Fourier, Symmetry, and Reconstruction

Cardano, Ferrari, and categorical Fourier duality

Henry Shin
Refacte

Revised November 3, 2026

Abstract

A useful substitution changes what we have to calculate. We derive root coordinates that respond simply to permutation, then use their products and powers to turn cubic and quartic equations into smaller equations with compatible radical choices. The same operations give a solvable family in every degree and a finite-field inversion algorithm. Further studies change the working objects again: symmetry quotients reduce integrals, modules turn monodromy into multiplication, and a Fourier–Mukai transform converts bundle operations into local algebra and an algorithm on rank and degree. Separate arithmetic studies treat point counts, periods, and cancellation estimates. The last section asks what is needed to compose such transforms and to extend them to higher coefficients. Each construction is judged by the calculation it enables and by the hypotheses that make the return to the original problem valid.

Contents

1 Change what the equation asks us to calculate 3

2 Make the coefficient construction uniform 4

3 Design coordinates with a simple update rule 7

4 Centering: the first substitution is a projection 11

5 Cardano: separate the cyclic motion 12

6 Ferrari: reconstruct a root from three character lines 15

7 Why Ferrari's perfect square appears 18

8 Two quartics worked through 20

9 Cardano in every degree: the dihedral family 21

10 Read the radical choices in the covering tower 24

11 The same Fourier projectors split an integral 26

12 A circle: Fourier transformation beyond eigenvalues 29

13 Fourier–Mukai: reconstructing global bundles 32

14 Another cubic: from Fourier projectors to counting and periods 38

15 A quantitative shadow: where the square-root bound comes from 42

16 Kernel composition and reference theorems 44

1 Change what the equation asks us to calculate

Why does Cardano look for a root in the form $u+v$? Why does Ferrari introduce a cubic in order to solve a quartic? A useful explanation should predict the substitutions, identify the information their auxiliary equations retain, and tell us how to recover the roots without making incompatible choices. Once it does this, it should also help us recognize other problems to which the same method applies.

The first change is to stop calculating individual roots directly. We seek combinations of roots whose response to a permutation is easy to describe. A cyclic permutation should merely multiply a chosen combination by a scalar. That requirement determines its coefficients. Products and powers of the resulting quantities then remove the scalar factors and produce expressions that can be read from the given polynomial. Returning to the roots is a separate, explicit inversion step.

The operations matter as much as the components. If u and v have inverse characters, their product is invariant. If u has a character of order n, its n-th power is invariant. These facts predict which auxiliary quantities can be expressed in the coefficients. Their actual values and their compatibility come from the multiplication of the original algebra. Cardano's quadratic and Ferrari's cubic then arise as equations for the remaining symmetry orbits.

The companion exposition *What a Boundary Map Can Prove* follows an incidence complex through counting laws, geometric inequalities, and extension data.

The main algebraic route has three jobs: construct useful coordinates, calculate the smaller equations they satisfy, and recover all the original roots while respecting the relations between the coordinates. Sections on centering, the cubic, and the quartic complete those jobs. The high-degree example then tests exactly which restriction makes the method continue to work.

The other studies are independent uses of a similar design question. A symmetry quotient replaces an integral by lower-genus integrals. A module turns an invertible matrix into multiplication by a variable. A bundle transform can turn tensor products into truncated-polynomial calculations. These are different operations with different hypotheses; one is not a required ladder to the next. The later geometric and higher-coefficient constructions are developed when their own questions need them.

1.1 What is being reconstructed?

Let K be a field of characteristic zero and introduce independent root coordinates $r_1, \dots, r_n$. The symmetric group acts on $K[r_1, \dots, r_n]$ by permuting them. The invariant ring is generated by the elementary symmetric functions $e_1, \dots, e_n$, and

$$\prod_{i=1}^{n}(X - r_i) = X^n - e_1X^{n-1} + e_2X^{n-2} - \cdots + (-1)^n e_n.$$

Thus coefficient space is the quotient of ordered-root space by the permutation action. The inverse problem is to lift a coefficient point through that quotient.

Working with formal roots does not assume that the equation has already been solved. We construct identities on the universal root space, express their invariant parts in the coefficients, and solve the resulting reconstruction equations. The identities then specialize to particular polynomials, including ones with smaller Galois groups.

Three operations must be distinguished. An invertible coordinate change retains all information. Passing to invariants forgets distinctions between points in the same orbit. A radical choice recovers part of what was forgotten. Centering changes neither the splitting field nor the Galois

group; choosing a resolvent root can select an intermediate field. Later, passing to eigenvalues or to a Grothendieck group will forget different information, namely extensions. The appropriate reconstruction data change with the question.

The universal groups for the cubic and quartic are S_3 and S_4. A particular separable polynomial has a Galois group acting through a subgroup, so some field degrees may drop. Roots of unity are adjoined when needed for character decomposition. Polynomial identities remain valid at repeated roots; the interpretation as unramified covers will be restricted to the distinct-root locus.

1.2 Choose a calculation to complete

For polynomial equations, begin with the coefficient requirement in the projector calculation, then centering, and the cubic and quartic. These proofs can be carried out in universal root coordinates using permutation, multiplication, and orbit sums. The coefficient-space construction explains how the same operations are assembled before a numerical ordering of roots has been chosen; it is not an additional premise for the polynomial identities.

The dihedral family and its finite-field inversion give a second complete algebraic task. For other inputs, choose the operation to be simplified: split an integral, calculate a matrix action and its tensor product, or reduce a bundle problem. The point-count and cancellation studies ask how character calculations can control an entire family of arithmetic quantities. The final section makes the composition and higher-coefficient questions explicit.

Each return map and each theorem used in these studies has its own hypotheses. A successful change of working objects earns a calculation; it does not make unrelated hypotheses disappear.

2 Make the coefficient construction uniform

This section addresses a further question about the polynomial calculations: can the working objects be constructed uniformly from coefficients, rather than after choosing an ordering of numerical roots? The answer uses a finite algebra and its covers. The direct identities later in the text only require universal root coordinates; the construction here explains their behaviour as the coefficients vary. Introducing the coordinate ring of formal roots does not assume that their numerical values are known.

2.1 The equation already defines an object to reconstruct

Let k be a field of characteristic zero, and consider the universal monic polynomial

$$f(T) = T^n + c_1 T^{n-1} + \cdots + c_n.$$

Write

$$\mathcal{C}_n = \operatorname{Spec} k[c_1, \ldots, c_n, \Delta^{-1}], \qquad \mathcal{E} = \mathcal{O}_{\mathcal{C}_n}[T]/(f),$$

where Δ is the discriminant, and denote the class of T by α. The algebra $\mathcal{E}$ is finite étale of rank n. Its relative spectrum $Z = \operatorname{Spec}_{\mathcal{C}_n} \mathcal{E}$ parametrizes a polynomial together with one of its roots. The element α is its root observable: its value at such a point is precisely that chosen root.

There is also a space P of ordered factorizations

$$f(T) = (T - r_1) \cdots (T - r_n).$$

Construct it by adjoining formal variables r_i, imposing the coefficient identities in this equation, and inverting the discriminant. Equivalently, it is the space of ordered distinct n-tuples mapping to their elementary symmetric functions. The map

$$P \longrightarrow \mathcal{C}_n$$

is an S_n-torsor: every geometric fibre consists of the $n!$ orderings, freely and transitively permuted by S_n. A polynomial with a smaller arithmetic Galois group gives a possibly disconnected specialization of this torsor; the construction does not assume that its Galois group is the full symmetric group.

Over $\mathbb{C}$, this is a covering space. Following a loop in coefficient space permutes the roots. An intermediate quotient of P remembers only some of the labels and restricts which permutations remain possible. Algebraically, the same operation is finite étale descent.

2.2 Choose a cover on which Fourier decomposition applies

For three and four roots there is an especially useful normal abelian subgroup:

$$A = C_3 \triangleleft S_3, \qquad A = V_4 \triangleleft S_4.$$

Both act *regularly* on the root labels: once a starting label is chosen, every other label is reached by exactly one element of A. These groups therefore simultaneously give an abelian symmetry and account for every root. This is the reason for choosing them.

Put $Y = P/A$. The remaining map $P \to Y$ is an A-torsor. The map $Y \to \mathcal{C}_n$ has group S_n/A: it remembers an orientation in the cubic case and an ordering of the three pairings in the quartic case. The respective degrees are 2 and 6. Choosing just one quartic pairing is a different, degree-three intermediate cover, used later for Ferrari's shorter calculation.

Because A acts regularly on the root labels, pulling the root cover Z back to Y identifies it with $P \to Y$, after fixing the starting label. Thus we have reached an algebra in which the original root observable itself is subject to an abelian action. Nothing has yet suggested a sum of radicals. That sum will be forced by decomposing this observable.

2.3 A reconstruction theorem for the torsor algebra

Assume that k contains the character values of A; for C_3, adjoin a primitive cube root of unity. Let $\pi : P \to Y$ and $\mathcal{B} = \pi_* \mathcal{O}_P$.

Proposition 2.1 (The cover as a character-graded algebra)**.** *There is a decomposition*

$$\mathcal{B} = \bigoplus_{\chi \in \widehat{A}} \mathcal{L}_\chi, \qquad e_\chi(b) = \frac{1}{|A|} \sum_{a \in A} \chi(a)^{-1} a(b),$$

where every $\mathcal{L}_\chi$ is a line bundle, $\mathcal{L}_1 = \mathcal{O}_Y$, and multiplication induces isomorphisms

$$\mathcal{L}_\chi \otimes_{\mathcal{O}_Y} \mathcal{L}_\psi \xrightarrow{\sim} \mathcal{L}_{\chi\psi}.$$

The line bundles together with these multiplication maps reconstruct P as $\operatorname{Spec}_Y \mathcal{B}$*. Every root observable is reconstructed by summing its character projections.*

Proof. Étale locally on Y, the torsor is the constant set A, and $\mathcal{B}$ is the algebra of functions on A. Character orthogonality decomposes that algebra into the lines spanned by characters. The product of two character functions is their product character, so every displayed multiplication map is an isomorphism. These statements descend to Y, since the projectors and multiplication are intrinsic. The formula $\sum_\chi e_\chi = 1$ reconstructs every section. Finally, the relative spectrum of a finite algebra recovers its finite cover. □

The multiplication is essential: the direct sum of the lines alone does not determine the cover. Its multiplication records the equations for invariant powers and mixed products, including the compatibility of radical choices. This is already a categorical construction with computational content: the symmetric monoidal Fourier equivalence transports a commutative algebra object to a character-graded commutative algebra. The next section proves the underlying projector identities in detail.

2.4 Why the root is a sum of two or three terms

First remove the invariant component of the root observable:

$$r = \alpha - \frac{1}{n}\operatorname{Tr}_{\mathcal{E}/\mathcal{O}_{\mathcal{C}_n}}(\alpha).$$

After pulling back to Y, its trivial-character projection is zero. Indeed, averaging over A averages over all root labels, because the action is regular. Centering is therefore the projector onto the trace-zero part, before it is a substitution in a polynomial.

For the cubic, precisely two character lines remain. Set

$$u = e_\chi(r), \qquad v = e_{\chi^{-1}}(r),$$

where χ has order three. The identity resolution in Proposition 2.1 gives

$$\boxed{r = u + v.}$$

This constructs Cardano's substitution uniformly on coefficient space. The same two quantities can be discovered directly by asking for combinations of the roots that acquire a scalar factor under cycling. Here they arise as projections of the distinguished root observable. Multiplication predicts the quantities available downstairs: u^3, v^3, and uv. The remaining involution exchanges u and v, so the unordered pair $\{u^3, v^3\}$ satisfies a quadratic over coefficient space. The later coefficient calculation determines that quadratic and the mixed-product constraint.

For the quartic, the three nontrivial characters of V_4 remain. Write their projections as a, b, c. Then

$$r = a + b + c.$$

With the conventional normalization $u = 2a$, $v = 2b$, $w = 2c$, this becomes

$$\boxed{r = (u + v + w)/2.}$$

The other roots arise by the three allowed double sign changes. The squares u^2, v^2, w^2 are invariant under V_4; the quotient S_3 permutes them, predicting their cubic resolvent. Since the product of the three nontrivial characters is trivial, uvw is also invariant. It retains the sign compatibility that three independently extracted square roots would forget.

Thus the group action determines the number and form of the summands, multiplication determines the invariant powers and their compatibility, and descent determines the degrees of the resolvents. The coefficients determine their numerical values. The following calculations make those values explicit.

The construction has now answered its uniformity question. The next sections calculate the same projectors and their products directly. Their correctness can be checked by finite sums and polynomial identities. The cover description supplies an additional interpretation of the choices and their monodromy; the explicit calculation supplies their values.

3 Design coordinates with a simple update rule

3.1 Character projectors

First consider quantities $r_0, \dots, r_{n-1}$ that a move σ cycles, with indices read modulo n. We want a combination $u = \sum_i c_i r_i$ whose update is multiplication by one number: $\sigma(u) = \lambda u$. Comparing the coefficient of r_i gives

$$c_{i-1} = \lambda c_i, \qquad c_i = \lambda^{-i} c_0.$$

After a full cycle the coefficient must return, so a nonzero choice requires $\lambda^n = 1$. With $c_0 = 1/n$, the required combination is therefore

$$u = \frac{1}{n} \sum_{i=0}^{n-1} \lambda^{-i} r_i.$$

The Fourier weights have been derived from the operation we wanted. We now prove that these combinations separate all the components, and that the same construction works for any finite abelian group. The assumptions below provide the scalar roots of unity and make division by the group order legitimate.

Let A be a finite abelian group, and let k be a field of characteristic zero containing all roots of unity whose orders divide the exponent of A. Write

$$\widehat{A} = \mathrm{Hom}(A, k^\times)$$

for its character group. These hypotheses ensure that $\widehat{A}$ has $|A|$ elements and that its characters separate the elements of A. If A acts linearly on a vector space V, define

$$V_\chi = \{v \in V : a(v) = \chi(a)v \text{ for every } a \in A\}, \qquad e_\chi(v) = \frac{1}{|A|} \sum_{a \in A} \chi(a)^{-1} a(v).$$

The operator e_χ extracts the part of v with transformation rule χ.

