limits on modules, hence also on algebras. The fixed-point comparison after that operation reduces to $B \simeq (\prod_Q B)^{hQ}$, which holds because the translation action groupoid on Q is contractible. Finally if $D \otimes_A M \simeq 0$, then tensoring with B makes the finite product of copies of $B \otimes_A M$ zero. Since Q is nonempty and B is faithful, $M \simeq 0$. □

8.4 The stable Galois category and its range

Mathew's finite-cover construction makes the preceding restricted classification intrinsic to a stable homotopy theory [8, Definitions 6.1–6.2 and Theorem 6.5]. The Boolean algebra of idempotents of $\pi_0 \operatorname{End}_{\mathcal{C}}(\mathbf{1})$ plays the role of the Pierce components of the unit. Given a commutative algebra object C, its localization $C[e^{-1}]$ at an idempotent of the endomorphism ring of the free rank-one

C-module is the corresponding direct factor in its module category. The localization can be constructed using the telescope of multiplication by e. It is idempotent under relative tensor product and corresponds to the projection onto the summand where e acts as the identity.

An algebra $D \in \mathrm{CAlg}(\mathcal{C})$ is called a finite cover if some descendable C makes $C \otimes D$ a finite product of the algebras $C[e^{-1}]$, with independently chosen idempotents for the factors. A weak finite cover is defined in the same way with $C \otimes -$ conservative and limit-preserving. The idempotents need not partition the unit because several sheets may meet the same component. After a finite Boolean refinement into pairwise disjoint clopens, the product describes a finite number of copies on each clopen component.

The opposite of either finite-cover category is a Galois category in the general sense allowing several components, classified by a profinite groupoid [8, Theorem 5.36]. In particular the covers have discrete mapping spaces in their ambient ∞-category. The two categories agree when the unit is compact. For $\mathcal{C} = \mathrm{Mod}_A$ this compactness is automatic, since mapping out of A is computed on the underlying spectra and commutes with filtered colimits. For a connective A, Mathew's comparison theorem identifies these covers with the finite étale algebras of $\pi_0 A$, through their spectral étale lifts [8, Theorem 6.17]. This assertion requires more than the lifting theorem of Chapter 7. It says the tensor-categorical cover condition detects exactly the objects which the lifting theorem starts with in the connective case.

The effect of the definition on torsors can be checked in one direction directly. A faithful Rognes extension $A \to B$ is dualizable and has $B \otimes_A B \simeq \prod_G B$. Thus B is its own weak finite-cover splitting witness. In Mod_A, compactness of A upgrades this witness to descendability. Conversely, a G-torsor in the weak Galois category is a faithful G-Galois extension, and in the strong category it is a descendable one [8, Proposition 6.13 and Corollary 6.15]. A local trivialization checks the tensor-square comparison and faithfulness by conservativity. The local tensor-square map is the map from a trivial torsor to its product with the group. To check fixed points in the weak case, tensoring by the witness commutes with the fixed-point limit. In the strong case, recover the torsor algebra by Amitsur descent along the witness, then commute its fixed-point limit with that Amitsur limit. At every Amitsur stage the torsor stays trivial and its fixed points are the stage algebra, since the translation action groupoid of the regular action is contractible. The Amitsur limit of these stage algebras recovers the unit.

8.5 Periodic rings and higher local systems

Rognes proves that periodic complexification $KO \to KU$, with the complex-conjugation action, is a faithful C_2-Galois extension [7, Proposition 5.3.1]. Its

faithfulness can also be seen from the cofiber sequence

$$\Sigma KO \xrightarrow{\eta} KO \longrightarrow KU.$$

If $KU \otimes_{KO} M$ vanishes, the image of the stable Hopf element η acts invertibly as a degree-one endomorphism of M. A sufficiently high power of η is zero, so this invertibility forces $M \simeq 0$. The map on π_0 is the identity of $\mathbf{Z}$, while $\pi_1 KO \cong \mathbf{Z}/2$ and $\pi_1 KU = 0$. Thus this extension fails spectral flatness, yet gives module and algebra descent by Theorem 8.3. It lies in the stable finite-cover category. This explains why the connective hypothesis of the comparison with geometric finite étale algebras cannot simply be removed.

The finite Galois groupoid of a stable category can miss homotopy visible to its module objects. Consider $X = K(\mathbf{Z}, 2)$, an algebraically closed characteristic-zero field k, and the pointwise tensor category $\mathcal{C} = \mathrm{Fun}(X, \mathrm{Mod}_{Hk})$. Its commutative algebras are functors to CAlg_{Hk}. Evaluation preserves finite products, localization at idempotents, and the thick tensor ideal operations defining descendability. Every finite cover of $\mathcal{C}$ therefore takes values in finite covers of Hk, which are finite products of Hk. Its image lands in the maximal groupoid of these products, a 1-truncated space, and X is simply connected. Thus each such cover is constant up to equivalence. Conversely a constant product is already split by the unit. Maps between the constant products are the opposites of ordinary maps of finite sets, since X is connected and the value-wise mapping spaces are discrete. The finite Galois groupoid of $\mathcal{C}$ is trivial.

For the fiber $M = Hk \oplus \Sigma Hk$, its endomorphism spectrum decomposes into two copies of Hk, one of ΣHk, and one of $\Sigma^{-1} Hk$. The space $\mathrm{Aut}(M)$ is a union of components of its endomorphism space, so its identity component has $\pi_1 \cong (k, +)$ and no higher homotopy. The simply connected cover of $B\,\mathrm{Aut}(M)$ is therefore $K(k, 2)$. A nonzero additive homomorphism $\mathbf{Z} \to k$ gives a map

$$K(\mathbf{Z}, 2) \longrightarrow K(k, 2) \longrightarrow B\,\mathrm{Aut}(M)$$

which is nonzero on π_2. It classifies a nonconstant Hk-module local system. Its higher transport is invisible to the stable finite Galois groupoid, as happens for the fundamental progroupoid and the full shape in the space-valued theory.

Part IV

Reflections and factorization

CHAPTER 9

Reflective Galois structures and factorization

In Borceux and Janelidze's categorical Galois theorem, triviality is measured by a slice adjunction and splitting is found by changing the base [1, Chapter 5]. The higher analogue needs two independent ingredients. A slice adjunction specifies the chosen class of trivial objects. Effective descent says which objects and maps over a cover glue across its entire Čech nerve. Separating these ingredients makes it possible to use the same argument with rings, group objects, or spaces while stating their individual descent hypotheses explicitly.

9.1 The adjunction seen from a varying base

Let $L\colon \mathcal{C} \rightleftarrows \mathcal{D}\colon R$ be an adjunction of ∞-categories with unit η and counit ϵ. Assume both categories admit pullbacks. Over $B \in \mathcal{C}$, use the functors

$$L_B\colon \mathcal{C}_{/B} \longrightarrow \mathcal{D}_{/LB}, \qquad L_B(X \xrightarrow{f} B) = (LX \xrightarrow{Lf} LB),$$
$$R_B\colon \mathcal{D}_{/LB} \longrightarrow \mathcal{C}_{/B}, \quad R_B(D \xrightarrow{d} LB) = (B \times_{RLB} RD \longrightarrow B).$$

The pullback uses η_B. All slices contain the full mapping spaces over the indicated bases.

Proposition 9.1 (Slice adjunction and base change)**.** *There are adjunctions $L_B \dashv R_B$. For every $h\colon B' \to B$ there is a natural equivalence*

$$h^* R_B \simeq R_{B'}(Lh)^*.$$

These equivalences are coherent for chains of base changes. If Σ_h denotes postcomposition, there is also an equivalence $L_B \Sigma_h \simeq \Sigma_{Lh} L_{B'}$.

Proof. Over B, a map from X to $B \times_{RLB} RD$ is the same as a map $X \to RD$ with a specified identification of the composite to RLB with $\eta_B f$. The original adjunction identifies the space of such maps with the homotopy fiber

$$\mathrm{hofib}_{Lf}\left(\mathrm{Map}_{\mathcal{D}}(LX, D) \longrightarrow \mathrm{Map}_{\mathcal{D}}(LX, LB)\right).$$

This is the desired mapping space in $\mathcal{D}_{/LB}$. The unit at f has underlying map $X \to B \times_{RLB} RLX$ induced by (f, η_X). To check base change, use the fact that R preserves pullbacks and substitute

$$R(LB' \times_{LB} D) \simeq RLB' \times_{RLB} RD$$

in the formula for $R_{B'}$. Naturality of η and pasting of pullbacks identify the result with $B' \times_{RLB} RD$. Iterated identifications all arise from the limit of the same diagram, so their coherence follows from the universal cones. The assertion about postcomposition follows directly from functoriality of L. □

For a restricted Galois structure, let $\mathcal{E}$ in $\mathcal{C}$ and $\mathcal{F}$ in $\mathcal{D}$ be classes of maps containing equivalences and closed under composition and pullback. Write $\mathcal{C}_{\mathcal{E}/B}$ for the full subcategory of $\mathcal{C}_{/B}$ spanned by maps in $\mathcal{E}$, and similarly for $\mathcal{D}_{\mathcal{F}/LB}$. If L takes $\mathcal{E}$ to $\mathcal{F}$ and R takes $\mathcal{F}$ to $\mathcal{E}$, the displayed slice adjunction and its base-change equivalences restrict to these full subcategories. The full-subcategory condition matters. Maps between the selected objects over a base retain their higher morphisms whether or not their underlying arrows belong to $\mathcal{E}$.

Definition 9.2 (Admissibility and splitting)**.** The structure is admissible when the restricted R_B is fully faithful at each base. A map over B is *trivial* if the unit of the restricted slice adjunction at that map is an equivalence. A map over B is *split by* $p\colon U \to B$ if its pullback to U is trivial. It is a *covering* relative to a specified class of descent maps if one of those maps splits it.

Full faithfulness is equivalent to invertibility of the slice counit. It makes the full subcategory $\mathrm{Triv}(B)$ of trivial maps the essential image of R_B. A map is trivial exactly when its unit square

$$\begin{array}{ccc} X & \xrightarrow{\eta_X} & RLX \\ \downarrow & & \downarrow \\ B & \xrightarrow{\eta_B} & RLB \end{array}$$

is a pullback. Proposition 9.1 shows that trivial maps are stable under base change. Their defining square is computed using the global unit, but the proof of stability uses the right slice adjoints. There is no need to assume that L preserves all pullbacks.

Lemma 9.3 (A practical admissibility criterion)**.** *If R is fully faithful, the unrestricted structure is admissible precisely when L preserves every pullback of the form*

$$B \times_{RLB} RD$$

where the first arrow is η_B and the other is the image of $D \to LB$. In this assertion a preserved pullback has comparison map $L(B \times_{RLB} RD) \to D$.

Proof. Under the counit equivalences $LRD \simeq D$ and $LRLB \simeq LB$, the image of the cospan has first arrow equivalent to the identity of LB by the triangle identity. Its limit is D. The comparison is precisely the counit of $L_B \dashv R_B$. Invertibility of these counits for all B, D is the assertion. □

Admissibility can be checked with a smaller class $\mathcal{F}$ when the problem concerns only extensions. For instance, a reflection may be useful on slices of surjections even if its behavior on slices of all arrows has not been established. Effective descent on the selected slices requires both unrestricted descent and detection of the selected arrow condition, or a descent proof directly on the selected slices.

9.2 The comparison over one descent map

For $p\colon U \to B$ in $\mathcal{E}$, let $U_\bullet$ be its Čech nerve, with $U_k = U^{\times_B (k+1)}$. Use the pullback functors along its simplicial structure maps to form a cosimplicial diagram of restricted slices. Say that p is an effective descent map for $\mathcal{E}$ if the comparison

$$\mathcal{C}_{\mathcal{E}/B} \longrightarrow \lim_{[k]\in\Delta} \mathcal{C}_{\mathcal{E}/U_k}$$

is an equivalence. When every map is allowed and $\mathcal{C}$ is an ∞-topos, every effective epimorphism satisfies this condition by the van Kampen property of its Čech realization [2, Section 6.1]. More generally, if unrestricted descent holds and membership in $\mathcal{E}$ is detected after pullback along p, restricting the equivalence gives the indicated condition.

Theorem 9.4 (Relative higher Galois comparison)**.** *In an admissible structure, let $p\colon U \to B$ be an effective descent map for $\mathcal{E}$. The full subcategory $\mathrm{Split}(p)$ of maps split by p satisfies*

$$\mathrm{Split}(p) \simeq \lim_{[k]\in\Delta} \mathcal{D}_{\mathcal{F}/LU_k}.$$

The diagram on the right is formed with pullbacks along maps of $L(U_\bullet)$.

Proof. Restrict the unrestricted comparison for the selected slices to the full subcategories $\mathrm{Triv}(U_k)$. A point of the limit of the slices is a coherent family whose object in any degree k is equivalent to the pullback of the degree-zero object along any vertex. Consequently, if its object in degree zero is trivial, all its objects are trivial. Conversely a family consisting of trivial objects has a trivial degree-zero object. Limits of diagrams of full subcategories are full in the limit of the ambient diagrams. These observations give the full essential image of $\mathrm{Split}(p)$ and preserve its mapping spaces. Finally the equivalences $R_{U_k}\colon \mathcal{D}_{\mathcal{F}/LU_k} \simeq \mathrm{Triv}(U_k)$ respect restriction by Proposition 9.1. Taking their limit gives the result. □

At degree one the family has an equivalence between two pullbacks of a single object over LU. At degree two it gives a specified homotopy between composite gluing and direct gluing. The comparison retains the higher data as well. An ordinary category of strict descent isomorphisms is sufficient in

a nerve of an ordinary category, where the mapping spaces are discrete. In that case the degree-two relation is an equality and determines all the higher compatibility for an object. In categories with nondiscrete mapping spaces the totalization is intrinsic to the assertion.

There are two distinct possible simplifications of this comparison. If L preserves the pullbacks that form $U_\bullet$, the simplicial object $L(U_\bullet)$ is the Čech nerve of Lp. If it further presents an effective groupoid in a suitable category $\mathcal{D}$, the totalization can be expressed as its category of actions. If unrestricted descent also holds for Lp in $\mathcal{D}$, then the result is the entire unrestricted slice over LB. Each step uses a separate property. Applying a reflection termwise to a Čech nerve is meaningful even if its image fails the groupoid Segal conditions, and the comparison theorem still holds in the form given above.

Proposition 9.5 (The monadic comparison)**.** *Suppose that p^* on the selected slices is monadic and admits the restricted left adjoint Σ_p. Let $T = p^*\Sigma_p$ be its monad on $\mathcal{C}_{\mathcal{E}/U}$. If T sends trivial objects to trivial objects, it restricts to a monad T_{triv}, and*

$$\mathrm{Split}(p) \simeq \mathrm{Alg}_{T_{\mathrm{triv}}}(\mathrm{Triv}(U)).$$

Proof. The full subcategory of trivial objects contains every iterate of T on any one of its objects. Hence the unit and multiplication of the monad restrict to it with their coherences. A coherent algebra structure on such an object uses only these iterates and maps between them. Full faithfulness gives the same spaces for the operations, compatibility diagrams, and algebra morphisms in the smaller and larger categories. The restricted algebra category is thus the full subcategory of $\mathrm{Alg}_T(\mathcal{C}_{\mathcal{E}/U})$ on algebras whose underlying objects are trivial. By monadic comparison, underlying objects correspond to pullbacks of objects over B. This is exactly the splitting condition. □

The stability hypothesis is equivalent to requiring that $\Sigma_p R_U(D)$ be split by p for every allowed $D \to LU$. This recovers the extra condition in the monadic Galois-descent version of the ordinary theorem [1, Section 5.1]. The Čech comparison uses a different stated hypothesis, namely an effective descent comparison to the full cosimplicial diagram. This formulation avoids assuming that every monad arising from a chosen class already has an identified groupoid model.

Proposition 9.6 (Change of splitting cover)**.** *Suppose the specified descent maps are stable under pullback and composition. Coverings are stable under base change. If $q\colon V \to B$ is a specified descent map and q^*f is a covering over V, then f is a covering whenever f is allowed. If $V \to U$ is a map over B, then* $\mathrm{Split}(p) \subseteq \mathrm{Split}(V \to B)$.

