Quotients, Refinement, and Reconstruction

Comparing choices, changing coordinates, and detecting what remains

Revised November 3, 2026

Abstract

An invertible replacement proves equality of finite basis sizes. A relation between two groups gives another working object: quotient its two ambiguities, and the relation becomes a canonical isomorphism. We perform both calculations before isolating their reusable operations in an exact Mal'cev setting. The resulting comparison proves refinement and Jordan–Hölder beyond the vector-space case.

Subsequent questions require different changes of working objects. Changing a lift measures a coupling coefficient; successive kernels recover nilpotent chains by finite differences; counting the same flags in two orders proves associativity. When orbit membership is insufficient, an action groupoid retains the actual moves. Later chapters construct linear coefficients over a nonlinear base and state the additional foundational inputs those questions need. Each route identifies the operation gained, its proof, and the data it cannot determine.

How to work through the constructions

Start with a question and the operation that changes it. Section 1 uses invertible replacements to prove equality of finite basis sizes. It then constructs a group comparison directly: keep a subgroup of a product as a relation, calculate the ambiguity on each side, and obtain a map between the quotients. Both are complete arguments before the general categorical definitions enter.

The common comparison route asks what those calculations used. Matching equations, forgetting witnesses, and quotienting compatible relations become explicit operations. The Mal'cev condition captures the three-corner calculation already performed in groups. Sections 2 through 4 develop refinement and compare its group and module instances. This route earns a result applicable beyond scalar pivoting.

Other questions have their own entry points. In Section 9, try changing a vector to remove an operator's coupling term; the calculation motivates the marked data subsequently classified. In Section 10, calculate successive kernel sizes and recover chain counts by finite differences; the chapter then constructs the chains that justify the calculation. In Section 11, count submodules and flags to obtain a product and prove its associativity. These are different operations, with different hypotheses and payoffs.

The final part is a further choice, motivated by a changed question. Section 12 begins with the moves that an orbit set forgets and retains them as arrows. Section 13 proves why a full additive replacement of groups cannot work, then constructs linear coefficients over a retained base. The stated higher-categorical equivalences are imported inputs; their role is identified where they are used. They are not prerequisites for the earlier finite arguments.

The aim is to make each useful change of working objects inspectable: what prompted it, why it answers the question, what computation it makes possible, and what remains undecided. Original theorem labels and the dependency record retain the precise mathematical scope.

Contents

How to work through the constructions **2**

I Construct the moves and the comparison **6**

1 Exchange vectors and remove ambiguity **6**
1.1 Invert the move, then reuse it . . . 6
1.2 Remove ambiguity before choosing an output . . . 7
1.3 A precise license to reason with witnesses . . . 9
1.4 Relations compose by matching, then forgetting . . . 10
1.5 Exactness makes an identification into an object . . . 11

2 The three-corner rule and one comparison **12**
2.1 Why relative difference is enough . . . 13
2.2 Partial identifications and their supports . . . 14
2.3 One relation, two quotients, one object . . . 15
2.4 Which relations descend through a quotient? . . . 16

3 A quotient remembers which identifications were made **18**
3.1 The data in a normal series . . . 18
3.2 What a genuine refinement means . . . 19
3.3 The comparison rectangle . . . 21
3.4 Simplicity removes the remaining freedom . . . 23
3.5 Factor data becomes an invariant only after the comparison . . . 24
3.6 What is natural in the construction . . . 25
3.7 Dependency and scope check . . . 26

II Reuse the comparison and calculate its pages **27**

4 Groups, modules, and dimension as instances **27**
4.1 One verification for algebraic structures . . . 27
4.2 Reading the axioms in familiar examples . . . 28
4.3 Reading one cell as the familiar group butterfly . . . 29
4.4 A comparison that changes the order . . . 29
4.5 Finite basis size, with no dimension in the prerequisites . . . 30
4.6 What a factor list leaves unresolved . . . 31

5 Keeping the operations when we change the object **31**
5.1 Modules and their maps . . . 31
5.2 Preimages, quotients, and the local comparison . . . 32
5.3 Finite-dimensional tools, after dimension has been proved . . . 32
5.4 An operator is a polynomial action . . . 33
5.5 Simple layers and finite length . . . 34

6 Comparison pages with the operations retained **35**

6.1 The normality that makes the pages possible . . . 35
6.2 Terms, maps, and the next page . . . 36
6.3 A nonabelian crossing and the module specialization . . . 37

7 Filtered maps and their successive obstructions 38
7.1 A map seen one layer at a time . . . 39
7.2 The pages and their maps . . . 39
7.3 What eventually survives . . . 41
7.4 Two composition series inside an acyclic complex . . . 42
7.5 A calculation with nonzero differentials . . . 43
7.6 The butterfly is visible inside the differentials . . . 44
7.7 What this adds, and what it does not assert . . . 45

III Measure coupling and recover chains 46

8 What every additive measurement must retain 46
8.1 Splicing gives additivity before we name a group . . . 46
8.2 Constructing the universal recipient . . . 46
8.3 Why the chosen category matters . . . 47
8.4 A short bridge to Euler characteristic . . . 48
8.5 The information that remains to be recovered . . . 48

9 What attaches one layer to another 48
9.1 Keeping the endpoints fixed . . . 49
9.2 Change the quotient: pull back . . . 50
9.3 Change the subobject: push out . . . 50
9.4 Adding extensions by merging their kernels . . . 51
9.5 An extension that fits in one matrix . . . 53
9.6 A boundary map measures failure to lift . . . 54
9.7 An attachment can skip an adjacent layer . . . 55

10 Recovering a nilpotent operator from spaces of maps 56
10.1 A relation becomes a test object . . . 56
10.2 Which endpoints can be reached by a long chain? . . . 58
10.3 Lifting the endpoint filtration to a basis of chains . . . 58
10.4 Reading the operator from its measurements . . . 60

11 Associativity by counting filtrations 62
11.1 A setting where the counts are finite . . . 62
11.2 One flag, two intermediate objects . . . 62
11.3 The first product, computed completely . . . 64
11.4 Three layers and a second count . . . 64
11.5 What the multiplication remembers . . . 65

IV Retain the moves and construct linear coefficients 66

12 Transporter graphs and higher-quotient references **66**
12.1 Reference: homotopical algebraic objects . . . 66
12.2 The action groupoid and the role of associativity . . . 68
12.3 Extensions become fiber diagrams . . . 68
12.4 Reference: universal change of base . . . 69
12.5 Reference: families over BG . . . 70
12.6 Calculate the transporter graph . . . 71
12.7 Reference: higher normality and quotient data . . . 72

13 Linear operations over a retained base **73**
13.1 Why a full additive replacement cannot work . . . 74
13.2 Modules already arise as abelian objects over a group . . . 75
13.3 Reference: stabilization over a retained base . . . 76
13.4 Construct an attachment from relations . . . 77
13.5 A derived tensor product retains an equation . . . 79
13.6 Reference: cotangent objects and deformation theory . . . 79

A Dependencies, scope, and further reading **82**
A.1 What was proved from what . . . 82
A.2 Three different comparison claims . . . 83
A.3 The additional hypotheses behind reconstruction . . . 83
A.4 Where the higher foundations enter . . . 83
A.5 Audit and reproducibility . . . 84

References **84**

Part I
Construct the moves and the comparison

1 Exchange vectors and remove ambiguity

When we write $a = b + c$ and then $b = a - c$, something more organized than symbol manipulation is happening. We have applied an invertible map. When we then write $d(a - c) = da - dc$, we have used a map that preserves an operation. These two ideas—invertibility and preservation—will remain visible as the objects become more general.

1.1 Invert the move, then reuse it

Let $(V, +, 0)$ be an abelian group. For $c \in V$, translation is the map of sets $\tau_c : V \to V$, $x \mapsto x + c$. Associativity, the identity law, and the inverse law give

$$\tau_{-c}\tau_c(x) = (x + c) + (-c) = x + (c - c) = x.$$

The same calculation in the other order proves $\tau_{-c} = \tau_c^{-1}$. Thus $a = b + c$ means $a = \tau_c(b)$, and solving for b means applying the inverse. Translation need not preserve the zero element, so it is generally an isomorphism of sets, not an automorphism of abelian groups.

Associativity itself says that two composites are the same map:

$$\begin{array}{ccc} V \times V \times V & \xrightarrow{+ \times \mathrm{id}} & V \times V \\ {\scriptstyle \mathrm{id} \times +}\downarrow & & \downarrow{\scriptstyle +} \\ V \times V & \xrightarrow[+]{} & V. \end{array}$$

Here the usual canonical identification of the two parenthesizations of the triple product is understood. The diagram says that making one intermediate sum rather than the other does not change the result.

If W is another abelian group and $h : V \to W$ preserves addition, then $h(0) = 0$: the equality $h(0) = h(0)+h(0)$ can be translated by $-h(0)$. Also $h(-c) = -h(c)$, because $h(c)+h(-c) = h(0) = 0$. Consequently

$$h(a - c) = h(a) - h(c), \qquad h\tau_c = \tau_{h(c)}h.$$

For a module, multiplication by a scalar is an additive homomorphism, so distributivity is exactly this preservation statement. In a group, the same discussion uses multiplication and left or right translations, with the order retained. Its multiplication is a map of underlying sets; it is not generally a homomorphism from the product group.

There is also a structural explanation of the triangular substitutions used in basis exchange. Given an additive homomorphism $u : V \to W$, the *shear*

$$S_u : V \times W \longrightarrow V \times W, \qquad (v, w) \longmapsto (v, w + u(v))$$

has inverse S_{-u}. This follows by composing the maps, without counting any basis. Over a field, if $t_1 \neq 0$, the map

$$(x_1, \ldots, x_n) \longmapsto (t_1x_1, t_2x_1 + x_2, \ldots, t_nx_1 + x_n)$$

is a scaling in the first coordinate followed by the shear S_u on $k \times k^{n-1}$, where

$$u : k \longrightarrow k^{n-1}, \qquad u(z) = (t_2 t_1^{-1} z, \ldots, t_n t_1^{-1} z).$$

Both maps are invertible. The inverse of the composite sends y_1 to $t_1^{-1} y_1$ and y_i to $y_i - t_i t_1^{-1} y_1$ for $i > 1$. This establishes the elementary maneuver without presupposing invariance of dimension.

We can already finish the basis-size question using this operation. Suppose $v_1, \ldots, v_n$ is a basis and $w_1, \ldots, w_m$ is an independent list. Express w_1 in the v-basis. A nonzero coefficient gives a pivot; after permuting the v's, the invertible triangular substitution just constructed replaces v_1 by w_1 and leaves a basis. Express w_2 in that new basis. Some coefficient outside the fixed first vector must be nonzero, since otherwise w_2 would belong to the span of w_1. Use that coefficient as the next pivot, fixing w_1.

Continue in this way. After r steps the first r entries are $w_1, \ldots, w_r$, and the entire list is still a basis because every replacement was invertible. If $m > n$, then after n steps the next w_{n+1} would be a combination of the preceding w's, a contradiction. Hence $m \leq n$. For two bases, reverse their roles to obtain $n \leq m$. Thus $m = n$, without determinants, a dimension function, or a Jordan–Hölder theorem. The zero-space case gives two empty bases.

The useful object here was an invertible replacement, and the useful invariant was the basis property after each move. Later we obtain another proof from quotient comparisons. That proof will show how a similar comparison works when scalar pivots are unavailable; it is not needed to make this first argument complete.

1.2 Remove ambiguity before choosing an output

Before introducing the language for a common comparison, perform one complete comparison in groups. Let R be a subgroup of $A \times B$. Interpret $(a, b) \in R$ as saying that a is related to b. Put

$$A' = \{a : \text{some } b \text{ has } (a, b) \in R\}, \qquad B' = \{b : \text{some } a \text{ has } (a, b) \in R\},$$

and record the ambiguities at the identity:

$$N = \{a : (a, 1) \in R\}, \qquad M = \{b : (1, b) \in R\}.$$

The sets A', B' are subgroups because products and inverses can be taken coordinatewise in R. If $a \in A'$ and $n \in N$, choose $(a, b) \in R$; conjugating $(n, 1)$ by (a, b) proves $ana^{-1} \in N$. Thus N is normal in A', and the same calculation proves that M is normal in B'.

Try to send a to a related b. This need not give a map, because b may not be unique. But if $(a, b), (a, b') \in R$, their relative difference is $(1, b^{-1}b') \in R$, so $bM = b'M$. And if $a' = an$ with $n \in N$, then $(a, b)(n, 1) = (a', b) \in R$. The assignment

$$\theta : A'/N \longrightarrow B'/M, \qquad aN \longmapsto bM \quad \text{for any } (a, b) \in R$$

is therefore well defined. Multiplying two related pairs proves that it preserves multiplication. It is onto by the definition of B'. If $b \in M$, multiply (a, b) by $(1, b^{-1}) \in R$ to obtain $(a, 1) \in R$, so $a \in N$. Hence its kernel is trivial and it is injective. We have constructed the comparison isomorphism directly.

The decisive operation is relative difference. For example, three pairs $(a,b),(a',b),(a',b') \in R$ force the fourth because

$$(a,b)(a',b)^{-1}(a',b') = (a,b') \in R.$$

Thus a relation that need not be a function becomes a function after quotienting exactly the ambiguities on its two sides. This is a usable group calculation in its own right. The following definitions isolate the matching, forgetting, and quotient operations so that the same proof can be used for other algebraic objects. A reader concerned only with groups can keep the displayed formulas as the working model.

A *category* consists of objects, sets of maps between objects, identity maps, and associative composition. A diagram commutes when every indicated pair of composites with the same endpoints agrees. A map m is a *monomorphism* if $mu = mv$ implies $u = v$; a map e is an *epimorphism* if $ue = ve$ implies $u = v$. These are cancellation conditions on maps, not definitions in terms of elements.

A *subobject* of X is a monomorphism into X, with two such monomorphisms identified if an isomorphism over X relates them. We write $A \le B \le X$ when the map representing A factors through the one representing B. Products are characterized by their projections: a map to $X \times Y$ is uniquely a pair of maps to X and Y.

Definition 1.1 (Pullback)**.** For $f : X \to Z$ and $g : Y \to Z$, a pullback is an object P with maps $a : P \to X$, $b : P \to Y$ such that $fa = gb$, universal with that property: every pair $u : T \to X$, $v : T \to Y$ with $fu = gv$ has a unique map $t : T \to P$ satisfying $at = u$ and $bt = v$.

$$\begin{array}{ccc} P & \xrightarrow{\ b\ } & Y \\ {\scriptstyle a}\downarrow & & \downarrow{\scriptstyle g} \\ X & \xrightarrow[\ f\ ]{} & Z. \end{array}$$

We write $P = X \times_Z Y$.

In sets this is the set of pairs with the same image in Z. In groups and modules it has the same description with coordinatewise operations. The universal property retains the equation $fu = gv$, without selecting coordinates. A terminal object receives exactly one map from each object. An initial object has exactly one map to each object. An equalizer of $u, v : X \to Y$ is universal among maps $a : T \to X$ with $ua = va$. A finite-limit category has a terminal object, finite products, and equalizers, hence all the pullbacks used here: $X \times_Z Y$ is the equalizer of $f\pi_X, g\pi_Y : X \times Y \to Z$. The *kernel pair* of $f : X \to Y$ is $X \times_Y X \rightrightarrows X$: it records exactly which inputs f identifies.

A *coequalizer* of $u, v : K \rightrightarrows X$ is a map $q : X \to Q$ with $qu = qv$ such that every map $h : X \to Z$ satisfying $hu = hv$ factors uniquely as $h = \bar{h}q$. A *regular epimorphism* is a coequalizer of some pair. Every regular epimorphism is an epimorphism, by the uniqueness part of this property. We draw it as $X \twoheadrightarrow Q$.

Lemma 1.2 (The elementary descent rule)**.** *In a category with pullbacks, a regular epimorphism $q : X \twoheadrightarrow Q$ is the coequalizer of its kernel pair. It has the following lifting property against any monomorphism $m : A \hookrightarrow B$: in a commutative square*

$$\begin{array}{ccc} X & \xrightarrow{\ a\ } & A \\ {\scriptstyle q}\downarrow & \overset{d}{\nearrow} & \downarrow{\scriptstyle m} \\ Q & \xrightarrow[\ b\ ]{} & B, \end{array}$$

there is a unique d with $dq = a$ and $md = b$.

Proof. Choose a pair u, v of which q is a coequalizer. Since $qu = qv$, this pair factors through the kernel pair of q. A map equal on that kernel pair is therefore equal on u, v, and descends uniquely through q. This proves the first claim.

For the square, $mau = bqu = bqv = mav$, and monicity of m gives $au = av$. Hence $a = dq$ for a unique d. Now $mdq = ma = bq$, so $md = b$ because q is epic. The equation $dq = a$ also proves uniqueness. □

In particular, a regular epimorphism that is also a monomorphism is an isomorphism: apply the lifting property to its square with itself and identity maps. This argument uses neither a choice of representatives nor a splitting assumption.

Definition 1.3 (Regular category)**.** A category is *regular* if it has finite limits, every map factors as a regular epimorphism followed by a monomorphism, and pullbacks of regular epimorphisms are regular epimorphisms. The monomorphism in

$$X \twoheadrightarrow \mathrm{im}(f) \hookrightarrow Y$$

is called the *image* of $f : X \to Y$.

Lemma 1.4 (Image calculus)**.** *In a regular category, images are unique as subobjects, regular epimorphisms compose, and pulling back an image factorization gives an image factorization.*

Proof. Given two regular-epimorphism–monomorphism factorizations of a map, Lemma 1.2 gives a map between their middle objects in each direction. Its uniqueness says the two composites are identities. This proves uniqueness of images.

The lifting property in Lemma 1.2 is preserved under composition: first lift against the first regular epimorphism and then against the second. Factor their composite as mp, where p is a regular epimorphism and m is monic. The lifting property of the composite, applied against m, gives a right inverse of m. A monic map with a right inverse is an isomorphism. The composite is therefore a regular epimorphism.

Pull back the mono part of a factorization, and then pull back its regular epi part. The resulting square is a pullback of the composite, by the defining universal properties. Pullback stability makes the first part a regular epimorphism; pullbacks of monomorphisms are monic because maps into the pullback are determined by their two projections. This is the required image factorization. □

Images express what survives a map. Pullback stability says that this description remains valid after imposing a compatible condition at the target. It is the reason local witnesses can safely be used below.

1.3 A precise license to reason with witnesses

An arbitrary category may have too few maps from its terminal object to detect equality. In particular, a group has just one group homomorphism from the trivial group. Thus an argument about ordinary global points would miss nearly everything. We instead use *generalized elements*: an input x of X over a test object T means a map $x : T \to X$.

For a subobject $M \hookrightarrow X$, saying that x belongs to M means that x factors through that monomorphism. A *cover* below means a regular epimorphism, not an assumed topological covering.

Lemma 1.5 (Local witness calculus)**.** *Let $f : A \to X$ be a map in a regular category. A generalized element $x : T \to X$ belongs to* $\mathrm{im}(f)$ *if and only if there is a cover $e : T' \twoheadrightarrow T$ and a map $a : T' \to A$ such that $fa = xe$. Finite lists of such witnesses can be chosen over a common cover. They can be pulled back along any map of test objects.*

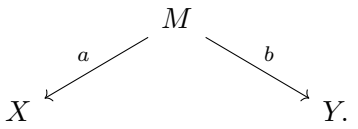
Proof. Factor $f = mp$ with $p : A \twoheadrightarrow I$ regular epic and $m : I \hookrightarrow X$ monic. If $x = mi$, pull back p along $i : T \to I$. The pullback cover has a map a with $fa = xe$. Conversely, such e, a give a square against m; Lemma 1.2 makes x factor through m.

Pulling back a cover gives a cover. To combine two choices over T, take their fiber product over T; either projection is a cover, and its composite to T is a cover by Lemma 1.4. Repeat for finitely many choices. Nested choices are handled by composing covers. □

This is the precise meaning of "there exists" in our relation calculations. A witness may exist only after a cover of the test object. Every proof using witnesses below permits that replacement. Equalities of subobjects are checked by inclusions in both directions: to check $M \leq N$, one may test the generalized element $M \hookrightarrow X$ itself. The method therefore proves actual categorical statements, not merely assertions about points.

1.4 Relations compose by matching, then forgetting

Definition 1.6 (Internal relation)**.** A relation $R : X \rightsquigarrow Y$ is a subobject of $X \times Y$. Choose a representative with projections



The pair (a, b) is jointly monic. We may write $x\,R\,y$ for membership in this subobject. Its opposite $R^{\mathrm{op}} : Y \rightsquigarrow X$ interchanges the projections. The identity relation Δ_X is the diagonal $X \hookrightarrow X \times X$.

If $S : Y \rightsquigarrow Z$ has carrier N, form $M \times_Y N$ and take its image in $X \times Z$. This is the composite SR. Thus

$$x\,(SR)\,z \quad \Longleftrightarrow \quad \text{locally there is } y \text{ with } x\,R\,y \text{ and } y\,S\,z.$$

Our convention is always that SR means first R, then S. The pullback retains the matching intermediate value; the image forgets which matching value was used.

Proposition 1.7 (The calculus of relations)**.** *In a regular category, relation composition is associative and has the diagonal relations as identities. It preserves inclusions, and*

$$(SR)^{\mathrm{op}} = R^{\mathrm{op}} S^{\mathrm{op}}.$$

Proof. For $R : X \rightsquigarrow Y$, $S : Y \rightsquigarrow Z$, and $T : Z \rightsquigarrow W$, membership in either $(TS)R$ or $T(SR)$ means that, over a cover, there are y, z with $x\,R\,y$, $y\,S\,z$, $z\,T\,w$. Lemma 1.5 combines any nested witnesses into one cover and also allows the same triple witness to be grouped in either order. This proves

equality of the two image subobjects. Equivalently, both are the image of $M_R \times_Y M_S \times_Z M_T$ in $X \times W$.

A diagonal intermediate relation forces the two matched coordinates to agree and introduces no further condition; hence the unit laws. An inclusion of relations sends every witness to a witness, proving monotonicity. Reversing a witness exchanges its endpoints and reverses the order of the two conditions, proving the last identity. □

This associativity has the same shape as elementary reassociation, but it applies to correspondences that may not be functions. Its proof identifies one object of complete witness data and shows that the order of forgetting intermediate information does not matter.

A map $f : X \to Y$ is identified with its graph relation, represented by $(\mathrm{id}_X, f) : X \hookrightarrow X \times Y$. The opposite graph is generally not a map: it remembers all possible preimages.

Lemma 1.8 (Graphs, kernels, and supports)**.** *For a map $f : X \to Y$,*

$$f^{\mathrm{op}} f = \mathrm{Eq}(f), \qquad f f^{\mathrm{op}} = \Delta_{\mathrm{im}(f)} \quad \text{as a relation on } Y.$$

Here $\mathrm{Eq}(f) = X \times_Y X$ is the kernel pair, and the second diagonal is supported on the image subobject of Y. Consequently

$$f f^{\mathrm{op}} f = f, \qquad f f^{\mathrm{op}} = \Delta_Y \text{ if } f \text{ is regular epic.}$$

Every relation with projections $a : M \to X$, $b : M \to Y$ satisfies

$$R = b a^{\mathrm{op}}.$$

Proof. The first composite relates x, x' exactly when $f(x) = f(x')$, which is the defining pullback. The second relates y, y' exactly when, locally, $y = f(x) = y'$ for some x; this is the image diagonal by Lemma 1.5. These descriptions give the two consequences. Finally ba^{op} is the image of (a, b) in $X \times Y$; that pair already represents the subobject R. □

Lemma 1.9 (When a relation is a map)**.** *A relation $R : X \rightsquigarrow Y$ is the graph of a unique map $f : X \to Y$ if*

$$\Delta_X \leq R^{\mathrm{op}} R, \qquad R R^{\mathrm{op}} \leq \Delta_Y.$$

The first condition says that every input has a locally available output; the second says that the output is unique.

Proof. Write $R = ba^{\mathrm{op}}$ with carrier M. The first condition says that the image of $a : M \to X$ is all of X, so a is regular epic. If $u, v : T \to M$ have $au = av$, then bu and bv are two outputs for the same input. The second condition gives $bu = bv$. Joint monicity of (a, b) gives $u = v$, so a is also monic and hence an isomorphism. Now $f = ba^{-1}$ has graph R. Its value is forced by that graph, which proves uniqueness. □

1.5 Exactness makes an identification into an object

A relation $E : X \rightsquigarrow X$ is *reflexive* if $\Delta_X \leq E$, *symmetric* if $E^{\mathrm{op}} = E$, and *transitive* if $EE \leq E$. An *internal equivalence relation* has all three properties. Reflexivity also gives $E \leq EE$, so an equivalence relation satisfies $EE = E$.

Kernel pairs are equivalence relations. For example, transitivity says that $f(x) = f(y)$ and $f(y) = f(z)$ imply $f(x) = f(z)$; this is equally an equality of maps over any test object.

Definition 1.10 (Barr exactness)**.** A regular category is *Barr-exact* if every internal equivalence relation is the kernel pair of a map. Equivalently, for every such E on X there is a regular epimorphism $q : X \twoheadrightarrow Q$ with $E = q^{\mathrm{op}}q$. We call Q the quotient X/E.

For the equivalence of these formulations, factor a map exhibiting the kernel pair into its image. The monomorphism part does not change its kernel pair, and the other part is regular epic. Lemma 1.2 then gives the quotient's full universal property: maps from X that identify the pairs in E descend uniquely to Q. Two such quotients are uniquely isomorphic over X, by applying this property twice. Exactness concerns these equivalence-relation quotients; it does not assert that the category is abelian or that its maps can be added.

A category is *pointed* if it has an object 0 that is both initial and terminal. There is then a zero map between every two objects, namely the composite through 0. A *kernel* of $f : X \to Y$ is the pullback of $0 \to Y$ along f:

$$\begin{array}{ccc} \ker(f) & \longrightarrow & 0 \\ \downarrow & & \downarrow \\ X & \xrightarrow[f]{} & Y. \end{array}$$

The left map is monic since $0 \to Y$ is monic. In a pointed category it is a distinguished subobject, but it is only one class of the whole kernel pair.

If $E = q^{\mathrm{op}}q$, its zero class is the relation $E0_X$, viewed as a subobject of X, where $0_X : 0 \to X$. Since $q0_X = 0_Q$,

$$E0_X = q^{\mathrm{op}}0_Q = \ker(q).$$

It is saturated: $E(E0_X) = E0_X$. We will retain both the full equivalence relation and its zero class. In groups the zero class determines the normal quotient, but we will not build that additional fact into the general framework.

The roles of the assumptions are now visible. Finite limits impose compatible equations. Images forget witnesses. Regularity makes that forgetting compatible with substitution. Exactness realizes consistent identifications as quotient objects. Pointedness selects a zero class. The next condition will make different identifications compatible with one another.

2 The three-corner rule and one comparison

Suppose a relation connects inputs in X with outputs in Y. Three corners of a rectangle may already force the fourth:

$$x\,R\,y, \qquad x'\,R\,y, \qquad x'\,R\,y' \quad \Longrightarrow \quad x\,R\,y'.$$

This property is called *difunctionality*. It says that the relation breaks into complete rectangular blocks. It does not say that each input has just one output. Several inputs and several outputs can belong to the same block.

The structural comparison is already suggested by this picture. Group the inputs according to which block they occupy, and group the outputs in the same way. The two resulting quotient objects should agree. We will construct that agreement, including its maps, before using it to compare any filtrations.

2.1 Why relative difference is enough

Definition 2.1 (Mal'cev category)**.** A regular category is a *Mal'cev category* if every reflexive internal relation is an equivalence relation.

This categorical condition has an elementary source. Suppose we have algebraic structures with a ternary term p satisfying

$$p(x, y, y) = x, \qquad p(x, x, y) = y.$$

If a compatible relation $R \subseteq X \times X$ is reflexive and $x\,R\,y$, its closure under p gives

$$p\big((y, y), (x, y), (x, x)\big) = (y, x) \in R.$$

Thus it is symmetric. If also $y\,R\,z$, closure gives

$$p\big((x, y), (y, y), (y, z)\big) = (x, z) \in R,$$

so it is transitive. The operation is applied coordinatewise to pairs. For groups $p(x, y, z) = xy^{-1}z$; for modules it is $x - y + z$. No commutativity of group multiplication was used. A term is an operation on the underlying algebra; in the group case it need not be a group homomorphism from X^3 to X.

The group calculation in Section 1 used precisely this move: take a relative difference, then reattach it at another point. In a module the corresponding calculation is $(x, y) - (x', y) + (x', y') = (x, y')$. The categorical condition records the rectangle-completion effect shared by these calculations. The next proposition proves that it supplies the same reusable rule even when there is no chosen formula for relative difference.

Proposition 2.2 (Three forms of the Mal'cev condition)**.** *For a regular category the following conditions are equivalent:*

1. *every reflexive relation is an equivalence relation;*
2. *equivalence relations on the same object commute under relation composition;*
3. *every relation $R : X \rightsquigarrow Y$ is difunctional, that is,*

$$RR^{\mathrm{op}}R = R.$$

Proof. Assume the first condition. If E, F are equivalence relations on X, their composite EF is reflexive, since both contain Δ_X. It is therefore symmetric, and

$$EF = (EF)^{\mathrm{op}} = FE.$$

This proves the second condition.