Proposition 3.1. *The operators e_χ are mutually orthogonal projectors with sum the identity. Consequently*

$$V = \bigoplus_{\chi \in \widehat{A}} V_\chi, \qquad v = \sum_{\chi \in \widehat{A}} e_\chi(v).$$

Proof. For a nontrivial character η, choose $b \in A$ with $\eta(b) \neq 1$. Multiplication by b permutes A, so

$$\sum_{a \in A} \eta(a) = \sum_{a \in A} \eta(ba) = \eta(b) \sum_{a \in A} \eta(a).$$

The sum is therefore zero. For the trivial character it is $|A|$. Reindexing the definition of e_χ gives $b(e_\chi(v)) = \chi(b)e_\chi(v)$, so its image lies in V_χ. On V_ψ, character orthogonality says that e_χ is the identity if $\chi = \psi$, and zero otherwise. Finally, applying the same summation argument on $\widehat{A}$ gives

$$\sum_{\chi\in\widehat{A}} \chi(a)^{-1} = |A|\delta_{a,1}.$$

Thus $\sum_\chi e_\chi = \mathrm{id}_V$, proving the claim. □

For $A = C_n = \langle\sigma\rangle$, choose a primitive n-th root of unity ζ, and suppose that σ cyclically permutes quantities $r_0, \ldots, r_{n-1}$. Their Fourier coordinates are

$$u_j = \frac{1}{n}\sum_{i=0}^{n-1} \zeta^{-ij} r_i, \qquad r_i = \sum_{j=0}^{n-1} \zeta^{ij} u_j.$$

Here $\sigma(u_j) = \zeta^j u_j$. The coordinate u_0 is the mean of the roots; centering removes it. The other coordinates isolate the remaining cyclic motions. The familiar sums of radicals will arise by applying this inverse transform after determining the u_j.

3.2 The tensor rule behind Fourier coordinates

The useful rule can be checked before introducing the category notation. If u transforms by α and v by β, applying a group element to their tensor gives

$$a(u \otimes v) = a(u) \otimes a(v) = \alpha(a)\beta(a)(u \otimes v).$$

Thus the characters multiply when the quantities are combined. We can predict the destination of a product by multiplying its labels. This is the operation the following equivalence must preserve.

Character decomposition is compatible with every equivariant linear map. Allowing arbitrary-dimensional vector spaces in both categories, it therefore defines an equivalence

$$\mathrm{Rep}_k(A) \simeq \mathrm{Vect}_k^{\widehat{A}}, \qquad V \longmapsto (V_\chi)_{\chi\in\widehat{A}}.$$

On the right are character-indexed families of vector spaces. The inverse takes their direct sum and lets $a \in A$ act on the χ-summand by $\chi(a)$. For an arbitrary finite-dimensional representation, taking the trace of the projectors recovers the numerical Fourier formula

$$\dim V_\chi = \frac{1}{|A|}\sum_{a\in A} \chi(a)^{-1}\,\mathrm{Tr}(a \mid V).$$

Thus the character transform records the dimensions of spaces that already exist before taking traces.

The tensor product carries the decisive additional information. A tensor $v \otimes w$, with $v \in V_\alpha$ and $w \in W_\beta$, transforms by $\alpha\beta$. Hence

$$(V \otimes_k W)_\chi = \bigoplus_{\alpha\beta=\chi} V_\alpha \otimes_k W_\beta.$$

The equivalence is therefore symmetric monoidal when the family side is equipped with convolution of its character indices. Its unit is k in the trivial-character component and zero elsewhere. Pointwise tensor product of the components would give a different operation and would lose the multiplication rule needed below. This is an elementary categorical form of Fourier theory: the transform retains spaces and maps, together with their tensor relations, rather than only numerical coefficients; see also Baez [4].

3.3 An algebra is more than its representation

Suppose now that R is a commutative k-algebra and that A acts by algebra automorphisms. Its Fourier decomposition satisfies

$$R = \bigoplus_{\chi \in \widehat{A}} R_\chi, \qquad R_\chi R_\psi \subseteq R_{\chi\psi}, \qquad R_1 = R^A.$$

Equivalently, R is a commutative algebra object in $\mathrm{Rep}_k(A)$, transported by the preceding equivalence to a character-graded algebra. Multiplication, its unit, and its associativity remain part of the data.

This distinction has an immediate arithmetic meaning. The extensions $\mathbb{Q}(\sqrt{2})$ and $\mathbb{Q}(\sqrt{3})$, with their order-two Galois actions, have isomorphic underlying representations: each is the sum of the trivial and sign representations. Their multiplication differs. A generator of the sign line squares to 2 in one algebra and to 3 in the other. The representation identifies the line in which a radical belongs; its multiplication specifies which radical extension it defines.

More generally, if u has character χ of order n, then u^n is invariant. If u and v have inverse characters, then uv is invariant. These statements predict the quantities to compute, but they do not prescribe their numerical values. Those values come from the original algebra, or from its coefficient functions. In Cardano's method this distinction becomes the difference between knowing that uv is invariant and determining that it equals a particular multiple of the cubic's linear coefficient.

The coefficient functions supply a second grading: ordinary polynomial degree. The coefficient of X^{n-j} is, up to sign, a homogeneous symmetric polynomial of degree j in the roots. After changing to Fourier coordinates it is still homogeneous of degree j, and every monomial that occurs must have trivial total character. Character and degree together can leave only a few possible monomials. This makes it possible to determine a coefficient by evaluating a small number of convenient configurations, instead of expanding a long product. The resulting numerical constants are normalization data for a structure already identified.

Taking an invariant power forgets a root-of-unity phase. A mixed product can retain the relative phase between two coordinates. Consequently separate equations for the powers do not generally suffice for reconstruction: their solutions must also satisfy the mixed-product equations. These compatibility conditions are the multiplication maps of the original algebra expressed in chosen coordinates.

When A is normal in a larger group G, the quotient G/A acts on R^A. It also reorganizes the character components: for $g \in G$,

$$g(R_\chi) = R_{g\cdot\chi}, \qquad (g \cdot \chi)(a) = \chi(g^{-1}ag).$$

Thus the remaining symmetry acts on the invariant quantities obtained by taking powers and compatible products. Resolvent equations describe this residual action. Their coefficients lie downstairs because they are symmetric expressions invariant under the remaining group.

3.4 From a cyclic Galois action to a radical

The reconstruction mechanism is particularly sharp for fields.

Proposition 3.2. *Let L/F be a cyclic Galois extension of degree n, with generator σ, and assume that F has characteristic zero and contains a primitive n-th root of unity ζ. There is a nonzero $\theta \in L$ such that*

$$\sigma(\theta) = \zeta\theta, \qquad \theta^n = a \in F^\times, \qquad L = F(\theta).$$

Proof. The independence of distinct field automorphisms implies that the operator

$$e = \frac{1}{n} \sum_{j=0}^{n-1} \zeta^{-j} \sigma^j$$

is nonzero. Choose $b \in L$ with $\theta = e(b) \neq 0$. The projector calculation gives $\sigma(\theta) = \zeta\theta$, whence $\theta^n \in F$. The conjugates $\theta, \zeta\theta, \ldots, \zeta^{n-1}\theta$ are distinct. Therefore $[F(\theta) : F] = n$, so $F(\theta) = L$. □

This is the Fourier mechanism in Kummer theory [2, 5]. The same argument shows that, for a finite abelian Galois group A over a coefficient field F containing the character values, every character space L_χ is a line over F. Indeed, its projector is nonzero, and the ratio of two nonzero elements of L_χ is invariant. Multiplication gives isomorphisms

$$L_\chi \otimes_F L_\psi \simeq L_{\chi\psi}.$$

Choosing a basis in each line expresses these maps as scalars. Associativity constrains those scalars and thereby constrains the choices of radicals. Changing bases changes their displayed values without changing the extension.

The proposition is a structural existence proof. Its definition of θ uses conjugates in L, which are not initially known when solving an equation. An explicit solution must still express suitable powers and products of these Fourier coordinates using the given coefficients. The low-dimensional invariant calculations for the cubic and quartic perform exactly this task.

3.5 Solvable towers and the scope of the construction

Let a separable polynomial over a characteristic-zero field K have splitting field L and solvable Galois group G. Adjoin enough roots of unity, for example those of order $|G|$, and put

$$K' = K(\mu_{|G|}), \qquad E = LK'.$$

The preliminary extension K'/K is itself radical: a primitive $|G|$-th root of unity is a $|G|$-th root of 1. The group $H = \mathrm{Gal}(E/K')$ is a solvable subgroup of G. Choose a subnormal series with cyclic prime-order quotients,

$$H = H_0 \triangleright H_1 \triangleright \cdots \triangleright H_m = 1.$$

Its fixed fields form a tower

$$K' = E^{H_0} \subset E^{H_1} \subset \cdots \subset E^{H_m} = E,$$

and each step is cyclic Galois. Applying Proposition 3.2 at each step realizes the tower by radicals. Normality is needed between consecutive groups; the individual H_i need not all be normal in H. This is the structural reason solvable groups lead to radical solutions [2].

For explicit formulas, one usually chooses a larger abelian normal subgroup when its character coordinates make the invariant algebra simple. The choices of subgroup, ordering, primitive roots of unity, and bases of character lines are not unique. Fourier decomposition makes their roles precise; it does not remove coordinate choices or automatically minimize the size of a formula.

Finally, two settings must remain distinct. In a universal calculation, roots are variable coordinates and one studies a polynomial ring, possibly modulo their sum, with a permutation action. Before inverting the discriminant, this root space includes ramification: its character components need not be line bundles over the invariant ring, and their multiplication maps need not

be isomorphisms. On the finite étale torsor locus, Proposition 2.1 proves that they are line bundles with isomorphic multiplication maps. In a particular Galois field extension, these become one-dimensional vector spaces with nonzero multiplication. Universal polynomial identities specialize to concrete equations even when roots coincide or the actual Galois group becomes smaller; statements about Galois degrees require the corresponding field hypotheses. The abstract group predicts possible intermediate symmetries. The algebra, its multiplication, and the given coefficients determine the equation being solved.

4 Centering: the first substitution is a projection

Let

$$f(X) = X^n + a_1 X^{n-1} + \cdots + a_n$$

have roots $\alpha_1, \ldots, \alpha_n$, counted with multiplicity. Their mean is

$$m = \frac{\alpha_1 + \cdots + \alpha_n}{n} = -\frac{a_1}{n}.$$

Writing $X = Y + m$ replaces the roots by $r_i = \alpha_i - m$, whose sum is zero. The coefficient of Y^{n-1} therefore vanishes.

This is the projection associated with the decomposition

$$K^n = K(1, \ldots, 1) \oplus W_n, \qquad W_n = \{(r_1, \ldots, r_n) : r_1 + \cdots + r_n = 0\}. \tag{1}$$

The first summand records the common translation of all roots; the second records their configuration relative to their mean. Both are preserved by permutations. Thus depressing a polynomial removes the trivial representation and leaves the standard representation. This projection is equivariant and uses a scalar already in the coefficient field. The splitting field and Galois group are unchanged; only the coordinates have been simplified.

There is also an intrinsic formulation. In the finite algebra $E = K[X]/(f)$, let α be the class of X. The trace of multiplication by α is $-a_1$. Replacing α by

$$\alpha - \frac{\mathrm{Tr}_{E/K}(\alpha)}{n}\, 1$$

makes its trace zero. This description does not require choosing or ordering the roots.

For later reference, the translations give

$$X^3 + AX^2 + BX + C \longmapsto Y^3 + pY + q, \qquad X = Y - A/3,$$
$$p = B - A^2/3, \qquad q = C - AB/3 + 2A^3/27;$$

and

$$X^4 + AX^3 + BX^2 + CX + D \longmapsto Y^4 + pY^2 + qY + r, \qquad X = Y - A/4,$$
$$p = B - 3A^2/8,$$
$$q = C - AB/2 + A^3/8,$$
$$r = D - AC/4 + A^2B/16 - 3A^4/256.$$

5 Cardano: separate the cyclic motion

Start with the coefficient data

$$g(Y) = Y^3 + pY + q$$

over a field of characteristic zero. For the identities, introduce three formal roots r_0, r_1, r_2 with sum zero and let σ cycle their labels. Work in their polynomial algebra B, imposing the coefficient identities of g. Define the orbit sum and orbit product by

$$\mathrm{Tr}_B(x) = x + \sigma(x) + \sigma^2(x), \qquad \mathrm{Norm}_B(x) = x\,\sigma(x)\,\sigma^2(x).$$

Writing $r = r_0$, we have $\mathrm{Tr}_B(r) = 0$ and $g(Y) = \mathrm{Norm}_B(Y - r)$. These are finite polynomial operations, and adjoining a primitive cube root ω makes the scalar-update coordinates available. The uniform cover in Section 2 identifies these same sums and products with relative trace and norm on the distinct-root locus; the polynomial derivation itself does not require that interpretation.

5.1 Why the root must be a sum of two components

The character decomposition of this algebra has three summands:

$$B = B_1 \oplus B_\chi \oplus B_{\chi^{-1}}, \qquad \chi(\sigma) = \omega.$$

The projector onto B_1 is normalized trace. Thus the trace-zero condition removes precisely that summand. Applying the identity decomposition to the actual element r gives

$$r = (e_1 + e_\chi + e_{\chi^{-1}})r = e_\chi r + e_{\chi^{-1}} r.$$

Define $u = e_\chi r$ and $v = e_{\chi^{-1}} r$. Then $r = u + v$ follows before any polynomial expansion. The two terms are prescribed projections of the root observable, and their transformation rules are part of their definition. This is why this particular decomposition is useful: multiplication will turn uv, u^3, and v^3 into invariant quantities.

To calculate the projectors, label the conjugates $r_j = \sigma^j r$, for $j = 0, 1, 2$. Proposition 3.1 gives

$$u = \frac{r_0 + \omega^2 r_1 + \omega r_2}{3}, \qquad v = \frac{r_0 + \omega r_1 + \omega^2 r_2}{3}. \tag{2}$$

The scalar-update requirement in the projector calculation selects exactly these weighted sums. Their transformation rules are

$$\sigma(u) = \omega u, \qquad \sigma(v) = \omega^{-1} v.$$

Applying σ^j to the decomposition of r therefore gives

$$r_j = \omega^j u + \omega^{-j} v, \qquad j = 0, 1, 2. \tag{3}$$

The inverse transform is the return map to the original roots. Centering leaves two scalar-update quantities; their multipliers have order three, so cubing removes the cyclic motion. Reversing the order of cycling interchanges the two quantities. Their cubes therefore form an unordered pair governed by a quadratic. What was a search for three individual roots has become a calculation of two powers and one compatibility relation.

When the coefficients vary over $\mathbb{C}$, this also describes analytic continuation. A loop on the orientation cover can rotate the three root sheets; its action on u, v is multiplication by opposite cube-root phases. Returning to the coefficient space permits orientation reversal and hence interchange of the two components. Cubes erase the rotation, and a mixed product preserves the relative phase needed for reconstruction. The topological monodromy is the same action used by the algebraic projectors.

5.2 Multiplication and trace determine the auxiliary equations

We now calculate the invariant data without substituting an ansatz into the cubic. In the character grading,

$$uv \in B_1, \qquad u^3, v^3 \in B_1, \qquad u^2, v^2, u^2v, uv^2 \text{ have nontrivial characters.}$$

Adding the three cyclic images kills every nontrivial component: its scalar multipliers sum to zero. On an invariant quantity, the orbit sum is three times that quantity. Expand the square and cube, then discard precisely the terms with nontrivial total character. This gives

$$\mathrm{Tr}_B(r^2) = 6uv, \qquad \mathrm{Tr}_B(r^3) = 3(u^3 + v^3).$$

The coefficient identities for the given polynomial g, equivalently the sums of powers of its three roots, are

$$\mathrm{Tr}_B(r^2) = -2p, \qquad \mathrm{Tr}_B(r^3) = -3q.$$

For example, the first is Newton's identity with $\mathrm{Tr}_B(r) = 0$; the second also follows by taking the trace of $r^3 + pr + q = 0$. Hence

$$uv = -\frac{p}{3}, \qquad u^3 + v^3 = -q. \tag{4}$$

The coefficients enter through trace and norm, while the character grading determines which products survive. The familiar cancellation conditions have emerged as identities of the algebra we constructed.

