Research manuscript | November 2026

From Representability to Higher Sheaf Cohomology

Mapping spaces, stable extensions, and descent

Anonymous

Prepared for Higher Categorical Semantics

Abstract

The passage from classical category theory to ∞-categories changes the data carried by a universal construction. A representative controls a space of maps, a short exact sequence has a space of realizations, and a sheaf glues along all coherences of a cover. This manuscript develops these changes through Jan R. Strooker's progression from representability and exactness to homological algebra and sheaf cohomology. It compares stable extension spaces with Yoneda extensions, identifies the classical right satellites of a left exact functor by an enriched Kan extension, and recovers derived global sections from localization of module-valued presheaves. A finite-window comparison theorem determines maps in complete stable categories from bounded Postnikov pieces. An amplitude theorem identifies the input window sufficient for computing a prescribed output window of a derived right adjoint. Explicit calculations for finite sites, cyclic actions, and a projective line exhibit the hypotheses and demonstrate where extension and descent data change the answer.

Keywords. ∞-category, exact ∞-category, stable category, Yoneda extension, satellite, t-structure, ∞-topos, hyperdescent, sheaf cohomology.

Authorship and preparation. The manuscript is prepared with AI assistance and formatted with LaTeX. The byline is left anonymous for editorial consideration. Mathematical claims are supported by proofs or attributed to the cited literature. The uploaded textbook is the source of the historical and thematic comparison.

Introduction and conventions

Strooker's *Introduction to Categories, Homological Algebra and Sheaf Cohomology* [38] links four subjects by a common method. First recognize a construction by its effect on morphisms. Then use its universal property to organize exact sequences, extensions, and the change of base for sheaves. The method survives the passage to higher categories once a morphism set is replaced by a mapping space and universal properties are required to be coherent in the test object. This manuscript follows that mathematical progression. The definitions, proofs, and calculations below are developed independently, with references to the classical text and to sources for the higher-categorical inputs.

Four levels of comparison. In an ordinary category, a cone is specified by maps and strict compatibility. In an ∞-category, a cone includes coherent compatibility, and two cones with the same individual maps can differ as coherent cones. In a stable ∞-category, the same principle gives an equivalence between the space of cofiber sequences with fixed outer terms A, B and $\mathrm{Map}(B, \Sigma A)$. In a sheaf ∞-topos, Čech gluing uses the entire cosimplicial diagram of intersections. With these distinctions in place, taking degree zero, taking homotopy groups, and passing to a heart each recovers a specific classical construction.

Main comparisons. Theorem 13.2 gives a mapping-space formula for an exact t-exact functor by combining an increasing connective filtration of the source with a decreasing coconnective tower of the target. The formula also gives a criterion for full faithfulness from extension maps on the heart. Theorem 14.1 specifies which layers of the source can affect a chosen range of

the output of an exact functor of finite one-sided cohomological amplitude. Theorem 11.1 compares the graded Kan extension with right derived functors by means of the tensor product of an extension functor and an additive functor. Theorem 17.1 constructs the comparison with classical derived sheaves for a fixed ring by applying exact sheafification before passing to the stable localization. The applications show how these results interact in concrete section calculations.

Foundational inputs. We use the model of ∞-categories by quasicategories, or any equivalent model in which the cited universal properties hold. The straightening theorem, accessible adjoint functor theorem, higher Giraud theorem, theory of exact ∞-categories, and the existence of suitable derived resolutions are established results. They are cited where used. The comparison arguments and their consequences are proved in the text. Examples involving topological local systems state the topological hypotheses under which the classification by the fundamental ∞-groupoid applies.

The bounded Postnikov inductions in Theorems 13.2 and 14.1 use the classical truncation and dévissage arguments for t-structures, developed in [41, §1.3] and carried into the stable setting in [19, §1.2.1]. Their role here is to organize coherent finite windows and to separate these finite arguments from the reconstruction hypotheses needed for unbounded objects. Failures of reconstruction and the distinction between hyperdescent and Postnikov completeness are addressed separately in [22, 7] and Chapter 27.

Size. Fix universes $\mathbb{U} \in \mathbb{V}$. The words small and locally small refer to $\mathbb{U}$ unless an explicit universe shift is made. A site is small. A category of all $\mathbb{U}$-small spaces, spectra, or modules is a $\mathbb{V}$-small object when a functor category requires it. For coends on a large abelian category, one can form the coend in the larger universe and then use the comparison theorem to recover a $\mathbb{U}$-small value under the stated hypotheses.

Grading. In a stable category, suspension is written Σ. Under the cochain convention it corresponds to [1], so $\pi_i X = H^{-i}(X)$. The subcategory $\mathcal{C}_{\geq n}$ consists of n-connective objects and $\mathcal{C}_{\leq n}$ of n-coconnective objects. The maps of truncations are $\tau_{\geq n}X \to X$ and $X \to \tau_{\leq n}X$. The symbol $\tau_{[a,b]}$ denotes the composite $\tau_{\leq b}\tau_{\geq a}$, canonically equivalent to $\tau_{\geq a}\tau_{\leq b}$ when $a \leq b$. The category and whether a homotopy group or object is meant will be clear from the domain of the notation.

Contents

Abstract i

Introduction and conventions ii

I Representability and adjunctions 1

1 Categories with mapping spaces 2
- 1.1 The nerve and its inner horns . . . 2
- 1.2 Equivalences, functors, and coherence . . . 4
- 1.3 Slices and homotopy fibers . . . 5
- 1.4 Localizing morphisms . . . 6
- 1.5 Two probes for coherent information . . . 7

2 Yoneda, weights, and universal properties 9
- 2.1 Representatives and their elements . . . 9
- 2.2 Density and the category of elements . . . 10
- 2.3 Weights and composition . . . 12
- 2.4 Presheaf evaluation and detection . . . 13

3 Limits and the shape of a diagram 16
- 3.1 Cones and their tests . . . 16
- 3.2 Changing the index . . . 17
- 3.3 Pointwise Kan extension . . . 19
- 3.4 A limit with a cycle . . . 20

4 Adjoints and coherent change of base 22
- 4.1 Adjunctions from representability . . . 22
- 4.2 Adjoints preserve the corresponding universal constructions . . . 23
- 4.3 Mates of a comparison square . . . 24
- 4.4 Reflection and localization . . . 26

4.5 Adjoint triples and parametrized objects 26

II Exactness and stabilization 28

5 Images, fibers, and pointed exactness 29
5.1 Monomorphisms in an ∞-topos 29
5.2 Zero objects and the data of a nullhomotopy . . . 31
5.3 Kernels and cokernels in the heart 32
5.4 Effective equivalence relations and higher quotients 33

6 Stable categories and exact diagrams 35
6.1 Stability 35
6.2 Mapping spectra 37
6.3 Triangles and the boundary map 38
6.4 Pasting and the octahedral comparison 39

7 Exact categories and derived enhancements 41
7.1 Admissibility and exact structure 41
7.2 Extension-closed subcategories 42
7.3 Stable hulls and exact flags 43
7.4 Grothendieck categories 44
7.5 Localization of complexes 45

8 t-structures and Postnikov layers 47
8.1 Truncations 47
8.2 The heart and exactness 49
8.3 Layer maps and reconstruction 50
8.4 Separatedness and generators of layers 51

9 Presentability and compact generators 52
9.1 Small generators in large categories 52
9.2 Morita reconstruction 53
9.3 Products, generators, and completeness 55
9.4 Perfect objects 56

III Extensions and derived functors 58

10 Extension spaces and extension spectra 59
10.1 The universal extension space 59
10.2 Classical extensions and their automorphisms . . 60
10.3 Iterated extensions 61
10.4 Boundary actions 62

11 Satellites as enriched Kan extensions 64
11.1 The graded extension category 64
11.2 The comparison with resolutions 65
11.3 The dual construction 66
11.4 An enriched lift and its relation to truncation . . 67

12 Derived functors and their models 69
12.1 Resolution and localization 69
12.2 Tensor and Hom 70
12.3 Acyclic resolutions and dimension shifting 71
12.4 Composition and a spectral sequence 72

13 Inverse towers and the comparison of maps 74
13.1 The fiber of the shift difference 74
13.2 Mapping by two truncations 75
13.3 Tests and counterexamples 78

14 Cohomological amplitude and output windows 80
14.1 An input interval for a finite output interval . . . 80
14.2 Extending bounds from the heart 82
14.3 Two finite tests 83
14.4 Composition and a relative form 83

IV Higher sheaves and their cohomology 85

15 Sites and sheaves of spaces 86
15.1 The topology of a sieve 86
15.2 Sheafification as a localization 87
15.3 Stalks and concrete sheaves 88

15.4 Hypercovers and Whitehead tests 89
15.5 Topos descent and objects over a base 90

16 Change of space and supports **91**
16.1 Geometric morphisms 91
16.2 Open base change 92
16.3 A slice formula 93
16.4 Open and closed decompositions 93

17 Derived sheaves of modules **96**
17.1 Presheaves, sections, and exact sheafification . . . 96
17.2 Exact localization at local quasi-isomorphisms . . 97
17.3 Hyperdescent and the comparison 98
17.4 Varying rings and change of coefficients 100

18 Global sections and acyclic sheaves **102**
18.1 Representing global sections 102
18.2 Flabby sheaves 103
18.3 The Godement construction 104
18.4 Soft and fine sheaves 105

19 Čech complexes, totalization, and convergence **107**
19.1 The two-open calculation 107
19.2 Cosimplicial filtration 108
19.3 A Postnikov calculation of hypercohomology . . . 110

20 Finite sites and explicit cohomological bounds **112**
20.1 Sheaves on an ordered space 112
20.2 A finite projective resolution 113
20.3 Sharp and nonconstant examples 115
20.4 Relative limits 115
20.5 The incidence presentation 116

21 Local systems and higher monodromy **117**
21.1 Transport along a fundamental groupoid 117
21.2 Constant coefficients and an ordinary case 118
21.3 Classification with a fixed fiber 119
21.4 The first obstruction in sections 120

22 Group actions and derived invariants 122
22.1 Two adjoints to a constant action 122
22.2 The cyclic calculation 123
22.3 Averaging and its derived meaning 124
22.4 Free groups and a finite cellular model 125
22.5 Comparison with section cohomology 125

23 A ringed calculation on the projective line 127
23.1 Two affines and the transition function 127
23.2 Base change without flatness of the base map . . 128
23.3 A window for a quasi-coherent complex 129
23.4 A boundary realized by a line bundle 130

24 Tensor, projection, and a duality calculation 132
24.1 The projection comparison 132
24.2 A trace on the projective line 134
24.3 A finite-base-change criterion 135

25 Torsors, delooping, and connecting classes 137
25.1 Classifying a local trivialization problem 137
25.2 Abelian coefficients as infinite loop objects 138
25.3 An obstruction is a fiber of a map 139
25.4 An internal viewpoint 140

26 Internal quantification and the shape of a topos 142
26.1 Slices as contexts 142
26.2 What constant coefficients see 143
26.3 A shape that requires finite quotients 144
26.4 Two truncation operations 145

27 Products and an unbounded sheaf example 147
27.1 A product of local epimorphisms 147
27.2 A sufficient bounded-product criterion 148
27.3 A tower with a nontrivial limit at an accumulation point 149
27.4 Where the amplitude estimate fails 151

28 A finite-site differential from an extension class 153

28.1 Constructing a coefficient with two layers 153
28.2 The value of derived sections 154
28.3 Identifying the page of the differential 155
28.4 A direct description of a representative 156

29 Descent for extension spaces **158**
29.1 Gluing the coefficient categories 158
29.2 Local extension data and comparison paths 159
29.3 Reconstructing a global map from local finite windows 161

30 Local-to-global Ext and classical recovery **163**
30.1 A sheaf of derived mapping objects 163
30.2 Concordance with the source 165
30.3 Three points of contact between the constructions 167

References **169**

Part I

Representability and adjunctions

This part develops the mapping-space tests used throughout the manuscript. The Yoneda embedding identifies natural transformations with elements of a presheaf, and the category of elements gives the weighted colimit formula. The comparison of pointwise Kan extensions reduces many changes of base to tests on comma categories. A limit over a circle and the two adjoints to a constant action give the first examples in which the index contributes cohomology.

Chapter 1

Categories with mapping spaces

1.1 The nerve and its inner horns

An ordinary category $\mathcal{C}$ is encoded by its nerve $N\mathcal{C}$. An n-simplex is a chain of n composable arrows. Its face maps compose adjacent arrows or omit a first or last object. Degeneracies insert identity arrows. This encoding identifies the usual category of small categories with the full subcategory of simplicial sets whose inner horns have unique fillers. An inner horn $\Lambda^n_i \subset \Delta^n$, with $0 < i < n$, is the union of all faces except the face opposite i. For $n = 2$ it consists of the two edges $0 \to 1 \to 2$. Its filler supplies the composite edge $0 \to 2$ and the triangle relating it to those two edges.

Definition 1.1. A quasicategory is a simplicial set in which every inner horn admits a filler. A space is modeled by a Kan complex, in which horns of all indices admit fillers. Write $\mathcal{S}$ for the ∞-category of spaces. A functor between quasicategories is a simplicial map.

For a quasicategory C, the filler for an inner 2-horn supplies a choice of composite. The inner 3-horn conditions give compatibility of such choices with triple composition. Higher conditions carry the remaining coherences. This description is a presentation of the structure. Statements about universal properties will be formulated using mapping spaces and are invariant under equivalence of presentations. We use the construction and comparison results for mapping spaces from [18, §1.2] and [21].

For example, a model of the right mapping space $\mathrm{Map}_C(x, y)$ takes n-simplices to maps $\Delta^{n+1} \to C$ that are constant at x on the face with vertices $0, \dots, n$ and carry the last vertex to y. The face and degeneracy maps act on $0, \dots, n$ and fix the last vertex. The inner horn axioms imply that this is a Kan complex. Its 0-simplices are arrows. Its paths identify homotopic arrows relative to the endpoints. Composition is represented by the extension of the two adjacent edges to a triangle, and it is functorial up to the coherences of C.

Proposition 1.2. *For an ordinary category $\mathcal{C}$, $\mathrm{Map}_{N\mathcal{C}}(x, y)$ is the discrete space on $\mathrm{Hom}_{\mathcal{C}}(x, y)$. The homotopy category of $N\mathcal{C}$ is $\mathcal{C}$.*

Proof. A simplex in the right mapping space is a chain

$$x \xrightarrow{\mathrm{id}} x \xrightarrow{\mathrm{id}} \cdots \xrightarrow{\mathrm{id}} x \xrightarrow{f} y.$$

It is determined by f, and every face and degeneracy leaves that arrow unchanged. Composition of components is the composition supplied by the nerve. This identifies both objects and hom-sets of the homotopy category. □

The homotopy category $\mathrm{h}\, C$ has the objects of C and hom-sets $\pi_0 \mathrm{Map}_C(x, y)$. A morphism of C is an equivalence when its component is an isomorphism in $\mathrm{h}\, C$. The core $C^{\simeq}$ is the space of objects and equivalences of C, with their higher homotopies. For instance, if $\mathcal{C}$ is an ordinary category, $(N\mathcal{C})^{\simeq}$ is the nerve of its maximal subgroupoid. The mapping space between two ∞-categories in Cat_∞ is $\mathrm{Fun}(C, D)^{\simeq}$. The functor category $\mathrm{Fun}(C, D)$ itself includes transformations whose components need not be equivalences.

Example 1.3. Let A be a nonzero abelian group and regard $K(A, 2)$ as an ∞-groupoid. At a chosen basepoint x, its endomorphism space is $\Omega_x K(A, 2) \simeq K(A, 1)$. Its homotopy category is equivalent to a point because $K(A, 2)$ is connected and simply connected. In particular the functor $K(A, 2) \to *$ induces an

equivalence of homotopy categories and loses the nonzero group $\pi_1 \operatorname{Map}(x, x)$. The discrete hom-sets in $\mathrm{h}\, C$ retain exactly the components of each mapping space.

1.2 Equivalences, functors, and coherence

A functor $F : C \to D$ is fully faithful when the induced map $\operatorname{Map}_C(x, y) \to \operatorname{Map}_D(Fx, Fy)$ is an equivalence for every pair. It is essentially surjective when every object of D is equivalent to one in the image. The equivalence criterion is full faithfulness and essential surjectivity, as in the ordinary case with the strengthened hypothesis on maps [18, §1.2.10]. One proof selects a representative of each object in the target and an equivalence to it. Full faithfulness identifies the mapping spaces for the selected representatives with the target mapping spaces. Transport of the composition operations along these equivalences supplies a functor inverse. The criterion itself is independent of the representatives.

An n-simplex of $\operatorname{Fun}(C, D)$ is a map $C \times \Delta^n \to D$. In particular, a natural transformation $F \to G$ is a functor on $C \times \Delta^1$ with given end restrictions. On an edge $x \to y$ it gives a square with a specified commutativity homotopy. Its value on the product of a composable chain with Δ^1 gives the compatibility of those homotopies. The space of transformations between F and G is $\operatorname{Map}_{\operatorname{Fun}(C,D)}(F, G)$. It generally has higher homotopy even when C has discrete mapping spaces.

Proposition 1.4. *A natural transformation in* $\operatorname{Fun}(C, D)$ *is an equivalence exactly when every component is an equivalence.*

Proof. Each evaluation functor preserves equivalences, giving necessity. For sufficiency, let J be the nerve of the groupoid on two objects and one isomorphism between them. Restriction along $\Delta^1 \to J$ identifies $\operatorname{Fun}(J, D)$ with the full subcategory of the arrow category spanned by equivalences. This is the universal

property of inverting the edge. A diagram of arrows whose values are equivalences therefore lifts, with contractible space of choices, from $C \to \mathrm{Fun}(\Delta^1, D)$ to $C \to \mathrm{Fun}(J, D)$. The exponential law turns this into an extension $J \to \mathrm{Fun}(C, D)$ of the given edge. It exhibits the edge as an equivalence. □

Example 1.5 (Actions as diagrams)**.** For a discrete group G, $BG = N(\mathbf{B}G)$ is the space given by the nerve of the one-object groupoid with automorphisms G. A functor $BG \to \mathcal{S}$ is an object with a coherent G-action. Each g acts by an equivalence. The pair (g, h) specifies the compatibility between the composite action and the action of gh, and higher simplices specify its further compatibilities. With target Set all mapping spaces are discrete, so the same functor is a usual G-set. With target spectra its limit is the homotopy fixed-point spectrum, calculated in Chapter 22.

1.3 Slices and homotopy fibers

The slice $C_{/z}$ is the category of arrows ending in z. One simplicial definition describes a simplex $\Delta^n \to C_{/z}$ as a map $\Delta^n \star \Delta^0 \to C$ with last vertex z. For two objects $f : x \to z$, $g : y \to z$, the test for a map $f \to g$ uses the homotopy fiber

$$\mathrm{Map}_{C_{/z}}(f, g) \simeq \mathrm{Map}_C(x, y) \times_{\mathrm{Map}_C(x,z)} \{f\}, \tag{1.1}$$

where the map out of $\mathrm{Map}_C(x, y)$ is postcomposition by g. An element consists of $h : x \to y$ and a path from gh to f. The higher simplices give homotopies of these pairs with the corresponding boundary fixed. The coslice $C_{z/}$ uses cones with first vertex z and a dual formula.

Calculation 1.6. In the ∞-category of spaces over S, take two objects $* \to S$ corresponding to points s, t. Their mapping space is the path space from t to s. If $S = K(A, 2)$ and $s = t$, this is $K(A, 1)$. Thus a map over S can carry coherence even when the underlying map $* \to *$ has no choices. If S is discrete, the answer is one point when $s = t$ and empty otherwise. The two answers recover the slice formulas in spaces and in sets.

An initial object i satisfies $\mathrm{Map}_C(i, c) \simeq *$ for every c, and a terminal object t satisfies $\mathrm{Map}_C(c, t) \simeq *$. When an initial object exists, the full subcategory of initial objects is a contractible space. Indeed, every map between two initial objects is an equivalence because both composites are homotopic to identities, and every mapping space among these objects is contractible. The same proof works for terminal objects. The identity of z is terminal in $C_{/z}$ by (1.1); the relevant fiber is a fiber of an identity map.

Example 1.7. Let $C = \mathcal{S}_*$ be pointed spaces. The one-point pointed space is both initial and terminal. For a based map $f : X \to Y$, the space of factorizations through the zero object with specified comparison to f is the space of paths in $\mathrm{Map}_{\mathcal{S}_*}(X, Y)$ from the constant map to f. In particular at the constant map the choices form $\Omega \mathrm{Map}_{\mathcal{S}_*}(X, Y)$. This is the form of a specified nullhomotopy that appears in a fiber or a cofiber square.

1.4 Localizing morphisms

Given a category C and a collection W of arrows, its localization $q : C \to C[W^{-1}]$ is characterized by

$$\mathrm{Fun}(C[W^{-1}], D) \simeq \mathrm{Fun}_W(C, D).$$

The right side is the full subcategory of functors carrying W to equivalences. In this universal property the objects and all transformation spaces are compared. Existence for small data follows by adjoining inverses using the pushout that replaces each selected walking arrow by the walking equivalence. For a large category the same formulation is used in an enlarged universe or in a suitably accessible presentation. Constructions of localization by simplicial methods are developed in [10, 6].

In the category of cochain complexes, take W to be quasi-isomorphisms. Its ordinary localization gives the usual derived category, and the ∞-categorical localization retains mapping spaces. Applying the homotopy category functor to the latter gives the ordinary localization by the same universal property

tested against ordinary categories. To see the scope of the statement, any functor into $N\mathcal{A}$ factors through $\mathrm{h}\, C$, because homotopic arrows have equal images in a category with discrete mapping spaces. Coherent commutativity then becomes ordinary commutativity. Thus the tests against $N\mathcal{A}$ impose exactly the ordinary localization property. In Chapter 12 the resulting higher localization is computed using resolutions.

Example 1.8 (Groupoid completion)**.** When W contains all edges, the localization is a space, denoted $|C|$. It is left adjoint to the inclusion $\mathcal{S} \subset \mathrm{Cat}_\infty$, by testing against spaces. For an ordinary category this is the classifying space of its nerve. Let C be the poset of nonempty proper subsets of $\{0, 1, 2\}$. Its nerve is the barycentric subdivision of the boundary of a triangle, so $|C| \simeq S^1$. Inverting the order relations creates an automorphism group $\mathbb{Z}$ at a basepoint even though each endomorphism in C was an identity. Higher-dimensional face posets give higher homotopy after the same construction.

1.5 Two probes for coherent information

Proposition 1.9. *Let K be a space and let C admit a cotensor y^K with K. If the constant K-diagram at y has limit, it represents this cotensor. Explicitly,*

$$\mathrm{Map}_C(x, y^K) \simeq \mathrm{Map}_{\mathcal{S}}\Big(K, \mathrm{Map}_C(x, y)\Big).$$

Proof. The mapping space into the limit is the limit of the corresponding constant diagram of spaces. A limit of the constant space M indexed by K is the space of coherent choices of a point of M for each point of K. Formally, maps from a test space T into this limit are maps from $T \times K$ into M, by the adjunction for the product and the expression of K as the colimit of the constant point diagram over K. The exponential law gives the stated equivalence. □

For $K = S^1$ the cotensor can detect loops in every $\mathrm{Map}_C(x, y)$. Take $y = K(A, 1)$ in spaces. Then evaluation at the basepoint of S^1 exhibits $\mathrm{Map}(S^1, K(A, 1))$ with fiber the discrete group A. The evaluation fibration has a section of constant loops, and $K(A, 1)$ has a model as a simplicial abelian group. In that model subtracting the constant loop at the evaluation point yields a product decomposition

$$\mathrm{Map}(S^1, K(A, 1)) \simeq K(A, 1) \times A.$$

The formula shows directly how an ordinary hom-set into a cotensor can recover a higher homotopy group.

A second probe is the arrow category. For maps $f : x \to y$ and $g : u \to v$ there is an equivalence

$$\mathrm{Map}_{\mathrm{Fun}(\Delta^1, C)}(f, g) \simeq \mathrm{Map}_C(x, u) \times_{\mathrm{Map}_C(x,v)} \mathrm{Map}_C(y, v).$$

The two maps into $\mathrm{Map}_C(x, v)$ compose with g and f, respectively. This gives a pullback of spaces, which retains the specified homotopy of the two composites. In pointed spaces the formula simplifies when two of the corners are zero, as in the following calculation.

Exercise 1.10. In pointed spaces take $f : X \to *$ and $g : * \to Y$. Prove that $\mathrm{Map}_{\mathrm{Fun}(\Delta^1, \mathcal{S}_*)}(f, g) \simeq \mathrm{Map}_{\mathcal{S}_*}(\Sigma X, Y)$. What is its set of components for $X = S^0$ with a basepoint and $Y = K(A, 1)$?

Solution. The two end mapping spaces are contractible and the two maps into $\mathrm{Map}_{\mathcal{S}_*}(X, Y)$ choose the zero map. Their pullback is its based loop space, which is the mapping space out of the suspension by the pointed tensor adjunction. The suspension of pointed S^0 is S^1, so the components are A. □

This computation already displays the mechanism behind stable exact sequences. A square with zero on the two outer edges may carry the boundary map of a cofiber sequence as its comparison homotopy. Working in the arrow category or the slice records that map before any passage to homotopy classes.

Chapter 2

Yoneda, weights, and universal properties

2.1 Representatives and their elements

For $c \in C$, write $yc = \mathrm{Map}_C(-, c)$ for the represented presheaf on C, and $\mathcal{P}(C) = \mathrm{Fun}(C^{\mathrm{op}}, \mathcal{S})$. The functor $y : C \to \mathcal{P}(C)$ is the Yoneda embedding. A presheaf F has a category of elements $\int_C F$ equipped with a right fibration over C. Its fiber over c is $F(c)$. Along $u : c \to d$ its transport is $F(u) : F(d) \to F(c)$. This is the right-fibration case of straightening and unstraightening [18, §2.2]. In the ordinary discrete case the objects are pairs (c, a) with $a \in F(c)$, and an arrow over $u : c \to d$ is permitted when $a = F(u)b$. In spaces this equality is replaced by the space of paths from a to $F(u)b$.

Theorem 2.1 (Higher Yoneda)**.** *Evaluation at* id_c *induces a natural equivalence*

$$\mathrm{Map}_{\mathcal{P}(C)}(yc, F) \simeq F(c).$$

In particular y is fully faithful, and the space of pairs consisting of a representative of a representable presheaf and a specified representing equivalence is contractible.

Proof. The right fibration for yc is $C_{/c} \to C$. A map over C from this right fibration to the one for F is a coherent section of the pullback of F along $C_{/c} \to C$. The section space is the limit of F over $(C_{/c})^{\mathrm{op}}$. In $C_{/c}$ the identity of c is terminal, so it is initial in the opposite, and evaluation there computes the limit. This

gives the equivalence and its naturality. For $F = yd$ it gives full faithfulness. The space of representatives with a fixed equivalence to F is a fiber of the inclusion of the core of the essential image. Full faithfulness and the existence of a representative make that fiber contractible. □

The dual statement uses the covariant represented functor $\mathrm{Map}_C(c, -)$. Its version of the category of elements is a left fibration. Keeping the variances explicit is useful. For example, a right adjoint to $F : C \to D$ represents the presheaf $c \mapsto \mathrm{Map}_D(Fc, d)$, while a left adjoint to F represents the covariant functor $c \mapsto \mathrm{Map}_D(d, Fc)$.

Example 2.2 (The free generator as a representative)**.** In Set the functor $\mathrm{Hom}(*, -)$ represents the underlying set functor. For left modules over a ring R, the analogous representative is R: evaluation at 1 gives $\mathrm{Hom}_R(R, M) \cong M$ as sets. For R a ring spectrum, the free module of rank one represents the underlying infinite loop space, and its suspended copies represent the spaces at every level of the underlying spectrum. These identifications are natural in the module and are compatible with the spectrum structure. Passing to π_0 recovers evaluation at the classical generator for discrete R-modules.

2.2 Density and the category of elements

Theorem 2.3 (Density)**.** *Every presheaf $F \in \mathcal{P}(C)$ is a canonical colimit of representables, indexed by its category of elements:*

$$F \simeq \operatorname*{colim}_{(c,a) \in \int_C F} yc.$$

Proof. The evaluation equivalence of Theorem 2.1 gives a coherent cocone $yc \to F$ for $(c, a) \in \int_C F$. Evaluate at d. The colimit of spaces is computed by taking the groupoid completion of the category of elements of the diagram $(c, a) \mapsto \mathrm{Map}_C(d, c)$.

Its objects are triples $(c, a, u : d \to c)$; they map to $F(d)$ by $F(u)a$. Over $b \in F(d)$ the category with specified comparison $b \simeq F(u)a$ has initial object (d, b, id_d). Indeed, a morphism from this object to the triple is precisely the arrow u and the given comparison, with contractible space of ways of presenting them. Equivalently the functor to $F(d)$ is left adjoint to the inclusion of the identity triples, with counit an equivalence. Groupoid completion therefore gives $F(d)$. Colimits of presheaves are pointwise, so the cocone is a colimit. □

Corollary 2.4 (Free cocompletion)**.** *If D has small colimits, restriction along y is an equivalence*

$$\mathrm{Fun}^{\mathrm{colim}}(\mathcal{P}(C), D) \simeq \mathrm{Fun}(C, D).$$

Proof. For $g : C \to D$, define $\widehat{g}(F) = \mathrm{colim}_{\int F} g(c)$. Its right adjoint sends d to $c \mapsto \mathrm{Map}_D(g(c), d)$: a map out of $\widehat{g}(F)$ is the limit of the spaces $\mathrm{Map}_D(g(c), d)$ over $(\int F)^{\mathrm{op}}$, and by Theorem 2.3 this is the space of maps from F to that presheaf. Hence $\widehat{g}$ preserves colimits. On yc the index category has terminal object id_c and gives $g(c)$. Density also shows that a colimit-preserving extension and its transformations are recovered from their values on the yc. □

The ordinary density formula is obtained with sets, an ordinary category of elements, and ordinary colimits. For comparison, consider the one-object category $\mathbf{B}G$. A presheaf F with value a point and trivial action has category of elements BG. The representable presheaf at the unique object is the free transitive right G-space. The colimit of the action diagram of this representable is the homotopy orbit of a free transitive action on G, and is contractible. The colimit of the constant point diagram over BG is BG. In the density formula the action on the summands therefore matters even though their labels c are all the same object.

2.3 Weights and composition

If D is cocomplete, it is tensored over spaces. Define

$$K \odot d := \operatorname*{colim}_K d.$$

It satisfies $\mathrm{Map}_D(K \odot d, e) \simeq \mathrm{Map}_{\mathcal{S}}(K, \mathrm{Map}_D(d, e))$. The extension in Corollary 2.4 is often written as a coend

$$F \otimes_C g = \int^{c \in C} F(c) \odot g(c). \tag{2.1}$$

For a bifunctor $H : C^{\mathrm{op}} \times C \to D$, the coend can be defined using the opposite of the twisted-arrow category. An arrow $s \to t$ is evaluated as $H(t, s)$. The maps of the diagram express the compatibility of the contravariant and covariant variables. For a space weight and a diagram, (2.1) is characterized by

$$\mathrm{Map}_D(F \otimes_C g, e) \simeq \mathrm{Map}_{\mathcal{P}(C)}\Big(F, \mathrm{Map}_D(g(-), e)\Big).$$

This recovers the category-of-elements formula because both sides represent the same functor of e. Ends use the twisted-arrow category without taking the opposite, evaluating H at (s, t). They compute spaces of natural transformations from the bifunctor $(s, t) \mapsto \mathrm{Map}_D(g(s), h(t))$.

Calculation 2.5 (Ordinary tensor relations)**.** For a small preadditive category $\mathcal{A}$, a right module $M : \mathcal{A}^{\mathrm{op}} \to \mathrm{Ab}$ and a left module $N : \mathcal{A} \to \mathrm{Ab}$ are additive functors. Their coend is the quotient of $\bigoplus_a M(a) \otimes_{\mathbb{Z}} N(a)$ by all relations

$$M(u)m \otimes n = m \otimes N(u)n,$$
$$u : a \to b, \quad m \in M(b), \quad n \in N(a).$$

A homomorphism out of the quotient is exactly a family of compatible bilinear pairings, which proves the formula. Setting $M = \mathrm{Hom}_{\mathcal{A}}(-, b)$, the map $[u \otimes n] \mapsto N(u)n$ is inverted by $x \mapsto [\mathrm{id}_b \otimes x]$. Each inverse identity follows from a tensor relation. This proof of the additive co-Yoneda formula will compute the satellites in Chapter 11.

Suppose $M : C^{\mathrm{op}} \times D \to \mathcal{S}$ and $N : D^{\mathrm{op}} \times E \to \mathcal{S}$ are space-valued profunctors. Their composite has value $\int^d M(c,d) \times N(d,e)$. The identity profunctor on C is $\mathrm{Map}_C(-,-)$. Inserting it on either side gives the original profunctor by Yoneda. For three profunctors both ways of composing are equivalent to one iterated coend in the two intermediate variables, since product with a fixed space preserves colimits and the universal properties of the two colimits commute. This produces the associativity equivalence. Coherence for four factors follows because the comparison maps all realize the same universal family of tensors. In particular, the associativity maps are specified by the construction and not selected independently for each tuple.

Example 2.6 (Profunctors from functors)**.** For $f : C \to D$, the profunctor $(c,d) \mapsto \mathrm{Map}_D(f(c), d)$ represents extension along f on a suitable variance. More explicitly, the value at d of a pointwise left Kan extension of $g : C \to E$ is the C^{op}-weighted colimit with weight $\mathrm{Map}_D(f(-), d)$. Its elements are arrows $f(c) \to d$ with coherences over C. When f is fully faithful and $d = f(c_0)$, Yoneda reduces the weight to yc_0; the resulting value is $g(c_0)$.

2.4 Presheaf evaluation and detection

Write $\mathrm{ev}_c : \mathcal{P}(C) \to \mathcal{S}$ for evaluation. It has a left adjoint $K \mapsto K \odot yc$ by Yoneda and the tensor formula. It also has a right adjoint

$$K \longmapsto \Big(d \longmapsto \mathrm{Map}_{\mathcal{S}}(\mathrm{Map}_C(c,d), K) \Big).$$

For the right adjoint, a map from a presheaf F is determined by the map $F(c) \to K$, since an arrow $c \to d$ acts on $F(d)$ by restriction. The coherent statement follows from right Kan extension along the inclusion of the vertex c into C^{op}. In the formula, $c \to d$ in C is $d \to c$ in C^{op}, the required direction for a right Kan extension.

The family of representables detects equivalences. If $u : F \to G$ induces equivalences on mapping spaces out of every yc, then u is pointwise an equivalence by Yoneda and is an equivalence by Proposition 1.4. This is a detection statement in the higher category. A map of presheaves of spaces inducing bijections merely on components of all evaluations need not be an equivalence. For example on the one-object terminal site, the map $K(A, 1) \to *$ induces such a bijection and has nontrivial π_1.

Proposition 2.7. *For a finite simplicial set K and $c \in C$, the presheaf $K \odot yc$ is compact in $\mathcal{P}(C)$: its mapping-space functor preserves filtered colimits.*

Proof. By adjunction the mapping-space functor is $F \mapsto \mathrm{Map}_{\mathcal{S}}(K, F(c))$. Evaluation preserves colimits. Mapping out of a finite simplicial set into spaces preserves filtered colimits, since a finite simplicial set is a finite colimit of simplices under skeletal attachment, and filtered colimits of spaces commute with finite limits. Applying this fact to the mapping-space functor proves the claim. □

In particular the yc are compact. Compactness and detection serve different purposes. Detection ensures that no nontrivial equivalence test is missed. Compactness ensures that the tests commute with the filtered colimits used to build larger objects. Both properties of representables will reappear in the stable Morita argument.

Exercise 2.8. Let C be the category with two objects $0, 1$ and one nonidentity arrow $0 \to 1$. Compute the category of elements of $y0$ and $y1$. For a space-valued presheaf $F = (F(1) \to F(0))$, give an explicit pushout description of its density colimit.