Assume the second. Give R a carrier with projections $a : M \to X$ and $b : M \to Y$. The kernel pairs

$$E_a = a^{\mathrm{op}}a, \qquad E_b = b^{\mathrm{op}}b$$

are equivalence relations on the *same* object M, so they commute. Using Lemma 1.8, we obtain

$$\begin{aligned} RR^{\mathrm{op}}R &= ba^{\mathrm{op}}ab^{\mathrm{op}}ba^{\mathrm{op}} \\ &= bE_aE_ba^{\mathrm{op}} = bE_bE_aa^{\mathrm{op}} \\ &= (bb^{\mathrm{op}}b)(a^{\mathrm{op}}aa^{\mathrm{op}}) = ba^{\mathrm{op}} = R. \end{aligned}$$

The equalities are typed on M; we have not interchanged unrelated relations on X and Y. This proves the third condition.

Finally assume the third and let R be reflexive. Monotonicity gives

$$R^{\mathrm{op}} = \Delta_X R^{\mathrm{op}} \Delta_X \leq R R^{\mathrm{op}} R = R.$$

Taking opposites gives the reverse inclusion, so R is symmetric. Also $\Delta_X \leq R^{\mathrm{op}}$, and hence $RR \leq RR^{\mathrm{op}}R = R$. Thus R is transitive and is an equivalence relation. □

The local-witness interpretation of $RR^{\mathrm{op}}R \leq R$ is exactly the three-corners implication above. The reverse inclusion always holds: a related pair can supply all three edges of a degenerate rectangle. The proposition explains why the rectangle-completion law holds simultaneously for every relation, rather than having to be verified anew for each quotient comparison.

2.2 Partial identifications and their supports

Relations used in a filtration may live on a subobject rather than all of the ambient object. We must record that support instead of silently assuming reflexivity everywhere.

Lemma 2.3 (Partial equivalence relations)**.** *Let $E : X \rightsquigarrow X$ be symmetric and transitive in a regular category. The two projections of its carrier have the same image $m : D \hookrightarrow X$. Restricted to D, the relation is an equivalence relation E_D. In a Barr-exact category there is a regular epimorphism $q : D \twoheadrightarrow Q$ with*

$$E = m q^{\mathrm{op}} q m^{\mathrm{op}}.$$

Conversely, every relation of this displayed form is symmetric and transitive, with support D.

Proof. Symmetry interchanges the two projections, so their images agree. Their factorizations through D give a relation $E_D \hookrightarrow D \times D$; the pair is still jointly monic. Both of its projections are regular epic. If x is an input in D, a cover supplies a witness $x\, E_D\, y$. Symmetry gives $y\, E_D\, x$, and transitivity gives $x\, E_D\, x$. Lemma 1.5 descends this membership, so E_D is reflexive. Symmetry and transitivity are inherited by restriction. Exactness gives $E_D = q^{\mathrm{op}} q$, proving the formula.

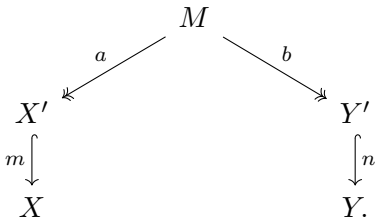
Conversely, $q^{\mathrm{op}} q$ is an equivalence relation on D. Its extension along m is symmetric; since $m^{\mathrm{op}} m = \Delta_D$, its square is itself. The projections onto D are regular epic because the relation is reflexive, so their images in X are D. □

We call Q the quotient of the partial equivalence relation. It is the quotient of its support, not a quotient of all of X. This distinction permits the same relation language to describe every step of a chain whose terms are only normal in their successors.

For example, a quotient map $q : A \twoheadrightarrow S$ with $A \hookrightarrow G$ determines a relation on G: two inputs are related precisely when they both lie in A and have the same image in S. Its support is A; in a pointed category its zero class is $\ker(q)$. Thus one partial relation remembers the numerator, the identifications, and the chosen quotient, up to its unique quotient isomorphism.

2.3 One relation, two quotients, one object

Let $R : X \rightsquigarrow Y$ have carrier M and factor its projections through their images:



The pair (a, b) is jointly monic. We use $\widehat{R} : X' \rightsquigarrow Y'$ for the relation on these supports, so $R = n\widehat{R}m^{\mathrm{op}}$. Both projections of $\widehat{R}$ are regular epic.

Theorem 2.4 (The universal support-quotient comparison)**.** *In a Barr-exact Mal'cev category, put*

$$E = \widehat{R}^{\mathrm{op}}\widehat{R} \quad \text{on } X', \qquad F = \widehat{R}\widehat{R}^{\mathrm{op}} \quad \text{on } Y'.$$

These are equivalence relations. For their regular quotient maps $q : X' \twoheadrightarrow Q = X'/E$ and $p : Y' \twoheadrightarrow P = Y'/F$, there is a unique isomorphism $\theta : Q \xrightarrow{\sim} P$ satisfying

$$\theta qa = pb.$$

In particular the two support quotients of any relation are canonically isomorphic.

Proof. Because a is regular epic, every input in X' has locally some $\widehat{R}$-related output. Using that same output twice proves $\Delta_{X'} \leq E$. Likewise $\Delta_{Y'} \leq F$. Both relations are symmetric by taking opposites. Difunctionality gives

$$E^2 = \widehat{R}^{\mathrm{op}}(\widehat{R}\widehat{R}^{\mathrm{op}}\widehat{R}) = E, \qquad F^2 = (\widehat{R}\widehat{R}^{\mathrm{op}}\widehat{R})\widehat{R}^{\mathrm{op}} = F.$$

Thus they are equivalence relations, and exactness supplies q, p.

We construct the map rather than merely matching equivalence classes. The map $pb : M \to P$ is equal on the kernel pair of a: if $au = av$, then bu and bv share a $\widehat{R}$-related input, so they are F-related and $pbu = pbv$. Lemma 1.2 gives a unique $g : X' \to P$ with $ga = pb$.

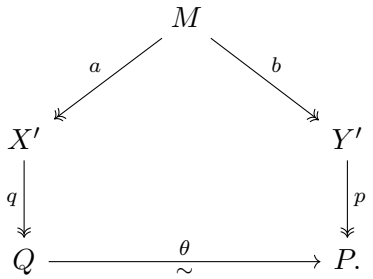
Next g is equal on E. An E-related pair is locally of the form $a(u), a(v)$ with $b(u) = b(v)$. Consequently $ga(u) = pb(u) = pb(v) = ga(v)$. These equalities hold after a cover, and covers are epimorphisms, so they hold on the relation itself. Hence g descends uniquely through q, giving $\theta : Q \to P$.

The same construction with the opposite relation gives $\psi : P \to Q$ such that $\psi pb = qa$. It follows that

$$\psi\theta qa = qa, \qquad \theta\psi pb = pb.$$

Both qa and pb are composites of regular epimorphisms, hence are epic. Cancellation proves $\psi\theta = \mathrm{id}_Q$ and $\theta\psi = \mathrm{id}_P$. Finally qa is epic, so the equation specifying θ proves its uniqueness. □

The construction is summarized by the commutative diagram



There is generally no map $X' \to Y'$ to put across the middle row. The relation supplies multiple possible outputs. The quotient removes precisely their ambiguity, after which the comparison becomes a uniquely determined isomorphism.

For a familiar model, take the additive group $\mathbb{Z}$ and the relation xRy meaning that $x - y$ is even. This is a subgroup of $\mathbb{Z} \times \mathbb{Z}$, with both supports equal to $\mathbb{Z}$. Two inputs share an output precisely when they have the same parity; the same is true of two outputs sharing an input. Thus both support quotients are $\mathbb{Z}/2\mathbb{Z}$, and the comparison sends the parity class of an input to that of any related output. There are infinitely many possible outputs before quotienting, but exactly one class afterwards. This illustrates the construction; the categorical theorem did not need a preferred output or a choice of representatives.

Here is another view of the same map, useful for computations. Its graph is

$$L = p\widehat{R}q^{\mathrm{op}} : Q \rightsquigarrow P.$$

Indeed,

$$L^{\mathrm{op}}L = q\widehat{R}^{\mathrm{op}}p^{\mathrm{op}}p\widehat{R}q^{\mathrm{op}} = qE^2q^{\mathrm{op}} = \Delta_Q,$$
$$LL^{\mathrm{op}} = p\widehat{R}q^{\mathrm{op}}q\widehat{R}^{\mathrm{op}}p^{\mathrm{op}} = pF^2p^{\mathrm{op}} = \Delta_P.$$

We used $q^{\mathrm{op}}q = E$, $p^{\mathrm{op}}p = F$, their idempotence, and $qq^{\mathrm{op}} = \Delta_Q$, $pp^{\mathrm{op}} = \Delta_P$. Lemma 1.9 makes L a map, and makes L^{op} its inverse. This is the same θ because its graph relates the quotient images of every pair in $\widehat{R}$.

This theorem is the comparison mechanism from which the refinement rectangle will be built. It is stronger than a statement about lengths: it constructs the actual isomorphism between the surviving objects. No filtration, simplicity assumption, or Jordan–Hölder theorem has entered its proof.

2.4 Which relations descend through a quotient?

There is one further compatibility to make explicit. If we identify inputs by a quotient map, a relation descends only when changing an input within its equivalence class does not change the relation. This is saturation.

Proposition 2.5 (Saturation and descent)**.** *Let $q : X \twoheadrightarrow Q$ and $p : Y \twoheadrightarrow P$ be regular epimorphisms, with kernel pairs $E = q^{\mathrm{op}}q$ and $F = p^{\mathrm{op}}p$. There is an order-preserving bijection*

$$\{\textit{relations } L : Q \rightsquigarrow P\} \longleftrightarrow \{R : X \rightsquigarrow Y \mid FRE = R\}.$$

It sends L to $p^{\mathrm{op}}Lq$, and its inverse sends R to pRq^{op}. The first construction is precisely inverse image of the relation along $q \times p$.

Proof. Inverse image consists of pairs (x, y) whose quotient images satisfy $q(x)\, L\, p(y)$, which is $p^{\mathrm{op}} L q$ by the graph formulas. Since $qq^{\mathrm{op}} = \Delta_Q$ and $pp^{\mathrm{op}} = \Delta_P$, the two constructions satisfy

$$p(p^{\mathrm{op}} L q) q^{\mathrm{op}} = L, \qquad p^{\mathrm{op}}(pRq^{\mathrm{op}})q = FRE.$$

An inverse-image relation is saturated because $Fp^{\mathrm{op}} = p^{\mathrm{op}}$ and $qE = q$. Thus the formulas are mutually inverse exactly on the stated saturated relations. Both are order preserving by monotonicity of relation composition. □

For a relation T on X and a single quotient q, saturation is $ETE = T$. Since E contains the diagonal, this implies separately $ET = T$ and $TE = T$. Consequently the correspondence respects composition of saturated relations:

$$(qSq^{\mathrm{op}})(qTq^{\mathrm{op}}) = qSETq^{\mathrm{op}} = qSTq^{\mathrm{op}}.$$

It also respects opposites. Thus a saturated symmetric transitive relation descends to a symmetric transitive relation, including its possibly proper support. Conversely, inverse image preserves these properties: for example

$$(q^{\mathrm{op}} L q)^2 = q^{\mathrm{op}} L^2 q \leq q^{\mathrm{op}} L q$$

when L is transitive. These statements hold for partial equivalence relations as well as for equivalence relations on all of X.

Corollary 2.6 (Quotients of quotients)**.** *In a Barr-exact category, fix $q : X \twoheadrightarrow Q$ with kernel pair E. Equivalence relations L on Q correspond to equivalence relations T on X containing E, by*

$$T = q^{\mathrm{op}} L q, \qquad L = qTq^{\mathrm{op}}.$$

If $t : X \twoheadrightarrow Z$ is the quotient of T, it factors uniquely as $t = uq$, where $u : Q \twoheadrightarrow Z$ is the quotient of L. Thus $(X/E)/L \cong X/T$, with the isomorphism determined by the quotient maps.

Proof. If T is an equivalence relation containing E, then $T \leq ETE \leq TTT = T$, so it is saturated. Its descent is reflexive because $q\Delta_X q^{\mathrm{op}} = \Delta_Q$, and is symmetric and transitive by the preceding formulas. Conversely, the inverse image of an equivalence relation is an equivalence relation containing $q^{\mathrm{op}}\Delta_Q q = E$.

Since $E \leq T$, the map t is equal on the kernel pair of q, so it descends uniquely to u. Factor u into its image. If its mono part were a proper subobject of Z, then $t = uq$ would factor through that same subobject, contradicting the lifting property of the regular epimorphism t. Hence u is regular epic. Its kernel pair is

$$u^{\mathrm{op}} u = qq^{\mathrm{op}} u^{\mathrm{op}} uqq^{\mathrm{op}} = qt^{\mathrm{op}} tq^{\mathrm{op}} = qTq^{\mathrm{op}} = L.$$

Lemma 1.2 makes it the quotient of L. The universal property proves the claimed canonical isomorphism. □

This last result needs exactness and regularity, but not the Mal'cev condition. Recording that distinction matters: Mal'cev is used to make *every* relation difunctional and hence to produce the universal comparison. The ordinary law for compatible iterated quotients has already been supplied by exactness. Later, in a higher setting, we will see why retaining such compatibility data remains essential.

The relation calculus and support comparison are classical. For context, see Knop, *Tensor envelopes of regular categories*, Section 5 and Appendices A–B. The arguments here have been written out so that the refinement theorem can depend on explicit constructions rather than on an unnamed isomorphism theorem.

3 A quotient remembers which identifications were made

We now have a language in which a quotient is an operation on a relation, and a comparison of quotients is an operation on a correspondence. We will use that language once, in a pointed Barr-exact Mal'cev category $\mathcal{C}$. Groups and modules will enter afterwards by verifying these axioms. Neither commutative addition of arrows nor an invariant called dimension is used in this chapter.

There are three stages. First we record a normal series without forgetting its quotient maps. Next we compare two such records by a rectangular array of correspondences. Finally simplicity prevents the comparison from breaking a factor into smaller nonzero factors. The last step is Jordan–Hölder; the rectangle, rather than a numerical count, does the work.

3.1 The data in a normal series

Definition 3.1 (Normal series with its quotient maps)**.** A finite normal series of X consists of a chain of subobjects and chosen regular epimorphisms

$$0 = X_0 \hookrightarrow X_1 \hookrightarrow \cdots \hookrightarrow X_n = X, \qquad q_i : X_i \twoheadrightarrow S_i,$$

such that the inclusion $X_{i-1} \hookrightarrow X_i$ is a kernel of q_i. The objects S_i are its *factors*. Repetitions of subobjects and zero factors are permitted. Two series are *equivalent* when their lists of factors are isomorphic after a permutation. We also use equivalence *after deleting zero factors*, stating explicitly when we do so.

The kernel square in this definition is

$$\begin{array}{ccc} X_{i-1} & \hookrightarrow & X_i \\ \downarrow & & \downarrow{\scriptstyle q_i} \\ 0 & \longrightarrow & S_i. \end{array} \tag{3.1a}$$

It is a pullback. In particular, the normality assertion is relative to X_i, not to all of X. For a group this allows a subgroup normal in its successor without demanding that it be normal in the ambient group.

We keep q_i because the kernel alone need not determine a quotient in this generality. A familiar implication from groups, such as "a regular epimorphism with zero kernel is an isomorphism", is not inserted as an unstated axiom. All our comparisons retain the kernel pair of the actual map, which is the full record of its identifications.

Here and below, composition of relations is function-like: SR means first R, then S. For a relation $r : X \rightsquigarrow X$ and a subobject $U \hookrightarrow X$, write rU for the image of the restriction of r whose input lies in U. Thus rX is its output support and $r0$ its zero class. A symmetric transitive relation is a

partial equivalence relation; it is an equivalence relation on its support, by Lemma 2.3. Its support quotient will be written

$$Q(r) = (rX)/r.$$

The slash is shorthand for the effective quotient of the restricted equivalence relation, not a new operation on objects without relations.

Lemma 3.2 (A series as a sequence of identifications)**.** *For $n \geq 1$, the data of Definition 3.1, up to the unique isomorphisms allowed for quotient objects, are equivalently a list of partial equivalence relations $r_1, \ldots, r_n$ on X satisfying*

$$r_1 0 = 0, \qquad r_i X = r_{i+1} 0 \ (1 \leq i < n), \qquad r_n X = X. \tag{1}$$

The associated factors are $Q(r_i)$. The empty series of 0 is treated separately; it has no endpoint relations to name.

Proof. Let $a_i : X_i \hookrightarrow X$ be the composite inclusion. Regarding arrows as their graphs, put

$$r_i = a_i q_i^{\mathrm{op}} q_i a_i^{\mathrm{op}}.$$

On X_i, this is the kernel pair of q_i; it relates precisely the two inputs that q_i identifies. Extending it along the monomorphism a_i makes it a partial equivalence relation on X. Its support is X_i. The zero arrow into X_i is sent to the zero arrow into S_i, so taking the zero class of that kernel pair gives precisely the pullback defining $\ker q_i$. Therefore

$$r_i X = X_i, \qquad r_i 0 = X_{i-1}.$$

These equalities give (1). Since q_i is a regular epimorphism, it is the effective quotient of its kernel pair, and $Q(r_i) \cong S_i$ canonically.

Conversely, set $X_i = r_i X$. The zero class $r_i 0$ is a subobject of $r_i X$. Thus the identities (1) make the X_i a chain, beginning with $X_0 = 0$ and ending with $X_n = X$. By exactness the restricted relation r_i has an effective quotient $q_i : X_i \twoheadrightarrow Q(r_i)$. The same kernel-pair calculation says that $\ker q_i = r_i 0 = X_{i-1}$. These two constructions are inverse because an effective quotient is determined, up to a unique compatible isomorphism, by its kernel pair. □

This change of language has retained more than the chain. It retains which pairs of points become equal at each step. That is the information needed to compare two quotient constructions in a category where the zero class alone may not determine the entire relation.

3.2 What a genuine refinement means

To refine a series, we refine one of its factors and pull that refinement back. Suppose $q : D \twoheadrightarrow Q$ is one quotient step, with kernel $K \hookrightarrow D$, and suppose

$$0 = Q_0 \hookrightarrow Q_1 \hookrightarrow \cdots \hookrightarrow Q_t = Q, \qquad e_j : Q_j \twoheadrightarrow Z_j, \quad \ker e_j = Q_{j-1},$$

is a normal series of its factor. Form $P_j = D \times_Q Q_j$. We obtain

$$\begin{array}{ccccccc} K = P_0 & \hookrightarrow & P_1 & \hookrightarrow & \cdots & \hookrightarrow & P_t = D \\ \downarrow & & \downarrow{\scriptstyle p_1} & & & & \downarrow{\scriptstyle q} \\ 0 = Q_0 & \hookrightarrow & Q_1 & \hookrightarrow & \cdots & \hookrightarrow & Q_t = Q. \end{array}$$

Each displayed comparison square is a pullback. Pullback stability of regular epimorphisms makes $p_j : P_j \twoheadrightarrow Q_j$ a regular epimorphism. The refined quotient map at the jth step is

$$e_j p_j : P_j \twoheadrightarrow Z_j. \tag{2}$$

Its kernel is P_{j-1}: its pullback over 0 is $P_j \times_{Q_j} Q_{j-1} = D \times_Q Q_{j-1}$. Regular epimorphisms are closed under composition, so this really is a normal series between K and D. Replacing any or all of the original steps in this way is what we mean by a *refinement*. For the empty series of the zero object we adopt the harmless convention that inserting a finite list of zero steps is a refinement. This convention concerns only how empty lists are represented; it introduces no nonzero factor.

It would be insufficient to insert arbitrary kernel subobjects between K and D while forgetting their relation to q. We need the new quotients to have arisen from the original quotient. The following saturation condition expresses exactly that requirement.

Proposition 3.3 (Saturation is descent of a refinement)**.** *Let r be a partial equivalence relation on X, let $D = rX$, and let $q : D \twoheadrightarrow Q(r)$ be its quotient. For $k \geq 1$, lists of partial equivalence relations $t_1, \ldots, t_k$ on X satisfying*

$$rt_j r = t_j \quad (1 \leq j \leq k), \tag{3}$$

$$t_1 0 = r0, \qquad t_j X = t_{j+1} 0 \ (1 \leq j < k), \qquad t_k X = rX \tag{4}$$

correspond to normal series of $Q(r)$ with exactly k steps. Under this correspondence, their support quotients $Q(t_j)$ are the factors of the pulled-back refinement of q.

Proof. The equality $rt_j r = t_j$ first implies that the support of t_j lies inside D, so we can work on D. Write $E = q^{\mathrm{op}} q$ for the kernel pair on D; it is the restriction of r. The saturated descent lemma, Proposition 2.5, gives mutually inverse operations

$$T_j = q\, t_j\, q^{\mathrm{op}}, \qquad t_j = q^{\mathrm{op}} T_j q.$$

The first is the direct image relation on $Q(r)$; the second is its pullback along $q \times q$. Their being inverse uses $qq^{\mathrm{op}} = 1_{Q(r)}$ and $Et_jE = t_j$. That lemma also proves that symmetry and transitivity descend, so each T_j is a partial equivalence relation.

We check the endpoints, since those checks are what make this a descent of *series*. Saturation and reflexivity of E on D give $t_j E = t_j = E t_j$. The inverse image of the zero subobject of $Q(r)$ is $E0$. Therefore, as subobjects obtained by regular images,

$$T_j 0 = q\, t_j E 0 = q\, t_j 0, \qquad T_j Q(r) = q(t_j D).$$

Now $q(E0) = 0$ and $qD = Q(r)$. Applying these formulas to (4) gives

$$T_1 0 = 0, \qquad T_j Q(r) = T_{j+1} 0, \qquad T_k Q(r) = Q(r).$$

Lemma 3.2 therefore supplies a normal series of $Q(r)$.

In the other direction, begin with a normal series of $Q(r)$ and its relations T_j. Pulling back the relations gives $t_j = q^{\mathrm{op}} T_j q$, which are saturated and have precisely the support and zero-class boundaries in (4). Concretely, if Q_j is the support of T_j, the support of t_j is $P_j = D \times_{Q(r)} Q_j$. If $e_j : Q_j \twoheadrightarrow Z_j$ is its factor quotient, the kernel pair of $e_j p_j$ on P_j is t_j. As $e_j p_j$ is a regular epimorphism, its codomain Z_j is the effective quotient of t_j. This gives the unique compatible isomorphism $Q(t_j) \cong Z_j$ and identifies the quotient maps as well as the objects. □

The proposition is a small descent theorem. The equality $rt_j r = t_j$ says that changing either input by an identification already imposed by r does not change the new relation. Exactly then the relation can be defined downstairs, independently of the choice of a representative. The proof above replaces that verbal test by pullbacks, images, and unique quotient maps.

3.3 The comparison rectangle

Let r, s be partial equivalence relations on the same object. Their composite rs need not be an equivalence relation on all of X. It is nevertheless a relation, so the Mal'cev theorem applies to it. This produces the identities that will make our rectangle fit together.

Lemma 3.4 (Two endpoint identities)**.** *For partial equivalence relations r, s on X,*

$$rsrs = rs, \tag{5}$$

and

$$rsr0 = rs0, \qquad rsrX = rsX. \tag{6}$$

Proof. The relation $P = rs$ is difunctional by Proposition 2.2, so $PP^{\mathrm{op}}P = P$. Symmetry and idempotence of r, s reduce the left side to

$$rssrrs = rsrs,$$

which proves (5).

Every subobject in a pointed category contains its zero arrow. In particular $0 \subseteq r0$ and $0 \subseteq s0$. Monotonicity of relational image and (5) now give

$$rs0 \subseteq rsr0 \subseteq rsrs0 = rs0.$$

Here the second inclusion comes from $r0 \subseteq rs0$, itself obtained by applying r to $0 \subseteq s0$. For supports, $sX \subseteq X$ gives

$$rsX = rsrsX \subseteq rsrX \subseteq rsX;$$

the final inclusion uses $rX \subseteq X$. Both chains give equality. □

Now take two normal series of X of positive lengths m, n, encoded by $r_1, \dots, r_m$ and $s_1, \dots, s_n$. At the crossing of the ith row and jth column, define

$$r_{ij} = r_i s_j r_i, \qquad s_{ji} = s_j r_i s_j, \qquad P_{ij} = r_i s_j. \tag{7}$$

The two cell relations are the two equivalence relations associated with the correspondence P_{ij}:

$$P_{ij} P_{ij}^{\mathrm{op}} = r_{ij}, \qquad P_{ij}^{\mathrm{op}} P_{ij} = s_{ji}.$$

Our composition convention matters here: s_{ji} is on the input support of P_{ij}, and r_{ij} on its output support.

Proposition 3.5 (Every row and column is a genuine refinement)**.** *For fixed i, the list $r_{i1}, \dots, r_{in}$ satisfies Proposition 3.3 for r_i. For fixed j, the list $s_{j1}, \dots, s_{jm}$ satisfies that proposition for s_j.*

Proof. We verify every condition for a row. Symmetry follows by reversing the word $r_i s_j r_i$. Idempotence follows from Lemma 3.4:

$$r_{ij} r_{ij} = r_i s_j r_i s_j r_i = r_i s_j r_i = r_{ij}.$$

Thus each cell is symmetric and transitive. Its saturation is

$$r_i r_{ij} r_i = r_i^2 s_j r_i^2 = r_{ij}.$$

For the first zero class, use $s_1 0 = 0$:

$$r_{i1} 0 = r_i s_1 r_i 0 = r_i s_1 0 = r_i 0.$$

For an intermediate join, use $s_j 0 = s_{j-1} X$:

$$\begin{aligned} r_{ij} 0 &= r_i s_j r_i 0 = r_i s_j 0 = r_i s_{j-1} X \\ &= r_i s_{j-1} r_i X = r_{i,j-1} X \qquad (2 \le j \le n). \end{aligned}$$

For the last support, use $s_n X = X$:

$$r_{in} X = r_i s_n r_i X = r_i s_n X = r_i X.$$

Every equality that removes the trailing r_i is one of (6). We have proved saturation and all boundary conditions, so Proposition 3.3 supplies a pulled-back refinement of the original quotient. Interchanging r and s proves the column assertion. □

The boundary identities are the gluing conditions. They say that the support of one cell is precisely the zero class required to start the next cell. Without these identities we would have an assortment of quotients, not a refinement of a filtration.

Proposition 3.6 (The two readings of a cell have the same quotient)**.** *For every (i, j) there is a canonical isomorphism*

$$Q(s_{ji}) \xrightarrow{\sim} Q(r_{ij}). \tag{8}$$

It is induced by the correspondence P_{ij} and characterized by its two projections, not chosen by comparing lengths or other invariants.

Proof. Let C_{ij} be the object representing $P_{ij} \hookrightarrow X \times X$. Factor its two projections through their image supports:

$$C_{ij} \xrightarrow{a_{ij}} L_{ij} \hookrightarrow X, \qquad C_{ij} \xrightarrow{b_{ij}} R_{ij} \hookrightarrow X.$$

Both a_{ij} and b_{ij} are regular epimorphisms. The input relation $P_{ij}^{\mathrm{op}} P_{ij}$ is s_{ji}, so $L_{ij} = s_{ji} X$; similarly $R_{ij} = r_{ij} X$. Write their effective quotient maps as $\lambda_{ij} : L_{ij} \twoheadrightarrow Q(s_{ji})$ and $\rho_{ij} : R_{ij} \twoheadrightarrow Q(r_{ij})$.

The support-quotient comparison theorem, Theorem 2.4, applies to P_{ij} and gives the unique isomorphism θ_{ij} for which

$$\begin{array}{ccccc} L_{ij} & \xleftarrow{a_{ij}} & C_{ij} & \xrightarrow{b_{ij}} & R_{ij} \\ {\scriptstyle \lambda_{ij}}\downarrow & & & & \downarrow{\scriptstyle \rho_{ij}} \\ Q(s_{ji}) & & \xrightarrow[\theta_{ij}]{\sim} & & Q(r_{ij}) \end{array}$$

commutes, meaning $\theta_{ij}\lambda_{ij}a_{ij} = \rho_{ij}b_{ij}$. The cited theorem proves existence and invertibility by effective quotients and difunctionality. Its uniqueness also follows directly because $\lambda_{ij}a_{ij}$ is an epimorphism. The displayed computation of the two kernel relations identifies its conclusion with (8). □

This is the categorical butterfly. It is the same correspondence read from its two ends, with each end quotiented by precisely the ambiguity introduced by that correspondence. No picture of four wings is needed to remember the construction: form P, compute $P^{\mathrm{op}}P$ and PP^{op}, and take their support quotients.

Theorem 3.7 (Schreier refinement in an exact Mal'cev category)**.** *In a pointed Barr-exact Mal'cev category, any two finite normal series of an object have equivalent refinements in the pullback sense specified above. The refinements may contain repeated subobjects and zero factors.*

Proof. Suppose first that both series have positive numbers m, n of steps. For the first series, concatenate the m rows supplied by Proposition 3.5. Their endpoints are the original successive subobjects, so they fit together into a normal series of X. For the second series, concatenate the n columns. Both refined lists have mn factors. Their indexing orders differ, but the transposition $(i, j) \leftrightarrow (j, i)$ pairs them. Proposition 3.6 gives the isomorphism for each paired cell.