There is a useful global presentation of the same mechanism. In the universal centered root-coordinate ring, put

$$s = uv, \qquad U = u^3, \qquad V = v^3.$$

These are fixed by the cyclic action, and their multiplication retains the relation $UV = s^3$.

Proposition 5.1. *The invariant rings of the cyclic and full permutation actions are*

$$k[u, v]^{C_3} = k[s, U, V]/(UV - s^3), \qquad k[u, v]^{S_3} = k[s, U + V]. \tag{5}$$

Proof. A monomial $u^i v^j$ is fixed by the cyclic action precisely when $i - j$ is divisible by 3. Removing a power of uv leaves a power of either u^3 or v^3. These monomials give the first presentation, with the single relation $UV = s^3$. The remaining involution interchanges U, V and fixes s. Its invariants are generated by $s, U + V, UV$, and the relation eliminates UV. □

Consequently the map from the cyclic quotient to the coefficient space is explicitly

$$(s, U, V) \longmapsto (p, q) = (-3s, -U - V), \qquad UV = s^3.$$

Recovering a point above (p, q) requires separating the unordered pair U, V, then lifting one cyclic character coordinate by a cube root. The invariant presentation extends as a polynomial identity across the discriminant locus; the torsor interpretation applies on the distinct-root locus.

5.3 The quadratic resolvent and compatible cube roots

The unordered pair $\{U, V\}$ has sum $-q$ and product $-p^3/27$. Therefore its members are the roots of

$$T^2 + qT - \frac{p^3}{27} = 0. \tag{6}$$

Writing

$$\delta = \frac{q^2}{4} + \frac{p^3}{27},$$

we obtain

$$U = -\frac{q}{2} + \sqrt{\delta}, \qquad V = -\frac{q}{2} - \sqrt{\delta}.$$

Suppose first that $U \neq 0$. Choose a cube root u of U, and then define

$$v = -\frac{p}{3u}. \tag{7}$$

The product relation ensures $v^3 = V$. Formula (3) now gives all three roots.

The instruction in (7) is essential. Independently chosen cube roots of U and V need not have product $-p/3$. Once u is chosen, v is determined. Replacing u by ωu replaces v by $\omega^{-1}v$, cycling the three answers.

If $U = 0$ but $V \neq 0$, interchange their roles. If both vanish, $p = q = 0$ and all three roots are zero. Thus the apparent division by u is a choice of coordinate chart, not a restriction on the method.

For the original cubic, the result is

$$X_j = -\frac{A}{3} + \omega^j u + \omega^{-j} v, \qquad j = 0, 1, 2, \tag{8}$$

with p, q as in Section 4 and the radicals chosen compatibly.

5.4 A worked cubic

Take $X^3 + 3X^2 - 6 = 0$. The translation $Y = X + 1$ gives

$$Y^3 - 3Y - 4 = 0.$$

The invariant equations are $uv = 1$ and $u^3 + v^3 = 4$. Thus

$$U, V = 2 \pm \sqrt{3}.$$

Taking $u = \sqrt[3]{2 + \sqrt{3}}$ to be the positive real cube root and setting $v = 1/u$ produces

$$X_0 = -1 + \sqrt[3]{2 + \sqrt{3}} + \sqrt[3]{2 - \sqrt{3}}.$$

The other roots are $-1 + \omega u + \omega^2 v$ and $-1 + \omega^2 u + \omega v$. The reciprocal relation between the real cube roots is the numerical form of the invariant constraint $uv = 1$.

5.5 What the square root remembers

The discriminant is

$$\Delta = -4p^3 - 27q^2 = -27(U - V)^2.$$

The alternating polynomial

$$d = (r_0 - r_1)(r_0 - r_2)(r_1 - r_2)$$

satisfies

$$d = 3(\omega - \omega^2)(U - V), \qquad d^2 = \Delta. \tag{9}$$

Even permutations preserve d, whereas odd permutations reverse its sign. A choice of square root of Δ distinguishes the two cyclic orientations of the three roots. The remaining cube-root choice selects a starting point in the chosen cyclic order.

There is a small but important field distinction. Cardano uses $\sqrt{\delta}$, while the alternating invariant uses $\sqrt{\Delta}$, and $\delta = -\Delta/108$. Over $k = K(\omega)$ these generate the same extension, because k contains $\sqrt{-3} = \omega - \omega^2$. They need not generate the same extension over K. For example, $Y^3 - 2$ has $\delta = 1$ but discriminant -108. The roots of unity are part of the Galois-theoretic construction.

6 Ferrari: reconstruct a root from three character lines

Apply the same construction to

$$g(Y) = Y^4 + pY^2 + qY + r,$$

using four formal roots whose sum is zero. Let the double transpositions permute their labels. They form V_4; each move preserves one pairing of the four labels and exchanges signs on the other two pairing coordinates. The quotient $S_4/V_4 \simeq S_3$ permutes the three pairings

$$12 \mid 34, \qquad 13 \mid 24, \qquad 14 \mid 23. \tag{10}$$

Thus the surjection

$$S_4 \longrightarrow S_3 \tag{11}$$

predicts a cubic resolvent before any auxiliary polynomial is chosen. We now construct the quantities that it will govern.

6.1 Find coordinates for the double sign changes

Let r_1 be one formal root, and let χ_1, χ_2, χ_3 be the three nontrivial characters of V_4. Averaging its four images is zero because the roots sum to zero. The projector identity therefore gives

$$r_1 = e_{\chi_1} r_1 + e_{\chi_2} r_1 + e_{\chi_3} r_1.$$

Define $u = 2e_{\chi_1} r_1$, $v = 2e_{\chi_2} r_1$, and $w = 2e_{\chi_3} r_1$. This normalization makes the coordinates sums of paired roots; the decomposition itself gives

$$r_1 = (u + v + w)/2.$$

To calculate the three quantities, order the remaining root labels by the double transpositions and apply the weighted sums. The result can then be checked directly under each of the three moves:

$$\begin{aligned} u &= \frac{r_1 + r_2 - r_3 - r_4}{2} = r_1 + r_2, \\ v &= \frac{r_1 - r_2 + r_3 - r_4}{2} = r_1 + r_3, \\ w &= \frac{r_1 - r_2 - r_3 + r_4}{2} = r_1 + r_4. \end{aligned} \tag{12}$$

Their transformation rules are completely explicit:

	u	v	w
$(12)(34)$	u	$-v$	$-w$
$(13)(24)$	$-u$	v	$-w$
$(14)(23)$	$-u$	$-v$	w

Each nontrivial group element reverses exactly two signs. Applying these three group elements to the decomposition of the root observable gives

$$\begin{aligned} r_1 &= (u + v + w)/2, \quad & r_2 &= (u - v - w)/2, \\ r_3 &= (-u + v - w)/2, \quad & r_4 &= (-u - v + w)/2. \end{aligned} \tag{13}$$

Thus the eventual sums of square roots are already dictated by Fourier inversion. What remains is to find the squares and retain the correct relation among their signs.

6.2 The invariants, including the information squares forget

Put

$$U = u^2, \qquad V = v^2, \qquad W = w^2, \qquad t = uvw.$$

The squares are attached naturally to the pairings. For example,

$$U = (r_1 + r_2)^2 = (r_3 + r_4)^2,$$

so U does not distinguish the two blocks of $12 \mid 34$. The product t, however, retains a compatibility condition that three unrelated squares would lose.

Proposition 6.1. *For the Klein four action,*

$$K[u, v, w]^{V_4} = K[U, V, W, t]/(t^2 - UVW). \tag{14}$$

The quotient S_4/V_4 permutes U, V, W and fixes t. If e_1, e_2, e_3 denote the elementary symmetric polynomials in U, V, W, then

$$K[u, v, w]^{S_4} = K[e_1, e_2, t], \qquad e_3 = t^2. \tag{15}$$

Proof. A monomial $u^i v^j w^\ell$ survives all the double sign changes exactly when i, j, ℓ have the same parity. If all are even, it is a monomial in U, V, W; if all are odd, it is uvw times such a monomial. This proves (14) and its single relation.

The transpositions $(23), (34)$ interchange u, v and v, w, respectively. The transposition (12) sends (u, v, w) to $(u, -w, -v)$. These generate S_4, so t is fixed and the squares are permuted. Every V_4-invariant has a unique expression $A(U, V, W) + tB(U, V, W)$; full invariance says that A, B are symmetric. Their dependence on e_3 can be replaced by dependence on t^2. □

Write $e_1 = U + V + W$ and $e_2 = UV + UW + VW$. Define trace here as the sum over the four double-sign patterns in the displayed table. A monomial with a nontrivial character cancels in this sum; an invariant monomial occurs four times. Apply this rule to the square, cube, and fourth power of $r_1 = (u + v + w)/2$:

$$\mathrm{Tr}(r_1^2) = e_1, \qquad \mathrm{Tr}(r_1^3) = 3t, \qquad \mathrm{Tr}(r_1^4) = \frac{e_1^2 + 4e_2}{4}.$$

For example, in the cube only the mixed term uvw has trivial character. In the fourth power, the surviving terms are the pure fourth powers and the products of two squares. Their coefficients give the last formula. The known quartic gives Newton's identities

$$\mathrm{Tr}(r_1^2) = -2p, \qquad \mathrm{Tr}(r_1^3) = -3q, \qquad \mathrm{Tr}(r_1^4) = 2p^2 - 4r.$$

Comparing the two descriptions yields

$$p = -\frac{e_1}{2}, \qquad q = -t, \qquad r = \frac{e_1^2 - 4e_2}{16}. \tag{16}$$

Equivalently,

$$U + V + W = -2p, \quad UV + UW + VW = p^2 - 4r, \quad UVW = q^2, \quad uvw = -q. \tag{17}$$

6.3 The resolvent cubic and coherent square roots

The quotient has permuted the three character squares while retaining their product constraint. The first three identities in (17) therefore say precisely that U, V, W are the roots of

$$R(Z) = Z^3 + 2pZ^2 + (p^2 - 4r)Z - q^2. \tag{18}$$

This cubic is now available to Cardano's method. Once all three of its roots have been ordered, choose square roots u, v, w satisfying

$$u^2 = U, \quad v^2 = V, \quad w^2 = W, \qquad uvw = -q. \tag{19}$$

Then (13) gives the four roots of the quartic.

For $q \neq 0$, all three squares are nonzero. One may choose u and v arbitrarily and set $w = -q/(uv)$. The relation $UVW = q^2$ gives $w^2 = W$. There are four compatible sign choices, corresponding exactly to V_4. An arbitrary choice among the eight sign patterns would forget the cubic invariant uvw and might change the sign of the coefficient q.

For $q = 0$, the quartic is biquadratic:

$$Y^4 + pY^2 + r = 0.$$

Solve $Z^2 + pZ + r = 0$ and take both square roots of each answer. In the invariant description, at least one of u, v, w vanishes, and the product condition is automatically satisfied. No division by a vanishing coordinate is needed.

The familiar alternative resolvent with roots

$$\theta_1 = r_1r_2 + r_3r_4, \quad \theta_2 = r_1r_3 + r_2r_4, \quad \theta_3 = r_1r_4 + r_2r_3$$

contains exactly the same information. Indeed $\theta_1 = p + U$, and similarly for the other two. Translating (18) by $\theta = Z + p$ gives

$$\theta^3 - p\theta^2 - 4r\theta + 4pr - q^2 = 0. \tag{20}$$

This is the standard pairing resolvent used to study quartic Galois groups; see Milne's treatment in Chapter 4 of [2].

The resolvent also preserves the discriminant. For example,

$$U - V = (r_2 - r_3)(r_1 - r_4),$$

and the analogous identities for $U - W$ and $V - W$ together account for all six root differences. Squaring and multiplying gives

$$\operatorname{disc} R = \operatorname{disc} g.$$

In particular, distinct quartic roots give three distinct pairing invariants. The resolvent does not accidentally merge two pairings away from the discriminant locus.

7 Why Ferrari's perfect square appears

Ferrari's practical method can work with one pairing. A quartic factored into two monic quadratics can be rewritten using the average and half-difference of those factors. This simple change makes their product a difference of squares. We will determine the coefficients of that new form directly.

The pairing is not initially known numerically. We use formal root labels to discover the shape of the factors, then obtain an equation for the pairing coordinate from the given coefficients. The construction works as a polynomial identity; on the distinct-root locus, the covers in Section 10 additionally describe which pairing and block each choice selects.

7.1 Average and difference expose the square

For the pairing $12 \mid 34$, put the first two roots in one factor and the last two in the other. Each factor is monic quadratic and their product is the given quartic. Exchanging the blocks exchanges the two factors. We can therefore look for a part fixed by this exchange and a part whose sign changes.

Write the two factors as

$$Q_-(Y) = (Y - r_1)(Y - r_2), \qquad Q_+(Y) = (Y - r_3)(Y - r_4).$$

Let ι exchange the two blocks, so that it exchanges Q_+ and Q_-. Their product can be written as the product of one factor and its exchanged value:

$$g(Y) = Q_+(Y)\,\iota(Q_+(Y)).$$

This product is unchanged when the blocks are exchanged. In the covering interpretation it is a norm. For the calculation here, only the exchange operation and ordinary polynomial multiplication are needed.

Apply the two character projectors of this involution:

$$A(Y) = \frac{Q_+(Y) + Q_-(Y)}{2}, \qquad B(Y) = \frac{Q_+(Y) - Q_-(Y)}{2}.$$

Then $\iota(A) = A$, $\iota(B) = -B$, and Fourier inversion gives $Q_+ = A + B$, $Q_- = A - B$. Multiplying these two reconstructed components gives

$$g(Y) = A(Y)^2 - B(Y)^2. \tag{21}$$

The difference of squares has been derived from a change of working quantities: average and half-difference in place of the two factors. The two-character projector language records exactly these two operations. The following coefficient comparison makes this form computable before the factors themselves are known.

Centering determines the shapes of the two components. Their average is monic with vanishing linear coefficient, whereas their difference has degree at most one. Hence

$$A(Y) = Y^2 + m, \qquad B(Y) = sY + d, \qquad s = r_1 + r_2 = -(r_3 + r_4).$$

The quantities s and d change sign together when the blocks are exchanged. In particular, $z = s^2$ is attached to the unordered pairing: it is precisely one root of R. Choosing a sign of s distinguishes the blocks whenever $s \neq 0$.