Solution. The category of elements of $y0$ is a point at id_0, and that of $y1$ is $C_{/1} \cong C$. The diagram consists of the fibers of elements $F(0), F(1)$ and the arrows over $0 \to 1$. Its weighted colimit is the pushout

$$(F(0) \odot y0) \amalg_{F(1) \odot y0} (F(1) \odot y1).$$

The two attaching maps come from restriction $F(1) \to F(0)$ and $y0 \to y1$. Evaluation at 1 gives $F(1)$, since $y0(1)$ is empty. Evaluation at 0 gives $F(0) \amalg_{F(1)} F(1) \simeq F(0)$, with the prescribed restriction. This verifies the formula pointwise. □

Chapter 3

Limits and the shape of a diagram

3.1 Cones and their tests

For a diagram $p : K \to C$, a cone under p is an extension to the simplicial join $K \star \Delta^0$. Its apex is the last vertex. The category $C_{p/}$ is the category of such extensions and their coherent transformations. A colimit is an initial object of $C_{p/}$. The limit is defined dually by a terminal object in the category of cones over p. Their mapping-space tests read

$$\mathrm{Map}_C(\operatorname*{colim}_K p, c) \simeq \lim_{k \in K^{\mathrm{op}}} \mathrm{Map}_C(p(k), c), \tag{3.1}$$

$$\mathrm{Map}_C(c, \lim_K p) \simeq \lim_{k \in K} \mathrm{Map}_C(c, p(k)). \tag{3.2}$$

A cocone $p \to \mathrm{const}(b)$ induces the transformation in (3.1). If it is an equivalence for each c, then maps from that cocone to any other cocone are fibers of these equivalences over the other cocone, and are contractible. The converse follows from the same fiber calculation. This also explains why coherent naturality is built into the universal map.

Taking K discrete gives sums and products. Taking K the span $a \leftarrow b \to c$ gives a pushout, with mapping space out the pullback of the three mapping spaces in (3.1). For a simplicial set of indices presented as a union $K_1 \cup_{K_0} K_2$, the mapping space of compatible cones is a pullback of the corresponding cone spaces. Under the hypotheses that the glued diagram and all three constituent diagrams have colimits, it follows that the glued colimit is the pushout of the two constituent colimits over the colimit on K_0. This is the form of calculation by attachment.

Calculation 3.1. In $\mathcal{S}$, the pullback of $* \to K(A,n) \leftarrow *$ at the same point is $K(A,n-1)$ for $n \geq 1$, taking $K(A,0) = A$ discrete. For $n = 1$, applying π_0 after taking the pullback gives A. Applying π_0 first gives a pullback of three singleton sets. Even a finite limit can use a higher coherence that π_0 discards.

Here are two sufficient constructions of finite limits. Products and pullbacks give finite limits by decomposition of a finite simplicial set into finite attachments. More explicitly a diagram on an n-simplex has limit its initial vertex, and in attaching a simplex along its boundary the new cone space is the pullback of the two cone spaces over the boundary cone space. Induction on dimension and the number of simplices completes the construction. Dually one obtains finite colimits from finite coproducts and pushouts. A finite category can have infinitely many nondegenerate simplices in its nerve. A finite group supplies such a case, with one nondegenerate simplex for every word of nonidentity elements of every length. Thus a category with finite limits need not have all limits indexed by finite categories.

3.2 Changing the index

For $u : I \to J$ and $j \in J$, write $j/u = I \times_J J_{j/}$ for the category of arrows $j \to u(i)$ and $u/j = I \times_J J_{/j}$ for the category of arrows $u(i) \to j$. The orientation is significant. The functor u is final if each $|j/u|$ is contractible, and initial if u^{op} is final. These are the hypotheses in the higher cofinality theorem [18, §4.1].

Proposition 3.2. *If $u : I \to J$ is final and the indicated colimits exist, then the canonical map*

$$\operatorname*{colim}_I pu \longrightarrow \operatorname*{colim}_J p$$

is an equivalence for every $p : J \to C$. The analogous assertion for limits uses an initial functor.

Proof. Compute the left Kan extension of the terminal weight $1_{I^{\mathrm{op}}}$ along u^{op}. Evaluated at j, it is the colimit of the point over

$(j/u)^{\mathrm{op}}$, which is $|j/u|$. Thus the natural augmentation to $1_{J^{\mathrm{op}}}$ is a pointwise equivalence. Tensoring this equivalence with p over J, the weighted colimit on the left identifies with $1_{I^{\mathrm{op}}} \otimes_I pu$. This follows from the adjunction of left Kan extension and restriction applied to the test presheaf $j \mapsto \mathrm{Map}_C(p(j), c)$. The right identifies with $1_{J^{\mathrm{op}}} \otimes_J p$. Duality gives the limit assertion. □

Example 3.3 (A directed set and a cover)**.** If P is a directed poset and $Q \subseteq P$ is cofinal in the sense that for every p there exists $q \geq p$, then p/Q is a nonempty directed poset. For any finite collection of its objects choose a common upper bound in P and then one in Q; the upper bound is still in p/Q. A nonempty directed poset has contractible nerve after groupoid completion, because a map from a finite simplicial set can be coned off by such an upper bound. It follows that directed subsequences compute the same colimits. In a tower, using the opposite poset gives the same inverse limit over a cofinal subsequence. Both operations also apply to a neighborhood basis at a point, ordered by reverse inclusion for a stalk colimit.

Example 3.4 (Finality is a test on whole comma categories)**.** Let I be the face poset of the boundary of a triangle and let $u : I \to *$. Each object of the target has a lift, and I is connected. The colimit of the constant point diagram over I is S^1, so u is not final. More generally the constant point diagram over J tests only $|I| \to |J|$. To test finality object by object, use the diagrams $j' \mapsto \mathrm{Map}_J(j, j')$. Their colimit over I is $|j/u|$, while their colimit over J is contractible because j/J has initial object id_j.

A related shape operation is interchange of limits. If $p : I \times J \to C$ admits the relevant limits, then $\lim_I \lim_J p \simeq \lim_{I \times J} p \simeq \lim_J \lim_I p$. Apply (3.2) and identify both sides as spaces of cones over the same diagram. The dual statement for colimits is proved by (3.1). To interchange a limit with a colimit requires additional hypotheses. In spaces, filtered colimits commute with finite limits. One proof uses Kan models: a finite set of simplices and their compatibility witnesses descends to a common stage of a filtered

diagram. Horn filling and tests by finite spheres yield equivalences for the canonical comparison. The ∞-categorical statement also holds for filtered ∞-categories by the corresponding finite extension property [18, §5.3].

3.3 Pointwise Kan extension

For $u : I \to J$ and $F : I \to D$, if D is cocomplete then

$$(\operatorname{Lan}_u F)(j) \simeq \operatorname*{colim}_{(i,\, u(i)\to j)\in u/j} F(i). \tag{3.3}$$

A morphism $j \to j'$ acts on this index by composition, so the colimits define a functor. A natural transformation from it to $G : J \to D$ is a coherent family of maps $F(i) \to G(j)$ over arrows $u(i) \to j$. Its value on each $\mathrm{id}_{u(i)}$ gives a transformation $F \to Gu$. Conversely such a transformation supplies the family by composing with $G(u(i) \to j)$. The correspondences are inverse on transformation spaces by the universal property of each colimit. Dually,

$$(\operatorname{Ran}_u F)(j) \simeq \operatorname*{lim}_{(i,\, j\to u(i))\in j/u} F(i) \tag{3.4}$$

when D is complete. For a target with only some limits or colimits, these formulas apply provided the individual slice limits or colimits exist.

Proposition 3.5. *If u is fully faithful, the unit $F \to u^* \operatorname{Lan}_u F$ and counit $u^* \operatorname{Ran}_u F \to F$ are equivalences whenever the indicated Kan extensions exist.*

Proof. For $j = u(i_0)$ full faithfulness identifies u/j with $I_{/i_0}$, whose terminal object is id_{i_0}. The colimit in (3.3) is then $F(i_0)$ with the specified structural map. For right Kan extension use the initial object of $I_{i_0/}$ and (3.4). □

Calculation 3.6 (Upper bounds)**.** A poset is an ordinary category with at most one arrow between any two objects. For a monotone diagram $p : K \to P$, the colimit exists exactly when

the set of all $p(k)$ has a least upper bound. All coherent choices reduce to the order relation by Proposition 1.2. If $i : P \to Q$ is monotone and Q is a complete lattice, a map $f : P \to Q$ has left Kan extension along i given by

$$q \longmapsto \bigvee_{p\colon i(p) \leq q} f(p).$$

The empty join is the least element. Right Kan extension uses the meet over p with $q \leq i(p)$. Taking $f = \mathrm{id}_P$ for a complete lattice P shows the familiar construction of a left or right adjoint to an inclusion when the respective floor or ceiling has a universal inequality. This formula is the discrete order-valued instance of the weighted construction.

3.4 A limit with a cycle

For a space K, the limit of a constant diagram at a spectrum E is the function spectrum

$$\lim_{K} E \simeq \mathrm{map}_{\mathrm{Sp}}(\Sigma^{\infty}_{+} K, E).$$

This follows by testing maps from an arbitrary spectrum and using the tensor of spectra over spaces. At $K = S^1$ the suspension spectrum of S^1 with a disjoint basepoint splits as $\mathbb{S} \vee \Sigma\mathbb{S}$ after choosing a point of the circle. Therefore the constant limit is $E \oplus \Sigma^{-1} E$. For $E = \mathrm{H}A$ this records one copy of A in cohomological degrees 0 and 1. If the underlying shape were replaced by a point, the second copy would disappear.

For an endomorphism $v : E \to E$ that is an equivalence, consider the corresponding $\mathbb{Z}$-action, using the positive generator to act by v. The limit over $B\mathbb{Z}$ is the homotopy equalizer of $1, v : E \rightrightarrows E$. A cone is a point equipped with a path comparing its value and its transport along the generator. One way to prove this for coherent diagrams is to use the free cellular model for the universal cover of S^1: it has one free orbit of vertices and one free orbit of edges. Mapping this cellular diagram equivariantly

into the diagram at E gives the homotopy equalizer. In spectra the equalizer is the fiber of $1 - v$, so its homotopy groups sit in short exact sequences with the kernels and cokernels of $1 - v$ in adjacent degrees. The same two-term calculation gives the first group cohomology of an infinite cyclic group.

Exercise 3.7. Let $C = [1]$ and let $i : * \to C$ select the initial vertex 0. Give $\mathrm{Lan}_i E$ and $\mathrm{Ran}_i E$ in a pointed complete and cocomplete category. Repeat for the final vertex 1.

Solution. For 0, the arrow $0 \to j$ exists uniquely at both $j = 0, 1$. Thus $\mathrm{Lan}_i E$ is $E \xrightarrow{1} E$. Arrows $j \to 0$ exist only at $j = 0$, so $\mathrm{Ran}_i E$ is $E \to 0$. For 1, $\mathrm{Lan}_i E$ is $0 \to E$ and $\mathrm{Ran}_i E$ is $E \xrightarrow{1} E$, with the same calculation applied to the reversed order. □

Chapter 4

Adjoints and coherent change of base

4.1 Adjunctions from representability

An adjunction $L : C \rightleftarrows D : R$ is a natural equivalence

$$\mathrm{Map}_D(Lc, d) \simeq \mathrm{Map}_C(c, Rd).$$

Naturality is in $C^{\mathrm{op}} \times D$. If L is given, let $P : D \to \mathcal{P}(C)$ take d to $\mathrm{Map}_D(L-, d)$. If each $P(d)$ is representable, P factors through the full essential image of the Yoneda embedding. Composing with its inverse equivalence to C yields R. The contractibility of the space of representing choices in Theorem 2.1 shows that the construction determines the adjunction, not merely a list of object assignments.

The identity of Lc corresponds to the unit $c \to RLc$, and the identity of Rd corresponds to the counit $LRd \to d$. For example, applying the inverse correspondence to a map $c \to Rd$ gives $Lc \to LRd \to d$. Taking $d = Lc$ and the unit gives the triangle homotopy at Lc. A dual argument gives the triangle at Rd. An adjunction defined by an equivalence of bifunctors includes coherence of these operations with all simplices of the functors [25, Chs. 4–5].

Proposition 4.1. *For $L \dashv R$, L is fully faithful exactly when the unit is an equivalence. The dual criterion for R uses the counit.*

Proof. The map $\mathrm{Map}_C(c, c') \to \mathrm{Map}_D(Lc, Lc')$ becomes post-composition by the unit $c' \to RLc'$ under adjunction. Yoneda identifies equivalence on these tests for every c with equivalence

of the unit. The right-adjoint assertion follows by taking opposites. □

Adjoints compose. For $L_1 \dashv R_1$ and $L_2 \dashv R_2$ with compatible domains, the equivalences of mapping spaces give $L_2 L_1 \dashv R_1 R_2$. Every parenthesization is induced from the same chain of mapping-space equivalences. As a consequence, if a family of pullback functors is coherent, the corresponding right adjoints have coherent composition comparisons in the reversed order. This accounts for the functoriality of direct and inverse images in Chapter 16.

4.2 Adjoints preserve the corresponding universal constructions

Proposition 4.2. *A left adjoint preserves all existing colimits, and a right adjoint preserves all existing limits.*

Proof. For $p : K \to C$ and $b = \operatorname{colim} p$, the functor of d represented covariantly by Lb is

$$\operatorname{Map}_D(Lb, d) \simeq \operatorname{Map}_C(b, Rd) \simeq \lim_{K^{\mathrm{op}}} \operatorname{Map}_D(Lp(-), d).$$

This proves that the induced cocone with apex Lb is colimiting. The limit statement follows by duality. The proof also shows that no assumption about existence of a colimit of Lp is needed beyond the source colimit and the adjunction. □

The accessible adjoint functor theorem for presentable ∞-categories has two forms [18, Cor. 5.5.2.9]. A functor between presentable categories is a left adjoint if and only if it preserves small colimits. It is a right adjoint if and only if it is accessible and preserves small limits. An accessible functor preserves κ-filtered colimits for some regular cardinal κ. Accessibility limits the size of the generating diagrams used in constructing a representative. Preservation of limits alone proves the compatibility

of a candidate with universal cones, whereas the smallness condition ensures that it has a representative in the category under consideration.

Example 4.3 (Modules and adjoints)**.** For a map of associative ring spectra $A \to B$, restriction on left modules is right adjoint to $B \otimes_A -$ and left adjoint to $\mathrm{map}_A(B, -)$ with its induced left B-action. For the first adjunction, a B-linear map from $B \otimes_A M$ is equivalently an A-linear map from M to the restricted target. For the second, evaluate an A-linear map on the unit of B; B-linearity then recovers the original map through the right action of B on B. These equivalences are natural in both variables by the tensor–Hom universal properties. When A, B are Eilenberg–Mac Lane rings the formulas are derived tensor and derived Hom. Restriction carries both adjoints because it preserves limits and colimits as tested on underlying spectra and is accessible.

Example 4.4 (Galois connections)**.** For posets P, Q, the adjunction condition is $L(p) \leq q$ exactly when $p \leq R(q)$. If P has all joins and L preserves them, define $R(q) = \bigvee\{p \mid L(p) \leq q\}$. Preservation gives $L(R(q)) \leq q$, so the membership condition is equivalent to $p \leq R(q)$. For a map of sets $f : X \to Y$, inverse image $f^{-1} : \mathcal{P}(Y) \to \mathcal{P}(X)$ has both adjoints. Its left adjoint sends S to the image $f(S)$. Its right adjoint sends S to $\{y \mid f^{-1}\{y\} \subseteq S\}$, including points with empty fiber. This is the posetal form of existential and universal quantification.

4.3 Mates of a comparison square

Let

$$\begin{array}{ccc} C' & \xrightarrow{u} & C \\ {\scriptstyle L'}\downarrow & & \downarrow{\scriptstyle L} \\ D' & \xrightarrow[v]{} & D \end{array}$$

be a square equipped with a transformation $\alpha : Lu \to vL'$, and suppose $L \dashv R$ and $L' \dashv R'$. Its right mate is the transformation

$$uR' \longrightarrow RLuR' \xrightarrow{R\alpha R'} RvL'R' \longrightarrow Rv. \tag{4.1}$$

The first and last maps use the unit of $L \dashv R$ and the counit of $L' \dashv R'$. Conversely, from $\beta : uR' \to Rv$ form

$$Lu \longrightarrow LuR'L' \xrightarrow{L\beta L'} LRvL' \longrightarrow vL'.$$

The two triangle identities show that these operations give inverse equivalences between the corresponding transformation spaces. Inserting the formulas for two pasted squares and cancelling adjacent unit–counit composites gives the mate of the pasted comparison. This identity holds at the level of spaces of choices. The operation is used in Beck–Chevalley comparisons.

An equivalence α need not have an equivalence as mate. The two transformations live between different pairs of functors, and the units and counits can change the pointwise result. A useful sufficient test applies to restrictions and pointwise right Kan extensions. If a comparison between the right Kan extensions is induced, at d', by an initial functor between the comma categories that occur in (3.4), then the mate is a pointwise equivalence by Proposition 3.2. In general its fiber or cofiber can record the missing part of the index.

Calculation 4.5 (A square without Beck–Chevalley)**.** Let $B = [1]$, $u : \{0\} \to B$, and $v : \{1\} \to B$. The pullback square of indexing categories has empty upper-left corner. For $E \in \mathcal{S}$, $\mathrm{Ran}_v E$ at 0 is E, since there is exactly one arrow $0 \to 1$. The right Kan extension along $\varnothing \to \{0\}$ of the restricted diagram has value $*$, the empty limit. Thus restriction of $\mathrm{Ran}_v E$ along u has comparison $E \to *$, which is an equivalence exactly when E is contractible. Pullback of the underlying categories ensures that the same objects are selected; the comma index for right Kan extension can include arrows leaving those objects.

4.4 Reflection and localization

A full reflective subcategory $D \subseteq C$ has a left adjoint $a : C \to D$ to the inclusion i. Its unit identifies ia as an idempotent monad up to coherent equivalence. For $z \in D$,

$$\mathrm{Map}_C(iax, iz) \simeq \mathrm{Map}_C(x, iz).$$

Call a map an a-equivalence when a sends it to an equivalence. Every unit $x \to iax$ is such a map. Conversely a functor $F : C \to E$ inverting all a-equivalences inverts every unit and hence identifies F with Fia. This proves the localization property on objects and transformation spaces. The counit $ai \simeq \mathrm{id}_D$ proves the inverse identity on restriction.

Example 4.6 (Components and discrete spaces)**.** The inclusion $\mathrm{Set} \to \mathcal{S}$ of discrete spaces is fully faithful, with left adjoint π_0. A map into a discrete space factors through components, and its mapping space is itself discrete. The localization is not left exact: the diagram in calculation 3.1 with $n = 1$ is a pullback with discrete nontrivial fiber, but the pullback of the reflected diagram is a point. Sheafification of spaces, by contrast, is a left exact localization. The finite-limit condition is what permits a reflector to support the internal pullback and descent calculations of an ∞-topos.

4.5 Adjoint triples and parametrized objects

Let $\mathcal{X}$ be a locally cartesian closed category with pullbacks. For a map $f : X \to Y$ the base-change functor $f^* : \mathcal{X}_{/Y} \to \mathcal{X}_{/X}$ has a left adjoint $f_!$ by postcomposition. The adjunction follows from the pullback universal property in the slice. Its right adjoint f_* is dependent product. The fiber of $f_*Z \to Y$ in spaces over a point y is the space of sections of $Z|_{X_y} \to X_y$. The fiber of $f_!Z \to Y$ is the total space $Z|_{X_y}$. In the case of families of sets, the two fibers

are the product and coproduct of the fibers over X_y. This also explains why changing a set-valued diagram to a space-valued one changes an ordinary product of fiber values to a limit over a possibly nondiscrete fiber.

Exercise 4.7. For $f : BG \to *$, identify $f_!, f^*, f_*$ on slices of spaces with the operations on G-actions.

Solution. Straightening over the groupoid BG identifies spaces over BG with G-spaces, and f^* is the constant action on a space. Its two adjoints are colimit and limit over BG, by (3.1) and (3.2). They are the homotopy orbit space and the homotopy fixed-point space. The first agrees with the total space of the corresponding bundle over BG. The second is its section space. □

Part II

Exactness and stabilization

The abelian image construction and the effective image in a topos have different universal properties. Stable fibers and cofibers give a third setting, where exact sequences are coherent diagrams. A t-structure places ordinary kernels and cokernels in the heart and supplies the Postnikov layers used in the comparison theorems. Compact generation relates these constructions to module categories through the endomorphism spectrum of a generator.

Chapter 5

Images, fibers, and pointed exactness

5.1 Monomorphisms in an ∞-topos

In the classical construction of images, a monomorphism is detected by the number of lifts of a test map. The higher formulation controls the space of lifts. A space is (-1)-truncated if it is empty or contractible. A map $m : E \to B$ in a category C with pullbacks is (-1)-truncated if each homotopy fiber of $\mathrm{Map}_C(T, m)$, for every T, is (-1)-truncated. Equivalently the relative diagonal

$$E \longrightarrow E \times_B E$$

is an equivalence. This follows on applying $\mathrm{Map}_C(T, -)$: the diagonal of a space into its product with itself is an equivalence precisely when the space is empty or contractible. We use monomorphism for a (-1)-truncated map. By Proposition 1.2, this recovers left cancellation in an ordinary category with pullbacks.

For example, in $\mathcal{S}$ a monomorphism is equivalent over its target to the inclusion of a union of components. A homotopy fiber is empty or contractible and emptiness is constant along any path in the base. The map restricts to an equivalence onto the union with nonempty fibers. Thus the monomorphisms into a connected space have precisely the empty and the full subobject as classes. If X is connected and noncontractible, a point $* \to X$ fails to be monic: its fiber at that point is ΩX. When X is discrete, monomorphisms again specify ordinary subsets.

Definition 5.1. In an ∞-topos $\mathcal{X}$, a map $e : U \to V$ is an

effective epimorphism if the colimit of its augmented Čech nerve is V. The n-simplices of this nerve before augmentation are $U^{\times_V (n+1)}$.

In spaces, effective epimorphisms are precisely the maps surjective on components. The realization of a Čech nerve is stable under pullback. Over a point v, the simplicial object is the nerve with n-simplices U_v^{n+1}. If U_v is nonempty, a point supplies an extra degeneracy contracting this nerve to a point. If U_v is empty, the realization is empty. Thus the augmentation is an equivalence precisely when all fibers are nonempty. In a general ∞-topos the corresponding condition is local nonemptiness and is encoded by the relative (-1)-truncation [18, §6.2]. The condition says more than an epimorphism test in a homotopy category. For $* \to S^1$ the effective epimorphism condition holds, yet distinct degree maps $S^1 \to S^1$ become equal after composing with the point in the homotopy category.

Proposition 5.2 (Image from relative truncation)**.** *In an ∞-topos, a map $f : X \to Y$ factors, up to contractible choice, as an effective epimorphism followed by a monomorphism. The factorization is preserved by base change and its monomorphism is initial among subobjects of Y through which f factors.*

Proof. In the slice over Y, take the (-1)-truncation $I = \tau_{\leq -1}^{/Y} X$. Its universal property provides the desired factorization and shows that the maps to any subobject of Y out of I and out of X are equivalent. The image can also be computed as the realization of the Čech nerve of f: effectivity of groupoid objects identifies the realization with a quotient for X, and its map to Y has diagonal an equivalence. These are standard ∞-topos axioms and consequences [18, §§6.1–6.2]. The map $X \to I$ is an effective epimorphism by effectivity. Pullbacks preserve the nerve and its realization, giving stability under base change. The universal property also gives uniqueness with a contractible space of choices. □

The construction reduces to the image of a function when

both objects are discrete. In the heart of a stable t-structure, the ordinary kernel–cokernel factorization gives a second image construction. These two uses of image apply in different ambient categories. The following pointed calculation explains which conditions are present before passing to a heart.

5.2 Zero objects and the data of a nullhomotopy

A pointed category C has a zero object 0, both initial and terminal. The zero map $0_{X,Y} : X \to Y$ is the composite through 0. Its specified factorization has contractible space of choices. For a map $f : X \to Y$, define

$$\mathrm{fib}(f) = X \times_Y 0, \qquad \mathrm{cofib}(f) = Y \amalg_X 0.$$

The first is a universal arrow into X with specified nullhomotopy of the composite with f. The second is a universal arrow out of Y with the dual property. For every test object T,

$$\mathrm{Map}_C(T, \mathrm{fib}\, f) \simeq \mathrm{fib}_0\Big(\mathrm{Map}_C(T, X) \to \mathrm{Map}_C(T, Y)\Big),$$
$$\mathrm{Map}_C(\mathrm{cofib}\, f, T) \simeq \mathrm{fib}_0\Big(\mathrm{Map}_C(Y, T) \to \mathrm{Map}_C(X, T)\Big).$$

Here a fiber includes a path from the image to the zero map. For a discrete category these are the familiar kernel and cokernel tests, since a fiber of discrete spaces is a set and its elements have no extra path choices.

Let $f = 0_{X,Y}$ in pointed spaces. The test for the fiber gives $X \times \Omega Y$, and the cofiber is $Y \vee \Sigma X$. In particular, the fiber of $0 \to Y$ is ΩY, which can be nontrivial. In an ordinary pointed category, the kernel of $0 \to Y$ is zero, because the only candidate map into 0 already composes to zero. Thus replacing the word kernel by fiber also changes what a zero-source example can detect.

Proposition 5.3. *In a stable ∞-category the monomorphisms in the mapping-space sense are exactly the equivalences.*

Proof. For a monomorphism f each zero fiber of $\mathrm{Map}(T, f)$ is nonempty because it contains the zero map. It is therefore contractible. The fiber formula and Yoneda give $\operatorname{fib} f \simeq 0$. In a stable category $\operatorname{cofib} f \simeq \Sigma \operatorname{fib} f$, so $\operatorname{cofib} f \simeq 0$, and f is an equivalence. An equivalence is a monomorphism by the relative diagonal test. □

For example, $\mathbb{Z} \xrightarrow{2} \mathbb{Z}$ is a monomorphism of abelian groups. In $\mathrm{D}(\mathbb{Z})$ its fiber is $(\mathbb{Z}/2)[-1]$, so it fails the higher test when the mapping spaces from all objects of the stable category are allowed. Restricting test objects to the heart recovers the monomorphism test in Ab. The shift in the fiber is exactly what the smaller family of tests ignores.

5.3 Kernels and cokernels in the heart

Let C be stable with a t-structure, and let $A, B \in C^{\heartsuit}$. For $f : A \to B$, the long exact sequence of homotopy objects gives

$$\pi_1 \operatorname{cofib}(f) = \ker(f), \quad \pi_0 \operatorname{cofib}(f) = \operatorname{coker}(f),$$
$$\pi_0 \operatorname{fib}(f) = \ker(f), \quad \pi_{-1} \operatorname{fib}(f) = \operatorname{coker}(f).$$

All other homotopy objects of these two objects are zero, and the objects are in the indicated finite windows. One can obtain the kernel and cokernel directly as $\tau_{\geq 0} \operatorname{fib} f$ and $\tau_{\leq 0} \operatorname{cofib} f$, respectively. These one-sided truncations already lie in the heart: since A, B lie in the heart, closure of $C_{\leq 0}$ under fibers gives $\operatorname{fib} f \in C_{\leq 0}$, while closure of $C_{\geq 0}$ under cofibers gives $\operatorname{cofib} f \in C_{\geq 0}$. The remaining truncation therefore places each object in both halves. Testing the first object against objects of the heart gives the ordinary kernel universal property by adjunction. Testing the second covariantly gives the cokernel property. The theorem that the heart is abelian supplies $\operatorname{coim}(f) \cong \operatorname{im}(f)$ in the heart [19, §1.2.1]; Chapter 8 spells out its use.

In the stable category the two-step operations behave differently. In the fiber sequence $\operatorname{fib} f \to X \to Y$, the cofiber of the

first arrow is Y. In the cofiber sequence $X \to Y \to \operatorname{cofib} f$, the fiber of the second arrow is X. Thus applying the formal stable version of coimage and image directly would give Y and X, respectively. A nontrivial stable factorization theory can instead use a t-structure, or distinguish admissible classes as in an exact ∞-category. Ordinary normality and conormality axioms for pointed categories, such as those used by Strooker [38, §2.5], are best compared with the abelian heart and with its selected exact sequences.

Calculation 5.4 (Two ways to take a pushout)**.** In Ab the pushout of $0 \leftarrow A \to 0$ is zero. In $\mathrm{D}(\mathbb{Z})$ the same span of heart objects has pushout $A[1]$. The cocone into an object T consists of the unique maps out of zero and a loop at the zero map $A \to T$. Its space is $\operatorname{Map}(A[1], T)$. For $T = B$ in the heart this loop space is contractible because $\operatorname{Map}(A, B)$ is discrete. Hence testing only in the heart again gives zero. The two answers implement different universal properties with different test objects.

5.4 Effective equivalence relations and higher quotients

The Čech nerve of $U \to V$ is a groupoid object. Its arrows join two lifts of the same point of V, with the path identifying their images included in the pullback. In an ∞-topos every groupoid object is effective: it is equivalent to the Čech nerve of its map from degree zero to its colimit. This is the higher analog of effectivity of equivalence relations in an ordinary topos. Ordinary equivalence relations require the arrow object to be a subobject of the product of the object of objects with itself. Higher groupoid objects admit automorphisms and higher isotropy, so their realizations can be classifying spaces rather than sets of orbits.

Example 5.5 (Quotient of a point by a group)**.** For a discrete group G, the nerve of its action on a point has n-simplices G^n. Its realization is BG. Its Čech description is the nerve of $* \to BG$,

since the n-fold path comparisons yield $(\Omega BG)^n \simeq G^n$. The set of orbits is one point, which is $\pi_0 BG$. The fundamental group records the stabilizer of that orbit. Taking this quotient in a topos of sheaves allows stabilizers that vary with the base.

Exercise 5.6. For a map of spaces $X \to Y$, show that the canonical image construction agrees with the colimit of its Čech nerve even when X and Y are disconnected. Use it to compute the image of an inclusion of a single point into $S^1 \amalg S^2$.

Solution. Pull the augmented nerve back along each point y. If the homotopy fiber X_y is nonempty, the extra degeneracy described above contracts the realized fiber. Otherwise it is empty. This defines the union of the components of Y met by X, with its inclusion. The image is whichever component contains the chosen point. In particular its homotopy type is that of the whole circle or sphere, even though its source is contractible. □

Chapter 6

Stable categories and exact diagrams

6.1 Stability

A stable ∞-category is pointed, has finite limits and colimits, and has the property that its pushout squares and pullback squares coincide. Equivalently it is pointed with finite limits and colimits and its suspension and loop functors are inverse equivalences [19, §1.1]. Suspension is the pushout $0 \amalg_X 0$, and looping is the pullback $0 \times_X 0$. Their adjunction can be read from the mapping-space tests

$$\mathrm{Map}(\Sigma X, Y) \simeq \Omega \,\mathrm{Map}(X, Y) \simeq \mathrm{Map}(X, \Omega Y).$$

A square realizing suspension becomes a pullback, so $X \simeq \Omega\Sigma X$. The dual square gives $\Sigma\Omega Y \simeq Y$. This also shows explicitly why suspension can be iterated in either direction.

The ∞-categories of spectra and of derived modules over a ring are basic examples. Categories of functors with stable target are stable, because their limits and colimits are computed pointwise. For a reflective localization $L : C \to D$, a sufficient condition for stability is that its full subcategory of local objects is closed under suspension in C. Local objects are already closed under limits; closure under suspension then makes this subcategory stable, and the reflection preserves finite colimits and is exact.

The localizations used below are reflections onto $\mathcal{K}^{\perp}$, where $\mathcal{K}$ is a stable colimit-closed subcategory and the reflection exists, as in the accessible presentable setting. For $Y \in \mathcal{K}^{\perp}$ and $K \in \mathcal{K}$, stability gives

$$\mathrm{Map}_C(K, \Sigma Y) \simeq \mathrm{Map}_C(\Omega K, Y) \simeq *.$$

Thus the local objects are closed under suspension. Their finite limits are calculated in C, and stability identifies their finite colimits with the corresponding finite limits. This supplies the finite-diagram test for the exactness of the reflection.

Proposition 6.1. *For a functor between stable ∞-categories, preservation of finite limits, preservation of finite colimits, and preservation of the zero object and cofiber sequences are equivalent. Such a functor is called exact.*

Proof. Preservation of either class of finite diagrams preserves the zero object and bicartesian squares. It therefore implies preservation of the other class using products or coproducts and pullbacks or pushouts. To prove the criterion by cofibers, a functor preserving zero and cofiber sequences preserves suspension, via the cofiber of $X \to 0$. The canonical maps comparing inverse suspensions are therefore equivalences as well. The equivalence of fiber with desuspension of cofiber shows preservation of fibers. Binary products and coproducts in the source and target are biproducts, as shown below. The split cofiber sequence $X \to X \oplus Y \to Y$ with its specified section and retraction is preserved. It identifies $F(X \oplus Y)$ with $FX \oplus FY$ by the splitting of a cofiber sequence. Now finite products and fibers suffice: a pullback of $X \to Z \leftarrow Y$ is the fiber of $(f, -g) : X \oplus Y \to Z$, and a functor preserving the biproducts preserves addition and negatives. This yields finite limits. □

Here is the biproduct argument used in the proof. The square defining the product $X \times Y$ has projections and the two maps from the factors given by identities and zeros. The split cofiber sequence $X \to X \amalg Y \to Y$ is obtained by pasting pushout squares, and the comparison $X \amalg Y \to X \times Y$ induces the identity on its fiber X and base Y. The latter maps form a fiber sequence by the product definition. Testing the two fiber sequences on mapping spaces shows the middle comparison is an equivalence: in spectra this is the five lemma, and in a stable category it follows by applying $\mathrm{Map}(\Sigma^n T, -)$ and the long exact sequence of pointed homotopy sets in each shift. We may therefore use

matrix notation for finite diagrams. The addition of two maps $u, v : X \to Y$ is the composite of the diagonal, $u \oplus v$, and the codiagonal. This agrees with addition from the infinite loop structure below.

6.2 Mapping spectra

For X, Y in a stable category C, the spaces $\mathrm{Map}_C(X, \Sigma^n Y)$, with structure equivalences

$$\mathrm{Map}_C(X, \Sigma^n Y) \simeq \Omega \,\mathrm{Map}_C(X, \Sigma^{n+1} Y),$$

form a spectrum $\mathrm{map}_C(X, Y)$, functorial contravariantly in X and covariantly in Y. Its homotopy groups satisfy

$$\pi_i \,\mathrm{map}_C(X, Y) = \mathrm{Hom}_{\mathrm{h}\, C}(\Sigma^i X, Y). \tag{6.1}$$

Indeed, looping the mapping spaces is the same as applying a suspension to their source, and the structure equivalences make the identification independent of the chosen level. Exactness in either variable follows from the mapping-space limit tests and the fact that all cofiber sequences are fiber sequences with their specified comparison data.

Composition of the mapping spectra gives C its canonical spectral enrichment [19, §1.4]. At the level of maps, this composition sends a map $\Sigma^i X \to Y$ and $\Sigma^j Y \to Z$ to the map from $\Sigma^{i+j} X$ obtained by suspending the first and composing with the second. The enrichment records all the coherent paths witnessing compatibility with identities, suspensions, and addition. In particular $\mathrm{h}\, C$ is additive. Its zero maps and addition agree with the operations defined by biproducts.

Example 6.2. For the sphere spectrum $\mathbb{S}$, $\mathrm{map}_{\mathrm{Sp}}(\mathbb{S}, E) \simeq E$. For a ring spectrum R, $\mathrm{map}_{\mathrm{Mod}_R}(R, E)$ is the underlying spectrum of E, with the action encoded by the endomorphism spectrum of the free module. For $R = \mathrm{H}\mathbb{Z}$ and discrete abelian groups A, B, the negative homotopy groups of $\mathrm{map}_{\mathrm{D}(\mathbb{Z})}(A, B)$ are $\mathrm{Ext}^n_{\mathbb{Z}}(A, B)$

in degree $-n$. In particular $\mathrm{Map}_{\mathrm{D}(\mathbb{Z})}(A, B)$ is discrete while the mapping spectrum can have nonzero negative homotopy groups. The spaces obtained after suspending B recover those groups.

6.3 Triangles and the boundary map

The bicartesian square associated to $A \to B$ defines $C = \mathrm{cofib}(A \to B)$. Pushing out along $B \to 0$ gives a boundary $C \to \Sigma A$. The map $B \to C$ is the fiber of the boundary, by the equivalence between pullbacks and pushouts. A cofiber sequence with its boundary is thus a coherent four-term diagram

$$A \longrightarrow B \longrightarrow C \longrightarrow \Sigma A.$$

The corresponding sequence in $\mathrm{h}\, C$ is a distinguished triangle. Under rotation the last map of the rotated triangle is minus the suspension of the first. This is the usual convention from cochain cones: in $\mathrm{Cone}(f)^n = B^n \oplus A^{n+1}$ the differential is the matrix with diagonal $d_B, -d_A$ and off-diagonal f. Reversing one edge of the suspension square accounts for the same minus sign in stable categories.

For a fixed test object T, applying $\mathrm{map}(T, -)$ gives a fiber sequence of spectra. The long exact sequence of π_i is consequently exact in all degrees. Boundary maps are induced by $C \to \Sigma A$, so the choice of coherent sequence specifies them. The dual statement for $\mathrm{map}(-, T)$ uses the rotated sequence and contravariant composition. This is the source of the two long exact extension sequences.