Proof. For a new base $B' \to B$, pull back a splitting map $U \to B$ and use stability of triviality. If $W \to V$ splits q^*f, its composite $W \to B$ is a specified

descent map that splits f. The refinement statement follows from the same stability of triviality. All assertions concern full subcategories and respect the coherent pullback functors between their comparisons. □

An allowed effective descent map p is *normal* for this structure when p^*p is trivial. This belongs to the relative terminology. A normal splitting cover is a descent map that splits itself, and its pullbacks are again normal whenever the descent class is stable under pullback. The definition has no dependence on a choice of geometric point or separable closure. Chapter 10 relates it to centrality in the ordinary group calculation.

9.3 Two factorization systems arising from truncation

Return to a reflection, using all arrows and identifying R with inclusion. Write L also for the endofunctor RL under this convention. Let $\mathcal{M}$ be the trivial maps and $\mathcal{W}$ the maps inverted by L.

Every orthogonality assertion below concerns a contractible space of fillers in every coherent commutative square. Equivalently, $a\colon A \to B$ is orthogonal to $m\colon X \to Y$ if the natural map

$$\operatorname{Map}(B, X) \longrightarrow \operatorname{Map}(A, X) \times_{\operatorname{Map}(A,Y)} \operatorname{Map}(B, Y)$$

is an equivalence. This definition retains the space of commutativity data of a square as well as the space of diagonals.

Theorem 9.7 (Admissible reflective factorization)**.** *For an unrestricted admissible reflection there is an orthogonal factorization system $(\mathcal{W}, \mathcal{M})$ with factorization*

$$A \longrightarrow B \times_{LB} LA \longrightarrow B \qquad (f\colon A \to B).$$

The factorizations are functorial and $\mathcal{W}$ is stable under pullback along $\mathcal{M}$.

Proof. The right arrow is $R_B(L_B f)$. Admissibility makes its L-image the arrow $LA \to LB$ up to equivalence, with counit inverse to the image of the left arrow by the triangle identity. Hence the left arrow lies in $\mathcal{W}$. To see orthogonality, take a lifting problem against $X = Y \times_{LY} LX \to Y$. For a fixed bottom arrow $B \to Y$, lifts amount to compatible maps $B \to LX$ over LY. For $a \in \mathcal{W}$, restriction along a is an equivalence on maps into either of the local objects LX, LY. It is therefore an equivalence on their homotopy fibers over the prescribed composite to LY, proving orthogonality. The class $\mathcal{W}$ is closed under retracts, and so is $\mathcal{M}$, since the natural comparison from an object to its slice reflection carries retract diagrams to retract diagrams. The usual orthogonality-retract argument now identifies each class with the full orthogonal of the other. Functoriality follows from the units and pullback universal properties.

For the stability statement let $m\colon C \to B$ be trivial and $a\colon A \to B$ inverted. Write $C \simeq B \times_{LB} D$ for $D \in \mathcal{D}_{/LB}$. Then $A \times_B C \simeq A \times_{LA} (LA \times_{LB} D)$. Applying admissibility at A and B gives the comparison $LA \times_{LB} D \to D$ under L, an equivalence. □

Stable units mean that L preserves pullbacks over every local object in the ambient category, including pullbacks in which neither of the objects over that base is local. This implies admissibility by the practical criterion. Left exactness requires finite-limit preservation over arbitrary bases and is stronger. If L is left exact, $\mathcal{W}$ is stable under arbitrary pullback. Conversely, if $\mathcal{C}$ has finite limits and $\mathcal{W}$ is pullback stable, L is left exact. Indeed the reflection fixes the terminal object, since it is local. For a pullback $A \times_B C$, the map to $A \times_{LB} C$ is a pullback of the diagonal of $B \to LB$. That diagonal lies in $\mathcal{W}$ by pullback stability and two-out-of-three, since either projection $B \times_{LB} B \to B$ lies in $\mathcal{W}$ and the composite with the diagonal is the identity. The map

$$A \times_{LB} C \longrightarrow LA \times_{LB} LC$$

is the composite of pullbacks of the two units, so is also in $\mathcal{W}$. The final object is local, since the right adjoint preserves limits. This proves preservation of pullbacks.

Corollary 9.8 (Left-exact localizations)**.** *Let $L\colon \mathcal{X} \to \mathcal{Y}$ be an accessible left-exact localization of an ∞-topos, and allow all arrows. For any effective epimorphism $p\colon U \to B$, a map over B is split by p if and only if it is already trivial.*

Proof. The localized category $\mathcal{Y}$ is an ∞-topos. Left exactness identifies $L(U_\bullet)$ with the Čech nerve of Lp. Since L is a left adjoint, it preserves the colimit of the original nerve, so Lp is also an effective epimorphism. The right-hand side of Theorem 9.4 is therefore equivalent to $\mathcal{Y}_{/LB}$ by effective descent there. Under that equivalence the trivial object $R_B(D)$ corresponds to the descent datum of D, since the base-change formula of Proposition 9.1 identifies its restriction in every degree. Thus the equivalence on slices realizes the inclusion of the trivial objects into the split objects as an equivalence. □

Definition 9.9 (Relative truncation)**.** For $n \geq -1$ and an ∞-topos $\mathcal{X}$, the relative n-truncation of $f\colon A \to B$ is the reflection of A into the full subcategory of n-truncated objects in $\mathcal{X}_{/B}$. Write its total object as $\tau^B_{\leq n} A$. A map is n-connected if its relative n-truncation is terminal in its target slice.

The pair of n-connected and n-truncated maps is an orthogonal factorization system in an ∞-topos, with both classes stable under pullback [2, Sections 5.5.6 and 6.5.1]. A lifting problem with an n-truncated right arrow over B depends only on the relative n-truncation of the left arrow by the slice universal property. For $A \to B$, the first map in the relative truncation factorization is n-connected,

and the second is n-truncated. The factorization and lifting conditions are often called the truncation modality [17]. The indexing here pairs n-connected maps with n-truncated maps, and an effective epimorphism is (-1)-connected.

Global n-truncation, with codomain the n-truncated objects, has stable units. Indeed, given $A \to D \leftarrow C$ with D n-truncated, factor the comparison

$$A \times_D C \longrightarrow \tau_{\leq n} A \times_D \tau_{\leq n} C$$

as two pullbacks of n-connected units. The target is n-truncated because the full subcategory of truncated objects is closed under limits, and the composite is n-connected. This identifies the comparison with a truncation. The reflective factorization of Theorem 9.7 therefore applies. Its intermediate object

$$B \times_{\tau_{\leq n} B} \tau_{\leq n} A$$

is constructed from the global truncations. Relative truncation uses $\tau^B_{\leq n} A$. They have different universal properties over the unreplaced B.

Example 9.10 (Relative components of paths)**.** Take $\mathcal{X} = \mathcal{S}$, $B = S^1$ and $A = *$ with the basepoint. The global 0-truncations of A and B coincide, so the global reflective factorization has intermediate object S^1. The fiber of $* \to S^1$ over the basepoint is $\Omega S^1 \simeq \mathbb{Z}$, where $\mathbb{Z}$ is discrete. Thus $* \to S^1$ is already 0-truncated over S^1, and the relative factorization has intermediate object $*$. Since this fiber has infinitely many components, the map cannot be 0-connected. Pullback to the basepoint also exhibits failure of pullback stability for the maps inverted by global 0-truncation.

9.4 Localizing the reflective pair

For global n-truncation, its trivial maps are the base changes of maps between n-truncated objects. Their fibers are n-truncated. When the base admits a cover by an n-truncated object, the converse holds after passing to that cover. This includes every base space for $n \geq 0$, since the coproduct of basepoints of its components gives a discrete effective epimorphism.

Proposition 9.11 (Local triviality and stable inversion)**.** *Fix $n \geq -1$ in an ∞-topos $\mathcal{X}$. Every n-connected arrow is inverted by global n-truncation after arbitrary pullback. Every arrow that becomes trivial for global n-truncation over an effective epimorphism is n-truncated. Over a base admitting an effective epimorphism from an n-truncated object, both converses hold.*

Proof. If f is n-connected, its pullbacks are n-connected. For any such pullback $a\colon X \to Y$ and any globally n-truncated object D, the projection $Y \times D \to Y$ is an n-truncated arrow. Orthogonality then identifies maps $X \to D$ with maps $Y \to D$. Global truncation thus inverts a.

For the reverse implication over a base of the indicated kind, write f as $A \xrightarrow{e} T \xrightarrow{m} B$ in the relative n-connected, n-truncated factorization. If every pullback of f is inverted by global truncation, the same holds for m since it holds for e. Pull back to an effective epimorphism $U \to B$ with U n-truncated. The pullback of m has n-truncated domain and codomain, so inversion by global truncation makes it an equivalence already. Then m is an equivalence by effective descent, proving f is n-connected.

Trivial arrows are n-truncated since relative truncation is stable under pullback. This property descends along effective epimorphisms by induction on the truncation of the relative diagonal, starting with descent of equivalences. Descent of equivalences is immediate from the slice comparison over the Čech nerve. Finally, if $U \to B$ is an effective epimorphism with U n-truncated and f is n-truncated, then $U \times_B A$ is n-truncated by composition with $U \to 1$. Both domain and base of the new arrow are local, so its unit square is a pullback. □

Over spaces the last result gives a homotopical form of monotone–light factorization. For $n = 0$, the left class consists of maps with nonempty connected homotopy fibers, and the right class consists of maps with discrete homotopy fibers. This concerns the homotopy theory of spaces. The original result concerns compact Hausdorff spaces and ordinary point inverses [1, Section 5.8]. Those settings have different base change, and passing from a topological map to its homotopy type need not preserve its point inverse. The corresponding constructions agree in the formal operation of testing connectedness after base change and testing the right class by local splitting.

CHAPTER 10

Central extensions and higher bands

Categorical centrality in the ordinary group calculation arises by splitting for abelianization. A higher central extension also has an interpretation as a principal bundle over a delooping, where its band keeps track of all the homotopy of the kernel. The constructions are compatible for discrete groups. To compare them in higher degrees, it is useful to prove the ordinary result first and to keep the higher notion of a band explicit.

10.1 The ordinary abelianization comparison

Use $E_{\mathrm{ab}} = E/[E,E]$ and write $a_E \colon E \to E_{\mathrm{ab}}$. For groups, admit surjections in both the source and target of the abelianization reflection. Its slice right adjoint over G takes a surjection $D \to G_{\mathrm{ab}}$ of abelian groups to $G \times_{G_{\mathrm{ab}}} D \to G$. The following computation works for this class and makes the admissibility criterion concrete.

Lemma 10.1 (Abelianization over a discrete base)**.** *Let $P = G \times_{G_{\mathrm{ab}}} D$, with $D \to G_{\mathrm{ab}}$ surjective and D abelian. Then*

$$[P,P] = [G,G] \times \{1\} \qquad \textit{and} \qquad P_{\mathrm{ab}} \cong D.$$

In particular, abelianization is admissible on these surjective slices.

Proof. The two projections from P are surjective. Every commutator in P has second coordinate 1 since D is abelian. Conversely, given $x, y \in G$, choose lifts $(x,d),(y,d') \in P$. Their commutator is $([x,y],1)$. Since such commutators generate $[G,G]$, this proves the first equality. The second projection has kernel $[G,G] \times \{1\}$ and is surjective, yielding the second assertion. This identifies the slice counit with an isomorphism. □

Every surjection of ordinary groups is an effective descent map for surjective arrows. One can see this directly by viewing a group as a reduced simplicial set satisfying the Segal conditions and the invertibility condition. A surjection $U \to G$ induces a pointwise surjection between their group nerves, since the map in degree r is $U^r \to G^r$. Ordinary effective descent for a surjection of sets comes from forming equivalence classes under its compatible gluing relation. Applied degree by degree, this gives descent for diagrams of sets, hence a descended simplicial set over the nerve of G. The Segal comparisons are maps over the products G^r. Their pullbacks to U^r are isomorphisms, so they are

isomorphisms. The same argument shows that the shear comparison (d_2, d_1) in degree two is an isomorphism, which recovers invertibility in the descended monoid. Surjectivity in degree one is also detected after pullback. Thus the descended object is an allowed group surjection, with morphisms descending by the same degreewise argument.

Theorem 10.2 (Trivial, split, and central group extensions)**.** *For a group surjection $e\colon E \to G$ with kernel K, the following statements hold.*

1. *The surjection is trivial for abelianization exactly when $K \cap [E, E] = 1$.*
2. *The surjection is split by some surjection exactly when K is central in E.*
3. *A central surjection is split by itself and thus is normal for this Galois structure.*

The reflection of e into central extensions over G is the quotient $E/[K, E] \to G$.

Proof. Consider the unit comparison

$$\phi\colon E \longrightarrow G \times_{G_{\mathrm{ab}}} E_{\mathrm{ab}}.$$

It is surjective. To lift (g, u), first choose x lifting g. The difference between u and $a_E(x)$ lies in the kernel of $E_{\mathrm{ab}} \to G_{\mathrm{ab}}$. This kernel is the image of K because the surjection $[E, E] \to [G, G]$ allows any representative mapping to $[G, G]$ to be corrected to an element of K without changing its abelianized class. Multiplying x by a lift of the difference in K gives the desired lift. The kernel of ϕ is visibly $K \cap [E, E]$. This proves the first statement. It also implies that any trivial extension is central, since $[K, E] \subseteq K \cap [E, E]$ by normality of the kernel.

Suppose $p\colon U \to G$ splits e. The kernel of $U \times_G E \to U$ is canonically K and is central in $U \times_G E$ by the first statement. The second projection $U \times_G E \to E$ is surjective, so each $y \in E$ is the image of an element commuting with every element of the lifted K. Therefore K is central in E.

Conversely, centrality makes the map

$$E \times K \longrightarrow E \times_G E, \qquad (x, k) \longmapsto (x, xk)$$

a group isomorphism over E. The projection $E \times K \to E$ is trivial by the first statement, since K is abelian and $[E \times K, E \times K] = [E, E] \times \{1\}$. Hence e is split by itself.

Finally, $[K, E]$ is a normal subgroup of E contained in K. In the quotient by this subgroup the image of K commutes with the image of E. Every map from e to a central extension over G kills each commutator $[k, x]$, since k maps into the central kernel of the target. It therefore factors uniquely through this quotient. This proves the universal property. □

This gives the ordinary relative comparison and its algebraic content [1, Section 5.2]. A general treatment of categorical centrality for admissible reflections appears in [19]. The split object in the sense of Chapter 9 is an extension that becomes trivial after a chosen change of base. An ordinary group splitting by a section is a separate property. For instance, a central extension that has a group-homomorphic section is isomorphic over its base to a direct product with its kernel, whereas the classification below also includes nonsplit group extensions.

10.2 Higher central extensions as pointed principal bundles

Work in an ∞-topos $\mathcal{X}$. Write BG for the delooping of a group object G. The group object and pointed connected object descriptions are equivalent via B and Ω [2, Section 6.1]. If A is a commutative group object, BA itself has a commutative group structure, so its delooping B^2A is defined. One way to construct this structure is to apply the bar construction to the grouplike commutative monoid structure on A. Realization preserves the finite products in the bar diagram, so the commutative structure survives and its underlying monoid is grouplike because it is connected.

Definition 10.3 (Framed central extension)**.** Fix a group object G and a commutative group object A. A framed central extension of G by A is a principal BA-bundle $P \to BG$ with a point over the basepoint of BG. Equivalences preserve the principal structure and the chosen point, including their coherences. The associated group extension is $A \to \Omega P \to G$.

The choice of a principal action is the band in this definition. The basepoint trivialization identifies the fiber over the basepoint with BA using that action. Taking loops then identifies the kernel with A. The structure is thus a higher analogue of prescribing an abelian coefficient group and keeping its identification with the central kernel fixed.

Theorem 10.4 (Classifying space of framed extensions)**.** *The space of framed central extensions of G by A is*

$$\operatorname{Map}_{\mathcal{X}_*}(BG, B^2A).$$

For the extension specified by $c\colon BG \to B^2A$, there is a pointed fiber sequence

$$BA \longrightarrow \operatorname{hofib}(c) \longrightarrow BG \xrightarrow{c} B^2A.$$

Its associated group map has π_0-surjective augmentation $\pi_0\Omega \operatorname{hofib}(c) \to \pi_0 G$, and the image of $\pi_0 A$ is central in $\pi_0\Omega \operatorname{hofib}(c)$.