7.2 The coefficients determine the reconstruction data

Only now do we use the known coefficients p, q, r. In the norm $(Y^2+m)^2-(sY+d)^2$, the coefficients of Y^2, Y, and 1 are respectively

$$2m - s^2 = p, \qquad -2sd = q, \qquad m^2 - d^2 = r.$$

For $q \neq 0$, every resolvent root z is nonzero. The first two equations then give

$$m = \frac{p+z}{2}, \qquad d = -\frac{q}{2s}.$$

The third is the remaining norm compatibility condition:

$$z(p+z)^2 - q^2 = 4rz,$$

which is exactly $R(z) = 0$. The resolvent therefore certifies the descent of the quadratic norm to the prescribed quartic. Its degree three records the three possible pairings; the square root of its chosen root records the two possible block orders.

The reconstructed factors are

$$Q_- = A - B = Y^2 - sY + b, \qquad Q_+ = A + B = Y^2 + sY + c,$$

where

$$b = \frac{p + z + q/s}{2}, \qquad c = \frac{p + z - q/s}{2}. \tag{22}$$

Their roots are

$$\frac{s \pm \sqrt{-2p - z - 2q/s}}{2}, \qquad \frac{-s \pm \sqrt{-2p - z + 2q/s}}{2}. \tag{23}$$

These last two square roots distinguish the roots within each block. Changing s to $-s$ exchanges the two blocks, while changing either remaining sign exchanges the two roots of that block. Undoing the initial centering gives the roots of the original quartic.

In Ferrari's familiar notation, the same norm reads

$$g(Y) = (Y^2 + m)^2 - \left(sY - \frac{q}{2s}\right)^2, \qquad m = \frac{p+z}{2}. \tag{24}$$

Eliminating $z = s^2 = 2m - p$ gives the usual auxiliary equation

$$q^2 = 4(2m - p)(m^2 - r).$$

The parameter m is the constant coefficient of the invariant projection $(Q_+ + Q_-)/2$. The linear square is the square of the anti-invariant projection $(Q_+ - Q_-)/2$. Both substitutions have thus been constructed by the cover and its character projectors.

If $q = 0$, the original quartic is biquadratic, and one solves $Z^2 + pZ + r = 0$ before taking square roots. The block cover and its norm still exist for distinct roots, but s can vanish and no longer distinguish its sheets. The formulas involving $1/s$ are a coordinate chart for the reconstruction, not a requirement that every quartic lie in that chart.

8 Two quartics worked through

8.1 An example in which the signs are visible

Consider

$$g(Y) = Y^4 - 10Y^2 - 4Y + 8.$$

Its resolvent is

$$R(Z) = Z^3 - 20Z^2 + 68Z - 16 = (Z - 4)(Z^2 - 16Z + 4),$$

so we may order the squares as

$$U = 4, \qquad V = 8 + 2\sqrt{15}, \qquad W = 8 - 2\sqrt{15}.$$

The choices

$$u = 2, \qquad v = \sqrt{5} + \sqrt{3}, \qquad w = \sqrt{5} - \sqrt{3}$$

have the required squares and satisfy $uvw = 4 = -q$. Formula (13) gives

$$1 + \sqrt{5}, \quad 1 - \sqrt{5}, \quad -1 + \sqrt{3}, \quad -1 - \sqrt{3}.$$

Alternatively, selecting only $z = 4$ and $s = 2$ gives

$$g(Y) = (Y^2 - 2Y - 4)(Y^2 + 2Y - 2).$$

The two presentations show the same geometry at different resolutions: the first orders all three pairings and reconstructs the four roots; the second chooses a single pairing and solves its two blocks.

8.2 An example with the full symmetric Galois group

The preceding example has extra arithmetic simplifications. For a quartic with full S_4 symmetry over $\mathbb{Q}$, take

$$g(Y) = Y^4 - Y - 1.$$

Its resolvent is

$$R(Z) = Z^3 + 4Z - 1.$$

Cardano gives its unique real root as

$$z = \sqrt[3]{\frac{1}{2} + \sqrt{\frac{283}{108}}} + \sqrt[3]{\frac{1}{2} - \sqrt{\frac{283}{108}}}, \tag{25}$$

where real cube roots are meant. Their product is $-4/3$, as required. Since $R(0) < 0 < R(1)$ and $R'(Z) = 3Z^2 + 4 > 0$ for real Z, we have $0 < z < 1$. Set $s = \sqrt{z} > 0$. Ferrari gives the real factorization

$$Y^4 - Y - 1 = \left(Y^2 - sY + \frac{z - 1/s}{2}\right)\left(Y^2 + sY + \frac{z + 1/s}{2}\right).$$

Its four roots are

$$\frac{s \pm \sqrt{-z + 2/s}}{2}, \qquad \frac{-s \pm i\sqrt{z + 2/s}}{2}.$$

This is a complete radical solution, with the choices specified rather than hidden in nested radical notation.

For completeness, its Galois group can be checked with little computation. Modulo 2, the polynomial becomes $Y^4 + Y + 1$. It has no linear factor, and it is not divisible by the only irreducible quadratic over $\mathbb{F}_2$, namely $Y^2 + Y + 1$; reduction by that quadratic leaves remainder 1. Hence the quartic is irreducible over $\mathbb{Q}$. The resolvent has no rational root and has discriminant -283, so its Galois group is S_3.

The quartic Galois group $G \subseteq S_4$ is transitive and maps onto this S_3. Thus $|G|$ is divisible by both 4 and 6. The only possibilities are 12 and 24. The unique index-two subgroup of S_4 is A_4, whose image on pairings is only C_3, so the order cannot be 12. Therefore $G = S_4$. The calculation above genuinely traverses the nonabelian symmetry of a general quartic.

9 Cardano in every degree: the dihedral family

Cardano's calculation extends to an infinite family once its symmetry is made explicit. The relevant feature of the cubic is a pair of Fourier coordinates carrying inverse characters. Keep that representation, replace the third roots of unity by the n-th roots of unity, and ask what equation the quotient produces.

9.1 The invariant ring determines the family

Let $n \geq 3$, and let k have characteristic zero and contain a primitive n-th root of unity ζ. On $k[u, v]$, consider

$$\rho(u, v) = (\zeta u, \zeta^{-1} v), \qquad \tau(u, v) = (v, u).$$

The relations $\rho^n = \tau^2 = 1$ and $\tau\rho\tau = \rho^{-1}$ identify the acting group as $\mathrm{Dih}_n = C_n \rtimes C_2$, of order $2n$. Put

$$s = uv, \qquad U = u^n, \qquad V = v^n, \qquad t = U + V.$$

The same invariant-monomial argument used for the cubic gives

$$k[u, v]^{C_n} = k[s, U, V]/(UV - s^n), \qquad k[u, v]^{\mathrm{Dih}_n} = k[s, t]. \tag{26}$$

Indeed, $u^i v^j$ is rotation-invariant precisely when n divides $i-j$. After factoring powers of uv, it is a power of u^n or v^n. The reflection exchanges U and V, leaving their sum and product; the latter is already s^n.

The quotient by the reflection alone has coordinates $x = u + v$ and $s = uv$. Therefore $u^n + v^n$ must be a polynomial in x, s. Its formula is forced by

$$(u+v)(u^{n-1}+v^{n-1}) = u^n + v^n + uv(u^{n-2}+v^{n-2}).$$

Define the Dickson polynomials by

$$D_0(X,a) = 2, \quad D_1(X,a) = X, \quad D_n(X,a) = XD_{n-1}(X,a) - aD_{n-2}(X,a). \tag{27}$$

The recurrence proves

$$D_n(u+v, uv) = u^n + v^n. \tag{28}$$

For example,

$$\begin{aligned} D_3(X,a) &= X^3 - 3aX, \\ D_4(X,a) &= X^4 - 4aX^2 + 2a^2, \\ D_5(X,a) &= X^5 - 5aX^3 + 5a^2X. \end{aligned}$$

Thus the equation arises by fixing the invariant $s = a$ and expressing the remaining invariant t on the reflection quotient. For $a = 1$, these polynomials satisfy $D_n(X,1) = 2T_n(X/2)$, where T_n is the Chebyshev polynomial. Their arithmetic applications are developed in [3].

There are two complementary descriptions of the information retained. The coordinate $x = u+v$, together with s, remembers the unordered pair $\{u, v\}$. The coordinates s, t remember its entire orbit under cyclic phase changes as well. Generically, the degree-n polynomial describes the n possible unordered pairs lying above a fixed dihedral orbit. Its degree therefore counts the remaining choices; its coefficients express the information common to every choice.

9.2 One quadratic and one radical in every degree

Fix $a \neq 0$. To solve $D_n(X,a) = t$, the invariant equations require

$$U + V = t, \qquad UV = a^n.$$

Hence

$$U, V = \frac{t \pm \sqrt{t^2 - 4a^n}}{2}.$$

Choose u with $u^n = U$, and set $v = a/u$. The product constraint then ensures $v^n = V$. Fourier reconstruction gives

$$X_j = \zeta^j u + \zeta^{-j}\frac{a}{u}, \qquad j = 0, \ldots, n-1. \tag{29}$$

These are all the roots, counted with multiplicity. One way to see the multiplicities as well is to form the monic product of the $X - X_j$: for independent u, v, its n distinct roots satisfy (28), so it equals $D_n(X, uv) - (u^n + v^n)$. This polynomial identity remains valid after specialization.

Cardano is exactly the case $n = 3$, with $a = -p/3$ and $t = -q$. The case $a = 0$, excluded only to allow division by u, is simply $X^n = t$.

For every n, the root vector in (29) has just two Fourier frequencies, $+1$ and -1. This is the structural restriction that makes high degree manageable. A general polynomial of degree n does not have this two-frequency form.

The geometric Galois group is equally explicit. With z transcendental and $a \in k^\times$ fixed, write

$$x = z + a/z, \qquad t = z^n + a^n/z^n.$$

The extension $k(z)/k(t)$ has degree $2n$. Its $2n$ automorphisms are generated by $z \mapsto \zeta z$ and $z \mapsto a/z$, so its group is Dih_n. It is the splitting field of $D_n(X, a) - t$: the roots are given by (29), and the first two recover z by a linear equation. Solving the equation reverses this quotient construction. A quadratic distinguishes u^n from v^n; an n-th root chooses a cyclic phase.

The familiar trigonometric solution has the same source. Take $a = 1$ and $t = 2\cos\theta$, where $0 < \theta < \pi$. Then $u^n = e^{i\theta}$ and $v = u^{-1}$, so the roots are

$$X_j = 2\cos\left(\frac{\theta + 2\pi j}{n}\right), \qquad j = 0, \dots, n-1.$$

All are real and distinct. A point on the unit circle projects to twice its real part under $z \mapsto z + z^{-1}$, while multiplication of angles upstairs becomes the Dickson polynomial downstairs. At $n = 3$, this is the trigonometric form of Cardano's method when all three roots are real. The radical and trigonometric formulas are two coordinate descriptions of the same covering construction. For higher n, they continue to solve the same dihedral family without a new substitution.

9.3 An irreducible quintic with an explicit solution

Consider

$$f(X) = X^5 - 5X^3 + 5X - 3.$$

Put

$$u = \sqrt[5]{\frac{3 + \sqrt{5}}{2}} > 0, \qquad \zeta = e^{2\pi i/5}.$$

All five roots are

$$X_j = \zeta^j u + \zeta^{-j} u^{-1}, \qquad j = 0, \dots, 4.$$

In particular, the real root is

$$\sqrt[5]{\frac{3 + \sqrt{5}}{2}} + \sqrt[5]{\frac{3 - \sqrt{5}}{2}}.$$

This is an irreducible quintic: the translate

$$f(X + 3) = X^5 + 15X^4 + 85X^3 + 225X^2 + 275X + 120$$

is Eisenstein at 5.

Here the arithmetic Galois group over $\mathbb{Q}$ is $C_5 \rtimes C_4$, rather than the geometric group Dih_5. This distinction records the action on roots of unity. To verify it, let $M = \mathbb{Q}(\zeta, u)$. Since $\sqrt{5} \in \mathbb{Q}(\zeta)$, its degree is at most 20; irreducibility of f forces $[M : \mathbb{Q}(\zeta)] = 5$, hence equality. The automorphisms

$$\rho : \ (\zeta, u) \longmapsto (\zeta, \zeta u), \qquad \tau : \ (\zeta, u) \longmapsto (\zeta^2, u^{-1})$$

act on root indices as $j \mapsto j + 1$ and $j \mapsto 3j$. The second is well-defined because $\zeta \mapsto \zeta^2$ sends $\sqrt{5}$ to $-\sqrt{5}$. These generate 20 distinct automorphisms, acting faithfully on the roots, and therefore identify M as their splitting field with the asserted group. Its solvability explains the radical expression; the general quintic has a different Galois group [2].

9.4 The same quotient in finite-field computation

The identity also gives an inversion algorithm for nonlinear maps. Work now over $\mathbb{F}_q$, with $a = 1$. The polynomial identity

$$D_m(D_n(X, 1), 1) = D_{mn}(X, 1)$$

follows by evaluating at $X = z + z^{-1}$.

For $x \in \mathbb{F}_q$, choose a root z of $Z^2 - xZ + 1$ in $\mathbb{F}_{q^2}$. Frobenius either fixes z or exchanges it with z^{-1}. Thus $z^{q-1} = 1$ or $z^{q+1} = 1$, and in either case $z^{q^2-1} = 1$. If

$$\gcd(n, q^2 - 1) = 1, \qquad mn \equiv 1 \pmod{q^2 - 1},$$

then

$$D_m(D_n(x, 1), 1) = z^{mn} + z^{-mn} = x.$$

Consequently D_n permutes $\mathbb{F}_q$, with inverse D_m. For example, over $\mathbb{F}_7$, the inverse of $D_5(X, 1) = X^5 + 2X^3 + 5X$ is $D_{29}(X, 1)$, since $5 \cdot 29 \equiv 1 \pmod{48}$. These are equalities of functions on the finite field. The full permutation criterion, including necessity, appears in [3].

The calculation takes place in cyclic multiplicative groups, where inversion means inverting an exponent. Passing through $z \mapsto z + z^{-1}$ transports that calculation to a polynomial map. The same environment that organizes Cardano's radicals therefore also organizes finite-field permutations.

10 Read the radical choices in the covering tower

The covers of Section 2 produced the reconstruction calculations. We can now identify exactly which sheet each radical choice selects and which information each map forgets. Let W_n now denote the affine space of centered ordered roots, and let W_n° be its open subset where no two roots coincide. Its quotient by S_n is the space of centered monic polynomials with distinct roots. A point of the quotient remembers the coefficients and forgets the ordering.

Over the distinct-root locus, the maps below are finite *étale* covers: algebraic counterparts of covering spaces, with no branching. Over $\mathbb{C}$, one can follow the roots continuously along paths in coefficient space. A loop can return the same polynomial while permuting its roots; this permutation is its monodromy.

10.1 The cubic tower

The cubic construction factors the ordered-root cover as

$$W_3^\circ \xrightarrow{C_3} W_3^\circ/C_3 \xrightarrow{C_2} W_3^\circ/S_3. \tag{30}$$

The arrows forget information. The first forgets which root is the starting point in a cyclic ordering. The second forgets the cyclic orientation. Solving moves in the reverse direction: a square root chooses the orientation, and a cube root chooses the starting point. Before restricting to distinct roots, the intermediate quotient has the explicit equation $UV = s^3$.