Lemma 6.3 (Split sequences)**.** *For a cofiber sequence $A \to E \xrightarrow{p} B \xrightarrow{\delta} \Sigma A$, a choice of section of p up to a specified comparison is equivalent to a choice of nullhomotopy of δ. Given such a choice, $E \simeq A \oplus B$ over B and under A, and the space of sections is a torsor for $\mathrm{Map}(B, A)$.*

Proof. Use $E = \mathrm{fib}(\delta)$. The fiber over id_B of $\mathrm{Map}(B, E) \to \mathrm{Map}(B, B)$ is the space of paths from δ to zero in $\mathrm{Map}(B, \Sigma A)$.

If a path is specified, translation by it identifies the space with $\Omega \operatorname{Map}(B, \Sigma A) \simeq \operatorname{Map}(B, A)$. A section s gives $(i, s) : A \oplus B \to E$, which induces identities on the fiber and the base of the projection to B. The fiber-sequence test gives the equivalence. □

6.4 Pasting and the octahedral comparison

Let $A \xrightarrow{f} B \xrightarrow{g} C$ be a coherent composable pair. There is a cofiber sequence

$$\operatorname{cofib}(f) \longrightarrow \operatorname{cofib}(gf) \longrightarrow \operatorname{cofib}(g). \tag{6.2}$$

To construct it, the first map is induced by g and the identity of A. Its cofiber is computed by pasting

$$(C \amalg_A 0) \amalg_{B \amalg_A 0} 0 \simeq C \amalg_B 0.$$

Indeed, mapping out sends this identity to the result of composing pullback squares of mapping spaces. Both descriptions supply the same coherent cone on $A \to B \to C$ with B sent to zero. The boundary of (6.2) is therefore specified by the original pair. Together with the three defining cofiber sequences this is the octahedral diagram in the homotopy category. The construction is functorial on coherent pairs. The octahedral axiom in the homotopy category retains the existence of the diagram, while the stable construction also retains its natural dependence on the input.

There is a related three-by-three formula. Given a square of maps, take the cofibers of its two rows and the induced map between them. The cofiber of this last map agrees with the result of first taking column cofibers and then taking the induced cofiber between them. In either order one is taking a finite colimit in two variables. Since finite colimits commute with finite colimits, (3.1) identifies both with the same total cofiber. The total fiber of the

square is its twice desuspension. When the square is bicartesian the total cofiber is zero, and conversely a vanishing total cofiber forces the comparison from the pushout to the final vertex to be an equivalence. This is a useful criterion for diagram exactness in module spectra.

Exercise 6.4. For $u : X \to Y$, compute the cofiber of the map $X \to X \oplus Y$ with components $(1, u)$. Determine the second map of the cofiber sequence as a row matrix.

Solution. The automorphism of $X \oplus Y$ with matrix

$$\begin{pmatrix} 1 & 0 \\ -u & 1 \end{pmatrix}$$

carries $(1, u)$ to the first inclusion. Its inverse has u in the lower left entry. The cofiber of the first inclusion is Y with its projection $p_2 = (0, 1)$. Precomposing this projection with the displayed automorphism gives

$$(0, 1) \begin{pmatrix} 1 & 0 \\ -u & 1 \end{pmatrix} = (-u, 1).$$

The boundary is zero because the sequence is split. □

Chapter 7

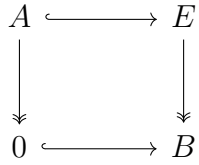
Exact categories and derived enhancements

7.1 Admissibility and exact structure

Strooker calls a pointed category exact when its mono–epi and kernel–cokernel axioms give a specified factorization theory [38, Def. 2.5.9]. In modern exact category theory one begins with an additive category and specifies which kernel–cokernel pairs are to count as exact. Every abelian category carries its usual exact structure. The category of finitely generated projective modules carries the split exact structure, even when it has very few kernels or cokernels as an ordinary category. The admissible pairs in such a category therefore record data beyond its underlying additive category.

An exact ∞-category has an additive underlying ∞-category and two compatible classes of arrows, called ingressives and egressives, abstracting admissible monomorphisms and epimorphisms [1, §3]. Additive here means that finite products and coproducts form biproducts, and that the commutative monoid structures on mapping spaces are grouplike. The classes contain equivalences and are closed under composition. Pushouts along ingressives and pullbacks along egressives exist and preserve their respective classes. The exactness axiom identifies the ambigressive pushout and pullback squares: in a square with horizontal ingressives and vertical egressives the two universal properties agree, with the corresponding existence statements for such squares. A short exact sequence is the corresponding square with zero in the lower-left

corner.



The classification of extension spaces depends on using these designated diagrams.

Example 7.1. If $\mathcal{A}$ is abelian and all monos and epis are designated, a bicartesian square with one corner zero is exactly a short exact sequence. In a stable category designate all arrows in both classes. Then the same axiom is stability, and a short exact sequence is a cofiber sequence. Its first arrow can have a nonzero fiber, as in Proposition 5.3; the admissibility is a condition of the chosen structure rather than the categorical monomorphism test.

Example 7.2 (The split structure)**.** In an additive category with biproducts the pair $A \to A \oplus B \to B$ is split exact. Pushout of the first map along $f : A \to A'$ is the first inclusion into $A' \oplus B$, with the induced map $(f, 1)$. This follows from the mapping-space test for the pushout. Pullback of the second map along $g : B' \to B$ is the second projection from $A \oplus B'$. Composition of split inclusions and of split projections again has the required biproduct decomposition. These computations underlie the split exact structure in the additive setting. If one works with retracts, the exact structure specifies which complements lie in the category.

7.2 Extension-closed subcategories

Proposition 7.3. *Let D be a stable category and $E \subseteq D$ a full additive subcategory closed under equivalences and extensions, in the sense that $A, B \in E$ and a cofiber sequence $A \to M \to B$ imply $M \in E$. Give E the ingressives whose cofibers belong to E and the egressives whose fibers belong to E. These classes supply the induced exact structure.*

Proof. For an ingressive $A \to M$ with cofiber B, its pushout along $A \to A'$ has cofiber B. The new middle object is an extension of B by A' and lies in E. The dual argument gives the pullback axiom for egressives. For two composable ingressives, (6.2) identifies the cofiber of the composite as an extension of their two cofibers. Hence it belongs to E, and the same holds for composites of egressives by duality. The other required zero and equivalence axioms follow from the stable square definitions. Because E is full, the mapping-space tests for the resulting universal squares restrict from D. Conversely, a required pullback or pushout in E agrees with the constructed one from D by uniqueness, so the ambigressive comparison follows. □

A heart in a t-structure is a full extension-closed subcategory. So is a finite window $D_{[a,b]}$, since membership in both bounds is closed under extensions. The exact structure on the window contains coherent extensions of its objects and usually has nondiscrete mapping spaces. The heart supplies ordinary short exact sequences. This allows bounded pieces of a stable category to have exactness independently of whether they admit all finite limits.

7.3 Stable hulls and exact flags

For a small exact ∞-category E, the stable hull $\mathrm{St}(E)$ comes with an exact functor from E and is universal for exact functors to stable categories:

$$\mathrm{Fun}^{\mathrm{ex}}(\mathrm{St}(E), D) \simeq \mathrm{Fun}^{\mathrm{ex}}(E, D).$$

We use Klemenc's stable-hull theorem [17]. For an ordinary exact category with its given structure the hull identifies with its bounded derived ∞-category. A conceptual construction starts with an additive presheaf embedding, formally stabilizes, and imposes that designated exact sequences become cofiber sequences. The universal property then controls the result. The theorem proves that this construction retains the exact category

in the sense needed to compare extensions. Both the exact structure and the localization in this statement are important.

Calculation 7.4 (A split hull and an abelian hull)**.** Let $R = \mathbb{Z}$. The bounded derived category of finite free modules with the split exact structure is obtained from bounded complexes of finite free modules by chain-homotopy equivalences. Every bounded acyclic complex in this exact structure has contracting homotopy: the split short exact sequences with the cycle objects identify it with a sum of two-term identity complexes. On the other hand the abelian category of all abelian groups uses all short exact sequences. A nonsplit extension such as $0 \to \mathbb{Z} \xrightarrow{2} \mathbb{Z} \to \mathbb{Z}/2 \to 0$ yields a genuine cofiber sequence with $\mathbb{Z}/2$ as a heart object. The complex $[\mathbb{Z} \xrightarrow{2} \mathbb{Z}]$ still represents $\mathbb{Z}/2$ in the bounded derived category of Ab. Its two terms are free, and its stable cofiber construction adds an object that was not a discrete object of the starting exact category of finite free modules.

The flag categories in the S-construction use a diagram on the arrows $i \leq j$ of $[n]$, with objects A_{ij}, subject to $A_{ii} \simeq 0$ and the exactness conditions on $A_{ij} \to A_{ik} \to A_{jk}$ for $i \leq j \leq k$. The objects A_{0i} form a flag. The remaining A_{ij} record its subquotients with their comparisons. In a stable category the subquotients are functorially cofibers, and the compatibility for triples follows from (6.2). For example, in S_3 a flag $A \to B \to C$ with its initial zero has three subquotients associated to this pair of maps. Their cofiber sequence is (6.2). The flag categories and their faces retain the same exactness data simultaneously across all n. Barwick uses the resulting simplicial constructions in the theory of algebraic K-theory [1]. We use the flag description in Chapter 10 to compare extension classes and their concatenation.

7.4 Grothendieck categories

A Grothendieck abelian category is a cocomplete abelian category with exact filtered colimits and a generator. Exact filtered colimits are the axiom AB5. A generator G has faithful $\mathrm{Hom}(G, -)$,

or equivalently every object is a quotient of a coproduct of copies of G. To see this, evaluate the map from the coproduct indexed by all $G \to X$. If its cokernel is nonzero, faithfulness detects the nonzero quotient map, contradicting its vanishing on all summands. The reverse implication follows by precomposing two maps with a chosen quotient presentation. Grothendieck categories have enough injectives [12, 30]. The generator need not be projective.

Examples include Mod_R for an ordinary ring, and sheaves of modules on a small ringed site. For modules, exactness of filtered colimits can be checked on underlying abelian groups by representing any element of a kernel at a finite stage. In sheaves, exact sheafification gives the same conclusion from pointwise exact filtered colimits of presheaves. The free module sheaves on representables supply a family of generators. Taking their coproduct gives a generator, because a transformation is zero if all of its evaluations on sections are zero.

The Gabriel–Popescu theorem gives an exact localization of an ordinary module category [11, 32]. More precisely, for a generator G of a Grothendieck category, let $S = \mathrm{End}(G)$. The functor $\mathrm{Hom}(G, -)$ has values in right S-modules by precomposition, is fully faithful, and has an exact left adjoint. One should retain the direction of exactness: the fully faithful right adjoint can have higher right derived functors. Exactness of the left adjoint yields a localization after passage to the derived categories, using the argument of Theorem 17.1. A compact generator in a stable category gives a different presentation in terms of its endomorphism spectrum, as proved in Chapter 9.

7.5 Localization of complexes

For an abelian category $\mathcal{A}$, a quasi-isomorphism is a map of complexes inducing isomorphisms on all cohomology objects. Its cone is acyclic. The derived ∞-category $\mathrm{D}(\mathcal{A})$ is the localization of complexes at quasi-isomorphisms. The operations of shift and

mapping cone descend to the stable category; for the existence and comparison with classical derived categories see [19, 29, 28]. The bounded derived category is generated under finite extensions and shifts by the heart. This follows from the canonical truncations of a complex, which replace a terminal term by a cycle or a quotient to place its cohomology in a selected range. It also shows why the bounded category of complexes of projective objects and the bounded derived category of all modules can differ: the former requires finite projective models. For $R = k[\epsilon]/(\epsilon^2)$, the residue module k has a free resolution with multiplication by ϵ in every positive degree. Its Ext groups into k are nonzero in arbitrarily high degrees, so k has no bounded projective model.

Exercise 7.5. Let R be a ring and consider the category of finitely generated projective R-modules with its split exact structure. Show that an extension in the abelian category of R-modules of two objects of this category has middle term in the category. Why does a similar assertion fail for the category of free modules of rank one?

Solution. The quotient term is projective, so the surjection has a section and the middle term is the direct sum of the two end terms. A finite direct sum of finitely generated projectives is again finitely generated projective. Even over a field, the extension of two rank-one vector spaces has rank two. The rank-one subcategory therefore lacks extension closure and the biproducts required for the same additive structure. □

Chapter 8

t-structures and Postnikov layers

8.1 Truncations

A t-structure on a stable category C consists of full subcategories $C_{\geq 0}, C_{\leq 0}$, closed under equivalence, with $C_{\geq 0}[1] \subseteq C_{\geq 0}$, $C_{\leq 0}[-1] \subseteq C_{\leq 0}$, and orthogonality

$$\mathrm{Map}_C(U, V) \simeq * \quad (U \in C_{\geq 1},\ V \in C_{\leq 0}).$$

Here $C_{\geq n} = \Sigma^n C_{\geq 0}$ and $C_{\leq n} = \Sigma^n C_{\leq 0}$. Every X has a cofiber sequence $U \to X \to V$ with $U \in C_{\geq 1}$, $V \in C_{\leq 0}$. Applying mapping spaces from $C_{\geq 1}$ and into $C_{\leq 0}$ identifies U, V as universal approximations. In particular there are adjunctions

$$C_{\geq n} \hookleftarrow C, \qquad C \hookleftarrow C_{\leq n}.$$

In the first display the inclusion is left adjoint and $\tau_{\geq n}$ is right adjoint. In the second the inclusion is right adjoint and $\tau_{\leq n}$ is left adjoint. Both adjunctions and the sequence

$$\tau_{\geq n+1} X \longrightarrow X \longrightarrow \tau_{\leq n} X \tag{8.1}$$

are functorial. The t-structure on $\mathrm{D}(\mathcal{A})$ is defined by vanishing of positive or negative cohomology with the convention $\pi_i X = H^{-i}(X)$. On spectra it is defined by vanishing of stable homotopy groups outside the indicated ranges.

Orthogonality also characterizes the two halves. If $\mathrm{Map}(U, X)$ is contractible for all $U \in C_{\geq n+1}$, then $\tau_{\geq n+1} X \simeq 0$, since its identity corresponds to a map into X. Thus $X \in C_{\leq n}$. Dually an object in $C_{\geq n}$ is characterized by vanishing against $C_{\leq n-1}$.

It follows from the mapping-space tests that both halves are closed under extensions. The connective half is closed under colimits that exist in C, while the coconnective half is closed under limits. These statements concern the indicated direction of closure. Inverse limits of connective objects can have lower homotopy groups, as computed in Chapter 13.

Lemma 8.1. *For $a \le b$, $\tau_{\ge a}\tau_{\le b}X \simeq \tau_{\le b}\tau_{\ge a}X$, naturally in X. Either expression belongs to $C_{\ge a} \cap C_{\le b}$, which we denote by $C_{[a,b]}$.*

Proof. If $Y \in C_{\ge a}$, the triangle $\tau_{\ge b+1}Y \to Y \to \tau_{\le b}Y$ has its first two terms in $C_{\ge a}$. This subcategory is closed under cofibers of maps between its objects, by its orthogonality to $C_{\le a-1}$ and closure under suspensions. Thus $\tau_{\le b}Y \in C_{\ge a}$. Dually $\tau_{\ge a}$ preserves $C_{\le b}$.

Set $f : \tau_{\ge a}X \to X$, whose cofiber is $\tau_{\le a-1}X$. Naturality gives a morphism of truncation triangles

$$\begin{array}{ccccc} \tau_{\ge b+1}\tau_{\ge a}X & \longrightarrow & \tau_{\ge a}X & \longrightarrow & \tau_{\le b}\tau_{\ge a}X \\ {\scriptstyle\lambda}\downarrow & & \downarrow{\scriptstyle f} & & \downarrow{\scriptstyle \tau_{\le b}f} \\ \tau_{\ge b+1}X & \longrightarrow & X & \longrightarrow & \tau_{\le b}X. \end{array}$$

The left vertical map λ is an equivalence by the nested right-adjoint universal properties. Indeed, for $U \in C_{\ge b+1} \subseteq C_{\ge a}$,

$$\begin{aligned} \mathrm{Map}_C(U, \tau_{\ge b+1}\tau_{\ge a}X) &\simeq \mathrm{Map}_C(U, \tau_{\ge a}X) \\ &\simeq \mathrm{Map}_C(U, X) \\ &\simeq \mathrm{Map}_C(U, \tau_{\ge b+1}X). \end{aligned}$$

These equivalences identify the map induced by λ, so Yoneda in $C_{\ge b+1}$ applies. The three-by-three cofiber comparison of Chapter 6 now identifies

$$\mathrm{cofib}(\tau_{\le b}f) \simeq \mathrm{cofib}(f) \simeq \tau_{\le a-1}X.$$

Since $\tau_{\le b}\tau_{\ge a}X \in C_{\ge a}$, this is the truncation triangle of $\tau_{\le b}X$ at degree a. Uniqueness gives the natural equivalence in the statement. Equivalently, testing the resulting triangle on $\mathrm{Map}(U, -)$, for $U \in C_{\ge a}$, makes its last term contractible and verifies the same connective-cover universal property. □

8.2 The heart and exactness

The heart is $C^\heartsuit = C_{\geq 0} \cap C_{\leq 0}$. Its mapping spaces are discrete: for A, B in the heart and $i > 0$, $\pi_i \operatorname{Map}(A, B) = \operatorname{Hom}_{\mathrm{h}\, C}(\Sigma^i A, B) = 0$. Translations by maps identify the homotopy groups in all components because the mapping space is an infinite loop space. The heart is an abelian category [19, Prop. 1.2.1.11]. Its kernels and cokernels are the truncations of fibers and cofibers in Chapter 5. For example, a morphism f in the heart is epic precisely when its cofiber lies in $C_{[1,1]}$. It is monic precisely when its cofiber lies in $C_{[0,0]}$. When the kernel and cokernel both vanish the cofiber vanishes.

Proposition 8.2. *A cofiber sequence $A \to E \to B$ with outer terms A, B in the heart yields a short exact sequence there. Conversely a short exact sequence in the heart determines a cofiber sequence in C, with the indicated maps.*

Proof. Extension closure places E in the heart. The long exact sequence of homotopy objects reduces to $0 \to A \to E \to B \to 0$. Conversely, take the cofiber C' of the monomorphism $A \to E$. It lies in $C_{[0,0]}$ and its cokernel in the heart is B. The second map factors through C' by the cokernel property, and gives an equivalence $C' \simeq B$. The composite to be nullhomotopic lies in a discrete mapping space, so the specified nullhomotopy has contractible space of choices when the composite is zero. This gives the coherent sequence. □

A functor $F : C \to D$ between stable t-categories is t-exact if it preserves both halves of the t-structure. For an exact t-exact functor, applying it to (8.1) produces the truncation triangle of FX, by the universal property. Hence F commutes with both truncations and induces an exact functor of hearts by Proposition 8.2. Conversely an exact functor can have nonzero amplitude even when its restriction to an appropriate class of discrete objects resembles an ordinary additive functor. For example $M \mapsto (\mathbb{Z}/2) \otimes^{\mathbf{L}}_{\mathbb{Z}} M$ sends $\mathbb{Z}/2$ to an object with homotopy in degrees $0, 1$. The positive group is $\operatorname{Tor}^{\mathbb{Z}}_1(\mathbb{Z}/2, \mathbb{Z}/2)$.

8.3 Layer maps and reconstruction

Write $\pi_n X = \Sigma^{-n}\tau_{[n,n]}X$, an object of $C^{\heartsuit}$. For X and consecutive coconnective truncations there is a cofiber sequence

$$\Sigma^n \pi_n X \longrightarrow \tau_{\leq n}X \longrightarrow \tau_{\leq n-1}X \xrightarrow{k_n} \Sigma^{n+1}\pi_n X. \tag{8.2}$$

Indeed, apply (8.1) with $n-1$ to $\tau_{\leq n}X$ and use Lemma 8.1. The map k_n is the Postnikov extension class. If the preceding truncation and the new layer are specified, $\tau_{\leq n}X$ is recovered as the fiber of k_n. In particular specifying every $\pi_n X$ alone generally does not specify X.

Example 8.3. For $R = k[\epsilon]/(\epsilon^2)$, a free resolution of k has one copy of R in each nonnegative homological degree and all differentials multiplication by ϵ. Exactness follows because the kernel and image of multiplication by ϵ both equal (ϵ). The complex of maps into k has zero differentials, so $\mathrm{Ext}^2_R(k,k) \cong k$. Let $u : k \to k[2]$ represent a nonzero class. Its fiber X lies in $C_{[0,1]}$ with $\pi_0 X = \pi_1 X = k$. The split object $k \oplus k[1]$ has the same homotopy objects. In (8.2) for $n = 1$, the map of X is u while that of the split object is zero. Naturality of truncation forces any equivalence of these objects to identify their layer maps up to automorphisms of the layers, which cannot carry a nonzero class to zero.

We call C left complete for this t-structure when every canonical map $X \to \lim_{n\to+\infty} \tau_{\leq n}X$ is an equivalence, and right complete when every $\mathrm{colim}_{m\to+\infty} \tau_{\geq -m}X \to X$ is an equivalence. The definitions include existence of the indicated limits and colimits. The morphisms of the first tower lower the upper bound and those of the second filtration lower the lower bound. A bounded object satisfies both reconstruction statements because its diagrams become constant. For a general object, the statements are properties of the t-structure and of the limit and colimit operations. Neeman gives Grothendieck abelian examples with failure of left completeness of the derived category [22]. The heart and its filtered colimits are therefore insufficient to infer this property.

8.4 Separatedness and generators of layers

Left separatedness means $\bigcap_n C_{\geq n} = 0$; right separatedness means $\bigcap_n C_{\leq n} = 0$. Left completeness implies the first: if all coconnective truncations are zero, the limit is zero. Right completeness implies the second by the dual argument. If both separatedness conditions hold, the π_n detect equivalences. To see this, apply them to the cofiber of a map. If $\pi_n X = 0$ for every n, induction on length using (8.2) gives $\tau_{[a,b]} X = 0$ for all finite windows. For fixed b, this says all connective truncations of $\tau_{\leq b} X$ are zero, so right separatedness forces $\tau_{\leq b} X = 0$. Left separatedness then gives $X = 0$. The proof separates the detection of zero from construction of a limit.

Exercise 8.4. Let $u : A \to B[n]$ for heart objects A, B and $n \geq 2$. Locate the homotopy objects of $\operatorname{fib} u$ and identify its nontrivial Postnikov boundary.

Solution. Rotate the fiber sequence to $B[n-1] \to \operatorname{fib} u \to A \xrightarrow{u} B[n]$. The long exact sequence gives $\pi_0 = A$ and $\pi_{n-1} = B$, with the intervening and remaining layers zero. Its truncation in degree $n-2$ is A. Hence the boundary in (8.2) for $n-1$ is the specified u. This example explains how an n-extension can appear as a two-layer object. □

Chapter 9

Presentability and compact generators

9.1 Small generators in large categories

An ∞-category is presentable when it has small colimits and is accessible. Accessibility means that for some regular cardinal κ there is an essentially small full subcategory of κ-compact objects generating the category under κ-filtered colimits. An object K is κ-compact if $\mathrm{Map}(K, -)$ preserves such colimits. In particular compact without a cardinal modifier means ω-compact. Another characterization of presentable categories is as accessible localizations of presheaf categories [18, §5.5]. The presheaf category is generated by the compact representables of Chapter 2. A localization can change compactness because the right adjoint identifying local objects may fail to preserve the filtered colimits against which compactness is tested.

For a stable category, a family $\{G_i\}$ detects equivalences if the family of mapping spectra $\mathrm{map}(G_i, -)$ does. It suffices to detect zero objects by applying the spectra to a cofiber. In terms of the homotopy category, this requires vanishing of $\mathrm{Hom}(\Sigma^n G_i, X)$ for all i, n to imply $X \simeq 0$. If C is presentable and the G_i are compact, this detection condition is equivalent to generation under shifts, cofibers, and small coproducts. Indeed, the localizing subcategory they generate has a right adjoint to its inclusion, by the presentable localization theorem. The cofiber of its counit at X is right orthogonal to the generators and hence zero. Conversely the class of sources T with $\mathrm{map}(T, X) \simeq 0$ is closed under those

constructions; if it includes every object then id_X vanishes.

The relationship to the classical generator involves all the shifts. A projective generator in a module category tests exact sequences in the heart and can be the single compact generator of its derived category. A generator of a Grothendieck category can require resolutions to make the same test. Its derived mapping functor may contain higher Ext groups even when the classical Hom functor is faithful. A derived compact generator also need not belong to the heart. Taking a shifted generator changes its placement but leaves detection by all its shifts unchanged.

9.2 Morita reconstruction

Theorem 9.1 (Spectral Morita, single generator)**.** *Let C be a presentable stable ∞-category with a compact generator G. If $E = \mathrm{map}_C(G, G)$ has multiplication by composition, the functor $\mathrm{map}_C(G, -)$ is an equivalence from C to right E-modules.*

Proof. Precomposition gives the right E-action. The left E-object G defines a left adjoint $M \mapsto M \otimes_E G$ by the tensor product of the right module and the left object in C. It can be constructed by realizing the two-sided bar construction in C. The mapping-spectrum adjunction follows by testing the bar construction and using its balancing relations. The right adjoint is exact. Compactness of G implies that its mapping spectrum preserves filtered colimits, since this can be checked on all levels using compactness of shifts of G. A coproduct is a filtered colimit of finite coproducts, so exactness and compactness give preservation of small coproducts. In stable categories an exact functor preserving coproducts preserves small colimits. This follows by forming simplicial replacements from coproducts, then realizing them by their sequential skeletal filtrations of finite colimit attachments. Sequential colimits of spectra or objects of a stable category are cofibers of the difference between the identity and the shift on the coproduct. The same constructions give preservation of the colimits in question.

The unit of the adjunction is an equivalence on the free right module E, since $E \otimes_E G \simeq G$. The class on which the unit is an equivalence is closed under small colimits and shifts, since both functors preserve them. The free module generates the module category under these operations: the underlying spectra are generated by spheres, and the free–forgetful bar construction resolves any module by free modules on spectra. Thus the unit is an equivalence everywhere. The right adjoint is conservative by the detection hypothesis on G. Applying it to the counit and using the triangle identities now gives an equivalence, so the counit is an equivalence as well. □

This is the stable ∞-categorical form of a standard Morita theorem [26, 19]. The spectrum E, including its multiplication, is necessary in the conclusion. Taking only $\pi_0 E$ would forget morphisms between shifts of G. For $C = \mathrm{Sp}$ and $G = \mathbb{S}$, $E = \mathbb{S}$ while $\pi_0 E = \mathbb{Z}$. The resulting ordinary ring models the heart of spectra, but the module category over $\mathrm{H}\mathbb{Z}$ differs from spectra. For example $\pi_1 \mathbb{S} \cong \mathbb{Z}/2$, whereas $\pi_1 \mathrm{H}\mathbb{Z} = 0$.

Calculation 9.2 (The arrow category)**.** For a commutative ring R, set $C = \mathrm{Fun}([1], \mathrm{D}(R))$. Its objects are coherent arrows. The left adjoints to evaluation at $0, 1$ send R to

$$P_0 = (R \xrightarrow{1} R), \qquad P_1 = (0 \to R).$$

They are compact, and the spectra of maps out of them are the underlying derived modules at $0, 1$, respectively. Thus $G = P_0 \oplus P_1$ is a compact generator. Its endomorphism spectrum is discrete because each $\mathrm{map}(P_i, P_j)$ evaluates to R or zero in degree zero. In the order (P_0, P_1) the endomorphism ring is

$$\begin{pmatrix} R & R \\ 0 & R \end{pmatrix}.$$

Here row j, column i is $\mathrm{map}(P_i, P_j)$, so the column indexes the source and the row the target. Indeed $\mathrm{map}(P_0, P_1) \simeq P_1(0) = 0$, while $\mathrm{map}(P_1, P_0) \simeq P_0(1) = R$. The Morita equivalence gives

the arrow category as derived right modules over this triangular ring. A right module decomposes using the two diagonal idempotents. Right multiplication by the upper-right matrix unit goes from the first component to the second, which recovers the displayed orientation of [1].

9.3 Products, generators, and completeness

A useful sufficient condition for a Grothendieck category $\mathcal{A}$ to have a left complete derived category is exactness of countable products. For complexes, termwise products then preserve quasi-isomorphisms. They compute products in the derived category by deriving the product adjunction against the exact diagonal functor. Consequently each H^q commutes with those products. The inverse limit of a sequence of complexes is the fiber of $1 - \text{shift}$ on the derived product, as proved in Chapter 13. For a tower of coconnective truncations of X, each π_q is eventually constant. On its product $1 - \text{shift}$ is surjective and has kernel the stabilized value. Applying the long exact sequence to the fiber shows $X \to \lim_n \tau_{\le n} X$ induces isomorphisms on every π_q. In $\mathrm{D}(\mathcal{A})$ the groups detect equivalences by the definition of quasi-isomorphism. Filtered colimits are exact by AB5, so filtered colimits of complexes compute the derived ones. On the connective reconstruction sequence each π_q also stabilizes, proving right completeness in this case.

Module categories have exact products as well as exact filtered colimits. Product exactness uses the fact that an element of a product of targets of surjections has a lift in every coordinate and hence in the product. For sheaves, epimorphisms are locally surjective. Infinitely many local lifting conditions can require infinitely many neighborhoods whose intersection is not open. Thus the element argument for products of modules does not immediately prove exactness of sheaf products. This is one source of the care needed with unbounded derived sheaves.

Example 9.3 (A product of stable categories)**.** Let $C = \prod_{i \in I} \mathrm{D}(R)$ for an infinite discrete set I. The objects R_i supported at one coordinate form a compact generating family. A compact object X has finite support: write X as the filtered colimit of its restrictions to finite subsets, extended by zero. Compactness factors the identity of X through one such restriction in the homotopy category. All components outside that finite subset are therefore zero. Every component of X is compact in $\mathrm{D}(R)$, since testing against diagrams supported at the component reduces compactness to that category. Conversely an object of finite support with compact components is compact, using finite products of mapping spaces and their commutation with filtered colimits. It follows that this category has no single compact generator if $R \neq 0$, although its heart has an ordinary generator with nonzero component at every coordinate.

9.4 Perfect objects

For a ring spectrum E, a perfect module is in the idempotent-complete stable subcategory generated by E using finite colimits and shifts. Every perfect module is compact. Conversely a compact module is perfect: the full subcategory of perfect modules is essentially small in the chosen universe and its Ind-completion maps to Mod_E by the colimit-preserving extension. On perfect modules and their filtered colimits this extension is fully faithful by compactness. Its essential image is closed under finite colimits and filtered colimits, and hence under arbitrary small colimits, since coproducts and realizations can be built from those operations. It contains E and therefore all modules by the generation used in Theorem 9.1. A compact object of an Ind-completion is a retract of an object at a stage of its filtered presentation. Since the subcategory is already idempotent complete, this proves the characterization. For an ordinary ring, perfect complexes are precisely the objects represented by bounded complexes of finitely generated projectives. This finite description is the source of a finite base-change calculation in Chapter 23.

Exercise 9.4. Compute a compact generator of $\mathrm{Fun}([2], \mathrm{D}(R))$ and the orientation of its endomorphism ring for commutative R.

Solution. Take $G = P_0 \oplus P_1 \oplus P_2$, where P_i is R with identity maps on indices at least i and zero below i. The evaluation adjunction shows $\mathrm{map}(P_i, P_j) \simeq P_j(i)$, equal to R if $j \leq i$ and zero otherwise. In the matrix of endomorphisms row j, column i therefore permits R if $j \leq i$. Composition multiplies entries along the unique chain of arrows. This is the upper triangular 3×3 matrix ring over R. □

Part III

Extensions and derived functors

A boundary map classifies a short stable extension. Its suspensions give the extension spectrum, whose groups recover the classical Yoneda operations. The additive coend for satellites and the derived localization of complexes each express a universal extension of a functor. The final two chapters use towers to identify which bounded Postnikov pieces determine maps and which layers can affect a fixed range of derived output.

Chapter 10

Extension spaces and extension spectra

10.1 The universal extension space

Fix objects A, B in a stable ∞-category C. Let $\mathrm{Seq}_C(B, A)$ be the space of cofiber sequences $A \to E \to B$ with specified identifications of their outer terms with A, B. The space keeps equivalences between such sequences respecting the identifications and all their higher comparisons. Equivalently it is the fiber of the outer-term functor on the cores over (A, B).

Theorem 10.1. *The boundary of a cofiber sequence gives an equivalence*

$$\mathrm{Seq}_C(B, A) \simeq \mathrm{Map}_C(B, \Sigma A).$$

It is functorial in A by pushout and contravariantly functorial in B by pullback.

Proof. For a map $\delta : B \to \Sigma A$ take its fiber $E \to B$. The fiber of $E \to B$ is canonically $\Omega\Sigma A \simeq A$, by the pullback square defining E and pasting. Thus $A \to E \to B$ is a cofiber sequence, with boundary the given δ. Starting with a cofiber sequence, the same pasting identifies its middle term with the fiber of the boundary. These constructions are functorial operations on diagrams with their comparison homotopies, since limits and colimits are functorial when the shapes and cone data are fixed. The two equivalences are compatible with all maps of diagrams, yielding inverse equivalences of spaces. Under a map $A \to A'$, the pushout of the first edge keeps cofiber B; the map of cofiber sequences acts on the boundary by postcomposition with $\Sigma A \to$

$\Sigma A'$. A pullback along $B' \to B$ acts by precomposition, using the equivalent fiber presentation. This proves both naturalities. □

For any stable C and integer n, put

$$E^n_C(B, A) = \mathrm{Map}_C(B, \Sigma^n A).$$

The maps $E^n_C \simeq \Omega E^{n+1}_C$, based at the zero arrow, exhibit these spaces as the spaces of the mapping spectrum, with the nth space indexed by $n \geq 0$. In particular

$$\pi_i(E^n_C(B, A), 0) \cong \mathrm{Hom}_{\mathrm{h}\, C}(B, \Sigma^{n-i} A) \qquad (i \geq 0). \tag{10.1}$$

For $i > 0$ and a point other than zero, translation in the infinite loop space gives an identification with the same group. When A, B lie in a heart, all the groups with $i > n$ vanish by orthogonality. Thus the higher extension space in this case is an n-type. The ordinary Yoneda extension group will be its set of components when the ambient category is the bounded derived category of the heart.

10.2 Classical extensions and their automorphisms

A short exact sequence $0 \to A \xrightarrow{i} E \xrightarrow{p} B \to 0$ in an abelian category has automorphisms fixing A, B canonically identified with $\mathrm{Hom}(B, A)$. If v is such an automorphism, $v - \mathrm{id}_E$ vanishes on A and has image in $\ker p = A$, so $v - \mathrm{id}_E = ihp$ for a unique $h : B \to A$. Conversely $\mathrm{id}_E + ihp$ has inverse $\mathrm{id}_E - ihp$ since $pi = 0$, and multiplication corresponds to addition of h. This is the ordinary formula for the fundamental group in (10.1) at $n = 1$. The equivalence of this extension groupoid with the stable sequence space follows from Proposition 8.2 and the fact that the mapping spaces in the heart are discrete.

The sum of two extensions of B by A is formed by taking their direct sum, pulling back along the diagonal $B \to B \oplus B$, and

pushing out along the codiagonal $A \oplus A \to A$. By functoriality of Theorem 10.1, the corresponding boundary is

$$B \xrightarrow{\Delta} B \oplus B \xrightarrow{\delta_1 \oplus \delta_2} \Sigma A \oplus \Sigma A \xrightarrow{\nabla} \Sigma A.$$

This is addition in the mapping spectrum. In particular the Baer operation on components is induced by a specified operation on the extension space. The split sequence is its zero point, and its automorphism group is a loop group at that point.

Calculation 10.2. For $A = B = \mathbb{Z}/p$ in $\mathrm{D}(\mathbb{Z})$, p prime, the two-term projective resolution with differential p gives $\mathrm{Hom}(A, B) \cong \mathrm{Ext}^1_{\mathbb{Z}}(A, B) \cong \mathbb{F}_p$ and $\mathrm{Ext}^{n \geq 2} = 0$. The short extension space has p connected components, each equivalent to $B(\mathbb{F}_p)$. The extension $0 \to \mathbb{Z}/p \to \mathbb{Z}/p^2 \to \mathbb{Z}/p \to 0$ is a nonzero component. The 2-extension space is connected with $\pi_1 = \pi_2 = \mathbb{F}_p$ and no further groups. In this example the two groups split at the level of the complex of maps computed by the resolution, whose differential is zero, so the space is $K(\mathbb{F}_p, 1) \times K(\mathbb{F}_p, 2)$. The vanishing of the group of classes of 2-extensions leaves a nontrivial space of comparisons.