Proof. The universal principal BA-bundle is $1 \to B^2A$. Its Čech nerve is the bar construction of BA. Pullback classifies principal bundles, with mapping spaces,

by the principal-bundle comparison, applied to BA [4]. At the basepoint of BG a lift to the pullback of the universal bundle is exactly a specified identification between the classifying map at that basepoint and the basepoint of B^2A. This gives the asserted equivalence of spaces and the fiber sequence.

The fibers of $P \to BG$ are locally BA, which is connected. Relative connectedness can be checked after an effective epimorphism by relative truncation and effective descent. Consequently $P \to BG$ is 0-connected and P is connected, since BG is connected. Taking loops preserves the pointed pullback describing the fiber and gives the group kernel A. The long exact homotopy sequence, or its internal sheaf form, gives the asserted surjectivity because BA has terminal π_0.

To see centrality, loop the principal action $BA \times P \to P$ and use preservation of products by loops. The resulting map $A \times \Omega P \to \Omega P$ is a morphism of group objects and its restrictions to the two factors are the kernel inclusion and the identity. The images of those factors commute on π_0, since the source is a product of groups there. The principal action records a choice of coherence beyond this equality on components. □

For a general higher group G, the kernel of $\pi_0 A \to \pi_0 \Omega P$ is the image of the connecting homomorphism $\pi_1 G \to \pi_0 A$. The homotopy fiber assertion incorporates this correction to the component-level group sequence.

Proposition 10.5 (Recognition from a coherent action)**.** *Let $p\colon E \to G$ be a map of group objects surjective on the sheaves of components. Suppose a coherent action of A on E over G is supplied, with the simplicial action diagram $A^\bullet \times E$ formed in group objects and with A carrying its specified commutative group structure. If the action shear $A \times E \to E \times_G E$ is an equivalence, the resulting pointed group extension has a framed central structure.*

Proof. Deloop each term of the action diagram. The equivalence between group objects and pointed connected objects preserves products, so the resulting simplicial action diagram is $BA^\bullet \times BE$. The map $BE \to BG$ is 0-connected, since the long exact sequence and surjectivity of $\pi_0 p$ give connected homotopy fibers at the basepoint, and the basepoint $1 \to BG$ is an effective epimorphism. In particular $BE \times_{BG} BE$ is connected. The delooped shear $BA \times BE \to BE \times_{BG} BE$ is therefore a map between pointed connected objects. Its loop comparison is the given equivalence, so it is an equivalence as well. This proves that $BE \to BG$ is a principal BA-bundle, with the frame given by the basepoint of BE. □

Corollary 10.6 (Ordinary central extensions and their automorphisms)**.** *For ordinary groups G and ordinary abelian groups A regarded as discrete objects of $\mathcal{S}$, framed extensions give exactly the ordinary central extensions of G with kernel identified with A. Their isomorphism classes are $H^2(G; A)$ for the trivial action,*

and the automorphism group of an extension with kernel and quotient fixed is $H^1(G; A)$.

Proof. The fiber sequence $BA \to P \to BG$ shows that P is a connected 1-type. Thus ΩP is an ordinary group. Its surjectivity, injectivity of the kernel, and centrality follow from the theorem and the long exact sequence. Conversely an ordinary central extension $A \to E \to G$ yields the group homomorphism $A \times E \to E$ by multiplication of the central inclusion with the identity. Delooping this homomorphism gives a left BA-action on BE over BG. The map

$$BA \times BE \longrightarrow BE \times_{BG} BE$$

is an equivalence. Both sides are pointed and connected, the right side by surjectivity of $E \to G$, and looping identifies the map with the ordinary torsor comparison $A \times E \to E \times_G E$. The latter is bijective by uniqueness of the kernel element relating two lifts. Also $BE \to BG$ is an effective epimorphism since both spaces are connected. This gives the principal bundle with its point. The constructions are inverse on objects and isomorphisms, using the equivalence between groups and pointed connected 1-types. The central action in this degree is forced by the kernel inclusion and the group multiplication, so it adds no independent data.

Finally $B^2A = K(A, 2)$. Homotopy groups of $\mathrm{Map}_*(BG, K(A, 2))$ are the reduced cohomology groups in degrees $2 - i$ for $0 \le i \le 2$. Here degree zero agrees with unreduced cohomology in positive cohomological degree, and $\tilde{H}^0(BG; A) = 0$ because BG is connected. Translation by an extension class identifies its based automorphism group with the first homotopy group at the zero class. This gives the two identifications. □

One can read this classification directly in a group bar resolution. A section of sets $s\colon G \to E$ with $s(1) = 1$ determines $\alpha(g, h) \in A$ by

$$s(g)s(h) = \alpha(g, h)s(gh).$$

Associativity and centrality give

$$\alpha(g, h) + \alpha(gh, k) = \alpha(h, k) + \alpha(g, hk).$$

Conversely a normalized solution defines the multiplication $(a, g)(b, h) = (a+b+\alpha(g, h), gh)$ on $A \times G$. The identity and inverses follow from the normalization and the cocycle equality, and its projection has central kernel A. Changing the section modifies α by the group coboundary of a normalized function $G \to A$. The bar nerve of G has r-simplices G^r, so its ordinary normalized cochain differential in degree two is exactly that equation. This comparison identifies the group cocycle with the classifying map and explains which operation in the proof is replaced by coherent descent in higher degrees.

10.3 Higher cohomology and the role of trivialization

If A_0 is an ordinary abelian group object in an ∞-topos, let $K(A_0, r)$ denote its Eilenberg–Mac Lane object in degree r. For the commutative group object $A = K(A_0, n)$ the classifying space in Theorem 10.4 is

$$\mathrm{Map}_{\mathcal{X}_*}(BG, K(A_0, n+2)).$$

The shifts come from using one delooping to pass from a group object to its classifying object and a second to classify the principal action on that object. In spaces, components are $H^{n+2}(BG; A_0)$ and higher homotopy groups are the reduced cohomology in the corresponding lower degrees. In a general topos, the pointed mapping space is the homotopy fiber of restriction along $1 \to BG$ over the zero section. This description records the chosen trivialization over the basepoint even when cohomology of the terminal object is nontrivial.

A variable coefficient object can instead be specified over BG. Let $\mathcal{A}$ be an abelian group object in the 0-truncated part of $\mathcal{X}_{/BG}$ and choose an identification of its restriction at the basepoint with A_0. Eilenberg–Mac Lane objects can be formed in this slice. A principal $K_{/BG}(\mathcal{A}, n+1)$-bundle in the slice is classified by a section of $K_{/BG}(\mathcal{A}, n+2)$. The proof is the same principal-bundle comparison, now with terminal object $BG \to BG$. Taking the homotopy fiber of restriction to the basepoint over the zero section adds a point in the pulled-back torsor. Looping the resulting pointed total object gives a group map with kernel $K(A_0, n)$. The action over BG is a relative action and the coefficient object may vary along loops. For the constant coefficient object it is the global principal action in the theorem. Thus the slice construction also tracks coefficient twisting without imposing centrality on an arbitrary variable band.

Example 10.7 (An extension with a higher obstruction class)**.** In spaces take $G = \Omega S^3$ and $A = K(\mathbb{Z}, 1)$. Let $c\colon S^3 \to B^2A = K(\mathbb{Z}, 3)$ represent the integral generator. The map on π_3 induced by c is an isomorphism, so the boundary $\pi_3(S^3) \to \pi_2(BA)$ in the fiber sequence $BA \to P \to S^3$ is an isomorphism. This gives $\pi_2(P) = 0$. For the zero class the total object is $BA \times S^3$ and has $\pi_2 \cong \mathbb{Z}$. Both looped extensions have trivial component groups. The principal class distinguishes them and is detected by higher homotopy.

Normalization of the band is essential to the comparison with abelianization. For ordinary central group extensions the conjugation action on A is trivial, with no higher choices of coherence. In the higher case it is possible to ask for a neutralizable coefficient system, a chosen null homotopy of an action, or a specified principal BA-bundle. Definition 10.3 chooses the principal-bundle condition. The classifying theorem concerns that structure and its morphisms.

10.4 Abelianization after passing to group objects in spaces

Two extensions of the abelianization functor illustrate why the band construction has an independent role. The first applies ordinary abelianization to π_0. The second uses the left adjoint to the forgetful functor from infinite loop spaces to group objects in spaces. The first is a reflection onto a discrete subcategory. The second has a counit that can change the positive homotopy groups of a connective spectrum.

Proposition 10.8 (A discrete reflection on higher group objects)**.** *The functor from group objects in spaces to ordinary abelian groups given by $G \mapsto (\pi_0 G)_{\mathrm{ab}}$ is an admissible reflection when the allowed maps are surjective on π_0 in the source and surjective in the target. A map split for this structure by an allowed effective descent map has 0-truncated fibers.*

Proof. Maps into a discrete abelian group factor through π_0 and abelianization by the two respective universal properties. This proves the reflection assertion. Given an allowed abelian group surjection $D \to (\pi_0 G)_{\mathrm{ab}}$, the underlying space of the slice right adjoint is the pullback $G \times_{(\pi_0 G)_{\mathrm{ab}}} D$. Since its base is discrete, its component set is $\pi_0 G \times_{(\pi_0 G)_{\mathrm{ab}}} D$. Lemma 10.1 then computes the counit. Closure of the allowed maps under pullback follows from stability of effective epimorphisms of spaces. Surjectivity of the projection to G follows from the same stability. The sliced adjunction is admissible.

Every trivial map is a pullback from a map of discrete abelian groups, so has discrete fibers. If an allowed map becomes trivial along a descent map $U \to G$, the underlying map of spaces $U \to G$ is an effective epimorphism since it is surjective on components. Truncation of fibers descends along such a map by Chapter 9. □

In fact, in this structure each allowed base change map is effective for the restricted Čech comparison. In the functor ∞-topos of simplicial spaces, the map of group nerves associated to $U \to G$ is an effective epimorphism. This holds pointwise because the map in simplicial degree r is $U^r \to G^r$. Descent in that functor category provides a simplicial object over the group nerve of G. The degree-zero condition is unchanged under pullback since both group nerves have terminal object in degree zero. The Segal comparisons over G^r and the invertibility comparison over G^2 become equivalences when pulled back along the effective epimorphisms $U^r \to G^r$ and $U^2 \to G^2$, respectively. Here the invertibility comparison is the shear (d_2, d_1), transported along the shear equivalence for the base group nerve. They are therefore equivalences prior to pullback. The descended object is a group nerve, and the required surjectivity on π_0 in degree one is detected after base change. The whole comparison is an equivalence on the selected full subcategories, with all spaces of morphisms.

Consequently a higher principal extension with a nondiscrete kernel, such as the preceding example, may have no interpretation as a covering for this

discrete Galois structure. Its fibers already prevent that comparison. The discrete reflection retains the ordinary group case proved above and provides one admissible choice when the intended fibers have no higher homotopy.

Proposition 10.9 (Infinite-loop abelianization in spaces)**.** *Write U for the forgetful functor from commutative group objects to group objects in spaces. Under the equivalence between commutative group objects and connective spectra, a left adjoint to U is*

$$F(G) = \Sigma^{-1}\Sigma^{\infty}BG.$$

Here Σ^{∞} is reduced suspension spectrum. The zeroth homotopy group of $F(G)$ is $(\pi_0 G)_{\mathrm{ab}}$. The counit is noninvertible at $H(\mathbb{Z}^2)$, so full faithfulness fails for U.

Proof. Under infinite-loop recognition, a connective spectrum V gives the infinite loop space $\Omega^{\infty}V$, with delooping $\Omega^{\infty}\Sigma V$ [6, Remark 5.2.6.26]. For pointed connected BG, the adjunction between reduced suspension spectrum and infinite loop space gives

$$\begin{aligned}\mathrm{Map}_{\mathrm{Grp}(\mathcal{S})}(G, U(V)) &\simeq \mathrm{Map}_{\mathcal{S}_*}(BG, \Omega^{\infty}\Sigma V)\\ &\simeq \mathrm{Map}_{\mathrm{Sp}}(\Sigma^{\infty}BG, \Sigma V)\\ &\simeq \mathrm{Map}_{\mathrm{Sp}}(\Sigma^{-1}\Sigma^{\infty}BG, V).\end{aligned}$$

The spectrum in the last source is connective because BG is connected. Its zeroth homotopy group is the degree-one stable homotopy group of BG, and stable Hurewicz identifies this with $H_1(BG;\mathbb{Z}) \cong (\pi_0 G)_{\mathrm{ab}}$.

For a counterexample to full faithfulness, use the connective Eilenberg–Mac Lane spectrum $H(\mathbb{Z}^2)$. Its underlying group object is discrete $\mathbb{Z}^2$, with delooping $S^1 \times S^1$. There is a stable splitting

$$\Sigma^{\infty}(S^1 \times S^1) \simeq \Sigma^{\infty}S^1 \oplus \Sigma^{\infty}S^1 \oplus \Sigma^{\infty}S^2.$$

Indeed the stable cofiber of the wedge inclusion is the suspension spectrum of the smash product, and the two projections split the wedge inclusion in spectra. Thus $F(U(H(\mathbb{Z}^2)))$ has a sphere spectrum suspended once as a summand. It has nonzero first homotopy group, while $H(\mathbb{Z}^2)$ has zero first homotopy group. The counit cannot be an equivalence, contrary to what full faithfulness would imply. □

The two constructions agree on the zeroth homotopy group and differ even when the input is discrete and already abelian. The higher band theorem uses delooping and effective groupoid descent directly, so its classifying space requires neither an admissible infinite-loop reflection nor a presentation by fiberwise discrete data.

Part V

Directed Galois theories

CHAPTER 11

Directed symmetries and generalized descent

A passage from a splitting object to its groupoid of changes of trivialization discards no information when all transports are invertible. There are also settings where a comparison has an intrinsic direction. A specialization of points can induce a map of fibers which is far from an equivalence. The directed analog of a Galois action is then a functor on a category with genuine noninvertible morphisms. This chapter distinguishes that construction from the generalized descent on precategories in [1, Chapter 7], and gives a higher form of both constructions.

11.1 Complete Segal objects and groupoid objects

Let $\mathcal{B}$ be an ∞-topos. A simplicial object $C_\bullet$ in $\mathcal{B}$ is a Segal object if the canonical map

$$C_n \longrightarrow C_1 \times_{C_0} C_1 \times_{C_0} \cdots \times_{C_0} C_1$$

is an equivalence for $n \geq 2$. The maps in each pullback are the target and the source of successive arrows. Here d_1 is the source and d_0 the target. The simplicial identities give identity and composition laws with all their coherences. In particular choosing a representative composition at degree two gives a contractible space of choices, and degree three provides the associativity homotopy for the compatible choices.

There is an object C_1^{eq} of equivalences. A generalized arrow over $T \in \mathcal{B}$ belongs to this subobject when it is invertible in the category encoded by the Segal space $\mathrm{Map}_{\mathcal{B}}(T, C_\bullet)$. Equivalence is an intrinsic property of an arrow. The object can be built by taking the object of coherent inverse data for the walking isomorphism. The space of choices of coherent inverses to an equivalence is contractible. Completeness means that the degeneracy map specifying identities induces an equivalence $C_0 \simeq C_1^{\mathrm{eq}}$. In contexts T this says that equivalences of objects are modeled by paths in the object space. Complete Segal objects are one model of small ∞-categories internal to $\mathcal{B}$ [11].

Remark 11.1. An internal groupoid object can have different object and arrow objects, as happens for a nontrivial Čech nerve. Such a groupoid object is a Segal object with all arrows invertible. It usually fails completeness in the complete Segal sense above. A complete Segal object with all arrows invertible

is equivalent to a constant simplicial object. An effective groupoid object in an ∞-topos is a presentation of its quotient, while a constant simplicial object on the quotient is already complete.

Working contextually is important. Global elements of objects of an ∞-topos need not detect equivalences. For example, a statement about fibers over all $c : T \to C_0$ can be tested at the generic point $c = \mathrm{id}_{C_0}$ in context C_0. A statement about fibers over all $1 \to C_0$ can hold vacuously while the assertion over C_0 fails. The recovery of internal constructions from their underlying global categories requires its own hypotheses.