For a particular separable cubic with splitting field L over $k = K(\omega)$ and $\mathrm{Gal}(L/k) = S_3$, the corresponding field tower is

$$k \subset k(\sqrt{\Delta}) = L^{C_3} \subset L,$$

with degrees 2 and 3. If the actual Galois group is smaller, these degrees can drop. The universal construction and its polynomial identities remain valid.

10.2 The quartic tower and a distinction worth keeping

The quartic has the tower

$$W_4^\circ \xrightarrow{V_4} W_4^\circ/V_4 \xrightarrow{S_3} W_4^\circ/S_4. \tag{31}$$

The intermediate space remembers an ordering of the three pairings, together with the invariants already specified by the quartic. Its generic fiber over coefficient space has six points, one for each ordering of the three resolvent roots. Over it, the four remaining choices are the compatible sign patterns in (19).

If a quartic has splitting field L/K and Galois group S_4, then

$$K \subset L^{V_4} \subset L$$

has degrees 6 and 4. The field L^{V_4} is the *splitting field* of the resolvent cubic. It is generally not the degree-three field obtained by adjoining just one resolvent root.

Indeed, one pairing is stabilized by a dihedral subgroup $D_8 \subset S_4$ of order eight. Selecting one resolvent root follows the degree-three cover associated with this stabilizer. Choosing $s = \sqrt{z}$ then distinguishes its two blocks, and solving their two quadratics distinguishes the roots inside each block. Generically this route has successive degrees $3, 2, 2, 2$. Solving the whole resolvent first instead gives degrees $6, 2, 2$. Both recover the same splitting field.

The stabilizer chain for the one-pairing route is explicitly

$$S_4 \supset D_8 \supset \langle (12), (34) \rangle \supset \langle (34) \rangle \supset \{1\}.$$

The group $\langle (12), (34) \rangle$ preserves each of the two blocks separately. Although it is abstractly a Klein four group, it is different from the normal double-transposition subgroup V_4. Adjoining the difference $r_1 - r_2$ leaves only the swap (34); adjoining $r_3 - r_4$ removes that last ambiguity.

The distinction explains why the practical one-resolvent-root algorithm and the normal-subgroup tower organize their square roots differently. A choice of one pairing does not yet reduce the symmetry to V_4; an ordering of all three pairings does.

These degree lists describe intermediate fields inside the quartic splitting field. A literal implementation of Cardano may temporarily adjoin a cube root of unity outside that field. In particular, the degree-three step in the first list is not itself being identified with a pure cube-root adjunction.

10.3 Radical extraction as trivializing a torsor

There is a precise cohomological formulation of the remaining ambiguity. For a field F of characteristic not dividing m, the Kummer sequence gives

$$H^1_{\text{ét}}(\operatorname{Spec} F, \mu_m) \simeq F^\times/(F^\times)^m. \tag{32}$$

Here μ_m is the group of m-th roots of unity. A nonzero equation $z^m = a$ defines a μ_m-torsor: over an algebraic closure its solutions form a set on which μ_m acts freely and transitively. Its cohomology class is the class of a modulo m-th powers. Adjoining a solution makes the class trivial. See [5, 6] for the general statements.

For Cardano, after choosing the orientation and adjoining ω, the remaining equation $u^3 = U$ represents the class $[U]$. The equation for v contributes no independent choice because $UV = (-p/3)^3$. When $p \neq 0$, it follows that $[V] = [U]^{-1}$. Reversing the orientation interchanges these classes,

matching the inversion action in $S_3 \simeq C_3 \rtimes C_2$. If $p = 0$, one uses the nonzero coordinate instead of writing a class for zero.

For Ferrari with $q \neq 0$, work over the splitting field F of the resolvent and order its roots. The classes $[U], [V], [W]$ in $F^\times/(F^\times)^2$ satisfy

$$[U][V][W] = [q^2] = 1.$$

At most two independent square classes remain. The equation $uvw = -q$ fixes their compatibility. This is the cohomological expression of the four-element V_4-fiber.

The torsor language identifies exactly what is being trivialized and why one radical is dependent on the others. The invariant equations supply explicit representatives of the classes, so the description remains computational.

Repeated roots lie on the discriminant locus, where the covering maps branch. The polynomial identities and reconstruction formulas still hold, with multiplicities and the zero-coordinate cases handled as above; the unramified-cover interpretation is reserved for the complement of that locus.

10.4 Character lines in families

The field calculations can also vary with parameters. A character component over a field is a vector space, often a line. Over a parameter space, the corresponding component can vary as a line bundle. Multiplication between character components becomes multiplication between line bundles, and choosing a radical becomes a problem of choosing compatible local trivializations.

For example, when n is invertible on a scheme S, a μ_n-torsor is equivalently described by a line bundle together with a trivialization of its n-th tensor power. Locally it has the form

$$z^n = a, \qquad a \in \mathcal{O}_S^\times.$$

Changing the local generator changes a by an n-th power. Thus the familiar ambiguity of a radical is the coordinate expression of descent data. The Kummer sequence records this construction cohomologically; over a field K, it gives

$$H^1_{\text{ét}}(K, \mu_n) \simeq K^\times/(K^\times)^n.$$

This is the geometric version of the eigenline whose n-th power is invariant. The line bundle and its multiplication retain information that disappears if one records only a local equation [5].

11 The same Fourier projectors split an integral

Ferrari's calculation used the characters of the Klein four group to separate four root coordinates. The same operation can separate the differentials on a curve. In a suitable family, it turns a genus-two integration problem into two elliptic ones, and the resulting decomposition also controls point counts over finite fields.

11.1 Two quotient curves from one equation

Choose $a, b \in \mathbb{C}$ such that

$$ab(a-1)(b-1)(a-b) \neq 0, \qquad f(u) = (u-1)(u-a)(u-b),$$

and let C be the smooth projective curve with affine equation

$$C: \qquad y^2 = f(x^2). \tag{33}$$

The map $x : C \to \mathbb{P}^1$ has degree two and branches at the six distinct roots of $f(x^2)$. It is unramified at infinity. Riemann–Hurwitz therefore gives

$$2g(C) - 2 = 2(-2) + 6 = 2, \qquad g(C) = 2.$$

The involutions

$$\sigma(x, y) = (-x, y), \qquad h(x, y) = (x, -y)$$

commute and generate a Klein four group. The second is the hyperelliptic involution. The fixed rational functions for σ are generated by x^2, y; those for σh are generated by x^2, xy. They give the quotient curves and maps

$$E_+ : \quad v^2 = f(u), \qquad \pi_+(x, y) = (u, v) = (x^2, y),$$
$$E_- : \quad w^2 = uf(u), \qquad \pi_-(x, y) = (u, w) = (x^2, xy).$$

Both maps have degree two and extend to the smooth projective models. The cubic defining E_+ and the quartic defining E_- have distinct roots, so both curves have genus one. We take the point at infinity as origin on E_+, and $(0, 0)$ as origin on E_-. These quotient constructions are the classical degree-two elliptic subcovers of an even sextic [13].

11.2 The decomposition reduces the integrals

The two forms

$$\omega_0 = \frac{dx}{y}, \qquad \omega_1 = \frac{x\,dx}{y}$$

are a basis of $H^0(C, \Omega^1_C)$. At a finite branch point, y is a local parameter and the equation shows that dx/y is regular. At infinity put $t = 1/x$ and $Y = t^3 y$. Then

$$Y^2 = (1 - t^2)(1 - at^2)(1 - bt^2), \qquad \omega_0 = -\frac{t\,dt}{Y}, \quad \omega_1 = -\frac{dt}{Y}.$$

Thus both forms are regular there as well. They are linearly independent, and the genus calculation shows that they exhaust the holomorphic differentials.

Now $\sigma^*\omega_0 = -\omega_0$ and $\sigma^*\omega_1 = \omega_1$. The character projectors

$$\frac{1 - \sigma^*}{2}, \qquad \frac{1 + \sigma^*}{2}$$

separate these two lines. More concretely, the quotient maps give

$$\pi_-^*\left(\frac{du}{w}\right) = 2\omega_0, \qquad \pi_+^*\left(\frac{du}{v}\right) = 2\omega_1. \tag{34}$$

The factors of two come simply from $d(x^2) = 2x\,dx$. The equalities are first checked where the displayed coordinates are valid and then hold globally as identities of holomorphic forms.

For any piecewise smooth path γ on C, and any constants $A, B \in \mathbb{C}$, it follows that

$$\boxed{\int_\gamma \frac{(A + Bx)\,dx}{y} = \frac{A}{2}\int_{\pi_-(\gamma)} \frac{du}{w} + \frac{B}{2}\int_{\pi_+(\gamma)} \frac{du}{v}.} \tag{35}$$

The integration problems on the right are elliptic. Specifying the paths on the curves also specifies the necessary continuation of the square roots. The decomposition works for every holomorphic differential in this family, including its periods around closed paths.

This is the same reason for using Fourier coordinates as in Ferrari's formula: a symmetry separates a coupled problem into components with simpler reconstruction rules. Here the components live on quotient curves, so the gain is a reduction in geometric complexity.

11.3 What the periods reconstruct

The decomposition extends to rational cohomology:

$$H^1(C,\mathbb{Q}) = \pi_+^* H^1(E_+,\mathbb{Q}) \oplus \pi_-^* H^1(E_-,\mathbb{Q}). \tag{36}$$

Indeed, the pullbacks are injective because their trace maps compose with them to give multiplication by two. Their images have opposite σ-eigenvalues and each has dimension two. Since $\dim_{\mathbb{Q}} H^1(C,\mathbb{Q}) = 4$, they account for the entire space. Equivalently, one can apply the preceding differential calculation to the holomorphic and antiholomorphic parts of cohomology.

To retain the periods together with their lattice, consider the Jacobian

$$J(C) = H^0(C,\Omega^1_C)^*/H_1(C,\mathbb{Z}),$$

where a cycle acts on differentials by integration. This complex torus is also the algebraic variety $\mathrm{Pic}^0(C)$ of degree-zero line bundles. A map of curves induces pullback on line bundles and a norm map in the other direction. For the present quotients, write

$$T = (\pi_{+*}, \pi_{-*}) : J(C) \longrightarrow E_+ \times E_-, \qquad S(P,Q) = \pi_+^* P + \pi_-^* Q.$$

Here $\mathrm{Pic}^0(E_\pm)$ is identified with $E_\pm$ using the chosen origins.

Proposition 11.1. *The maps S and T are isogenies of degree four, and*

$$T \circ S = [2] \times [2].$$

Proof. For each degree-two quotient, pullback followed by norm, $\pi_{\pm *}\pi_\pm^*$, is multiplication by two. The off-diagonal maps vanish. For example, π_{+*} is unchanged by σ, whereas σ acts on the image of π_-^* by negation: on E_- it sends (u,w) to $(u,-w)$, which is $[-1]$ for the chosen origin. Hence $\pi_{+*}\pi_-^*$ equals its own negative. Its image lies in a finite group of two-torsion points; a homomorphism from a connected elliptic curve to that group is zero. The other off-diagonal map is treated in the same way.

Thus $TS = [2] \times [2]$. This composition is an isogeny, so S and T, between abelian surfaces of the same dimension, are themselves isogenies. Under the canonical principal polarizations of the Jacobians and the product polarization, norm and pullback are dual. Consequently S and T have equal degrees. Their product is the degree of $[2]$ on an abelian surface, namely $2^4 = 16$, so each degree is four. □

An isogeny has finite kernel; it need not be an isomorphism. Thus the elliptic factors reconstruct the Jacobian up to finite information in the period lattice. The curve itself has not become a product. This distinction makes visible the information that the rational character projectors, with their denominators, do not retain integrally.

11.4 The same decomposition counts points

The equations and maps also make sense over an odd finite field $\mathbb{F}_q$, provided that $ab(a-1)(b-1)(a-b) \neq 0$ there. They give the exact identity

$$\boxed{\#C(\mathbb{F}_q) = \#E_+(\mathbb{F}_q) + \#E_-(\mathbb{F}_q) - (q+1).} \tag{37}$$

One can check this directly, without a cohomological trace theorem. Let χ be the quadratic character, extended by $\chi(0) = 0$. The number of solutions of $x^2 = u$ is $1 + \chi(u)$, so

$$\sum_{x \in \mathbb{F}_q} \chi(f(x^2)) = \sum_{u \in \mathbb{F}_q} \chi(f(u)) + \sum_{u \in \mathbb{F}_q} \chi(uf(u)).$$

The affine point counts follow by adding q. The smooth projective models of C and E_- each have two rational points at infinity, because their defining polynomials are monic of even degree; E_+ has one. Adding these points proves (37).

For example, take $a = 2$, $b = 3$ over $\mathbb{F}_{13}$. The two character sums on the right are -6 and 1, respectively. Hence

$$\#E_+(\mathbb{F}_{13}) = 8, \qquad \#E_-(\mathbb{F}_{13}) = 16, \qquad \#C(\mathbb{F}_{13}) = 10 = 8 + 16 - 14.$$

The cohomological explanation is that the decomposition of H^1, now in ℓ-adic cohomology for $\ell \neq \mathrm{char}(\mathbb{F}_q)$, respects Frobenius. Therefore the numerator polynomials of the zeta functions factor as $L_C(T) = L_{E_+}(T)L_{E_-}(T)$. The integral reduction and the point-count identity are two consequences of the same decomposition.

The integral-splitting problem is now complete: the maps and transformed differentials give explicit lower-genus integrals, and the same decomposition controls the stated point counts. A different problem begins in the next section. Given an invertible matrix, can we change its working description so that tensor products and specialization become explicit module operations? The circle provides a finite calculation of those operations. Readers interested only in the integral reduction do not need that further construction.

12 A circle: Fourier transformation beyond eigenvalues

Start with an invertible matrix and the operations we want to perform on it: find fixed vectors, take tensor products, and vary a scalar parameter. Replace the matrix by multiplication by one variable on a module. A Jordan block then becomes a relation involving a power of that variable, and tensor product becomes a concrete substitution in two variables. The circle interprets this calculation as transport. We retain a two-term complex precisely when specializing the parameter would otherwise give the wrong answer.

12.1 Monodromy as a module on the character space

Fix a base point and a positive orientation on S^1. A finite-dimensional complex local system is a vector space V transported consistently along paths; transport around the positive loop is an invertible map $T : V \to V$. Cutting the circle at the base point makes the transport trivial, and gluing its two ends by T reconstructs the local system. A morphism is a linear map commuting with monodromy. Thus the category $\mathrm{Loc}_{\mathrm{fd}}(S^1, \mathbb{C})$ is the category of such pairs (V, T). Since $\pi_1(S^1) = \mathbb{Z}$, no finite-order relation forces T to be diagonalizable.