10.3 Iterated extensions

In an abelian category an n-extension of B by A is an exact sequence

$$0 \to A \to E_{n-1} \to \cdots \to E_0 \to B \to 0.$$

Maps between such sequences fix the outer terms. For $n \geq 2$ the arrows can fail to be isomorphisms. The classical equivalence relation is generated by these arrows. The associated extension space is the groupoid completion of the extension category, so its components are the equivalence classes for exactly this relation. Taking merely the core of this category would retain only isomorphisms of the entire sequence, and in general produces too many components.

To compare a representative with the derived category, write it as a splice of short exact sequences by factoring through its intermediate images. Each short sequence gives a boundary of degree one. Compose these with appropriate suspensions to obtain $B \to \Sigma^n A$. For an extension of B by A of degree n and an extension of A by D of degree m, the boundary of the splice is

$$B \longrightarrow \Sigma^n A \longrightarrow \Sigma^{n+m} D.$$

The second map is the n-fold suspension of the boundary of the second extension. Fixing this convention specifies the order of Yoneda products. Maps of sequences preserve the intermediate image factorizations in the abelian case, hence preserve the boundary up to the natural homotopy. The same construction with compatible flags gives the comparison of higher extension spaces.

The space-level theorem of Retakh for ordinary exact categories, and its ∞-categorical version of Børve and Trygsland [24, 2], identifies extension spaces in an exact ∞-category with those in its stable hull. In the latter, the n-extension space is $\mathrm{Map}(B, \Sigma^n A)$. In particular, for an ordinary abelian $\mathcal{A}$, the components recover $\mathrm{Ext}^n_{\mathcal{A}}(B, A)$ in the Yoneda sense from $\mathrm{D}^b(\mathcal{A})$. The theorem uses the exact structure in constructing higher extensions. One cannot apply it to a category with arbitrary flags for which the comparison squares have not been declared exact. In a stable category the theorem includes the loop equivalences of the mapping spectrum; in an abelian category it relates them to the homotopy of an ordinary category of sequences.

10.4 Boundary actions

Let $A \to E \to B \xrightarrow{\delta} \Sigma A$ be an extension. Applying $\mathrm{map}(T, -)$ gives a fiber sequence and hence a long exact sequence. The boundary

$$\mathrm{Hom}(T, \Sigma^n B) \longrightarrow \mathrm{Hom}(T, \Sigma^{n+1} A)$$

sends α to $(\Sigma^n \delta)\alpha$. Under the comparison with classical extensions this is splicing with the fixed short extension. Applying $\mathrm{map}(-, T)$ to the same sequence gives the contravariant long exact sequence. Its boundary is composition on the other side, with the sign fixed by the choice of rotation. In particular the two long exact sequences in classical Ext arise from exactness of one bifunctor in each variable.

Calculation 10.3 (A nonzero splice)**.** Let $R = k[\epsilon]/(\epsilon^2)$ and write e for the class of

$$0 \longrightarrow k \xrightarrow{1 \mapsto \epsilon} R \longrightarrow k \longrightarrow 0.$$

This class is nonzero because a section of the quotient would send $1 \in k$ to an element annihilated by ϵ and congruent to 1 modulo ϵ, which is impossible. Apply $\mathrm{Ext}^*_R(-, k)$ to the sequence. For $q \geq 1$, the two groups $\mathrm{Ext}^q_R(R, k)$ and $\mathrm{Ext}^{q+1}_R(R, k)$ vanish. The intervening boundary is therefore an isomorphism $\mathrm{Ext}^q_R(k, k) \to \mathrm{Ext}^{q+1}_R(k, k)$. By the product description it is splicing with e, up to the sign of the boundary convention. In particular $e^2 \neq 0$. Iterating proves that every $\mathrm{Ext}^q_R(k, k)$ is generated by e^q as a k-vector space. The class e^2 is a concrete choice for the map u of Example 8.3.

Exercise 10.4. Let k be a field and A, B vector spaces regarded as heart objects in $\mathrm{D}(k)$. Identify $E^n_{\mathrm{D}(k)}(B, A)$ for $n \geq 0$. How does its answer reconcile the vanishing of positive Ext with nontrivial spaces of equivalences?

Solution. Every vector space is projective and $\mathrm{RHom}_k(B, A)$ is the Eilenberg–Mac Lane spectrum of $\mathrm{Hom}_k(B, A)$. Thus $E^n \simeq K(\mathrm{Hom}_k(B, A), n)$. For $n > 0$ it has one component and can have nonzero π_n. In particular a split short sequence has automorphisms even though every short sequence has the same extension class. For higher n, the group appears as a higher homotopy group of the space of sequences. □

Chapter 11

Satellites as enriched Kan extensions

11.1 The graded extension category

Strooker's construction of satellites starts with the category built from Yoneda extensions and extends an additive functor along its degree-zero inclusion [38, §3.2]. A graded formulation makes the weighting objects explicit. Let $\mathcal{A}$ be an abelian category. Define a preadditive category $\mathcal{G}$ with objects (A, r), $A \in \mathcal{A}$, $r \in \mathbb{Z}$, and with morphism groups

$$\mathcal{G}((A,r),(B,s)) = \begin{cases} \operatorname{Ext}^{s-r}_{\mathcal{A}}(A,B), & s \geq r, \\ 0, & s < r. \end{cases}$$

Composition in nonnegative degree is the Yoneda product, with degree-zero composition given by pushout, pullback, and ordinary composition. Identities are the usual degree-zero identities. Associativity follows from splicing, or from composition of shifted maps in the bounded derived category. Let $i : \mathcal{A} \to \mathcal{G}$ be the full degree-zero inclusion.

For an additive $T : \mathcal{A} \to \mathrm{Ab}$ the enriched left Kan extension is

$$(\operatorname{Lan}_i^{\mathrm{Ab}} T)(B,n) = \int^{A \in \mathcal{A}} \operatorname{Ext}^n_{\mathcal{A}}(A,B) \otimes_{\mathbb{Z}} T(A) \quad (n \geq 0). \tag{11.1}$$

The expression in degree $n < 0$ is zero. Here and below the coend is formed in the larger universe when necessary. To verify the formula, the endomorphism and composition actions in $\mathcal{G}$ on its hom groups act on these tensor quotients. A transformation from this functor to $U : \mathcal{G} \to \mathrm{Ab}$ is determined by its values on $[\mathrm{id}_A \otimes x]$ at $(A, 0)$. Given an additive transformation $T \to Ui$,

the family of maps on tensors is $[\eta \otimes x] \mapsto U(\eta)(x')$, where x' is the image of x. It respects the relations in calculation 2.5, giving the inverse to restriction. This is the enriched Kan adjunction. Degree zero is T by the additive co-Yoneda formula.

A short exact sequence $A \to E \to B$ gives a degree-one arrow $(B, n) \to (A, n+1)$ in $\mathcal{G}$. The Kan extension carries it to the connecting map given on a tensor by composing its extension class with the weight. This construction imposes functoriality under diagrams of short exact sequences. When the hypotheses for the next theorem hold, the connecting maps also give the classical long exact sequences.

11.2 The comparison with resolutions

Theorem 11.1 (Satellites and right derivation)**.** *Let $\mathcal{A}$ have enough injectives and $T : \mathcal{A} \to \mathrm{Ab}$ be additive and left exact. For $n \geq 0$ the natural pairing obtained by applying $\mathrm{R}T$ to the Yoneda extension gives an isomorphism*

$$\int^{A \in \mathcal{A}} \mathrm{Ext}^n_{\mathcal{A}}(A, B) \otimes_{\mathbb{Z}} T(A) \xrightarrow{\cong} R^n T(B).$$

It takes the action of each short extension in (11.1) to the corresponding connecting morphism of right derived functors.

Proof. A class $\eta : A \to B[n]$ in the derived category induces $\mathrm{R}T(A) \to \mathrm{R}T(B)[n]$. Applying H^0 gives $T(A) \to R^n T(B)$. This uses left exactness to identify $H^0 \mathrm{R}T(A) = T(A)$. The pairing is bilinear and balanced because $\mathrm{R}T$ is functorial and additive, so it defines the displayed map. Composition of extension maps gives compatibility with boundaries.

Write $M_n(B) = \mathrm{Ext}^n_{\mathcal{A}}(-, B)$, a right $\mathcal{A}$-module. Choose a monomorphism $B \to I$ with I injective and let C be its cokernel. At $n = 1$, the long exact Ext sequence gives a pointwise right exact sequence

$$\mathrm{Hom}_{\mathcal{A}}(-, I) \longrightarrow \mathrm{Hom}_{\mathcal{A}}(-, C) \longrightarrow M_1(B) \longrightarrow 0.$$

The coend functor $M \mapsto M \otimes_{\mathcal{A}} T$ preserves cokernels, since it is a colimit of bilinear tensor constructions or equivalently a left adjoint. It converts the first two terms by co-Yoneda to $T(I)$ and $T(C)$. The resulting cokernel is $R^1T(B)$ by the derived sequence and the vanishing of $R^1T(I)$. The map is the specified one because the surjection from $\mathrm{Hom}_{\mathcal{A}}(-, C)$ is splicing with the fixed short extension, and the derived sequence gives the same boundary. For $n > 1$, the boundary identifies $M_{n-1}(C) \simeq M_n(B)$ pointwise. The corresponding derived boundary identifies $R^{n-1}T(C) \simeq R^nT(B)$. Induction along these identifications proves the result in all degrees. For $n = 0$ the claim is co-Yoneda. The diagram of comparisons for any other short extension commutes by the pairing, so the result is natural and independent of the choice of I. □

This proof uses effacement in the second variable at exactly one place: $\mathrm{Ext}^n_{\mathcal{A}}(A, I) = 0$ for $n > 0$. No projectivity of A is needed. Once the coend is computed in the larger universe, the isomorphism with $R^nT(B)$ gives a value in the original universe. The universal extension by a Kan formula can therefore be used without a small skeleton of all objects of $\mathcal{A}$.

Corollary 11.2. *For $T = \mathrm{Hom}_{\mathcal{A}}(X, -)$, the natural isomorphism is the usual evaluation of the Yoneda product:*

$$\int^A \mathrm{Ext}^n_{\mathcal{A}}(A, B) \otimes_{\mathbb{Z}} \mathrm{Hom}_{\mathcal{A}}(X, A) \cong \mathrm{Ext}^n_{\mathcal{A}}(X, B).$$

Proof. This is also right co-Yoneda. The tensor $\eta \otimes h$ maps to $h^*\eta$, and the inverse sends γ to $\gamma \otimes \mathrm{id}_X$. The tensor relation for h gives the inverse identity on generators. □

11.3 The dual construction

For an additive $T : \mathcal{A} \to \mathrm{Ab}$ and $n \geq 0$, the enriched right Kan extension at $(B, -n)$ is

$$\mathrm{Nat}_{\mathrm{add}}\Big(\mathrm{Ext}^n_{\mathcal{A}}(B, -), T\Big).$$

A map $U \to \mathrm{Ran}_i^{\mathrm{Ab}} T$ is obtained from $Ui \to T$ by applying U to arrows $(B, -n) \to (A, 0)$ and then evaluating the given map. Conversely evaluating at identities gives the restriction. This is the dual universal property to (11.1). It has value $T(B)$ at $n = 0$ by Yoneda.

Proposition 11.3. *If $\mathcal{A}$ has enough projectives and T is right exact, the right Kan value at $(B, -n)$ identifies naturally with $L_n T(B)$.*

Proof. For a short exact sequence $K \to P \to B$ with P projective there is a pointwise right exact sequence

$$\mathrm{Hom}_{\mathcal{A}}(P, -) \to \mathrm{Hom}_{\mathcal{A}}(K, -) \to \mathrm{Ext}^1_{\mathcal{A}}(B, -) \to 0.$$

The functor of transformations into T takes this sequence to a kernel description. By Yoneda the kernel is that of $T(K) \to T(P)$, which is $L_1 T(B)$ by the long exact sequence. For $n > 1$, $\mathrm{Ext}^n_{\mathcal{A}}(B, -) \cong \mathrm{Ext}^{n-1}_{\mathcal{A}}(K, -)$ and $L_n T(B) \cong L_{n-1} T(K)$. The natural comparison uses the pairing with $\mathrm{L}T$ applied to $B \to A[n]$, then identifies $H^{-n}\mathrm{L}T(B) \to H^0 \mathrm{L}T(A)$. Thus the dimension shifting maps agree and induction proves the claim. At degree zero use right exactness and Yoneda. □

11.4 An enriched lift and its relation to truncation

Let $i : \mathcal{B} \to C$ now be a fully faithful inclusion of a small spectral category of objects into a stable category, and let $F : \mathcal{B} \to D$ be a spectral functor into a presentable stable category. In this case the pointwise spectral left Kan extension is

$$X \longmapsto \int^{B \in \mathcal{B}} \mathrm{map}_C(iB, X) \otimes F(B). \tag{11.2}$$

Here $\otimes$ is the tensor of D over spectra and the coend is spectral. It can be built by the bar construction whose degree k takes the

coproduct over $B_0, \dots, B_k$ of the tensors of $F(B_0)$, the successive mapping spectra, and the terminal weight. Faces compose mapping spectra or act on F, and degeneracies insert units. A map out of its realization is a coherent balanced pairing. Taking the representable weight at iB gives $F(B)$ by enriched Yoneda, proved by the extra degeneracy supplied by id_B. This proves the pointwise Kan property in the same manner as the abelian quotient formula.

The formulas (11.1) and (11.2) impose different colimit operations. Taking homotopy groups of a spectral coend generally produces a spectral sequence with higher homological terms. For example the derived balanced tensor of $H(\mathbb{Z}/p)$ with itself over $H\mathbb{Z}$ has $\pi_1 = \mathbb{Z}/p$, whereas the ordinary balanced tensor of the two discrete modules is $\mathbb{Z}/p$ in degree zero. Thus the enriched Kan computation of Theorem 11.1 uses an ordinary additive coend with an already graded extension weight. It is not obtained by simply applying π_0 termwise to (11.2).

Calculation 11.4. Take $T = \mathrm{Hom}_{\mathbb{Z}}(\mathbb{Z}/p, -)$. The resolution $\mathbb{Z} \xrightarrow{p} \mathbb{Z} \to \mathbb{Z}/p$ gives

$$R^0T(B) = B[p], \qquad R^1T(B) = B/pB, \qquad R^{q\geq 2}T(B) = 0.$$

For $0 \to \mathbb{Z} \xrightarrow{p} \mathbb{Z} \to \mathbb{Z}/p \to 0$, the boundary $T(\mathbb{Z}/p) \to R^1T(\mathbb{Z})$ is an isomorphism: both adjacent degree-zero groups on $\mathbb{Z}$ vanish and the map on R^1 induced by p is zero. The Kan extension has exactly this boundary. Its weight sends $\mathrm{id}_{\mathbb{Z}/p}$ to the class of this short exact sequence, which is the generator of $\mathrm{Ext}^1_{\mathbb{Z}}(\mathbb{Z}/p, \mathbb{Z})$. The computation identifies the sign-free generator convention by taking the chosen injection to be p.

Exercise 11.5. Let T be exact. Prove directly from (11.1), assuming enough injectives, that its positive right satellites vanish.

Solution. The first coend is $\mathrm{coker}(T(I) \to T(C))$ for $B \to I \to C$, and the map is epic by exactness. All higher coends reduce to this case by the isomorphisms $M_n(B) \simeq M_{n-1}(C)$. The degree-zero coend is T itself. □

Chapter 12

Derived functors and their models

12.1 Resolution and localization

Let $\mathcal{A}, \mathcal{B}$ be Grothendieck abelian categories and $T : \mathcal{A} \to \mathcal{B}$ be additive. A complex I is K-injective when the complex of maps from each acyclic complex into I is acyclic. This is equivalent to vanishing of all homotopy classes from acyclic complexes into shifts of I. The criterion implies that maps into I carry quasi-isomorphisms in the source to quasi-isomorphisms of mapping complexes: apply the complex of maps to the cone of a quasi-isomorphism. Every complex of a Grothendieck category has a quasi-isomorphism into a K-injective complex, and these can be chosen with injective terms [28, 31]. The bounded-below case uses ordinary bounded-below complexes of injectives. Their K-injectivity follows by inductively solving the chain-homotopy equation into the injective terms, starting at the lowest degree.

These statements give a model for the derived ∞-category. First use the stable category of complexes with their mapping complexes and chain homotopies. Its homotopy category is the homotopy category of complexes. The full subcategory of K-injective objects is reflective in this model: given $X \to I_X$ as above, mapping it into every K-injective object induces equivalences on mapping spaces by the defining test. The unit exhibits the reflector and gives the localization at quasi-isomorphisms. Hence a map in $\mathrm{D}(\mathcal{A})$ can be calculated into a K-injective representative, with the whole mapping complex and not just its group of components.

Write q for the localization and i for the inclusion of these local

objects. Termwise T is a functor on complexes with homotopies because it is additive; it respects the complex of maps, including its differential. The right derived functor on the ∞-category is represented by $I \mapsto T(I)$ on the local objects, followed by localization in $\mathcal{B}$. This is an exact functor $\mathrm{R}T : \mathrm{D}(\mathcal{A}) \to \mathrm{D}(\mathcal{B})$: local complexes are closed under shifts and cones, and T preserves the cone formula with its specified differentials. On $A \in \mathcal{A}$ viewed as a complex in degree zero it yields $T(I_A^\bullet)$ for an injective resolution.

Proposition 12.1 (The localization comparison)**.** *Let $q : C \to D$ be a reflective localization with fully faithful right adjoint i, and let $F : C \to E$. The unit gives a transformation $F \to Fiq$, exhibiting Fi as the left Kan extension of F along q.*

Proof. For $G : D \to E$, a map $F \to Gq$ restricts along i to $Fi \to Gqi \simeq G$. Conversely a map $Fi \to G$ composes with $F \to Fiq$ to give $F \to Gq$. Naturality for the unit and the triangle identities identify the two composites with the identities on transformation spaces. This is the defining Kan property, and the same argument applies after any postcomposition. □

Apply the proposition with F the termwise functor followed by localization in $\mathcal{B}$. The direction of the comparison $T \to (\mathrm{R}T)q$ explains why the right derived functor in this convention is a left Kan extension along the localization. For a coreflective localization using K-projective models, the arrow of comparison reverses and the corresponding left derived functor is a right Kan extension. The word right in right derived refers to the resolution and the degrees of the classical derived groups, whereas left in the displayed Kan property refers to the direction of the transformation.

12.2 Tensor and Hom

Let R be an ordinary ring. A complex of right R-modules is K-flat if tensoring it with an acyclic complex of left modules produces an

acyclic complex. Replacing one factor by such a resolution gives $M \otimes^{\mathbf{L}}_{R} N$. For bounded-above complexes one can use bounded-above complexes of flat modules, with totalization by sums. For unbounded complexes one uses K-flat resolutions rather than arbitrary termwise flat replacements [28, 29]. Derived Hom can be computed by a K-projective source or a K-injective target. These choices implement the derived tensor–Hom adjunction. Under $\mathrm{D}(R) \simeq \mathrm{Mod}_{HR}$ they are the tensor and internal map operations of HR-modules, with the appropriate ring actions or bimodule structure [26, 27].

Calculation 12.2. The complex $P = [\mathbb{Z} \xrightarrow{p} \mathbb{Z}]$, in cochain degrees $-1, 0$, resolves $\mathbb{Z}/p$. Tensoring P with $\mathbb{Z}/p$ gives zero differential. Thus

$$(\mathbb{Z}/p) \otimes^{\mathbf{L}}_{\mathbb{Z}} (\mathbb{Z}/p) \simeq (\mathbb{Z}/p) \oplus (\mathbb{Z}/p)[1].$$

On the other hand mapping from P into $\mathbb{Z}/p$ gives a complex in degrees $0, 1$, also with zero differential. Thus

$$\mathrm{RHom}_{\mathbb{Z}}(\mathbb{Z}/p, \mathbb{Z}/p) \simeq (\mathbb{Z}/p) \oplus (\mathbb{Z}/p)[-1].$$

The homotopy shift is opposite in the two formulas. In homological grading Tor_1 is the positive homotopy group of the first, and Ext^1 is the negative group of the second. The underived tensor and underived Hom agree as abelian groups here, while the derived objects occupy different windows.

12.3 Acyclic resolutions and dimension shifting

Assume T is left exact. An object J is T-acyclic if $R^q T(J) = 0$ for $q > 0$. A bounded-below resolution $A \to J^\bullet$ by T-acyclics computes $\mathrm{R}T(A)$ via $T(J^\bullet)$. One comparison uses the double complex of resolutions, filtered in the two directions, or proceeds by successive truncation of the resolution and the connecting maps. For the latter, write $C^0 = A$, $C^{n+1} = \mathrm{coker}(C^n \to J^n)$. The terms fit

into short exact sequences $0 \to C^n \to J^n \to C^{n+1} \to 0$. Acyclicity gives $R^{q+1}T(C^n) \simeq R^qT(C^{n+1})$ for $q > 0$, and $R^1T(C^n) = \operatorname{coker}(T(J^n) \to T(C^{n+1}))$. Left exactness identifies the kernel of $T(J^{n+1}) \to T(J^{n+2})$ with $T(C^{n+1})$. Together these identities identify the qth cohomology of $T(J^\bullet)$ with $R^qT(A)$ for every q. The comparison into the derived functor is obtained by resolving $J^\bullet$, and these identifications show that it is a quasi-isomorphism. This is the method used for flabby, soft, and fine sheaves below.

There is also a characterization by effaceable connected sequences. A covariant cohomological δ-functor is a sequence of additive functors together with the natural boundaries making each long sequence exact. If for every B and $n > 0$ there is a monomorphism $B \to B'$ sending $T^n(B)$ to zero, then positive degrees are effaceable. The universal δ-functor criterion [40, Ch. 2] identifies an effaceable extension of a left exact T with its right derived sequence. The proof inducts using the exactness around $T^n(B)$: after choosing an effacement, a class is the boundary of a class in $T^{n-1}(B'/B)$, so a transformation in degree $n-1$ determines its required value. Exactness makes it independent of the lift, and comparison of two effacements by a common pushout makes it independent of the chosen monomorphism. For injective resolutions $B \to I$ is a single effacement valid in every positive degree. The theorem of Chapter 11 realizes the same boundary mechanism by tensoring a weight.

12.4 Composition and a spectral sequence

Suppose $S : \mathcal{B} \to \mathcal{D}$ and $T : \mathcal{A} \to \mathcal{B}$ are left exact functors of Grothendieck categories. Assume T sends injectives to S-acyclics. For $A \in \mathcal{A}$, take a bounded-below injective resolution I. The complex $T(I)$ is bounded below and termwise S-acyclic, so the natural comparison $\mathrm{R}S(\mathrm{R}T(A)) \simeq \mathrm{R}(ST)(A)$ follows from the acyclic-complex calculation. One can obtain the latter calculation by resolving all terms and filtering a first-quadrant

double complex, whose column spectral sequence collapses by acyclicity. The filtration of $\mathrm{R}T(A)$ by its cohomological layers induces a spectral sequence

$$E_2^{p,q} = R^p S(R^q T(A)) \implies R^{p+q}(ST)(A). \tag{12.1}$$

Here $p, q \geq 0$. The layer $R^q T(A)[-q]$ contributes in total cohomological degree $p + q$. The connecting maps for successive truncations form the exact couple. With this indexing the first potentially nonzero differential on the layer groups is d_2, of bidegree $(2, -1)$, giving the displayed page convention. Only finitely many layers can affect a fixed total degree since $p, q \geq 0$, so the filtration in that degree converges. This is the classical Grothendieck spectral sequence. More generally the exact couple of a tower can require a convergence argument, as in Chapters 13 and 19.

Exercise 12.3. For an ordinary ring map $R \to S$, show why extension of scalars on derived categories can be composed along $R \to S \to T$ even when neither underived extension of scalars is exact.

Solution. Derived extension of scalars is left adjoint to restriction. The composite of the two restrictions is restriction along $R \to T$, strictly on ordinary modules and coherently on derived categories. Both the iterated derived extension and the direct derived extension are left adjoint to it. Uniqueness of the adjoint gives the comparison $T \otimes^{\mathbf{L}}_S (S \otimes^{\mathbf{L}}_R M) \simeq T \otimes^{\mathbf{L}}_R M$, coherent under another composition. Flatness is needed only to replace a derived tensor by an underived one in a chosen model. □

Chapter 13

Inverse towers and the comparison of maps

13.1 The fiber of the shift difference

Let C be stable with countable products and let $X_0 \leftarrow X_1 \leftarrow \cdots$ be a tower, with transition $u_n : X_{n+1} \to X_n$. The shift s on $P = \prod_{n \geq 0} X_n$ has nth coordinate $u_n \, \mathrm{pr}_{n+1}$. The limit is given by the fiber sequence

$$\lim_n X_n \longrightarrow P \xrightarrow{1-s} P. \tag{13.1}$$

To see this, the nerve of the tower index is the free category generated by the consecutive edges. Its spine is the union of those edges along their endpoints. Restriction of coherent diagrams and coherent cones to the spine is an equivalence, by attaching the composite simplices along inner horns. Equivalently, maps to a tower consist of a map to each X_n with a comparison path for each consecutive edge, with further compositional choices contractible. Thus the limit is the homotopy equalizer of id_P, s. The equalizer is the fiber of their difference in a stable category. This proves (13.1) as a comparison of natural cones.

For a tower of spectra, products commute with all homotopy groups by representability of the groups by shifted spheres. For any countable tower A_n of abelian groups, $\lim A_n$ is the kernel of $1 - s$ on the product and $\lim^1 A_n$ is its cokernel. The long exact sequence of (13.1) gives

$$0 \to \lim\nolimits^1_n \pi_{q+1} X_n \longrightarrow \pi_q \lim_n X_n \longrightarrow \lim_n \pi_q X_n \to 0. \tag{13.2}$$

It holds also in $\mathrm{D}(R)$ with its standard t-structure, by exactness of products of R-modules. For an eventually constant tower

of groups (eventually all transitions are isomorphisms) $1 - s$ is surjective. To solve $(1 - s)x = b$, choose a coordinate beyond which all transitions are isomorphisms, solve forwards using their inverses, and solve backwards along the finitely many earlier arrows. Its kernel identifies with the stabilized group. Thus if every π_q is eventually constant, the limit has those stabilized homotopy groups, with no $\lim^1$ terms. The eventual stage can depend on q.

Calculation 13.1 (A nonzero derived inverse limit)**.** For a prime p, take the tower $\mathbb{Z} \xleftarrow{p} \mathbb{Z} \xleftarrow{p} \cdots$. The limit is zero: the first coordinate is divisible by every power of p, and each later coordinate is determined by it. The formula

$$(b_n)_{n\geq 0} \longmapsto \sum_{n\geq 0} p^n b_n \pmod{\mathbb{Z}}$$

gives a surjection from the product to $\mathbb{Z}_p/\mathbb{Z}$. It annihilates every sequence $b_n = a_n - pa_{n+1}$, since its sum is a_0. Conversely, if the sum is an integer m, set

$$a_n = p^{-n}\left(m - \sum_{i<n} p^i b_i\right).$$

The expression in parentheses is an integer divisible by p^n. Thus $a_n \in \mathbb{Z}$ and $a_n - pa_{n+1} = b_n$. This proves $\lim^1 \mathbb{Z} \cong \mathbb{Z}_p/\mathbb{Z}$ for the indicated tower. Applying (13.2) to the corresponding tower of Eilenberg–Mac Lane spectra gives $H(\mathbb{Z}_p/\mathbb{Z})[-1]$. Every term is connective, while the limit has a nonzero group in degree -1.

13.2 Mapping by two truncations

For $m, n \geq 0$, put $W_{m,n}X = \tau_{[-m,n]}X$. For fixed Y, if $Y \simeq \lim_n \tau_{\leq n}Y$, then $\mathrm{Map}_C(X, Y) \simeq \lim_n \mathrm{Map}_C(X, \tau_{\leq n}Y)$. For fixed X, if $X \simeq \operatorname{colim}_m \tau_{\geq -m}X$, then $\mathrm{Map}_C(X, Y) \simeq \lim_m \mathrm{Map}_C(\tau_{\geq -m}X, Y)$. Both formulas follow directly from representability. They can be combined with two adjunctions to compute each entry in a finite window.

Theorem 13.2 (Finite-window comparison)**.** *Let C be stable with a t-structure. If X admits connective reconstruction and Y admits coconnective reconstruction, there is a natural equivalence*

$$\mathrm{Map}_C(X,Y) \simeq \lim_{(m,n)\in(\mathbb{N}^{\mathrm{op}})^2} \mathrm{Map}_C(W_{m,n}X, W_{m,n}Y). \tag{13.3}$$

The transitions on the right apply the appropriate truncation to both entries. In particular, let $F : C \to D$ be exact and t-exact, with C, D left and right complete. If for all $A, B \in C^{\heartsuit}$ and $q \geq 0$ the maps

$$\mathrm{Hom}_{\mathrm{h}\,C}(A, B[q]) \longrightarrow \mathrm{Hom}_{\mathrm{h}\,D}(FA, FB[q]) \tag{13.4}$$

are isomorphisms, then F is fully faithful.

Proof. The two reconstruction hypotheses identify the left side of (13.3) with the limit of $\mathrm{Map}_C(\tau_{\geq -m}X, \tau_{\leq n}Y)$. Since the target is coconnective in degree n, the adjunction for $\tau_{\leq n}$ allows the source to be replaced by $W_{m,n}X$. This source is connective in degree $-m$ by Lemma 8.1, so the other adjunction allows the target to be replaced by $\tau_{\geq -m}\tau_{\leq n}Y \simeq W_{m,n}Y$. Naturality of both adjunctions and of their truncation comparisons identifies the tower transitions. For example lowering n composes with the target truncation and then uses its universal property on the source. Increasing the lower bound precomposes with a connective cover and then factors through the connective cover of the target. These are the maps that result from applying the indicated truncations to both entries.

For the full faithfulness assertion, the map of mapping spectra for heart objects is an equivalence. The hypothesis gives the isomorphism on homotopy groups in degrees $-q \leq 0$ by (6.1). The positive groups vanish by orthogonality. For fixed A in the heart, exactness of mapping spectra in the second variable and the Postnikov triangle show by induction that $\mathrm{map}_C(A, Y) \to \mathrm{map}_D(FA, FY)$ is an equivalence for every bounded Y. The induction uses the same comparison for a shift of a heart object and the shorter truncation. Next fix a bounded Y and apply the same induction to the first variable, using contravariant exactness.

This gives equivalence of mapping spectra for any pair of bounded objects. Finally F commutes with the truncations by t-exactness, so it induces a pointwise equivalence of the two diagrams in (13.3). Taking limits gives the claimed full faithfulness. □

The inductions on bounded objects are the classical t-structure dévissage of [41, §1.3]. Formula (13.3) packages the two truncation directions in a coherent diagram of mapping spaces and makes explicit which reconstruction hypotheses pass from bounded objects to unbounded ones.

The proof uses all extension degrees in the heart, even when one is interested only in ordinary morphisms between bounded objects. A map between bounded objects can compare two distinct layers, and their relative shift uses a positive group in (13.4). The heart equivalence itself checks only $q = 0$.

Corollary 13.3. *Under the hypotheses of Theorem 13.2, suppose the functor of hearts is essentially surjective. If F preserves limits of coconnective reconstruction towers used to assemble bounded-below objects and colimits of connective reconstruction filtrations, then F is an equivalence. In particular preservation of all sequential limits and colimits is sufficient.*

Proof. Every bounded object of D lifts by induction on the number of degrees in its window. For the induction step, use (8.2). Lift the preceding truncation and the last layer using the induction hypothesis and essential surjectivity on the heart. Full faithfulness on bounded objects lifts the boundary between them. The fiber of the lifted boundary maps to the desired object and belongs to the required window by extension closure. Thus F gives an equivalence of bounded subcategories.

For any $Y \in D$, the objects $W_{m,n}Y$ form a diagram covariant as m increases and contravariant as n increases. The bounded equivalence lifts this entire diagram, including its coherences. First take the inverse limit in n for each m and then the colimit in m. The corresponding operations in D give $\tau_{\geq -m}Y$ and then Y. The stated preservation identifies the image of the constructed

object with Y. With full faithfulness this proves the equivalence. □

In this last proof it suffices that F preserve the limits and colimits of the lifted diagrams, which are compatible finite-window towers and filtrations. The simpler sufficient condition of sequential preservation does not require any identification of a lifted tower with the truncation tower of its limit before applying F.

13.3 Tests and counterexamples

For the standard t-structure on $\mathrm{D}(R)$ the hypotheses of reconstruction hold by exactness of products and filtered colimits, as in Chapter 9. In spectra the same conclusion follows from (13.2) and the preservation of homotopy groups by filtered colimits. The two categories $\mathrm{D}(\mathbb{Z})$ and Sp have equivalent hearts, both Ab. The forgetful functor from $H\mathbb{Z}$-modules to spectra is exact and t-exact and realizes the equivalence of hearts. But it is not fully faithful. In $\mathrm{D}(\mathbb{Z})$, $\mathrm{Hom}(\mathbb{F}_2, \mathbb{F}_2[2]) = 0$ by the length-one resolution. In Sp, the stable mod-2 operation Sq^2 is a nonzero map $H\mathbb{F}_2 \to \Sigma^2 H\mathbb{F}_2$ [14]. Nonvanishing is detected on the generator $x \in H^2(\mathbb{CP}^\infty; \mathbb{F}_2)$, since $\mathrm{Sq}^2 x = x^2 \neq 0$. Thus (13.4) fails in degree two despite agreement in degree zero.

For a bounded example of the mapping formula, let $X = A[i]$ and $Y = B[j]$ for heart objects. The tower is eventually constant once both degrees lie in the interval $[-m, n]$. Its mapping space has components $\mathrm{Hom}(A, B[j-i])$. The high homotopy groups reduce to $\mathrm{Hom}(A, B[j-i-q])$ for $q > 0$, and vanish when the difference is negative. This recovers the extension spectrum calculation from a single finite rectangle.

Exercise 13.4. Suppose a tower of spectra has $\pi_q X_n = 0$ for $q < 0$, and every transition on π_0 is surjective. Show that its limit is connective. Explain why this assertion does not follow from connectivity of the terms alone.

Solution. For any tower of surjections the shift difference on the product is surjective: choose x_0, and solve for x_{n+1} lifting $x_n - b_n$. Hence $\lim^1 \pi_0 X_n = 0$. In (13.2), the only possible group in negative degree is that term at $q = -1$. In calculation 13.1 the transitions on π_0 are not surjective, and the derived limit contributes in degree -1. □

Chapter 14

Cohomological amplitude and output windows

14.1 An input interval for a finite output interval

Let C, D be stable t-categories. Say that an exact functor $F : C \to D$ has amplitude bounded by $[-d, 0]$, for $d \geq 0$, when it takes $C_{\geq 0}$ into $D_{\geq -d}$ and $C_{\leq 0}$ into $D_{\leq 0}$. This is a one-sided cohomological bound in homotopy indexing. In cochain indexing a heart object can have cohomology in degrees 0 through d. For example derived invariants for a group of cohomological dimension d have this form under suitable hypotheses. The bound is about whole halves of a t-structure. A check restricted to heart objects needs an additional argument for unbounded inputs.

Theorem 14.1 (Amplitude and windows)**.** *If F has amplitude bounded by $[-d, 0]$, then for all $X \in C$ and all integers $a \leq b$ there is a natural equivalence*

$$\tau_{[a,b]} F(X) \simeq \tau_{[a,b]} F\Big(\tau_{[a,b+d]} X\Big). \tag{14.1}$$

The equivalence is induced by the span from $F(\tau_{\geq a} X)$ into $F(X)$ and into $F(\tau_{[a,b+d]} X)$.

Proof. The cofiber of $F(\tau_{\geq a} X) \to F(X)$ is $F(\tau_{\leq a-1} X)$, which lies in $D_{\leq a-1}$. Mapping from a test object in $D_{\geq a}$ into this triangle shows that applying $\tau_{\geq a}$ to the first arrow gives an equivalence. The fiber of $F(\tau_{\geq a} X) \to F(\tau_{[a,b+d]} X)$ is $F(\tau_{\geq b+d+1} X)$ by the truncation triangle and Lemma 8.1. By the lower bound it lies in $D_{\geq b+1}$. Mapping from the triangle into an object in $D_{\leq b}$

shows that applying $\tau_{\leq b}$ to the second arrow gives an equivalence. Applying the full window to both arrows and using commutation of truncations proves (14.1). Every map used is a functorial truncation comparison, so the equivalence is natural in X. □

The proof is a finite truncation argument using the standard orthogonality and adjunctions of a t-structure [41, §1.3]. It applies to each object without a completeness hypothesis once the amplitude bounds hold on both entire halves; extending bounds verified only on the heart is the separate issue below.