11.2 Coherent presheaves and direction

Definition 11.2. A right fibration over an internal category $C_\bullet$ is a map $q : E_\bullet \to C_\bullet$ of complete Segal objects such that each terminal-vertex map gives a pullback square

$$\begin{array}{ccc} E_n & \longrightarrow & E_0 \\ \downarrow & & \downarrow \\ C_n & \longrightarrow & C_0. \end{array}$$

A left fibration uses the initial vertex instead. Write $\mathrm{RFib}_{\mathcal{B}}(C)$ and $\mathrm{LFib}_{\mathcal{B}}(C)$ for the respective full subcategories over C.

For a generalized edge $f : x \to y$ with parameters in T, the right-fibration square in degree one identifies a lift ending at an element of E_y with that element. Taking the source of such lifts defines a map $f^* : E_y \to E_x$ in $\mathcal{B}_{/T}$. In degree two, lifting the two-simplex ending at z identifies $(gf)^*$ with f^*g^* on the fiber over z. Higher terminal-vertex squares provide the compatible homotopies for longer strings. A left fibration instead gives an operation $f_! : E_x \to E_y$. The object $E_\bullet$ stores the transport as a coherent system, even when no levelwise representatives of its homotopies have been chosen.

There is a large internal universe $\Omega_{\mathcal{B}}$ of objects in slices of $\mathcal{B}$, taking T to $\mathcal{B}_{/T}$ with reindexing by pullback, in a suitable enlargement of universes. The internal Grothendieck construction identifies the internal category of small right fibrations with the internal functor category from C^{op} to $\Omega_{\mathcal{B}}$. The left-fibration version is obtained by taking the opposites of the total category and the base. This theorem and internal Yoneda are developed in [11, §§4.5–4.7]. For a contextual point $c : T \to C_0$, the representable right fibration is the slice $C_{/c} \to C|_T$, formed in $\mathcal{B}_{/T}$. The Yoneda comparison evaluates a map of right fibrations on the identity of c. The lift of an arrow into c then recovers its value by pullback transport. Coherence makes this a universal description of the mapping object, including in families over T.

Externally the constructions specialize to straightening in spaces,

$$\mathrm{RFib}_{\mathcal{S}}(C) \simeq \mathrm{Fun}(C^{\mathrm{op}}, \mathcal{S}), \qquad \mathrm{LFib}_{\mathcal{S}}(C) \simeq \mathrm{Fun}(C, \mathcal{S}).$$

We use the right-hand convention for ordinary presheaves and the second convention for covariant representations. An equivalence $C \to D$ induces equivalent categories of either sort. In the special case of a groupoid, inversion identifies the two conventions. For a general C, reversing C has mathematical content, and the relevant variance has to be part of any classification theorem [2, §2.2].

Proposition 11.3 (Inverting the transport)**.** *Let W be a set of arrows of a small ∞-category C. Restriction along $C \to C[W^{-1}]$ identifies*

$$\mathrm{Fun}(C[W^{-1}]^{\mathrm{op}}, \mathcal{S})$$

with the full subcategory of right fibrations on C in which transport along W is by equivalences. The corresponding statement holds for left fibrations.

Proof. The universal property of localization applies to C^{op} and the reversals of the arrows in W. It identifies its entire functor category into $\mathcal{S}$ with the indicated full subcategory of $\mathrm{Fun}(C^{\mathrm{op}}, \mathcal{S})$. Under straightening, sending a reversed arrow to an equivalence is precisely the transport condition. The covariant proof uses C itself. □

For W consisting of all arrows, $C[W^{-1}]$ is the groupoidification C^{gpd}. It represents the subtheory with invertible transport. In particular it does not generally classify all right or all left fibrations on C.

11.3 Two factorization systems

The factorization developed in [1, §7.3] first sends a functor to its essential image, then embeds that image fully faithfully. This has a direct higher analog. For an ∞-category B, write $B^{\simeq}$ for its maximal subgroupoid.

Proposition 11.4. *Every functor $f : A \to B$ factors as an essentially surjective functor followed by a fully faithful functor. These classes form an orthogonal factorization system on Cat_∞.*

Proof. Choose the full subcategory $I \subseteq B$ on the union of the components of $B^{\simeq}$ meeting $f(A^{\simeq})$. It contains the image of f, every object is equivalent to one from A, and inclusion is fully faithful.

For orthogonality, take an essentially surjective $e : A \to B$ and a fully faithful $m : X \to Y$. Replace X over Y by its full essential image using the preceding construction and full faithfulness. Given a square with e on the left and m on the right, the map $B \to Y$ sends every object into this full subcategory since every object is equivalent to one in the image of A. A full subcategory is determined by its allowed vertices and includes every simplex of Y with those vertices in a quasicategory model. The map $B \to Y$ thus lifts. Postcomposition with a full

inclusion gives a fully faithful functor $\mathrm{Fun}(B, X) \to \mathrm{Fun}(B, Y)$. On cores the space of a lift with its prescribed identification is empty or contractible. The existence just proved, and the same full faithfulness for A, show that imposing the rest of the square leaves a contractible space of lifts. Closure under retracts follows from the intrinsic definitions. Existence and orthogonality give the factorization system. □

A second factorization has a different right class. Given $f : A \to B$, let $(b \downarrow f)$ be the comma category of objects $(a, b \to f(a))$, with arrows over b. Write $|D|$ for the groupoidification of D. We call f final if every $|b \downarrow f|$ is contractible. By the higher form of Quillen's Theorem A, this is equivalent to f preserving colimit computations by restriction [2, §4.1.3]. An initial functor is one whose opposite is final. Its test uses the comma categories of arrows $f(a) \to b$.

Theorem 11.5 (Comprehensive factorization)**.** *There is an orthogonal factorization system on Cat_∞ with left class the final functors and right class the right fibrations, taken up to equivalence of arrows. The right-fibration factor of $f : A \to B$ is classified by the presheaf*

$$L_f = \mathrm{Lan}_{f^{\mathrm{op}}} \mathbf{1}_A, \qquad L_f(b) \simeq |b \downarrow f|.$$

Dually, the initial functors and the left fibrations form an orthogonal factorization system.

Proof. Here $\mathbf{1}_A$ is the terminal presheaf. The pointwise formula for left Kan extension has indexing category $(b \downarrow f)^{\mathrm{op}}$. The colimit of the constant point diagram is its groupoidification, canonically equivalent to that of the opposite category. In these terms the action of L_f is precomposition on comma categories. Let $p : E_f \to B$ be the right fibration classified by L_f. The Kan unit $\mathbf{1}_A \to (f^{\mathrm{op}})^* L_f$ corresponds to a section of $f^* E_f$ over A, giving $j : A \to E_f$ with $pj = f$. Under the comma description, $j(a)$ is the point represented by $(a, \mathrm{id}_{f(a)})$.

To prove finality of j, fix $e = (b, z)$ in E_f. A lift of $\beta : b \to f(a)$ from e to $j(a)$ is specified by a path between z and the transported point $L_f(\beta)(j(a))$ in $L_f(b)$. This description is natural in (a, β), since an arrow of this comma category supplies a homotopy relating the two transported points. Consequently $(e \downarrow j)$ is classified over $(b \downarrow f)$ by the family of path spaces $\mathrm{Map}_{L_f(b)}(z, \lambda(a, \beta))$, where $\lambda : (b \downarrow f) \to |b \downarrow f|$ is the localization map. This family sends every arrow to an equivalence. Its total groupoidification is the colimit of the family in spaces by straightening and the universal property of a colimit. Universality of colimits in spaces computes this colimit by pulling the colimit of the points $\lambda(a, \beta)$ in $\mathcal{S}_{/L_f(b)}$ back along z. That colimit is $\mathrm{id}_{L_f(b)}$, so the pullback is contractible. This is the comma criterion for j.

For orthogonality, pull a right fibration back along the bottom map $B \to C$ of a square and let the resulting fibration over B correspond to F. A top map lifting a given arrow $e : A \to B$ gives a section over A, equivalently an element

of $\mathrm{Map}(\mathbf{1}_A, (e^{\mathrm{op}})^* F)$. The space of fillers is the fiber of the restriction map from $\mathrm{Map}(\mathbf{1}_B, F)$ over this element. The left Kan extension adjunction identifies restriction with precomposition by $L_e \to \mathbf{1}_B$. This arrow is an equivalence when e is final, so the fiber is contractible. Finally final functors and right fibrations are each closed under retracts. For the right class this follows directly from the defining pullback condition, or from its lifting tests against terminal-vertex inclusions. Thus the factorizations and mutual orthogonality identify both classes. Taking opposites proves the dual assertion. □

The distinction between the two factorizations is visible on a category with indefinitely composable arrows. Take $\mathbb{Z}$ with its usual order, viewed as an ∞-category. The inclusion of the even integers is final since the even integers above each given b form a poset with a least element. It is fully faithful and fails essential surjectivity. Conversely $\mathbb{Z}_{\mathrm{disc}} \to \mathbb{Z}$ is essentially surjective, whereas $(b \downarrow f)$ is a discrete infinite ray for this f and has noncontractible groupoidification. Even a groupoidification equivalence need not be final. The inclusion of 0 in $\mathbb{Z}$ induces an equivalence on groupoidifications because $\mathbb{Z}$ is filtered and nonempty, yet $(1 \downarrow f)$ is empty.

The internal versions of the comprehensive factorization and of the comma criterion are established in [11, §4.4]. The tests there use every context T and take groupoidification internally. Applying the preceding external criterion only to global points of B would not give that result.

11.4 Descent along an augmented diagram

The precategories of [1, §§7.1–7.5] are diagrams retaining a small amount of categorical incidence data, before enforcing the Segal pullback conditions. Their coefficient presheaves are formed by a two-dimensional limit. The corresponding higher construction can start from an arbitrary indexing ∞-category. This preserves the distinction the original book makes between a limit of change-of-base categories and the category of ordinary directed presheaves on an internal category.

Let I be small, let $D : I^{\triangleright} \to \mathcal{A}$ be an augmented diagram with cone point a, and let $F, G : \mathcal{A}^{\mathrm{op}} \to \mathrm{Cat}_\infty$ be contravariant coefficient systems. The natural transformations between them are homotopy coherent. We call the augmentation effective for G when the canonical comparison

$$G(a) \longrightarrow \lim_{i \in I^{\mathrm{op}}} G(D_i)$$

is an equivalence. An object of the limit includes objects in the coefficient categories and equivalences relating every pullback of them, together with higher coherences. The maps between the objects of a compatible family are likewise compared in the entire mapping spaces.

Theorem 11.6 (Relative splitting along a diagram)**.** *Suppose $F \to G$ is pointwise fully faithful, the augmentation D is effective for G, and there is an object i_0 such that every i admits an arrow $i \to i_0$ in I. The choices of those arrows are immaterial. Let $\mathrm{Split}_F(D) \subseteq G(a)$ consist of those x whose restriction to D_{i_0} belongs to the essential image of $F(D_{i_0})$. Then comparison induces an equivalence*

$$\mathrm{Split}_F(D) \simeq \lim_{i \in I^{\mathrm{op}}} F(D_i).$$

Proof. A limit of full inclusions is fully faithful. Its image consists of those coherent families whose entries lie in the respective essential images. One way to see this last assertion is to replace each $F(D_i)$ by the full subcategory of $G(D_i)$ on its image. Naturality ensures that each reindexing map restricts to the relevant full subcategories. In the model of a limit as the category of Cartesian sections of the associated Cartesian fibration over I, the section belongs to the subfibration precisely when its objects do. Fullness in each fiber implies that every compatibility equivalence and coherent simplex of that section belongs to the subfibration as well.

For a section $x_\bullet$, membership of x_{i_0} implies membership of x_i because an arrow $i \to i_0$ yields $x_i \simeq G(D_i \to D_{i_0})(x_{i_0})$. Naturality of $F \to G$ places this pullback in the image of $F(D_i)$. The converse follows by evaluation at i_0. This identifies the limit of the $F(D_i)$ with the pullback of

$$F(D_{i_0}) \longrightarrow G(D_{i_0}) \quad \text{along} \quad \lim G(D_i) \longrightarrow G(D_{i_0}).$$

Using effectivity, this pullback is the defining full subcategory of $G(a)$. □

Taking I to be the shape of the precategory in [1, §7.2], each of its object, arrow, and composable-pair indices maps to the object index. In the ordinary category-valued case the coherent pullback equivalences are isomorphisms in the coefficient categories, and the limit computes the pseudolimit used in the book. The relative splitting criterion is then exactly the requirement at the object index in [1, Theorem 7.5.2]. The higher theorem allows the whole simplicial shape and the resulting mapping spaces, with no reduction to comparison isomorphisms on a truncated shape.

If a diagram has several distinguished objects collectively receiving arrows from each index, the same argument imposes the condition at all distinguished objects. If $D_\bullet$ is simplicial, $I = \Delta^{\mathrm{op}}$ and $i_0 = [0]$ can be used. In that case a vertex inclusion in Δ supplies the arrow required by the proof. For the Čech nerve of an effective epimorphism in an ∞-topos, the slice coefficient system $G(T) = \mathcal{B}_{/T}$ has effective descent. The theorem then restricts ordinary descent to a full class of objects that is stable under reindexing. Its naturality requirement on $F \to G$ is explicit. It does not imply that every adjunction on slice categories determines such a subsystem.

There is a precise interpretation of the slice construction on an arbitrary simplicial object $D_\bullet$. Every colimit in an ∞-topos is van Kampen [2, §6.1.3], so

$$\lim_{[n]\in\Delta} \mathcal{B}_{/D_n} \simeq \mathcal{B}_{/|D_\bullet|_{\mathcal{B}}}, \qquad |D_\bullet|_{\mathcal{B}} = \operatorname*{colim}_{[n]\in\Delta^{\mathrm{op}}} D_n.$$

Indeed a compatible family amounts to a map of simplicial objects $E_\bullet \to D_\bullet$ for which the naturality square for every simplicial operator is Cartesian. If $D_\bullet$ is Segal, this includes both vertex conditions, so transports from one endpoint to the other are equivalences. Over a complete Segal space the realization $|D_\bullet|_{\mathcal{S}}$ models the groupoidification. It counts simplices as coherent paths in a space, and in that target every edge becomes invertible. Thus these particular limit coefficients generalize the precategorical construction, while noninvertible directed transport requires the left- or right-fibration construction.

11.5 Descent of directed representations

Two useful forms of directed gluing keep track of which category contains the colimit. The first has its presentation in Cat_∞, where arrows remain directed. The second is indexed by objects in a base topos and glues all simplicial degrees using pullback.

Proposition 11.7 (Gluing over categorical colimits)**.** *Let $C = \mathrm{colim}_{i\in I} C_i$ in Cat_∞, with all C_i small, and let $\mathcal{E}$ be an ∞-category in the coefficient universe. Then*

$$\mathrm{Fun}(C, \mathcal{E}) \simeq \lim_{i\in I^{\mathrm{op}}} \mathrm{Fun}(C_i, \mathcal{E}).$$

If a collection W_i of arrows is specified in each C_i compatibly with the diagram, the equivalence restricts to functors inverting the union W of their images in C. The contravariant version follows by taking opposites of the indexing categories C_i.

Proof. For any ∞-category K of the appropriate size, $K \times -$ preserves colimits, as its right adjoint is $\mathrm{Fun}(K, -)$. Mapping from K to the comparison functor therefore yields

$$\mathrm{Map}(K \times C, \mathcal{E}) \simeq \lim_i \mathrm{Map}(K \times C_i, \mathcal{E}).$$

Testing against all K proves the equivalence of ∞-categories. Inverting an image arrow is a pointwise property of an object of the functor category. A full subcategory of the limit defined by these conditions consequently corresponds to the specified full subcategory on the left. The operation $C \mapsto C^{\mathrm{op}}$ is an autoequivalence of Cat_∞ and preserves the colimit in question. □

Proposition 11.8 (Changing the base of an internal category)**.** *Let $X \in \mathcal{B}$, let $U \to X$ be an effective epimorphism, and let C be a small internal ∞-category in*

$\mathcal{B}_{/X}$. For $U_n = U^{\times_X (n+1)}$ let C_{U_n} be the pullback internal category. Then there is an equivalence

$$\mathrm{RFib}_{\mathcal{B}_{/X}}(C) \simeq \lim_{[n]\in\Delta} \mathrm{RFib}_{\mathcal{B}_{/U_n}}(C_{U_n}).$$

The left-fibration version is obtained by the same argument with initial vertices.