The rank-one characters are parametrized by $\lambda \in \mathbb{C}^\times$, their monodromy eigenvalue. Put

$$R = \mathbb{C}[z, z^{-1}], \qquad \mathbb{G}_m = \operatorname{Spec} R.$$

An invertible operator on a finite-dimensional vector space defines an R-module by specifying the action of z. Such modules are precisely the finite-length R-modules, or equivalently the coherent sheaves on $\mathbb{G}_m$ with finite support. Here a sheaf is described completely by its module; no gluing calculation is hidden in the notation.

For the Fourier convention used below, let $\mathcal{S}(V, T)$ be V with z acting as T^{-1}. This defines an exact equivalence

$$\mathcal{S} : \mathrm{Loc}_{\mathrm{fd}}(S^1, \mathbb{C}) \xrightarrow{\sim} \mathrm{Coh}_{\mathrm{fin}}(\mathbb{G}_m). \tag{38}$$

The inverse takes multiplication by z^{-1} as monodromy. Both constructions keep all intertwining maps and exact sequences. A rank-one system with monodromy λ becomes the skyscraper module $R/(z - \lambda^{-1})$: a point on the character space. The inverse appears because the transform pairs the input with a character, rather than its inverse.

12.2 The universal kernel and its two-term calculation

The module description is also obtained by an integral-transform construction. Let $\mathcal{U}$ be the local system of free rank-one R-modules on S^1 whose monodromy is multiplication by z. It packages every rank-one character at once: specialization at $z = a$ gives the character with monodromy a. For a local system $\mathcal{L}_T$, form

$$\mathcal{M}(\mathcal{L}_T) = R\Gamma(S^1, \mathcal{L}_T \otimes_{\mathbb{C}} \mathcal{U}).$$

This is the circle version of a Mellin transform: tensor with a universal character and then take cohomology. The corresponding construction on complex algebraic varieties, and its relation to Fourier–Mukai theory, is developed in [15].

Use one vertex and one positively oriented edge for the circle. On its universal cover, the boundary of a lifted edge is its terminal vertex minus its initial vertex. For coefficients with monodromy zT, the cellular cochain differential is therefore $zT - I$. Consequently

$$\mathcal{M}(\mathcal{L}_T) \simeq C_T := [R \otimes_{\mathbb{C}} V \xrightarrow{zT-I} R \otimes_{\mathbb{C}} V], \tag{39}$$

where the two terms have cohomological degrees $0, 1$.

Proposition 12.1. *There is a natural isomorphism*

$$\mathcal{M}(\mathcal{L}_T) \simeq \mathcal{S}(V, T)[-1].$$

Thus $\mathcal{M}[1]$, restricted to finite-dimensional local systems, is the exact equivalence (38).

Proof. The determinant $\det(zT - I)$ is a nonzero Laurent polynomial. Hence $zT - I$ is invertible over $\mathbb{C}(z)$, and is injective on the free R-module $R \otimes V$. In its cokernel, the relation $zTv = v$ says that z acts as T^{-1}. Explicitly, the map

$$R \otimes V \longrightarrow \mathcal{S}(V, T), \qquad \sum_j z^j v_j \longmapsto \sum_j T^{-j} v_j$$

annihilates the image of $zT - I$. Every class reduces uniquely to a constant vector, so it identifies the cokernel with $\mathcal{S}(V, T)$. The complex has cohomology only in degree one, which accounts for the shift $[-1]$. These maps commute with every intertwiner. □

Here the categorical transform is an equivalence on the stated source and target, not an assertion about all constructible sheaves on all spaces. Its kernel, cohomology, and reconstruction can all be calculated from one matrix.

12.3 Thickened points record what a list of characters loses

Compare the monodromies

$$T_0 = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \qquad T_1 = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}.$$

They have the same eigenvalues and the same trace of every integral power. Nevertheless,

$$\dim H^0(S^1, \mathcal{L}_{T_0}) = 2, \qquad \dim H^0(S^1, \mathcal{L}_{T_1}) = 1,$$

because a global flat section is precisely a vector fixed by monodromy. Their spectral modules are

$$\mathcal{S}(V, T_0) \simeq k(1) \oplus k(1), \qquad \mathcal{S}(V, T_1) \simeq R/(z-1)^2,$$

where $k(1) = R/(z-1)$. The second is a length-two thickening of the point 1: multiplication by $z-1$ is nonzero, but its square vanishes. That nilpotent action records the nonsplit extension between the two character lines. It changes an observable cohomology group even though every character trace agrees.

12.4 Tensor product becomes convolution

Let $m : \mathbb{G}_m \times \mathbb{G}_m \to \mathbb{G}_m$ be multiplication. For finite-length modules, define convolution by

$$M \star N = m_*(M \boxtimes N).$$

Concretely, its underlying vector space is $M \otimes_{\mathbb{C}} N$, and z acts as $z_M \otimes z_N$. The supports are finite, so this operation stays within finite-length coherent sheaves and uses ordinary pushforward. Therefore

$$\mathcal{S}(V \otimes W, T \otimes U) \simeq \mathcal{S}(V, T) \star \mathcal{S}(W, U). \tag{40}$$

The unit is the point module $k(1)$. With these operations, (38) is an exact symmetric monoidal equivalence. It preserves extensions and multiplication, as the finite character decomposition did earlier.

For unipotent monodromy, the logarithm $N = \log T$ is a finite polynomial in $T - I$. In the formal coordinate $s = -\log z$ at 1, the spectral action is $s = N$. Multiplication of spectral parameters becomes addition of these coordinates, so tensor product is governed by

$$N_{V \otimes W} = N_V \otimes I + I \otimes N_W.$$

For two nonsplit rank-two systems, the calculation becomes

$$\mathbb{C}[x, y]/(x^2, y^2), \qquad s \cdot f = (x + y)f.$$

There is a length-three chain $1 \mapsto x + y \mapsto 2xy \mapsto 0$, and the independent vector $x - y$ is killed by s. The tensor product therefore has unipotent Jordan blocks of sizes three and one. Below, Fourier–Mukai duality will turn this same calculation into a decomposition of vector bundles on an elliptic curve. The passage to a logarithmic coordinate uses characteristic zero.

The companion *When a Formal Logarithm Fails* studies the obstruction to a universal logarithmic coordinate in characteristic three: the quadratic correction succeeds, and the next extension is blocked by a nonzero spectral sequence differential.

12.5 Why taking a fibre must retain a complex

The trivial rank-one local system gives an especially small example:

$$\mathcal{M}(\mathcal{L}_I) = [R \xrightarrow{z-1} R] \simeq k(1)[-1]. \tag{41}$$

Specialize the free complex at $z = a$. It becomes

$$[\mathbb{C} \xrightarrow{a-1} \mathbb{C}].$$

For $a \neq 1$, its cohomology vanishes. At $a = 1$, its differential is zero, so its H^0 and H^1 are both $\mathbb{C}$, as required for the ordinary cohomology of a circle. This operation is the derived fibre

$$\mathcal{M}(\mathcal{L}_I) \otimes^L_R k(a).$$

The ordinary fibre of the single module $k(1)[-1]$ would retain only degree one. The missing degree zero is supplied by $\mathrm{Tor}^R_1(k(1), k(1)) \simeq \mathbb{C}$, followed by the shift. Keeping a free resolution before specialization retains precisely the information needed at the exceptional character.

The module and the free two-term complex have answered different parts of the matrix problem. The module makes tensor calculations explicit; the complex gives the correct exceptional specialization. These are completed calculations. The next section addresses holomorphic bundles, with different input objects and a different parameter space. Rank-one complex local systems on an elliptic curve are parametrized by $(\mathbb{C}^\times)^2$, whereas its degree-zero holomorphic line bundles form $\mathrm{Pic}^0(E)$. A theorem is needed to transfer the bundle operations to that latter space; the circle computation alone does not supply it.

13 Fourier–Mukai: reconstructing global bundles

This section separates two tasks. First, construct a calculus that decomposes a tensor product of nilpotent operators into explicit chains, and a terminating arithmetic rule for choosing transformations of rank and degree. Second, state the geometric theorems that let these calculations act on bundles over an elliptic curve. The chain construction in Section 13.3 is complete on its own. Its bundle application needs the Fourier–Mukai equivalence; the rank-and-degree procedure also needs its stability theorem. The preliminary bundle examples specify why those imported operations must retain extensions.

13.1 A question that requires extension data

Before constructing a transform, it is useful to see what the transformed data must distinguish. On a smooth complex elliptic curve E, compare the split bundle $\mathcal{O}_E \oplus \mathcal{O}_E$ with a nonsplit extension

$$0 \longrightarrow \mathcal{O}_E \longrightarrow F_2 \longrightarrow \mathcal{O}_E \longrightarrow 0. \tag{42}$$

Both have two trivial line bundles as successive quotients, and both have class $2[\mathcal{O}_E]$ in the ordinary Grothendieck group $K_0(E)$. Nevertheless, they have different numbers of global sections.

The beginning of the cohomology exact sequence is

$$0 \longrightarrow \mathbb{C} \longrightarrow H^0(E, F_2) \longrightarrow \mathbb{C} \xrightarrow{\delta} \mathbb{C} \longrightarrow H^1(E, F_2) \longrightarrow \mathbb{C} \longrightarrow 0.$$

The connecting map δ takes 1 to the extension class in $\mathrm{Ext}^1(\mathcal{O}_E, \mathcal{O}_E) = H^1(E, \mathcal{O}_E) \simeq \mathbb{C}$. The extension is nonsplit, so δ is a nonzero map between one-dimensional spaces and is therefore an isomorphism. Exactness gives

$$h^0(E, F_2) = h^1(E, F_2) = 1, \qquad h^0(E, \mathcal{O}_E \oplus \mathcal{O}_E) = 2. \tag{43}$$

Here h^i denotes the dimension of H^i. A record of the two constituent line bundles cannot recover this elementary invariant. A useful Fourier transform must retain how the constituents are attached. The distinction will become a small piece of geometry on the dual curve: a doubled point with a nonzero nilpotent direction, rather than two independent copies of a point.

13.2 A character space carrying sheaves

The circle calculation replaced a list of eigenvalues by a sheaf that retained nilpotent structure. For holomorphic bundles, a geometric transform must likewise preserve extensions and their variation. Its dual parameter space will now consist of line bundles, and its kernel will be their universal family. Let E be a smooth complex elliptic curve with origin 0, and let

$$\widehat{E} = \mathrm{Pic}^0(E)$$

be its dual elliptic curve. A point of $\widehat{E}$ specifies a degree-zero line bundle on E. Tensor product supplies the group law on this parameter space, just as multiplication supplies the group law on ordinary characters.

The universal family is the normalized Poincaré line bundle $\mathcal{P}$ on $E \times \widehat{E}$. It is the geometric kernel replacing the scalar pairing $\chi(x)$. With p_E and $p_{\widehat{E}}$ the two projections, define

$$\Phi_{\mathcal{P}}(F) = Rp_{\widehat{E}*}(Lp_E^* F \otimes^L \mathcal{P}). \tag{44}$$

The construction has the same three operations as an integral transform: move the input to the product, multiply by the kernel, and integrate along the original variable. Here integration is derived pushforward, which retains all cohomological degrees rather than only ordinary sections.

Mukai's theorem says that this functor is an equivalence

$$D^b \mathrm{Coh}(E) \simeq D^b \mathrm{Coh}(\widehat{E}).$$

After identifying the elliptic curve with its dual by the origin, and using the compatible normalized Poincaré kernel, its elementary transforms are

$$\Phi(k(x)) \simeq \mathcal{P}_x, \qquad \Phi(\mathcal{P}_x) \simeq k(-x)[-1], \qquad \Phi^2 \simeq (-1_E)^*[-1]. \tag{45}$$

Here $k(x)$ is the skyscraper sheaf at x. A point mass becomes a character-like line bundle; applying the transform twice gives reflection, together with a cohomological shift. These are precise categorical counterparts of elementary Fourier identities [8].

The shift matters. For instance, a nontrivial degree-zero line bundle L has no cohomology: a nonzero section would trivialize it, and Serre duality gives the same conclusion for $H^1(E, L)$. The trivial bundle has both H^0 and H^1. Derived pushforward organizes this jump into the object $k(0)[-1]$. Its derived fibre at the origin has cohomology in degrees zero and one: the skyscraper sheaf has a nonzero Tor_1 with itself. Replacing it by a list of ordinary fibre dimensions would lose the gluing and derived structure at the exceptional point.

If $m : E \times E \to E$ is addition, define convolution by

$$F \star G = Rm_*(F \boxtimes^L G).$$

The multiplicativity of the Poincaré kernel yields

$$\Phi(F \star G) \simeq \Phi(F) \otimes^L \Phi(G). \tag{46}$$

Thus convolution becomes tensor product as a functorial identity. It applies to objects and maps, including extension data, rather than merely to numerical functions [8].

13.3 A chain calculus for tensor products

The calculation to perform is independent of elliptic curves. Given positive integers r, s, form

$$M_{r,s} = \mathbb{C}[x, y]/(x^r, y^s), \qquad E(f) = (x + y)f.$$

We want a basis made of chains $v, Ev, E^2v, \ldots$, with the length of each chain explicit. This computes a tensor product of nilpotent operators: multiplication by x acts on the first factor, multiplication by y on the second, and their sum acts on the tensor product. We will construct the basis, rather than infer its existence from a representation-theoretic decomposition.

Build an operation that runs a chain backwards. First work in one factor, $\mathbb{C}[x]/(x^r)$, with the forward move $E_r x^i = x^{i+1}$. Seek a backward move $F_r x^i = c_i x^{i-1}$. The endpoint rules require $c_0 = c_r = 0$: the bottom cannot move down, and the zero relation $x^r = 0$ must remain zero. We want the difference between a forward–backward move and a backward–forward move to depend only on the current degree, decreasing by two after each forward step. Thus write

$$c_{i+1} - c_i = A - 2i \qquad (0 \le i < r).$$

Summing over all r steps and using $c_r - c_0 = 0$ forces $A = r - 1$. Summing only the first i steps now determines every coefficient:

$$c_i = \sum_{j=0}^{i-1} (r - 1 - 2j) = i(r - i).$$

The backward move has been derived from its endpoint and commutator requirements. On the two factors, add their backward moves and their degree weights. This gives

$$F(x^i y^j) = i(r - i)x^{i-1}y^j + j(s - j)x^i y^{j-1}, \qquad H(x^i y^j) = \big(2(i + j) - (r + s - 2)\big)x^i y^j.$$

Terms with a zero coefficient are omitted. The endpoint rules make these well-defined on the truncated polynomial space. Moves on different factors commute, so the one-factor calculation gives

$$FE - EF = -H, \qquad HE - EH = 2E.$$

These two identities are all the operator algebra the proof needs. The chosen drop by two is a normalization; its purpose is the reverse-step formula that will determine the chain lengths below.