The same proof gives a shifted version. If $F(C_{\geq 0}) \subset D_{\geq \ell}$ and $F(C_{\leq 0}) \subset D_{\leq u}$ for $\ell \leq u$, the interval is $[a-u, b-\ell]$, since one can apply the theorem to $\Sigma^{-u}F$ with $d = u - \ell$. The shifts express the minimal distances from the tested output to the relevant input layers. When the extremal bounds are realized on heart objects the interval is sharp in the following sense.

Proposition 14.2 (Sharpness on a layer)**.** *Assume there is $A \in C^{\heartsuit}$ with $\pi_{-d}F(A) \neq 0$. The upper endpoint $b+d$ in (14.1) cannot be lowered uniformly. If there is $B \in C^{\heartsuit}$ with $\pi_0 F(B) \neq 0$, the lower endpoint a cannot be raised uniformly.*

Proof. Take $X = A[b+d]$. Removing the highest input degree in the interval removes the whole object, while $\pi_b F(X) = \pi_{-d}F(A) \neq 0$. For the other endpoint take $X = B[a]$ and use $\pi_a F(X) = \pi_0 F(B)$. □

A window is an object with its extension data. Even when the target is a derived category of modules, selecting only the groups in the source interval loses the connecting maps. This can affect a single output group. Consider $R = k[x,y]$ and $k = R/(x,y)$. The Koszul resolution gives $\operatorname{Ext}^i_R(k,k) \cong k, k^2, k$ in degrees $0, 1, 2$ and zero above. Let $u : k \to k[2]$ be the class of its top degree and let $X = \operatorname{fib} u$. Then $\pi_0 X = \pi_1 X = k$. Put $F = \mathrm{RHom}_R(k, -)$ with values in $\mathrm{D}(k)$, using the derived k-linear module structure, or use its underlying $\mathrm{D}(\mathbb{Z})$-valued functor. The resolution of length two yields the bound $[-2, 0]$ on both halves. Applying F to

$$k[1] \longrightarrow X \longrightarrow k \xrightarrow{u} k[2]$$

shows that the map $\pi_0 F(k) \to \pi_0 F(k[2])$ is an isomorphism: it sends the identity to u. The preceding group $\pi_1 F(k)$ is zero. The resulting long exact sequence gives $\pi_0 F(X) \cong \pi_0 F(k[1]) \cong k^2$. For the split object the output group in degree zero is $k^2 \oplus k$. Here the upper layers are the same, but the boundary in the finite input interval changes the answer.

14.2 Extending bounds from the heart

Lemma 14.3. *Let $D = \mathrm{Sp}$ or $D = \mathrm{D}(R)$ with its standard t-structure. Suppose $F : C \to D$ is exact, preserves the limits of the coconnective towers of connective objects, and every connective X is the limit of its tower in C. If $F(A) \in D_{\geq -d}$ for every $A \in C^{\heartsuit}$, then $F(C_{\geq 0}) \subset D_{\geq -d}$.*

Proof. For $X \in C_{\geq 0}$, each $\tau_{\leq n} X$, $n \geq 0$, is built from its finitely many layers in degrees $0, \dots, n$. By the lower bound on the heart and extension closure, $F(\tau_{\leq n} X) \in D_{\geq -d}$. The fiber of its transition to $F(\tau_{\leq n-1} X)$ is $F(\pi_n X)[n] \in D_{\geq n-d}$. For each fixed degree q, the transition on π_q is an isomorphism for $n > q + d$, by the exact sequence. In particular the tower of π_{q+1}'s is eventually constant, and (13.2) gives $\pi_q F(X) = \lim_n \pi_q F(\tau_{\leq n} X)$. This group is zero for $q < -d$. The stated detection of connective objects by their groups in D proves the bound. □

For the upper bound, let F be right adjoint to a functor $L : D \to C$ preserving connective objects. By suspension, L carries $D_{\geq 1}$ into $C_{\geq 1}$ when L is exact. For $X \in C_{\leq 0}$ and $U \in D_{\geq 1}$, $\mathrm{Map}_D(U, F(X)) \simeq \mathrm{Map}_C(LU, X) \simeq *$. Orthogonality gives $F(X) \in D_{\leq 0}$. Thus a derived right adjoint often has the upper bound immediately, while the lower bound comes from a cohomological dimension argument such as Lemma 14.3. This proof also indicates where the hypotheses enter: without eventual stabilization, a tower of uniformly bounded-below objects can have a lower homotopy group in the limit.

14.3 Two finite tests

First take $R = \mathbb{Z}$, $M = \mathbb{Z}/p$, and $F = \mathrm{RHom}_{\mathbb{Z}}(M, -)$. The resolution $[\mathbb{Z} \xrightarrow{p} \mathbb{Z}]$ identifies $F(X)$ with $\mathrm{fib}(p : X \to X)$. For every q there is an exact sequence

$$0 \longrightarrow \pi_{q+1}X/p\pi_{q+1}X \longrightarrow \pi_q F(X) \longrightarrow (\pi_q X)[p] \longrightarrow 0.$$

It directly proves the amplitude $[-1, 0]$. The input interval for degree q is $[q, q + 1]$; the two end groups in the short exact sequence are computable from the homotopy objects of X. The short exact sequence follows from a coherent fiber, and its middle group is specified by the whole two-layer window.

Second take $F : \mathrm{Fun}([1], \mathrm{D}(R)) \to \mathrm{D}(R)$ to be evaluation at 1. It has bound $d = 0$, since truncations are pointwise. In contrast, the limit of a diagram of shape $B\mathbb{Z}$ is the fiber of $1 - v$ from Chapter 3. It has bound $d = 1$ by the same long exact calculation, and this bound is sharp on a nonzero trivial action. The two index categories have different coherent cone requirements. In the first case the actual limit is evaluation at 0, which also has bound 0. In the second, a cycle forces a path comparison between a value and its transport.

14.4 Composition and a relative form

If exact functors $F : C \to D$, $G : D \to E$ have amplitude bounds $[-d, 0]$ and $[-e, 0]$, the composite has bound $[-(d+e), 0]$. This follows by applying the two inclusions to shifted halves. Equivalently one may compute an output window $[a, b]$ of GF by first retaining the window $[a, b + e]$ of F, and then retaining $[a, b + e + d]$ in the original input. The comparison spans in Theorem 14.1 are natural, so these two computations identify the same window of the composite. No equality of the best possible bounds is asserted by this additive upper bound.

A relative version works for a square of exact t-categories and a comparison $F \to G$ between functors of the same bounded amplitude. Applying Theorem 14.1 to both sides identifies the map

between output windows with its computation on the same input window. Thus to verify that a comparison induces equivalences in degrees $[a, b]$ on every input, it suffices to verify it on inputs already in $[a, b + d]$. If the comparison is between exact functors and one has a natural transformation between them, one can use finite Postnikov induction to verify it on those bounded inputs by testing on the corresponding shifted heart objects. This last induction may require testing an extra endpoint if only partial comparison bounds are known. When the transformation is an equivalence on the entire output of each heart object, exactness gives the induction without an endpoint adjustment.

Exercise 14.4. Suppose F has amplitude bounded by $[\ell, u]$ with $\ell = -2, u = 1$. Which input interval suffices for output $[-1, 2]$? Give a test that would prove its upper endpoint sharp.

Solution. The interval is $[-2, 4]$ by the shifted formula. To test the upper endpoint, find a heart object A with $\pi_{-2}F(A) \neq 0$. Then $A[4]$ has zero smaller input window, but its output has nonzero π_2. The lower endpoint is tested by a heart object with nonzero $\pi_1 F(A)$ shifted to degree -2. □

Part IV

Higher sheaves and their cohomology

Covering sieves, hypercovers, and exact sheafification give the passage from presheaves to derived coefficients. Open decompositions provide local section and extension formulas. Finite posets, higher monodromy, and a relative projective line then test these formulas with explicit cochains. A family of spheres accumulating at one point supplies a sheaf-theoretic test of the completeness hypothesis, and local-to-global Ext connects the calculations back to the classical extension groups.

Chapter 15

Sites and sheaves of spaces

15.1 The topology of a sieve

Let C be a small ordinary category. A sieve S on U is a subpresheaf of the ordinary representable yU closed under precomposition. A Grothendieck topology specifies covering sieves, requiring the maximal sieve to cover, stability under pullback of sieves, and transitivity of covering. Regard a sieve as a presheaf of discrete spaces. A space-valued presheaf F is a sheaf if, for every covering sieve,

$$F(U) \simeq \mathrm{Map}_{\mathcal{P}(C)}(yU, F) \longrightarrow \mathrm{Map}_{\mathcal{P}(C)}(S, F) \quad \text{is an equivalence.} \tag{15.1}$$

By Theorem 2.3 the right side is the limit of the values $F(V)$ over the opposite of the category of arrows $V \to U$ in the sieve. Its objects are local sections and its paths and higher cells enforce coherent compatibility. For discrete F, a limit of discrete spaces is discrete and calculated as the ordinary set of matching families. Thus the 0-truncated sheaves are precisely ordinary sheaves of sets. One can see the reverse implication by testing truncation on the sheafifications of representables, which detect all evaluations.

In the site of open subsets of a topological space, the topology uses ordinary open covers. For a family $\{U_i\}$ covering U, the Čech nerve of $\coprod_i yU_i \to yU$ has degree n equal to the coproduct of representables of the $(n+1)$-fold intersections. In presheaves its realization is the generated sieve. Indeed the map of presheaves is pointwise a map of discrete sets, and its Čech realization at each test object is the ordinary image, by Proposition 5.2. The same statement works in a site with the required fiber products for the covering family. Mapping into F gives the coherent Čech

formula

$$F(U) \simeq \operatorname{Tot}\left([n] \longmapsto \prod_{i_0,\dots,i_n} F(U_{i_0} \cap \cdots \cap U_{i_n}) \right). \tag{15.2}$$

Here Tot means the limit over Δ, with all faces and degeneracies. In a general site the sieve version requires no fiber products in C.

15.2 Sheafification as a localization

Write $\mathcal{X} = \operatorname{Sh}(C; \mathcal{S})$. This is the accessible left exact localization of $\mathcal{P}(C)$ at the covering sieve inclusions [18, §6.2.2]. Its reflector a is sheafification. The adjective left exact means that a preserves finite limits. The inclusion preserves all limits, because being local is tested by mapping-space limits. Colimits in $\mathcal{X}$ are formed by taking a presheaf colimit and applying a, as follows from its adjunction. In particular the image in $\mathcal{X}$ of a covering sieve is the entire sheafified representable.

One construction starts with matching-family operators. For a presheaf F, the first operator at U takes a colimit of $\operatorname{Map}(S, F)$ over covering sieves ordered by refinement, using restriction along an inclusion of smaller sieves. Intersections of covering sieves are covering, so the refinement indexing is filtered. Filtered colimits of spaces commute with finite limits and each $\operatorname{Map}(S, -)$ preserves limits. This is the source of left exactness of the operator. Iterating the operator at an adequate regular cardinal yields sheafification; its accessible localization property and convergence are standard [18, §6.2.2]. In the set-valued theory two applications suffice. The transfinite formulation avoids a bound on the number of coherences that must be imposed for arbitrary space-valued data.

Example 15.1 (Sections of a constant sheaf)**.** Let X be the disjoint union of two nonempty connected spaces. The constant presheaf on a set A has value A even on the empty open set. Its sheafification is the sheaf of locally constant maps to the

discrete space A. A section over all of X is a pair in $A \times A$. One proves this by taking the sheaf of germs: local sections of the presheaf are constant, and specifying their compatible germs is precisely local constancy. For the empty open, the sheaf has a single section, enforced by the empty cover.

15.3 Stalks and concrete sheaves

For a topological space X and $x \in X$, the stalk of a presheaf is the filtered colimit over open neighborhoods of x, ordered by reverse inclusion. For a sheaf of sets F the stalk F_x consists of germs of sections. The disjoint union of all F_x maps to X and has a topology with basic opens

$$\{\, [s]_x \mid x \in U \,\}, \qquad s \in F(U).$$

If two such basic opens meet, equality of germs gives an open neighborhood where the sections agree, so these sets are a basis. Projection restricts to a homeomorphism on each such basic open onto U. Conversely the local sections of a local homeomorphism over X form a sheaf. A continuous section of the germ space is locally one of the basic sections, and compatibility glues it uniquely to a section of F. These operations identify ordinary sheaves of sets with spaces étalé over X. This recovers the concrete sheaf construction in [38, §§4.2–4.5].

A sheaf of spaces also has a stalk computed by this neighborhood colimit. Stalks are inverse image functors of points of the sheaf ∞-topos. The functors preserve finite limits because the colimits are filtered, and they preserve colimits in the sheaf category by the inverse image adjunction. On finite truncations, a map of sheaves is an equivalence when it induces equivalences on all stalks. For 0-truncated sheaves this follows by local lifting and local equality of sections. For higher finite truncations, homotopy fibers and repeated diagonals reduce the assertion by induction to sheaves of sets. Thus the stalks test all homotopy sheaves. For unbounded sheaves of spaces the use of the stalks as a test for equivalence itself requires a completeness condition.

15.4 Hypercovers and Whitehead tests

A hypercover of U is an augmented simplicial diagram $V_\bullet \to U$ in which V_n covers the nth relative matching object for each $n \geq 0$. Matching objects are obtained by right Kan extension of the diagram through degrees $< n$, so the matching map checks that compatible lower faces can be filled locally. A Čech nerve is a hypercover with each positive-degree matching map an equivalence. General hypercovers allow a new covering step for each degree. For sites it is common to take V_n to be a coproduct of representables and to use locally surjective matching maps.

Hypercompletion $\mathcal{X}^\wedge$ is the localization at the infinitely connective maps of $\mathcal{X}$, those that induce equivalences on all finite relative truncations. It is left exact. An object is hypercomplete if it is local for this class, and equivalently if it satisfies descent for hypercovers in a suitable site presentation [18, 8, §6.5]. Hypercover descent is again the totalization formula, now for all the levels of $V_\bullet$. The equivalence between hyperdescent and the Whitehead localization in the cited theorem explains why stalkwise equivalence suffices for hypercomplete sheaves on a topological space: it gives isomorphism on the finite local homotopy invariants, and the localization inverts that comparison. For categories of spectra one obtains the corresponding statement by testing all infinite loop space levels.

A finite truncation F is hypercomplete. For a map whose truncations are equivalences, every mapping-space test into a finite truncation factors through the finite truncated comparison. Hypercomplete objects are closed under limits. It follows that an object reconstructed as the limit of its Postnikov truncations is hypercomplete. The implication in the other direction requires further hypotheses. Hypercompletion is defined by a class of local equivalences; Postnikov completeness asserts a limit reconstruction inside the chosen category. In stable sheaf categories these distinctions are treated explicitly in [7, 20]. Finite-window arguments below state which reconstruction is used.

15.5 Topos descent and objects over a base

An ∞-topos can be characterized by the higher Giraud axioms [18, Thm. 6.1.0.6]. In one formulation it is presentable, colimits are universal, coproducts are disjoint, and groupoid objects are effective. Universal colimits are those stable under pullback. Effectivity identifies a quotient with the object whose Čech nerve is the given groupoid object, as in Chapter 5. A consequence is descent for slices along an effective epimorphism $U \to V$:

$$\mathcal{X}_{/V} \simeq \lim_{[n]\in\Delta} \mathcal{X}_{/U^{\times V^{(n+1)}}}.$$

On the right the transition functors are pullbacks. An object over V pulls back to a compatible family. Conversely a compatible family gives a groupoid diagram of its total objects. Its colimit over the groupoid nerve maps to V, and universal colimits give the original family after pulling back. Mapping spaces descend by the same pullback and colimit universal properties. This result identifies descent as a gluing operation on entire slice categories. For sets it reduces to gluing families over the fibers of a surjection.

Exercise 15.2. For a discrete group G, apply slice descent to $* \to BG$ in spaces. Identify the limit of categories on the right.

Solution. The n-fold index is a discrete set G^n and the slice of spaces over it is $\mathrm{Fun}(G^n, \mathcal{S})$. A compatible object is a space with a coherent G-action, with transition coherences recorded by faces and degeneracies. Straightening over BG identifies the limit with $\mathrm{Fun}(BG, \mathcal{S})$ (using inversion to pass from the opposite groupoid if needed). Its underlying bundle is the total object over BG. □

Chapter 16

Change of space and supports

16.1 Geometric morphisms

A geometric morphism $f : \mathcal{X} \to \mathcal{Y}$ of ∞-topoi is specified by a left exact left adjoint $f^* : \mathcal{Y} \to \mathcal{X}$ with right adjoint f_*. On topological spaces, for $f : X \to Y$, the direct image on a sheaf F is

$$(f_*F)(V) = F(f^{-1}V).$$

Indeed preimages preserve unions and finite intersections, so the descent of F gives the descent of f_*F. The inverse image is obtained by sheafifying the presheaf given at $U \subset X$ by the colimit of $G(V)$ over $V \supseteq f(U)$, ordered by reverse inclusion. This is the left Kan extension along the map of open indices in the appropriate variance. The indexing categories of such neighborhoods are filtered. Finite limits commute with those colimits, and sheafification is left exact, so f^* is left exact. Adjunction is the Kan adjunction followed by sheafification. In the set-valued case its stalk at x is $G_{f(x)}$, which also follows by restricting to neighborhoods.

For continuous maps $X \xrightarrow{f} Y \xrightarrow{g} Z$, direct images compose by inverse images of opens. Uniqueness of left adjoints identifies $(gf)^* \simeq f^*g^*$ and gives coherent composition for longer chains, as in Chapter 4. This is the higher-categorical form of the comparisons of direct and inverse image in [38, §4.7]. In stable sheaf categories the direct image is exact as a functor of stable categories: it preserves limits and hence finite colimits by Proposition 6.1. Its induced degree-zero functor on ordinary

coefficients can have higher derived functors. These are measured by the nonzero homotopy sheaves of the resulting stable object.

16.2 Open base change

Let $j : V \hookrightarrow Y$ be an open inclusion and $j' : X_V = f^{-1}V \hookrightarrow X$. Write $f_V : X_V \to V$ for the restriction. There is a canonical equivalence

$$j^* f_* F \simeq (f_V)_* (j')^* F. \tag{16.1}$$

Both sides evaluate on an open $W \subset V$ as $F(f^{-1}W)$. The restriction maps are the same, so this is an equivalence of sheaves, and the comparison is the mate of the equality of the two composites of inverse images. The argument works for sheaves with stable coefficients when F assigns its derived sections to opens. For ordinary module sheaves the derived direct image has this same formula with $\mathrm{R}\Gamma(f^{-1}W, F)$ as the derived value. It can be seen by using an injective resolution and the fact that restriction along an open inclusion preserves injectives, since its left adjoint, extension by zero, is exact. Restriction therefore commutes with this derived operation.

Calculation 16.1 (A base-change failure for a closed point)**.** Let $j : (0, \infty) \hookrightarrow \mathbb{R}$ be the open inclusion, $i : \{0\} \hookrightarrow \mathbb{R}$, and $A \neq 0$ a constant abelian coefficient. The geometric pullback of j along i is empty. The right side of the base-change comparison for this pullback has zero stable value. The left side $i^* \mathrm{R} j_* \underline{A}$ has $\pi_0 = A$, since every sufficiently small interval about 0 meets $(0, \infty)$ in a connected interval and the transition maps on its sections are identities. The stalk colimit computes π_0 of this value. Thus the comparison fails. The topological square is Cartesian, while the neighborhoods used in the direct image still see the open half-line.

16.3 A slice formula

Let $f : X \to Y$ be a map of objects in an ∞-topos $\mathcal{X}$, viewed now through the slice base-change adjoints. Suppose $g : Y' \to Y$, $X' = Y' \times_Y X$, and $f' : X' \to Y'$, $g' : X' \to X$ are the other edges. In slices there is always a Beck–Chevalley equivalence

$$g^* f_* Z \simeq f'_* (g')^* Z \qquad (Z \in \mathcal{X}_{/X}).$$

To verify it, let $W \to Y'$ be a test object. By the adjunctions of Chapter 4,

$$\begin{aligned} \mathrm{Map}_{/Y'}(W, g^* f_* Z) &\simeq \mathrm{Map}_{/Y}(g_! W, f_* Z) \\ &\simeq \mathrm{Map}_{/X}(f^* g_! W, Z), \\ \mathrm{Map}_{/Y'}(W, f'_* (g')^* Z) &\simeq \mathrm{Map}_{/X'}((f')^* W, (g')^* Z) \\ &\simeq \mathrm{Map}_{/X}(g'_! (f')^* W, Z). \end{aligned}$$

The objects $f^* g_! W$ and $g'_! (f')^* W$ are both $X \times_Y W$ over X. Hence the two functors of W agree. This formula concerns dependent products of objects in slices. Calculation 16.1 concerns geometric direct images of sheaves along maps of spaces. The dependent product argument applies when the given square is a square of these slice constructions.

16.4 Open and closed decompositions

Fix a topological space X, open subset U, and closed complement Z, with inclusions $j : U \to X$ and $i : Z \to X$. Work with derived sheaves of modules over a fixed ordinary ring. The stalk functors detect equivalences in this model, as proved in Chapter 17. The functor j^* has a fully faithful left adjoint $j_!$, extension by zero, and a fully faithful right adjoint $\mathrm{R}j_*$. The first has the prescribed stalk on U and zero stalk on Z, and is exact by the stalk criterion for ordinary sheaves. The second is fully faithful by (16.1) applied to U. The functor i^* has fully faithful right adjoint i_*, exact for a closed embedding: at points of Z its stalks are those of the

original sheaf, and outside Z there is a neighborhood disjoint from Z. This also shows i_* preserves colimits. Its right adjoint in the derived setting is denoted $i^!$.

Proposition 16.2 (The two support sequences)**.** *There are functorial cofiber sequences*

$$j_!j^*F \longrightarrow F \longrightarrow i_*i^*F, \qquad i_*i^!F \longrightarrow F \longrightarrow \mathrm{R}j_*j^*F.$$

Proof. The first comparison has composite zero because $i^*j_! = 0$. Its cofiber map is an equivalence on U, where both cofiber and i_*i^*F vanish, and on Z, where both identify on stalks with F. Thus it is an equivalence. In particular any sheaf restricting to zero on U is in the image of i_*. The fiber K of the second comparison restricts to zero on U, since $j^*\mathrm{R}j_* \simeq \mathrm{id}$. For any B on Z the space of maps from i_*B into $\mathrm{R}j_*j^*F$ is contractible, by $j^*i_*B = 0$ and the adjunction for j_*. Thus maps from i_*B into K agree with those into F. Since K is in the essential image of i_*, its inverse image represents $i^!F$ by the universal property. This gives the second sequence. □

Applying derived global sections gives the triangle

$$\mathrm{R}\Gamma_Z(X, F) \longrightarrow \mathrm{R}\Gamma(X, F) \longrightarrow \mathrm{R}\Gamma(U, j^*F)$$

where the first term is $\mathrm{R}\Gamma(Z, i^!F)$. Indeed the second support sequence initially gives $\mathrm{R}\Gamma(X, i_*i^!F)$, and composition of the section and direct-image right adjoints gives

$$\mathrm{R}\Gamma(X, i_*i^!F) \simeq \mathrm{R}\Gamma(Z, i^!F).$$

In degree zero for an ordinary sheaf it consists of sections whose restriction to U vanishes. The higher terms measure the obstruction to solving that vanishing problem in resolutions. The triangle is the cohomological exactness axiom for the open complement.

Calculation 16.3 (A local contribution at a point)**.** Let $X = \mathbb{R}^2$, $Z = \{0\}$, and use the constant sheaf with coefficients A an abelian group. The classical comparison with singular cohomology applies

to these locally contractible paracompact spaces [4, 15]. The derived global sections on X are A in degree zero and on $U = \mathbb{R}^2 - \{0\}$ they have groups A in degrees $0, 1$. Restriction is the identity on H^0. The long exact support sequence then gives $H^q_Z(X, A) = 0$ for $q \neq 2$ and $H^2_Z(X, A) = A$. The boundary from $H^1(U, A)$ supplies this isomorphism. Replacing the open complement by a contractible model would remove this local contribution.

Exercise 16.4. If $X = U \amalg Z$ is the disjoint union of two open and closed subsets, show that the two support sequences split and identify all four adjoints for the two inclusions.

Solution. The empty overlap in descent identifies the sheaf category on X with the product of the categories on the two pieces. Both pullbacks are projections. Each projection has both adjoints given by inserting zero in the other coordinate, since maps to and from zero are contractible in stable categories. Therefore $j_! \simeq \mathrm{R}j_*$ and $i^* \simeq i^!$ in this case. The two triangles are the canonical split product triangles. □

Chapter 17

Derived sheaves of modules

17.1 Presheaves, sections, and exact sheafification

Fix an ordinary ring R and a small ordinary site C. Write $\mathcal{B} = \mathrm{Fun}(C^{\mathrm{op}}, \mathrm{Mod}_R)$ for ordinary presheaves of left R-modules and $\mathcal{A} = \mathrm{Sh}(C; \mathrm{Mod}_R)$ for ordinary sheaves. The inclusion $i : \mathcal{A} \to \mathcal{B}$ has a left adjoint a, sheafification. For modules the underlying abelian group of the sheafification is the sheafification of the underlying groups. Indeed operations descend along finite products, which sheafification preserves, and the extension of each action map is determined by the adjunction. The functor a is additive, left exact by the finite-limit property of sheafification, and right exact because it is a left adjoint. Therefore it is exact. This can also be checked for a topological space on stalks, where the sheafification map is an isomorphism and stalks are exact filtered colimits.

For $U \in C$, let $R[yU]$ denote the presheaf freely generated by the ordinary representable on U. Thus at V it is the direct sum of copies of R indexed by $V \to U$. A map of presheaves from $R[yU]$ to F is determined by its value on id_U, and a map of sheaves from $aR[yU]$ is determined by the same rule using the adjunction. These presheaves are projective generators of $\mathcal{B}$: evaluation is exact and detects zero morphisms. The sheafified objects generate $\mathcal{A}$. Their projectivity can fail because sheaf epimorphisms give local rather than global lifts of sections. Cocompleteness and exact filtered colimits of $\mathcal{A}$ follow by forming the colimits in $\mathcal{B}$

and applying a. Thus the resolutions used in Chapter 12 exist in both categories.

17.2 Exact localization at local quasi-isomorphisms

Call a map of presheaf complexes a local quasi-isomorphism if it induces isomorphisms after sheafifying the presheaves of cohomology groups. In the category of presheaves with values in $\mathrm{D}(R)$, this means that all associated homotopy sheaves of the map are isomorphic. The equivalence

$$\mathrm{D}(\mathcal{B}) \simeq \mathrm{Fun}(C^{\mathrm{op}}, \mathrm{D}(R)) \tag{17.1}$$

uses the projective or injective model structure on diagrams of complexes, with weak equivalences the objectwise quasi-isomorphisms [6, 18]. The strictification theorem for diagrams in a combinatorial model category gives the underlying ∞-category of coherent diagrams on the ordinary index category. On objects, evaluation of the complex of presheaves gives the displayed diagram. This theorem is what identifies the mapping spaces of coherent diagrams as well.

Theorem 17.1 (Derived sheafification)**.** *The functor* $\mathrm{D}(a) : \mathrm{D}(\mathcal{B}) \to \mathrm{D}(\mathcal{A})$ *is an exact reflective localization, with fully faithful right adjoint* $\mathrm{R}i$*. Under* (17.1) *the localization inverts exactly the local quasi-isomorphisms. For* $K \in \mathrm{D}(\mathcal{A})$*, the diagram of its right adjoint values is* $U \mapsto \mathrm{R}\Gamma(U, K)$*.*

Proof. For a K-injective complex I in $\mathcal{A}$, the complex iI is K-injective in $\mathcal{B}$. In fact if P is acyclic in $\mathcal{B}$, then aP is acyclic by exactness, and the complex of maps from P to iI is the complex of maps from aP to I, which is acyclic. The same adjunction of mapping complexes now gives the derived adjunction $\mathrm{D}(a) \dashv \mathrm{R}i$. The counit on I is computed termwise by $aiI \simeq I$. Hence it is an equivalence, which makes $\mathrm{R}i$ fully faithful. Exactness of a gives $H^q(aP) = aH^q(P)$, so the inverted arrows are exactly those

specified. Finally $(iI)(U) = I(U)$ computes $\mathrm{R}\Gamma(U, K)$ because the section functor is the ordinary right adjoint evaluated by $\mathrm{Hom}_{\mathcal{A}}(aR[yU], -)$. This identifies the diagram with the derived sections on every U and its restriction maps. □

For a topological space a map of complexes of sheaves is a quasi-isomorphism if and only if it is so on all stalks. Indeed a map of sheaves of modules is an isomorphism if and only if all its stalks are, and stalks commute with cohomology by exactness. Therefore the stable category $\mathrm{D}(\mathrm{Sh}(X; \mathrm{Mod}_R))$ admits the stalk detection used in Chapter 16. This argument depends on working in the localized derived category.

17.3 Hyperdescent and the comparison

Proposition 17.2. *Under* (17.1)*, the full subcategory given by* Ri *is the category of* $\mathrm{D}(R)$*-valued presheaves satisfying hypercover descent.*

Proof. Let $V_\bullet \to U$ be a hypercover built from coproducts of representables. The augmented normalized complex of free presheaf modules on it is a local quasi-isomorphism onto $R[yU]$. This is the linearization of the local equivalence of the realization of a hypercover with its target [8]. A direct verification uses the matching conditions: for each finite collection of simplices of a section of the simplicial diagram, a cone with a local apex can be chosen after a cover, by first lifting the apex and then filling the prescribed boundary of each cone simplex in order of dimension. At each finite step refine the cover to satisfy the finitely many conditions. A finite normalized cycle is then locally the boundary of the sum of its cone simplices. In degree zero the same argument applies to the augmentation kernel, and the augmentation is locally surjective by the covering of U. After sheafification the augmented complex is exact. This proves the asserted quasi-isomorphism.

Mapping out of that augmentation into $\mathrm{R}iK$ is an equivalence by Theorem 17.1. The mapping spectrum out of a realization is the totalization of the level mapping spectra, by (3.1) and shifts. The left term is the section value at U, while the terms on the right are products of section values on the representables in V_n. Thus the presheaf satisfies hyperdescent.

Conversely, suppose F satisfies this descent and consider the unit $F \to \mathrm{R}i\,\mathrm{D}(a)F$. Its cofiber Q satisfies hyperdescent because in stable categories totalizations, being limits, preserve finite fiber and cofiber sequences. The unit is a local quasi-isomorphism, so all homotopy sheaves of Q vanish. Each presheaf $\Omega^\infty\Sigma^n Q$ (using the underlying spectrum) satisfies hyperdescent, since it is obtained by limit-preserving functors. Its associated component sheaf is terminal and its higher homotopy sheaves are trivial. Translation in the group-like section spaces identifies the tests at arbitrary local basepoints with the test at zero. By hypercomplete Whitehead, each of these sheaves is terminal. All the section spectra of Q are therefore zero. Hence the unit is an equivalence and F belongs to the essential image. □

This proof specifies the relation to Čech descent. Every derived sheaf diagram satisfies Čech descent. A diagram with only Čech descent has the comparison to the hyperdescent localization, and the comparison is an equivalence exactly when the additional hypercover tests hold. A sheaf of ordinary modules A, placed in cochain degree zero in $\mathcal{A}$, corresponds to the diagram $U \mapsto \mathrm{R}\Gamma(U, A)$, whose section cohomology may be nonzero in positive degrees. The presheaf $U \mapsto H(A(U))$ can fail even the stable Čech test. Both diagrams have the same associated homotopy sheaves, which by itself tests equivalence only after the appropriate localization.

Here homotopy sheaves are the sheafifications of the homotopy groups of the section diagrams. The sheaf t-structure on $\mathrm{D}(\mathcal{A})$ is transported to this hyperdescent model through $\mathrm{R}i$; truncating an object there can change its sections beyond the objectwise truncation of the old section spectra. Hypercomplete Whitehead detects a local object whose associated homotopy sheaves van-

ish. It does not assert that homotopy sheaves commute with infinite limits. In Chapter 27, a product of positive shifts acquires a nonzero degree-zero homotopy sheaf, and the Postnikov comparison consequently fails this Whitehead test.

17.4 Varying rings and change of coefficients

For a sheaf of ordinary rings $\mathcal{O}_X$ on a topological space, one can sheafify presheaves of $\mathcal{O}_X$-modules by sheafifying the action as well. Stalkwise this gives the usual $(\mathcal{O}_X)_x$-modules. The adjunction is again exact on the left, and the derived localization argument of Theorem 17.1 applies to the abelian presheaf and sheaf categories. The identification (17.1) with a single fixed target $\mathrm{D}(R)$ was for constant coefficients. For varying coefficients the module at U is over $\mathcal{O}_X(U)$, with restriction involving the maps of rings. A fixed-target diagram of plain modules does not specify these actions.

For a morphism of ringed spaces $f : (X, \mathcal{O}_X) \to (Y, \mathcal{O}_Y)$, the inverse image of underlying sheaves f^{-1} is exact, but extension of scalars can require derivation. The left adjoint to $\mathrm{R}f_*$ on derived modules is

$$\mathrm{L}f^*(M) = \mathcal{O}_X \otimes^{\mathbf{L}}_{f^{-1}\mathcal{O}_Y} f^{-1}M.$$

The formula follows by the two adjunctions for inverse image and extension of scalars, after taking K-flat and K-injective models [34]. Flatness of $\mathcal{O}_X$ over $f^{-1}\mathcal{O}_Y$ suffices to use the underived formula on discrete modules. For the ring map $\mathbb{Z} \to \mathbb{Z}/p$ on one-point spaces and $M = \mathbb{Z}/p$, calculation 12.2 gives an extra homotopy group in degree 1. The geometric inverse image on underlying sheaves is exact in this example; the additional group is caused entirely by the scalar extension.

Exercise 17.3. For a point with coefficient ring R, determine the functors in Theorem 17.1. For an infinite discrete space with the same coefficients, identify the product and coproduct descriptions of its sheaf category.

Solution. A sheaf on a point is specified by the module on the point, with zero on the empty open. The derived section diagram is the chosen complex on the point and zero on the empty open. For an infinite discrete space X, the point inclusions give a cover by pairwise disjoint opens. A sheaf is specified by its values on these points, and its value on U is the product of the values on the points in U. The stable category is $\prod_{x \in X} \mathrm{D}(R)$ with coordinatewise products and coproducts. Global sections is the product functor; the categorical coproduct of a family of sheaves is computed at each point as a direct sum. □

Chapter 18

Global sections and acyclic sheaves

18.1 Representing global sections

For a topological space X and an ordinary ring R, let R_X denote the constant sheaf of left R-modules. A map $R_X \to A$ is determined by the image of 1, a global section of A. The same statement holds over an open set. The adjunction between the constant sheaf functor and global sections extends to derived categories:

$$\Delta_X : \mathrm{D}(R) \rightleftarrows \mathrm{D}(\mathrm{Sh}(X; \mathrm{Mod}_R)) : \mathrm{R}\Gamma(X, -).$$

Here Δ_X is formed by sheafifying a constant presheaf complex. It is exact on ordinary modules, because this can be checked on its stalks, and it preserves colimits. The adjunction on complexes follows from the underived one using a K-injective target. For the rank-one free module R, the right adjoint has underlying spectrum $\mathrm{map}(R_X, -)$. Hence

$$H^q(X, A) = \pi_{-q}\mathrm{R}\Gamma(X, A) \qquad (A \in \mathrm{Sh}(X; \mathrm{Mod}_R))$$

is classical sheaf cohomology. This also follows directly by applying Γ to an injective resolution. Every cofiber sequence of derived sheaves yields a long exact sequence of cohomology, as derived global sections is exact by Proposition 6.1. For the point, the right adjoint is evaluation and discrete coefficients have only degree-zero cohomology.

A sheaf A is acyclic for global sections if $H^q(X, A) = 0$ for $q > 0$. For such A, the natural arrow from its underived sections

in degree zero to $\mathrm{R}\Gamma(X, A)$ is an equivalence. It suffices for computation that the terms of a bounded-below resolution be acyclic for the selected section functor, by the proof in Chapter 12. The following classes of sheaves give concrete resolutions.

18.2 Flabby sheaves

A sheaf of modules A on X is flabby (also called flasque) if $A(V) \to A(U)$ is surjective for all $U \subseteq V$ open. This is a stronger lifting condition than the local one in a sheaf epimorphism. It is inherited by restrictions to opens. The key lifting lemma is global.