Proof. Effective descent in the topos gives $\mathcal{B}_{/X} \simeq \lim \mathcal{B}_{/U_n}$. Taking simplicial diagrams commutes with this limit, and taking arrows with target the prescribed C also commutes with the limit. This glues the underlying map of simplicial objects, functorially on morphisms. An arrow in $\mathcal{B}_{/X}$ is an equivalence if its pullback to U is an equivalence. To see this, its pullbacks to each U_n are equivalences by further reindexing, and the limit of slice categories detects equivalences on entries.

The pullback functors between all these slices preserve limits. Consequently the Segal equations for the glued source can be checked after pullback to U, where they hold. The same reasoning applies to the completeness comparison by computing the space of coherent inverse data using limits. Finally the terminal-vertex pullback conditions amount to equivalences with finite pullbacks and hold after reindexing to U. The glued map is therefore a right fibration. This restricts the equivalence on underlying maps to precisely the subcategories in the statement. □

CHAPTER 12

Exit paths and constructible monodromy

A covering of a connected space transports its fibers along paths, with inverse paths producing inverse equivalences. A constructible object on a stratified space has locally constant fibers on each stratum, together with comparisons across incident strata. An exit path starts in a lower stratum and may move toward higher strata. The reverse path can violate the stratification condition. This asymmetry supplies the directed information in constructible monodromy.

12.1 The topological classification

Give a poset P its Alexandrov topology of upward-closed subsets. A P-stratified space is a continuous map $s : X \to P$, and $X_p = s^{-1}(p)$. An exit simplex $\sigma : |\Delta^n| \to X$ is a singular simplex for which there is a chain $p_0 \leq \cdots \leq p_n$ in P such that $\sigma(t_0, \ldots, t_n)$ belongs to X_{p_i} when i is the largest index for which $t_i > 0$. Face and degeneracy operations respect this condition. A conically stratified space admits open neighborhoods of points $x \in X_p$ given by stratified charts $Z \times C(L)$, where $Z \times \{\text{vertex}\}$ lies in X_p and the link L is stratified by the elements above p. Here $C(L)$ is the open cone, with a basis of vertex neighborhoods given by the cones of uniformly bounded radius. Conicality ensures that the exit simplicial set is an ∞-category, denoted $\mathrm{Exit}_P(X)$. We use the conical chart and exit-path conventions in [12, §2.1].

The homotopy type $\Pi_\infty X_p$ is the fiber of $\mathrm{Exit}_P(X) \to P$ over p. The inclusion of any individual stratum in the exit category is fully faithful. Indeed a simplex whose vertices are all over p has constant stratum label, so lies entirely in X_p. All endomorphisms in $\mathrm{Exit}_P(X)$ are equivalences for the same reason. More generally the functor to P is conservative. For a one-stratum space the simplicial model is the ordinary singular complex, with no directed constraints.

A hypersheaf of spaces on X is a sheaf which satisfies descent for hypercovers. Call it P-hyperconstructible if its restriction, as a hypersheaf, to every X_p is locally constant. This is a property, so the hyperconstructible objects form a full subcategory. In particular the construction does not prescribe a single common space as the value of all stalks. Each locally constant restriction can carry its own monodromy.

Theorem 12.1 (Topological exodromy, space coefficients)**.** *If $X \to P$ is conically stratified and each X_p admits a basis of weakly contractible open neighborhoods,*

there is an equivalence

$$\mathrm{Cons}^{\mathrm{hyp}}_{P}(X;\mathcal{S}) \simeq \mathrm{Fun}(\mathrm{Exit}_{P}(X),\mathcal{S}).$$

Evaluation of a representation at x corresponds to the stalk at x.

This is the space-coefficient case of the theorem of Porta and Teyssier [12, Theorem 5.4.1 and Corollary 6.2.5]. Their coefficient-general statement assumes joint conservativity of the hyper-restrictions to the strata. For spaces this conservativity follows from the stalk criterion for hypersheaves. The stalks of the restriction to a stratum are the stalks on X at those same points, and the strata contain every point. For the full theorem the extra conservativity condition is substantive. The statement above uses hyperdescent throughout. Passing to ordinary sheaves requires a comparison with hypercompletion in the setting at hand.

The proof architecture of the cited theorem clarifies the role of the hypotheses. The exodromy adjunction reconstructs a hypersheaf from a representation by assigning to an open U the limit over its exit-path category. Conical charts give cofinal tests for computing its inverse image on a stratum. In these tests the restriction of both functors and of the adjunction transformations agrees with the corresponding monodromy adjunction on the stratum. For a locally weakly contractible stratum, that adjunction is an equivalence on local systems. Hyper-restriction detects the unit on constructible hypersheaves, and evaluation at every point detects the counit on representations. This is where the theorem uses the local geometry beyond just the existence of a topological stratification [12, §§4–5].

By the convention of Chapter 11, the representations in Theorem 12.1 are classified by left fibrations over $\mathrm{Exit}_{P}(X)$. If one uses entrance paths and denotes the resulting opposite category by $\mathrm{Ent}_{P}(X)$, the same data are ordinary presheaves on $\mathrm{Ent}_{P}(X)$. This change of orientation applies to the indexing category. The stalk maps attached to exit paths are still covariant.

Proposition 12.2 (Sections and open restriction)**.** *Under the hypotheses of Theorem 12.1, let F be the hypersheaf corresponding to $M : \mathrm{Exit}_{P}(X) \to \mathcal{S}$. For an open $U \subseteq X$ with the induced stratification,*

$$F(U) \simeq \lim_{\mathrm{Exit}_{P}(U)} M|_{\mathrm{Exit}_{P}(U)}.$$

The limits retain every compatibility with the noninvertible arrows of the exit category.

Proof. The strata of an open subspace inherit the required local bases and conical charts. Naturality of exodromy for open inclusions identifies hypersheaf restriction with precomposition by $\mathrm{Exit}_{P}(U) \to \mathrm{Exit}_{P}(X)$ [12, Proposition 6.2.3]. Sections are maps from the terminal hypersheaf. The terminal hypersheaf is

constructible, and under an equivalence of categories it corresponds to the terminal functor, of constant value $*$. The mapping space out of that functor is the limit being asserted, by the defining adjunction of the constant-diagram functor with the limit functor. □

12.2 A conical calculation with higher stalks

Let M be the plane with the discrete closed subset $\mathbb{Z} \times \{0\}$ removed, and let $L = S^2$. Both are locally contractible manifolds, M with infinitely many punctures. On $X = M \times C(L)$, put the lower stratum $Z = M \times \{\text{vertex}\}$ and the upper stratum $U = M \times L \times \mathbb{R}_{>0}$. The lower stratum has conical charts obtained by taking disks in M times cones of uniformly bounded radius. The upper stratum is open and is an ordinary manifold. This gives explicit conical and weak local contractibility hypotheses for exodromy.

Write $K = \Pi_\infty(M)$ and $T = \Pi_\infty(L)$. The exit category of the cone is $T^{\triangleleft}$, the category obtained by adjoining an initial object to T. To justify this, the nonvertex subcategory is full and equivalent to T by product with $\mathbb{R}_{>0}$, while the vertex is initial by the cone lemma [12, Lemma 2.1.8]. An initial vertex determines, up to contractible choice, a coherent cone on the inclusion of the nonvertex subcategory. The resulting functor from $T^{\triangleleft}$ is fully faithful. The mapping spaces out of the initial vertex are contractible and there are no arrows from a nonvertex object to the vertex, since such arrows decrease the stratum label. Mapping spaces between nonvertex objects agree by fullness. The functor is therefore an equivalence. Exit simplices of a product with an unstratified space decompose into a singular simplex in that space and an exit simplex in the stratified factor. It follows that

$$\mathrm{Exit}_P(X) \simeq K \times T^{\triangleleft}.$$

Write $\pi : K \times T \to K$ for projection. A representation is equivalently the full coherent data

$$A : K \to \mathcal{S}, \qquad B : K \times T \to \mathcal{S}, \qquad \eta : \pi^* A \longrightarrow B.$$

Indeed currying gives a K-indexed family of diagrams on $T^{\triangleleft}$. A diagram on $T^{\triangleleft}$ is a diagram on T equipped with a cone, which gives precisely the displayed data in each fiber. Natural transformations are compatible maps of the two diagrams and the comparison; the compatibility square is specified in the functor ∞-category. At a given k, adjunction also writes the comparison as $A(k) \to \lim_{t \in T} B(k,t)$. This last description shows why keeping only the monodromy on each stratum can lose a comparison even when both restrictions are known exactly.

The inclusion of K at the initial vertex is initial. For an object (k,t) of the entire exit category, the category of arrows from (k', vertex) to it is equivalent to $K_{/k}$, which has a terminal object. Therefore the global sections

of the hypersheaf corresponding to (A, B, η) are $\lim_K A$ by Proposition 12.2 and the initiality criterion. The comparison η still controls the behavior on neighborhoods intersecting the upper stratum.

For example, take A and B to be constant of value the Eilenberg–Mac Lane space $Q = K(\mathbb{Z}, 2)$. The identity comparison and the comparison constant at the base point give inequivalent constructible hypersheaves with equivalent restrictions to both strata. A natural equivalence between the two representations would give a commutative square up to homotopy on any radial exit arrow, with equivalences on its vertical sides. On $\pi_2(Q)$, one horizontal side is the identity and the other is zero. No such square can have invertible vertical sides.

There is also a link contribution in the direct image j_*B for $j : U \hookrightarrow X$. In the representation category j_* corresponds to right Kan extension along $K \times T \hookrightarrow K \times T^{\triangleleft}$. This can be seen directly from $j_*B(V) = B(V \cap U)$ and the section formula. At (k, vertex), the comma category for right Kan extension is equivalent to $K_{k/} \times T$. In this product, restricting a limit to the slice at id_k is allowed because the identity is initial in $K_{k/}$. Therefore the value on the lower stratum is $k \mapsto \lim_T B(k, -)$. For the constant diagram with value Q, this gives

$$\mathrm{Map}_{\mathcal{S}}(S^2, Q) \simeq Q \times \mathbb{Z}_{\mathrm{disc}}.$$

Evaluation at a chosen point of S^2 has fiber $\Omega^2 Q \simeq \mathbb{Z}_{\mathrm{disc}}$. Taking Q in its abelian group space model separates a map into its value at the chosen point and a pointed map by translation, proving the product decomposition. In particular the direct image has extra components on the lower stratum although the initial value Q is connected. The higher homotopy type of the link is used in this computation. Its fundamental group is trivial.

Inverting all exit paths in this example sends $K \times T^{\triangleleft}$ to K. Groupoidification preserves products for these factors, as can be seen either from the existence of an initial object in the cone or by taking iterated colimits of constant points. Accordingly a functor comes from ordinary monodromy on X precisely when η is an equivalence. In that case $B \simeq \pi^* A$ by η, which already implies the factorization. Conversely any representation from a groupoid sends the exit comparisons to equivalences. This agrees with the deformation retraction of X to the lower stratum.

12.3 A van Kampen principle with directed coefficients

The right colimit for local presentations of a stratified space is a colimit of exit categories. Let $\{U_\alpha\}$ be an open cover of a conically stratified X with locally weakly contractible strata, and let $V_\bullet \to X$ be its open Čech presentation. Its degree n can be read as the disjoint union of the intersections $U_{\alpha_0} \cap \cdots \cap U_{\alpha_n}$.

The exit-path van Kampen theorem gives

$$\mathrm{Exit}_P(X) \simeq \operatorname*{colim}_{[n]\in\Delta^{\mathrm{op}}} \mathrm{Exit}_P(V_n) \qquad \text{in } \mathrm{Cat}_\infty .$$

For the application here this follows from the colimit-preserving exit-path functor on the hypersheaf topos, obtained by extending its values on opens [12, §4]. Indeed the coproducts of intersections give the Čech nerve in that topos, and its colimit is terminal because the cover map is an effective epimorphism. The exit functor sends this colimit to the displayed colimit. It sends a coproduct of the opens as hypersheaves to a coproduct of their exit categories, which agrees with the exit category of the disjoint union.

Proposition 11.7 now identifies constructible hypersheaves with coherent families of representations on these local exit categories. This operation compares local assignments on intersections by equivalences of representations. Their values along individual exit arrows inside an open set can be arbitrary maps. The equivalence relating two restrictions of a representation on an overlap and the map that the representation assigns to an exit arrow inside an open carry different parts of the data.

Suppose the strata are coarsened by a map of posets $Q \to P$, and both stratifications on X are conical with locally weakly contractible strata. Let $v : \mathrm{Exit}_Q(X) \to \mathrm{Exit}_P(X)$ be the comparison. Under these hypotheses v is the localization at the arrows whose images are equivalences [12, Corollary 6.2.7]. Therefore a Q-hyperconstructible hypersheaf is P-hyperconstructible if and only if the corresponding representation sends every such arrow to an equivalence. One implication follows by restriction from a P-representation. In the other direction the universal property of localization descends the representation to $\mathrm{Exit}_P(X)$, and naturality of exodromy identifies its pullback with the original hypersheaf. The test is about coherent functors, so the comparison includes mapping spaces of classes that survive the coarsening. With the one-stratum coarsening, the test forces every transport to be invertible and recovers locally constant hypersheaves.

12.4 Recovering the directed category from its coefficients

Call an object a of a presentable ∞-category atomic if $\mathrm{Map}(a, -)$ preserves all small colimits. This includes preservation of the empty colimit and is stronger than compactness, which asks only for preservation of filtered colimits. An atomic object is one on which it is possible to test arbitrary colimits by mapping out of it.

Proposition 12.3. *Let C be a small ∞-category admitting a conservative functor to a poset. There is a canonical equivalence*

$$C \simeq (\mathrm{Fun}(C, \mathcal{S})_{\mathrm{at}})^{\mathrm{op}} ,$$

where the subscript denotes the full subcategory of atomic objects.

Proof. Covariant representables $h_c = \mathrm{Map}_C(c, -)$ are atomic since mapping out of h_c is evaluation and colimits are pointwise. Conversely every functor $F : C \to \mathcal{S}$ is a colimit of covariant representables over its category of elements, by density. For F atomic, apply $\mathrm{Map}(F, -)$ to one such colimit presentation of F. The identity of F is in the image on components of some stage, since π_0 preserves colimits of spaces and every element of a colimit of sets comes from a stage. Thus F is a retract of a representable, with maps $F \to h_c \to F$ whose composite is homotopic to the identity.

Every endomorphism of an object of C maps to an identity in the poset and is an equivalence by conservativity. By Yoneda the induced endomorphism of h_c is also an equivalence. With one composite the identity and the other invertible, the two maps in the retract diagram are equivalences. Every atom is therefore representable. Covariant Yoneda is fully faithful on C^{op}, giving the claimed variance and equivalence. □

Combining this proposition with Theorem 12.1 recovers the exit category as the opposite of the atomic constructible hypersheaves. Under this identification the atom associated to x corepresents the stalk at x. A map of atoms has variance reversed relative to an exit path. This follows from the natural transformation between evaluation functors and the contravariance of corepresenting objects. In particular the directed category can be recovered from the intrinsic colimits and arrows of the coefficient category, without a choice of a base point. If C does not admit the conservativity assumption, the general statement from the proof recovers the idempotent completion, since atoms are then retracts of representables [2, §5.1.6].

Atomic generation is also a route to exodromy for stratified ∞-topoi whose geometric origin need not be conical. Take a geometric morphism to $\mathrm{Fun}(P, \mathcal{S})$, define the stratum topoi by geometric base change along the point topoi at p, and define constructibility by locally constant restrictions. The formulation of [13, Definition 2.2.10] assumes atomic generation of the constructible subcategory, its closure under limits and colimits in the ambient topos, and preservation of limits by the inverse image of the stratification. The first property identifies the constructible category with a presheaf category by the free colimit completion of its atoms. The limit compatibility gives the left adjoint to stratification pullback on constructible coefficients, and this adjoint carries atoms to atoms because its right adjoint preserves colimits. Taking opposites then supplies the functor from the abstract exit category to P. Thus a theorem proving these properties for a given stratified topos supplies a classifier without assuming that its simplices are represented by topological exit paths.