Choose the start of each chain by cancellation. Put $h = \min(r, s)$. For $0 \le k < h$, seek a homogeneous vector

$$v_k = \sum_{j=0}^{k} a_j x^j y^{k-j}, \qquad a_0 = 1,$$

with $Fv_k = 0$. Comparing the coefficient of $x^j y^{k-j-1}$ gives

$$a_{j+1}(j+1)(r-j-1) + a_j(k-j)(s-k+j) = 0.$$

Thus the desired vector is constructed by the finite recurrence

$$a_{j+1} = -a_j \frac{(k-j)(s-k+j)}{(j+1)(r-j-1)}, \qquad 0 \le j < k.$$

Its denominators are nonzero: $j < k < r$. The normalization makes $v_k \neq 0$, since its y^k-coefficient is one. We have changed a search for a Jordan basis into a triangular coefficient calculation.

Determine the length, without guessing a normal form. Set $m_k = r + s - 2 - 2k$. Since $Hv_k = -m_k v_k$, the two operator identities give, by induction on a,

$$FE^a v_k = a(m_k - a + 1)E^{a-1}v_k \qquad (a \ge 1).$$

Multiplication by $x + y$ is nilpotent because it raises total degree. Let q be the largest exponent with $E^q v_k \neq 0$. Applying the identity to $E^{q+1}v_k = 0$ yields

$$0 = (q+1)(m_k - q)E^q v_k.$$

Over $\mathbb{C}$, this forces $q = m_k$. The constructed chain therefore has exactly

$$v_k, Ev_k, \ldots, E^{m_k}v_k, \qquad m_k + 1 = r + s - 1 - 2k$$

nonzero members.

Verify that the chains form a basis. In a relation of total degree d, the possible chain vectors are $E^{d-k}v_k$. Choose the least k whose coefficient is nonzero and apply F^{d-k}. Every term with a larger index is killed: it has fewer than $d-k$ forward steps before reaching its vector annihilated by F. The chosen term becomes a nonzero scalar times v_k, since all factors in the displayed backward-step formula are nonzero along its chain. This is a contradiction. Homogeneous components separate relations of different degrees, so all the chain vectors are independent. Finally,

$$\sum_{k=0}^{h-1}(r + s - 1 - 2k) = h(r + s - h) = rs = \dim M_{r,s}.$$

They form a basis, and with t acting as E the complete answer is

$$M_{r,s} \simeq \bigoplus_{k=0}^{h-1} \mathbb{C}[t]/(t^{r+s-1-2k}).$$

For example, when $(r, s) = (2, 3)$, the recurrence gives $v_0 = 1$ and $v_1 = y - 2x$. The two chains are

$$1 \longmapsto x + y \longmapsto 2xy + y^2 \longmapsto 3xy^2 \longmapsto 0,$$

$$y - 2x \longmapsto y^2 - xy \longmapsto 0.$$

Their lengths are four and two. The coefficients also show exactly why the proof uses characteristic zero. For instance, in characteristic two, multiplication by $x + y$ on $\mathbb{F}_2[x, y]/(x^2, y^2)$ has square zero; the characteristic-zero lengths three and one no longer apply.

Transport the answer to elliptic bundles, using a theorem. Call a vector bundle on E unipotent if it has a filtration with successive quotients $\mathcal{O}_E$. The geometric input is that $\Psi = \Phi[1]$ restricts to the following equivalence and tensor rule:

$$\Psi : \mathrm{Unip}(E) \xrightarrow{\sim} \mathrm{Coh}_0(\widehat{E}), \qquad \Psi(F \otimes G) \simeq \Psi(F) \star \Psi(G), \tag{47}$$

where the target consists of coherent sheaves supported at the origin. This is a consequence of the elliptic Fourier–Mukai theorem, not of the chain calculation; the precise compatibility is discussed by Schröer [7]. The formal neighbourhood of the origin has ring $\mathbb{C}[[t]]$, so its finite-length modules are nilpotent operators. Their indecomposables and inverse transforms satisfy

$$\Psi(F_r) \simeq \mathbb{C}[[t]]/(t^r), \qquad F_1 = \mathcal{O}_E. \tag{48}$$

Over $\mathbb{C}$, a formal logarithm changes the group law of $\widehat{E}$ to addition. Convolution then acts on the tensor product of two modules by $t = x + y$. Consequently the chain algorithm above computes $\Psi(F_r \otimes F_s)$. For $r = s = 2$, it gives lengths three and one, and the inverse equivalence gives

$$\boxed{F_2 \otimes F_2 \simeq F_3 \oplus \mathcal{O}_E.} \tag{49}$$

For arbitrary positive r, s, the constructed basis proves

$$F_r \otimes F_s \simeq \bigoplus_{i=1}^{\min(r,s)} F_{r+s+1-2i}. \tag{50}$$

The local calculation supplies the summands and their lengths. The geometric theorem supplies the right to transport them to global bundles. Both pieces are necessary, and their jobs are different.

13.4 Rank and degree: a terminating reduction

The next calculation chooses a sequence of operations that reduces a stable bundle to a line bundle. For a nonzero vector bundle V, write $(r, d) = (\operatorname{rank} V, \deg V)$ and $\mu(V) = d/r$. Stability means that every nonzero proper subbundle of smaller rank has strictly smaller slope. We assume $\gcd(r, d) = 1$, and keep the actual bundle together with its pair of integers throughout the calculation.

There are two available operations. Tensoring by the specified line bundle $\mathcal{O}_E(m \cdot 0)$ changes (r, d) to $(r, d + mr)$ and preserves stability. The normalized Fourier–Mukai transform has numerical action

$$(r, d) \longmapsto (d, -r). \tag{51}$$

To verify this action, use the numerical Grothendieck group, recorded by rank and degree with basis $[\mathcal{O}_E]$, $[k(0)]$. Riemann–Roch gives its Euler pairing

$$\chi(U, V) = \operatorname{rank}(U) \deg(V) - \deg(U) \operatorname{rank}(V).$$

An equivalence preserves the pairing. The identities $\Phi(\mathcal{O}_E) = k(0)[-1]$ and $\Phi(k(0)) = \mathcal{O}_E$ therefore send the basis vectors to $(0, -1)$ and $(1, 0)$, proving the formula. For complexes these are alternating ranks and degrees; the minus sign records the shift.

Specify when a numerical step is an operation on bundles. For a stable bundle of positive degree, the elliptic transform is a stable vector bundle in cohomological degree zero. For a stable bundle of negative degree, its only nonzero cohomology sheaf is in degree one and is stable [14]. These are geometric inputs to the algorithm.

Part of the positive-degree assertion has a short explanation. For every degree-zero line bundle P, slope comparison and Serre duality give

$$H^1(E, V \otimes P) \simeq \operatorname{Hom}(V \otimes P, \mathcal{O}_E)^* = 0.$$

Riemann–Roch then gives $h^0(E, V \otimes P) = d$. Cohomology and base change assemble these spaces into a rank-d bundle. Its stability is the additional theorem. Integer arithmetic alone cannot establish that assertion.

Choose the next operation from the remainder. At the current pair (r, d), take

$$q = \left\lfloor \frac{d}{r} \right\rfloor, \qquad m = -q, \qquad e = d + mr, \qquad 0 \le e < r.$$

Tensor the current bundle by $\mathcal{O}_E(m \cdot 0)$, and record this exact twist. If $e = 0$, stop. Otherwise apply Φ, record it, and repeat with

$$(r, d) \longmapsto (r, e) \longmapsto (e, -r).$$

The input to every forward transform has positive degree, so the geometric theorem guarantees that every current object is a stable bundle. Both operations preserve the greatest common divisor. Every transform strictly decreases the positive rank. Hence the procedure terminates. At termination, $\gcd(r, 0) = 1$ forces $r = 1$, and the current object is a degree-zero line bundle L.

This is a complete rule for selecting the operations, including negative initial degrees. It specifies a finite word in twists and transforms from (r, d), with no search over bundles or possible decompositions. Evaluating a Fourier–Mukai transform on a presented bundle is still a geometric computation; the arithmetic schedule does not replace it.

Run the rule on rank five and degree two. The nonzero steps are

$$\begin{aligned}
(5, 2) &\xrightarrow{\Phi} (2, -5), \\
(2, -5) &\xrightarrow{\otimes \mathcal{O}_E(3 \cdot 0)} (2, 1), \\
(2, 1) &\xrightarrow{\Phi} (1, -2), \\
(1, -2) &\xrightarrow{\otimes \mathcal{O}_E(2 \cdot 0)} (1, 0).
\end{aligned}$$

Thus

$$\mathcal{T}(V) = \Phi\big(\Phi(V) \otimes \mathcal{O}_E(3 \cdot 0)\big) \otimes \mathcal{O}_E(2 \cdot 0) = L.$$

The line bundle itself must be retained. The final numerical pair $(1, 0)$ cannot distinguish the different points of $\operatorname{Pic}^0(E)$.

Reverse the recorded word. Undo a twist by tensoring with its inverse, and undo a transform using

$$\Phi^{-1} \simeq (-1_E)^* \Phi[1].$$

After reversing the twist immediately following a forward transform, the bundle has the negative degree that transform produced. The negative-degree theorem places its transform in degree one; the shift in Φ^{-1} cancels that degree. Thus reversing the word stays within stable bundles and recovers the original object, not merely its numerical class.

Conversely, starting with any degree-zero line bundle and reversing the fixed word for a given coprime pair produces a stable bundle with that pair. The forward and reverse functors are inverse. This gives a bijection between the isomorphism classes of stable bundles of that rank and degree and $\operatorname{Pic}^0(E)$, relative to the recorded choices of origin, kernel and twists. The terminating remainder calculation chooses the word; the elliptic transform theorem makes the word invertible on the stated objects.

14 Another cubic: from Fourier projectors to counting and periods

The same projectors can do something quite different from extracting roots: they can count solutions. This opens a second geometric branch, in which the spectral operator is Frobenius and its traces count solutions over finite fields. We will begin with a five-term calculation, recover the number of points on a curve over every finite extension, and then identify a categorical statement whose numerical consequence is a character-sum identity.

14.1 A projector that counts cube roots

Consider the smooth projective cubic

$$E: \quad X^3 + Y^3 = Z^3$$

over $\mathbb{F}_7$. It has genus one and a rational point. Let $\omega = e^{2\pi i/3}$, choose the cubic character $\chi : \mathbb{F}_7^\times \to \mu_3$ by $\chi(3) = \omega$, and put $\chi(0) = 0$. The element 3 generates $\mathbb{F}_7^\times$; in the order $1, 2, 3, 4, 5, 6$, the character values are

$$1, \quad \omega^2, \quad \omega, \quad \omega, \quad \omega^2, \quad 1.$$

For every $t \in \mathbb{F}_7$, including zero,

$$A(t) := \#\{x \in \mathbb{F}_7 : x^3 = t\} = 1 + \chi(t) + \chi(t)^2. \tag{52}$$

For nonzero t, the right side is three when t is a cube and zero otherwise. At zero it is one. This is character orthogonality serving as an exact indicator, rather than as a change of coordinates.

The affine point count is the convolution

$$N_{\text{aff}} = \sum_{t \in \mathbb{F}_7} A(t)A(1-t).$$

Define the Jacobi sum

$$J(\rho, \eta) = \sum_{t \in \mathbb{F}_7} \rho(t)\eta(1-t),$$

where the nontrivial characters are extended by zero. The terms linear in a nontrivial character sum to zero. Also

$$J(\chi, \chi^{-1}) = -\chi(-1) = -1 :$$

indeed, $t \mapsto t/(1-t)$ runs through all field elements except -1, with the zero term contributing nothing. Writing $J = J(\chi, \chi)$, expansion therefore gives

$$N_{\text{aff}} = 7 + J + \overline{J} - 2.$$

At infinity, $Z = 0$, there are three points: their ratios X/Y are the three cube roots of -1. Thus

$$\#E(\mathbb{F}_7) = 8 + J + \overline{J}. \tag{53}$$

The correction at infinity matters: a character sum initially counts points in an affine chart, not automatically on its projective completion.

Only $t = 2, 3, 4, 5, 6$ contribute to J. Their contributions are, respectively,

$$\omega^2, \quad 1, \quad \omega^2, \quad 1, \quad \omega^2.$$

Consequently

$$J = 2 + 3\omega^2 = -1 - 3\omega, \qquad J + \overline{J} = 1, \qquad J\overline{J} = 7. \tag{54}$$

Equation (53) gives nine points: six affine points and three at infinity.

14.2 Why this calculation controls every extension field

Fix a prime $\ell \neq 7$ and use ℓ-adic cohomology with coefficients enlarged to contain the character values. Geometric Frobenius is normalized here to act by 7 on $H^2(E_{\overline{\mathbb{F}}_7}, \overline{\mathbb{Q}}_\ell)$. The trace formula and the cohomology of a smooth projective genus-one curve give

$$\#E(\mathbb{F}_{7^r}) = 1 + 7^r - \mathrm{Tr}\left(\mathrm{Fr}^r \mid H^1(E_{\overline{\mathbb{F}}_7}, \overline{\mathbb{Q}}_\ell)\right). \tag{55}$$

Here H^1 is two-dimensional and the determinant of Frobenius on it is 7. Since the point count is nine, its trace is -1. Its characteristic polynomial is therefore

$$T^2 + T + 7.$$

By (54), the two roots are

$$\alpha = -J = 1 + 3\omega, \qquad \beta = -\overline{J} = 1 + 3\omega^2. \tag{56}$$

The general Jacobi-sum description of Fermat curves explains this identification through the action

$$(a, b)[X : Y : Z] = [aX : bY : Z], \qquad (a, b) \in \mu_3^2.$$

Its two nonzero character components in H^1 are lines. Frobenius preserves them, and the corresponding scalars are the two negatives of the Jacobi sums. The character-projector calculation has found the spectrum of a geometric operator [16].

Taking powers of this one operator now gives

$$\boxed{\#E(\mathbb{F}_{7^r}) = 7^r + 1 - \alpha^r - \beta^r.} \tag{57}$$

No diagonalization choices are needed for computation. Put $s_r = \alpha^r + \beta^r$. Then

$$s_0 = 2, \qquad s_1 = -1, \qquad s_r = -s_{r-1} - 7s_{r-2} \quad (r \geq 2).$$

This gives $s_2 = -13$, $s_3 = 20$, and $s_4 = 71$, hence

$$\#E(\mathbb{F}_7) = 9, \quad \#E(\mathbb{F}_{49}) = 63, \quad \#E(\mathbb{F}_{343}) = 324, \quad \#E(\mathbb{F}_{2401}) = 2331.$$

Direct finite-field enumeration confirms these four values. The recurrence itself determines infinitely many further counts. Equivalently, the zeta function is

$$Z(E, T) := \exp\left(\sum_{r \geq 1} \#E(\mathbb{F}_{7^r}) \frac{T^r}{r}\right) = \frac{1 + T + 7T^2}{(1 - T)(1 - 7T)}. \tag{58}$$

Moreover $|\alpha| = |\beta| = \sqrt{7}$, so

$$|\#E(\mathbb{F}_{7^r}) - (7^r + 1)| \leq 2\,7^{r/2}. \tag{59}$$

For this special curve the norm identity is already visible in the five-term sum. Cohomology explains why the same two numbers govern all extension fields, and places the calculation inside a theory that applies to curves without this convenient character formula.