Lemma 18.1. *For a short exact sequence $0 \to A \to B \to C \to 0$ of sheaves of modules with A flabby, $B(U) \to C(U)$ is surjective for every open U. If B is also flabby, so is C.*

Proof. Fix $s \in C(U)$. Consider lifts over open subsets of U, ordered by extension of the lift. A chain of lifts glues over its union by the sheaf condition. Choose a maximal lift t over $W \subseteq U$. If a point of U lies outside W, the local surjectivity of the given epimorphism gives a lift t' over a neighborhood $V \subset U$ of that point. The difference of the lifts on $V \cap W$ is in $A(V \cap W)$. It extends to $A(V)$ by flabbiness. Adjust t' by this extension and glue to t, contradicting maximality. Thus $W = U$. If B is flabby, a section of C on $U \subset V$ lifts to $B(U)$, extends to $B(V)$, and maps to an extension in $C(V)$. □

An injective sheaf of modules is flabby. To prove this, let $j_U : U \hookrightarrow X$ and $j_V : V \hookrightarrow X$ be open inclusions. The comparison $j_{U!}R_U \to j_{V!}R_V$ is a monomorphism, by its stalks. Mapping it into the injective sheaf gives a surjection, which identifies by adjunction with restriction from V to U. The argument also works for a varying sheaf of rings $\mathcal{O}_X$ by using its restrictions as free rank-one modules. Now resolve a flabby sheaf A by injectives. At each step embed the current flabby cokernel into a flabby injective. Its new cokernel is flabby by Lemma 18.1, and the

sequence remains exact on sections. Applying this at every step shows $H^{q>0}(U, A) = 0$. This proves flabby acyclicity on each open [33].

18.3 The Godement construction

For a sheaf A of modules, form the sheaf

$$G(A)(U) = \prod_{x \in U} A_x.$$

It assigns an arbitrary family of germs to each open set. Restrictions are coordinate projections, so it is flabby. A section gives its family of germs, which embeds A into $G(A)$. If a section has all germs zero, its restrictions vanish on neighborhoods of all points and the sheaf property gives its vanishing. The same formula gives a sheaf of $\mathcal{O}_X$-modules using the maps $\mathcal{O}_X(U) \to (\mathcal{O}_X)_x$.

Set $C^0 = A$, $I^n = G(C^n)$, and $C^{n+1} = \operatorname{coker}(C^n \to I^n)$. The differential $I^n \to I^{n+1}$ is the quotient map followed by the germ map of C^{n+1}. The resulting augmented complex is exact by construction and is functorial in A. Its terms are flabby, so $\Gamma(U, I^\bullet)$ computes the derived sections of A for every open U. This gives a model compatible with all restrictions at once. The formula uses products of stalk modules, which are taken in ordinary modules and are exact. The sheaf-theoretic cokernels are still formed by sheafification and can differ from pointwise cokernels. Tracking this distinction is necessary when evaluating the resolution on an open.

Example 18.2 (Acyclic on one space and on all opens)**.** For an injective A, $H^{q>0}(X, A) = 0$ follows from the definition of derived functors. The flabby argument yields the stronger statement on every open U and makes it independent of the initial injective model. For a sheaf known only to be X-acyclic, that statement on opens is a separate property. An open inclusion j computes $(\mathrm{R}j_*A)(V)$ from sections on $V \cap U$, so acyclicity on all these opens is the property needed to replace the derived direct image by an underived one.

18.4 Soft and fine sheaves

For a Hausdorff space X, a sheaf A is soft when sections over a closed subspace Z, meaning sections of $i^{-1}A$ on Z, extend to X. On paracompact Hausdorff spaces soft sheaves of abelian groups are acyclic for global sections [4, 16]. This proof uses local extension and a locally finite shrinking of a cover to combine extensions over closed sets. In contrast, flabbiness is a lifting property for opens and its proof above uses a maximal open lift. The paracompactness assumption for soft acyclicity supplies the refinement step.

A sheaf of abelian groups is fine if, for each locally finite open cover U_i, the identity is a locally finite sum of endomorphisms ϕ_i supported in U_i. For a paracompact Hausdorff space this gives acyclicity. The usual Čech contraction for a cover admitting these endomorphisms takes a p-cochain c to the $(p-1)$-cochain

$$(hc)_{i_0 \dots i_{p-1}} = \sum_j \phi_j(c_{j i_0 \dots i_{p-1}}).$$

Each summand is initially defined on the overlap with U_j, and its support permits extension by zero to the required intersection. The sum is locally finite. The alternating Čech differential satisfies $dh + hd = \mathrm{id}$ in positive degrees, since the term omitting the j index gives $\sum_j \phi_j = \mathrm{id}$ and the other terms cancel. Refinement by locally finite covers and the comparison of Čech and derived cohomology on paracompact Hausdorff spaces give the acyclicity assertion [4]. On a smooth paracompact manifold, sheaves of smooth modules are fine using smooth partitions of unity.

Calculation 18.3 (A fine resolution on a circle)**.** On a smooth paracompact manifold the de Rham complex resolves the constant real sheaf by the Poincaré lemma. Its terms are fine, so it computes $\mathrm{R}\Gamma(X, \underline{\mathbb{R}})$ via global differential forms. On a circle with periodic coordinate t of period 1, a 1-form is $g(t)\,dt$. It is exact if and only if $\int_0^1 g(t)\,dt = 0$. Necessity follows by integrating a derivative of a periodic function. For sufficiency $t \mapsto \int_0^t g(s)\,ds$ is a periodic primitive. Thus $H^0(S^1; \mathbb{R}) \cong \mathbb{R}$ from the constant

functions and $H^1(S^1;\mathbb{R}) \cong \mathbb{R}$ from the integral. Higher cohomology vanishes because there are no forms in higher degrees. This resolves the constant sheaf using a sheaf of real vector spaces that supports local multiplication by functions.

Exercise 18.4. Show that a skyscraper sheaf $i_{x*}M$ at a point is flabby. Describe the restrictions of its global sections when x is a nonclosed point.

Solution. On U the value is M if $x \in U$ and zero otherwise. For an inclusion the restriction is either the identity or the unique surjection onto zero. This proves flabbiness without a closedness assumption. If x is nonclosed, a point $y \in \overline{\{x\}}$ has every neighborhood containing x; the stalk at y is therefore M. The sheaf can have nonzero stalks at specializations of x. Extension by zero on the complement of a closed subspace in Chapter 16 uses closedness to obtain its stated stalk formula. □

Chapter 19

Čech complexes, totalization, and convergence

19.1 The two-open calculation

For two opens U, V covering X, sheaf descent in a stable coefficient category gives

$$F(X) \longrightarrow F(U) \oplus F(V) \xrightarrow{r_U - r_V} F(U \cap V) \tag{19.1}$$

as a fiber sequence. This follows from the pullback of the two section maps, since a pullback in a stable category is the fiber of their difference on the biproduct. For a derived sheaf F, the terms are its derived sections. The long exact sequence of π_i is Mayer–Vietoris. For an ordinary sheaf, $F(U)$ means $\mathrm{R}\Gamma(U, F)$ in this formula, and the complex of ordinary sections computes the fiber if all three opens are acyclic for the coefficient sheaf.

For a finite open cover $\{U_i\}_{i=1}^{r}$, let $U_I = \bigcap_{i \in I} U_i$ for each nonempty subset I. The diagram of yU_I with inclusion maps has colimit the covering sieve in presheaves of spaces. Evaluated on a test open W, the nonempty subindex of those I for which $W \subseteq U_I$ is the poset of nonempty subsets of $\{i : W \subseteq U_i\}$, with reversed order. It has an initial object and thus contractible groupoid completion. If the set is empty the colimit is empty. This proves the identification with the sieve. Consequently the derived value on X is the limit of the diagram of the $F(U_I)$ with maps by restriction. If every nonempty finite intersection is acyclic for an ordinary coefficient A, this limit is the alternating

complex with term

$$C^s = \prod_{i_0<\cdots<i_s} A(U_{\{i_0,\dots,i_s\}})$$

in cochain degree s, using restriction on each face. One obtains this complex either by a projective resolution of the finite subset diagram as in Chapter 20, or by the alternating augmented chain resolution of the cover. In the latter proof each test open in the sieve gives the augmented chain complex of a simplex on its set of eligible indices. Its contraction proves exactness after sheafification. Applying derived maps into A gives the claimed finite complex, since the intersections are acyclic.

Calculation 19.1 (Failure of the ordinary section complex)**.** Consider the cover of S^2 by $U = S^2 - \{N\}$ and $V = S^2 - \{S\}$, with constant abelian coefficient $A \neq 0$. The opens U, V are contractible, and $U \cap V$ deformation retracts to S^1. The ordinary sections on each nonempty open are A, so the two-open complex has only A in cohomology degree zero. In (19.1) the overlap has a further degree-one cohomology group A, which gives $H^2(S^2, A) = A$ by the boundary. All of these section cohomology calculations agree with singular cohomology on the spaces in question. This example isolates the acyclicity hypothesis in the ordinary Čech calculation.

19.2 Cosimplicial filtration

For a cosimplicial spectrum $C^\bullet$, let $\mathrm{Tot}_s C^\bullet$ be the limit through degree s (equivalently the totalization against the s-skeleton of the cosimplicial simplex). The partial totalizations form a tower whose limit is $\mathrm{Tot}\, C^\bullet$. Let $M^s C$ be the degree-s matching object defined from the codegeneracies to lower degrees, and $N^s C = \mathrm{fib}(C^s \to M^s C)$. There is a fiber sequence

$$\Omega^s N^s C \longrightarrow \mathrm{Tot}_s C^\bullet \longrightarrow \mathrm{Tot}_{s-1} C^\bullet. \tag{19.2}$$

The first step $\mathrm{Tot}_0 = C^0$ uses $\mathrm{Tot}_{-1} = 0$. For $s > 0$, extending the data from lower degrees amounts to assigning the s-simplex with

its boundary given by cofaces and its images under codegeneracies given by the matching object. Taking the fiber over the zero lower data gives a map from $\Delta^s/\partial\Delta^s \simeq S^s$ into the matching fiber, which is $\Omega^s N^s C$. One may carry out this computation in a Reedy fibrant model for the cosimplicial spectrum so that all the pullbacks are the specified homotopy pullbacks. This is the totalization tower of Bousfield and Kan [3].

Applying homotopy groups to (19.2) gives the exact couple of the tower. In the indexing with abutment $\pi_{t-s}\operatorname{Tot} C^\bullet$, the first term is $E_1^{s,t} = \pi_t N^s C$. The split codegeneracy relations identify it with the normalized cochains of the cosimplicial abelian group $\pi_t C^\bullet$. The first differential is the alternating sum of the cofaces: they are the attaching maps for the oriented boundary of the s-simplex. Consequently

$$E_2^{s,t} = H^s(\pi_t C^\bullet), \qquad d_r : E_r^{s,t} \longrightarrow E_r^{s+r,t+r-1}. \tag{19.3}$$

This indexing keeps the degree change of a differential equal to -1 in $t-s$. The spectral sequence is initially associated with the tower. Identification of its stable filtration with the homotopy groups of the totalization uses convergence.

Proposition 19.2. *If there exists N such that every C^s is N-coconnective, then the tower spectral sequence is strongly convergent in each total degree, with a finite filtration on that group.*

Proof. Both the matching object and its fiber are N-coconnective, since that bound is closed under limits. Thus $\Omega^s N^s C$ has no homotopy above $N - s$. For each fixed q, the transitions on π_q are isomorphisms for $s > N - q + 1$ by the long exact sequence of (19.2). The same holds for π_{q+1} with its corresponding bound. The Milnor sequence identifies $\pi_q \operatorname{Tot} C^\bullet$ with the stabilized $\pi_q \operatorname{Tot}_s C^\bullet$. The filtration by kernels of restriction to the partial totalizations therefore terminates in degree q. On each finite stage the exact couple gives the associated graded, and at any bidegree on a fixed total line outgoing and incoming differentials involve only finitely many possible columns because of the uniform bound and $s \geq 0$. This gives the finite filtration and its associated graded from (19.3). □

For a hypercover $U_\bullet \to X$ and a derived sheaf K, apply this tower to the section spectra on U_s (products if U_s is a union of representables). The totalization is $\mathrm{R}\Gamma(X, K)$ by Proposition 17.2. An ordinary sheaf placed in degree zero gives $N = 0$ because derived sections are coconnective for such a sheaf, and products of spectra preserve this bound. The E_2-page may have several nonzero rows indexed by the cohomology of the opens in the hypercover. When every component of every U_s is acyclic, the single row gives the ordinary normalized Čech or hypercover cochain calculation. For a general unbounded K, convergence can be proved by showing stabilization in a given total degree or by a separate cohomological dimension bound. The existence of the totalization by itself gives neither condition.

19.3 A Postnikov calculation of hypercohomology

The tower in the preceding section uses cosimplicial degree. Another tower uses the Postnikov layers of K and the exact functor $F = \mathrm{R}\Gamma(X, -)$. Let $H^q(K)$ denote the ordinary cohomology sheaf of a derived complex. For bounded K, applying F to each successive layer triangle gives the spectral sequence

$$E_2^{p,q} = H^p(X, H^q(K)) \implies H^{p+q}\mathrm{R}\Gamma(X, K).$$

The shift convention is $H^q(K)[-q]$ for the layer. This is the same exact-couple construction as (12.1). Its differentials encode the boundaries of the layer tower. If F has amplitude $[-d, 0]$, Theorem 14.1 shows that to compute H^j of the abutment it suffices to use input cohomology degrees $j - d, \dots, j$, with their connecting data. For a finite range $H^a, \dots, H^b$ in cochain indexing, the input range becomes $a - d, \dots, b$, as follows by replacing homotopy index q with cochain index $-q$. The finite reduction computes the output as a derived object in its range and thus retains group extensions between adjacent filtration pieces. If only individual E_∞ groups are used, these extensions remain to be determined.

Exercise 19.3. For a three-open cover by U_0, U_1, U_2 acyclic on every intersection for A, write the cochain groups and both differentials.

Solution. The groups are $C^0 = \prod_i A(U_i)$, $C^1 = \prod_{i<j} A(U_i \cap U_j)$, and $C^2 = A(U_0 \cap U_1 \cap U_2)$. On $U_i \cap U_j$, $(dc)_{ij} = c_j|_{ij} - c_i|_{ij}$. On the triple intersection, $(dc)_{012} = c_{12}|_{012} - c_{02}|_{012} + c_{01}|_{012}$. The composites vanish by cancellation of each restriction. The finite intersection descent computation identifies the resulting cohomology with the sheaf cohomology. □

Chapter 20

Finite sites and explicit cohomological bounds

20.1 Sheaves on an ordered space

Let P be a finite poset with the Alexandrov topology of upper subsets. Its minimal open at p is $U_p = \{q : p \leq q\}$. Every open covering U_p contains one member containing p, which then contains U_p. Thus the topology restricted to the basis of minimal opens has only maximal covering sieves. Restriction to this basis identifies sheaves with diagrams on P, with $p \leq q$ giving the restriction from U_p to U_q. In stable coefficients D with limits, the inverse takes a diagram $M : P \to D$ to

$$F(U) = \lim_{p \in U} M(p).$$

On U_p the index has initial object p, so the formula gives $M(p)$. To see descent, write each open as a union of its minimal opens, and use the density of this basis: the covering data for a test minimal open is the full nerve of those covering opens that contain it. This index is nonempty for each such test and its augmented Čech nerve is contractible. The limit formula consequently satisfies descent for any cover. For modules the abelian equivalence between sheaves and diagrams is exact, and (17.1) gives its derived form. Global sections in this model is the coherent limit over P.

20.2 A finite projective resolution

For a commutative ordinary ring R, let P_p denote the free R-diagram on the evaluation at p. Its value at q is R if $p \leq q$, and zero otherwise. Maps out of it are identified with evaluation at p on the free rank-one module. The diagrams P_p are projective. Let $\underline{R}$ denote the constant diagram. It represents the ordinary limit because a map from it specifies compatible images of its generators at all p.

Theorem 20.1. *If the longest strict chain in P has $d+1$ vertices, the constant diagram $\underline{R}$ has a projective resolution of length d with term*

$$B_n = \bigoplus_{p_0 < \cdots < p_n} P_{p_n}.$$

For R-module diagrams the derived limit has amplitude bounded by $[-d, 0]$. Its cohomology on an ordinary diagram M is computed by the cochain groups

$$C^n(P; M) = \prod_{p_0 < \cdots < p_n} M(p_n).$$

Proof. For $n \geq 1$, on a summand for $p_0 < \cdots < p_n$, the face d_i, $0 \leq i < n$, omits p_i and uses the identity of P_{p_n}. The face d_n uses the comparison $P_{p_n} \to P_{p_{n-1}}$ induced by $p_{n-1} < p_n$. Thus

$$d\Big|_{B_n} = \sum_{i=0}^{n} (-1)^i d_i : B_n \longrightarrow B_{n-1}.$$

The augmentation $\epsilon : B_0 \to \underline{R}$ uses the canonical map $P_p \to \underline{R}$. At a test q, this becomes the augmented simplicial chain complex of $N(P_{\leq q})$: its basis consists of chains $c = (p_0 < \cdots < p_n \leq q)$. Define

$$h_n(c) = \begin{cases} (-1)^{n+1}(p_0 < \cdots < p_n < q), & p_n < q, \\ 0, & p_n = q, \end{cases}$$

and let $s_{-1} : R \to B_0(q)$ send 1 to (q). If $p_n < q$, the last face of the appended chain has total sign $(-1)^{n+1}(-1)^{n+1} = 1$

and returns c. For $n \geq 1$ and $0 \leq i \leq n$, the other terms of $dh_n(c)$ and $h_{n-1}d(c)$ cancel because their signs are $(-1)^{n+1+i}$ and $(-1)^{n+i}$. If $p_n = q$ and $n \geq 1$, only the last face of c contributes to $h_{n-1}d(c)$, with total sign $(-1)^n(-1)^n = 1$. In degree zero, $dh_0(p) = (p) - (q)$, with value zero when $p = q$. Hence

$$dh_n + h_{n-1}d = \mathrm{id} \quad (n \geq 1), \qquad dh_0 + s_{-1}\epsilon = \mathrm{id} \quad (n = 0),$$

and $\epsilon s_{-1} = \mathrm{id}_R$. The augmented complex is therefore objectwise exact, proving that it is a projective resolution.

Applying maps into M gives the stated cochain groups. For a chain $p_0 < \cdots < p_{n+1}$, the full differential is

$$\begin{aligned}(dc)(p_0, \ldots, p_{n+1}) = &\sum_{i=0}^{n} (-1)^i c(p_0, \ldots, \widehat{p_i}, \ldots, p_{n+1}) \\ &+ (-1)^{n+1} M(p_n \to p_{n+1})\, c(p_0, \ldots, p_n).\end{aligned}$$

The first $n + 1$ terms already take values in $M(p_{n+1})$; the last face requires the displayed structure map.

For a complex of diagrams the same bounded projective resolution computes the derived Hom from $\underline{R}$ and thus the derived limit. Filter the total mapping complex by the finitely many resolution degrees. Its graded terms are finite products of evaluations of the input shifted by $[-n]$ for $0 \leq n \leq d$. For a connective input these lie in degrees at least $-d$, and for a coconnective input in degrees at most zero. Extension closure proves the amplitude bound for all inputs. □

For the empty poset the limit in a stable category is zero, so the cohomological bound is automatic; the theorem above uses nonempty P to specify d. For a poset with an initial object, limit is evaluation there and the bound can be improved to 0. The length d is therefore a uniform bound obtained without testing the shape for a particular simplification.

20.3 Sharp and nonconstant examples

For a constant diagram $\underline{M}$, the differential of $C^\bullet(P;\underline{M})$ is the usual simplicial coboundary of NP, with constant coefficients M. Take P to be the poset of nonempty proper subsets of a set of $d+2$ elements. Its order complex is the barycentric subdivision of the boundary of a $(d+1)$-simplex, which is S^d. The largest chain has $d+1$ vertices. For $d \geq 1$, a nonzero coefficient gives nonzero $H^d(P;\underline{M})$. Therefore d is the best bound on derived limit in this family, and the upper endpoint of the corresponding input window in Theorem 14.1 is sharp. A shifted constant diagram realizes this endpoint in any chosen output degree.

Now take P with three vertices u, v, w, strict relations $u < w$, $v < w$, and no further nonidentity relations. The nerve is contractible because w is terminal. For the diagram $M(u) = M(v) = 0$, $M(w) = A$, the derived limit is the pullback $0 \times_A 0 \simeq A[-1]$ in $\mathrm{D}(R)$. Thus a constant-coefficient test can miss the best bound for all coefficients. The projective resolution computes the same group: $C^0 = A$ and $C^1 = A \oplus A$, with diagonal differential up to the common orientation sign. The kernel is zero and the cokernel is A. This is an example in which the value at a terminal vertex gives no shortcut for the limit.

20.4 Relative limits

Let $f : P \to Q$ be monotone, hence continuous for the upper topologies. For $q \in Q$, put

$$P_q = \{p \in P : q \leq f(p)\}.$$

This is an upper subset of P. The direct image on $U_q \subset Q$ has value $\lim_{p \in P_q} M(p)$, by the formula for the preimage. This is also the comma category q/f for the pointwise right Kan extension. For every q with $P_q \neq \varnothing$, let d_q be its maximal chain length minus one. Then the amplitude of the derived direct image is bounded by $[-\max_q d_q, 0]$ when at least one P_q is nonempty, by

Theorem 20.1 and the pointwise t-structure of the target. Empty fibers contribute zero. The resulting output window estimate is uniform over Q. At a chosen q the potentially smaller d_q controls the section value there.

The ordinary set-theoretic fiber $f^{-1}\{q\}$ can give a different calculation. If $P = \{p\}$, $Q = [1]$, and $f(p) = 1$, then $P_0 = \{p\}$ while $f^{-1}\{0\}$ is empty. The direct image at 0 sees $M(p)$. This is the finite ordered version of a closed-point base-change failure and agrees with the comma-category test in Chapter 4.

20.5 The incidence presentation

For finite P, the direct sum $G = \bigoplus_{p \in P} P_p$ is a compact generator in the stable diagram category. Its endomorphisms are discrete, with matrix entry from P_p to P_q equal to R if $q \leq p$, and zero otherwise. Composition uses the unique comparable chain of morphisms in the poset. Thus Theorem 9.1 identifies the diagram category with derived right modules over the incidence matrix ring. Under this equivalence the module for $\underline{R}$ has one copy of R at each idempotent with all structure maps identities. The resolution of Theorem 20.1 is its finite projective resolution by idempotent summands. In particular the derived limit is a derived Hom from a perfect object. It preserves coproducts in this finite case, a property that can fail for limits over an infinite site.

Exercise 20.2. Let $P = [1]$. Use Theorem 20.1 to compute the derived limit of $M_0 \xrightarrow{u} M_1$ and compare the result with the initial-object formula.

Solution. The two terms are $M_0 \oplus M_1$ and M_1, with differential $(x, y) \mapsto y - u(x)$. The map has section $z \mapsto (0, z)$. Its fiber is M_0, with graph inclusion $x \mapsto (x, u(x))$. This gives the limit even for complexes, by the split sequence of Lemma 6.3, and agrees with evaluation at 0. □

Chapter 21

Local systems and higher monodromy

21.1 Transport along a fundamental groupoid

Let X be a locally weakly contractible space, such as a manifold or a finite simplicial complex with its usual topology. Its fundamental ∞-groupoid $\Pi_\infty X$ is the homotopy type of the singular complex. Let D be a presentable ∞-category. A locally constant D-valued hypersheaf is one that becomes a constant hypersheaf on the members of an open cover. The higher monodromy theorem identifies their category with $\mathrm{Fun}(\Pi_\infty X, D)$ under these hypotheses [15, 13]. In this identification a diagram M gives the section formula

$$F_M(U) \simeq \lim_{\Pi_\infty U} M|_{\Pi_\infty U}. \tag{21.1}$$

The result for space-valued coefficients is a higher form of the classification of locally constant sheaves by covering spaces. For general presentable coefficients the cited theorem uses the hyperdescent version.

Here is the mechanism for the formula on a good open cover, whose nonempty finite intersections are contractible. The homotopy type of X is the colimit of the Čech diagram of the homotopy types of the cover; this is the topological hypercover comparison [9]. Functors out of a colimit of groupoids give the limit of the functor categories by the exponential law and the mapping-space test. On a contractible open a local system is determined by its fiber, and all its mapping data agree with those of the chosen coefficient category. The identifications on overlaps

and their homotopies give an object of the limit, and hence a functor on $\Pi_\infty X$. Taking compatible sections of that functor over U is exactly the limit in (21.1). The theorem extends this computation to the stated locally weakly contractible spaces.

21.2 Constant coefficients and an ordinary case

If $D = \mathrm{Set}$ and X is connected and has a basepoint x, the equivalences in D are ordinary bijections. A locally constant sheaf of sets is specified by its fiber A and an action of $\pi_1(X, x)$ on A. Its global sections are the invariant elements. In fact the functor on the fundamental groupoid factors through the ordinary 1-groupoid, because all transformation spaces in the target have no positive homotopy. The stable case can use higher homotopy of the automorphism space of its fiber.

If $D = \mathrm{D}(R)$ and M is the constant functor at an object E, then the global sections are cotensor by the space $\Pi_\infty X$. For $X = S^d$ with $d \geq 1$, the splitting of the suspension spectrum with a disjoint basepoint gives

$$\mathrm{R}\Gamma(S^d, \Delta_{S^d} E) \simeq E \oplus E[-d].$$

This formula also follows by repeatedly covering the sphere by two contractible opens. The common overlap has the homotopy type of the sphere of one lower dimension, and the two maps from E to its sections have the same constant component. Canceling the split identity component of the difference map in (19.1) gives the desuspended reduced component. The starting sphere S^0 is the two-point product. On an ordinary coefficient $E = A$ this computes H^0 and H^d with value A. On an unbounded E it specifies the whole derived object and shows directly the amplitude bound on these constant systems.

21.3 Classification with a fixed fiber

Let $M \in D$, and let $\mathrm{Aut}_D(M)$ be the union of the components of $\mathrm{Map}_D(M, M)$ that are equivalences. It is a group-like space under composition. The full component of $D^{\simeq}$ containing M is a delooping $B\,\mathrm{Aut}_D(M)$: it is connected, with loop space the automorphism space. A diagram on a connected ∞-groupoid whose fiber is equivalent to M factors through this component. With a specified identification of one fiber, its classifying space is the based mapping space

$$\mathrm{Map}_*(\Pi_\infty X, B\,\mathrm{Aut}_D(M)).$$

Without that identification it is the unbased mapping space. The distinction is the action of automorphisms of the chosen fiber on the based classifications. This follows by taking the fiber of evaluation at x on unbased maps. In the discrete case $\mathrm{Aut}(M)$ is a group and the based description recovers group homomorphisms on π_1.

Calculation 21.1 (A spectral local system on a simply connected space)**.** Let $k = \mathbb{F}_2$, $D = \mathrm{D}(k)$, and $M = k \oplus k[1]$. The only invertible component of $\mathrm{Map}_D(M, M)$ is the identity component, because a degree-zero map is a diagonal pair of scalars and the only nonzero scalar in k is 1. At that component the positive homotopy groups are

$$\pi_1\,\mathrm{Aut}_D(M) = \mathrm{Hom}_{\mathrm{h}\,D}(M[1], M) \cong k, \qquad \pi_{q>1}\,\mathrm{Aut}_D(M) = 0.$$

In degree one the generator is the map from the shifted first summand of M to the second summand, both $k[1]$. Hence $\pi_2 B\,\mathrm{Aut}_D(M) \cong k$. There is a trivial and a nontrivial local system with this fiber on S^2. They remain distinct after forgetting the identification of a chosen fiber, since $\pi_0\,\mathrm{Aut}_D(M)$ is trivial.

Their derived section objects can be compared explicitly. Cover S^2 by the complements of two opposite points. Both opens are contractible and their overlap has homotopy type S^1. Fix a

trivialization on each open. The nontrivial gluing comparison on the overlap is a loop in $\mathrm{Aut}(M)$ representing the nonzero class above. Under the splitting of overlap sections $M \oplus M[-1]$, the two restriction maps from M have the forms $(1, 0)$ and $(1, v)$, with $v : M \to M[-1]$ the adjoint of the loop class $M[1] \to M$. Indeed the value of the gluing at the chosen point is the identity, and its reduced component is the loop in the mapping space. Translating that loop by the identity yields the indicated based map. This calculation is the tensor adjunction of Chapter 2 applied to S^1.

Taking the difference of these maps as in (19.1), the first component splits as a difference of two identity maps. The fiber of the full difference is therefore $\mathrm{fib}(v)$. Write the shifted target explicitly as $M[-1] = k[-1] \oplus k$. The map

$$v : k \oplus k[1] \longrightarrow k[-1] \oplus k$$

is the identity from the first source summand k to the second target summand $k = (k[1])[-1]$, and is zero on all other entries. Its fiber therefore retains the unused source $k[1]$ and the loop $\Omega(k[-1]) = k[-2]$ of the unused target. Thus

$$\mathrm{R}\Gamma(S^2, F_{\mathrm{nontriv}}) \simeq k[1] \oplus k[-2].$$

For the trivial gluing $v = 0$ and the answer is $M \oplus M[-2]$. The two additional groups of the trivial system disappear for the nontrivial system. This is visible in the coherent overlap map although all ordinary monodromy on $\pi_1(S^2)$ is trivial and the two homotopy sheaves of both systems are constant.

21.4 The first obstruction in sections

A coherent section of a local system with fiber M is a point of a limit and can be lifted through the Postnikov tower of the coefficient objects. In a stable coefficient category, a system with two successive homotopy layers $A, B[1]$ is classified by a boundary $A \to B[2]$ in the diagram category. Applying global

sections gives a boundary on the global derived objects. A section of the lower truncation lifts precisely when its image under this boundary is nullhomotopic, and its space of lifts is the fiber of the corresponding map of mapping spaces. The proof is the fiber description of Theorem 10.1 applied to the section functor. In calculation 21.1, the nontrivial loop in the overlap makes this boundary an isomorphism on one pair of groups. It gives a geometric source of a nonzero layer differential even though the stalkwise coefficient complexes split.

Exercise 21.2. Let $X = S^1$ and $E \in \mathrm{D}(k)$, with an automorphism $u : E \to E$. Compute the global sections when $u - 1$ is an equivalence. For $E = k^2$ and $u = \left(\begin{smallmatrix} 1 & 1 \\ 0 & 1 \end{smallmatrix}\right)$, compute the two cohomology groups of its local system.

Solution. The limit over $B\mathbb{Z}$ is $\mathrm{fib}(1 - u)$, so it is zero when $u - 1$ is an equivalence. In the given matrix, $u - 1$ has one-dimensional kernel and one-dimensional cokernel. The fiber has $H^0 \cong k$ from the kernel and $H^1 \cong k$ from the cokernel, with all other groups zero. The ordinary invariant sections compute just the kernel. □

Chapter 22

Group actions and derived invariants

22.1 Two adjoints to a constant action

For a discrete group G and a stable presentable category D, the functor $\Delta : D \to \mathrm{Fun}(BG, D)$ supplies the trivial action. Its left adjoint is $M \mapsto M_{hG} = \mathrm{colim}_{BG} M$ and its right adjoint is $M \mapsto M^{hG} = \lim_{BG} M$. These adjunctions are the universal properties of coherent cocones and cones. When $D = \mathrm{D}(R)$ for a commutative ordinary ring, $\mathrm{Fun}(BG, D)$ identifies with $\mathrm{D}(R[G])$: the ordinary diagram category of R-modules on $\mathbf{B}G$ is the category of left $R[G]$-modules and its derived category identifies with the coherent diagrams by (17.1). Under that identification the functors are

$$M_{hG} \simeq R \otimes^{\mathbf{L}}_{R[G]} M, \qquad M^{hG} \simeq \mathrm{RHom}_{R[G]}(R, M),$$

with R the trivial module. The equalities follow from the tensor and Hom adjunctions, since the derived constant action is restriction along the augmentation. In particular for ordinary modules, homotopy orbits give group homology and homotopy fixed points give group cohomology. The latter is $H^q(G; M) = \pi_{-q} M^{hG}$.

The canonical simplicial free resolution of the trivial module is the bar resolution [5]. In degree n it is free on n-tuples $(g_1, \ldots, g_n)$ of group elements, and the differential takes alternating faces using multiplication of adjacent entries and the action on the first term. After forgetting the G-action the associated augmented chain complex is contractible: the homogeneous model has chains in G^{n+1}, and prepending the identity gives an extra degeneracy. Thus it is a projective resolution. Applying $\mathrm{Hom}(-, M)$

gives the inhomogeneous cochain complex of group cohomology. For spectra one obtains the homotopy fixed-point tower by the corresponding cosimplicial bar of the action, with degree n the product of underlying spectra indexed by G^n. The identities for the action produce its cofaces and codegeneracies. This is a special case of totalization descent over $* \to BG$.

22.2 The cyclic calculation

Let $G = C_m = \langle t \mid t^m = 1 \rangle$, and put $S = R[C_m]$, $N = 1 + t + \cdots + t^{m-1}$. There is a free resolution of the trivial module

$$\cdots \xrightarrow{N} S \xrightarrow{t-1} S \xrightarrow{N} S \xrightarrow{t-1} S \xrightarrow{\epsilon} R \to 0. \tag{22.1}$$

For exactness, write an element of S by its m coefficients. The kernel of $t-1$ consists of constant coefficient vectors, and every such vector is in the image of N. The kernel of N consists of vectors whose coefficient sum is zero. It equals the image of $t-1$, since any vector of sum zero is a cyclic difference of an m-tuple, obtained by successive partial sums. It is also the kernel of augmentation. The two maps compose to zero because $(t-1)N = t^m - 1 = 0$. Thus (22.1) is exact for any commutative R.

For an ordinary left module M, the cochain differential begins with $t - 1 : M \to M$, then $N : M \to M$, and alternates. Hence

$$\begin{aligned} H^0(C_m; M) &= \ker(t-1), \\ H^{2i+1}(C_m; M) &= \ker N/(t-1)M \qquad (i \geq 0), \\ H^{2i}(C_m; M) &= \ker(t-1)/NM \qquad (i \geq 1). \end{aligned}$$

For the trivial module $M = \mathbb{Z}$, these give $\mathbb{Z}$ in degree zero, zero in all positive odd degrees, and $\mathbb{Z}/m$ in positive even degrees. For the sign action of C_2 on $\mathbb{Z}$, $t - 1 = -2$ and $N = 0$. The degree-zero and positive even groups are zero, and every positive odd group is $\mathbb{Z}/2$. The same cyclic group thus gives two patterns depending on the coefficient transport.

When $m = p$ is prime and $R = M = \mathbb{F}_p$ with trivial action, both differentials vanish and every H^q is $\mathbb{F}_p$. Consequently the right adjoint $M \mapsto M^{hC_p}$ has no uniform lower bound on the images of connective objects. In particular there is no finite input-window formula of the form (14.1) for this right adjoint. The obstruction appears on one object of its heart.

22.3 Averaging and its derived meaning

If G is finite and $|G|$ is invertible in R, the idempotent

$$e = \frac{1}{|G|} \sum_{g \in G} g$$

acts on M with image the invariants. Multiplication on both sides by each g leaves e unchanged. Thus the augmentation $R[G] \to R$ has an $R[G]$-linear section $1 \mapsto e$, making the trivial module projective. It is also a projective right module by the same identity. Derived invariants and derived coinvariants on complexes are calculated termwise. The norm, normalized by $|G|$, gives the corresponding isomorphism of underived invariants and coinvariants: the class of m maps to em and an invariant x maps back to its class. These inverse identities follow from the action relations. The map extends to complexes termwise and provides the coherent equivalence between the two derived functors. The bound on the right adjoint is $d = 0$ in this case.

The projectivity argument uses the coefficient ring and cannot be inferred merely from finiteness of the category $\mathbf{B}G$. Its nerve has simplices in all degrees, and (22.1) exhibits the positive cohomology that may remain. Finite posets and finite groups therefore provide different geometric tests for cohomological amplitude.

22.4 Free groups and a finite cellular model

For a free group $G = F_r$ on $g_1, \dots, g_r$ with $r \geq 1$, the universal cover of the wedge of r circles is a tree. It has free G-orbits consisting of one vertex orbit and r oriented edge orbits. Its augmented cellular chains give the projective resolution

$$0 \longrightarrow R[G]^r \xrightarrow{(g_i - 1)_i} R[G] \longrightarrow R \longrightarrow 0.$$

Contractibility of the tree proves exactness of the augmented chains. Applying the mapping functor into an action gives a two-term derived complex whose underlying arrow is $M \to M^r$ with coordinates $g_i - 1$. On ordinary modules, H^0 is its kernel, H^1 its cokernel, and the higher groups vanish. For complexes, this resolution gives a fiber formula and an amplitude bound $[-1, 0]$. A trivial nonzero module has $H^1 = M^r$, so the bound is sharp. This recovers the $\mathbb{Z}$-action calculation when $r = 1$.