12.5 The profinite directed case

The framework of Barwick, Glasman and Haine supplies an arithmetic counterpart [10]. For a coherent scheme Y, in particular a quasicompact quasiseparated scheme, their category $\mathrm{Gal}(Y)$ has geometric points as objects and étale specializations as morphisms. The construction equips it with a profinite categorical structure. A specialization carries incidence information, whereas endomorphisms at a geometric point account for the absolute Galois action on its residue-field data. Keeping both features is necessary to reconstruct the associated constructible coefficients.

For space coefficients in the sense of [10, Corollary 0.5.2], the theorem takes the form

$$\mathrm{Cons}(Y_{\text{ét}}; \mathcal{S}_\pi) \simeq \mathrm{Fun}_{\mathrm{cts}}(\mathrm{Gal}(Y), \mathcal{S}_\pi),$$

where on the left the constructibility condition is that of the stratified étale topos, and $\mathcal{S}_\pi$ denotes the π-finite spaces. Such spaces are truncated, have finitely many components, and have finite homotopy groups in each component. The right side uses the profinite stratified shape and its continuous representation category. When a field is viewed as a one-point scheme, a chosen separable closure supplies a single geometric base point up to isomorphism, and the loop automorphisms in this classification recover its absolute Galois group with its profinite structure. General schemes may also have noninvertible specializations between points.

The classifying pro-space of the profinite stratified shape recovers the ordinary profinite shape for spectral stratified topoi [10, Example 0.4.8]. This forgetful operation is the arithmetic counterpart of discarding directed transport by groupoidifying an exit category. The arithmetic theorem retains a continuity condition on representations, while Theorem 12.1 uses an ordinary small exit ∞-category without an added profinite topology. The comparison thus relates the two generalizations through the role of coherence and direction, with the appropriate topology determined by the coefficient problem.

Part VI

Higher topoi and presentation

CHAPTER 13

Geometric morphisms and the classification of sections

The final chapter of Borceux and Janelidze passes from objects in a topos to morphisms between topoi. Their descent groupoid is a kernel construction in a two-dimensional category. It records identifications between geometric morphisms and thereby differs from the kernel pair of a map of objects in a fixed topos. In higher dimensions this distinction persists. The kernel of an étale geometric morphism can be computed from an internal Čech nerve. An external atlas of a topos requires separate information about its overlaps. We first prove the representation theorem for étale geometric morphisms, then use it to identify both the scope and the limits of the internal construction. The classical background is [1, §§7.7–7.10].

13.1 The category of geometric morphisms

An ∞-topos is a presentable ∞-category admitting a presentation by an accessible left exact localization of an ∞-category of presheaves of spaces. Write $\mathcal{P}(C) = \mathrm{Fun}(C^{\mathrm{op}}, \mathcal{S})$ for the presheaves on a small ∞-category. A geometric morphism $f\colon \mathcal{Y} \to \mathcal{X}$ is specified by an inverse image $f^*\colon \mathcal{X} \to \mathcal{Y}$. It preserves finite limits and small colimits, so the adjoint functor theorem supplies its right adjoint f_*. Thus a morphism of topoi reverses the direction of the functor that performs substitution. In particular a localization $L\colon \mathcal{X} \to \mathcal{X}'$ gives a geometric morphism $\mathcal{X}' \to \mathcal{X}$ when L is left exact. In the inverse image convention, a morphism from f to g in the category of geometric morphisms is a natural transformation $f^* \to g^*$. Passing to mates reverses this direction if one works with direct images.

For each pair $\mathcal{Y}, \mathcal{X}$, put

$$\mathrm{Geom}(\mathcal{Y}, \mathcal{X}) = \mathrm{Fun}^{\mathrm{lex,L}}(\mathcal{X}, \mathcal{Y}).$$

We take a full subcategory of the indicated functor category, with all natural transformations between geometric inverse images. Composition of functors and their transformations yields the $(\infty, 2)$-category of geometric morphisms in this convention. Its underlying $(\infty, 1)$-category, written $\mathbf{Top}_\infty$, keeps the maximal subgroupoids of the hom ∞-categories. A comparison saying that two diagrams of inverse images are identified in $\mathbf{Top}_\infty$ means that their comparison transformation is invertible. The noninvertible transformations can still be

useful. For instance, a classifier for a theory can have models as the objects of a hom category and homomorphisms as its noninvertible transformations. The core records the space of models and their equivalences.

The terminal ∞-topos in $\mathbf{Top}_\infty$ is $\mathcal{S}$. For each $\mathcal{X}$, the inverse image of its structure map to $\mathcal{S}$ sends a space K to $\Delta_{\mathcal{X}} K = \operatorname{colim}_K 1_{\mathcal{X}}$. Indeed, $\mathcal{S}$ is the free cocompletion of the terminal category, and the requirement to preserve the terminal object specifies its image. The resulting functor preserves finite limits by the finite limit extension theorem for presheaves [2, Prop. 6.1.5.2]. Its right adjoint is global sections $\operatorname{Map}_{\mathcal{X}}(1_{\mathcal{X}}, -)$. The entire hom category $\operatorname{Geom}(\mathcal{X}, \mathcal{S})$ is contractible. This argument uses the terminal-object condition on the inverse image, so it applies equally when maps of topoi carry higher homotopies.

An inverse image preserves effective epimorphisms because it preserves Čech nerves by left exactness and preserves their realizations. It also preserves monomorphisms by the diagonal test. An image factorization is an effective epimorphism followed by a monomorphism, so an inverse image preserves images as well. In particular the family of images of generating covers transported by the inverse image remains jointly covering. This elementary property is what links geometric morphisms to a site presentation.

13.2 The universal property of a site

Suppose C has a Grothendieck topology. A sieve R on c is regarded as a monomorphism $R \hookrightarrow y(c)$ in $\mathcal{P}(C)$. Sheafification $\mathcal{P}(C) \to \operatorname{Shv}(C; \mathcal{S})$ is the localization inverting covering sieve inclusions. If a family $\{c_i \to c\}$ generates R, the image of $\coprod_i y(c_i) \to y(c)$ is R. The sheaf condition at R for F is the equivalence

$$F(c) \xrightarrow{\sim} \operatorname{Map}_{\mathcal{P}(C)}(R, F).$$

It imposes a coherent limit condition on the whole sieve. When there are pullbacks of the covering arrows, the condition also admits a Čech form by replacing R by the realization of the Čech nerve of $\coprod_i y(c_i) \to R$. This identification takes place already in presheaves, where effective epimorphisms and colimits are calculated pointwise.

Theorem 13.1 (Site description of inverse images)**.** *Let C be a small ∞-category with finite limits and a Grothendieck topology. Let $j \colon C \to \operatorname{Shv}(C; \mathcal{S})$ be Yoneda followed by sheafification. For every ∞-topos $\mathcal{Y}$, precomposition by j gives an equivalence of $\operatorname{Geom}(\mathcal{Y}, \operatorname{Shv}(C; \mathcal{S}))$ with the full subcategory of $\operatorname{Fun}(C, \mathcal{Y})$ on functors A satisfying the following conditions.*

1. *The functor A preserves finite limits.*
2. *For every family $\{c_i \to c\}$ generating a covering sieve, the map $\coprod_i A(c_i) \to A(c)$ is an effective epimorphism.*

Proof. The free cocompletion property identifies A with its colimit-preserving extension $\overline{A}\colon \mathcal{P}(C) \to \mathcal{Y}$. Under this identification the extension is left exact precisely when A is left exact. Necessity uses the left exactness of Yoneda. For sufficiency apply the finite limit extension theorem [2, Prop. 6.1.5.2] to the finitely complete domain C.

The extension then takes the image factorization

$$\coprod_i y(c_i) \longrightarrow R \longrightarrow y(c)$$

to an image factorization. Thus the condition on the generating family is equivalent to $\overline{A}$ inverting the covering sieve. By the universal property of localization, $\overline{A}$ factors if and only if all these inclusions are inverted. The descended functor is left exact since it is the restriction of $\overline{A}$ along the right adjoint inclusion of sheaves, which preserves limits. It preserves colimits by the localization property or by taking the corresponding right adjoint on presheaves, whose image lies in sheaves. These constructions are equivalences of full functor subcategories and preserve natural transformations. This proves the equivalence including its morphisms [2, Prop. 6.2.3.20]. □

For a presentation by an arbitrary accessible left exact localization $L\colon \mathcal{P}(C) \to \mathcal{X}$, the first part of the argument survives without any topology on C. Restriction to C identifies inverse images with those functors for which $\overline{A}$ is left exact and sends the L-equivalences to equivalences. When L is displayed by a set of generators, testing those generators suffices since $\overline{A}$ preserves colimits. An arbitrary localization can impose relations beyond those specified by covering sieves. Thus the general inverse image theorem starts with the actual localization, which may include a condition at a map between presheaves that is not a sieve inclusion.

For an ordinary site (C, J) this construction provides a canonical lift of its ordinary topos. The 0-truncated objects of $\mathrm{Shv}(C; \mathcal{S})$ are exactly its sheaves of sets. To see this, evaluation on the sheafified representables tests whether a sheaf takes 0-truncated values. Conversely, mapping into a sheaf taking such values can be tested on a colimit of sheafified representables and then gives a limit of 0-truncated spaces. The canonical lift depends only on the ordinary topos. In the terminology of higher topos theory it is the 1-localic lift, and inverse images from an ∞-topos into it are detected by the corresponding functors on 0-truncated objects [2, Prop. 6.4.5.7]. One may therefore extend an ordinary geometric morphism between Grothendieck topoi to the 1-localic lifts. Recovering an arbitrary ∞-topos from its 0-truncated objects is a stronger assertion and generally fails.

13.3 Étale maps and dependent substitution

For a map $u\colon U \to V$ in an ∞-topos $\mathcal{X}$, base change has adjoints on both sides

$$u_! \dashv u^* \dashv u_* \qquad \left(u^*\colon \mathcal{X}_{/V} \longrightarrow \mathcal{X}_{/U}\right).$$

The first adjoint composes an arrow to U with u. The second exists because pullback preserves colimits. Pullback preserves finite limits as well. The resulting map of topoi $\mathcal{X}_{/U} \to \mathcal{X}_{/V}$ is étale. In particular $\pi_U\colon \mathcal{X}_{/U} \to \mathcal{X}$ denotes the map obtained from $U \to 1_{\mathcal{X}}$. More generally a geometric morphism to $\mathcal{X}$ is called étale when it is equivalent over $\mathcal{X}$ to some π_U. In this use of the word U may have higher homotopy in its fibers.

Composition and pullback satisfy a relative projection formula. For $A \to B$ in $\mathcal{X}_{/V}$ and $D \to u^*B$ in $\mathcal{X}_{/U}$, there is a natural equivalence

$$u_!(u^*A \times_{u^*B} D) \simeq A \times_B u_!D.$$

Indeed the underlying object on either side is $A \times_B D$ in $\mathcal{X}$, with its composite map to V. If a square in $\mathcal{X}$ is a pullback and the parallel horizontal maps are $g'\colon U' \to U$ and $g\colon V' \to V$, then $g^*u_! \simeq u'_!g'^*$ by the same calculation. For the right adjoints the comparison $g^*u_* \simeq u'_*g'^*$ follows by testing against $T \in \mathcal{X}_{/V'}$ and using $u^*g_!T \simeq g'_!(u'^*T)$. These identities describe the two ways to integrate a family after change of base. They extend Frobenius identities on subobjects, where the relative (-1)-truncation of dependent sum is read as existential quantification.

Theorem 13.2 (The slice classifies sections)**.** *For $U \in \mathcal{X}$, the category of geometric morphisms $\mathcal{Y} \to \mathcal{X}_{/U}$ over the hom category $\mathrm{Geom}(\mathcal{Y}, \mathcal{X})$ is represented by the left fibration classified by*

$$\mathrm{Geom}(\mathcal{Y}, \mathcal{X}) \longrightarrow \mathcal{S}, \qquad f \longmapsto \mathrm{Map}_{\mathcal{Y}}(1_{\mathcal{Y}}, f^*U).$$

In particular, a lift of f to a geometric morphism $\mathcal{Y} \to \mathcal{X}_{/U}$ over $\mathcal{X}$ amounts to a section of f^*U, with all paths between sections.

Proof. The object π_U^*U is $U \times U \to U$ and carries the diagonal section δ over U. A geometric morphism h from $\mathcal{Y}$ to $\mathcal{X}_{/U}$ gives the pair $(h^*\pi_U^*, h^*\delta)$. Conversely, a pair (f^*, s) with $s\colon 1_{\mathcal{Y}} \to f^*U$ induces the functor

$$(A \xrightarrow{a} U) \longmapsto 1_{\mathcal{Y}} \times_{f^*U} f^*A.$$

Here the first step applies f^* in the slice. The second step is pullback along s. Both steps preserve small colimits and finite limits. For the first step, finite limits follow from left exactness and colimits in a slice are computed in its ambient category. For the second step, colimits follow from universality and limits

from the right adjoint description of pullback as right adjoint to composition. Consequently this formula defines a geometric inverse image.

For any $A \to U$ there is a pullback square in the slice expressing it as $\pi_U^* A \times_{\pi_U^* U} 1_{\mathcal{X}_{/U}}$, with the right-hand map determined by δ. An inverse image preserves this square. Applying the construction to $(h^*\pi_U^*, h^*\delta)$ therefore recovers h^* naturally on objects and maps. The reverse composite recovers (f^*, s) by evaluating on objects $C \times U$ and then on the diagonal. These natural comparisons also hold for transformations. In detail, a transformation between pairs consists of a transformation $\alpha\colon f^* \to g^*$ and a specified path $\alpha_U s \simeq t$. It induces the transformation of the pullback formulas by the universal property of the pullback. Restriction to π_U^* and evaluation at δ recovers the original pair of data. Since these operations preserve parametrized families of transformations, the comparison is an equivalence of ∞-categories. The category of these pairs is the Grothendieck construction of the displayed space-valued functor. See also [2, Prop. 6.3.5.5 and Cor. 6.3.5.6]. □

The section classifier gives two useful refinements. First, if $g\colon \mathcal{Y} \to \mathcal{X}$ is geometric, then base change of π_U has source $\mathcal{Y}_{/g^*U}$. A test geometric morphism $f\colon \mathcal{Z} \to \mathcal{Y}$ lifts to that slice exactly by specifying a section of f^*g^*U, which also describes a compatible lift of gf to $\mathcal{X}_{/U}$. Second, there is an equivalence of ∞-categories

$$\mathcal{X} \simeq (\mathbf{Top}_\infty/\mathcal{X})_{\mathrm{et}}, \qquad U \longmapsto \pi_U.$$

Essential surjectivity is the definition of étale. For full faithfulness use the section classifier with $\mathcal{Y} = \mathcal{X}_{/V}$ and $f = \pi_V$. A section of $V \times U \to V$ in that slice is precisely a map $V \to U$. Spaces of such sections compute the relative mapping spaces. This result identifies a whole category of geometric tests over $\mathcal{X}$ with the objects of $\mathcal{X}$, without choosing a site or a point.

13.4 Internal classifying objects and external classifiers

An internal group G in $\mathcal{X}$ admits a delooping $BG \in \mathcal{X}$. For every $f\colon \mathcal{Y} \to \mathcal{X}$, $f^*BG \simeq B(f^*G)$ because the bar construction uses finite products and geometric realization. Combining the slice classifier with torsor classification gives the space of lifts of f through $\mathcal{X}_{/BG}$ as the space of f^*G-torsors in $\mathcal{Y}$. The universal torsor there is classified by the diagonal section of $\pi_{BG}^* BG$. For an arrow $V \to 1$ in $\mathcal{X}$ the same rule describes G-torsors in the slice $\mathcal{X}_{/V}$ by maps $V \to BG$. For a general geometric f, its inverse image replaces every coefficient object before the section is taken.

The object BG represents torsors on objects of $\mathcal{X}$. The slice $\mathcal{X}_{/BG}$ represents the choice of a torsor over a geometric map into $\mathcal{X}$. In both statements the group is fixed in the base. The category of G-objects includes arbitrary equivariant maps and will figure in the presentation of the next chapter.

CHAPTER 14

Groupoid presentations of higher topoi

A presentation gives two kinds of information. It identifies a domain on which objects can be described locally, and it says how to identify their pullbacks on overlaps. The ordinary Joyal–Tierney theorem provides both pieces using locales. The higher construction below presents slices by étale geometric morphisms and keeps track of the type of each overlap. The difference becomes visible even when the atlas itself is the topos of spaces.