An empty series can occur only for $X = 0$. Every subobject of 0 is isomorphic to 0: a monomorphism $A \to 0$ has its unique map $0 \to A$ as a right inverse, and monicity makes it a left inverse as well. A regular epimorphism out of 0 also has zero codomain, since it is at once monic and a regular epimorphism. Thus every factor of every series of 0 is zero. We may insert zero steps to give both series the same positive number of zero factors, or treat two empty lists as already equivalent. These are refinements by the zero object's zero series. This resolves the degenerate case without assuming a nonempty comparison rectangle. □

The theorem does not produce a single nested chain containing the terms of both original chains. Two different lines in a two-dimensional vector space already prevent that: neither line contains the other. It produces one refinement of each chain and a matching between their quotient factors. This distinction matters even in the most elementary example.

3.4 Simplicity removes the remaining freedom

Definition 3.8 (Quotient-simple object)**.** A nonzero object S is *quotient-simple* if every regular epimorphism $e : S \twoheadrightarrow T$ is either an isomorphism or has zero codomain. A composition series is a finite normal series all of whose factors are quotient-simple.

The condition concerns the arrow e. The weaker phrase "its codomain is isomorphic to S or to 0" would not exclude a proper self-quotient. The definition above does exclude it. In the group and module instances, it agrees with the usual definition of a nontrivial simple object.

Lemma 3.9 (A simple factor cannot be broken into smaller factors)**.** *Every finite normal series of a quotient-simple object S has exactly one nonzero factor, isomorphic to S. Its other factors are zero.*

Proof. The last quotient is a regular epimorphism $e : S \twoheadrightarrow T$. If $T = 0$, its kernel is all of S, so the preceding subobject is S. Remove this last, zero-factor step and consider the shorter series. Because

the series is finite and begins at $0 \neq S$, this removal process must eventually reach a last step whose preceding subobject is proper in S.

For that step the quotient cannot have zero codomain: a map to 0 has kernel S. Simplicity therefore makes the quotient map an isomorphism. Its kernel is 0, so every preceding subobject of the chain is 0. Every regular quotient of 0 is 0, as checked in the proof of Theorem 3.7. We have one quotient isomorphic to S, zero quotients before it, and the zero quotients stripped from its end. □

Theorem 3.10 (Jordan–Hölder)**.** *Let $\mathcal{C}$ be a pointed Barr-exact Mal'cev category. If an object X has two finite composition series, then their factor lists agree, including multiplicity, up to permutation and isomorphism. In particular, the two series have the same number of factors.*

Proof. Apply Theorem 3.7 to obtain equivalent refinements. Each refinement is obtained from normal series of the original factors, pulled back along the original quotient maps. By Lemma 3.9, each such factor series contributes exactly its original simple factor and some zero factors. Proposition 3.3 identifies these with the factors of the pulled-back refinement.

Consequently, deleting zero factors from the first refined list recovers the first original list, and deleting them from the second recovers the second original list. An isomorphism cannot identify a nonzero object with 0, so the matching between refined lists restricts to a matching between their nonzero lists. This is the desired permutation and set of factor isomorphisms. The equality of lengths is a consequence of that bijection, not an input to its construction. □

No existence assertion is hidden in this theorem. We assumed two finite composition series and compared them. A category may contain objects with no finite composition series. Even when maximal or minimal subobjects can be selected, a termination argument is still required before one can assert existence. The finite sequences used here are already supplied, so neither an ascending nor a descending chain condition enters the uniqueness proof.

3.5 Factor data becomes an invariant only after the comparison

Now, and only now, a numerical or algebraic record of the factors becomes legitimate. Call an object *composition-finite* when it has a finite composition series, allowing 0 with its empty series. Write $\mu(X)$ for the finite multiset of isomorphism classes of its simple factors. Theorem 3.10 proves that $\mu(X)$ is independent of the chosen series. It does not reconstruct X: distinct extension data can give the same multiset.

Proposition 3.11 (Splicing an extension)**.** *Suppose $e : E \twoheadrightarrow Q$ is a regular epimorphism with kernel $K \hookrightarrow E$. If K and Q are composition-finite, then so is E, and*

$$\mu(E) = \mu(K) + \mu(Q). \tag{9}$$

Here $+$ means disjoint union of finite multisets, including multiplicity.

Proof. Choose composition series for K and Q. Pull the series of Q back along e. The construction (2) gives a normal series from K to E whose factors are exactly those of Q. Prepend the chosen composition series from 0 to K. Every factor in the resulting series is simple, so it is a composition series of E. Its factor list is the concatenation of the two chosen lists. Jordan–Hölder makes that calculation independent of all the choices and proves (9). □

This proposition establishes closure under these extensions. It does not assert, without an additional argument, that every subobject or every regular quotient of a composition-finite object is composition-finite in our weakest categorical setting. The direction used in the proof is precisely the direction stated.

For the next statement assume that the isomorphism classes of simple objects under discussion form a set Σ, for example by working in an essentially small category. The free commutative monoid $\mathbb{N}^{(\Sigma)}$ consists of finitely supported functions $\Sigma \to \mathbb{N}$, with pointwise addition. Thus $\mu(X)$ is an element of this monoid.

Corollary 3.12 (The universal additive record)**.** *Let M be a commutative monoid. An arbitrary assignment $w : \Sigma \to M$ extends uniquely to an isomorphism invariant v of composition-finite objects satisfying*

$$v(0) = 0, \qquad v(E) = v(K) + v(Q)$$

for every kernel–regular-quotient sequence $K \hookrightarrow E \twoheadrightarrow Q$ whose terms are composition-finite. The formula is

$$v(X) = \sum_{[S]\in\Sigma} \mu(X)_{[S]}\, w([S]).$$

Proof. Only finitely many summands are nonzero. Jordan–Hölder makes the formula well defined, and Proposition 3.11 gives the required additivity. For uniqueness, every intermediate term of a composition series is composition-finite, because its initial segment is itself a composition series. Repeatedly applying additivity to these successive quotient maps forces $v(X)$ to equal the displayed sum. □

Taking every $w([S]) = 1$ defines length. Taking values in a group gives the corresponding universal group of additive invariants. These are consequences of factor comparison; using them to justify that comparison would reverse the dependency and create circular reasoning. The order here is deliberate: quotient construction, refinement, factor matching, then invariant.

3.6 What is natural in the construction

Proposition 3.13 (Transport by a regular functor)**.** *Let $F : \mathcal{C} \to \mathcal{D}$ be a functor between pointed Barr-exact Mal'cev categories that preserves finite limits, regular epimorphisms, and a zero object. Then F transports the relation rectangle, its refinements, and its cell isomorphisms to the corresponding constructions in $\mathcal{D}$, up to their canonical structural isomorphisms.*

Proof. Finite-limit preservation gives preservation of monomorphisms, products, kernel pairs, kernels, and pullbacks. If a map is factored as a regular epimorphism followed by a monomorphism, its image under F is again such a factorization. Uniqueness of image factorizations therefore shows that F preserves images. Relational composition is an image of a pullback projection, so $F(SR)$ identifies with $F(S)F(R)$. Converse relations are preserved by preservation of the product symmetry.

If q is the effective quotient of an equivalence relation E, then $F(q)$ is a regular epimorphism and its kernel pair is $F(E)$. Hence it is the effective quotient of $F(E)$. These observations preserve every operation in (7), the support and zero-class equations, and the pullback squares defining refinements. Finally the image of the commuting diagram in Proposition 3.6 satisfies the same characterization as the cell comparison in $\mathcal{D}$. Its uniqueness identifies the two isomorphisms. □

A regular functor need not preserve simple objects. For instance, forgetting a module's scalar action can reveal additive subgroups that were not submodules. The proposition transports the comparison construction; a statement about composition series additionally requires that the relevant simple factors remain nonzero and quotient-simple. This distinction prevents a change of framework from silently forgetting part of the theorem's information.

3.7 Dependency and scope check

The ordinary categorical proof is now complete. The foundation needed finite limits and regular images for relation composition, exactness for the effective quotients, and the Mal'cev condition for difunctionality. Pointedness supplied zero classes and kernel boundaries. Saturated descent made each row a refinement of the *chosen* quotient map. The support-quotient comparison matched each cell. Simplicity then removed the zero cells. No invariant length, dimension, Grothendieck group, or pre-existing Jordan–Hölder theorem entered the chain of dependencies.

The comparison pattern is the classical categorical refinement argument associated with Lambek; Knop's Appendix B2 gives its exact Mal'cev formulation and emphasizes the saturation requirement. The proof here has supplied the descent, boundary, and quotient-map checks explicitly; it is not a claim of a new theorem. The detailed primary reference is Friedrich Knop, *Tensor envelopes of regular categories*, Appendix B2.

What has become general is a construction with stated hypotheses, not the assertion that every category has composition factors. The later homotopical chapters will ask which quotient and refinement data survive when equalities become coherent identifications. The present theorem is available there whenever its ordinary exact Mal'cev hypotheses are actually verified; it does not, by terminology alone, assert a Jordan–Hölder theorem for arbitrary higher objects.

Part II
Reuse the comparison and calculate its pages

4 Groups, modules, and dimension as instances

The comparison theorem has been proved before choosing a multiplication table, a scalar field, or a basis. To use it, we must now check its hypotheses. This is an essential part of the argument: a category name does not by itself verify an image factorization or construct a quotient.

4.1 One verification for algebraic structures

An algebraic theory specifies operations of finite arity and equations between expressions built from them. A model is a set with those operations satisfying the equations; a morphism preserves every operation. A subalgebra is a subset closed under them. An internal relation on a model A is a subalgebra of $A \times A$. An equivalence relation of this kind is a *congruence*: equivalent inputs give equivalent outputs for every operation.

Proposition 4.1. *The category of models of an algebraic theory is Barr-exact. If its one-element model is also initial, it is pointed. If the theory has a ternary term p satisfying*

$$p(x, y, y) = x, \qquad p(x, x, z) = z,$$

then it is a Mal'cev category. Under these last two conditions the general comparison theorem applies to its finite composition series.

Proof. Products carry the operations coordinatewise. An equalizer is the subset where two homomorphisms agree, closed under every operation. These give finite limits, with the one-element algebra terminal. The image of a homomorphism is a subalgebra, so every map factors as a surjection followed by an inclusion. A pullback of a surjection is surjective: an element of the other object has a lift of its image, giving a pair in the pullback.

Let $q : A \to B$ be surjective. A homomorphism $h : A \to C$ constant on its fibers defines $\bar{h}(b) = h(a)$ for any a with $q(a) = b$. This is independent of a, and choosing lifts of the finitely many inputs of an operation proves that $\bar{h}$ preserves that operation. Surjectivity proves uniqueness. Thus q is the coequalizer of its kernel pair.

Conversely a regular epimorphism is surjective. Factor a coequalizer $q : A \to B$ through its set-theoretic image as $q = me$. If q coequalizes u, v, injectivity of m shows $eu = ev$. Its universal property supplies $d : B \to \mathrm{im}(q)$ with $dq = e$. Then $mdq = q$, so $md = \mathrm{id}_B$ because coequalizers are epimorphisms. Also $dme = e$, and the surjectivity of e gives $dm = \mathrm{id}_{\mathrm{im}(q)}$. Hence m is an isomorphism. We have established regularity with exactly the stated quotient maps. Monomorphisms are precisely the injective homomorphisms: injectivity gives left cancellation, while factoring a monomorphism into its image makes the surjective part both monic and regular epic, hence an isomorphism by Lemma 1.2. Thus the internal subobjects and relations really are the subalgebras used in the next step.

For a congruence E on A, form the set A/E and define every operation on classes by applying it to representatives. Closure of E under the operation makes this well-defined; the equations descend

because they hold on representatives. The quotient map has kernel pair exactly E and has the same coequalizer property just proved. Thus every internal equivalence relation is effective, proving Barr-exactness.

Finally let $R \subseteq A \times A$ be reflexive and an internal relation. It is closed under the term p. If xRy, apply p to $(y,y),(x,y),(x,x)$ to obtain (y,x), proving symmetry. If xRy and yRz, apply p to $(x,y),(y,y),(y,z)$ to obtain (x,z), proving transitivity. Thus every reflexive internal relation is an equivalence relation. The pointed assertion is precisely the extra initial-object hypothesis. □

The displayed terms are operations in a presentation. For a nonabelian group, $G^3 \to G$, $(x,y,z) \mapsto xy^{-1}z$, need not be a group homomorphism. What we used is that a subalgebra is closed under terms, including that term. This distinction lets algebraic expressions justify a categorical property without confusing them with arrows in the category.

4.2 Reading the axioms in familiar examples

For groups use multiplication, inverse and identity, with their axioms. The one-element group is initial and terminal, and $p(x,y,z) = xy^{-1}z$ has the required identities. For left modules over a fixed unital ring R, use addition, negative, zero, and an operation $x \mapsto rx$ for every $r \in R$, with the module equations. The zero module is initial and terminal, and $p(x,y,z) = x - y + z$ works. Neither argument uses a dimension, a composition series, or Jordan–Hölder.

We check that the categorical quotient data are the familiar ones. In a group G, the kernel N of a homomorphism is normal, since $q(gng^{-1}) = q(g)q(n)q(g)^{-1} = 1$. Conversely, for a normal subgroup N, the relation $x \sim y$ if $x^{-1}y \in N$ is an equivalence relation compatible with multiplication and inverses. For example, if $x' = xn$ and $y' = ym$, then $x'y' = xy(y^{-1}ny)m$ differs from xy by an element of N. The congruence quotient just constructed is the usual group G/N.

More generally, if E is a group congruence, its identity class N is a normal subgroup: closure gives products and inverses, and conjugating a pair $(n,1) \in E$ by $(g,g) \in E$ gives $(gng^{-1},1) \in E$. Translating (x,y) by (x^{-1},x^{-1}) shows xEy exactly when $x^{-1}y \in N$. Thus normal subgroups and these quotient relations agree.

For modules, a congruence has zero class a submodule N, and translating pairs gives $x \sim y$ exactly when $y - x \in N$. The quotient operations are $(x+N)+(y+N) = x+y+N$ and $r(x+N) = rx+N$. Consequently all isomorphisms supplied by the categorical theorem are R-linear. A quotient-simple group is exactly a nontrivial group with no proper nontrivial normal subgroup; a quotient-simple module is exactly a nonzero module with no proper nonzero submodule.

We therefore obtain both classical theorems from Theorem 3.10: two composition series of a group have the same group factors, and two of an R-module have the same module factors. This is one specialization step for each category, not a second proof of the comparison theorem.

The same verification covers groups with specified endomorphism operators: include those operators and their preservation equations in the theory. Subobjects and quotients retain them. It also covers Lie algebras over a field, using the additive term $x - y + z$, and nonunital associative algebras over a commutative ring. Their kernels are respectively ideals and two-sided ideals. For Lie algebras our categorical use of simple includes a one-dimensional abelian algebra. Unital rings with identity-preserving maps do not have a zero object and are not a direct instance of this pointed statement.

4.3 Reading one cell as the familiar group butterfly

Take $A_1 \trianglelefteq A_2 \leq G$ and $B_1 \trianglelefteq B_2 \leq G$. Let r, s be the partial equivalence relations on G supplied by $A_2 \twoheadrightarrow A_2/A_1$ and $B_2 \twoheadrightarrow B_2/B_1$. Put

$$C = A_2 \cap B_2, \qquad U = A_1 \cap B_2, \qquad W = A_2 \cap B_1.$$

Applying s and then r to a subset means taking the appropriate congruence classes inside B_2 and then inside A_2. In particular, $sG = B_2$ and $s0 = B_1$, so the endpoint identities give

$$(rsr)G = rsG = A_1C, \qquad (rsr)0 = rs0 = A_1W.$$

Here 0 denotes the identity subgroup. Interchanging the two relations gives support B_1C and identity class B_1U for srs. Since a group congruence is determined by its identity class, the two cell quotients are

$$Q(rsr) = \frac{A_1C}{A_1W}, \qquad Q(srs) = \frac{B_1C}{B_1U}.$$

We can see their common quotient and every normality assertion directly. The groups U, W are normal in C: conjugation by an element of C preserves A_1 or B_1, respectively, and remains in the other ambient subgroup. Consequently $UW = WU$ is a normal subgroup of C. The product A_1C is a subgroup because $A_1 \trianglelefteq A_2$ and $C \leq A_2$. In A_1C/A_1, the image of C is the whole group and the image of $W \trianglelefteq C$ is normal. Its inverse image is A_1W, proving $A_1W \trianglelefteq A_1C$. The same argument proves $B_1U \trianglelefteq B_1C$.

Now use the two homomorphisms

$$C \longrightarrow A_1C/A_1W, \quad c \longmapsto c(A_1W), \qquad C \longrightarrow B_1C/B_1U, \quad c \longmapsto c(B_1U).$$

Both are onto, since the extra factor A_1 or B_1 is killed. The first kernel is $C \cap A_1W = UW$: if $c = aw$ lies in C, then $a = cw^{-1} \in A_1 \cap C = U$; the converse inclusion is immediate. The second kernel is $C \cap B_1U = WU = UW$, by the same calculation. The effective quotient property therefore gives

$$\begin{array}{ccccc} & & C & & \\ & \swarrow & \downarrow & \searrow & \\ A_1C/A_1W & \xleftarrow{\sim} & C/UW & \xrightarrow{\sim} & B_1C/B_1U. \end{array}$$

These isomorphisms give the usual butterfly, with its maps specified. They are also the previously constructed cell comparison: for every element of the correspondence rs, its two ends can be represented as bc and ac, with $b \in B_1$, $a \in A_1$ and $c \in C$. Both quotient maps then use the same class of c. The characterization in Proposition 3.6 identifies the maps. Thus the subgroup products are the concrete supports and identity classes of one categorical construction.

4.4 A comparison that changes the order

In $G = C_6 = \mathbb{Z}/6\mathbb{Z}$ put $A_1 = \{0, 3\}$ and $B_1 = \{0, 2, 4\}$. Reduction modulo 3 has kernel A_1 and reduction modulo 2 has kernel B_1. Hence

$$0 < A_1 < G, \qquad 0 < B_1 < G$$

have factors (C_2, C_3) and (C_3, C_2). Every nonzero element of C_2 or C_3 generates the group, so the factors are simple. The relation comparison becomes the following concrete rectangle:

	$B_1/0$	G/B_1
$A_1/0$	0	C_2
G/A_1	C_3	0

Indeed $A_1 \cap B_1 = 0$ and $A_1 + B_1 = G$, since $3 + 4 = 1$ modulo 6. The two nonzero comparisons are induced by $a \mapsto a + B_1$ and $b \mapsto b + A_1$. Each is onto by the sum equality and injective by the intersection equality. These are actual isomorphisms, not just equalities of cardinalities. A refined chain need not contain both A_1 and B_1; the matching occurs between its factors.

Normality is always in the next term. In A_4, for example, $1 < V_4 < A_4$ has both terms normal in the ambient group, but V_4 is not simple. Insert one of its order-two subgroups to obtain a composition series. That subgroup is normal in the abelian group V_4 and is not normal in A_4: conjugation by a three-cycle permutes the three double transpositions. The lattice of subgroups normal in the ambient group describes a different filtration problem.

4.5 Finite basis size, with no dimension in the prerequisites

Let k be a field. A list $v_1, \dots, v_n$ is a basis if it spans the vector space V and every relation $\sum_i a_i v_i = 0$ has all $a_i = 0$. The module construction above provides subspaces, quotient spaces and their linear maps before any notion of dimension is introduced.

Theorem 4.2 (Invariance of finite basis size)**.** *If V has finite bases $v_1, \dots, v_n$ and $w_1, \dots, w_m$, then $n = m$.*

Proof. Put $V_i = \operatorname{span}(v_1, \dots, v_i)$, with $V_0 = 0$. Each inclusion is strict: otherwise v_i would be a linear combination of the preceding vectors, contradicting independence. The map

$$k \longrightarrow V_i/V_{i-1}, \qquad a \longmapsto av_i + V_{i-1}$$

is onto by the definition of span. If its image is zero, a relation with coefficient a on v_i shows $a = 0$. It is therefore an isomorphism of k-modules. The module k is simple: a submodule containing $a \neq 0$ contains $a^{-1}a = 1$ and every scalar multiple of 1. The basis thus supplies a composition series with n factors isomorphic to k. The other basis supplies one with m factors. Apply the general Jordan–Hölder theorem in k-modules. If $V = 0$, independence forces both lists empty, giving the same conclusion. □

The direct exchange proof in Section 1 already established this equality by successive invertible replacements. The proof here obtains it from comparisons between quotient layers. The extra work has a different payoff: it explains why the comparison survives in algebraic settings without scalar pivots. Neither argument uses a pre-existing dimension function. We may define $\dim_k V$ using their common conclusion, then develop rank–nullity and other finite-dimensional operations without circularity.

4.6 What a factor list leaves unresolved

Both C_4 and $C_2 \times C_2$ have factors C_2, C_2. Their exact sequences can be written

$$0 \longrightarrow C_2 \xrightarrow{a \mapsto 2a} C_4 \longrightarrow C_2 \longrightarrow 0, \qquad 0 \longrightarrow C_2 \longrightarrow C_2 \times C_2 \longrightarrow C_2 \longrightarrow 0,$$

where the last two maps in the second sequence are first-factor inclusion and second-factor projection. The first quotient is reduction modulo 2. The displayed kernels and surjections verify exactness directly. Yet C_4 has an element of order four, whereas every element of $C_2 \times C_2$ is killed by multiplication by two. They are not isomorphic. The general theorem compares the simple pieces. It does not split the extension joining them. The following parts develop measurements and constructions that retain more of this attachment information.

5 Keeping the operations when we change the object

The common comparison has produced intrinsic simple factors. Its group example also showed a limit: the factors of C_4 and $C_2 \times C_2$ agree, although the groups do not. We now develop that question in a setting where the attachments can be calculated explicitly. The objects will be modules, and our main examples will be vector spaces equipped with an operator. We retain the operator in every map, just as the dimension argument retained scalar multiplication in every quotient.

5.1 Modules and their maps

A ring R has an abelian group under addition, an associative multiplication with identity 1, and the two distributive laws. Multiplication need not be commutative. A left R-module is an abelian group M with operations rx for $r \in R$, $x \in M$, satisfying

$$r(x+y) = rx + ry, \quad (r+s)x = rx + sx, \quad r(sx) = (rs)x, \quad 1x = x.$$

The same cancellation arguments as for vector spaces give $r0 = 0$, $0x = 0$ and $(-1)x = -x$. A submodule contains zero and is closed under addition and all these scalar operations. A module map preserves addition and satisfies $f(rx) = rf(x)$. In this part all unspecified maps between modules are module maps. We work over a fixed ring whenever forming extensions or comparing modules.

Intersections of submodules are submodules. Their sum $U + W$ consists of the elements $u + w$; distributivity proves that it is a submodule, and it is the smallest one containing both. The direct sum of two modules is the set of pairs with coordinatewise addition and scalar action. Each module identity holds in both coordinates. Its inclusions and projections are module maps. A map out of $A \oplus B$ is determined by its restrictions to $(a, 0)$ and $(0, b)$, since $(a, b) = (a, 0) + (0, b)$.

For a submodule $N \subseteq M$, the additive quotient M/N is already defined by the group construction. Set $r(x + N) = rx + N$. Replacing x by $x + n$ changes rx by $rn \in N$, so this action is well-defined; the module identities descend on representatives. The projection is a module map. If a module map $f : M \to X$ kills N, its unique additive factorization $\bar{f} : M/N \to X$ is a module map, because

$$\bar{f}(r(x+N)) = f(rx) = rf(x) = r\bar{f}(x+N).$$

Thus the quotient universal property retains every operation required here.

Kernels, images and preimages of submodules are submodules: apply a map to a sum or a scalar multiple, and use the corresponding closure in the target. An image of a submodule is the image of the map restricted to it. Consequently the first isomorphism theorem supplies a module isomorphism

$$M/\ker f \xrightarrow{\sim} \operatorname{im} f.$$

The inverse of a bijective module map is a module map: writing $y = f(x)$ gives $f^{-1}(ry) = rx$, and the analogous calculation handles sums. All the short exact sequences below mean injectivity, surjectivity, and image equal to kernel for these maps.

5.2 Preimages, quotients, and the local comparison

For a surjection $q : M \to Q$ with kernel N, inverse image and direct image give an inclusion-preserving correspondence between submodules of Q and submodules of M containing N. Indeed, for a submodule H of M,

$$q^{-1}(q(H)) = N + H : \quad q(x) = q(h) \iff x - h \in N.$$

Every quotient map and induced isomorphism in this correspondence is linear over R. For $N \subseteq L \subseteq M$ this also proves $(M/N)/(L/N) \cong M/L$ by $(x + N) + (L/N) \mapsto x + L$.

The kernel calculation at the heart of the butterfly remains valid: if $A_1 \subseteq A_2$ and $B_1 \subseteq B_2$ are submodules of M, write $C = A_2 \cap B_2$, $U = A_1 \cap B_2$, $W = A_2 \cap B_1$. The map

$$C \longrightarrow (A_1 + C)/(A_1 + W), \qquad c \longmapsto c + (A_1 + W)$$

is onto, and its kernel is $U + W$. To check the nontrivial inclusion, if $c = a + w$ with $c, w \in C$ and $a \in A_1$, then $a = c - w \in U$. The reverse inclusion follows by substitution. Exchanging A and B gives the same kernel and the other wing. This is a concrete calculation of the common support quotient, retaining the R-action. It is available independently of any theorem about composition series.

5.3 Finite-dimensional tools, after dimension has been proved

The later operator examples use subspace bases and lifts. We establish their existence here, after Theorem 4.2; none is being fed back into the earlier proof of dimension.

Lemma 5.1. *If a vector space V has a finite basis of size n, every subspace has a finite basis. Every independent list in V extends to a basis of V. For a subspace W, a basis of W followed by lifts of a basis of V/W is a basis of V.*

Proof. Start with any finite independent list. Pass through a fixed finite spanning list of V, appending a vector exactly when it is outside the span of the vectors already present. Each append preserves independence: a relation with nonzero coefficient on the new vector would put it in the old span. At the end the list spans V, so it is a basis. The already proved invariance of basis size shows that the original independent list had at most n members. An arbitrary independent set cannot contain $n + 1$ distinct vectors, so it too is finite and the same conclusion applies.

To construct a basis of W, start with the empty list and, as long as it does not span W, choose an element outside its span. This cannot produce $n+1$ independent vectors. Thus it stops after finitely many steps with a basis of W. This requires only finitely many choices, bounded by n.

The images in V/W of a finite basis of V span the quotient. Applying the append procedure starting with the empty list to those images gives a quotient basis. Choose one lift of each of its finitely many members. The lifts together with a basis of W span V: subtract the lifted expression of a vector's quotient class, leaving an element of W. In a relation among this list, projection to V/W first forces all lift coefficients to vanish; independence in W then forces the remaining coefficients to vanish. □

In particular, $\dim V = \dim W + \dim(V/W)$. Applying the first isomorphism theorem to a linear map gives $\dim V = \dim \ker f + \dim \operatorname{im} f$. These are consequences of the dimension theorem already proved, and will be used only from here on. The same basis-lifting proof gives a vector-space section of any surjective linear map onto a finite-dimensional space. Such a section need not preserve a further module or operator structure.

5.4 An operator is a polynomial action

For a field k, $k[t]$ consists of finite formal sums $\sum_{i\geq 0} a_i t^i$. Addition is coefficientwise and multiplication is convolution: the coefficient of t^r in a product is $\sum_{i+j=r} a_i b_j$. All sums involved are finite. Associativity holds because either parenthesization of a triple product has coefficient $\sum_{i+j+\ell=r} a_i b_j c_\ell$; distributivity is the corresponding distributivity of finite sums. The constant polynomial 1 is the identity.

Given a k-linear endomorphism T of V, define

$$p(t)v = p(T)v = \sum_i a_i T^i v, \qquad T^0 = \mathrm{id}.$$

Induction gives $T^i T^j = T^{i+j}$, so convolution proves $p(T)q(T) = (pq)(T)$. The other module axioms follow from linearity. Conversely, a $k[t]$-module has its underlying k-vector space and the linear operator $T(v) = tv$; the module axioms recover the displayed polynomial action. These constructions are inverse.

A linear map $f : (V,T) \to (V',T')$ is a module map exactly when $fT = T'f$. Necessity uses the polynomial t; sufficiency follows by induction on powers and addition. Submodules are exactly the T-stable subspaces. Thus the square

$$\begin{array}{ccc} V & \xrightarrow{T} & V \\ {\scriptstyle f}\downarrow & & \downarrow{\scriptstyle f} \\ V' & \xrightarrow[T']{} & V' \end{array}$$

records the extra condition that distinguishes an operator-preserving map from an arbitrary linear map.

For $\lambda \in k$, let S_λ be k with t acting by λ. It is simple as a module because its only k-subspaces are zero and itself. It is also $k[t]/(t-\lambda)$: evaluation $p \mapsto p(\lambda)$ is onto, and the identity $t^i - \lambda^i = (t-\lambda)\sum_{j=0}^{i-1} t^{i-1-j}\lambda^j$ shows that its kernel is $(t-\lambda)$.

We also use $P_r = k[t]/(t^r)$ for $r \geq 1$. Its vectors have unique representatives of degree less than r, by separating the terms of degree at least r. Thus $1, t, \ldots, t^{r-1}$ is a vector-space basis, and multiplication by t shifts along a nilpotent chain. No polynomial factorization theorem is required for these modules.

5.5 Simple layers and finite length

A module is simple if it is nonzero and has only zero and itself as submodules. A composition series is a finite strictly increasing submodule filtration from zero to the whole module with simple successive quotients. A module has finite length if it admits such a series; at this stage that phrase is an existence condition. The zero module has the series with one term and no factors.