22.5 Comparison with section cohomology

The formula for sections of a local system on X uses the homotopy type $\Pi_\infty X$. If X is connected and aspherical, this homotopy type is $B\pi_1 X$, and the formulas above apply directly to the diagram of its fibers. For $X = S^2$ the fundamental group is trivial and the space is not aspherical. The cohomology of a constant coefficient has a degree-two term, as in Chapter 21. The map to $B\pi_1 X$ forgets the cell whose cochain produces that term. These two computations show precisely when derived group invariants model global sections of local systems.

Exercise 22.1. Let M be the regular module $R[C_m]$ for the cyclic group. Verify from the periodic formulas that $H^{q>0}(C_m; M) = 0$, even when m is zero in R.

Solution. For the regular action $\ker(t-1) = RM_0$ with $M_0 = N$, and the image NM is the same set of constant coefficient vectors. The kernel of N is the image of $t-1$, by the partial-sum calculation in (22.1). Both positive-degree quotients vanish. This is acyclicity of a specific coefficient module and holds independently of the amplitude on all modules. □

Chapter 23

A ringed calculation on the projective line

23.1 Two affines and the transition function

Let A be a commutative ring and $X = \mathbb{P}^1_A$. Write the homogeneous coordinates as x_0, x_1, and use $U_0 = D_+(x_0)$, $U_1 = D_+(x_1)$. They are affine with rings $A[t]$ and $A[t^{-1}]$, respectively, where $t = x_1/x_0$. The overlap has ring $A[t, t^{-1}]$. The line bundle $\mathcal{O}_X(n)$ is trivial on each open with local frames x_0^n, x_1^n, and the second frame is t^n times the first on the overlap. This comparison specifies the gluing by an invertible function. For $n < 0$ the frame notation refers to powers of the dual line, and the same comparison holds.

Each of the two opens and their overlap is affine. Quasi-coherent modules on affines have no higher cohomology [35]. Thus Chapter 19 gives a two-term complex in degrees $0, 1$ for $\mathrm{R}\Gamma(X, \mathcal{O}_X(n))$:

$$A[t] \oplus A[t^{-1}] \xrightarrow{d_n} A[t, t^{-1}], \qquad d_n(a, b) = a - t^n b. \tag{23.1}$$

The sign is the difference of restrictions in the first trivialization. This is an equivalence in $\mathrm{D}(A)$, not just a computation of dimensions. It follows from the pullback of the derived section objects, all three of which are degree-zero modules here.

Calculation 23.1 (Cohomology over the base ring)**.** Every Laurent monomial t^i with $i \geq 0$ is in the image of the first component of (23.1), and every t^i with $i \leq n$ is in the image of the second.

Monomials with $0 \le i \le n$ occur in both images and each gives a kernel summand of rank one. Monomials with $n < i < 0$ occur in neither image and each gives a cokernel summand. The other exponents give isomorphism summands. Therefore

$$H^0(X,\mathcal{O}(n)) \cong \begin{cases} A^{n+1}, & n \ge 0, \\ 0, & n < 0, \end{cases}$$

$$H^1(X,\mathcal{O}(n)) \cong \begin{cases} A^{-n-1}, & n \le -2, \\ 0, & n \ge -1. \end{cases}$$

The bases are the common exponents $0,\dots,n$ and the missing exponents $n+1,\dots,-1$, respectively. Higher cohomology vanishes. This is the case of the general projective-space formula in [36]. In (23.1), the splitting by Laurent exponent explicitly removes acyclic summands and gives the indicated free cohomology module in its degree.

For clarity about the splitting, fix i. The term in degree 1 is At^i. The degree-zero term has one generator if exactly one inequality $i \ge 0, i \le n$ holds, two if both hold, and zero if neither holds. In the first case the differential is 1 or -1, and the summand contracts. In the second it is the row $(1,-1)$, which splits onto the target and has diagonal kernel. In the third the summand is one copy of A in degree 1. Taking the direct sum of these identifications defines a quasi-isomorphism and even a chain homotopy equivalence to a bounded finite free complex. The maps use only $0, 1, -1$ and powers of the chosen variable, so they are compatible with a change of A.

23.2 Base change without flatness of the base map

Let $A \to B$ be any homomorphism of commutative rings and $X_B = \mathbb{P}^1_B$. There is a comparison

$$B \otimes^{\mathbf{L}}_A \mathrm{R}\Gamma(X,\mathcal{O}_X(n)) \longrightarrow \mathrm{R}\Gamma(X_B,\mathcal{O}_{X_B}(n)). \tag{23.2}$$

Using (23.1), its source is computed by tensoring both terms with B because the terms are free A-modules and the complex is bounded. The tensor yields the corresponding polynomial and Laurent polynomial modules over B, with the same transition function t^n. On the base-changed affine cover this complex is the target by the same descent computation. The comparison agrees on the affine terms with the usual tensor map and on their difference with the induced map of fibers, so (23.2) is an equivalence. The proof works for an arbitrary base ring map because it uses the flatness of the displayed model over A. The derived inverse image of the line bundle itself is the usual line bundle on X_B, since it is locally free over the source structure sheaf.

23.3 A window for a quasi-coherent complex

We work on the Zariski site. Write $\mathrm{D}(\mathcal{O}_X)$ for the derived category of all sheaves of $\mathcal{O}_X$-modules, and $\mathrm{D}_{\mathrm{QCoh}}(\mathcal{O}_X)$ for its full stable subcategory on complexes with quasi-coherent cohomology sheaves. Its t-structure is the restriction of the sheaf t-structure, since truncation preserves quasi-coherent cohomology. On an affine scheme $\operatorname{Spec} R$, derived global sections give an equivalence

$$\mathrm{R}\Gamma(\operatorname{Spec} R, -) : \mathrm{D}_{\mathrm{QCoh}}(\mathcal{O}_{\operatorname{Spec} R}) \simeq \mathrm{D}(R).$$

This includes unbounded complexes and is t-exact. The precise input is [37, Lemma 36.3.5]; its proof uses affine vanishing for quasi-coherent sheaves [35] together with the unbounded truncation bounds of Lemma 36.3.4 in the same section.

For $X = \mathbb{P}^1_A$, the standard opens U_0, U_1 and their intersection are affine. These open immersions are quasi-compact, since X is separated; restriction is exact and preserves quasi-coherent cohomology. For $K \in \mathrm{D}_{\mathrm{QCoh}}(\mathcal{O}_X)$, the two-open descent sequence identifies $\mathrm{R}\Gamma(X, K)$ with the fiber of the difference map from the

two affine section objects to the overlap section object. Thus the unbounded affine theorem applies to each of the three terms.

If K is 0-connective in homotopy indexing, all three affine section objects are 0-connective. Their fiber is (-1)-connective by the exact sequence. If K is 0-coconnective, all three are 0-coconnective, and the fiber is 0-coconnective by closure under limits. Hence $\mathrm{R}\Gamma(X,-)$ on $\mathrm{D}_{\mathrm{QCoh}}(\mathcal{O}_X)$ has amplitude bounded by $[-1,0]$. For $A \neq 0$, $\mathcal{O}(-2)$ has $\pi_{-1}\mathrm{R}\Gamma(X,\mathcal{O}(-2)) \cong A$, so the bound is sharp. The upper endpoint is attained by $\mathcal{O}$.

With cochain indexing, the cohomology in degree j of $\mathrm{R}\Gamma(X,K)$ is determined by the two input degrees $j-1, j$. In this special amplitude one gets a short exact sequence

$$0 \to H^1(X, H^{j-1}K) \to H^j\mathrm{R}\Gamma(X,K) \to H^0(X,H^jK) \to 0. \tag{23.3}$$

To prove it, apply the window equivalence for output degree $-j$ and use the two-layer cofiber sequence in homotopy indexing

$$(H^{j-1}K)[-j+1] \longrightarrow \tau_{[-j,-j+1]}K \longrightarrow (H^jK)[-j].$$

After applying $\mathrm{R}\Gamma(X,-)$, the neighboring terms in the long exact sequence are $H^{-1}(X,H^jK)=0$ and $H^2(X,H^{j-1}K)=0$. This gives (23.3) with its stated arrow directions. The groups at the ends are explicit derived-section groups of ordinary quasi-coherent sheaves. The short exact sequence is functorial and its class can retain the connecting information of K.

23.4 A boundary realized by a line bundle

The Euler sequence on $\mathbb{P}^1_A$ in the required twist takes the form

$$0 \longrightarrow \mathcal{O}(-2) \longrightarrow \mathcal{O}(-1)^{\oplus 2} \xrightarrow{(x_0,x_1)} \mathcal{O} \longrightarrow 0, \tag{23.4}$$

with first map the column $(-x_1, x_0)$. Exactness is checked on each of the two affines. On U_0 the row is $(1,t)$, and its kernel

is generated by $(-t, 1)$; on U_1 the row is $(t^{-1}, 1)$ with the corresponding generator. The comparison of these kernels is the transition of $\mathcal{O}(-2)$, so the sequence is exact globally. Apply derived sections. The middle term is acyclic with zero derived value by calculation 23.1. The boundary $H^0(X, \mathcal{O}) \to H^1(X, \mathcal{O}(-2))$ is consequently an isomorphism $A \to A$. For $A \neq 0$ the extension is nonsplit, since a splitting would make this boundary zero. This gives a concrete instance of the identification of boundaries with extension classes in Chapters 10 and 11. Its nonzero value is detected by one step of Čech descent.

Exercise 23.2. Track the boundary in (23.4) using the cover. Which missing exponent in (23.1) for $n = -2$ represents the boundary of the section 1?

Solution. Choose the local lifts $(x_0^{-1}, 0)$ on U_0 and $(0, x_1^{-1})$ on U_1. On the overlap their difference lies in the kernel of the row. Express it as $(-x_1, x_0)$ times a section of $\mathcal{O}(-2)$. The coefficient in the frame x_0^{-2} is $-t^{-1}$ for the stated difference of lifts. Up to the chosen sign it is the only missing exponent -1. This gives the generator of $H^1(X, \mathcal{O}(-2))$, agreeing with the derived boundary. □

Chapter 24

Tensor, projection, and a duality calculation

24.1 The projection comparison

Let $f : (X, \mathcal{O}_X) \to (Y, \mathcal{O}_Y)$ be a morphism of ringed spaces with commutative rings. The derived inverse image $\mathrm{L}f^*$ is symmetric monoidal and has right adjoint $\mathrm{R}f_*$. Write $L = \mathrm{L}f^*$ and $R_f = \mathrm{R}f_*$ in this chapter. A right adjoint to a strong monoidal functor is lax monoidal: the comparison $R_f E \otimes R_f F \to R_f(E \otimes F)$ is adjoint to the tensor of the two counits. The projection comparison

$$R_f(E) \otimes M \longrightarrow R_f(E \otimes LM) \tag{24.1}$$

is adjoint to $LR_f(E) \otimes LM \to E \otimes LM$, again using the counit. All tensors in this chapter are derived. The definition uses no finite-dimensional assumption.

An object M is dualizable if it has a dual $M^\vee$, evaluation $M^\vee \otimes M \to \mathbb{1}$, and coevaluation $\mathbb{1} \to M \otimes M^\vee$ satisfying the triangle identities. For such M, tensoring with M is both a left and a right adjoint to tensoring with $M^\vee$ in the corresponding orientations. Perfect complexes of sheaves are dualizable in the derived tensor category: locally choose a bounded finite projective complex and its dual complex, and use the evaluation and coevaluation for those finite projectives. The canonical internal derived Hom into the unit gives the dual globally because the comparisons are local equivalences. A strong monoidal functor preserves this dual data.

Proposition 24.1 (Projection for a dualizable coefficient)**.** *If M is dualizable on Y,* (24.1) *is an equivalence for every E.*

Proof. For a test object T on Y, use the duality and adjunction to obtain

$$\begin{aligned}\operatorname{Map}(T, R_f E \otimes M) &\simeq \operatorname{Map}(T \otimes M^\vee, R_f E)\\ &\simeq \operatorname{Map}(LT \otimes L(M^\vee), E)\\ &\simeq \operatorname{Map}(LT, E \otimes LM)\\ &\simeq \operatorname{Map}(T, R_f(E \otimes LM)).\end{aligned}$$

Tracing the equivalences sends a map $T \to R_f E \otimes M$ to its composite with (24.1). This follows by expanding the two duality adjunctions, whose consecutive evaluation and coevaluation cancel by the triangle identity. The equivalence is natural in T, so Yoneda proves the claim. □

For example, for $p : X \to \operatorname{Spec} A$ and a perfect A-complex P, this identifies $\mathrm{R}\Gamma(X, E) \otimes^{\mathbf{L}}_A P$ with $\mathrm{R}\Gamma(X, E \otimes^{\mathbf{L}}_A p^* P)$ in the quasi-coherent situation. In the two-affine computation of Chapter 23 one can also see it by applying the exact tensor functor to a finite fiber diagram. This direct proof extends to nonperfect P when all affine terms commute with the base change in question. The dualizability proof is independent of a finite cover and therefore applies to arbitrary E and f in its setting.

Calculation 24.2 (An infinite product test)**.** Take X to be a countably infinite discrete space and $f : X \to *$, with a nonzero commutative coefficient ring A. Let E have value A at every point, and let $M = \bigoplus_{j \geq 0} A$ be a free module of countable rank. The projection comparison in degree zero is

$$\Big(\prod_{i \geq 0} A\Big) \otimes_A \Big(\bigoplus_{j \geq 0} A\Big) \longrightarrow \prod_{i \geq 0} \Big(\bigoplus_{j \geq 0} A\Big).$$

All modules in the displayed tensor with M can be computed underived because M is free, and products in A-modules are exact. The image consists of families whose entries have a common finite set of supported j-coordinates. The family whose ith entry is the basis vector at $j = i$ is outside the image. Thus this case of (24.1) fails. The calculation identifies the obstruction as the difference between an infinite product followed by a coproduct and the reverse order.

24.2 A trace on the projective line

Return to $X = \mathbb{P}^1_A$. The class of t^{-1} in the cokernel of d_{-2} in (23.1) is a basis of $H^1(X, \mathcal{O}(-2))$. Use it to fix the trace

$$\mathrm{tr} : \mathrm{R}\Gamma(X, \mathcal{O}(-2)) \xrightarrow{\simeq} A[-1]$$

that sends this class to 1 in cochain degree 1. The product of line bundles $\mathcal{O}(n) \otimes \mathcal{O}(-n-2) \simeq \mathcal{O}(-2)$ and lax monoidality of sections give a pairing

$$\mathrm{R}\Gamma(X, \mathcal{O}(n)) \otimes^{\mathbf{L}}_A \mathrm{R}\Gamma(X, \mathcal{O}(-n-2)) \longrightarrow A[-1]. \tag{24.2}$$

It can be computed on the Čech complexes: multiply a global section on the first factor by a representative of an overlap 1-cochain on the second, using the first trivialization. Boundaries map to boundaries because the first factor is a matching pair of local sections. Thus the operation agrees with multiplication in derived sections.

Theorem 24.3 (Duality for the twists on $\mathbb{P}^1$)**.** *The adjoint of* (24.2) *induces a natural equivalence*

$$\mathrm{R}\Gamma(X, \mathcal{O}(-n-2))[1] \simeq \mathrm{RHom}_A(\mathrm{R}\Gamma(X, \mathcal{O}(n)), A).$$

With the specified trace the equivalence commutes with arbitrary changes of A.

Proof. For $n \geq 0$, the first factor has its only group in degree zero, with basis t^i for $0 \leq i \leq n$. The second factor has its only group in degree one, with basis of missing exponents t^j for $-n-1 \leq j \leq -1$. The trace of $t^i t^j$ is 1 when $i + j = -1$ and zero otherwise, since all other exponents are boundaries for the twist -2. Thus the pairing matrix is a permutation matrix. The adjoint is an isomorphism of the only possible free cohomology modules and hence an equivalence. For $n = -1$ both objects are zero. For $n \leq -2$, exchange the factors and use the same computation for $-n-2 \geq 0$. The derived symmetry has no extra sign on these two ordinary cohomology groups because their

degrees are $0, 1$. This proves the equivalence in all cases. The pairing is defined by multiplication and by the Laurent monomial t^{-1}; both are preserved by base change. The two complexes are perfect, and duality of finite projective complexes commutes with derived extension of scalars. With (23.2) this gives compatibility over B for every $A \to B$. □

This calculation is the line-bundle instance of Serre duality on a relative projective line, with $\mathcal{O}(-2)$ the relative dualizing line bundle. The direct proof identifies the trace by a single exponent and specifies the derived degree shift. For instance, replacing $A[-1]$ in the trace by A in degree zero would make its map on the computed H^1 group zero. The derived shift is therefore part of the universal pairing.

24.3 A finite-base-change criterion

The proof of (23.2) extends to a finite affine open cover of an A-scheme for which every nonempty intersection is affine. Let K have quasi-coherent cohomology. Suppose that on every intersection the derived section object commutes with a given base change $A \to B$, with the comparisons compatible on restrictions. Then the global base-change comparison is an equivalence. In fact the section object is the finite limit over nonempty intersections by Chapter 19. Extension of scalars $B \otimes^{\mathbf{L}}_{A} -$ is exact and thus commutes with finite limits. Applying it to the intersection diagram and using the assumed equivalences gives the diagram for the base-changed scheme and its sections. This criterion isolates the calculation on affines. It applies in particular when their section objects are flat models with the stated identification after base change. For infinite covers one would also have to check that tensor preserves the totalization used, a condition that cannot be deduced from exactness.

Exercise 24.4. For $n = 2$, write the matrix of the pairing (24.2) in the bases $1, t, t^2$ and t^{-3}, t^{-2}, t^{-1}. What does it become after a quotient $A \to A/I$?

Solution. The matrix has ones on the anti-diagonal and zeros elsewhere, since $t^i t^j = t^{-1}$ exactly when $j = -i - 1$. Reduction to A/I gives the same anti-diagonal matrix over the quotient. It remains invertible, reflecting the base-change compatibility proved by the derived complexes. $\square$

Chapter 25

Torsors, delooping, and connecting classes

25.1 Classifying a local trivialization problem

Let $\mathcal{X}$ be an ∞-topos and G a group object. Its bar construction has n-simplices G^n, with the usual multiplication and unit maps, and its realization is a pointed connected object BG. The canonical map $1 \to BG$ is an effective epimorphism, and its Čech nerve is the bar construction by effectivity. In particular $1 \times_{BG} 1 \simeq G$. For $T \in \mathcal{X}$, a G-torsor over T is an object $P \to T$ with a G-action over T such that $P \to T$ is an effective epimorphism and the comparison $P \times G \to P \times_T P$ is an equivalence. The latter map takes a point and a group element to the point and its translate, with higher coherence in the action [23].

Proposition 25.1 (Torsor classification)**.** *Pullback of* $1 \to BG$ *along* $T \to BG$ *identifies the space of* G*-torsors over* T *with* $\mathrm{Map}_{\mathcal{X}}(T, BG)$*.*

Proof. Pullback preserves effective epimorphisms and their Čech nerves, so it gives a torsor with the stated action. Conversely, for a torsor the Čech nerve of $P \to T$ is $P \times G^{\bullet}$. This follows in degree one by the torsor comparison, and in all degrees by the Segal description as iterated fiber products of the degree-one arrows. Forgetting the point of P gives a map of simplicial objects $P \times G^{\bullet} \to G^{\bullet}$, and its realization is $T \to BG$. The square with P over 1 in degree zero becomes a pullback after pulling back to $P \to T$, where the torsor is trivialized by its diagonal point.

Pullback along an effective epimorphism detects equivalences by slice descent, so the square is a pullback. The constructions and these comparisons are functorial in torsor equivalences and in homotopies of the classifying map. They are therefore inverse equivalences on spaces. □

For a sheaf of ordinary groups, the components give the classical pointed set $H^1(T, G)$. A choice of local trivializations of a torsor yields transition functions g_{ij} on overlaps, satisfying $g_{ij}g_{jk} = g_{ik}$. Changing the trivializations by h_i gives the corresponding gauge transformations of g_{ij}. The ordinary groupoid of cocycles and gauge maps is the restriction of Proposition 25.1 to such a cover. Its automorphisms can be nontrivial even for a single cohomology class. For an abelian group object the operation of tensoring torsors is induced by the addition on BG, and H^1 is a group.

Calculation 25.2 (Automorphisms of a torsor on a circle)**.** Take $T = S^1$ in spaces and G a discrete group. The free loop space $\mathrm{Map}(S^1, BG)$ fits into the evaluation fibration with fiber G at a point. The monodromy of the evaluation fibration along $h \in G$ conjugates g to hgh^{-1}, by transport of a based loop along h. Therefore the free loop space is the groupoid completion of the conjugation action groupoid. It decomposes as

$$\mathrm{Map}(S^1, BG) \simeq \coprod_{[g]} B\, C_G(g),$$

where $[g]$ ranges over conjugacy classes and $C_G(g)$ is the centralizer. For $G = S_3$ the three components correspond to the identity, transpositions, and 3-cycles. The automorphism groups are S_3, C_2, and C_3, respectively. This calculation also classifies locally constant G-torsors on the ordinary circle via Chapter 21.

25.2 Abelian coefficients as infinite loop objects

For a sheaf of abelian groups A on a small ordinary site, its degree-zero derived sheaf determines a spectrum object. In the model of

Chapter 17 its value at U is the spectrum underlying $\mathrm{R}\Gamma(U, A)$. For $n \geq 0$ let $K(A, n)$ be its nth infinite loop sheaf, with value $\Omega^{\infty}\big(\mathrm{R}\Gamma(U, A)[n]\big)$. It is n-truncated: for $q > n$ the groups of a section space vanish because the derived sections of a heart object are coconnective. The associated homotopy sheaves are A in degree n and zero in the other positive degrees, with terminal component sheaf for $n > 0$ and component sheaf A for $n = 0$. Indeed the higher derived section presheaves in positive cochain degrees become zero upon sheafification, by applying the counit of Theorem 17.1. The loop comparison $\Omega K(A, n+1) \simeq K(A, n)$ follows from the stable shift. These are the Eilenberg–Mac Lane objects of the sheaf ∞-topos.

When the site has terminal object X, the section-space description gives

$$\pi_i \operatorname{Map}_{\mathcal{X}}(1, K(A, n)) \cong H^{n-i}(X, A), \qquad 0 \leq i \leq n. \tag{25.1}$$

The basepoint is the zero section for $i > 0$. The general slice version uses T and the pullback of A to the slice. At $n = 1$, $K(A, 1) = BA$ and (25.1) agrees with Proposition 25.1. At $n = 0$, it recovers ordinary sections. At $n = 2$, the space has the degree-two classes, degree-one comparisons, and degree-zero automorphisms appearing one dimension higher. Taking only components recovers the ordinary cohomology groups in each degree.

25.3 An obstruction is a fiber of a map

A short exact sequence $0 \to A \to E \to B \to 0$ gives a map $B \to A[1]$ in the derived category. After applying $K(-, n)$ in its spectrum construction one gets a fiber sequence of pointed sheaves

$$K(E, n) \longrightarrow K(B, n) \longrightarrow K(A, n+1).$$

The functor taking the nth infinite loop sheaf preserves limits, which gives the sequence and its basepoint. Hence a class in $H^n(X, B)$ lifts precisely when its boundary in $H^{n+1}(X, A)$ is zero.

With a chosen representative of the class, the space of lifts is the space of nullhomotopies of its composite to $K(A, n+1)$ on global sections. If that space is nonempty it is a torsor for the loop space of the target mapping space. In particular its component-level choices are acted on by $H^n(X, A)$. This is the space-level version of the connecting homomorphism in the classical cohomology sequence. The same argument works for any map of coefficients whose desired lifts are modeled by a homotopy fiber.

Calculation 25.3 (The cyclic Bockstein)**.** Work in the topos of C_m-sets, or its higher version of C_m-spaces, with the trivial abelian coefficient sequence

$$0 \longrightarrow \mathbb{Z} \xrightarrow{m} \mathbb{Z} \longrightarrow \mathbb{Z}/m \longrightarrow 0.$$

Derived global sections on module objects here are the invariants of Chapter 22. By the cyclic calculation, $H^1(C_m; \mathbb{Z}) = 0$ and multiplication by m on $H^2(C_m; \mathbb{Z}) \cong \mathbb{Z}/m$ is zero. Thus the boundary $H^1(C_m; \mathbb{Z}/m) \to H^2(C_m; \mathbb{Z})$ is an isomorphism. In the normalized cochain model, take the generator $c(t^i) = i \bmod m$, $0 \le i < m$, and lift it to $\widetilde{c}(t^i) = i$. The coboundary is

$$d\widetilde{c}(t^i, t^j) = i + j - ((i + j) \bmod m).$$

Dividing by m gives the integer-valued 2-cocycle of the carry in adding i, j. This represents the boundary class. It is nonzero by the isomorphism just proved. The fiber description thus detects failure of the degree-one class to lift to an integral degree-one class. The cochain and derived extension constructions give the same obstruction.

25.4 An internal viewpoint

The universal torsor $1 \to BG$ is locally trivial in the sense of the effective epimorphism itself. Given $T \to BG$, the comparison with a chosen trivial torsor is the space of sections of $P \to T$, obtained by pulling back the point $1 \to BG$. Its local existence

is the relative (-1)-truncation of that object, while a global trivialization is a point of its section space. The image factorization of Chapter 5 distinguishes these two tests. It is possible for local existence to hold over all of T while the global section space is empty, as happens for a nontrivial ordinary covering torsor. The same formulation applies to nullhomotopies of higher classifying maps, whose fibers can be higher spaces rather than sets.

Exercise 25.4. Let A be an abelian group and $X = S^d$ with $d \geq 1$. Compute the homotopy groups of $\mathrm{Map}(X, K(A, n))$ for $n \geq d$ and compare the groups in degree $n - d$ and degree n.

Solution. By the splitting of constant section spectra on S^d, the space is $K(A, n) \times K(A, n - d)$, using a chosen basepoint of the sphere. The degree-n group comes from the ordinary sections of A, and the degree-$(n - d)$ group comes from the top cohomology of the sphere. If $n = d$, the latter group indexes the connected components. This is (25.1) with the cohomology of the sphere. □

Chapter 26

Internal quantification and the shape of a topos

26.1 Slices as contexts

In the internal language of an ∞-topos $\mathcal{X}$, a context is an object T, and a family over that context is an object of $\mathcal{X}_{/T}$. Substitution along a map $f : S \to T$ is pullback. Composition of substitutions has the coherent comparison of iterated pullback. For an object $E \to S$, the dependent sum $\sum_f E$ and dependent product $\prod_f E$ over T are the adjoints $f_!$ and f_* of Chapter 4. The first keeps the total object and composes its structure map with f. The second represents coherent sections of the family along fibers of f. Its universal property is

$$\mathrm{Map}_{/T}(W, \prod_f E) \simeq \mathrm{Map}_{/S}(W \times_T S, E).$$

This formula is valid for arbitrary test contexts $W \to T$. It encodes how a term of the dependent product can be evaluated after any substitution. The slice Beck–Chevalley proof in Chapter 16 is precisely the compatibility of this construction with substitution.

A subobject $A \hookrightarrow S$ is a proposition in context S: its fibers are empty or contractible. The existential quantification along f is the image of $\sum_f A \to T$, or relative (-1)-truncation of that sum. The universal quantification is $\prod_f A \to T$. It is already a subobject, since a right adjoint preserves the pullback and diagonal that test monomorphisms. For a subobject $B \hookrightarrow T$, there are equivalences of the corresponding conditions

$$\exists_f A \subseteq B \iff A \subseteq f^* B, \qquad f^* B \subseteq A \iff B \subseteq \forall_f A.$$

The first uses image minimality and the sum adjunction; the second is the product adjunction restricted to subobjects. Thus the ordinary logic of subobjects is the proposition-valued part of the slice operations. The comparisons on arbitrary families remain spaces of maps.

Example 26.1 (Local existence and a section)**.** Let $\mathcal{X} = \mathrm{Sh}(S^1; \mathcal{S})$ and let P be the sheaf of local lifts into the twofold covering $S^1 \to S^1$. Its map to 1 has full image because the covering admits local lifts. A global section would give a continuous section of the covering by the germ-space correspondence. On fundamental groups this would right-invert multiplication by $2 : \mathbb{Z} \to \mathbb{Z}$, which is impossible. Thus the quantified proposition of existence holds locally over all the base, while the section test fails. The classifying map of Chapter 25 records the obstruction.

The diagonal of a family encodes equality in the internal language. For two sections x, y over T, pulling the relative diagonal back along (x, y) gives the family of paths between them. Its sections are the homotopies between x, y over T. Iterated diagonals similarly encode higher equality. This is the categorical form of identity types in homotopy type theory [39]. For a (-1)-truncated family equality of inhabitants is contractible when they exist. For a locally constant family $K(A, 1)$, the equality of a section with itself has a nontrivial local loop group A. The distinction is already present before global sections are taken.

26.2 What constant coefficients see

Let $p : \mathcal{X} \to \mathcal{S}$ be the terminal geometric morphism, with constant object functor p^* and global sections p_*. The shape of $\mathcal{X}$ is the pro-space that represents the functor

$$K \longmapsto p_* p^* K.$$

Here pro-representation means a cofiltered diagram of spaces S_i with $p_* p^* K \simeq \mathrm{colim}_i \mathrm{Map}(S_i, K)$, with the colimit over the

opposite of the diagram index. The general theory and existence are developed in [18, §7.1.6]. When this functor is represented by one space S, the shape can be taken to be S. The section functor for arbitrary coefficients still depends on the coefficient object. Shape gives the calculation on constant coefficients and on the corresponding constant Eilenberg–Mac Lane objects.

For $\mathcal{X} = \mathcal{P}(C)$, the inverse image is the constant presheaf because it preserves colimits and finite limits pointwise, and its right adjoint is limit over C^{op}. The constant diagram test gives

$$p_* p^* K \simeq \mathrm{Map}_{\mathcal{S}}(|C^{\mathrm{op}}|, K) \simeq \mathrm{Map}_{\mathcal{S}}(|C|, K).$$

The opposite has the same groupoid completion up to the canonical inversion equivalence in spaces. Thus the shape of a finite poset topos is its order complex, including the sphere examples of Chapter 20. The 0-truncated version of the calculation sees only its connected components, while constant $K(A, n)$ probes higher cohomology.

26.3 A shape that requires finite quotients

Let $X = \{0, 1\}^{\mathbb{N}}$ with the product topology. Its clopen cylinders form a basis, and every clopen subset is compact. Let $\mathcal{P}$ range over finite clopen partitions of X, ordered by refinement. Write $Q_{\mathcal{P}}$ for the finite set of parts, with a map $Q_{\mathcal{P}'} \to Q_{\mathcal{P}}$ for each refinement. Then for a space K the constant-object section functor is

$$\Gamma(X, p^* K) \simeq \operatorname*{colim}_{\mathcal{P}} \mathrm{Map}(Q_{\mathcal{P}}, K). \tag{26.1}$$

To verify it, work on the basis of compact clopens and define the same formula for each such clopen U. Partitions of U refine every open cover of U: first extract a finite clopen subcover by compactness, then take all nonempty intersections of its members and complements. For a finite disjoint cover, partitions subordinate to that cover form a final subset of the refinement poset.

The formula then turns the cover into a product, because filtered colimits of spaces commute with finite products. Consequently it satisfies the sheaf condition on the basis. The map from the constant presheaf K to this sheaf exhibits sheafification: a compatible family of maps from the presheaf into any sheaf extends to each partition by the sheaf's product identification on its parts, and restriction along refinements gives the same map. This gives the adjunction on transformation spaces and proves (26.1).

The functor of K in (26.1) is not representable by a single space. To prove this, take a countable family of discrete spaces $K_i = \{0, 1\}$. A section of p^*K_i is a locally constant function $X \to \{0, 1\}$. The coordinate projections $x \mapsto x_i$ give a family of such sections. It does not arise from a section of $p^*(\prod_i K_i)$: the target $\prod_i K_i$ as a product of 0-truncated spaces is a discrete space, and the function determined by the coordinates is the identity on the underlying set of X. It is not locally constant on any nonempty cylinder. Thus the comparison to the product of the separate section spaces fails surjectivity on components. Every represented functor $\mathrm{Map}(S, -)$ preserves products, which proves the assertion. The cofiltered system of finite partitions does represent the functor in the pro-sense, by the explicit formula.

26.4 Two truncation operations

A finite quotient $Q_{\mathcal{P}}$ in the preceding calculation approximates the shape by changing the indexing space. A Postnikov truncation $\tau_{\leq n}K$ changes the coefficient space. The resulting section formulas have different limits. For example on the Cantor space the product failure uses a sequence of discrete coefficients, so it is present even when every coefficient is already 0-truncated. On the face poset of the boundary of a simplex, replacing the coefficient by its component set can remove the top cohomology even though the shape itself is represented by a finite complex. Keeping the variables separate makes these two computations compatible with the same Yoneda principle.

Exercise 26.2. Compute $\exists_f$ and $\forall_f$ on subsets for the projection $f : \{0,1\} \times T \to T$, where T is a set. Then compute the dependent product of a family of spaces E over $\{0,1\} \times T$.

Solution. Writing a subset as a pair (A_0, A_1) of subsets of T, the existential quantification is $A_0 \cup A_1$ and the universal quantification is $A_0 \cap A_1$. Over a point t, the dependent product of spaces is $E_{(0,t)} \times E_{(1,t)}$. For subobjects the product is inhabited precisely when both factors are inhabited, recovering the intersection formula. This is the discrete-base test of the general slice adjunction. □

Chapter 27

Products and an unbounded sheaf example

27.1 A product of local epimorphisms

Let $X = \mathbb{R}$, $U_n = (-1/n, 1/n)$, and $V = \mathbb{R} - \{0\}$ for $n \geq 1$. With integer coefficients, write $\mathbb{Z}_U$ for the constant sheaf on an open U. Put

$$F_n = j_{U_n!}\mathbb{Z}_{U_n} \oplus j_{V!}\mathbb{Z}_V \quad \text{and} \quad q_n : F_n \longrightarrow \mathbb{Z}_X$$

where q_n is the sum of the two extension-by-zero counits. It is an epimorphism. Indeed, U_n and V cover X, and at a point in either open the corresponding summand maps isomorphically on the stalk. Let $q = \prod_n q_n$ be the product in ordinary sheaves. We compute its cokernel.

For a connected open interval W containing 0, the V-summand of $F_n(W)$ has image zero in $\mathbb{Z}_X(W)$: its image is zero at 0, and every section of the constant sheaf is constant on W. If W contains a point outside U_n, the same reasoning at that point shows that the U_n-summand has image zero. If $W \subseteq U_n$, every integer lifts through the U_n-summand. Consequently the image of $q(W)$ in $\mathbb{Z}^{\mathbb{N}}$ consists of the sequences supported on the finitely many n with $W \subseteq U_n$. Passing to the stalk at 0 allows every finite support and no infinite support. The target stalk is $\prod_{n\geq 1}\mathbb{Z}$, because products of ordinary sheaves are computed on sections and the restriction maps on the indicated connected intervals are identities. Thus

$$(\operatorname{coker} q)_0 \cong \left(\prod_{n\geq 1}\mathbb{Z}\right) \Big/ \left(\bigoplus_{n\geq 1}\mathbb{Z}\right). \tag{27.1}$$

At each $x \neq 0$ choose a neighborhood inside V; all coordinates lift at once using the V-summand. So the cokernel has zero restriction to V and is the skyscraper at 0 with the group (27.1).

This calculation gives failure of exactness of countable products even on $\mathbb{R}$. If $K_n = \ker q_n$, the long exact sequence for the right derived functor of product, taken on the category of countable families of sheaves, has an injection $\operatorname{coker} q \hookrightarrow R^1 \prod_n K_n$. The relevant part of the sequence follows by taking the ordinary product on the category of complexes of injective families, or from the cofiber sequence of the derived product. Thus the difference has a derived group. There is no contradiction with exactness of products in ordinary modules: the individual coordinate lifts have neighborhoods that shrink with n.

27.2 A sufficient bounded-product criterion

The following separates exactness of products from the completeness it is often used to prove.