14.1 The descent formula for a diagram of slices

A colimit in an ∞-topos is van Kampen in the following sense. For a small diagram $U\colon I \to \mathcal{X}$ with colimit V, restriction along the colimit cone gives an equivalence of ∞-categories

$$\mathcal{X}_{/V} \;\simeq\; \lim_{i \in I^{\mathrm{op}}} \mathcal{X}_{/U_i}. \tag{14.1}$$

Transition functors on the right are pullbacks. Equivalently, a Cartesian family $A_i \to U_i$, meaning one whose transport squares are pullbacks, has a realization $\operatorname{colim}_i A_i \to V$ whose pullbacks are the given family. This is the descent form of the higher Giraud axioms [2, Thm. 6.1.3.9 and Prop. 6.1.3.10]. The realization description also shows why preserving pullbacks by colimits of slices is a stronger condition than mere universality of colimits. Universality gives the recovery of a slice object from an already existing pullback family. Descent asserts that a compatible family given independently of a global object descends.

Theorem 14.1 (Colimits of slices as higher topoi)**.** *With the preceding hypotheses, there is an equivalence*

$$\operatorname*{\mathbf{Top}_\infty\text{-}colim}_{i \in I} \mathcal{X}_{/U_i} \;\simeq\; \mathcal{X}_{/V}.$$

For every ∞-topos $\mathcal{Z}$, the induced comparison

$$\mathrm{Geom}(\mathcal{X}_{/V}, \mathcal{Z}) \;\longrightarrow\; \lim_{i \in I^{\mathrm{op}}} \mathrm{Geom}(\mathcal{X}_{/U_i}, \mathcal{Z})$$

is an equivalence of ∞-categories.

Proof. Apply $\mathrm{Fun}(\mathcal{Z}, -)$ to (14.1). The limit of functor categories classifies a coherent family of functors to the slices. It remains to restrict to geometric

inverse images. The projections from the limit preserve finite limits and small colimits, since all the transition functors do. For example, a diagram of compatible families has its componentwise colimit formed in each slice. A transition functor carries each component colimit to the corresponding colimit, so the colimit inherits the required compatibility and satisfies the universal property. The same argument applies to finite limits. Joint conservativity of projections then says that a functor into the limit is left exact and colimit-preserving exactly when each component is. Natural transformations are restricted through the same equivalence of functor categories. Taking cores yields the mapping space universal property of the colimit of topoi. This is the fixed base case of the more general colimit theorem for étale geometric morphisms [2, Thm. 6.3.5.13]. □

14.2 From a groupoid to its quotient

Let $G_\bullet \colon \Delta^{\mathrm{op}} \to \mathcal{X}$ be a groupoid object, in the Segal sense with all arrows invertible. Put $Q = |G_\bullet|$. The Segal identities and invertibility concern an internal diagram in $\mathcal{X}$. They imply that $G_\bullet$ is the Čech nerve of $G_0 \to Q$. We recall a proof using the van Kampen property because its argument makes explicit the role of invertibility.

Proposition 14.2 (Effectiveness of internal groupoids)**.** *The map $G_0 \to Q$ is an effective epimorphism, and there are compatible equivalences*

$$G_n \simeq \underbrace{G_0 \times_Q \cdots \times_Q G_0}_{n+1}.$$

Proof. Define a simplicial object $E_\bullet$ by $E_n = G_{n+1}$, where the simplicial operators retain the first vertex and act on the remaining vertices. In degree n, its map to G_n forgets that first vertex. There is also an augmentation $E_\bullet \to G_0$ given by the retained vertex. Duplicating the retained vertex is an extra degeneracy contracting this augmentation, so $|E_\bullet| \simeq G_0$.

The map $E_\bullet \to G_\bullet$ is Cartesian. To check this, first pull it back along a vertex of a simplex G_n. Over the first vertex of the remaining chain, the assertion is a Segal identity since the extra data is the one incoming arrow. Over any other vertex, the incoming arrow can be transported forward along the chosen segment of that chain. Composition by an invertible arrow is an equivalence on the corresponding space of arrows, with inverse given by composition by its inverse. This argument is functorial on generalized points and proves the pullback identity for any chosen vertex. To check an arbitrary simplicial operator $G_n \to G_m$, evaluate the two candidate pullbacks over a vertex of G_m. Both identify with the same pullback of $E_0 \to G_0$ along the composite vertex of G_n.

Descent now identifies E_0 with the pullback of $|E_\bullet|$ along $G_0 \to Q$. Its two projections are the endpoints of the arrow in G_1. Hence $G_1 \simeq G_0 \times_Q$

G_0. Compatibility with longer chains follows by the Segal maps and their composition coherences. The augmented Čech nerve therefore realizes to Q, proving effectiveness. This is the shift argument of [15, §6.10] in the notation of internal groupoids. □

A $G_\bullet$-object is a Cartesian transformation $A_\bullet \to G_\bullet$. It is given first by an object $A_0 \to G_0$, then by transport of this object along arrows of G_1, with composition homotopies on G_2 and their higher compatibility. Equivariant morphisms are maps of such transformations. The corresponding ∞-category will be denoted $\mathrm{Act}_{\mathcal{X}}(G_\bullet)$. Applying (14.1) gives

$$\mathrm{Act}_{\mathcal{X}}(G_\bullet) \;\simeq\; \lim_{[n]\in\Delta} \mathcal{X}_{/G_n} \;\simeq\; \mathcal{X}_{/Q}. \tag{14.2}$$

The inverse sends $A \to Q$ to the family $A \times_Q G_n$. The forward map sends a Cartesian family to its realization. Under Theorem 14.1, the right-hand side is also the colimit of the simplicial diagram of étale geometric morphisms $\mathcal{X}_{/G_n}$. The two variance conventions agree because the geometric morphisms point in the direction opposite to their inverse image functors.

For the special case of $U \to 1_{\mathcal{X}}$ effective epi, let $G_n = U^{n+1}$. Then $Q \simeq 1_{\mathcal{X}}$ and (14.2) recovers $\mathcal{X}$ itself from the Čech diagram of $\mathcal{X}_{/U} \to \mathcal{X}$. Indeed Theorem 13.2 identifies

$$\underbrace{\mathcal{X}_{/U} \times_{\mathcal{X}} \cdots \times_{\mathcal{X}} \mathcal{X}_{/U}}_{n+1} \simeq \mathcal{X}_{/U^{n+1}},$$

where the fiber products on the left are in $\mathbf{Top}_\infty$. The inverse image of this projection is conservative since evaluation of a Cartesian family at degree zero detects equivalences. Conversely, if π_U^* is conservative, it reflects the equivalence obtained by pulling back $\mathrm{im}(U \to 1) \to 1$ along U. This image must be all of 1, so $U \to 1$ is an effective epimorphism. For étale geometric morphisms, surjectivity in the sense of conservative inverse image thus agrees with being induced by an effective epimorphism of objects. All the formulas hold in a slice over any specified base object.

Presentations here are invariant under refinement in a concrete sense. If $W \to U \to 1_{\mathcal{X}}$ are effective epimorphisms, their Čech nerves have a map $W^{\bullet+1} \to U^{\bullet+1}$. Pullback on actions is an equivalence because under (14.2) it identifies with the identity functor of $\mathcal{X}$. Two covers $U \to 1_{\mathcal{X}}$ and $V \to 1_{\mathcal{X}}$ admit a common refinement by $U \times V$. Both projections are covers by pullback stability. There are typically many different simplicial diagrams presenting the same topos.

14.3 The classical localic theorem

For a locale L, let $\mathcal{O}(L)$ be its frame. A map of locales $f\colon L \to M$ is open when $f^*\colon \mathcal{O}(M) \to \mathcal{O}(L)$ admits a left adjoint $f_!$ satisfying the Frobenius identity

$$f_!(a \wedge f^* b) = f_!(a) \wedge b.$$

An étale locale over M is one locally isomorphic over M to open sublocales of M. Its sheaf of sections is an ordinary sheaf of sets on M. An open localic groupoid consists of locales of objects and arrows, groupoid operations in locales, and open source and target maps.

The Joyal–Tierney theorem provides, for every ordinary Grothendieck topos $\mathcal{E}$, an open localic groupoid G with $\mathcal{E}$ equivalent to the category of étale G-objects. These are étale locales over G_0 with an action along G_1, or equivalently sheaves of sets on G_0 with the corresponding equivariance. An ingredient of the theorem supplies a locale L with an open effective-descent geometric morphism $\mathrm{Shv}(L;\mathrm{Set}) \to \mathcal{E}$. Take its Čech nerve by pseudopullbacks of ordinary topoi. The map from a localic topos to an arbitrary Grothendieck topos is localic in the relative sense. Indeed its source has subterminal generators, and these are subobjects of the inverse image of the terminal object. Base change preserves relative localicity. Consequently the iterated overlaps are localic over the localic source and therefore localic. Openness under base change gives the openness of the groupoid. Effective descent identifies $\mathcal{E}$ with the category of sheaves on G_0 endowed with the required action. The existence of the localic cover and the effective-descent theorem are substantive parts of the classical result [14, Chs. VII–VIII]; see also [1, Thms. 7.10.1 and 7.10.7].

There are two uses of étaleness here. In the ordinary localic result, étale locales carry the sheaves on which the groupoid acts. The atlas of the target topos is supplied as an open geometric morphism of effective descent. In the higher relative result, an atlas $\mathcal{X}_{/U} \to \mathcal{X}$ is itself an étale geometric morphism. Its source need not be localic. The theory of internal groupoid actions proves the relative presentation once a cover is chosen. In the classical theorem, the existence of an external atlas with localic overlaps is additional content.

14.4 When a higher étale atlas has localic overlaps

Call an ∞-topos 0-localic if it is equivalent to $\mathrm{Shv}(L;\mathcal{S})$ for a locale L. Its locale is recovered from the subobjects of 1 [2, §7.3.2]. For a space A, the subobjects of the terminal object in $\mathcal{S}_{/A}$ are indexed by subsets of $\pi_0 A$. To see this, a subobject has fiber either empty or contractible, and this property is invariant under transport along paths. Conversely each union of path components gives such a subobject.

Lemma 14.3. *The ∞-topos $\mathcal{S}_{/A}$ is 0-localic exactly when A is 0-truncated.*

Proof. For a discrete space A, a family over A is a sheaf of spaces on the discrete locale with set of points A. For the converse, the reconstructed locale would have frame the power set of $\pi_0 A$. Over each nonempty connected component A_0 the corresponding open slice would then be equivalent to $\mathcal{S}$. Such an equivalence must identify the constant-object functor $\mathcal{S} \to \mathcal{S}_{/A_0}$ with an equivalence, since that functor is the unique geometric inverse image out of $\mathcal{S}$ up to equivalence. For $a \in A_0$, the object $\{a\} \to A_0$ has nonempty fiber at a. A constant object with nonempty fiber has a global section. A section of $\{a\} \to A_0$ makes id_{A_0} homotopic to the constant map at a, so A_0 is contractible. This holds for each connected component, giving a discrete homotopy type. □

Theorem 14.4 (Localic étale Čech presentations over spaces)**.** *Let B be a small space. The ∞-topos $\mathcal{S}_{/B}$ admits a surjective étale geometric morphism for which both the source and the double overlap are 0-localic if and only if B is 1-truncated. In that case every higher overlap of the chosen atlas can also be taken 0-localic.*

Proof. Write the covering source as $\mathcal{S}_{/U}$, with map to $\mathcal{S}_{/B}$ induced by $U \to B$. Surjectivity implies $U \to B$ is effective epi, hence surjective on path components. The source is 0-localic only if U is discrete by Lemma 14.3. For a discrete U, the space $U \times_B U$ is the disjoint union, over pairs $u, v \in U$, of spaces of paths in B from the image of u to the image of v. Since the double overlap is $\mathcal{S}_{/(U \times_B U)}$, its localicity forces these path spaces to be discrete. Using $u = v$ and one point over each component shows that each loop space of B is 0-truncated. Equivalently, B is 1-truncated.

Conversely, take U discrete with one chosen point of B in each component. This is an effective epimorphism, including the empty case. Its higher fiber products over B have discrete homotopy types, since a tuple of points is accompanied by a tuple of paths in the 1-truncated space B. Lemma 14.3 identifies all their slices with 0-localic ∞-topoi. □

For an infinite discrete group G and $B = BG$, the atlas coming from $* \to BG$ has nth term $\mathcal{S}_{/G^n}$. Equation (14.2) gives its category of actions as $\mathrm{Fun}(BG, \mathcal{S})$. Its 0-truncated objects are ordinary G-sets. Taking instead $B = K(\mathbb{Z}, 2)$ yields the atlas $\mathcal{S} \to \mathcal{S}_{/B}$ whose first overlap is $\mathcal{S}_{/K(\mathbb{Z},1)}$. That overlap is 1-localic since it is a presheaf category on an ordinary groupoid, and it fails to be 0-localic by Lemma 14.3. Moreover $\mathcal{S}_{/B}$ and $\mathcal{S}$ have the same ordinary topos of 0-truncated objects. Straightening identifies the first ordinary topos with set-valued functors on B, which factor through the fundamental groupoid. For $K(\mathbb{Z}, 2)$ that groupoid is contractible. The two higher topoi are inequivalent by Lemma 14.3. Theorem 14.4 concerns the specific class of surjective étale geometric morphisms and does not assume a theorem about all other notions of higher atlas.

14.5 Hyperdescent and the role of truncation

In any ∞-topos an augmented simplicial object $V_\bullet \to V$ is a hypercover if each map to its relative matching object is an effective epimorphism. In degree zero the matching object is V. In degree one it is $V_0 \times_V V_0$. In higher degrees it is the object of compatible boundaries relative to V. Each matching condition asks for local existence of a lift. When the augmentation is a Čech nerve, there is a stronger identity. The levels are specified iterated pullbacks and the geometric realization of a cover is exactly V.

The realization of an arbitrary hypercover maps to V by an ∞-connective morphism [2, Lem. 6.5.3.11]. This means its relative n-truncation is terminal for every finite n. Applying (14.1) to $V_\bullet$ identifies $\lim_\Delta \mathcal{X}_{/V_n}$ with $\mathcal{X}_{/|V_\bullet|}$. Consequently the restriction comparison for all objects over V is an equivalence precisely when $|V_\bullet| \to V$ is an equivalence. Indeed, if the restriction g^* for a map g is an equivalence, so is its left adjoint $g_!$. Applying $g_!$ to the terminal object in the source slice gives g as an object of the target slice, so g is terminal there. Conversely an equivalence g induces an equivalence of slices.

An ∞-topos is hypercomplete if every ∞-connective map is an equivalence. Hypercompletion is the left exact localization which forces these equivalences. Each ∞-connective morphism gives a constant simplicial hypercover in its target slice, according to the equivalent iterated-diagonal test for ∞-connectivity [2, Lem. 6.5.3.5]. Thus effectiveness of all hypercovers is equivalent to hypercompleteness [2, Thm. 6.5.3.12]. A bounded truncated object already satisfies the extra locality condition. Mapping into an n-truncated object cannot distinguish an ∞-connective map from an equivalence, by orthogonality to n-truncated maps. In particular, passing to ordinary sheaves of sets cannot test whether a given ∞-topos is hypercomplete.

A concrete instance of the distinction is the canonical 1-localic lift of continuous $\mathbb{Z}_p$-sets. One can present it by the category of transitive quotients $\mathbb{Z}_p/p^n\mathbb{Z}_p$ with its atomic topology. This is an ∞-topos of ordinary sheaves of spaces in the sense of covering sieves, and it fails to be hypercomplete by Wieland's example [2, Warning 7.2.2.31]. The site is ordinary and its underlying ordinary topos already has a classical group-theoretic presentation. Neither fact removes the hyperdescent distinction. The higher quotient (14.2) uses an effective groupoid object, so it works in that topos before hypercompletion as well as in hypercomplete examples.