Lemma 5.2. *A nonzero module map between simple modules is an isomorphism.*

Proof. Its image is nonzero, hence is the whole target. Its kernel cannot be the whole source, hence is zero. The kernel test gives injectivity, and the inverse of a bijective module map preserves the module operations. □

Here a k-algebra is a ring R with a unital map from k into its center. Restricting scalars makes each R-module a k-vector space, and every R-submodule is a k-subspace. Every finite-dimensional module over a k-algebra has a composition series. If it is nonzero, choose a proper submodule of largest possible vector-space dimension; zero is one candidate. The possible dimensions form a finite set by Lemma 5.1. Its quotient is simple: an intermediate nonzero proper submodule would lift to a larger proper submodule. Repeat inside the chosen proper submodule. Dimensions strictly decrease, so the process ends. This proves existence without using uniqueness of composition factors.

We will also need closure of finite-length modules under submodules and quotients over a general ring. Let $0 = M_0 \subsetneq \cdots \subsetneq M_n = M$ be a composition series and $N \subseteq M$ a submodule. Set $N_i = N \cap M_i$ and $Q_i = (M_i + N)/N$. For each i, inclusion and projection give a short exact sequence

$$0 \longrightarrow N_i/N_{i-1} \longrightarrow M_i/M_{i-1} \longrightarrow Q_i/Q_{i-1} \longrightarrow 0. \tag{10}$$

The first map sends $n + N_{i-1}$ to $n + M_{i-1}$ and is injective. The second sends $m + M_{i-1}$ to $(m + N) + Q_{i-1}$ and is onto. Its kernel consists of $m = a + n$ with $a \in M_{i-1}$ and $n \in N$; then $n = m - a \in N_i$, proving equality with the first image. The simple middle module forces precisely one outer module to be zero and the other to be isomorphic to it. Deleting repetitions from the N_i and Q_i therefore gives composition series for N and M/N.

Conversely, if A and B have composition series and $0 \to A \to E \xrightarrow{q} B \to 0$ is short exact, identify A with its image in E. Append to its series the preimages under q of a composition series of B. The quotient correspondence proves strictness, and the maps

$$q^{-1}(B_j)/q^{-1}(B_{j-1}) \longrightarrow B_j/B_{j-1}, \qquad e + q^{-1}(B_{j-1}) \longmapsto q(e) + B_{j-1}$$

are isomorphisms by their kernels and surjectivity. This constructs a composition series for E. Its factors are exactly the two given lists together. No uniqueness theorem is used in this construction.

We now have the objects and operations needed for an independent spectral comparison of two module filtrations. Length and multiplicities are already intrinsic by Theorem 3.10. The following page calculations will compare the factors directly, without using that theorem as a premise.

6 Comparison pages with the operations retained

The relation rectangle gives one proof for every pointed Barr-exact Mal'cev category. There is another useful way to watch the comparison proceed: let simple factors persist through a sequence of pages until maps pair them. We can carry this calculation out simultaneously for groups and modules, without pretending that arbitrary group maps behave like linear maps.

For this chapter a group with operators means a group equipped with a fixed family of endomorphisms. Subgroups and normal subgroups under discussion must be invariant under all those endomorphisms, and maps commute with them. Modules are included by their additive groups and the operators $x \mapsto rx$ for every scalar r. Quotient simplicity is simplicity with this structure retained. Ordinary groups are the case of an empty family of extra operators. This is a concrete common instance of the preceding framework; the page construction below is asserted for this instance, with composition series as input.

6.1 The normality that makes the pages possible

Write two composition series as

$$1 = A_0 \trianglelefteq A_1 \trianglelefteq \cdots \trianglelefteq A_m = G, \qquad 1 = B_0 \trianglelefteq B_1 \trianglelefteq \cdots \trianglelefteq B_n = G.$$

We use multiplicative notation, including for the identity subgroup 1. Extend both filtrations constantly below zero and beyond their last indices. Put $S_i = A_i/A_{i-1}$ and $T_j = B_j/B_{j-1}$ for the real indices $1 \le i \le m$, $1 \le j \le n$.

Lemma 6.1. *A subnormal subgroup of a simple group with operators is either trivial or the whole group. For every integer k and every real index j,*

$$B_{j-1}(A_k \cap B_j) \quad \text{is either } B_{j-1} \text{ or } B_j.$$

In particular this product is an invariant normal subgroup of B_j. Similarly the image of $A_i \cap B_k$ in S_i is either trivial or all of S_i.

Proof. Take a finite normal chain from a proper subnormal subgroup H to a simple group S. The term immediately before the first occurrence of S is proper and normal in S, so it is trivial. The earlier terms, including H, must be trivial. This proves the first claim with invariance under the operators included.

Intersecting a normal chain with a fixed subgroup preserves adjacent normality: conjugating an element of $U \cap K$ by an element of $V \cap K$ stays in both U and K when $U \trianglelefteq V$. Taking images under a homomorphism also preserves adjacent normality, since $f(v)f(u)f(v)^{-1} = f(vuv^{-1})$. Intersections and images here remain invariant under every operator. Thus $A_k \cap B_j$ is subnormal in B_j, by intersecting the tail of the A-chain with B_j; the constant extensions at the endpoints cause no change to this assertion. Its image in the simple factor T_j is trivial or everything. The preimage of that image is exactly $B_{j-1}(A_k \cap B_j)$, proving the stated alternatives. The same argument with the roles exchanged proves the last assertion. □

This check is essential. A nonzero homomorphism between simple groups need not be onto: an inclusion $C_2 \hookrightarrow A_5$ is an immediate counterexample. The maps constructed below have subnormal, and in fact normal, images for a specified reason. That is what makes the cancellation argument valid.

6.2 Terms, maps, and the next page

For $r \geq 1$ define two rows

$$U_r(i) = \frac{A_i \cap B_{n+i-r}}{A_{i-1} \cap B_{n+i-r}}, \qquad 1 \leq i \leq m, \tag{11}$$

$$V_r(j) = \frac{B_j}{B_{j-1}(A_{j+r-n-1} \cap B_j)}, \qquad 1 \leq j \leq n. \tag{12}$$

The same formulas give trivial groups outside the real index ranges. Every source denominator is invariant and normal in its numerator by intersection. Every target denominator is invariant and normal by Lemma 6.1. Consequently these are actual quotient groups with all their operators, not just sets of cosets.

For $j = n + i - r$, inclusion followed by quotient defines

$$d_{r,i} : U_r(i) \longrightarrow V_r(j), \qquad [x] \longmapsto [x]. \tag{13}$$

Its source denominator is $A_{i-1} \cap B_j$, and this is contained in the target denominator $B_{j-1}(A_{i-1} \cap B_j)$. Thus the map is well-defined, is a homomorphism, and preserves the operators. Each source has one designated target; each target has one designated incoming source.

Proposition 6.2 (The page transition)**.** *With $j = n + i - r$, the inclusion and quotient maps give canonical isomorphisms*

$$U_{r+1}(i) \cong \ker d_{r,i}, \qquad V_{r+1}(j) \cong V_r(j)/\operatorname{im} d_{r,i}.$$

The image in the second formula is an invariant normal subgroup.

Proof. Put $H = A_i \cap B_j$, $D = A_{i-1} \cap B_j$, and $K = B_{j-1}$. Thus the map in question is $H/D \to B_j/KD$. The subgroup D is normal in H, K is normal in B_j, and both KD and KH are normal in B_j by Lemma 6.1. If j is outside the real target range, the constant extension gives $B_{j-1} = B_j$, so these normality assertions hold automatically. This also covers transitions whose designated source or target lies outside the initial rows.

We first compute its kernel. Since $D \leq H$,

$$H \cap KD = (H \cap K)D.$$

Indeed, if $h = kd \in H$ with $k \in K$ and $d \in D$, then $k = hd^{-1} \in H \cap K$. The reverse inclusion follows by multiplying elements already in H and in KD. Restricting the quotient map to $H \cap K$ therefore gives a surjection onto the kernel with kernel $D \cap K$. Hence

$$\ker d_{r,i} \cong \frac{H \cap K}{D \cap K} = \frac{A_i \cap B_{j-1}}{A_{i-1} \cap B_{j-1}} = U_{r+1}(i).$$

This isomorphism is induced by inclusion; all its operations are retained.

The image is KH/KD, a normal subgroup of B_j/KD. The next target projection has precisely this kernel, so the compatible iterated quotient map gives

$$\frac{V_r(j)}{\operatorname{im} d_{r,i}} \cong \frac{B_j}{KH} = \frac{B_j}{B_{j-1}(A_i \cap B_j)} = V_{r+1}(j).$$

The last equality uses $i = j + r - n$. These are explicit maps between quotients, so their inverses preserve the operators as well. □

The transition can be displayed as two short exact sequences:

$$1 \longrightarrow U_{r+1}(i) \hookrightarrow U_r(i) \twoheadrightarrow \operatorname{im} d_{r,i} \longrightarrow 1,$$

$$1 \longrightarrow \operatorname{im} d_{r,i} \hookrightarrow V_r(j) \twoheadrightarrow V_{r+1}(j) \longrightarrow 1.$$

The first row retains the source part not yet seen in the target. The second removes the target part accounted for by that map. This is the two-row page mechanism $E_{r+1} \cong H(E_r, d_r)$, with the group quotients justified by the normality proof. We need no addition of maps to perform this particular calculation.

Theorem 6.3 (Every layer eventually finds its partner)**.** *The pages above give a bijection between the simple factors S_i and T_j, together with isomorphisms preserving all the operators. The proof uses neither Schreier refinement nor Jordan–Hölder as an input.*

Proof. The source $U_r(i)$ is isomorphic to the image of $A_i \cap B_{n+i-r}$ in S_i, by the kernel of the restricted quotient map. By Lemma 6.1 it is trivial or isomorphic to S_i. The target denominator in (12) is either B_{j-1} or B_j, so $V_r(j)$ is trivial or isomorphic to T_j.

The image of $d_{r,i}$ is normal by Proposition 6.2. A nontrivial map therefore surjects onto its simple target. Its kernel is an invariant normal subgroup of its simple source; it cannot be the whole source, so it is trivial. Consequently every nontrivial differential is an isomorphism. The transition then removes precisely that source and target. When the differential is trivial, the transition retains both terms by its canonical isomorphisms.

On the first page, $n + i - 1 \geq n$ for every real source index, so $U_1(i) = S_i$. For the targets $j - n \leq 0$, giving $V_1(j) = T_j$. If $G \neq 1$, then $m, n \geq 1$. On every page $r \geq m + n$ we have $n + i - r \leq 0$ and $j + r - n - 1 \geq m$ at real indices. Every term is therefore trivial. Each initial factor must have disappeared exactly once, by an isomorphism differential, and each such differential removes one source and one target. Following the transition isomorphisms to that page gives the required pairwise factor isomorphisms. For $G = 1$, both lists are empty and the conclusion is immediate. □

Notice the dependency: simplicity is part of the supplied filtrations; invariance of their lengths is the conclusion. No page count assumes that the two rows start with the same number of terms.

6.3 A nonabelian crossing and the module specialization

Let $G = A_5 \times C_2$ and use the two series with intermediate subgroups $A_5 \times 1$ and $1 \times C_2$. Their first-page factors have opposite orders. The first differential at $(i, j) = (1, 2)$ identifies $A_5 \times 1$ with $G/(1 \times C_2)$. It removes the two A_5 factors. The second page has no nontrivial differential, but the process has not terminated. On the third page, at $(i, j) = (2, 1)$, the actual differential is the identity on $1 \times C_2$, since both $U_3(2)$ and $V_3(1)$ are that group. Its source is identified with the original factor $G/(A_5 \times 1)$ by the restricted quotient map $1 \times C_2 \to G/(A_5 \times 1)$. Thus it pairs the two remaining C_2 factors. The fourth page is trivial. A quiet page need not be the last page.

For completeness, the simplicity of A_5 used only in this illustration can be checked by its conjugacy classes. Its nonidentity elements are the 15 double transpositions, 20 three-cycles, and 24 five-cycles. Their centralizers in A_5 have orders $4, 3, 5$, respectively: commuting permutations permute cycles

of equal length; on a three- or five-cycle they are powers of that cycle, with the odd transposition on the two unused letters excluded in the three-cycle case; the even permutations commuting with a double transposition form its four-element Klein group. Counting conjugates by cosets of the centralizer therefore gives class sizes $15, 20, 12, 12$. A normal subgroup is a union of these classes and the identity, and its order divides 60 because its cosets partition A_5. None of the proper nontrivial sums of 1 and a selection from $15, 20, 12, 12$ divides 60. Hence no proper nontrivial normal subgroup exists. This finite check is not a prerequisite of the page theorem.

For modules the products in (12) are sums, and the first row of (11) is a subquotient of modules. Put $F_pX = A_{p-n}$, $F_pY = B_p$, and take $f = \mathrm{id}_G : X \to Y$ on the underlying module. The page formulas in the next chapter recover exactly $U_r(i)$ at $p = n + i$ and $V_r(j)$ at $q = j$, with the differential landing at $q = p - r$. The next chapter proves that calculus for an *arbitrary filtered module map*, where the pages may have nonsimple terms and nonzero limiting kernel or cokernel. The present calculation explains why its composition-series cancellation can also be carried out for groups with operators.

The scope is precise. We proved the relation refinement theorem for all pointed Barr-exact Mal'cev categories. We proved this explicit page calculation for composition series of groups with operators. We have not asserted that every filtered group complex supports the ordinary abelian spectral-sequence machinery. General nonabelian exact-couple theories require additional constructions; one primary reference is Grandis, *Homotopy spectral sequences*, Sections 2–4.

7 Filtered maps and their successive obstructions

The common Jordan–Hölder theorem is already complete. A reader following only that theorem may pass to Section 8; the purpose here is to extend the page construction to arbitrary filtered module maps. We now develop the module page calculus beyond composition-series input. Its simplest acyclic example will recover the preceding comparison. We can put the two copies of an object into the exact complex

$$0 \longrightarrow M \xrightarrow{\mathrm{id}} M \longrightarrow 0,$$

give its two terms different filtrations, and calculate what happens to their layers. The complex has no homology: the identity has zero kernel and zero cokernel. Its initial layers need not cancel immediately, because they occur at different filtration levels. A spectral sequence records the successive cancellations.

When the layers are simple, every nonzero cancellation map is an isomorphism. Since nothing survives, these maps pair every factor on one side with exactly one factor on the other. This will prove the module Jordan–Hölder theorem without assuming composition length, composition multiplicities, or their invariance.

We shall construct all the pages needed for this argument. The setting is left modules over a fixed ring with identity. Only submodules, quotient modules, and module homomorphisms are used. The argument does not invoke the earlier categorical Jordan–Hölder theorem. The preceding chapter has already supplied the group and operator version of the particular simple-factor calculation. Here additive module structure permits arbitrary filtered maps, nonsimple layers, and nonzero limiting homology.

7.1 A map seen one layer at a time

Let $f : X \to Y$ be a module homomorphism. Give both modules finite increasing filtrations indexed by the integers. Thus

$$F_pX \subseteq F_{p+1}X, \qquad F_pY \subseteq F_{p+1}Y,$$

the filtrations are zero for sufficiently negative p, and they equal the whole modules for sufficiently positive p. Assume that f preserves them:

$$f(F_pX) \subseteq F_pY \quad \text{for every } p.$$

One may regard $X \xrightarrow{f} Y$ as a chain complex in degrees 1 and 0. Its homology is

$$H_1 = \ker f, \qquad H_0 = Y/f(X).$$

There are no other terms or differentials, so the condition that two successive differentials compose to zero is automatic.

At the first level of resolution, keep only the associated layers:

$$X_0(p) = F_pX/F_{p-1}X, \qquad Y_0(p) = F_pY/F_{p-1}Y.$$

Filtration preservation gives a well-defined module map

$$d_{0,p} : X_0(p) \longrightarrow Y_0(p), \qquad [x] \longmapsto [f(x)].$$

If x is changed by an element of $F_{p-1}X$, its image changes by an element of $F_{p-1}Y$. This is the entire well-definedness check.

Taking the kernel and cokernel of these maps retains the source classes whose images have fallen to a lower layer, and the target classes not yet accounted for by images. We then ask whether the retained source classes reach a target one layer lower, two layers lower, and so on. The page number measures how far down the filtration the map is being examined.

7.2 The pages and their maps

For $r \geq 1$ and integers p, q, define

$$X_r(p) = \frac{F_pX \cap f^{-1}(F_{p-r}Y)}{F_{p-1}X \cap f^{-1}(F_{p-r}Y)}, \tag{14}$$

$$Y_r(q) = \frac{F_qY}{F_{q-1}Y + \big(F_qY \cap f(F_{q+r-1}X)\big)}. \tag{15}$$

All the displayed sets are submodules. In the first formula the denominator lies in the numerator because $F_{p-1}X \subseteq F_pX$; in the second both summands of the denominator lie in F_qY.

The numerator of $X_r(p)$ consists of representatives in source layer p whose image has reached target level $p - r$ or lower. Its denominator identifies representatives that were already in a lower source layer and satisfy the same condition. On the target side, $Y_r(q)$ discards lower target layers and images already detected from source levels up to $q + r - 1$.

More precisely, $[x] \mapsto x + F_{p-1}X$ embeds $X_r(p)$ into $F_pX/F_{p-1}X$: the map on its numerator has exactly the displayed denominator as kernel. The target $Y_r(q)$ is a quotient of $F_qY/F_{q-1}Y$. Thus a

surviving source entry is retained as a submodule of its original layer, while a surviving target entry is retained as a quotient of its original layer.

There is a module map

$$d_{r,p} : X_r(p) \longrightarrow Y_r(p-r), \qquad [x] \longmapsto [f(x)]. \tag{16}$$

To check it, put $q = p - r$. A numerator representative has image in F_qY. Changing it by $z \in F_{p-1}X \cap f^{-1}(F_qY)$ changes the image by

$$f(z) \in F_qY \cap f(F_{p-1}X) = F_qY \cap f(F_{q+r-1}X),$$

which is in the target denominator. The formula is therefore independent of representatives and is linear.

For each page, the arrows go from the source row to the target row:

$$\begin{array}{ll} X_r(p) & \text{source degree 1} \\ \big\downarrow{\scriptstyle d_{r,p}} & \\ Y_r(p-r) & \text{target degree 0.} \end{array}$$

There are no arrows out of the target row or into the source row. Hence the square of the differential is zero.

Proposition 7.1 (The next page is homology)**.** *There are canonical isomorphisms*

$$X_{r+1}(p) \cong \ker\big(d_{r,p} : X_r(p) \to Y_r(p-r)\big),$$
$$Y_{r+1}(q) \cong Y_r(q)/\operatorname{im}\big(d_{r,q+r} : X_r(q+r) \to Y_r(q)\big)$$

for every $r \geq 0$.

Proof. For $r = 0$, the kernel of $d_{0,p}$ consists of the classes represented by $x \in F_pX$ with $f(x) \in F_{p-1}Y$. Since $f(F_{p-1}X) \subseteq F_{p-1}Y$, its denominator is $F_{p-1}X$. This is exactly $X_1(p)$. Quotienting $Y_0(q)$ by the image of $d_{0,q}$ gives

$$F_qY/\big(F_{q-1}Y + f(F_qX)\big) = Y_1(q).$$

Now take $r \geq 1$ and put $q = p - r$. The numerator for $X_{r+1}(p)$ is contained in the numerator for $X_r(p)$. Its intersection with the old denominator is

$$\big(F_pX \cap f^{-1}(F_{q-1}Y)\big) \cap \big(F_{p-1}X \cap f^{-1}(F_qY)\big) = F_{p-1}X \cap f^{-1}(F_{q-1}Y).$$

That is precisely the new denominator. Thus inclusion of representatives induces an injection $X_{r+1}(p) \hookrightarrow X_r(p)$. Every class in its image maps to zero under $d_{r,p}$, because its image under f lies in $F_{q-1}Y$.

Conversely, suppose $[x] \in X_r(p)$ maps to zero. Then

$$f(x) \in F_{q-1}Y + \big(F_qY \cap f(F_{p-1}X)\big).$$

Choose $z \in F_{p-1}X$ with $f(z) \in F_qY$ such that $f(x) - f(z) \in F_{q-1}Y$. The element z is in the denominator of $X_r(p)$, so $x - z$ represents the same class there. But $x - z$ is in the numerator of

$X_{r+1}(p)$. This proves that the injection identifies $X_{r+1}(p)$ with the kernel. The choice of z proves that a suitable representative exists; the canonical map itself is induced by inclusion and does not depend on that choice.

On the target side, the denominators increase with r, giving a natural surjection $Y_r(q) \twoheadrightarrow Y_{r+1}(q)$. The additional elements killed are the classes of

$$F_qY \cap f(F_{q+r}X).$$

Every such element is $f(x)$ for an $x \in F_{q+r}X$ whose image lies in F_qY, so it comes from $X_r(q+r)$. Conversely, every image of $d_{r,q+r}$ has this form. The kernel of the surjection is therefore exactly that image, proving the asserted cokernel formula. □

This calculation is a spectral sequence, in the ordinary algebraic sense. To see the conventional indexing, set

$$E_r^{p,q} = \begin{cases} X_r(p), & p+q = 1, \\ Y_r(p), & p+q = 0, \\ 0, & \text{otherwise.} \end{cases}$$

Then d_r has bidegree $(-r, r-1)$, the convention for an increasing filtration on a chain complex, and Proposition 7.1 says exactly

$$E_{r+1} = H(E_r, d_r).$$

The explicit formulas establish this relation; no general spectral sequence existence theorem is being assumed. The support lies on the two total-degree diagonals $p+q = 1$ and $p+q = 0$. The index q may be negative; no first-quadrant assumption is needed. The finite filtration support will supply the convergence bound.

The construction also respects maps. Suppose

$$\begin{array}{ccc} X & \xrightarrow{f} & Y \\ {\scriptstyle a}\downarrow & & \downarrow{\scriptstyle b} \\ X' & \xrightarrow[f']{} & Y' \end{array}$$

commutes and a, b preserve filtrations. The formula $[x] \mapsto [a(x)]$ defines $X_r(p) \to X'_r(p)$: the equality $f'a = bf$ carries both numerator and denominator into their primed versions. Similarly $[y] \mapsto [b(y)]$ defines $Y_r(q) \to Y'_r(q)$, since $b(f(F_jX)) = f'(a(F_jX)) \subseteq f'(F_jX')$. These maps commute with d_r by the same equality $f'a = bf$. The source inclusions and target quotient maps defining the next page commute with a, b as well, so the identifications with homology are natural. Thus every page and differential belongs to the given filtered map, rather than to choices of bases.

7.3 What eventually survives

Proposition 7.2 (Convergence in this two-term setting)**.** *For sufficiently large r, the pages stabilize to*

$$X_r(p) = \frac{F_pX \cap \ker f}{F_{p-1}X \cap \ker f},$$
$$Y_r(q) = \frac{F_qY}{F_{q-1}Y + (F_qY \cap f(X))}.$$

The first is the associated layer of the induced filtration on $\ker f$*; the second is canonically the associated layer of the quotient filtration on* $Y/f(X)$*. In particular, if* f *is an isomorphism, every page entry is eventually zero.*

Proof. Fix p. For sufficiently large r, $F_{p-r}Y = 0$, so the two inverse images in (14) become $\ker f$. Fix q. For sufficiently large r, $F_{q+r-1}X = X$, giving the second formula from (15).

Let $\pi : Y \to Y/f(X)$ be the quotient map. The quotient filtration is $F_q(Y/f(X)) = \pi(F_qY)$. The map

$$F_qY \longrightarrow \pi(F_qY)/\pi(F_{q-1}Y), \qquad y \longmapsto \pi(y) + \pi(F_{q-1}Y)$$

is surjective. An element $y \in F_qY$ is in its kernel exactly when $y = y_0 + f(x)$ for $y_0 \in F_{q-1}Y$. Then $f(x) = y - y_0 \in F_qY$, so the kernel is $F_{q-1}Y + (F_qY \cap f(X))$, as asserted.

Only finitely many filtration columns are nonzero on page 0. Proposition 7.1 shows that a zero entry stays zero. We may therefore take a single sufficiently large r for all relevant columns. If f is an isomorphism, its kernel and cokernel are zero, so all the stabilized entries are zero. □

Notice what the proof used: bounded filtrations, kernels, images, and explicit quotient maps. It did not count dimensions or composition factors. This is the point that makes a noncircular application to Jordan–Hölder possible.

7.4 Two composition series inside an acyclic complex

Suppose that a module M has two composition series

$$0 = A_0 \subsetneq A_1 \subsetneq \cdots \subsetneq A_m = M, \qquad 0 = B_0 \subsetneq B_1 \subsetneq \cdots \subsetneq B_n = M.$$

Thus each A_i/A_{i-1} and each B_j/B_{j-1} is a nonzero simple module. The numbers m and n are not assumed equal.

Extend the notation by $A_i = 0$ for $i \le 0$, $A_i = M$ for $i \ge m$, and similarly for B. Take $X = Y = M$ and $f = \mathrm{id}_M$, with

$$F_pX = A_{p-n}, \qquad F_pY = B_p. \tag{17}$$

The entire B filtration occurs before the nonzero stages of the A filtration. This makes the identity filtration preserving: if $p \le n$, then $F_pX = 0$; if $p > n$, then $F_pY = M$. We have a filtered complex even though the inverse identity need not preserve these two filtrations.

Its nonzero page-0 entries are precisely

$$X_0(n+i) = A_i/A_{i-1} \quad (1 \le i \le m), \qquad Y_0(j) = B_j/B_{j-1} \quad (1 \le j \le n).$$

Their column ranges are disjoint, so $d_0 = 0$ and $E_1 = E_0$. The positions can be summarized as follows:

row	nonzero columns	entry
source	$p = n + i, \quad 1 \le i \le m$	A_i/A_{i-1}
target	$p = j, \quad 1 \le j \le n$	B_j/B_{j-1}

Each d_r travels from a source column to the target column r places to its left.

Theorem 7.3 (Module Jordan–Hölder by spectral cancellation)**.** *The two composition series have equal length,* $m = n$*, and their factors can be paired by module isomorphisms. The nonzero differentials in the spectral sequence of* (17) *supply such a pairing.*

Proof. Recall first why a nonzero map $u : S \to T$ between simple modules is an isomorphism. Its kernel is a submodule of S, and cannot be all of S when $u \neq 0$, so it is zero. Its image is a nonzero submodule of T, so it is T. The first isomorphism theorem now gives the assertion.

Initially every nonzero page entry is simple. Inductively suppose this is true on page r. A differential is either zero, leaving its source and target unchanged on the next page, or is an isomorphism between two simple entries, making both zero on the next page. There is only one possible outgoing differential from each source position and one possible incoming differential to each target position on a fixed page. There are no incoming differentials to source positions and no outgoing differentials from target positions. Proposition 7.1 therefore proves the induction.

Keep each original factor as a label on its position until that position vanishes. A zero differential carries the label forward by the canonical kernel or quotient identification. A nonzero differential pairs its two labels by an isomorphism and removes them. Removed positions stay zero, so no label is used twice. Proposition 7.2 applies to $f = \mathrm{id}_M$ and says that eventually every position is zero. Consequently every original label is paired exactly once, one A factor with one B factor. This is a bijection of their finite index sets, proving $m = n$ and the required isomorphisms.

If $M = 0$, both composition series have no factors and the conclusion holds directly. □

For $M \neq 0$, termination also has an explicit bound: $E_{m+n} = 0$. Indeed, in a source column $p = n+i$, formula (14) then uses $F_{p-(m+n)}Y = B_{i-m} = 0$. In a target column $j \geq 1$, formula (15) uses $F_{j+(m+n)-1}X = A_{j+m-1} = M$, so its denominator is all of B_j. These bounds use the lengths of the two given lists, without assuming that those lengths agree or are intrinsic to M.

This proof obtains an actual pairing. An equality of alternating sums would give less information unless one had already established the independence of simple classes. No such independence, or Grothendieck group calculation, is needed here.

7.5 A calculation with nonzero differentials

Let $M = ke_1 \oplus ke_2$, with

$$A_1 = ke_1, \qquad B_1 = ke_2, \qquad A_2 = B_2 = M.$$

Take $n = 2$ in (17). Thus the source enters at columns $3, 4$, while the target enters at columns $1, 2$:

	1	2	3	4
X_1	0	0	ke_1	M/ke_1
Y_1	ke_2	M/ke_2	0	0.

The only potentially nonzero d_1 is

$$d_{1,3} : ke_1 \longrightarrow M/ke_2, \qquad ae_1 \longmapsto ae_1 + ke_2.$$

It is an isomorphism: every target class has such a representative, and the only vector common to the two coordinate lines is zero. Those two positions disappear. Page 2 therefore has just

$$X_2(4) = M/ke_1, \qquad Y_2(1) = ke_2.$$

There is no target in column 2, so $d_{2,4} = 0$; there is no source in column 3, so no d_2 reaches column 1. Page 3 has the same two surviving factors, but its source formula gives a useful representative:

$$X_3(4) = \frac{M \cap ke_2}{ke_1 \cap ke_2} = ke_2, \qquad Y_3(1) = \frac{ke_2}{ke_2 \cap ke_1} = ke_2.$$

The identification of this $X_3(4)$ with $X_2(4)$ sends ae_2 to $ae_2 + ke_1$. Now

$$d_{3,4} : ke_2 \xrightarrow{\mathrm{id}} ke_2$$

removes the remaining pair, and $E_4 = 0$. In particular, the zero differential d_2 did not mean that the calculation had stabilized: d_3 was still nonzero. Equality of two successive pages does not by itself imply that every later differential vanishes.