Proposition 27.1. *Let C be stable with a left separated t-structure and countable products. Suppose that for some $e \geq 0$ a countable product of objects in $C_{\geq a}$ is in $C_{\geq a-e}$ for every a. Then the canonical map $X \to \lim_n \tau_{\leq n} X$ is an equivalence for every X.*

Proof. By exactness of limits and the truncation sequences, the fiber of the comparison is $\lim_n \tau_{\geq n+1} X$. For any m, its limit is the same on the tail $n \geq m$ by initiality of the tail inclusion for limits. Both products in the shift-difference fiber formula for that tail belong to $C_{\geq m+1-e}$. The fiber of a map between two objects in $C_{\geq r}$ lies in $C_{\geq r-1}$, using the cofiber sequence and extension closure of the connective half. Therefore the limit belongs to $C_{\geq m-e}$ for every m. Left separatedness forces it to vanish, proving the comparison. □

For $\mathrm{D}(R)$ one can take $e = 0$ by exactness of ordinary module products. If ordinary products have higher derived functors, this does not settle whether a finite uniform e exists. The proposition uses only that bound on the entire connective halves. In particular the cokernel example above alone does not decide left completeness of the whole derived category on $\mathbb{R}$.

27.3 A tower with a nontrivial limit at an accumulation point

Let k be a nonzero field. Form X as the one-point compactification of the topological disjoint union of S^n for $n \geq 1$, writing ∞ for the added point. Every S^n is open and closed. The opens

$$W_N = \{\infty\} \cup \coprod_{n \geq N} S^n, \qquad N \geq 1,$$

form a neighborhood basis of ∞: a compact subset of the disjoint union is contained in a finite union of its components, since those components are an open cover. Let A_n be the extension by zero of the constant k-sheaf on S^n. In the derived category of k-sheaves on X set

$$K = \bigoplus_{n \geq 1} A_n[n].$$

The coproduct can be computed termwise on these complexes, because coproducts of ordinary sheaves are exact as filtered colimits of finite sums. Consequently $\pi_n K = A_n$ for $n \geq 1$, and all other homotopy sheaves vanish. Its coconnective truncation in degree $m \geq 1$ is the finite sum of the terms with $1 \leq n \leq m$. The inverse limit of this tower of finite sums with projections is the product

$$\lim_m \tau_{\leq m} K \simeq P := \prod_{n \geq 1} A_n[n] \tag{27.2}$$

in the derived category. This follows by mapping from a test object and comparing the product of all mapping spaces with the limit of its finite subproducts.

The stalk of K at ∞ is zero, because all summands have zero stalk there and stalks preserve coproducts. For the right side of (27.2), evaluate derived sections on W_N. The section functor is a right adjoint and hence preserves the product. For $n < N$ the restricted coefficient is zero. For $n \geq N$ it is extended from a clopen S^n, so its derived sections are $\mathrm{R}\Gamma(S^n, k)[n]$. That object has $\pi_0 = k$ from $H^n(S^n, k)$. Therefore

$$\pi_0 \mathrm{R}\Gamma(W_N, P) \cong \prod_{n \geq N} k,$$

and restriction from W_N to W_{N+1} deletes the Nth factor. To identify this with a homotopy-sheaf stalk, let G be any object of $\mathrm{D}(\mathrm{Sh}(X; \mathrm{Mod}_k))$ and choose a K-injective complex I of sheaves representing G. By Theorem 17.1, $\mathrm{R}iG$ is represented by iI, its value on U computes $\mathrm{R}\Gamma(U, G)$, and the counit is represented by $a(iI) = I$. Exactness of a therefore gives

$$\begin{aligned} \pi_0 G = H^0(I) &\cong a\Big(U \mapsto H^0(I(U))\Big) \\ &\cong a\Big(U \mapsto \pi_0 \mathrm{R}\Gamma(U, G)\Big). \end{aligned}$$

Sheafification preserves stalks, so the stalk of $\pi_0 G$ is the filtered colimit of these section groups over neighborhoods. Applying this to $G = P$ shows that at ∞ its degree-zero homotopy-sheaf stalk is

$$\operatorname*{colim}_N \prod_{n \geq N} k \cong \Big(\prod_{n \geq 1} k\Big) / \Big(\bigoplus_{n \geq 1} k\Big).$$

The quotient is nonzero, for instance by the constant sequence of ones. The map $K \to P$ fails to be an equivalence, as detected by this stalk. Every bounded piece is in the hyperdescent model of Proposition 17.2, and so is its limit.

The truncations $\tau_{\leq m}$ in (27.2) are those of the sheaf t-structure. They are not the objectwise Postnikov truncations of the section spectra $U \mapsto \mathrm{R}\Gamma(U, K)$. In particular, each factor $A_n[n]$ has zero degree-zero homotopy sheaf, whereas the categorical product P has the nonzero stalk just computed. Infinite products in the hyperdescent category thus need not commute with taking

homotopy sheaves. For $H = \mathrm{fib}(K \to P)$, the long exact sequence gives

$$(\pi_{-1}H)_\infty \cong (\pi_0 P)_\infty \cong (\prod_{n\geq 1} k)/(\bigoplus_{n\geq 1} k),$$

since $\pi_0 K = \pi_{-1} K = 0$. Consequently this fiber does not satisfy the vanishing hypothesis of hypercomplete Whitehead in Proposition 17.2. Hyperdescent and reconstruction by this sheaf Postnikov tower impose distinct conditions, compatible with the general distinctions in [22, 7].

27.4 Where the amplitude estimate fails

At the added point there can be no uniform finite bound on the degrees of derived sections on the chosen neighborhood basis. For any N and any proposed bound d, select $n \geq N$ with $n > d$. The ordinary coefficient A_n restricted to W_N has nonzero $H^n(W_N, A_n)$ from its sphere. Consequently a uniform application of Theorem 14.1 to these neighborhood section functors is unavailable. In the product example (27.2), the shifts bring cohomology in arbitrarily high neighborhood degrees down to degree zero. Each summand vanishes on sufficiently small neighborhoods, but the product has no common neighborhood of vanishing. This accounts for the difference between taking the stalk before and after the product.

If the dimensions of the clopen spheres are bounded, the same family of shifted constant coefficients with shifts tending to $+\infty$ cannot give this particular discrepancy in degree zero: for all but finitely many summands the corresponding group on every W_N is zero. After shrinking past the finitely many exceptions the product group in degree zero vanishes. This observation is about the displayed family of coefficients; a general completeness assertion requires an estimate on all relevant coefficient objects.

Exercise 27.2. For a sequence B_n of abelian groups, identify $\mathrm{colim}_N \prod_{n\geq N} B_n$ with a quotient of $\prod_n B_n$. Compare the result

with $\lim^1$ in calculation 13.1.

Solution. The map from the full product is surjective, since a tail sequence can be extended by zeros in the omitted coordinates. Its kernel consists of sequences vanishing on a tail, namely the direct sum. In calculation 13.1 the quotient is instead the cokernel of a shift-difference map in a derived inverse limit. Both groups use an infinite product, but the present quotient comes from the ordinary filtered colimit of stalk neighborhoods after taking derived sections. □

Chapter 28

A finite-site differential from an extension class

28.1 Constructing a coefficient with two layers

Fix a field k and $d \geq 2$. Let P_d be the poset of nonempty proper subsets of a set of $d+2$ elements, with the upper topology of Chapter 20. Write $C = \mathrm{Fun}(P_d, \mathrm{D}(k))$. For finite-dimensional k-vector spaces A, B, let $\underline{A}, \underline{B}$ be the constant diagrams. The projective resolution of the constant module applied with free generators in A gives

$$\mathrm{Hom}_{\mathrm{h}\,C}(\underline{A}, \underline{B}[d]) \cong \mathrm{Hom}_k(A, B) \otimes_k H^d(S^d; k) \cong \mathrm{Hom}_k(A, B) \tag{28.1}$$

after choosing the generator of $H^d(S^d; k)$ from an orientation of the order complex. Indeed the cochain complex of maps out of the projective resolution is the simplicial cochain complex of the order complex with coefficient $\mathrm{Hom}_k(A, B)$, since the target is constant. The direct sum of finitely many copies of the resolution for k is a projective resolution for $\underline{A}$. The isomorphism in (28.1) is independent of a basis of A, B, by the naturality of the coefficient cochain complex.

For a map $u : A \to B$ choose its corresponding point $\widehat{u} : \underline{A} \to \underline{B}[d]$ in C and set $K_u = \mathrm{fib}(\widehat{u})$. Since $d \geq 2$, the cofiber sequence

$$\underline{B}[d-1] \longrightarrow K_u \longrightarrow \underline{A} \xrightarrow{\widehat{u}} \underline{B}[d] \tag{28.2}$$

has two nonzero homotopy layers, $\pi_0 K_u = \underline{A}$ and $\pi_{d-1} K_u = \underline{B}$. These are independent of u. At every vertex of P_d, the map $\widehat{u}$ is

homotopic to zero because $\mathrm{Hom}_{\mathrm{h\,D}(k)}(A, B[d]) = 0$. Hence each stalk coefficient splits as $A \oplus B[d-1]$. The class in (28.1) belongs to the diagram and records coherent comparison over the whole site.

28.2 The value of derived sections

Write $F = \lim_{P_d} : C \to \mathrm{D}(k)$. With the chosen orientation,

$$F(\underline{A}) \simeq A \oplus A[-d], \qquad F(\underline{B}[d]) \simeq B[d] \oplus B.$$

These formulas are the constant diagram calculation for the sphere. Since k is a field, there are no degree-zero maps in $\mathrm{D}(k)$ between the displayed shifted summands of different homotopy degrees. Thus the only potentially nonzero matrix entry of $F(\widehat{u})$ is $A \to B$ in degree zero. To compute it, use the description

$$\pi_0 F(\underline{B}[d]) \cong \mathrm{Hom}_{\mathrm{h}\,C}(\underline{k}, \underline{B}[d]).$$

For a vector $a \in A$, the corresponding element of $\pi_0 F(\underline{A})$ is the constant map $\underline{k} \to \underline{A}$ with value a. Composing with $\widehat{u}$ corresponds under (28.1) to $u(a)$ times the orientation generator. This identifies the matrix entry with u itself.

Theorem 28.1 (Sections of the two-layer coefficient)**.** *The section spectrum of K_u, viewed in $\mathrm{D}(k)$, is*

$$F(K_u) \simeq A[-d] \oplus B[d-1] \oplus \mathrm{fib}(u : A \to B). \tag{28.3}$$

Its only possibly nonzero homotopy groups, in degrees $d-1, 0, -1, -d$, are respectively

$$B, \quad \ker u, \quad \mathrm{coker}\, u, \quad A.$$

Proof. Apply the exact right adjoint F to (28.2). In the displayed decompositions $F(\widehat{u})$ is the block $A \to B$ and zero on all other components, as just proved. The fiber of this block diagonal map is the sum of the fiber of u, the unused source $A[-d]$, and $B[d-1]$,

the desuspension of the unused target $B[d]$. This proves (28.3). The long exact sequence of the fiber of a map of degree-zero modules gives its two groups, $\ker u$ in degree zero and $\operatorname{coker} u$ in degree -1. All four degrees in the statement are distinct since $d \geq 2$. □

In particular for $A = B = k$, $u = 0$ gives four one-dimensional groups while $u = 1$ leaves only those in degrees $d-1$ and $-d$. This proves directly that identical cohomology sheaves and identical stalkwise complexes can give different section groups. The relative degree of the two layers and the dimension of the site place the discrepancy in the same two groups for every d.

28.3 Identifying the page of the differential

The Postnikov spectral sequence of Chapter 19 has two rows, at $q = 0$ and $q = 1 - d$ in cochain grading. Both have nonzero columns only at $p = 0, d$. The entries are

q	$p = 0$	$p = d$
0	A	A
$1 - d$	B	B

with the second entry in each row identified using the orientation. A differential of bidegree $(r, 1 - r)$ can connect these entries only if $r = d$, with source $(0, 0)$ and target $(d, 1 - d)$. The exact couple identifies this transgression with the connecting map for the two-layer triangle. The map from the bottom truncation to the shift of the top layer is $\widehat{u}$, and the section comparison on the selected groups is u. Thus d_d is u up to the sign of the orientation of the exact couple. One can check the identification by splicing the successive truncation triangles: in the intervening degrees the layers vanish, so their maps are equivalences, and the only boundary is the boundary of (28.2). Its kernel and cokernel on that page agree with the corresponding groups of Theorem 28.1.

For $d = 2$ this is the first possible differential, d_2, in hypercohomology. For larger d, all earlier pages keep the four entries. The example realizes a prescribed matrix as that later differential by choosing the extension class $\widehat{u}$. It uses an ordinary finite site, so the effect does not require an unbounded convergence phenomenon. In particular the amplitude bound in Theorem 20.1 gives a finite input interval for each group and (28.2) specifies the extension information to retain.

28.4 A direct description of a representative

Choose an oriented top cochain ω on the order complex representing the generator, and choose a chain of length $d + 1$ with its orientation. For instance, a cochain taking value 1 on one oriented top simplex and 0 on the other oriented top simplices evaluates to 1 on an oriented fundamental chain after the appropriate sign choice. There are no $(d + 1)$-simplices, so it is closed. Sending the P_{p_d}-summand of the degree-d resolution of $\underline{A}$ to $\underline{B}$ by the coefficient $\omega(p_0, \dots, p_d)u$ gives a representative of $\widehat{u}$ on that projective resolution. Changing ω by a coboundary changes the representative by a chain homotopy. One can take its mapping fiber as an ordinary complex of diagrams and compute sections from Theorem 20.1. This recovers the derived object of (28.3) by a finite cochain calculation. The geometric description by the order complex fixes the generator without a choice of cochain model, while the resolution gives an explicit model for computation.

Exercise 28.2. For $d = 3$, let $A = k^4$, $B = k^5$, and let u have rank 2. List the nonzero groups of the section object and their cohomological degrees.

Solution. The homotopy groups are k^5 in degree 2, k^2 in degree 0, k^3 in degree -1, and k^4 in degree -3. Changing sign to cohomological degree gives $-2, 0, 1, 3$, respectively. The ranks at

$0, 1$ use the kernel and cokernel of the d_3 differential. The two outer groups are the terms untouched by that differential. $\square$

Chapter 29

Descent for extension spaces

29.1 Gluing the coefficient categories

Let $X = U \cup V$ with U, V open and $W = U \cap V$. Denote by $D(Y)$ the derived ∞-category of modules on an open Y, using either a fixed ordinary coefficient ring or the restriction of a fixed sheaf of rings on X. All restrictions here are exact and t-exact. There is an equivalence

$$D(X) \simeq D(U) \times_{D(W)} D(V). \tag{29.1}$$

On the right, an object consists of A_U, A_V and a specified equivalence between their restrictions to W. The pullback is taken in Cat_∞. This means that the comparison is also part of the data of a map between two such objects.

Given gluing data with common overlap value A_W, form in $D(X)$ the pullback

$$(j_U)_* A_U \times_{(j_W)_* A_W} (j_V)_* A_V. \tag{29.2}$$

Here the direct images are derived and the maps are the restriction adjunctions with the given comparison. Restricting (29.2) to U commutes with the pullback. By open base change, the third term becomes the direct image of A_W along $W \subset U$, which identifies with the middle term after restriction. The first term becomes A_U. The pullback therefore restricts to A_U, and similarly to A_V, with the specified overlap comparison. Conversely, for $A \in D(X)$ the canonical map to (29.2) for its restrictions is an equivalence on U and V by this computation. Those restrictions

detect equivalences since every stalk is in one of the two opens. Both constructions are functorial and exhibit (29.1).

In particular if $A, B \in D(X)$, the mapping space from B to A is the pullback of the two mapping spaces on U, V over the one on W. This can also be obtained directly by mapping out of B into the pullback (29.2) and using the adjunctions for the three open inclusions. Applying this argument after shifting A proves the pullback of mapping spectra.

29.2 Local extension data and comparison paths

Theorem 29.1. *For fixed $A, B \in D(X)$ and $n \geq 0$, the space $E^n_X(B, A) = \mathrm{Map}_{D(X)}(B, A[n])$ is the homotopy pullback*

$$E^n_U(B_U, A_U) \times_{E^n_W(B_W, A_W)} E^n_V(B_V, A_V).$$

For $n = 1$ it is the space of short stable extensions with the designated outer terms. For ordinary heart objects, its components for all n are the Yoneda extension classes in the corresponding sheaf categories.

Proof. Suspension commutes with the restrictions. Apply the mapping-space calculation following (29.2) to $A[n]$. For $n = 1$, Theorem 10.1 identifies the mapping spaces with extension spaces naturally in the restrictions. For heart objects, Proposition 8.2 and the bounded derived comparison of Chapter 10 identify the components. The morphisms from a heart object to a finite shift of another can be computed in the bounded subcategory of the derived category by the usual injective or Yoneda comparison. □

Two specified local extension points glue exactly when the two points they induce on W can be joined by a path. The fiber over those local points is the space of such paths. After choosing one path, this fiber is equivalent to $\Omega E^n_W(B_W, A_W)$. For $n \geq 1$ the latter is $E^{n-1}_W(B_W, A_W)$ based at zero after translating. If one

specifies only the components on U, V, automorphisms of the two local points can act on the components of the gluing path space. That effect is computed by the long exact sequence of the mapping spectra.

In fact the pullback of mapping spectra is the fiber of a difference, so for ordinary A, B it gives the exact sequence

$$\begin{aligned}&\operatorname{Ext}_U^{n-1}(B_U, A_U) \oplus \operatorname{Ext}_V^{n-1}(B_V, A_V)\\ &\quad\longrightarrow \operatorname{Ext}_W^{n-1}(B_W, A_W) \longrightarrow \operatorname{Ext}_X^{n}(B, A)\\ &\quad\longrightarrow \operatorname{Ext}_U^{n}(B_U, A_U) \oplus \operatorname{Ext}_V^{n}(B_V, A_V) \longrightarrow \operatorname{Ext}_W^{n}(B_W, A_W).\end{aligned}$$

The middle boundary takes an overlap comparison to an extension class with locally zero restrictions. Its ambiguity is the image from the two local groups, by exactness. In the case $n = 1$, this says a globally nonsplit short extension may be assembled from two split ones, with their difference on the overlap a map $B_W \to A_W$.

Calculation 29.2 (The gluing of a split extension)**.** Use the projective line and its two affine opens, with $B = \mathcal{O}$ and $A = \mathcal{O}(-2)$. A map $B \to A$ is a section of A, and B is locally free of rank one. On either affine the group $\operatorname{Ext}^1(B, A)$ is zero, since the derived Hom is the derived sections of A and the latter is acyclic there. The global Ext^1 group is therefore the cokernel of the difference of the two Hom groups into the overlap Hom group. In the first frame, the latter is $A_{\text{base}}[t, t^{-1}]$, where A_{base} denotes the base ring. The cokernel is free of rank one on t^{-1} by (23.1). Gluing the local split extensions by the triangular automorphism with off-diagonal entry t^{-1} gives that generator. The Euler sequence of (23.4) has the same local splittings with off-diagonal difference $\pm t^{-1}$, as in its boundary calculation. This realizes the long exact sequence by an explicit change of comparison between split local objects.

29.3 Reconstructing a global map from local finite windows

Suppose X has an open cover by U, V such that the derived categories on U, V are left complete in their standard t-structures. Then $D(X)$ is left complete. Indeed restriction along each open has both adjoints, so preserves limits. It is t-exact, so it sends the canonical map $K \to \lim_n \tau_{\le n} K$ to the corresponding comparison for its restriction. The two maps are equivalences by the local hypotheses. The restrictions jointly detect equivalences, either by (29.1) or by testing stalks. This proves the claim. The same proof works for right completeness using colimits and connective truncations. In the present derived sheaf case right completeness also follows from exactness of filtered colimits.

Under these two local completeness assumptions the finite-window mapping formula of Theorem 13.2 can be applied on U, V, and on W. For the last assertion, completeness restricts to an open: extend an object by $j_!$ to the larger open, apply its canonical reconstruction, and restrict back, using $j^* j_! \simeq \mathrm{id}$ and preservation of the required limits, colimits, and truncations by j^*. Since the gluing uses a limit of mapping spaces, the limits in the window indices commute with it. Consequently global maps can be computed by gluing the local finite-window mapping spaces in each window and taking the resulting inverse limit. This reproduces the global formula but uses only the indicated local completeness checks.

For the space of Chapter 27, no open neighborhood of the accumulation point can satisfy the analogous left completeness test. It contains some W_N, and the same direct sum and product argument on that tail exhibits failure of reconstruction there. Completeness on such a larger neighborhood would pass to W_N by the preceding restriction argument. Thus the local criterion also explains why a cover avoiding the problem cannot be chosen in that example.

Exercise 29.3. Suppose E^1_U and E^1_V each have one component,

and $\mathrm{Hom}_U(B_U, A_U)$ and $\mathrm{Hom}_V(B_V, A_V)$ both vanish. Identify the set of components of E^1_X, using a chosen pair of split local extensions.

Solution. The two local components are represented by the split extensions. Their overlap points are equivalent, so the path fiber is nonempty and its components form a torsor for $\mathrm{Hom}_W(B_W, A_W)$. The two vanishing local Hom groups mean there are no component identifications coming from local automorphisms. Choosing the identity comparison of the split restrictions fixes the origin, giving the group $\mathrm{Hom}_W(B_W, A_W)$ as the set of global extension classes. This is also the $n = 1$ case of the exact sequence above. □

Chapter 30

Local-to-global Ext and classical recovery

30.1 A sheaf of derived mapping objects

Let $\mathcal{O}$ be a sheaf of commutative rings on a topological space X and let A, B be ordinary $\mathcal{O}$-modules. Choose an injective resolution $A \to I^\bullet$ in $\mathcal{O}$-modules. For an open U define the complex

$$J^\bullet(U) = \operatorname{Hom}_{\mathcal{O}|_U}(B|_U, I^\bullet|_U).$$

The terms form sheaves, since maps of sheaves glue uniquely on open covers. Restrictions of injectives are injective because extension by zero is an exact left adjoint. Thus $J^\bullet(U)$ calculates the derived Hom on U, naturally in U. The sheaf cohomology objects $H^q(J^\bullet)$ are written $\mathcal{E}xt^q_{\mathcal{O}}(B, A)$. They are sheafifications of the presheaves of local Ext groups. In degree zero this is the ordinary sheaf $\mathcal{H}om_{\mathcal{O}}(B, A)$.

Each term J^q is flabby. To extend a map $B|_U \to I^q|_U$ to $V \supseteq U$, view it by adjunction as a map from $j_!(B|_U)$ into $I^q|_V$, where $j : U \to V$. The natural map $j_!(B|_U) \to B|_V$ is monic on stalks. Injectivity of $I^q|_V$ gives an extension across it. This proves the restriction surjection. Consequently $J^\bullet$ is a bounded-below complex of section-acyclics. Such a complex computes $\mathrm{R}\Gamma$ by its ordinary section complex [42, Lemma 13.16.7]. One can verify the comparison with a bounded-below Cartan–Eilenberg injective double resolution: the column cohomology after sections vanishes in positive resolution degrees, and the first-quadrant totalization has only finitely many terms in each total degree. The resulting

comparison is therefore a quasi-isomorphism. That complex is also the injective-resolution model for $\mathrm{RHom}_{\mathcal{O}}(B, A)$. Thus the stable object $J^\bullet$ supplies local derived maps, with global derived Hom as its section object. On an overlap this construction yields exactly the derived mapping diagram of Chapter 29.

Theorem 30.1 (Local-to-global Ext)**.** *For ordinary $\mathcal{O}$-modules A, B there is a first-quadrant spectral sequence*

$$E_2^{p,q} = H^p(X, \mathcal{E}xt^q_{\mathcal{O}}(B, A)) \implies \mathrm{Ext}^{p+q}_{\mathcal{O}}(B, A). \tag{30.1}$$

It is induced by the cohomological truncations of $J^\bullet$.

Proof. The qth layer is $\mathcal{E}xt^q(B, A)[-q]$, and applying the exact derived section functor to the tower gives the exact couple in Chapter 19. The cohomology of the layers on X gives the displayed page. The groups vanish for $q < 0$ since the resolution starts in degree zero, and for $p < 0$ by the cohomological grading of derived sections of a sheaf. For total degree r the layers of cohomological degree $> r$ have sections with no cohomology in degrees $\leq r$. One can see this directly from bounded-below injective models of the cohomological lower truncation. Thus finitely many layers suffice for each total degree, and the exact couple converges to the cohomology of $\mathrm{R}\Gamma(X, J^\bullet)$, identified above with global Ext. □

When B is finite locally free, $\mathcal{H}om(B, -) \simeq B^\vee \otimes -$ is exact on sheaves by its local description. Thus all rows $q > 0$ vanish and (30.1) identifies $\mathrm{Ext}^p_{\mathcal{O}}(B, A)$ with $H^p(X, B^\vee \otimes A)$. More generally the local derived Hom can be computed from a bounded locally free resolution of B when it exists. On an open U where a term is finite free, the internal sheaf functor $\mathcal{H}om_{\mathcal{O}_U}(\mathcal{O}_U^{\oplus r}, -) \simeq (-)^{\oplus r}$ is exact on all sheaves of modules. Exactness is local, so a finite locally free term has the same property. For a bounded locally free resolution, the finite filtration of its internal Hom complex then shows that Hom into any acyclic sheaf complex is acyclic. The subsequent operation of taking sections is derived separately. This gives the same derived object as resolving the target. The

ordinary Ext functor from the source text is recovered from its global groups, while the sheaf model records the local calculation and its restriction maps.

Calculation 30.2 (An Ext group supported on a divisor)**.** Let $X = \mathbb{P}^1_A$ and $D \subset X$ be the divisor $x_1 = 0$, with A any nonzero commutative ring. The exact sequence

$$0 \longrightarrow \mathcal{O}(-1) \xrightarrow{x_1} \mathcal{O} \longrightarrow \mathcal{O}_D \longrightarrow 0$$

is locally the resolution by the nonzero divisor t on U_0 and by a unit on U_1. The derived internal Hom from $\mathcal{O}_D$ to $\mathcal{O}$ is the complex $\mathcal{O} \xrightarrow{x_1} \mathcal{O}(1)$ in degrees $0, 1$. Its only cohomology sheaf is $\mathcal{O}_D(1)$ in degree one. The divisor is $\operatorname{Spec} A$ and x_0 trivializes its twist; its pushforward is supported on D and has sections A with no higher cohomology. Thus (30.1) gives $\operatorname{Ext}^1_{\mathcal{O}}(\mathcal{O}_D, \mathcal{O}) \cong A$, with all other groups zero. Independently apply $\mathrm{R}\Gamma$ to the two-term internal Hom complex. The cohomology here is sheaf cohomology in the category of all $\mathcal{O}_X$-modules on the Zariski site. Affine acyclicity of quasi-coherent modules in that category [35] identifies it with the two-open Čech calculation of Chapter 23. In particular $\mathcal{O}$ and $\mathcal{O}(1)$ have no positive cohomology, so their bounded two-term complex computes derived sections termwise by the same acyclic-complex argument. The section arrow is $A \to A^2$ taking 1 to the coordinate x_1, whose cokernel is the x_0 coordinate. The two computations identify the same extension group through its local support.

30.2 Concordance with the source

The table locates the principal classical constructions in Strooker's text [38] and the corresponding higher calculations. Section numbers in the left column refer to the source. The right column points to the proofs and examples in this manuscript.

Strooker	Higher formulation and recovery
§§1.1–1.3 Categories and functors	Quasicategories, cores, and coherent transformations. The nerve recovers ordinary composition and discrete hom-sets (Chapter 1).
§§1.4–1.6 Representability, products, limits	The Yoneda equivalence, density by right fibrations, and cones tested by mapping spaces (Chapters 2 and 3).
§§1.7–1.9 Adjoints, order, preservation	Adjoint representability, mates, and cofinality. Restriction to posets recovers joins, meets, and Galois connections (Chapters 3 and 4).
§§2.1–2.5 Monos, zeros, additivity, exactness	Relative images, stable fibers and cofibers, and admissible sequences. Truncation gives the kernels and cokernels in the heart (Chapters 5 to 8).
§§2.6–2.9 Exact sequences, generators	Functorial cofiber sequences, exact functors, Grothendieck categories, and compact spectral generators (Chapters 6, 7 and 9).
§3.1 Yoneda extensions	Extensions with fixed endpoints form spaces, and their homotopy groups recover shifted Ext groups (Chapter 10).
§§3.2–3.4 Satellites and derivation	Weighted enriched Kan extensions, K-injective models, and the comparisons with long exact sequences (Chapters 11 and 12).
§§4.1–4.5 Concrete sheaves and sheafification	Germ spaces for discrete sheaves, covering-sieve localization, and hypercover descent (Chapters 15 and 17).
§§4.6–4.9 Base space and module sheaves	Geometric adjunctions, open base change, supports, and exact localization of derived module presheaves (Chapters 16 and 17).

Strooker	Higher formulation and recovery
§§4.10–4.12 Cohomology and acyclicity	Stable sections, flabby resolutions, fine and soft coefficients, and totalization (Chapters 18 and 19).

30.3 Three points of contact between the constructions

First, the graded Kan formula is determined by the ability to move an extension class along a test map. The tensor relation in (11.1) enforces exactly this action. An injective effacement turns the extension weight into a cokernel of two represented functors, so the comparison with the resolution applies without choosing generators of an Ext group. For local-to-global Ext the coefficient sheaves $\mathcal{E}xt^q$ provide a different weighting of the local calculations, and the exact couple describes the remaining gluing.

Second, cohomological amplitude is a bound on a stable functor, so it can be used before taking homotopy groups. In the finite-site example it follows from an explicit projective resolution of the constant diagram. On the projective line it follows from a two-open fiber sequence and affine exactness. In both cases the theorem applies to the coherent input window. The finite-site differential of Theorem 28.1 shows how a nonzero boundary within that window changes the sections.

Third, the universal properties require their own limiting operations. A finite-window test for maps uses reconstruction of the source and target, whereas a hypercover tests a sheaf against a family of geometric covering diagrams. The accumulation example gives a concrete distinction: the section spectra satisfy hyperdescent and the Postnikov comparison has a nonzero degree-zero stalk in its target. The left completeness hypothesis in the mapping theorem and the local equivalence condition in derived sheafification thus serve separate mathematical roles.

Exercise 30.3. Let A be an ordinary sheaf of R-modules on a topological space, with R commutative. Apply Theorem 30.1 with $B = R_X$ and identify the answer with one of the first constructions of Chapter 18.

Solution. The rank-one free sheaf has exact internal Hom $\mathcal{H}om(R_X, -) \simeq \mathrm{id}$, so the only row is $H^p(X, A)$ at $q = 0$. It follows that $\mathrm{Ext}^p(R_X, A) = H^p(X, A)$. At the level of the derived object, the internal Hom into the resolution $I^\bullet$ is $I^\bullet$ itself, and global derived Hom is the derived section object. This is the representation of sections by the constant rank-one sheaf. □

References

[1] C. Barwick, *On the algebraic K-theory of higher categories*. Journal of Topology **9** (2016), 245–347. arXiv:1204.3607.

[2] E. D. Børve and P. Trygsland, *A theorem of Retakh for exact ∞-categories and higher extension functors*. arXiv:2110.05138.

[3] A. K. Bousfield and D. M. Kan, *Homotopy Limits, Completions and Localizations*. Lecture Notes in Mathematics 304, Springer, 1972. Springer edition.

[4] G. E. Bredon, *Sheaf Theory*, 2nd ed. Graduate Texts in Mathematics 170, Springer, 1997. Springer edition.

[5] K. S. Brown, *Cohomology of Groups*. Graduate Texts in Mathematics 87, Springer, 1982. Springer edition.

[6] D.-C. Cisinski, *Higher Categories and Homotopical Algebra*. Cambridge Studies in Advanced Mathematics 180, Cambridge University Press, 2019. DOI: 10.1017/9781108588737.

[7] D. Clausen and A. Mathew, *Hyperdescent and étale K-theory*. Inventiones Mathematicae **225** (2021), 981–1076. arXiv:1905.06611.

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

[9] D. Dugger and D. C. Isaksen, *Hypercovers in topology*. arXiv:math/0111287.

[10] W. G. Dwyer and D. M. Kan, *Simplicial localizations of categories*. Journal of Pure and Applied Algebra **17** (1980), 267–284. Journal article.

[11] P. Gabriel, *Des catégories abéliennes*. Bulletin de la Société Mathématique de France **90** (1962), 323–448. Numdam edition.

[12] A. Grothendieck, *Sur quelques points d'algèbre homologique*. Tôhoku Mathematical Journal **9** (1957), 119–221. Journal article.

[13] P. J. Haine, M. Porta, and J.-B. Teyssier, *The homotopy-invariance of constructible sheaves*. Homology, Homotopy and Applications **25**, no. 2 (2023), 97–128. Journal article; arXiv:2010.06473.

[14] A. Hatcher, *Algebraic Topology*. Cambridge University Press, 2002, §4.L. Steenrod squares and powers.

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

[16] M. Kashiwara and P. Schapira, *Sheaves on Manifolds*. Grundlehren der mathematischen Wissenschaften 292, Springer, 1990. Springer edition.

[17] J. Klemenc, *The stable hull of an exact ∞-category*. Homology, Homotopy and Applications **24** (2022), 195–220. arXiv:2010.04957.

[18] J. Lurie, *Higher Topos Theory*. Annals of Mathematics Studies 170, Princeton University Press, 2009. Author's text.

[19] J. Lurie, *Higher Algebra*. 2017. Author's text.

[20] J. Lurie, *Spectral Algebraic Geometry*. 2018. Author's text.

[21] J. Lurie, *Kerodon*. Online reference. Kerodon.

[22] A. Neeman, *Non-left-complete derived categories*. Mathematical Research Letters **18**, no. 5 (2011), 827–832. Journal article; arXiv:1103.5539.

[23] T. Nikolaus, U. Schreiber, and D. Stevenson, *Principal ∞-bundles: General theory*. Journal of Homotopy and Related Structures **10**, no. 4 (2015), 749–801. Journal article; arXiv:1207.0248.

[24] V. S. Retakh, *Homotopic properties of categories of extensions*. Russian Mathematical Surveys **41**, no. 6 (1986), 217–218. Journal article.

[25] E. Riehl and D. Verity, *Elements of ∞-Category Theory*. Cambridge Studies in Advanced Mathematics 194, Cambridge University Press, 2022. DOI: 10.1017/9781108936880.

[26] S. Schwede and B. Shipley, *Stable model categories are categories of modules*. Topology **42** (2003), 103–153. arXiv:math/0108143.

[27] B. Shipley, *H$\mathbb{Z}$-algebra spectra are differential graded algebras*. American Journal of Mathematics **129** (2007), 351–379. arXiv:math/0209215.

[28] N. Spaltenstein, *Resolutions of unbounded complexes*. Compositio Mathematica **65** (1988), 121–154. Numdam edition.

[29] The Stacks Project Authors, *Derived Categories*. Tag 05QI.

[30] The Stacks Project Authors, *The Stacks Project*. Injectives in Grothendieck categories, Theorem 19.11.7. Tag 079H.

[31] The Stacks Project Authors, *Existence of K-injective complexes for Grothendieck abelian categories*. Tag 079P.

[32] The Stacks Project Authors, *The Gabriel–Popescu theorem*. Tag 0F5R.

[33] The Stacks Project Authors, *The Stacks Project*. Flasque sheaves, Lemma 20.12.3 (acyclicity). Tag 09SY.

[34] The Stacks Project Authors, *The Stacks Project*. Derived pullback, Lemma 20.27.1. Tag 06YJ.

[35] The Stacks Project Authors, *The Stacks Project*. Čech cohomology of quasi-coherent sheaves, Lemma 30.2.2 (affine vanishing). Tag 01XB.

[36] The Stacks Project Authors, *Cohomology of projective space*. Tag 01XS.

[37] The Stacks Project Authors, *The Stacks Project*. Derived category of quasi-coherent modules, §36.3. Lemma 36.3.4 (Tag 08D4); Lemma 36.3.5 (Tag 06Z0).

[38] J. R. Strooker, *Introduction to Categories, Homological Algebra and Sheaf Cohomology*. Cambridge Tracts in Mathematics 73, Cambridge University Press, 1978. The source for this manuscript is the supplied PDF of the digitally printed edition. Publisher's front matter.

[39] The Univalent Foundations Program, *Homotopy Type Theory: Univalent Foundations of Mathematics*. Institute for Advanced Study, 2013. Open edition.

[40] C. A. Weibel, *An Introduction to Homological Algebra*. Cambridge Studies in Advanced Mathematics 38, Cambridge University Press, 1994. DOI: 10.1017/CBO9781139644136.

[41] A. A. Beilinson, J. Bernstein, and P. Deligne, *Faisceaux pervers*. In *Analyse et topologie sur les espaces singuliers I*, Astérisque **100**, Société Mathématique de France, 1982, 5–171. Author's text, §1.3.

[42] The Stacks Project Authors, *The Stacks Project*. Higher derived functors, Lemma 13.16.7 (Leray's acyclicity lemma). Tag 015E.