Proposition 14.5 (Relative hypercompletion)**.** *Let $L\colon \mathcal{X} \to \mathcal{X}^\wedge$ be hypercompletion. For every $U \in \mathcal{X}$, the induced functor*

$$L_{/U}\colon \mathcal{X}_{/U} \longrightarrow (\mathcal{X}^\wedge)_{/LU}$$

is the hypercompletion of the slice. In particular, applying L to a hypercover

$V_\bullet \to V$ gives a slice comparison

$$(\mathcal{X}^\wedge)_{/LV} \;\simeq\; \lim_{[n]\in\Delta} (\mathcal{X}^\wedge)_{/LV_n}.$$

Proof. Let $i\colon \mathcal{X}^\wedge \hookrightarrow \mathcal{X}$ be the fully faithful inclusion. For $T \to LU$, the right adjoint of $L_{/U}$ is the object $U \times_{iLU} iT \to U$. Indeed, a map into this pullback over U amounts under $L \dashv i$ to a map into T over LU. The counit is invertible by left exactness of L and the counit equivalence $LiT \simeq T$. It follows that $L_{/U}$ is a localization. A morphism in the slice is inverted precisely when its underlying morphism in $\mathcal{X}$ is inverted by L. This is the class of ∞-connective maps by the strong saturation theorem [2, Prop. 6.5.2.8]. Connectivity in the slice agrees with the connectivity of the underlying map. For example, the slice over its codomain in $\mathcal{X}_{/U}$ identifies with the slice over that same codomain in $\mathcal{X}$, so the relative truncation tests are identical. This identifies the localization as the desired hypercompletion.

Finally, L preserves the realization of $V_\bullet$ and sends its ∞-connective augmentation to an equivalence. Apply (14.1) inside $\mathcal{X}^\wedge$. The result can also be read as descent for relative hypercomplete objects. The base object in the formula is LV, even when V itself has not been assumed hypercomplete. □

In the $\mathbb{Z}_p$ example the constant hypercover of a nonequivalence supplied by infinite connectivity has a slice comparison that fails before hypercompletion. The proposition makes the comparison an equivalence after applying L to its entire augmentation. The higher Čech theorem continues to apply to actual groupoid quotients in the original topos.

CHAPTER 15

Reconstruction and comparison

A classification theorem identifies a category of forms or coverings from a symmetry object and chosen coefficients. A reconstruction theorem asks what symmetry can be read back from that category. This direction is sensitive to the fiber functor, the availability of representing objects, and the tests used to define equivalence. We finish by comparing the different reconstruction problems.

15.1 Recovering the space of fibers

Let B be a small space and put $\mathcal{L}(B) = \mathrm{Fun}(B, \mathcal{S})$. Write $\mathrm{ev}_b : \mathcal{L}(B) \to \mathcal{S}$ for evaluation. A path from b to c induces a natural equivalence $\mathrm{ev}_b \to \mathrm{ev}_c$. The space of such natural equivalences remembers all paths and their homotopies.

Theorem 15.1 (Reconstruction from the evaluations)**.** *The functor*

$$B \longrightarrow \mathrm{Fun}(\mathcal{L}(B), \mathcal{S}), \qquad b \longmapsto \mathrm{ev}_b$$

is fully faithful. In particular, as group objects in spaces,

$$\mathrm{Aut}(\mathrm{ev}_b) \simeq \Omega_b B.$$

These assertions concern natural transformations in the full functor category.

Proof. For $b \in B$, the covariant representable $R_b = \mathrm{Map}_B(b, -)$ is an object of $\mathcal{L}(B)$. The covariant version of Yoneda gives $\mathrm{ev}_b \simeq \mathrm{Map}_{\mathcal{L}(B)}(R_b, -)$. A natural transformation from evaluation at b to evaluation at c is, by a second application of Yoneda, a point of $\mathrm{Map}_{\mathcal{L}(B)}(R_c, R_b) \simeq R_b(c)$. The equivalence is natural in b, c, and it respects composition because the two appearances of the contravariant representing-object assignment reverse order twice. Equivalently, a path $u : b \to c$ acts on each F by $F(u)$, which agrees with this identification. Every arrow in B is invertible, so the equivalence of mapping spaces identifies the full space of natural transformations between the two evaluations with its subspace of equivalences. □

This proof explains the regular-action formula. On writing $B = BG$, the representable R_* is the regular left G-space. Its endomorphisms are right multiplications, which compose in the opposite order. Passing from the representing object to the evaluation functor reverses the order again. The group on the

right side of Theorem 15.1 is therefore compatible with covariant monodromy. For the groupoid B, the pointed reconstruction only uses one evaluation, while the full family of evaluations retains the components and the paths between them.

Proposition 15.2 (Reconstruction from all geometric points)**.** *The core of the category of colimit-preserving, finite-limit-preserving functors $\mathcal{S}_{/B} \to \mathcal{S}$ is equivalent to B. Under straightening, its objects are the evaluations of Theorem 15.1.*

Proof. The universal property of the slice topos, as in Chapter 13, identifies geometric morphisms $\mathcal{T} \to \mathcal{S}_{/B}$ with pairs consisting of a geometric morphism $a : \mathcal{T} \to \mathcal{S}$ and a section of a^*B. For $\mathcal{T} = \mathcal{S}$, the space of choices of a is contractible. A colimit-preserving functor from spaces is determined by the image of the point, and left exactness forces that image to be terminal. The section space of B is B. The inverse image associated to b pulls back a family along $b : * \to B$, and that is evaluation under straightening. □

The distinction between this result and a general theorem on points of topoi is substantial. The slice of spaces has specially chosen geometric points, determined by the family classifier B. A general higher topos has a shape defined by constant-object tests. That definition does not assert that its points are the components of its shape, or that its points distinguish all objects.

15.2 What truncation preserves

The preceding calculation is also valid with truncated coefficients. For an integer $n \geq 0$, let $\mathcal{L}_n(B) = \mathrm{Fun}(B, \mathcal{S}_{\leq n})$. Denote the evaluation functor on $\mathcal{L}_n(B)$ by ev_b^n, giving its output as a space.

Proposition 15.3 (Truncated reconstruction)**.** *There is a natural equivalence*

$$\mathrm{Map}_{\mathrm{Fun}(\mathcal{L}_n(B),\mathcal{S})}(\mathrm{ev}_b^n, \mathrm{ev}_c^n) \simeq \tau_{\leq n} \mathrm{Map}_B(b, c).$$

Every component of the left side consists of equivalences. Restricted evaluation consequently defines a fully faithful functor $\tau_{\leq n+1}B \to \mathrm{Fun}(\mathcal{L}_n(B), \mathcal{S})$.

Proof. Objectwise truncation is left adjoint to the inclusion $\mathcal{L}_n(B) \hookrightarrow \mathcal{L}(B)$. Hence ev_b^n is represented by $R_b^n = \tau_{\leq n} R_b$ in $\mathcal{L}_n(B)$. The two Yoneda identifications give $\mathrm{Map}(\mathrm{ev}_b^n, \mathrm{ev}_c^n) \simeq \mathrm{Map}_{\mathcal{L}_n(B)}(R_c^n, R_b^n) \simeq R_b^n(c)$. The components of $R_b^n(c)$ come from actual paths in B, so all corresponding natural transformations are equivalences. The full subcategory spanned by the evaluations is thus an $(n+1)$-truncated groupoid. The evaluation functor factors through $\tau_{\leq n+1}B$ by its truncation universal property. The mapping-space description identifies the factorization as fully faithful. □

In particular, set-valued monodromy recovers the fundamental groupoid and its groups of loops at chosen points. A space of coefficients truncated in a higher degree records one additional degree in the classifying groupoid. The regular representable need not lie in the subcategory of finite coefficients. The proof above therefore gives no reconstruction of the uncompleted base from finite tests. The pro-space mapping formula and finite-range hypotheses from Chapters 4 and 5 are the relevant comparison in that setting.

For infinite discrete groups, the loss can already occur with set coefficients. The additive group $\mathbb{Q}$ has no nontrivial finite quotient. If a map $\mathbb{Q} \to H$ has finite image of order dividing m, write any x as my. Its image is the mth power of the image of y, hence the identity. Any action of $\mathbb{Q}$ on a finite set is trivial, since it factors through a finite permutation group. In contrast the regular action recovers $\mathbb{Q}$ itself by Proposition 15.3. The choice of coefficient category therefore changes the reconstruction even for an ordinary group.

15.3 A comparison of the classification results

The constructions used throughout the book can be read in the following order. Choose a base and the objects considered over it. Specify the splitting condition. Determine the overlap diagram under which such objects can descend. Then decide whether the overlap diagram has a smaller groupoid, pro-object, or directed category as a presentation. The outputs below indicate which of these steps require hypotheses.

Input	Key hypothesis	Classification
An internal cover $U \to B$	Effective epimorphism in an ∞-topos	Objects over B as cartesian actions of the overlap groupoid. Chapter 3.
An adjunction $L \dashv R$ and a cover $U \to B$	Admissibility in slices and effective descent for the chosen maps	Objects split by U as a limit of slices over the image overlap diagram. Additional exactness makes that diagram a groupoid. Chapter 9.
A space, a pro-space, or the shape of a topos	Representable base for ordinary straightening or the coefficient and locality hypotheses of higher Galois theory	Monodromy on the space, stage representations, or locally constant objects, respectively. Chapters 4–5.
A spectral algebra $R \to S$	The flat or stable descent hypotheses appropriate to the objects; a Galois tensor identity for action comparison	Amitsur descent and, when the tensor identity holds, coherent semilinear actions. Chapters 7–8.

Input	Key hypothesis	Classification
A stratified base	The constructibility and geometric hypotheses of the corresponding exodromy theorem	Stratumwise systems with transport along exit paths or an étale profinite category. Chapter 12.
A topos and a geometric atlas	Effectivity of the chosen presentation in the specified category of topoi or slices	A category of sheaves or descent objects associated to that presentation. Chapter 14.

These implications have several useful reversals and limitations. A trivializing cover gives a subgroup correspondence if the classification of intermediate objects imposes connectedness, pointing, and the appropriate truncation. Dropping these conditions leads to actions on families, with a different notion of map. The totalization formula for a relative reflection requires compatible coefficient categories. Promoting a general image simplicial diagram to a groupoid requires a check on the Segal pullbacks. A theorem on profinite completion concerns the tests retained by that completion. Finally, specialization maps for a stratified space can fail to be invertible, so passing to the groupoid completion has the effect of requiring additional inverses from its representations.

The common result is a homotopy-coherent descent statement whose strength is measured by its mapping spaces. The classical correspondence can then be obtained by restriction, truncation, or reversal of variance in those cases where the intermediate mathematical objects have the requisite form. The higher theorems keep the coherent data visible and make it possible to say exactly which features of the original symmetry object can be recovered.

Bibliography

[1] Francis Borceux and George Janelidze. *Galois Theories*. Cambridge Studies in Advanced Mathematics 72. Cambridge University Press, 2001. Publisher edition.

[2] Jacob Lurie. *Higher Topos Theory*. Annals of Mathematics Studies 170. Princeton University Press, 2009. arXiv:math/0608040.

[3] Marc Hoyois. Higher Galois theory. *Journal of Pure and Applied Algebra* 222 (2018), 1859–1877. arXiv:1506.07155.

[4] Thomas Nikolaus, Urs Schreiber, and Danny Stevenson. Principal infinity-bundles—General theory. *Journal of Homotopy and Related Structures* 10 (2015), 749–801. arXiv:1207.0248.

[5] Jacob Lurie. *Derived Algebraic Geometry VII: Spectral Schemes*, 2011. Author's manuscript.

[6] Jacob Lurie. *Higher Algebra*, 2017. Author's manuscript.

[7] John Rognes. Galois extensions of structured ring spectra. *Memoirs of the American Mathematical Society* 192 (2008), no. 898. arXiv:math/0502183.

[8] Akhil Mathew. The Galois group of a stable homotopy theory. *Advances in Mathematics* 291 (2016), 403–541. arXiv:1404.2156.

[9] The Stacks Project Authors. *The Stacks Project*. Tag 0244 and Tag 023R.

[10] Clark Barwick, Saul Glasman, and Peter Haine. *Exodromy*. arXiv:1807.03281.

[11] Louis Martini. Yoneda's lemma for internal higher categories. arXiv:2103.17141.

[12] Mauro Porta and Jean-Baptiste Teyssier. Topological exodromy with coefficients. arXiv:2211.05004.

[13] Peter J. Haine, Mauro Porta, and Jean-Baptiste Teyssier. Exodromy beyond conicality. arXiv:2401.12825.

[14] André Joyal and Myles Tierney. An extension of the Galois theory of Grothendieck. *Memoirs of the American Mathematical Society* 51 (1984), no. 309. DOI 10.1090/memo/0309.

[15] Charles Rezk. *Toposes and homotopy toposes*, v0.15, 2010. Author's manuscript.

[16] Daniel Dugger, Sharon Hollander, and Daniel C. Isaksen. Hypercovers and simplicial presheaves. *Mathematical Proceedings of the Cambridge Philosophical Society* 136 (2004), 9–51. arXiv:math/0205027.

[17] Egbert Rijke, Michael Shulman, and Bas Spitters. Modalities in homotopy type theory. *Logical Methods in Computer Science* 16 (2020), 2:2. arXiv:1706.07526.

[18] Jacob Lurie. *Kerodon*, §3.5. Truncations and Postnikov towers.

[19] George Janelidze and G. Max Kelly. Galois theory and a general notion of central extension. *Journal of Pure and Applied Algebra* 97 (1994), 135–161. DOI 10.1016/0022-4049(94)90057-4.

Subject index

abelianization
  discrete reflection, 73
  infinite-loop, 74
  ordinary, 67
action
  homotopy-coherent, 7
admissibility, 60
Amitsur complex
  coherent group action, 52
atomic object, 88

band, 69
Barr–Beck theorem, 16
Beck–Chevalley comparison, 95

Cantor space, 34
Čech nerve, 4, 14
central extension, 67
  framed, 69
classifying topos, 96
cohomology
  central extensions, 72
  continuous, 27
colimit
  van Kampen, 3
complete Segal nerve, 30
complete Segal object, 76
comprehensive factorization, 79
constant object, 29
covering
  pointed, 10
  regular, 13
covering space, 35

delooping, 6
descendability, 51
descent
  coherent, 2
  effective, 14
  faithfully flat, 49
  mapping spaces, 15
  monadic, 62

effective epimorphism, 4, 14
Eilenberg–Mac Lane space, 38
étale algebra
  spectral, 49
étale geometric morphism, 95
exit-path category, 84
exodromy, 84

factorization system
  reflective, 63
fiber functor, 104
finite cover
  weak, 55
fundamental group
  fiber functor, 35

Galois category
  stable homotopy theory, 55
Galois correspondence
  infinite, 26
Galois extension
  spectral, 52
Galois groupoid
  Pierce, 47
Galois structure, 59
geometric morphism, 92
Grothendieck topology, 93
group object, 6
groupoid
  change of presentation, 17
groupoid object
  actions, 98
  effectiveness, 98

homotopy fixed points, 9

intermediate algebra, 54
homotopy orbits, 9
hypercompleteness, 19
hypercompletion, 35, 102
hypercover, 19, 102

idempotent
spectral localization, 43
ind-completion, 25
inverse image, 92

Joyal–Tierney theorem, 100

left fibration, 77
local coefficients, 27, 38, 72
local system
stable modules, 56
localic groupoid, 100
localic topos, 100
locally constant object, 30
finite coefficient, 32
loop object, 6

modality
truncation, 64
monad
descent, 16
monodromy, 9
continuous, 22
obstruction theory, 37
monotone–light factorization, 66

nilpotence
tensor, 51

orthogonality, 63

π-finite space, 22
Pierce sheaf, 43
Pierce spectrum, 43
Postnikov tower, 37
pro-space, 22
profinite completion, 22
profinite exodromy, 90
profinite group
closed subgroup, 26
continuous action, 25
protruncation, 22

quotient
delooped, 13

reconstruction, 104
reflection
Galois, 59
relative splitting, 80
representation
at a stage, 23
right fibration, 77

shape, 29
protruncated, 35
relative, 29
sheafification, 93
split algebra
Pierce, 44
split object, 18
splitting
relative, 60
stable units, 64
Stone space
relative Pierce adjunction, 44
subgroup correspondence, 10

topos
classifying, 33
locally connected, 32
shape, 29
torsor
classification, 7
frames, 18
truncated map, 15
truncation
relative, 64

universes, 2

van Kampen colimit, 97