The cancellations occur at different distances:

$$ke_1 \xrightarrow[\sim]{d_{1,3}} M/ke_2 \qquad \text{columns } 3 \to 2$$

$$ke_2 \xrightarrow[\sim]{d_{3,4}} ke_2 \qquad \text{columns } 4 \to 1.$$

This last picture records the two nonzero arrows on their respective pages; it is not a single page. The source label ke_2 in the second row is the page-3 representative just calculated.

For comparison, if $A_1 = B_1 = ke_1$, then $d_1 = 0$ and the two nonzero maps occur on page 2:

$$d_{2,3} : ke_1 \xrightarrow{\sim} ke_1, \qquad d_{2,4} : M/ke_1 \xrightarrow{\sim} M/ke_1.$$

The page of cancellation records relative filtration position; it is not an invariant of the unfiltered module. One may also give M the operator $Te_1 = 0$, $Te_2 = e_2$ and regard it as a $k[t]$-module. The same crossed filtration calculation then pairs the simple type with eigenvalue 0 on page 1, and the different simple type with eigenvalue 1 on page 3. The argument matches module types, as well as the number of terms.

7.6 The butterfly is visible inside the differentials

The construction has a precise relation to the earlier refinement rectangle. Take any two finite filtrations A and B of a module M; their factors need not be simple. Use the same shifted identity complex. For $1 \le i \le m$ and $1 \le j \le n$, put

$$r = n + i - j, \qquad p = n + i.$$

Then $r \ge 1$, and the explicit pages become

$$X_r(n+i) = \frac{A_i \cap B_j}{A_{i-1} \cap B_j},$$
$$Y_r(j) = \frac{B_j}{B_{j-1} + (A_{i-1} \cap B_j)}.$$

The differential is induced by the inclusion $A_i \cap B_j \hookrightarrow B_j$. Its image is canonically

$$\operatorname{im} d_{n+i-j,n+i} \cong Q_{ij} := \frac{A_i \cap B_j}{(A_{i-1} \cap B_j) + (A_i \cap B_{j-1})}. \tag{18}$$

Here is the kernel check. If $x \in A_i \cap B_j$ has zero target class, write $x = b + z$ with $b \in B_{j-1}$ and $z \in A_{i-1} \cap B_j$. Since both x and z lie in A_i, so does $b = x - z$; hence $b \in A_i \cap B_{j-1}$. The converse is immediate. Thus the kernel of the map from $A_i \cap B_j$ to its image is exactly the denominator in (18), and the first isomorphism theorem applies.

The same Q_{ij} is what appears on taking associated layers in either order. To state this precisely, filter A_i/A_{i-1} by

$$\overline{B}_j = \frac{A_{i-1} + (A_i \cap B_j)}{A_{i-1}}.$$

The map $A_i \cap B_j \to \overline{B}_j/\overline{B}_{j-1}$ is surjective. Its kernel is $(A_{i-1} \cap B_j) + (A_i \cap B_{j-1})$: if $x = a + y$ with $a \in A_{i-1}$ and $y \in A_i \cap B_{j-1}$, then $a = x - y \in B_j$; the reverse inclusion is immediate. Consequently

$$\operatorname{gr}_j^B \operatorname{gr}_i^A M \cong Q_{ij} \cong \operatorname{gr}_i^A \operatorname{gr}_j^B M.$$

The second isomorphism follows by the same calculation after interchanging A and B. These are isomorphisms of individual indexed subquotients; no simultaneous choice of a basis or a splitting of the filtrations has been made.

Thus the static comparison rectangle and the spectral calculation share the same subquotients. The rectangle displays all Q_{ij} together. The spectral sequence encounters them in order of the distance $r = n + i - j$, as images of actual differentials. The kernel and quotient calculations proved above explain the connection. Calling the rectangle alone a spectral sequence would omit the page transitions and the differentials that we have now supplied.

7.7 What this adds, and what it does not assert

The initial framework is a filtered map. Its successive pages separate what the map already accounts for from what remains. The stable pages retain the layers of its kernel and cokernel. When the map is an identity, these stable pages are zero; when the initial layers are simple, the necessary cancellations become isomorphisms. Jordan–Hölder emerges from those two facts together.

A single filtration on a module concentrated in degree zero has no nonzero differentials. Its spectral sequence leaves its own associated graded object in place and does not by itself compare a second filtration. The two-term identity complex is what creates a comparison with nonzero differentials. Likewise, the bare statement that an alternating sum vanishes would not supply our pairing proof.

With finite filtrations, the stabilized pages describe associated layers of homology. They do not, without further data, specify how those layers are attached. That limitation returns us to the questions about extensions and reconstruction in the following chapters. Spectral sequences organize successive information; they do not make the extension problem disappear.

For the general theory, the Stacks Project constructs filtered-complex spectral sequences and proves finite-filtration convergence in Section 12.24, especially Lemmas 12.24.2 and 12.24.11. Its convention is cohomological with decreasing filtrations, whereas ours is homological with increasing filtrations. The two-filtration interchange is explicitly identified with Zassenhaus's lemma in Deligne's *Théorie de Hodge II*, §1.2.1. These references place the constructions in their classical setting; the proof in this chapter is given directly, and no originality claim is made for this presentation.

Part III
Measure coupling and recover chains

8 What every additive measurement must retain

The comparison of composition series gives more than an integer. For each isomorphism class of simple modules S, let $m_S(M)$ be the number of its occurrences in a composition series of M. The common Jordan–Hölder theorem (Theorem 3.10) makes this well-defined. Only finitely many of these numbers are nonzero for a fixed finite-length module M. The universal monoid statement was proved in Corollary 3.12. We now construct its group-valued recipient explicitly for modules, including the size and closure checks needed to use it.

8.1 Splicing gives additivity before we name a group

The last construction in Section 5 spliced composition series across a short exact sequence. Uniqueness now turns that particular construction into an intrinsic assertion:

$$0 \to A \to E \to B \to 0 \quad \Longrightarrow \quad m_S(E) = m_S(A) + m_S(B) \quad \text{for every simple } S. \tag{19}$$

The existence arguments there also show that the category of finite-length R-modules is closed under submodules, quotients and extensions. Every object subsequently used in such a sequence stays inside that category.

An additive measurement is an assignment $a(M)$ in an abelian group D, constant on isomorphic modules, such that $a(E) = a(A) + a(B)$ for every short exact sequence of finite-length modules. The zero sequence forces $a(0) = 0$. Repeatedly using the short exact sequences of a composition series gives

$$a(M) = \sum_S m_S(M) a(S). \tag{20}$$

Thus a choice of values on simple modules determines every additive measurement. Conversely any such choice defines an additive measurement by this formula and (19). This already answers the universal question in concrete terms.

8.2 Constructing the universal recipient

For completeness we construct the group that receives every additive measurement. Take one formal generator e_M for each isomorphism class of finite-length R-modules. These classes form a set: every simple module is cyclic, since the map $R \to S$, $r \mapsto rx$, has nonzero image for $x \neq 0$; inductively every finite-length module is finitely generated by generators of a submodule and lifts of generators of a quotient. Such a module is a quotient of some R^n, and the submodules of all these R^n form a set.

Let F be the abelian group of finitely supported integer combinations of the symbols e_M, with coefficientwise addition. Let J be the subgroup generated by

$$e_E - e_A - e_B \quad \text{for all } 0 \to A \to E \to B \to 0.$$

Define $K_0(R\text{-fl}) = F/J$, and write $[M]$ for the image of e_M. Here R-fl means finite-length left R-modules. We specify that category because enlarging it can introduce new exact sequences and therefore new relations.

Given an additive measurement a into D, the formula $\sum n_M e_M \mapsto \sum n_M a(M)$ is a homomorphism $F \to D$. It kills J exactly by additivity, so the quotient universal property gives a unique homomorphism $\bar{a} : K_0 \to D$ with $\bar{a}([M]) = a(M)$. Conversely such a homomorphism defines an additive measurement. This is the universal property of the Grothendieck group.

Theorem 8.1. *The homomorphism*

$$\bigoplus_{[S]\ simple} \mathbb{Z}e_S \longrightarrow K_0(R\text{-fl}), \qquad e_S \longmapsto [S]$$

is an isomorphism. Its inverse sends $[M]$ to $\sum_S m_S(M)e_S$.

Proof. Additivity of every m_S shows that the proposed inverse kills each defining relation and is therefore well-defined. A simple module has one factor, itself, so the composite on each e_S is the identity. For the other composite, the defining relations applied to a composition series give $[M] = \sum_S m_S(M)[S]$. Since the classes $[M]$ generate K_0, this composite too is the identity. □

The direct sum in the theorem means finite support, even if there are infinitely many simple isomorphism classes. The theorem proves exactly which information all additive measurements retain. It says nothing about measurements that fail additivity on short exact sequences.

Length is the homomorphism sending every $[S]$ to 1. For finite-dimensional vector spaces there is one simple class $[k]$, so $K_0 \cong \mathbb{Z}$ and $[V] = (\dim_k V)[k]$. For finite-dimensional modules over a k-algebra, dimension instead assigns the weight $\dim_k S$ to a simple class. To use the same proof in that category, observe that submodules and quotients are finite-dimensional by Lemma 5.1, extensions have a basis made from a submodule basis and quotient lifts, and all these modules have composition series by Section 5.5.

8.3 Why the chosen category matters

Let $S = k[t]/(t)$. In the category of finite-length $k[t]$-modules, $[S]$ is a nonzero generator of the free group in the theorem. If we instead impose exact-sequence relations for all finitely generated $k[t]$-modules, the sequence

$$0 \longrightarrow k[t] \xrightarrow{t} k[t] \longrightarrow S \longrightarrow 0$$

forces $[S] = 0$. Multiplication by t is injective by inspection of coefficients, and its image is exactly the polynomials with zero constant term, proving exactness. The module $k[t]$ does not have finite length: the strictly descending submodules (t^n) would each force a strictly smaller nonnegative length if it did, by (19). Thus this new relation was not available in the previous category.

The invariant depends on which objects and exact sequences we permit. Keeping that context visible prevents a numerical formula from silently changing its meaning.

8.4 A short bridge to Euler characteristic

Let $C^\bullet$ be a bounded complex of finite-length modules: it has maps $d^i : C^i \to C^{i+1}$ with $d^{i+1}d^i = 0$, and only finitely many C^i are nonzero. Put

$$Z^i = \ker d^i, \quad B^i = \operatorname{im} d^{i-1}, \quad H^i = Z^i/B^i.$$

The complex condition gives $B^i \subseteq Z^i$. Kernels, images and quotients are finite-length by the results already proved. The two short exact sequences

$$0 \to Z^i \to C^i \xrightarrow{d^i} B^{i+1} \to 0, \qquad 0 \to B^i \to Z^i \to H^i \to 0$$

give $[C^i] = [B^i] + [H^i] + [B^{i+1}]$. Alternating addition cancels every boundary class with its next occurrence. Boundedness makes all the sums finite, and proves

$$\sum_i (-1)^i [C^i] = \sum_i (-1)^i [H^i]. \tag{21}$$

Applying any additive measurement gives its numerical version. In particular, vector-space dimension gives the usual Euler characteristic formula.

This identity follows directly from the exact-sequence relations. We have introduced it only after the earlier comparison. In particular, we did not infer Jordan–Hölder from an equality in a Grothendieck group whose simple generators had already been justified by Jordan–Hölder. The preceding spectral argument compared actual simple modules and actual differentials.

8.5 The information that remains to be recovered

If $0 \to A \to E \to B \to 0$ is exact, its class is always $[A] + [B]$, whether or not the sequence splits. Every additive measurement therefore agrees on E and $A \oplus B$. This is an exact limitation of additivity. The next chapter retains the maps of the extension, and measures how the middle object can depend on them.

9 What attaches one layer to another

Take a linear operator on a two-dimensional vector space with basis u, v:

$$Tu = \mu u, \qquad Tv = \lambda v + au.$$

Can a change of the second vector remove the coupling coefficient a? Try $v' = v + bu$ and compute before invoking any classification:

$$Tv' = \lambda v' + \big(a + b(\mu - \lambda)\big)u.$$

If $\mu \neq \lambda$, choose $b = -a/(\mu - \lambda)$ and the two motions separate. If $\mu = \lambda$, this change cannot alter a at all. The same calculation produces both a coordinate repair and a precise limit on that repair. Section 9.5 will explain which data are being kept fixed, and what changes when that requirement is relaxed.

For finite-length modules, an exact sequence

$$0 \longrightarrow A \xrightarrow{i} E \xrightarrow{q} B \longrightarrow 0$$

always gives $[E] = [A] + [B]$, whether or not E is a direct sum. To learn more, we must retain the way A sits in E and the way E maps onto B. The maps are part of the object we now study.

Throughout this section R is a ring with identity, modules are left R-modules, and every map between modules is R-linear. Direct sums carry componentwise operations. An *extension of B by A* means the whole displayed short exact sequence, with A and B fixed. We will first construct extensions, then calculate a family of them, and finally use them to detect when a map cannot be lifted.

9.1 Keeping the endpoints fixed

Two extensions (i, E, q) and (i', E', q') are *equivalent* if there is a map $h : E \to E'$ making this diagram commute:

$$\begin{array}{ccccccccc} 0 & \longrightarrow & A & \xrightarrow{i} & E & \xrightarrow{q} & B & \longrightarrow & 0 \\ & & \| & & \downarrow{\scriptstyle h} & & \| & & \\ 0 & \longrightarrow & A & \xrightarrow{i'} & E' & \xrightarrow{q'} & B & \longrightarrow & 0. \end{array}$$

Thus $hi = i'$ and $q'h = q$. No separate invertibility assumption on h is needed.

Lemma 9.1. *Every such h is an isomorphism.*

Proof. If $h(e) = 0$, then $q(e) = q'h(e) = 0$, so $e = i(a)$ for some $a \in A$. Now $0 = hi(a) = i'(a)$, and injectivity of i' gives $a = 0$ and $e = 0$. Hence h is injective.

Given $e' \in E'$, choose $e \in E$ with $q(e) = q'(e')$. This is possible because q is surjective. The difference $e' - h(e)$ lies in $\ker q'$, so it is $i'(a)$ for some $a \in A$. Then $h(e + i(a)) = e'$. Thus h is surjective. The inverse of a bijective linear map is linear: apply h to the proposed addition and scalar identities and use its injectivity. □

Identity maps, compositions, and inverses now show that equivalence is indeed an equivalence relation. The word *fixed* matters: an isomorphism between the middle modules that changes an endpoint identification need not give an equivalence of extensions.

The *split extension* is

$$0 \longrightarrow A \xrightarrow{a \mapsto (a,0)} A \oplus B \xrightarrow{(a,b) \mapsto b} B \longrightarrow 0.$$

An extension is called split when it is equivalent to this one.

Proposition 9.2. *An extension is split if and only if there is a linear map $s : B \to E$ with $qs = \mathrm{id}_B$.*

Proof. Such a map is a *section* of q. Given a section, the map

$$A \oplus B \longrightarrow E, \qquad (a, b) \longmapsto i(a) + s(b)$$

commutes with both endpoint maps and is therefore an isomorphism by Lemma 9.1. Conversely, an equivalence with the split extension transports its section $b \mapsto (0, b)$ to a section of q. □

A choice of a lift of each individual element of B is insufficient. The lifts must form a linear map. Failure of that compatibility is precisely what a nonsplit extension records.

9.2 Change the quotient: pull back

Suppose we want to compare an extension with a map $f : P \to B$. We need an object that simultaneously records an element of E and an element of P which have the same image in B. Define

$$f^*E = E \times_B P = \{(e, p) \in E \oplus P : q(e) = f(p)\}.$$

This is a submodule: it is the kernel of the linear map $(e, p) \mapsto q(e) - f(p)$. Put

$$j(a) = (i(a), 0), \qquad \pi(e, p) = p, \qquad u(e, p) = e.$$

These maps form the diagram

$$\begin{array}{ccccccccc} 0 & \longrightarrow & A & \xrightarrow{j} & f^*E & \xrightarrow{\pi} & P & \longrightarrow & 0 \\ & & \| & & \downarrow{\scriptstyle u} & & \downarrow{\scriptstyle f} & & \\ 0 & \longrightarrow & A & \xrightarrow{i} & E & \xrightarrow{q} & B & \longrightarrow & 0. \end{array}$$

The left square commutes because $uj = i$. The right one commutes because membership in f^*E says exactly that $qu = f\pi$.

The top row is exact. The map j is injective because i is. If $\pi(e, p) = 0$, then $p = 0$ and $q(e) = 0$, so $(e, p) = (i(a), 0) = j(a)$ for a unique a. Conversely every $j(a)$ maps to zero. Finally, for any $p \in P$, surjectivity of q supplies e with $q(e) = f(p)$, so $\pi(e, p) = p$. No assumption that f is injective or surjective was used.

The right square is a pullback in the precise universal sense. If maps $r : T \to E$ and $s : T \to P$ satisfy $qr = fs$, the unique map to f^*E with coordinates r and s is

$$T \longrightarrow f^*E, \qquad t \longmapsto (r(t), s(t)).$$

The compatibility puts its image in f^*E; coordinatewise operations make it linear; its two coordinates force uniqueness.

9.3 Change the subobject: push out

Now suppose instead that $g : A \to A'$ is given. We want to replace the embedded A by A' while preserving the quotient B. In $A' \oplus E$ we therefore identify $(g(a), 0)$ with $(0, i(a))$. The required relation submodule is

$$L_g = \{(g(a), -i(a)) : a \in A\}, \qquad g_*E = (A' \oplus E)/L_g.$$

Here L_g is the image of a linear map, so it is a submodule and the quotient is defined. Brackets below denote classes modulo L_g. Define

$$i_g(a') = [a', 0], \qquad u_g(e) = [0, e], \qquad q_g([a', e]) = q(e).$$

The last map is well-defined because $q(i(a)) = 0$. The diagram is

$$\begin{array}{ccccccccc} 0 & \longrightarrow & A & \xrightarrow{i} & E & \xrightarrow{q} & B & \longrightarrow & 0 \\ & & \downarrow{\scriptstyle g} & & \downarrow{\scriptstyle u_g} & & \| & & \\ 0 & \longrightarrow & A' & \xrightarrow{i_g} & g_*E & \xrightarrow{q_g} & B & \longrightarrow & 0. \end{array}$$

The left square commutes because $[0, i(a)] = [g(a), 0]$. The right square commutes by the formula for q_g.

The bottom row is exact. If $[a', 0] = 0$, then $(a', 0) = (g(a), -i(a))$ for some a; injectivity of i gives $a = 0$ and then $a' = 0$. If $q_g([a', e]) = 0$, write $e = i(a)$. The relation gives

$$[a', e] = [a', i(a)] = [a' + g(a), 0],$$

so the kernel is the image of i_g. That image is visibly killed by q_g, and q_g is surjective because q is. Again no injectivity or surjectivity of g was required.

The left square is a pushout. Given maps $r : A' \to T$ and $s : E \to T$ with $rg = si$, the formula

$$[a', e] \longmapsto r(a') + s(e)$$

defines a linear map $g_*E \to T$: the map on the direct sum kills $(g(a), -i(a))$. It has the prescribed composites with i_g and u_g. Uniqueness follows because every class is $[a', 0] + [0, e]$.

These two constructions have opposite directions for a reason. A map *into* the quotient permits us to restrict the extension over a new domain. A map *out of* the subobject permits us to impose a new identification in the kernel. Their universal properties specify exactly which compatible maps survive the construction.

9.4 Adding extensions by merging their kernels

Suppose

$$\xi_j : \quad 0 \longrightarrow A \xrightarrow{i_j} E_j \xrightarrow{q_j} B \longrightarrow 0 \qquad (j = 1, 2)$$

are two extensions with the same endpoints. Taking their direct sum gives two copies of B and two copies of A. We first require the quotient coordinates to agree, and then merge the kernel coordinates by addition. Thus we pull back along the diagonal $B \to B \oplus B$ and push out along the sum $A \oplus A \to A$.

Here is the resulting object in one formula. Let

$$T_2 = \{(e_1, e_2) : q_1(e_1) = q_2(e_2)\}, \qquad L_2 = \{(i_1(a), -i_2(a)) : a \in A\},$$

and set $Q_2 = T_2/L_2$. Both are submodules of the relevant direct sum, and $L_2 \subseteq T_2$. Define

$$i_+(a) = [i_1(a), 0] = [0, i_2(a)], \qquad q_+([e_1, e_2]) = q_1(e_1) = q_2(e_2).$$

The equality defining i_+ is exactly the relation imposed by L_2. The map i_+ is injective: if $(i_1(a), 0) = (i_1(c), -i_2(c))$, then $c = 0$ and $a = 0$. Every pair in the kernel of q_+ has the form $(i_1(a), i_2(b))$, and its class is $i_+(a + b)$. Surjectivity follows by lifting a given $b \in B$ through both q_1 and q_2. Hence

$$\xi_1 + \xi_2 : \quad 0 \longrightarrow A \xrightarrow{i_+} Q_2 \xrightarrow{q_+} B \longrightarrow 0$$

is an extension. This operation is called the *Baer sum*.

To check the asserted pullback–pushout description, the pullback is T_2, with kernel $A \oplus A$. Its pushout along addition is

$$\frac{A \oplus T_2}{\{(a + b, (-i_1(a), -i_2(b))) : a, b \in A\}}.$$

It maps to Q_2 by

$$[c, (e_1, e_2)] \longmapsto [e_1 + i_1(c), e_2].$$

A displayed relation maps to $[i_1(b), -i_2(b)] = 0$. Its inverse sends $[e_1, e_2]$ to $[0, (e_1, e_2)]$; the relation defining L_2 maps to a pushout relation with $b = -a$. The two composites are identities, since $[c, (e_1, e_2)] = [0, (e_1 + i_1(c), e_2)]$ in the pushout. This verifies the construction, including its endpoint maps.

Proposition 9.3. *Baer sum makes the equivalence classes of extensions of B by A into an abelian group. Its zero is the split extension. An inverse of (i, E, q) is $(-i, E, q)$.*

Proof. First the operation is well-defined on classes. Equivalences $h_j : E_j \to E'_j$ send T_2 into T'_2 and L_2 into L'_2 coordinatewise. They induce a map of extensions, which is an isomorphism by Lemma 9.1.

For associativity, form the common-image tuples for any $n \geq 1$:

$$T_n = \{(e_1, \dots, e_n) : q_1(e_1) = \cdots = q_n(e_n)\},$$

and the relation submodule

$$L_n = \{(i_1(a_1), \dots, i_n(a_n)) : a_1 + \cdots + a_n = 0\}.$$

The quotient $Q_n = T_n/L_n$ is an extension of B by A. Its quotient map takes the common image; its kernel map sends a to $[i_1(a), 0, \dots, 0]$. This kernel map is injective, because a tuple of that form belongs to L_n only if $a = 0$. A tuple in the kernel of the quotient map has entries $i_j(a_j)$; modulo L_n it equals the kernel image of $\sum_j a_j$. Finally, a given element of B can be lifted in each coordinate. These statements prove exactness.

In the following calculation, $E_1 + E_2$ denotes the middle module of the Baer sum, with its displayed endpoint maps. For $n = 3$ the map

$$T_3 \longrightarrow (E_1 + E_2) + E_3, \qquad (e_1, e_2, e_3) \longmapsto [[e_1, e_2], e_3]$$

is linear. If $a_1 + a_2 + a_3 = 0$, its value on $(i_1(a_1), i_2(a_2), i_3(a_3))$ is the class of $(i_{12}(a_1 + a_2), i_3(a_3))$, which is zero by the outer sum relation. Here i_{12} denotes the kernel inclusion for $E_1 + E_2$. Thus it factors through Q_3. The induced map is the identity on the fixed endpoints, so it is an equivalence. The map

$$(e_1, e_2, e_3) \longmapsto [e_1, [e_2, e_3]]$$

likewise identifies Q_3 with $E_1 + (E_2 + E_3)$: a kernel tuple maps to the relation determined by $a_1 + (a_2 + a_3) = 0$. Both bracketings therefore represent the same extension class. This proves associativity without assuming it for the new operation.

Swapping the two coordinates identifies $E_1 + E_2$ with $E_2 + E_1$. It preserves the kernel inclusion because $[i_1(a), 0] = [0, i_2(a)]$, and preserves the common quotient image. This proves commutativity.

For the split extension $S = A \oplus B$, the formula

$$E + S \longrightarrow E, \qquad [e, (a, b)] \longmapsto e + i(a)$$

is defined on pairs with $q(e) = b$. The relation $(i(c), (-c, 0))$ maps to zero, and both endpoint maps are preserved. Lemma 9.1 makes it an equivalence, proving the identity law.

Finally, in $E+(-E)$, where $-E$ denotes $(-i, E, q)$, the relation submodule consists of pairs $(i(a), i(a))$. The linear map $e \mapsto [e, e]$ kills $i(A)$, so it factors uniquely through $E/i(A)$. The quotient map identifies $E/i(A)$ with B; hence we obtain a linear map $s : B \to E + (-E)$ with $q_+ s = \mathrm{id}_B$. By Proposition 9.2, this sum is split. The inverse law follows. □

These equivalence classes form a set. Indeed, choosing a set-theoretic lift of every element of B identifies the underlying set of a middle module with $A \times B$, by $(a, b) \mapsto i(a) + s(b)$; this uses only a choice of representatives, not a linear section. Module operations and endpoint maps on this fixed set range over sets of functions. Thus the extension classes can be represented by a set of such data. For arbitrary modules this set-theoretic observation uses the usual axiom of choice; none of the finite-filtration comparison proofs needs it.

We denote this abelian group by $\mathrm{Ext}^1_R(B, A)$. The order of the arguments records the sequence: A is the kernel and B the quotient. The superscript 1 refers to a single extension step. No higher extension theory is needed for what follows. These are classical constructions; the notation and the general pullback–pushout description are also recorded in the Stacks Project, Section 12.6.

The new associativity has a concrete source. Three extensions can be added by retaining triples with a common image and then imposing the single condition that the three kernel coordinates sum to zero. Either pairwise construction gives that same object up to the specified isomorphism. The quotient has removed the dependence on how the addition was parenthesized.

9.5 An extension that fits in one matrix

Let k be a field. For $\lambda \in k$, let S_λ be the one-dimensional $k[t]$-module on which t acts by multiplication by λ. It is simple because a nonzero vector spans its underlying one-dimensional space. Consider

$$0 \longrightarrow S_\mu \xrightarrow{i} E \xrightarrow{q} S_\lambda \longrightarrow 0.$$

Write T for the action of t. Choose $u = i(1)$ and choose $v \in E$ with $q(v) = 1$. These form a k-basis: every $e \in E$ can first have its quotient coordinate removed by subtracting $q(e)v$, leaving an element of $i(S_\mu)$; and applying q to $cu + dv = 0$ gives $d = 0$, followed by $c = 0$ since i is injective.

The endpoint maps commute with the operators, so

$$Tu = \mu u, \qquad Tv = \lambda v + au$$

for a unique $a \in k$. In this basis,

$$[T]_{(u,v)} = \begin{pmatrix} \mu & a \\ 0 & \lambda \end{pmatrix}.$$

Conversely, for any a, these formulas define a linear operator on $ku \oplus kv$, hence a $k[t]$-module E_a, and an extension with $i(1) = u$, $q(u) = 0$, and $q(v) = 1$. Thus every extension under discussion has this form.

The choice of u is fixed by the marked endpoint, whereas another lift of 1 has the form $v' = v + bu$. Its coefficient is found by one calculation:

$$T(v + bu) = \lambda(v + bu) + \big(a + b(\mu - \lambda)\big)u.$$

The change of lift acts on the attachment coefficient by

$$a \longmapsto a + b(\mu - \lambda).$$

Proposition 9.4. *The preceding coefficient gives an isomorphism of abelian groups*

$$\operatorname{Ext}^1_{k[t]}(S_\lambda, S_\mu) \cong k/(\mu - \lambda)k.$$

In particular, extensions split when $\lambda \neq \mu$. When $\lambda = \mu$, each $a \in k$ gives a distinct marked extension class, and its class is zero exactly when $a = 0$.

Proof. Changing the lift changes a by an element of $(\mu - \lambda)k$, so its coset is independent of that choice. An equivalence fixes u and sends a lift of 1 to a lift of 1; thus it preserves the same coset. Conversely, if $a' = a + b(\mu - \lambda)$, the map $E_{a'} \to E_a$ sending $u' \mapsto u$ and $v' \mapsto v + bu$ commutes with T by the displayed calculation. It therefore commutes with every polynomial in T, is $k[t]$-linear, and is an equivalence of extensions. The classification is consequently bijective.

To compute its addition, take E_a and E_c with chosen vectors u_1, v_1 and u_2, v_2. In their Baer sum put

$$u_+ = [u_1, 0] = [0, u_2], \qquad v_+ = [v_1, v_2].$$

The first vector is the marked kernel image of 1 and the second maps to 1 in the quotient. The same spanning and independence argument used above makes them a basis. Their operator satisfies

$$Tu_+ = \mu u_+, \qquad Tv_+ = \lambda v_+ + [au_1, cu_2] = \lambda v_+ + (a + c)u_+.$$

Thus Baer sum corresponds to addition of coefficient classes.

The split extension has coefficient zero. If $\mu - \lambda \neq 0$, it is invertible in k, so $(\mu - \lambda)k = k$ and the quotient group is zero. If $\mu = \lambda$, the ideal is zero, giving the stated result. □

When the eigenvalues differ, the explicit section can be read off as well. Choose $b = -a/(\mu - \lambda)$. Then $v' = v + bu$ satisfies $Tv' = \lambda v'$, and the map $S_\lambda \to E$, $c \mapsto cv'$, is $k[t]$-linear and is a section of q.

When the eigenvalues agree, a nonzero a cannot be removed by changing the lift. Yet all nonzero a give isomorphic *middle modules*. Indeed, put $w = au$; then

$$Tw = \lambda w, \qquad Tv = \lambda v + w,$$

which has the same matrix for every nonzero a. This rescaling changes the marked kernel vector, so it is not an equivalence keeping the endpoints fixed. The $a = 0$ middle module is different: $T - \lambda\mathrm{id}$ is zero there, whereas it is nonzero for $a \neq 0$, a property preserved by every intertwining isomorphism.

For $\lambda = \mu = 1$, the cases $a = 0$ and $a = 1$ give the identity operator and the triangular shear $(x, y) \mapsto (x + y, y)$. Both layers carry the identity action. The coefficient a now records explicitly the attachment that their common list of layers forgot.

There are therefore two different classification questions. One asks for the middle operator up to change of basis. The other remembers a specified subobject and a specified quotient map. The latter contains more information, which is why it supports the extension group just constructed.

9.6 A boundary map measures failure to lift

Return to an arbitrary extension

$$\xi: \quad 0 \longrightarrow A \xrightarrow{i} E \xrightarrow{q} B \longrightarrow 0.$$

Given $f : P \to B$, we can now form a precise obstruction to finding $\widetilde{f} : P \to E$ with $q\widetilde{f} = f$: take the pullback extension and its class. Define

$$\delta_\xi : \mathrm{Hom}_R(P, B) \longrightarrow \mathrm{Ext}^1_R(P, A), \qquad f \longmapsto [f^*E].$$

If $\widetilde{f}$ exists, then $p \mapsto (\widetilde{f}(p), p)$ is a section of the pullback extension. Conversely, a section of its projection to P has exactly this form, and its first coordinate is a lift. Proposition 9.2 therefore proves

$$\delta_\xi(f) = 0 \quad \Longleftrightarrow \quad f \text{ has a linear lift through } q.$$

This is a homomorphism of abelian groups. To verify the assertion, take $f, g : P \to B$. The map

$$(f^*E) + (g^*E) \longrightarrow (f + g)^*E, \qquad [(e_1, p), (e_2, p)] \longmapsto (e_1 + e_2, p)$$

is well-defined: the target condition is $q(e_1+e_2) = f(p)+g(p)$, and a Baer relation $((i(a), 0), (-i(a), 0))$ maps to zero. It is linear and preserves both endpoint maps. Lemma 9.1 makes it an equivalence, proving $\delta_\xi(f + g) = \delta_\xi(f) + \delta_\xi(g)$. The pullback along the zero map has section $p \mapsto (0, p)$, so $\delta_\xi(0) = 0$.

In particular, we have proved exactness of the sequence of abelian groups

$$0 \longrightarrow \mathrm{Hom}_R(P, A) \xrightarrow{i\circ -} \mathrm{Hom}_R(P, E) \xrightarrow{q\circ -} \mathrm{Hom}_R(P, B) \xrightarrow{\delta_\xi} \mathrm{Ext}^1_R(P, A).$$

Here exactness is asserted at the three Hom groups. The first map is injective because i is injective. A map $h : P \to E$ is killed by $q \circ -$ exactly when its image lies in $i(A)$; it then factors uniquely as $h = i\bar{h}$, with $\bar{h}$ linear because i is an injective linear map. Exactness at $\mathrm{Hom}_R(P, B)$ is the lifting criterion proved above. No surjectivity onto $\mathrm{Ext}^1_R(P, A)$ is claimed.

For example, if $P = B$ and $f = \mathrm{id}_B$, its pullback is equivalent to the original extension by $e \mapsto (e, q(e))$. Thus $\delta_\xi(\mathrm{id}_B) = [\xi]$: lifting the identity is exactly splitting the extension. In the equal-eigenvalue example, this class is the coefficient a. A section would require an eigenvector above 1; the calculation shows exactly when such a vector exists.

There is also a useful contrast. If $P = R$, every map $f : R \to B$ lifts: choose e with $q(e) = f(1)$ and define $\widetilde{f}(r) = re$. This is linear and satisfies $q\widetilde{f}(r) = rf(1) = f(r)$. The obstruction depends on the object from which we test the quotient. A test object with relations can detect constraints that a single free generator cannot.

9.7 An attachment can skip an adjacent layer

It is tempting to recover a filtered object by listing its simple layers and the extension between every adjacent pair. Even that list can miss information. The following example identifies the loss without using any classification theorem.

On k^3 fix the flag

$$0 = F_0 \subset F_1 = ke_1 \subset F_2 = ke_1 + ke_2 \subset F_3 = k^3.$$

Compare the zero operator with the operator N given by

$$N(e_1) = N(e_2) = 0, \qquad N(e_3) = e_1.$$

Both preserve every F_i, so these are filtrations by $k[t]$-submodules. Each of the three successive one-dimensional quotients has zero induced operator in both examples, hence is S_0.

Now retain the adjacent two-layer pieces as well. On $F_2/F_0 = F_2$ the operator is zero in both examples. On F_3/F_1 it is also zero, because $N(F_3) \subseteq F_1$. Thus the marked extensions

$$0 \to F_1/F_0 \to F_2/F_0 \to F_2/F_1 \to 0, \qquad 0 \to F_2/F_1 \to F_3/F_1 \to F_3/F_2 \to 0$$

are split and agree, using the indicated coordinate identifications, for the two operators. Nevertheless the full modules are not isomorphic. An isomorphism intertwining N with zero would satisfy $hN = 0$; injectivity of h would force $N = 0$, contrary to $N(e_3) = e_1$.

The map from the third layer into the first is invisible in either adjacent two-layer quotient. It remains visible if we retain the whole extension

$$0 \longrightarrow F_2 \longrightarrow F_3 \longrightarrow F_3/F_2 \longrightarrow 0$$

with its maps: in the N example every lift of the quotient generator has the form $e_3 + x$ with $x \in F_2$, and $N(e_3 + x) = e_1 \neq 0$, so no $k[t]$-linear section exists. For the zero operator the evident section does exist.

The example specifies what the proposed summary discarded. Complete successive extension data includes each already assembled module and the maps attaching the next layer. Replacing those data by isolated adjacent two-layer summaries can erase a longer attachment. We now seek observations that retain enough information to reconstruct all finite-dimensional nilpotent operators.

10 Recovering a nilpotent operator from spaces of maps

A nilpotent operator can be described by the motion of vectors along chains that eventually reach zero. We want to find those chains from measurements that can be computed directly, rather than assume a normal form in order to read them off.

For a first example, choose chains of lengths $3, 2, 2$: the operator sends each vector to its predecessor, and sends the first vector in each chain to zero. The seven chain vectors form a basis. The successive kernel sizes, starting with $\ker N^0 = 0$, are

$$(d_0, d_1, d_2, d_3, d_4) = (0, 3, 6, 7, 7).$$

Their first differences are $3, 3, 1, 0$: a chain contributes one new kernel vector at step r exactly when its length is at least r. Subtract once more to obtain $0, 2, 1$, the numbers of chains of lengths $1, 2, 3$. Thus a classification question has suggested a calculation with finite differences. The proof below must still earn the reverse direction: construct chains for an arbitrary nilpotent operator from its kernels and images, and show that these measurements determine their lengths.

Our objects will be finite-dimensional vector spaces over an arbitrary field k, equipped with a nilpotent operator. The answer will use only kernels, images, quotients, and finite bases. In particular, we will prove the required decomposition into chains. No classification of modules over a polynomial ring, or Jordan normal form theorem, is assumed.

10.1 A relation becomes a test object

Let V be a finite-dimensional k-vector space and let $N : V \to V$ be linear. Recall how N gives V the structure of a $k[t]$-module M:

$$\left(\sum_i a_i t^i\right) \cdot v = \sum_i a_i N^i v, \qquad N^0 = \mathrm{id}_V.$$

The identity $N^iN^j = N^{i+j}$ ensures that multiplication of polynomials agrees with composition of their actions. A linear map $F : V \to V'$ is a $k[t]$-module map precisely when $FN = N'F$: necessity follows by considering t, and sufficiency follows by induction for powers of t and then by linearity for polynomials.

Assume henceforth that $N^R = 0$ for some integer $R \geq 1$. This is the meaning of nilpotence. We allow $V = 0$.

For every $r \geq 1$, introduce the module

$$P_r = k[t]/(t^r).$$

Its elements have unique representatives $a_0 + a_1t + \cdots + a_{r-1}t^{r-1}$: remove all terms divisible by t^r, and observe that a polynomial of degree less than r divisible by t^r is zero. Thus $1, t, \ldots, t^{r-1}$, understood as residue classes, form a k-basis. Multiplication by t runs along this basis and then gives zero.

The module P_r has one generator with one relation, $t^r \cdot 1 = 0$. A map out of P_r asks whether a proposed image of the generator satisfies that relation.

Proposition 10.1 (A kernel represented by a module)**.** *Evaluation at* 1 *gives a* k*-linear isomorphism*

$$\mathrm{ev}_1 : \mathrm{Hom}_{k[t]}(P_r, M) \xrightarrow{\sim} \ker N^r, \qquad f \longmapsto f(1).$$

Its inverse sends $v \in \ker N^r$ *to the map*

$$f_v : P_r \longrightarrow M, \qquad \overline{p(t)} \longmapsto p(N)v.$$

Proof. If f is a module map, then $N^r f(1) = f(t^r \cdot 1) = 0$. Moreover, $f(\overline{p(t)}) = p(N)f(1)$, so $f(1)$ determines f.

Conversely, if $N^r v = 0$ and $p - p' = t^r q$, then $(p - p')(N)v = q(N)N^r v = 0$. Thus f_v is well-defined on residue classes. It is additive and k-linear, and

$$f_v(\overline{a(t)p(t)}) = a(N)p(N)v = a(t) \cdot f_v(\overline{p(t)}),$$

so it is a module map. The two constructions are inverse. Addition and scalar multiplication of module maps are defined pointwise; the module identities verify that these remain module maps. Evaluation and its inverse respect these operations. □

This isomorphism respects maps between the objects being measured. If $F : M \to M'$ is a module map, then $FN^r = (N')^r F$, so F sends $\ker N^r$ into $\ker(N')^r$. The square

$$\begin{array}{ccc} \mathrm{Hom}_{k[t]}(P_r, M) & \xrightarrow[\sim]{\mathrm{ev}_1} & \ker N^r \\ \downarrow{\scriptstyle f \mapsto F \circ f} & & \downarrow{\scriptstyle F|_{\ker N^r}} \\ \mathrm{Hom}_{k[t]}(P_r, M') & \xrightarrow[\sim]{\mathrm{ev}_1} & \ker (N')^r \end{array}$$

commutes because both paths send f to $F(f(1))$. This compatibility is what it means here for the identification to be *natural* in M.

There is also compatibility as we enlarge the test object. The quotient map $q_r : P_{r+1} \to P_r$ sends $\overline{p}$ to $\overline{p}$; it is well-defined because $(t^{r+1}) \subseteq (t^r)$. Precomposition by q_r corresponds, under evaluation at 1, to the inclusion $\ker N^r \subseteq \ker N^{r+1}$. The growing sequence of kernels is thus also a growing sequence of spaces of maps from explicit test objects.

10.2 Which endpoints can be reached by a long chain?

Set

$$K_r = \ker N^r, \qquad d_r = \dim K_r \quad (r \geq 0).$$

Since $N^0 = \mathrm{id}_V$, $K_0 = 0$ and $d_0 = 0$. For $r \geq R$, we have $K_r = V$. The numbers d_r are now well-defined by the dimension theorem already proved. By Proposition 10.1, they are precisely the dimensions of the spaces of maps from the P_r.

To understand these numbers, look at a second filtration, this time inside $\ker N$:

$$B_r = \ker N \cap \operatorname{im} N^{r-1} \qquad (r \geq 1).$$

Every B_r is a subspace, and

$$\ker N = B_1 \supseteq B_2 \supseteq \cdots \supseteq B_R \supseteq B_{R+1} = 0.$$

An element of B_r is killed by N and can be reached by applying N $r-1$ times. It can therefore be the endpoint of a chain of length r. Membership in B_{r+1} says that it can be reached by an even longer chain.

There is a short exact sequence of vector spaces

$$0 \longrightarrow K_{r-1} \longrightarrow K_r \xrightarrow{N^{r-1}} B_r \longrightarrow 0. \tag{22}$$

The middle map has the stated target because $N^r x = 0$ for $x \in K_r$. Its kernel is K_{r-1}. It is surjective: if $y \in B_r$, write $y = N^{r-1}x$; then $0 = Ny = N^r x$, so $x \in K_r$. The first arrow is inclusion, and the end arrows are the unique maps to or from the zero space.

Finite-dimensional dimension additivity applied to this sequence gives

$$\dim B_r = d_r - d_{r-1}. \tag{23}$$

Consequently the successive increments in the sequence (d_r) are nonnegative and nonincreasing. Define

$$E_r = B_r / B_{r+1}, \qquad b_r = \dim E_r. \tag{24}$$

Taking dimensions in the quotient gives

$$\boxed{b_r = 2d_r - d_{r-1} - d_{r+1}.} \tag{25}$$

This equation has been proved before any choice of chains. In particular, its right side is nonnegative because it is the dimension of an actual quotient space. We next prove why that quotient counts chains of exactly length r.

10.3 Lifting the endpoint filtration to a basis of chains

For each r, choose vectors

$$u_{r,1}, \ldots, u_{r,b_r} \in B_r$$

whose residue classes form a basis of B_r/B_{r+1}. Such vectors exist by finite-dimensional basis extension, proved after the dimension theorem. Equivalently, start from the empty basis of $B_{R+1} = 0$ and extend a basis successively as we pass to $B_R, B_{R-1}, \ldots, B_1$.

It follows that, for every $h \geq 1$, the vectors

$$\{u_{r,j} : h \leq r \leq R,\ 1 \leq j \leq b_r\} \quad \text{form a basis of } B_h. \tag{26}$$

Here is the elementary reason. If a subspace $W \subseteq U$ has a basis and we lift a basis of U/W, the two lists together span U: subtract a combination of the lifts to leave an element of W. They are independent: reduce a relation modulo W first, and then use independence in W. Applying this observation down the filtration proves (26).

Because $u_{r,j} \in \operatorname{im} N^{r-1}$, choose $v_{r,j} \in V$ with

$$N^{r-1} v_{r,j} = u_{r,j}.$$

Since $u_{r,j} \in \ker N$, we also have $N^r v_{r,j} = 0$. For $1 \leq h \leq r$, write

$$w_{r,j,h} = N^{r-h} v_{r,j}.$$

The index h is the height above zero: $w_{r,j,1} = u_{r,j}$, $N w_{r,j,1} = 0$, and $N w_{r,j,h} = w_{r,j,h-1}$ for $h > 1$. For example, a chain of length three has the form

$$w_{3,j,3} \xmapsto{N} w_{3,j,2} \xmapsto{N} w_{3,j,1} \xmapsto{N} 0.$$

The nodes in this last diagram are vectors, and the arrows show the action of N on them; it is not a diagram of vector spaces and linear maps.

Theorem 10.2 (A basis of nilpotent chains)**.** *The vectors $w_{r,j,h}$, for $1 \leq r \leq R$, $1 \leq j \leq b_r$ and $1 \leq h \leq r$, form a basis of V. Consequently there is a $k[t]$-module isomorphism*

$$M \cong \bigoplus_{r=1}^{R} P_r^{\oplus b_r}.$$

Proof. First prove independence. Suppose that a linear combination of these vectors is zero. If any coefficient is nonzero, let H be the largest height h at which a nonzero coefficient occurs. Apply N^{H-1} to the relation. Every term of height less than H vanishes, and a term $w_{r,j,H}$ becomes $u_{r,j}$. We obtain a relation among the vectors $u_{r,j}$ with $r \geq H$. These vectors are independent by (26). Thus every coefficient at height H is zero, contradicting its choice. No nonzero relation exists.

For spanning, prove by induction on h that K_h lies in the span of the chosen chain vectors. The case $h = 0$ is immediate. Given $x \in K_h$, the vector $N^{h-1}x$ lies in B_h. By (26), write

$$N^{h-1}x = \sum_{r \geq h,\, j} c_{r,j} u_{r,j}.$$

Each $u_{r,j}$ in this sum has the chosen lift $w_{r,j,h}$ at height h. Therefore

$$x - \sum_{r \geq h,\, j} c_{r,j} w_{r,j,h} \in K_{h-1}.$$

The induction hypothesis puts this difference in the required span, and so puts x there. Since $K_R = V$, spanning follows.

For the module statement, on the copy of P_r indexed by (r, j) define

$$\overline{p(t)} \longmapsto p(N) v_{r,j}.$$

This is a well-defined module map because $N^r v_{r,j} = 0$, by the same argument as Proposition 10.1. Combining these maps gives a map from the displayed finite direct sum to M. It sends the standard basis vectors $1, t, \dots, t^{r-1}$ in each summand to the corresponding chain, in reverse height order. The basis statement makes this map bijective; its inverse respects the module operations because the map is bijective and respects those operations. It is therefore a module isomorphism. □

The choices in this proof are visible. We selected representatives of quotient bases and then lifted them through powers of N. Different choices can give different chain bases. The spaces B_r, the quotients E_r, and their dimensions b_r depend only on N.

In particular, the construction proves the precise interpretation

$$\boxed{\dim \frac{\ker N \cap \operatorname{im} N^{r-1}}{\ker N \cap \operatorname{im} N^r} = \text{the number of chains of length exactly } r.}$$

The quotient removes endpoints which can be lifted farther. A basis of what remains supplies the new chains of length r in the construction.

10.4 Reading the operator from its measurements

In one chain of length s, the kernel of N^r is spanned by the bottom $\min(r, s)$ vectors. The direct-sum decomposition therefore gives

$$d_r = \sum_{s=1}^{R} b_s \min(r, s), \qquad d_r - d_{r-1} = \sum_{s \geq r} b_s. \tag{27}$$

The increment counts chains of length at least r. Subtracting successive increments gives exactly the second-difference formula (25), which we had already obtained from canonical subquotients.

Corollary 10.3 (Reconstruction by a finite nilpotent profile)**.** *Two finite-dimensional nilpotent $k[t]$-modules M and M' are isomorphic if and only if*

$$\dim_k \operatorname{Hom}_{k[t]}(P_r, M) = \dim_k \operatorname{Hom}_{k[t]}(P_r, M') \qquad \text{for every } r \geq 1.$$

If a common nilpotence bound R is known, it suffices to use $1 \leq r \leq R$.

Proof. A module isomorphism identifies the spaces of maps by postcomposition, so it preserves all the displayed dimensions. Conversely, equality of the dimensions means equality of the d_r by Proposition 10.1. The formula $b_r = 2d_r - d_{r-1} - d_{r+1}$ then gives the same chain multiplicities for both modules. Theorem 10.2 identifies each with the same direct sum of the P_r, so they are isomorphic.

For a known common bound R, put $d_0 = 0$ and $d_{R+1} = d_R$; nilpotence justifies the latter equality. The first R measurements therefore determine every possibly nonzero b_r. □

Thus the dimensions of these particular spaces of maps determine the isomorphism class in this particular category. They do not choose an isomorphism or a preferred basis. The proof succeeds because we established a decomposition theorem and then showed that the measurements recover every multiplicity in that decomposition.

Here are the first distinctions. Write J_s for multiplication by t on P_s; as an operator, it is a single nilpotent chain of length s. In dimension three we obtain

Operator	(d_1, d_2, d_3)	(b_1, b_2, b_3)
Zero on k^3	$(3, 3, 3)$	$(3, 0, 0)$
$J_2 \oplus J_1$	$(2, 3, 3)$	$(1, 1, 0)$
J_3	$(1, 2, 3)$	$(0, 0, 1)$

The skipped-attachment operator from Section 9.7, with $N(e_3) = e_1$ and $N(e_1) = N(e_2) = 0$, is the middle row: (e_3, e_1) is a chain of length two and e_2 is a chain of length one. Its first measurement is $d_1 = 2$, whereas the zero operator has $d_1 = 3$. Thus even maps from P_1 recover the distinction that both adjacent two-layer summaries missed.

All three have three simple factors isomorphic to $S_0 = P_1 = k[t]/(t)$. Indeed, in a chain the subspaces spanned by its first h bottom vectors are submodules, and each successive quotient is a copy of S_0. Concatenating these filtrations across the direct sum gives a composition series. The space S_0 is simple because it is one-dimensional over k and every submodule is a k-subspace. Consequently all three examples have the same class $3[S_0]$ in the Grothendieck group. Their spaces of maps from P_1, P_2, P_3 distinguish them.

This does not contradict the universality of the additive invariant: these dimensions of spaces of maps need not be additive on short exact sequences. For example,

$$0 \longrightarrow P_1 \xrightarrow{1 \mapsto t} P_2 \xrightarrow{\overline{p} \mapsto \overline{p}} P_1 \longrightarrow 0$$

is exact. The first map identifies P_1 with the submodule kt of P_2, and the second map is reduction modulo t, with kernel kt. Yet $\dim \mathrm{Hom}_{k[t]}(P_1, P_2) = 1$, whereas the two endpoint measurements sum to $1 + 1 = 2$. A module map $P_1 \to P_2$ must send 1 into kt, so its composition with $P_2 \to P_1$ is zero. In particular, the identity map of the quotient P_1 cannot lift to P_2. The failure of additivity is the lifting obstruction from the preceding chapter, appearing in an explicit measurement.

For a slightly larger calculation, suppose measurements give

$$(d_0, d_1, d_2, d_3, d_4, d_5) = (0, 3, 5, 6, 7, 7)$$

for an operator with $N^4 = 0$. The increments are $3, 2, 1, 1, 0$. Their successive differences give

$$(b_1, b_2, b_3, b_4) = (1, 1, 0, 1).$$

Hence $M \cong P_1 \oplus P_2 \oplus P_4$. The dimensions add to $1 + 2 + 4 = 7$, and the kernel dimensions of this sum, calculated as $\min(r, 1) + \min(r, 2) + \min(r, 4)$, recover the original measurements.

The scope matters. These same tests do not reconstruct an arbitrary operator. If t acts as the identity on a nonzero vector space, every N^r is the identity and every $\mathrm{Hom}_{k[t]}(P_r, M)$ is zero. Its profile is therefore the same as that of the zero module. The nilpotence assumption ensured that our tests eventually reach all of M; without it, an entire part of the object can remain undetected.

We have now proved a concrete reconstruction result. Composition factors identify the simple layers. The canonical subquotients E_r measure how far their endpoints can be lifted along the operator. The test modules P_r turn those lifting depths into dimensions of spaces of maps, and the chain-basis theorem shows that this information is sufficient to recover the finite-dimensional nilpotent object up to isomorphism.

11 Associativity by counting filtrations

An extension records an attachment. We can also ask how many attachments of one type occur inside a fixed middle object. This question leads to a new multiplication. Its associativity will follow from the correspondence between submodules of a quotient and submodules containing its kernel. The quotient calculation that began the exposition will thus return as the mechanism behind an algebra law.

11.1 A setting where the counts are finite

Fix a finite field k with q elements. We work with finite-dimensional $k[t]$-modules on which t acts nilpotently. Equivalently, our objects are pairs (V, N) with $N^s = 0$ for some s, and our maps are linear maps commuting with the operators. A submodule is exactly an N-invariant subspace. Throughout this chapter an isomorphism must respect N.

Submodules and quotients stay in this class because the same power of N is zero on them. Extensions stay in the class too. Indeed, in an exact sequence

$$0 \longrightarrow B \longrightarrow E \longrightarrow A \longrightarrow 0,$$

if $t^a A = 0$ and $t^b B = 0$, then $t^a E \subseteq B$ and hence $t^{a+b} E = 0$. Also $\dim E = \dim A + \dim B$, by the proved dimension formula for an exact sequence. These are closure checks, not extra assumptions about the category.

For isomorphism classes A, B, E, define the nonnegative integer

$$g^E_{A,B} = \#\{U \subseteq E : \ U \text{ is a submodule}, \ U \cong B, \ E/U \cong A\}. \tag{28}$$

The order matters: B is the lower layer and A is the quotient above it. This counts actual submodules of E. It does not count chosen isomorphisms $B \to U$, chosen quotient identifications, or marked extension classes. An isomorphism $E \to E'$ gives a bijection on these sets by sending U to its image, so the count depends only on the three isomorphism classes.

The set in (28) is finite because E has finitely many elements. For fixed A, B, only finitely many classes E can occur. They have the fixed dimension $d = \dim A + \dim B$. Choosing a basis identifies such an operator with a d by d matrix over k, and there are q^{d^2} matrices altogether. Each isomorphism class has a representative among them. This argument establishes finiteness without requiring a classification of nilpotent operators.

Let $\mathcal{H}$ be the rational vector space with one basis symbol u_M for each isomorphism class in our collection. Explicitly, its vectors are finite rational linear combinations of these symbols. Define a bilinear multiplication by

$$u_A u_B = \sum_{[E]} g^E_{A,B} u_E. \tag{29}$$

The preceding finiteness checks make this an element of $\mathcal{H}$. It is the untwisted Hall multiplication for this collection of modules. We now prove its algebra laws.

11.2 One flag, two intermediate objects

Theorem 11.1 (Associativity from a flag bijection)**.** *The multiplication* (29) *is associative, with identity* u_0.

Proof. Fix A, B, C, E. Count the set of flags

$$0 \subseteq U \subseteq V \subseteq E, \qquad U \cong C, \quad V/U \cong B, \quad E/V \cong A. \tag{30}$$

Inclusions may be equal when one of the specified factors is zero.

First fix U. If $E/U \cong D$, there are $g^E_{D,C}$ choices of U. For each such U, the possible submodules $V/U \subseteq E/U$ with the required factors are counted by $g^D_{A,B}$. The quotient correspondence gives a bijection

$$\{V : U \subseteq V \subseteq E\} \longleftrightarrow \{W : W \subseteq E/U\}, \quad V \longmapsto V/U, \quad W \longmapsto \pi^{-1}(W).$$

It also gives $(E/U)/(V/U) \cong E/V$, by the map $(e + U) + (V/U) \mapsto e + V$. Thus the number of flags is

$$\sum_{[D]} g^D_{A,B} g^E_{D,C}.$$

This is the coefficient of u_E in $(u_A u_B)u_C$.

Alternatively fix V. If $V \cong D$, there are $g^E_{A,D}$ choices of V, and then $g^D_{B,C}$ choices of $U \subseteq V$. The same flags are therefore counted by

$$\sum_{[D]} g^D_{B,C} g^E_{A,D},$$

the coefficient of u_E in $u_A(u_B u_C)$. No choice of an isomorphism to a representative D affects either count: every such isomorphism induces a bijection of the relevant submodule sets. Both counting procedures identify their data with precisely the flags in (30). Equality of all coefficients proves associativity on basis vectors, and bilinearity proves it on $\mathcal{H}$.

For the unit, the only submodule of E isomorphic to zero is 0, and the only submodule with zero quotient is E. Consequently $g^E_{A,0} = g^E_{0,A}$ is 1 if $E \cong A$ and 0 otherwise. Hence $u_A u_0 = u_0 u_A = u_A$. □

The two readings of the flag use different intermediate objects:

$$\begin{array}{ccccccccc}
0 & \longrightarrow & U & \hookrightarrow & V & \twoheadrightarrow & V/U & \longrightarrow & 0 \\
 & & \| & & \downarrow & & \downarrow & & \\
0 & \longrightarrow & U & \hookrightarrow & E & \twoheadrightarrow & E/U & \longrightarrow & 0.
\end{array}$$

Both rows are exact. The right vertical map sends $v + U$ to $v + U$; it is injective since its zero class means $v \in U$. The square on the right commutes on each representative v. The quotient of that vertical inclusion is E/V. Thus the intermediate object is either V, carrying the lower two layers, or E/U, carrying the upper two layers.

At the beginning of the exposition, associativity of the group operation was an axiom permitting us to compose translations. Here associativity of a newly defined operation is a theorem. The theorem is proved by a bijection: the two bracketings organize exactly the same data. This is a structural explanation with an explicit inverse correspondence, not an appeal to a resemblance between diagrams.

11.3 The first product, computed completely

Write $S = k[t]/(t)$ and $J_r = k[t]/(t^r)$. Thus J_r is the single nilpotent chain of length r constructed in the preceding chapter. There are two isomorphism classes of two-dimensional nilpotent modules: $S \oplus S$ and J_2. One can see this directly. If $N = 0$, the first description holds. If $N \neq 0$, a two-dimensional nilpotent operator has a chain $Nv \neq 0$, $N^2v = 0$, and (Nv, v) is a basis. For completeness, one may obtain such a chain by taking the last two nonzero vectors in any maximal chain; independence of a chain of length three would contradict dimension two. The chain basis gives J_2.

For $S \oplus S$, every line is a submodule of type S with quotient S. There are $q + 1$ lines: the q lines spanned by $(1, a)$ for $a \in k$, and the line spanned by $(0, 1)$. In J_2, every invariant line has zero restricted operator, since a nilpotent scalar is zero in a field. It must therefore lie in $\ker N$, which is itself one line. There is exactly one choice. Hence

$$u_S^2 = (q + 1)u_{S\oplus S} + u_{J_2}. \tag{31}$$

The nonsplit middle object appears once in this product, although the marked extension calculation gave $q - 1$ nonzero classes when the endpoints are identified with S. The reason is the definition of the count. Equation (31) counts submodules of a fixed middle object. Marked extension equivalence keeps endpoint identifications; our Hall basis keeps only the middle object's isomorphism class. Neither convention can silently replace the other.

11.4 Three layers and a second count

The preceding reconstruction theorem gives exactly three nilpotent types in dimension three:

$$S^{\oplus 3}, \qquad J_2 \oplus S, \qquad J_3.$$

The coefficient of u_E in u_S^3 counts full invariant flags $0 \subset U \subset V \subset E$, with $\dim U = 1$ and $\dim V = 2$. Every factor is S, because a one-dimensional nilpotent operator is zero.

For $E = S^{\oplus 3}$ all flags are invariant. There are $(q^3 - 1)/(q - 1) = q^2 + q + 1$ choices of U: the nonzero vectors partition into sets of $q - 1$ vectors spanning the same line. For each U, the planes containing it correspond to the $q + 1$ lines in E/U. This gives $(q^2 + q + 1)(q + 1)$ flags.

For $E = J_2 \oplus S$, put $L = \operatorname{im} N$ and $K = \ker N$. Then $L \subset K$, $\dim L = 1$, and $\dim K = 2$. An invariant line U lies in K, giving $q+1$ choices. An invariant plane V contains L: the induced operator on the one-dimensional quotient E/V is nilpotent and hence zero, so $N(E) \subseteq V$. Conversely, every plane containing L is invariant. If $U = L$, there are $q + 1$ possible planes through it. For each of the other q lines $U \subset K$, the only plane containing both U and L is K. The count is $2q + 1$.

For J_3, an invariant line must equal $\ker N$, and an invariant plane must equal $\operatorname{im} N$ by the same quotient argument and their dimensions. These two subspaces form a flag, so the count is 1. We obtain

$$u_S^3 = (q + 1)(q^2 + q + 1)u_{S^{\oplus 3}} + (2q + 1)u_{J_2\oplus S} + u_{J_3}. \tag{32}$$

Associativity says that either order of the two successive Hall multiplications gives this count. The flag description explains why before any structure constants are expanded.

11.5 What the multiplication remembers

Every nonzero coefficient in $u_A u_B$ comes from a short exact sequence, so it satisfies

$$[E] = [A] + [B] \quad \text{in } K_0.$$

This makes Hall multiplication graded by the additive invariant: the degree of u_M is $[M]$. In our nilpotent collection the only simple module is S. Indeed, if a nonzero nilpotent module is simple, its nonzero kernel of t is the whole module; then every vector line is a submodule, forcing dimension one. Writing $\mathcal{N}_q$ for this nilpotent collection, we have $K_0(\mathcal{N}_q) \cong \mathbb{Z}$, and this grading is ordinary dimension. Here $K_0(\mathcal{N}_q)$ is formed using generators and exact-sequence relations inside $\mathcal{N}_q$, as in Section 8; the closure checks at the beginning of this chapter make that construction applicable.

Within a single degree, the Hall algebra retains different basis vectors for different isomorphism classes. The terms in (31) all have degree two, but the multiplication distinguishes the split object from the chain and records their different numbers of subobjects. This statement is a grading statement; it does not claim that K_0 is a quotient algebra of $\mathcal{H}$.

We have reached three distinct forms of information. Composition multiplicities retain what is additive across exact sequences. Marked extensions retain a specified two-layer attachment. Hall coefficients retain counts of subobjects and their quotients, and their associativity packages the compatibility of longer flags. In the nilpotent setting, the spaces of maps studied earlier recover the whole isomorphism class. Each construction answers a precise question about what the preceding description left unresolved.

Hall algebras can be developed much further, including settings where geometric invariants replace finite counts. That further theory requires additional constructions. The result established here is the finite counting algebra and its associativity, proved directly from flags.

Part IV
Retain the moves and construct linear coefficients

12 Transporter graphs and higher-quotient references

Consider an action of a group G on a set X. The orbit set answers which points can be carried to one another. It does not record which moves carry them. For example, let the two-element group act trivially on a one-point set. Its orbit set is a singleton, exactly as for the trivial group. But there are two moves from the point to itself in the first action, and only one in the second.

Retain a move as an arrow: an arrow $x \to y$ is a group element g with $gx = y$. Arrows compose by multiplication and reverse by inversion. The resulting groupoid distinguishes the two examples immediately. At a point x, its loops form the stabilizer of x. This is the concrete reason to change the quotient object here: the question now asks for moves and their composition, rather than only orbit membership. The later homotopy-coherent construction generalizes this change and specifies which further compatibility it retains.

The change of foundations in this section is explicit. We use spaces, homotopy limits, and a small number of stated theorems about ∞-categories. Their foundations are not proved here. A result marked *Framework theorem, quoted* is an external input; the propositions following it are deductions with their own proofs. None of these inputs is used in the earlier proof of Jordan–Hölder. Conversely, invoking the word "higher" will not be used to justify an unproved quotient law.

12.1 Reference: homotopical algebraic objects

Write $\mathcal{S}$ for the ∞-category of spaces, or homotopy types. A set is a discrete space. A space can also have paths, homotopies between paths, and further coherences. The mapping object $\mathrm{Map}(X, Y)$ is consequently a space: its points are maps, its paths are homotopies, and so on. An equivalence is the appropriate replacement for an isomorphism. Limits and colimits in this setting mean homotopy-coherent limits and colimits.

An ordinary one-sorted algebraic theory $\mathbb{T}$ is a small category with finite products, generated under products by an object t. The objects t^n stand for n inputs. An arrow $t^n \to t$ stands for an n-ary operation, and composition records substitution. Equations between arrows record equations between operations. A model in a category with finite products is a finite-product-preserving functor out of $\mathbb{T}$.

For groups, multiplication is an arrow $t^2 \to t$, inverse is $t \to t$, and identity is $t^0 \to t$. Associativity is equality of the two arrows $t^3 \to t$ obtained by substitution. For left R-modules we include addition, zero, inverse, and one unary operation for each $r \in R$. The equation for multiplication of scalars is an equation between composites of these unary operations. Thus the presentation does not discard the ring when it passes to a common language.

Define

$$\mathcal{C}_{\mathbb{T}} = \mathrm{Fun}^{\times}(\mathbb{T}, \mathcal{S}). \tag{33}$$

Preservation of products now means preservation up to the specified coherent equivalences. This is the promised single construction: change the target of the theory from sets to spaces, retaining the operations and their compatibilities.

Proposition 12.1 (Recovering the ordinary algebra)**.** *The full subcategory of* 0*-truncated objects of* $\mathcal{C}_{\mathbb{T}}$ *is equivalent to the nerve of the ordinary category* $\mathrm{Fun}^{\times}(\mathbb{T}, \mathbf{Set})$.

Proof. An object M is 0-truncated when $\mathrm{Map}(P, M)$ is discrete for every P. If all $M(t^n)$ are discrete, the space of natural transformations into M is a limit of discrete mapping spaces and is therefore discrete. Thus such M is 0-truncated.

Conversely, the covariant representable $F = \mathrm{Hom}_{\mathbb{T}}(t, -)$ preserves products, so belongs to $\mathcal{C}_{\mathbb{T}}$. Yoneda identifies $\mathrm{Map}(F, M)$ with $M(t)$: a natural transformation is determined by its value on id_t, and the value on an arrow $a : t \to t^n$ is then $M(a)$ applied to that value. This is the space-valued Yoneda lemma, including its coherent mapping-space statement. If M is 0-truncated, $M(t)$ is discrete; product preservation gives $M(t^n) \simeq M(t)^n$, so all its values are discrete. Functors between ordinary discrete values have ordinary natural transformations and the ordinary product-preservation condition. This identifies both objects and maps of the displayed subcategory. □

Theorem 12.2 (Framework theorem, quoted: rigidification)**.** *For a small ordinary finite-product theory* $\mathbb{T}$*, the homotopy theory of strict simplicial* $\mathbb{T}$*-models, with objectwise weak equivalences, presents* $\mathcal{C}_{\mathbb{T}}$.

This is Lurie's Corollary 5.5.9.3, using Proposition 5.5.9.2, in [2]; see also [4]. This theorem is what permits us to compute with familiar simplicial algebra while interpreting the answer in (33). It is not a claim that all coherences can simply be omitted.

For group theory, the result is the homotopy theory of simplicial groups. Two further foundational inputs describe it. The Kan loop-group and classifying-space adjunction identifies this homotopy theory with pointed connected spaces [5, Chapter V, Proposition 6.3 and Corollary 6.4]. The loop–delooping recognition theorem identifies pointed connected spaces with grouplike E_1-spaces [3, Theorem 5.2.6.10]. An E_1 multiplication is associative with its coherent homotopies; "grouplike" says that its components form a group. The two quoted equivalences together supply the asserted identification; rigidification alone does not supply either comparison.

For module theory, the Dold–Kan theorem identifies simplicial R-modules with nonnegative chain complexes of R-modules; weak equivalences become quasi-isomorphisms [7, Tag 019G]. Localizing therefore gives connective derived R-modules. For a commutative ring R, this agrees with connective modules over its Eilenberg–Mac Lane ring spectrum HR [3, Theorem 7.1.2.13]. The HR-module structure is essential: an arbitrary spectrum, even a connective one, is not thereby an R-module.

The earlier theorem applies to the ordinary models recovered by Proposition 12.1, whenever their category satisfies its pointed exact Mal'cev hypotheses. Groups and modules are therefore instances inside a common homotopical construction. The exactness hypotheses do real work; product-preserving models in spaces alone do not imply a composition theorem for every higher object.

12.2 The action groupoid and the role of associativity

Let an ordinary group G act on the right on a set X. Form a groupoid whose objects are elements $x \in X$, and whose arrows are

$$(x, g) : x \longrightarrow xg.$$

Composition is

$$(x, g) \text{ followed by } (xg, h) = (x, gh).$$

It is defined because $(xg)h = x(gh)$; it is associative because multiplication in G is associative; identities and inverses follow from the corresponding group laws. This construction retains the actual witness g of an identification, not only the fact that two points lie in the same orbit.

Its nerve is the simplicial set

$$X_n = X \times G^n. \tag{34}$$

In degree n, the tuple $(x, g_1, \ldots, g_n)$ means the composable chain

$$x \longrightarrow xg_1 \longrightarrow xg_1g_2 \longrightarrow \cdots \longrightarrow xg_1 \cdots g_n.$$

The first face removes x by replacing it with xg_1; an interior face replaces g_i, g_{i+1} by their product; the last face removes g_n. A degeneracy inserts an identity element. For three consecutive arrows, the equality of the two ways to take interior faces is precisely $(g_1g_2)g_3 = g_1(g_2g_3)$. The unit identities similarly supply the face–degeneracy identities. The nerve is thus a way to organize, in one object, the algebraic compatibility behind the elementary maneuvers.

Definition 12.3. The homotopy quotient $X//G$ is the space obtained by realizing the nerve (34). For a point with the trivial action, write $BG = *//G$.

Proposition 12.4 (What an action quotient retains)**.** *The connected components of $X//G$ are the orbits of X. The component of x is equivalent to BG_x, where $G_x = \{g \in G : xg = x\}$. In particular a free action has homotopy quotient equivalent to its discrete orbit set.*

Proof. There is an arrow from x to y exactly when $y = xg$ for some g. Thus components are orbits. In a fixed orbit, the full subgroupoid on one representative x has automorphism group G_x. Its inclusion is fully faithful by construction and essentially surjective because every object in the orbit is isomorphic to x. It is therefore an equivalence of groupoids. Realization preserves such equivalences, giving BG_x. For a free action all stabilizers are trivial, so each component is contractible. □

For the trivial action on a point, the ordinary orbit set is a point, whereas the homotopy quotient is BG and retains every automorphism in G. For a subgroup $N \leq G$ acting on G by right translation, the action is free, so $G//N \simeq G/N$ as spaces. If N is normal, multiplication also descends to these cosets. The two assertions are different: freeness determines the homotopy type of the quotient; normality determines its induced group multiplication.

12.3 Extensions become fiber diagrams

An ordinary pullback asks for equal images. A homotopy pullback $X \times^h_Z Y$ retains a path between the images: a point is a pair (x, y) together with a path between $f(x)$ and $g(y)$. Paths between

such data must respect the specified comparison, with the corresponding higher coherences. In particular, a homotopy fiber over z remembers both a point x and a path between $f(x)$ and z. The next calculation is a case where this extra datum is a group element.

Proposition 12.5. *A short exact sequence of ordinary groups*

$$1 \longrightarrow N \longrightarrow G \xrightarrow{p} Q \longrightarrow 1$$

determines a homotopy pullback square

$$\begin{array}{ccc} BN & \longrightarrow & BG \\ \downarrow & & \downarrow{\scriptstyle Bp} \\ * & \longrightarrow & BQ. \end{array}$$

Proof. The homotopy fiber of the functor between one-object groupoids induced by p can be modeled by the groupoid with objects $q \in Q$ and arrows $g : q \to qp(g)$. Indeed its objects are an object of BG together with an isomorphism from the chosen object of BQ to its image, and its arrows are the compatible isomorphisms. This is the groupoid model of the homotopy pullback, retaining the isomorphism in the pullback datum. Surjectivity of p makes it transitive, and the automorphisms of $1 \in Q$ are exactly $\ker p = N$. The argument of Proposition 12.4 identifies this fiber with BN. The inclusion of its representative object gives the usual map $BN \to BG$. □

This square is a fiber description, not a claim about an ordinary set-theoretic fiber of the map of realizations. It also is not automatically a cofiber square. The distinctions matter because the homotopy theory of spaces is unstable.

12.4 Reference: universal change of base

An ∞-topos is a setting for spaces varying over a site; formally it is an accessible left exact localization of a presheaf ∞-category of spaces. The theory-model category $\mathcal{C}_{\mathbb{T}}$ need not be an ∞-topos. The effective-quotient results below apply in spaces or another specified ∞-topos; the subsequent descent equivalence is stated for spaces. We need the following two consequences of the foundational theory of an ∞-topos.

Theorem 12.6 (Framework theorem, quoted: effective quotients)**.** *In an ∞-topos, colimits are preserved by pullback. Every groupoid object is the Čech nerve of the map from its object of objects to its realization.*

These are the universality and effectivity assertions in [2, Theorem 6.1.0.6]. A Čech nerve of $U \to Y$ has degree n term $U \times_Y \cdots \times_Y U$ with $n+1$ factors. Effectivity says that the identification data in a groupoid object are precisely the data recovered by pulling back its quotient map against itself.

Proposition 12.7 (A quotient can be computed after base change)**.** *In an ∞-topos, let $U_\bullet$ be a groupoid object with realization Y, and let $f : Y' \to Y$. Then*

$$\left|Y' \times_Y U_\bullet\right| \simeq Y'.$$

Moreover $Y' \times_Y U_\bullet$ is the Čech nerve of $Y' \times_Y U_0 \to Y'$.

Proof. Pullback commutes with the realization colimit by Theorem 12.6, so its realization is $Y' \times_Y |U_\bullet| \simeq Y' \times_Y Y \simeq Y'$. Effectivity identifies U_n with the $(n+1)$-fold fiber product of U_0 over Y. Pulling back this description gives the corresponding $(n+1)$-fold fiber product of $Y' \times_Y U_0$ over Y', including its face and degeneracy maps. These identifications prove the nerve assertion. □

The usefulness is concrete. A quotient description can be transported to a new parameter space without recomputing its coherence relations. If we make two successive base changes, the two answers agree by the canonical associativity equivalence of fiber products. That equivalence is itself compatible with a third base change: both descriptions are limits of the same diagram. Universal properties account for the compatibility that would otherwise have to be checked separately in each example.

12.5 Reference: families over BG

Theorem 12.8 (Framework theorem, quoted: straightening for spaces)**.** *For a space B, taking homotopy fibers with transport gives an equivalence*

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

This is the groupoid case of straightening and unstraightening [2, §§2.2,3.2]. The symbol B on the right is regarded as an ∞-groupoid. Transport along an invertible path is an equivalence; transport along composable paths includes coherent composition data.

For $B = BG$, a functor $BG \to \mathcal{S}$ is a space with coherent G-action. A map $E \to BG$ determines the fiber X over the basepoint, and loops in BG transport X around the base. Conversely, the Grothendieck construction assembles the action into its total space $X/\!/G \to BG$. These are the actual forward and inverse constructions in this instance of the theorem. In identifying this functor notation with the usual left-action convention, we use the canonically equivalent opposite one-object model via inversion: the earlier right action becomes $g \cdot x = xg^{-1}$. Our concrete groupoid calculations continue to use the right-multiplication convention, with first g, then h, composed as gh. No change of coefficient structure is involved.

Corollary 12.9 (Equivariant information is descent information)**.** *The equivalence*

$$\mathcal{S}_{/BG} \simeq \mathrm{Fun}(BG, \mathcal{S})$$

identifies maps over BG with coherent equivariant maps. Its restriction along $BH \to BG$ for a subgroup H identifies homotopy pullback of families with restriction of their actions. Sections of $X/\!/G \to BG$ form the space of coherent fixed points of X.

Proof. The first assertion is the assertion about mapping spaces in the equivalence. Taking the fiber of a pulled-back family at a point of BH gives the same fiber as before, and its path transport is the original transport along the image path in BG. This is precisely restriction along the functor $BH \to BG$. Finally a section is a map over BG from the identity family $BG \to BG$. That family corresponds to the constant one-point functor. Maps from it to the action functor are compatible choices of a point, with the homotopies and coherences making it fixed. □

For a discrete G-set these are ordinary fixed points. For a general space the coherences are part of the answer. We have gained a theory of families, transport, and descent, not merely another notation for a set of orbits.

12.6 Calculate the transporter graph

For a finite group G and subgroups A, B, make a labelled graph with vertex set G. Choose generating sets for A and B. From a vertex g, allow the moves $g \mapsto a^{-1}g$ and $g \mapsto gb$ for a generator a or b, including inverses. A path records two words; its endpoint has the form $a^{-1}gb$. Thus every reachable vertex belongs to AgB, and every element of AgB is reached by such a path. Connected components are exactly double cosets. A traversal can store one parent move for each newly reached vertex, giving a certificate of its membership and an explicit word comparison.

We identify words that give the same pair of group elements; we are not computing the free path group of the labelled graph. At a chosen representative g, a closed comparison consists of pairs (a, b) with $a^{-1}gb = g$. Solving the equation gives $a = gbg^{-1}$, so these loops are parameterized by $A \cap gBg^{-1}$. Enumerating $a \in A$ and testing whether $g^{-1}ag \in B$ calculates this group with its multiplication inherited from A. After a component traversal has supplied its representative, both the list of vertices and the group of closed moves are known. One can check the count

$$|AgB| = \frac{|A|\,|B|}{|A \cap gBg^{-1}|}$$

directly by counting pairs of group elements. Every fiber of $(a, b) \mapsto a^{-1}gb$ has size $|A \cap gBg^{-1}|$.

For example take $G = S_3$ and $A = B = \langle(12)\rangle$. The component of 1 has the two vertices of A and a two-element loop group. For $g = (13)$, $gAg^{-1} = \langle(23)\rangle$, so the intersection is trivial and the component has four vertices. These two traversals exhaust the six-element input. A single ordinary intersection would have missed the second component entirely.

With an explicit multiplication table and membership tables for the subgroups, these are finite graph and lookup operations. With another encoding of a group, the cost of multiplication, enumeration and membership must also be counted. The construction does not give a polynomial-time algorithm for every succinct group representation.

The following theorem identifies the same output in homotopy language. The graph calculation provides the components and their loop groups; the language does not replace their computation.

Here is a useful calculation that prevents a misleading translation of the earlier subgroup diagrams.

Proposition 12.10 (Double cosets in a homotopy pullback)**.** *For subgroups $A, B \leq G$, there is an equivalence of spaces*

$$BA \times^{h}_{BG} BB \simeq \coprod_{[g] \in A\backslash G/B} B(A \cap gBg^{-1}).$$

The component selected by the identity element is $B(A \cap B)$.

Proof. A groupoid model for the homotopy pullback has one object for each $g \in G$. An arrow from g to g' is a pair $(a, b) \in A \times B$ such that $g' = a^{-1}gb$. First (a, b) and then (a', b') compose to (aa', bb'), since $a'^{-1}(a^{-1}gb)b' = (aa')^{-1}g(bb')$. This is the weak pullback: g records the comparison isomorphism in BG between the images of the two chosen objects.

Its orbits are the double cosets AgB. The stabilizer of g consists of pairs with $a = gbg^{-1}$ and is identified with $A \cap gBg^{-1}$ by projection to a. Apply the transitive-groupoid argument of Proposition 12.4 on each orbit. For $g = 1$ the stabilizer is $A \cap B$, as claimed. □

The classical intersection is present, but it is one component of a richer comparison object. Discarding the other components before doing base change would discard genuine information. This is one reason to keep pointings and transport explicit when moving a subgroup argument into homotopy theory.

In particular, a proper subgroup inclusion $H \hookrightarrow G$ does not give a monomorphism $BH \to BG$ in $\mathcal{S}$. Its homotopy fiber is the coset space, which has more than one point. A monomorphism of spaces has every fiber either empty or contractible, so this map is not one. The subgroup inclusion is monic in the category of discrete groups; its delooping has a different categorical type.

12.7 Reference: higher normality and quotient data

For a higher group map $H \to G$, a normality datum can be specified by a pointed connected space Q and a homotopy fiber square in the ambient ∞-category of spaces

$$\begin{array}{ccc} BH & \longrightarrow & BG \\ \downarrow & & \downarrow \\ * & \longrightarrow & Q. \end{array}$$

All displayed maps preserve the specified points. Requiring a pullback in spaces retains the entire homotopy fiber; taking a limit only inside the subcategory of pointed connected spaces could discard other components. The loop space ΩQ then plays the role of the quotient higher group. For an ordinary normal inclusion, $Q = B(G/H)$ supplies this datum. For general higher maps, the datum is additional structure and need not be unique.

The appropriate isomorphism laws consequently include this structure. Beardsley and Fox develop normality data for group objects in ∞-topoi. Their current formulation proves the first and second Noether theorems, exhibits failures of an unrestricted third theorem, and proves an iterated-quotient theorem with compatible data [6, §5, especially Theorem 5.14]. We import no general higher Jordan–Hölder theorem from that work.

One can already see why the ordinary definition of a filtration cannot be transferred without a decision about admissible arrows.

Proposition 12.11 (Literal subobject chains force discreteness)**.** *Suppose a group object in spaces admits a finite chain of categorical monomorphisms beginning at the trivial group and ending at G. Then the underlying space of G is discrete up to equivalence.*

Proof. Limits of product-preserving theory models are computed pointwise, because limits commute with products; evaluation on the generator detects equivalences because its powers give every object of the theory. Thus limits of group objects are computed on their underlying spaces, and equivalences are detected there. A map is monic precisely when its diagonal is an equivalence. Hence the forgetful functor detects these monomorphisms. Their composite gives a monomorphism $* \to G$ of spaces. Its fiber over the identity is ΩG, so that loop space is contractible. The identity component of G is therefore contractible. Translation by a point of a group object identifies its component with the identity component. All components are contractible, which is exactly the assertion that G is 0-truncated. □

Thus a theory of genuinely higher composition objects must specify more than the old finite chain of monomorphisms. Our common ordinary theorem is complete in its stated exact framework. The higher enlargement gives coherent quotients, families, and descent; any additional composition theorem must state and verify the admissible quotient and refinement conditions it uses. This is a precise boundary of the proved result, not a reason to discard the enlargement.

13 Linear operations over a retained base

The working problem in this section is to classify compatible couplings, and to produce a change of coordinates when two couplings agree. We first build a finite matrix calculus that does both jobs. The subsequent presentation by generators and relations reconstructs an attachment over an arbitrary ring. Stable categories and cotangent objects are retained as a reference layer: they organise wider comparisons, but they do not perform either of these calculations for us.

A calculus for two commuting operations. Let $N = k^m$ and $L = k^n$ over a field k. Suppose A_1, A_2 commute on N, and B_1, B_2 commute on L. We want two commuting operations on $N \oplus L$ that restrict to the given operations on N and induce the given operations on L. Their unknown parts are two rectangular matrices:

$$T_i = \begin{pmatrix} A_i & C_i \\ 0 & B_i \end{pmatrix}, \qquad i = 1, 2.$$

Multiplying the blocks, rather than naming a classification theory, gives the whole compatibility condition:

$$d_1(C_1, C_2) := A_1C_2 + C_1B_2 - A_2C_1 - C_2B_1 = 0.$$

The diagonal blocks already commute. Only this off-diagonal block can fail.

Change the lift of a vector of L by $S : L \to N$. The triangular coordinate change and its inverse are

$$P_S = \begin{pmatrix} 1 & S \\ 0 & 1 \end{pmatrix}, \qquad P_S^{-1} = P_{-S}.$$

Direct multiplication of $P_S^{-1}T_iP_S$ changes the coupling to

$$C_i' = C_i + A_iS - SB_i.$$

Put $d_0(S) = (A_1S - SB_1, A_2S - SB_2)$. Expanding once more yields

$$d_1d_0(S) = (A_1A_2 - A_2A_1)S - S(B_1B_2 - B_2B_1) = 0.$$

Thus the compatibility test is unchanged by every allowed coordinate move. We have derived a three-term linear calculation:

$$\mathrm{Hom}_k(L, N) \xrightarrow{d_0} \mathrm{Hom}_k(L, N)^2 \xrightarrow{d_1} \mathrm{Hom}_k(L, N).$$

It is a complex because the two actual operations commute, not because we decided in advance to encode the problem in cohomology.

The output is completely explicit. Compatible couplings are $\ker d_1$. Two are equivalent with both endpoint coordinates fixed exactly when their difference belongs to $\mathrm{im}\, d_0$. The automorphisms of a

fixed coupling are $\ker d_0$. Addition of triangular moves is $P_S P_R = P_{S+R}$, so these tests compose without another compatibility law. The quotient $\ker d_1 / \operatorname{im} d_0$ is the parameter space of marked couplings; it is not a classification after arbitrary changes of basis on N and L.

To compute, use the matrix units as bases of the three displayed spaces. Apply d_0 and d_1 to each unit, record their entries as columns, and row-reduce. A basis of $\ker d_1$ extended from a basis of $\operatorname{im} d_0$ supplies parameters for all classes. To compare two supplied couplings, solve $d_0(S) = C' - C$. A solution is the requested coordinate-change certificate; an inconsistent row is a certificate that the prescribed markings cannot be matched. This is a finite algorithm over k, with the usual additional bit-cost questions if field arithmetic itself has a cost.

For scalar endpoint operations put $\Delta_i = A_i - B_i$. Then

$$\Delta_1 C_2 - \Delta_2 C_1 = 0, \qquad (C_1, C_2) \longmapsto (C_1, C_2) + S(\Delta_1, \Delta_2).$$

If $(\Delta_1, \Delta_2) \neq (0, 0)$, the solution space is the line spanned by (Δ_1, Δ_2), and every coupling can be removed. If both differences vanish, every pair is compatible and no allowed move changes it. The two-parameter attachment space appears precisely at joint resonance. Merely knowing each diagonal operation separately would conceal the joint compatibility equation.

The same multiplication handles any finite list of commuting operations: use one C_i for each operation and one equation $d_1(C)_{ij} = 0$ for each pair. If the operators satisfy further polynomial relations, expand those relations in the block matrices as well. Pairwise commutativity alone does not impose an additional relation such as $T_1 T_2 = 0$. The required relations are part of the input.

The following older comparisons explain where this calculus sits in broader theories. They are secondary to the derived moves and tests.

13.1 Why a full additive replacement cannot work

Proposition 13.1. *There is no fully faithful functor from the category of groups to an additive category.*

Proof. Suppose $F : \mathbf{Grp} \to \mathcal{A}$ were such a functor. The trivial group has just one endomorphism. Thus $F1$ has one endomorphism, so its identity equals its zero endomorphism. It follows that $F1$ is a zero object: every map into or out of it is zero. Consequently F carries zero homomorphisms to zero morphisms, since these factor through 1.

Let $p : S_3 \to C_2$ be the sign quotient. There are four homomorphisms $C_2 \to S_3$: the trivial one and the three choices of a transposition as image of the generator. There are two homomorphisms $C_2 \to C_2$. Postcomposition with p sends only the trivial map to zero, and sends the three remaining maps to the identity. In $\mathcal{A}$, postcomposition with Fp is a homomorphism of abelian groups

$$\operatorname{Hom}_{\mathcal{A}}(FC_2, FS_3) \longrightarrow \operatorname{Hom}_{\mathcal{A}}(FC_2, FC_2).$$

Full faithfulness makes its kernel trivial, so it must be injective. But it has the four-element and two-element sets just described as its domain and codomain. This is a contradiction. □

The obstruction concerns a *full* replacement preserving all maps. It does not prohibit a linear invariant, nor an encoding with additional nonlinear structure. It explains why we retain the original base when constructing its linear theory. The homotopy category of a stable ∞-category is additive, so a fully faithful embedding of the nerve of all groups into a stable ∞-category would give the same contradiction after passage to homotopy categories.

13.2 Modules already arise as abelian objects over a group

Fix a group G. An object of $\mathbf{Grp}/G$ is a homomorphism $p : E \to G$. An abelian group object in this slice has, in addition, a zero section $s : G \to E$, fiberwise addition

$$\mu : E \times_G E \longrightarrow E,$$

and fiberwise inverse, satisfying the abelian group axioms. All these maps must themselves be group homomorphisms over G. The requirement relates the old group multiplication to the new fiberwise addition.

Theorem 13.2 (Modules over a group as linear objects)**.** *There is an equivalence of ordinary categories*

$$\mathrm{Ab}(\mathbf{Grp}/G) \simeq \mathbb{Z}[G]\text{-}\mathbf{Mod}.$$

Proof. Start with an abelian group object $p : E \to G$ and let $M = \ker p$. The fiberwise addition restricts to a second operation on M with the same identity as the inherited group multiplication, since s is a homomorphism and $s(1) = 1$. The fact that μ is a homomorphism gives the interchange law

$$\mu(aa', bb') = \mu(a, b)\mu(a', b') \qquad (a, a', b, b' \in M).$$

It follows that

$$\mu(a, b) = \mu(a, 1)\mu(1, b) = ab, \qquad \mu(a, b) = \mu(1, b)\mu(a, 1) = ba.$$

Thus M is abelian, and its two operations agree.

The section acts on M by conjugation:

$$g \cdot m = s(g)ms(g)^{-1}.$$

Normality of the kernel keeps this in M. Conjugation is an automorphism of its abelian group, and $s(gh) = s(g)s(h)$ gives $(gh)\cdot m = g\cdot(h\cdot m)$. Extending the action additively gives a left $\mathbb{Z}[G]$-module. The map

$$M \rtimes G \longrightarrow E, \qquad (m, g) \longmapsto ms(g)$$

is an isomorphism: its inverse sends e to $(es(p(e))^{-1}, p(e))$. The multiplication formula is

$$(m, g)(n, h) = (m + g \cdot n, gh). \tag{35}$$

Because μ is a homomorphism, its values on $(ms(g), ns(g))$ are obtained by multiplying its value on (m, n) with its value on $(s(g), s(g))$. The former is $m + n$ and the latter is $s(g)$, the zero in the fiber. Thus the fiberwise addition becomes $((m, g), (n, g)) \mapsto (m + n, g)$.

Conversely, let M be a left $\mathbb{Z}[G]$-module. Formula (35) is associative because the action is by automorphisms and respects multiplication in G. Its identity is $(0, 1)$ and the inverse of (m, g) is $(-g^{-1} \cdot m, g^{-1})$. Projection to G has the zero section $g \mapsto (0, g)$. Fiberwise addition is the formula above and fiberwise inverse is $(m, g) \mapsto (-m, g)$. To check addition is a group homomorphism, compare

$$(m + g \cdot m') + (n + g \cdot n') = (m + n) + g \cdot (m' + n').$$

This uses commutativity in M and additivity of the action. Fiberwise inverse is a homomorphism by the same additivity, and all abelian group axioms hold fiber by fiber.

A morphism of abelian group objects preserves p, s, and μ. Its restriction to M is additive and commutes with conjugation by $s(g)$, hence is $\mathbb{Z}[G]$-linear. Conversely a module map $u : M \to M'$ gives $(m, g) \mapsto (u(m), g)$, which preserves every structure map. These assignments are inverse on objects up to the displayed isomorphism and inverse on morphisms. This proves the equivalence. □

No representation has been chosen externally. The action appeared because a coefficient object over G must be transported along G. The theorem derives modules as internal linear objects over a nonlinear base. It does not identify a simple $\mathbb{Z}[G]$-module with a simple quotient group of G; these are different questions about the same base.

13.3 Reference: stabilization over a retained base

A spectrum object can be described by pointed objects $E_0, E_1, \dots$ and equivalences $E_n \simeq \Omega E_{n+1}$. The loop object is formed by a homotopy pullback at the distinguished point. Stabilization $\mathrm{Sp}(\mathcal{C})$ collects such objects in a suitable ∞-category $\mathcal{C}$ with finite limits. We use the standard theory in [3, §§1.4.2–1.4.3], including the standard t-structure on spectra and its heart **Ab**. The heart consists of the Eilenberg–Mac Lane spectra, having only degree-zero homotopy.

The terminology records a degree-by-degree separation. A spectrum is connective when $\pi_i = 0$ for $i < 0$, and coconnective when $\pi_i = 0$ for $i > 0$. The standard t-structure supplies compatible truncations to these ranges. Its heart is their intersection: only π_0 can remain, and it is an abelian group. Thus passing to the heart will recover ordinary linear coefficients from coefficients with higher degrees.

In the slice $\mathcal{S}_{/BG}$, a pointed object means a family $E \to BG$ with a section $BG \to E$; that section chooses the basepoint in each fiber. Looping is performed in each fiber. A spectrum object is consequently a family with compatible repeated deloopings, so the base and its transport remain part of the construction.

Theorem 13.3 (Stable coefficients on a groupoid)**.** *For an ordinary group G there is a natural equivalence*

$$\mathrm{Sp}(\mathcal{S}_{/BG}) \simeq \mathrm{Fun}(BG, \mathrm{Sp}).$$

*Its pointwise standard t-structure has heart $\mathbb{Z}[G]$-**Mod**.*

Proof. Straightening, Theorem 12.8, first identifies $\mathcal{S}_{/BG}$ with $\mathrm{Fun}(BG, \mathcal{S})$. Pointed objects and finite limits in a functor category are computed pointwise. Therefore a sequence of pointed functors with equivalences $E_n \simeq \Omega E_{n+1}$ is exactly a functor from BG to sequences of pointed spaces satisfying those equivalences. This identifies the objects, coherent maps, and their mapping spaces with $\mathrm{Fun}(BG, \mathrm{Sp})$.

Pointwise truncation gives the stated t-structure. An object in its heart assigns an Eilenberg–Mac Lane spectrum to the unique object of BG, hence an abelian group M. Its transport maps give a G-action by automorphisms of M. Mapping spaces between heart objects are discrete, so the coherent conditions reduce exactly to the ordinary action identities. Maps are additive maps commuting with the action. These are left $\mathbb{Z}[G]$-modules and their linear maps. □

This theorem adds higher coefficient groups and coherent transport while keeping BG in the construction. Its restriction to degree zero has the same module category as Theorem 13.2. The

two descriptions explain the relationship between ordinary internal linear objects and stable local coefficient systems.

The ambient homotopical category matters. We have stabilized $\mathcal{S}_{/BG}$, not the ordinary nerve of **Grp**. Indeed, in the ordinary pointed category of groups the pullback $1 \times_G 1$ is 1 for every G. Thus its ordinary loop construction is trivial, and spectrum objects in that unenhanced nerve are trivial. The homotopy fiber $* \times^h_{BG} *$ instead retains G. Homotopical enhancement must precede this stabilization.

13.4 Construct an attachment from relations

For this subsection R is an associative unital ring and modules are left modules. We use cochain complexes, whose differentials raise degree. The shift is $(C[1])^j = C^{j+1}$ with negated differential, so a module $N[1]$ has its only nonzero cohomology in degree -1.

We use the standard derived framework: localizing complexes at quasi-isomorphisms gives a stable ∞-category $\mathcal{D}(R)$; its heart is R-**Mod**; a fiber/cofiber sequence gives the long exact cohomology sequence; and a bounded projective resolution computes derived maps and derived tensor products. These are foundational inputs, not consequences of composition-factor uniqueness. References are [3, §1.3] and [7, Tags 06XP, 0643]. Write $[X, Y]$ for $\pi_0 \operatorname{Map}_{\mathcal{D}(R)}(X, Y)$ and set

$$\operatorname{Ext}^1_R(L, N) = [L, N[1]].$$

Theorem 13.4 (Extensions are connecting morphisms)**.** *The isomorphism classes of marked short exact sequences*

$$0 \longrightarrow N \xrightarrow{i} E \xrightarrow{p} L \longrightarrow 0$$

are naturally in bijection with $[L, N[1]]$. Here an isomorphism of marked extensions induces the identity on both N and L.

Proof. Choose a free surjection $\epsilon : P \to L$ and write $K = \ker \epsilon$. For a map $\phi : K \to N$, define the module by explicit relations:

$$E_\phi = (N \oplus P)/\{(\phi(k), -k) : k \in K\}.$$

The maps $n \mapsto [n, 0]$ and $[n, p] \mapsto \epsilon(p)$ give an attachment with the required endpoints. The first is injective: a relation with second coordinate zero has $k = 0$. If $\epsilon(p) = 0$, then $p \in K$ and $[n, p] = [n + \phi(p), 0]$, which proves exactness in the middle. The last map is onto because ϵ is onto.

Conversely, given a marked attachment, lift ϵ to $f : P \to E$ using freeness of P. For $k \in K$, $f(k)$ lies in the copy of N; define $\phi(k)$ by $i\phi(k) = f(k)$. The map

$$[n, p] \longmapsto i(n) + f(p)$$

is well defined, onto, and has zero kernel by the same two-coordinate calculation. It is an isomorphism preserving both markings.

Replacing f by $f + ia$, with $a : P \to N$, changes ϕ to $\phi + a|_K$. Conversely such a change gives the marked isomorphism

$$E_{\phi + a|_K} \longrightarrow E_\phi, \qquad [n, p] \longmapsto [n + a(p), p].$$

Any marked isomorphism gives this change after comparing the chosen lifts. Therefore the classification already constructed is

$$\operatorname{Hom}_R(K,N)/\operatorname{im}(\operatorname{Hom}_R(P,N)\to\operatorname{Hom}_R(K,N)).$$

To compare this output with $[L,N[1]]$, continue $P\to L$ to a free resolution. Its next free module maps onto K. A degree-one cocycle in the resulting Hom complex vanishes on the relations among those generators, hence factors uniquely through K. Its boundaries are exactly the restrictions $a|_K$. The resolution computation of derived maps therefore identifies $[L,N[1]]$ with the displayed quotient (with the chosen cochain sign convention). This proves the theorem while also constructing the module representing every class. No fiber universal property was used to recover its generators or relations. □

Proposition 13.5 (The automorphisms also survive)**.** *The automorphism group of a marked extension as above is naturally* $\operatorname{Hom}_R(L,N)$*. It agrees with the fundamental group of* $\operatorname{Map}(L,N[1])$ *at its class* δ*.*

Proof. For $u:L\to N$, the map 1_E+iup preserves both markings and has inverse 1_E-iup, since $pi=0$. Multiplication of these automorphisms corresponds to addition of u. Conversely, if f preserves the markings, $p(f-1)=0$, so $f-1$ factors uniquely as iv with $v:E\to N$. The equality $(f-1)i=0$ and monicity of i give $vi=0$. The quotient property of p makes $v=up$ for a unique $u:L\to N$.

In a stable category, looping the mapping space at zero identifies $\pi_1\operatorname{Map}(L,N[1])$ with $[L,N]$. Translation in this grouplike mapping space gives the same group at δ. The fully faithful inclusion of the heart identifies $[L,N]$ with $\operatorname{Hom}_R(L,N)$, proving the final assertion. □

The bijection also respects the Baer addition previously constructed. Naturality of the connecting map makes pullback along $L\to L\oplus L$ and pushout along $N\oplus N\to N$ act by precomposition and postcomposition. Applied to $\delta_1\oplus\delta_2$, their composite is $\delta_1+\delta_2$. Thus the derived notation Ext^1 identifies the same abelian group as the earlier extension construction, not only the same underlying set.

The factors N,L are the endpoints, and δ is the missing attachment. Its mapping space also retains the automorphisms of that attachment. This is a precise instance of gaining reconstruction data by retaining maps rather than just a list of objects.

Example 13.6 (The same two factors, different attachments)**.** Let $R=k[t]$ for a field k, and let $k=R/(t)$ as an R-module. The free resolution

$$0\longrightarrow R\xrightarrow{t}R\longrightarrow k\longrightarrow 0$$

is exact: multiplication by t is injective in the polynomial ring, and its image is the kernel of evaluation at 0. Applying $\operatorname{Hom}_R(-,k)$ makes the differential zero, hence $\operatorname{Ext}^1_R(k,k)\cong k$.

Every marked extension can also be described without derived notation. Let e be the specified image of 1 in its kernel and choose a lift f of 1 in its quotient. Then e,f span the underlying vector space: subtract a multiple of f from any element and the result is in ke. They are independent by applying the quotient map first, then using injectivity on the kernel. Necessarily

$$te=0,\qquad tf=\lambda e$$

for a unique $\lambda \in k$. Replacing f by $f + ce$ does not change λ. Conversely these formulas define an R-module extension for every λ. A map preserving the markings preserves this scalar, and the formulas give an isomorphism whenever the scalars agree. This is the concrete parameter detected by the derived calculation.

At $\lambda = 0$ the module is $k \oplus k$. At $\lambda \neq 0$ it is isomorphic, after forgetting the markings, to $R/(t^2)$: send 1 to f and t to λe. The map is surjective, has kernel (t^2), and the displayed independent vectors prove injectivity of the induced map. Both modules have two factors isomorphic to k, but their t-operators are different. The marked classification distinguishes all values of λ, while the middle-object classification merges all nonzero values. These classifications must not be conflated.

13.5 A derived tensor product retains an equation

For commutative R, the derived tensor product is denoted $M \otimes^{\mathbf{L}}_R N$. A projective resolution replaces one input before tensoring. This repairs the failure of ordinary tensoring to preserve exactness by retaining the resulting homology.

Proposition 13.7. *For $R = k[t]$ and $k = R/(t)$,*

$$H^j(k \otimes^{\mathbf{L}}_R k) = \begin{cases} k, & j = 0, -1, \\ 0, & \text{otherwise.} \end{cases}$$

The ordinary tensor product is only k.

Proof. Place the two free terms of the resolution in Example 13.6 in degrees $-1, 0$. Tensoring with k gives the complex

$$[\, k \xrightarrow{0} k \,]$$

in those degrees, because t acts as zero on k. Its cohomology is exactly the displayed answer. Ordinary tensoring gives $R/(t) \otimes_R R/(t) \cong R/(t)$: the map $a \otimes b \mapsto ab$ has inverse $c \mapsto c \otimes 1$. □

In affine geometry this is the derived self-intersection of the origin inside the affine line. The degree -1 copy of k retains the overlap of the two imposed copies of the equation $t = 0$. The statement about cohomology was proved directly by the resolution; the geometric name does not serve as its proof. The calculation is already beyond what a list of simple factors or an ordinary quotient alone records.

13.6 Reference: cotangent objects and deformation theory

There is a corresponding linear theory for a commutative algebra. For an ordinary A-module M, the split square-zero algebra $A \oplus M$ has multiplication

$$(a, m)(b, n) = (ab, an + bm).$$

Its ideal $0 \oplus M$ has square zero. A section of $A \oplus M \to A$ has the form $a \mapsto (a, d(a))$. Comparing products shows that it is a ring map exactly when $d(ab) = a\,d(b) + b\,d(a)$ and d is additive with $d(1) = 0$. Over a base ring R, being an R-algebra map adds $d(r) = 0$. Thus derivations arise as maps into linear objects over the algebra.

For the higher statement, the notation retains these markings:

$$\mathrm{Der}_A(B, M) := \mathrm{Map}_{(\mathrm{CAlg}_{A/})_{/B}}(B, B \oplus M).$$

The source is the identity map of the A-algebra B and the target is its split square-zero extension, with its projection to B. Thus this is the space of sections over B in A-algebras, including the specified coherences with the base and the projection.

Theorem 13.8 (Framework theorem, quoted: tangent modules and cotangents)**.** *For an E_∞-ring A, $\mathrm{Sp}(\mathrm{CAlg}_{/A}) \simeq \mathrm{Mod}_A$. For a map $A \to B$, derived A-derivations of B into a B-module M are represented by a cotangent module $L_{B/A}$:*

$$\mathrm{Der}_A(B, M) \simeq \mathrm{Map}_{\mathrm{Mod}_B}(L_{B/A}, M).$$

For $A \to B \to C$ there is a cofiber sequence

$$C \otimes_B L_{B/A} \longrightarrow L_{C/A} \longrightarrow L_{C/B}. \tag{36}$$

These are the tangent equivalence, representing property, and transitivity formalism of [3, §7.3; Corollaries 7.3.4.14 and 7.3.3.6]. Tensor products here are tensor products of module spectra, hence derived. In ordinary commutative algebra the corresponding cotangent complexes and transitivity triangle are constructed using simplicial resolutions [7, Tag 08QX].

The category $\mathrm{CAlg}_{/A}$ contains nonlinear algebras mapping to A. Its stabilization produces the linear coefficient category at A, just as spaces over BG produced spectrum-valued local coefficients. The representing property says exactly which measurements the cotangent module encodes: all derived derivations into any coefficient module.

Proposition 13.9 (An explicit obstruction to extending a derivation)**.** *Let $A \to B \to C$ be maps of E_∞-rings and M a C-module. A homotopy class of A-derivations $B \to M$ extends to an A-derivation $C \to M$ if and only if its image under the connecting homomorphism*

$$[C \otimes_B L_{B/A}, M] \longrightarrow [L_{C/B}, M[1]]$$

vanishes. The brackets denote morphisms in the homotopy category of C-modules.

Proof. Apply the mapping-spectrum functor into M to (36). Mapping out of a cofiber turns it into a fiber, giving the exact sequence

$$[L_{C/A}, M] \longrightarrow [C \otimes_B L_{B/A}, M] \longrightarrow [L_{C/B}, M[1]].$$

The first arrow is restriction of derivations: naturality of the representing property identifies the cotangent map with restriction, and extension of scalars identifies the middle term with $[L_{B/A}, M]$ as a B-module mapping group. Exactness says precisely that a class lies in the image of restriction if and only if its connecting class is zero. □

This is a theorem about existence of an extension of a homotopy class. It does not assert uniqueness. The full mapping fiber retains the choices of extension, their paths, and their higher compatibilities. The transitivity triangle organizes all three algebras in one statement; its connecting arrow identifies the obstruction rather than merely asserting that some obstruction ought to exist.

For a final explicit computation, use ordinary commutative algebras over a field k of characteristic zero. This choice also permits comparison with commutative differential graded algebras and E_∞-Hk-algebras [3, Proposition 7.1.4.11]; no equivalence between those models in arbitrary positive characteristic is being assumed. The calculation below is in ordinary derived cotangent theory. Under the rational comparison, its cotangent object corresponds to the spectral one because both corepresent the same derived-derivation functor, with the algebra and coefficient module retained.

Theorem 13.10 (Framework theorem, quoted: a hypersurface cotangent)**.** *If $P = k[x_1, \dots, x_n]$ and $A = P/(f)$ with f a nonzero divisor, the ordinary derived cotangent complex $L_{A/k}$ is represented in degrees $-1, 0$ by*

$$(f)/(f^2) \longrightarrow \Omega_{P/k} \otimes_P A, \qquad [f] \longmapsto df.$$

This regular-quotient computation follows from the cotangent transitivity triangle and the regular-sequence theorem [7, Tags 08SH, 08QX]. The ordinary module of differentials alone is its degree-zero cohomology; the complex retains the relation before taking that quotient.

Example 13.11 (The infinitesimal direction of the dual numbers)**.** Take $A = k[t]/(t^2)$. The polynomial t^2 is a nonzero divisor in $k[t]$, so Theorem 13.10 applies. Both terms are free rank-one A-modules, and the complex is

$$L_{A/k} \simeq \left[\, Ae \xrightarrow{e \mapsto 2t\,dt} A\,dt \,\right].$$

The characteristic assumption makes 2 invertible. Multiplication by $2t$ has image kt and kernel kt in A. Consequently

$$H^0(L_{A/k}) \cong k\,dt, \qquad H^{-1}(L_{A/k}) \cong k\,(te).$$

These assertions can be checked on $a + bt$: multiplication by $2t$ gives $2at$. The negative-degree term therefore survives; replacing the complex by its ordinary differentials loses it.

Let k be an A-module through $t \mapsto 0$. Since the displayed cotangent complex is a bounded complex of free modules, its derived maps into k are computed by the ordinary Hom complex. That complex has a copy of k in degrees $0, 1$ and zero differential. Thus

$$[L_{A/k}, k] \cong k, \qquad [L_{A/k}, k[1]] \cong k.$$

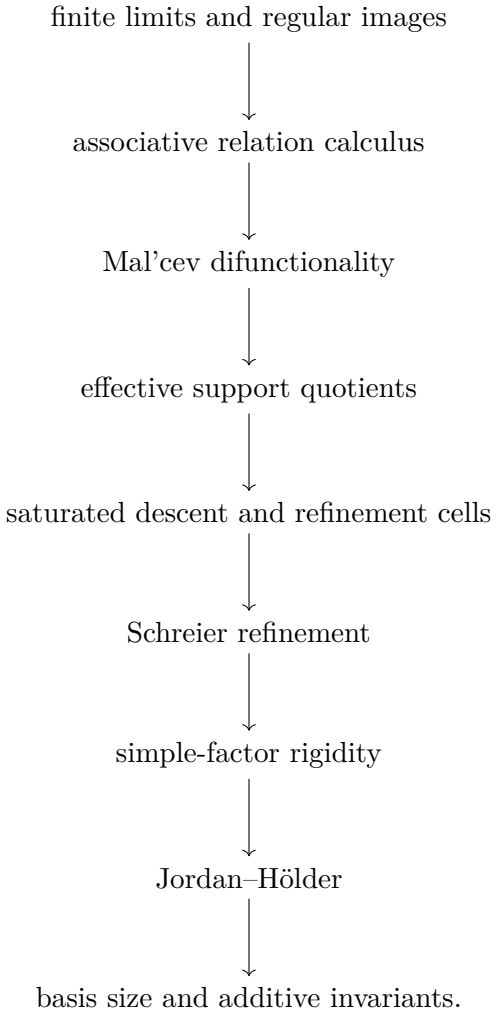
The first is the group of derivations into this coefficient module; the second is the next group available for obstruction calculations such as Proposition 13.9. A nonzero possible obstruction group does not by itself assert that a specified lifting problem has a nonzero obstruction: that requires computing its connecting class.

The progression is now visible. Exact quotients and refinement identify which simple pieces are invariant. Coherent quotients retain actions and transport. Linear objects over a base produce modules. Derived maps retain attachments and automorphisms; derived tensor products and cotangent complexes retain relations and obstruction classes. Each construction answers a new question about information the preceding one did not retain. The ordinary Jordan–Hölder theorem remains one proved consequence of the exact framework, while these higher and derived results enlarge the kinds of reconstruction that can be asked and answered.

A Dependencies, scope, and further reading

A.1 What was proved from what

There are two completed routes to the finite-basis conclusion. The first performs the invertible exchanges in Section 1. The second develops a comparison that also applies beyond vector spaces. Its dependency chain is:



Exactness enters at the effective quotients. Pointedness supplies the zero classes used to concatenate cells into a series. The Mal'cev condition is used where it is needed: to make every relation difunctional. Ordinary compatible iterated quotients require regularity and exactness, but do not require Mal'cev. The written proofs keep these dependencies separate.

The precise ordinary theorem is Theorem 3.10: in a pointed Barr-exact Mal'cev category, any two given finite composition series, retaining their regular quotient maps, have the same factors up to permutation and isomorphism. No global chain condition is needed to compare two finite series already supplied. This is a sufficient common framework, not a claim that these are the logically weakest possible axioms for every formulation of Jordan–Hölder.

The existence of a finite series is a separate issue. The basis flags provide the required series for a vector space with a given finite basis. Later, finite-dimensional module theory uses the now

established dimension function to prove existence of composition series. That later argument does not feed back into the proof of dimension. One can define K_0 by generators and exact-sequence relations before proving Jordan–Hölder. What requires the factor comparison is the identification of K_0 with the free abelian group on the simple isomorphism classes. Using that freeness to prove factor uniqueness would reverse the dependency.

The source of associativity changes as the constructions develop. For the original algebra it is an axiom equating two composites. For relation composition, Baer addition, and Hall multiplication, it is a theorem: both bracketings are obtained from the same complete data of witnesses, tuples, or flags. For coherent quotients, the same universal-property principle controls the higher compatibility of successive constructions. These assertions specify the maps or data that do the work; a resemblance between diagrams alone would not establish them.

A.2 Three different comparison claims

The scope of the page calculations is worth recording alongside the core theorem. The categorical rectangle works for all objects and finite normal series satisfying the exact Mal'cev hypotheses. The explicit nonabelian page construction in Section 6 uses composition-series input in groups with endomorphism operators; that simplicity is used to prove the normality of differential images. The filtered module calculus in Section 7 allows arbitrary finite filtrations and an arbitrary filtration-preserving module map. Its convergence records kernel and cokernel layers, and its staggered identity specialization pairs simple factors.

Thus the nonabelian page calculation and the module spectral calculation are compatible where both apply. Neither is asserted to provide a general spectral sequence for arbitrary filtered groups. The categorical theorem already supplies the unified Jordan–Hölder statement independently of those more specialized calculations.

A.3 The additional hypotheses behind reconstruction

The universal additive record applies to composition-finite objects, with the set-size qualification stated in Corollary 3.12. For finite-length modules the category is also closed under submodules and quotients, as explicitly proved. We did not assume that extra closure for every exact Mal'cev category. Additive records ignore whether an extension splits; the marked extension group retains its endpoint maps.

The probe theorem reconstructs finite-dimensional *nilpotent* $k[t]$-modules. Its second-difference formula is proved using canonical subquotients, and the chain decomposition is constructed before the reconstruction conclusion is drawn. The same tests fail for an invertible operator, as the example shows. Hall multiplication uses finite fields and finite-dimensional nilpotent modules so that every coefficient and every basis expansion involved is finite. Its associativity counts the same flags in two orders. Marked extension classes, isomorphism classes of middle objects, and actual submodules are kept distinct.

A.4 Where the higher foundations enter

The higher part has its own declared inputs. Rigidification identifies product-preserving theory models in spaces with a familiar simplicial presentation. The Kan loop-group equivalence, loop–

delooping recognition, and Dold–Kan explain the group and module instances. Effective groupoid quotients and universal colimits make homotopy quotients compatible with base change. Straightening interprets families over BG as coherent actions. Stabilization and the derived category supply mapping spectra, hearts, and fiber/cofiber sequences. The tangent/cotangent theorems supply linear coefficients for algebras, the representing object for derivations, and its transitivity sequence.

Those results are explicitly imported from the cited foundational literature; this exposition does not reprove higher topos theory or the construction of its model categories. Its deductions include the double-coset homotopy pullback, the coefficient category over BG, the marked extension classification and its automorphisms, and explicit derived tensor and cotangent calculations. The ordinary internal-linear equivalence $\mathrm{Ab}(\mathbf{Grp}/G) \simeq \mathbb{Z}[G]\text{-}\mathbf{Mod}$ is proved directly.

The term "higher" does not imply an unrestricted higher Jordan–Hölder theorem. Normality data and compatibility remain part of the problem, and literal subobject chains from the trivial higher group force discreteness by Proposition 12.11. What the higher framework contributes here is a precise enlargement of the retained data and several new reconstruction tools. Its ordinary discrete models recover the common theorem when its exact Mal'cev hypotheses are checked.

A.5 Audit and reproducibility

The public audit report inventories all theorem-like statements, their proof or imported-foundation status, the diagrams, references, and edition fingerprints. It distinguishes the full review of the preceding edition from the subsequent check of these small clarifications and the matching web, PDF, and source files. The ordinary core and the higher chapters received independent mathematical rereads; integration checks tracked changed action conventions, quotient maps, normality, degrees, and dependencies. Supplementary finite group and finite field calculations test examples of the actual maps and counts. Such computations do not establish the general theorems. Compilation and inspection of the rendered pages are separate presentation checks, not proof certification. No proof assistant formalization or historical priority claim is made.

The earlier edition is preserved separately. This edition reorganizes the core around relations and exactness, so that the group and module comparison theorems are genuine instances. It retains the explicit extension and reconstruction calculations, and places the new higher constructions after the common ordinary proof.

References

[1] Friedrich Knop, *Tensor envelopes of regular categories*, Advances in Mathematics **214** (2007), 571–617. Author preprint; Section 5 and Appendices A–B, especially Appendix B2.

[2] Jacob Lurie, *Higher Topos Theory*, Princeton University Press, 2009. Author PDF. See Corollary 5.5.9.3, Theorem 6.1.0.6, and Sections 2.2 and 3.2.

[3] Jacob Lurie, *Higher Algebra*, September 18, 2017. Author PDF. See Sections 1.3, 1.4.2–1.4.3, 5.2.6 and 7.3; Theorem 7.1.2.13; Proposition 7.1.4.11; Corollaries 7.3.3.6 and 7.3.4.14.

[4] Bernard Badzioch, *Algebraic theories in homotopy theory*, Annals of Mathematics **155** (2002), 895–913. Journal article; Author preprint.

[5] Paul G. Goerss and John F. Jardine, *Simplicial Homotopy Theory*, Progress in Mathematics **174**, Birkhäuser, 1999. Book PDF; Chapter V, Proposition 6.3 and Corollary 6.4.

[6] Jonathan Beardsley and Landon Fox, *Higher Groups and Higher Normality*, arXiv:2407.21210v3, August 13, 2026. In particular Definition 4.1 and Section 5, including the compatibility condition in Theorem 5.14.

[7] The Stacks Project Authors, *The Stacks Project*, online reference. The specific tags used are 010I (extensions), 012K (filtered complexes), 019G (Dold–Kan), 06XP (Ext), 0643 (projective resolutions), 08QX (cotangent transitivity), and 08SH (regular quotients).

[8] Pierre Deligne, *Théorie de Hodge II*, Publications Mathématiques de l'IHÉS **40** (1971), 5–57. Article PDF; Section 1.2.1 discusses interchange of associated gradings.

[9] Marco Grandis, *Homotopy spectral sequences*, arXiv:1007.0632, 2010. Sections 2–4 give a broader nonabelian exact-couple setting.