14.3 The same symmetry selects a period

Now consider the same equation over $\mathbb{C}$. In the chart $Z = 1$, the differential

$$\nu = \frac{dx}{y^2}, \qquad x^3 + y^3 = 1,$$

extends to a holomorphic differential on E. For example, near $y = 0$ the relation $x^2dx + y^2dy = 0$ rewrites it as $-dy/x^2$. Under $(x, y) \mapsto (ax, by)$, with $a^3 = b^3 = 1$, it satisfies

$$(a, b)^*\nu = ab\,\nu.$$

Thus this differential belongs to a definite character line, just as a Fourier coordinate does.

Let γ be the real arc from $(0, 1)$ to $(1, 0)$. Substituting $t = x^3$ gives the explicit integral

$$\int_\gamma \nu = \int_0^1 \frac{dx}{(1 - x^3)^{2/3}} = \frac{1}{3}\int_0^1 t^{-2/3}(1 - t)^{-2/3}\,dt = \frac{1}{3}B\left(\frac{1}{3}, \frac{1}{3}\right). \tag{60}$$

Both endpoint singularities in the last integral are integrable. To obtain a closed cycle, let $\sigma(x, y) = (\omega x, y)$ and $\tau(x, y) = (x, \omega y)$, and set

$$\delta = (1 - \sigma_*)(1 - \tau_*)\gamma.$$

Its boundary vanishes: τ fixes $(1, 0)$ and σ fixes $(0, 1)$. Since $\sigma^*\nu = \tau^*\nu = \omega\nu$, its period is

$$\int_\delta \nu = \frac{(1 - \omega)^2}{3}B\left(\frac{1}{3}, \frac{1}{3}\right). \tag{61}$$

The same algebraic family has therefore supplied a character sum in finite characteristic and a special-function integral over $\mathbb{C}$. These are different realizations of its symmetry, with different additional structures: Frobenius on one side, and the comparison between algebraic differential forms and topological cycles on the other.

14.4 Reference: a motivic lift of a finite identity

The finite identity has its own complete calculation. The second half of this subsection records a motivic theorem for reference; the identity alone does not prove that theorem. Over $\mathbb{F}_q$, choose a nontrivial additive character ψ, and for a nontrivial multiplicative character ρ put

$$G_\psi(\rho) = \sum_{u \neq 0} \rho(u)\psi(u).$$

Use the same definition of J, now summing over $\mathbb{F}_q$. If $\rho, \eta, \rho\eta$ are nontrivial, then

$$G_\psi(\rho)G_\psi(\eta) = J(\rho, \eta)G_\psi(\rho\eta). \tag{62}$$

To prove it, expand the product and partition by $s = u + v$. The contribution at $s = 0$ vanishes because $\rho\eta$ is nontrivial. For $s \neq 0$, set $u = st$, $v = s(1 - t)$: the inner sum becomes $(\rho\eta)(s)J(\rho, \eta)$, leaving exactly the Gauss sum on the right. This is a concrete convolution-to-product calculation.

There is an actual categorical lift. We state the two-character case of a theorem of Otsubo and Yamazaki [17]. Let $q = p^f$, let $d \mid q - 1$, and work in the tensor category of Chow motives in the covariant convention over $\mathbb{F}_q$, with coefficient field $\Lambda = \mathbb{Q}(\zeta_{pd})$. For the objects needed here, a motive

is specified by a smooth projective variety together with an idempotent algebraic correspondence. A finite group H acting on X supplies such a correspondence:

$$e_{\vartheta}^{X} = \frac{1}{|H|} \sum_{h \in H} \vartheta(h)^{-1} [\Gamma_h].$$

The graph $\Gamma_h \subset X \times X$ plays the role of the linear operator h. The image $h(X)^{\vartheta}$ is its character summand. This is the same projector formula, applied to geometry before taking a cohomology realization.

Let A_d be the smooth projective model of

$$u^q - u = v^d,$$

on which $\mathbb{F}_q \times \mu_d$ acts by $(c, a)(u, v) = (u + c, av)$. Let F_d be the smooth projective curve $U^d + V^d = W^d$, with the coordinatewise μ_d^2-action. For characters $\theta, \theta_1, \theta_2$ of μ_d, set

$$M_{\psi,\theta} = h(A_d)^{(\psi,\theta)}, \qquad N_{\theta_1,\theta_2} = h(F_d)^{(\theta_1,\theta_2)},$$

extending coefficients to Λ throughout.

Theorem 14.1 (Otsubo–Yamazaki, two-character case)**.** *If ψ, θ_1, θ_2, and $\theta_1\theta_2$ are nontrivial, then these character summands are invertible motives and there is an isomorphism*

$$M_{\psi,\theta_1} \otimes M_{\psi,\theta_2} \simeq M_{\psi,\theta_1\theta_2} \otimes N_{\theta_1,\theta_2}. \tag{63}$$

Here invertible means invertible for tensor product. Such an object's endomorphisms are scalars in Λ. To compare conventions, write

$$\rho_i(u) = \theta_i\left(u^{(q-1)/d}\right).$$

The Frobenius endomorphisms of the displayed factors are, respectively,

$$-G_\psi(\rho_1), \quad -G_\psi(\rho_2), \quad -G_\psi(\rho_1\rho_2), \quad -J(\rho_1, \rho_2).$$

The authors use signed Gauss and Jacobi sums, $-G$ and $-J$, which are the Frobenius scalars on these degree-one summands. Frobenius commutes with the motivic isomorphism, so equality of its scalar actions gives

$$(-G_\psi(\rho_1))(-G_\psi(\rho_2)) = (-G_\psi(\rho_1\rho_2))(-J(\rho_1, \rho_2)),$$

recovering (62). For our example take $q = 7$, $d = 3$, $\psi(u) = e^{2\pi i u/7}$, and choose θ so that $\theta(u^2) = \chi(u)$. Then $\theta_1 = \theta_2 = \theta$ satisfies every hypothesis.

The finite change of variables proves the Gauss–Jacobi identity above. The stronger motivic statement requires a construction of algebraic correspondences and a proof that they induce the asserted isomorphism; those are supplied by the cited work, not derived here. It can therefore be used as a theorem-dependent transport result, with its hypotheses and conventions retained. It is not an additional stage required to complete the character-sum calculation.

15 A quantitative shadow: where the square-root bound comes from

The Fermat cubic supplied a cancellation estimate through two explicitly computed eigenvalues. Geometric Fourier theory can produce such estimates when an elementary evaluation of the sum is unavailable. We will also see how local information determines the number of eigenvalues, and hence the constant in the estimate.

Let $\mathbb{F}_q$ have characteristic p, let $\psi : \mathbb{F}_q \to \mathbb{C}^\times$ be a nontrivial additive character, and take $a \in \mathbb{F}_q^\times$. The Kloosterman sum is

$$K(a) = \sum_{x \in \mathbb{F}_q^\times} \psi\Big(x + \frac{a}{x}\Big). \tag{64}$$

It has $q-1$ summands of absolute value one. The geometric result is

$$\boxed{|K(a)| \le 2\sqrt{q}.} \tag{65}$$

For example, over $\mathbb{F}_5$ with $\psi(u) = e^{2\pi i u/5}$, direct evaluation at $a = 2$ gives

$$K(2) = 2e^{4\pi i/5} + 2e^{6\pi i/5} = -1 - \sqrt{5}.$$

The objective is a uniform bound for every q and every nonzero a, without evaluating each sum separately.

15.1 The transform retains a family of cohomology groups

The first Fourier step is elementary. Define $g(0) = 0$ and $g(t) = \psi(1/t)$ for $t \neq 0$. For the unnormalized transform

$$\widehat{g}(a) = \sum_{t \in \mathbb{F}_q} g(t)\psi(at),$$

the substitution $x = at$ gives $\widehat{g}(a) = K(a)$ when $a \neq 0$.

Choose $\ell \neq p$, realize the character values in $\overline{\mathbb{Q}}_\ell$, and fix an embedding into $\mathbb{C}$ for absolute values. Use the Artin–Schreier sheaf $\mathcal{L}_\psi$ normalized so that geometric Frobenius has trace $\psi(u)$ at $u \in \mathbb{F}_q$. Let $j : \mathbb{G}_m \hookrightarrow \mathbb{A}^1_t$ and put

$$\mathcal{G} = j_!\mathcal{L}_\psi(1/t),$$

an ordinary sheaf placed in cohomological degree zero. Its trace function is g. The Fourier–Deligne transform on constructible derived categories is

$$\mathrm{FT}_\psi(\mathcal{G}) = Rp_{a!}\big(p_t^*\mathcal{G} \otimes \mathcal{L}_\psi(at)\big)[1], \tag{66}$$

where p_t, p_a project from $\mathbb{A}^1_t \times \mathbb{A}^1_a$. Tensoring multiplies the trace functions; compactly supported pushforward sums them along the t-line. The shift [1] changes the sign of the alternating trace. Thus the transformed object's trace at a is $-K(a)$ [10].

More precisely, for $a \neq 0$, its derived stalk is

$$R\Gamma_c\Big(\mathbb{G}_{m,\overline{\mathbb{F}}_q}, \mathcal{F}_a\Big)\,[1], \qquad \mathcal{F}_a = \mathcal{L}_\psi(x + a/x). \tag{67}$$

We will show that the unshifted compactly supported cohomology occurs only in degree one and has dimension two. The shift then places this vector space in degree zero. As a varies through $\mathbb{G}_m$, these spaces form the rank-two Kloosterman sheaf [9]. Its fibers retain Frobenius actions and monodromy, not merely the single number $K(a)$.

15.2 Two poles explain the dimension two

We compute on the x-line, with $a \neq 0$ fixed. The local calculation at each pole can be made explicit. Write the additive character as $\psi(u) = \exp(2\pi i \operatorname{Tr}(bu)/p)$ for some $b \in \mathbb{F}_q^\times$. Over the algebraic closure, its local Artin–Schreier cover at a simple pole has the form

$$y^p - y = c/t + \text{regular terms}, \qquad c \neq 0.$$

A Laurent series of the form $z^p - z$ with a pole has pole order divisible by p, so this equation defines a nontrivial degree-p extension of the local field. Its ramification is total: the residue field is algebraically closed. With the normalized valuation on the extension, $v(t) = p$, and comparison of the pole orders in the equation gives $v(y) = -1$. Thus $u = 1/y$ is a uniformizer.

The nonidentity automorphisms are $\sigma_d(y) = y + d$, where $d \in \mathbb{F}_p^\times$. Direct calculation gives

$$\sigma_d(u) - u = -\frac{d}{y(y+d)}, \qquad v\big(\sigma_d(u) - u\big) = 2.$$

By the definition of the lower ramification groups, this says $G_0 = G_1 = \mathbb{F}_p$ and $G_2 = 0$. A nontrivial one-dimensional character therefore has Swan conductor one. At 0, the phase has a simple pole with nonzero coefficient ba; at infinity it has a simple pole with nonzero coefficient b. Hence

$$\operatorname{Swan}_0(\mathcal{F}_a) = \operatorname{Swan}_\infty(\mathcal{F}_a) = 1.$$

The Grothendieck–Ogg–Shafarevich formula gives

$$\begin{aligned}\chi_c(\mathbb{G}_m, \mathcal{F}_a) &= \operatorname{rank}(\mathcal{F}_a)\,\chi_c(\mathbb{G}_m, \overline{\mathbb{Q}}_\ell) - \operatorname{Swan}_0(\mathcal{F}_a) - \operatorname{Swan}_\infty(\mathcal{F}_a) \\ &= 1 \cdot 0 - 1 - 1 = -2.\end{aligned}$$

The Euler characteristic of the constant sheaf on $\mathbb{G}_m$ is zero: removing two points from $\mathbb{P}^1$ subtracts two from its Euler characteristic two.

There are no nonzero compactly supported global sections of a lisse sheaf on this connected nonproper curve, so $H^0_c = 0$. Poincaré duality identifies H^2_c with geometric coinvariants, with a Tate twist. These vanish because the rank-one local system has nontrivial geometric monodromy. Consequently

$$\dim H^1_c(\mathbb{G}_m, \mathcal{F}_a) = 2. \tag{68}$$

These are the local-to-global dimension calculations behind the constant in (65); see [18], Section 4.

The two Swan conductors above belong to the input sheaf $\mathcal{F}_a$ on the x-line. The transformed rank-two sheaf on the a-line has different local behavior: it is tame at zero and has Swan conductor one at infinity. The transform changes the geometric object, not just its notation.

15.3 Weights explain the square root

The values of the rank-one Artin–Schreier sheaf have absolute value one over every finite extension: $\mathcal{F}_a$ is pure of weight zero. Deligne's theorem on weights bounds the absolute values of the eigenvalues on H^1_c by $\sqrt{q}$. Together with (68), this is already enough to prove the desired inequality.

In this case one has the stronger statement of purity in weight one. At both boundary points the rank-one wild inertia character is nontrivial, so there is no boundary contribution and

$$H^1_c(\mathbb{G}_m, \mathcal{F}_a) \xrightarrow{\sim} H^1(\mathbb{G}_m, \mathcal{F}_a).$$

Compactly supported cohomology has weights at most one; ordinary cohomology has weights at least one, by duality and the corresponding upper bound for the dual sheaf. Hence both Frobenius eigenvalues α_a, β_a have absolute value exactly $\sqrt{q}$. The trace formula gives

$$K(a) = -(\alpha_a + \beta_a), \qquad |K(a)| \le |\alpha_a| + |\beta_a| = 2\sqrt{q}.$$

The substantive inputs are the ramification calculation, the cohomological formulas, and the theorem on weights. The last step is an elementary inequality.

The mechanism also persists under extension of the base field. For $a \in \mathbb{F}_q^\times$, set

$$K_r(a) = \sum_{x \in \mathbb{F}_{q^r}^\times} \psi\Big(\mathrm{Tr}_{\mathbb{F}_{q^r}/\mathbb{F}_q}(x + a/x)\Big).$$

The same fiber gives

$$K_r(a) = -(\alpha_a^r + \beta_a^r), \qquad |K_r(a)| \le 2q^{r/2}. \tag{69}$$

As with the Fermat curve, an entire sequence of sums is governed by one finite-dimensional operator. Here Fourier transformation constructs those operators uniformly as a geometric family. Its rank, ramification, and weights make quantitative predictions that are unavailable from the list of summands alone.

16 Kernel composition and reference theorems

The main calculation in this section eliminates an intermediate variable from two successive integral transforms. It provides a composition rule, a unit and an associativity check under explicit geometric hypotheses. Sections 16.2–16.5 are a reference layer: they describe additional structures and imported theorems for higher coefficients. No proof of spectral representability, higher ambidexterity or the higher Kummer classification is supplied here. Those results belong to the background library, not to the chain or kernel calculi just constructed.

16.1 Compose kernels by eliminating the intermediate variable

Kernel composition is a symbolic calculation once its geometric operations are available. To make every step legitimate, let X, Y, Z be smooth proper varieties over a field and let $\mathcal{K}$ and $\mathcal{L}$ be perfect coherent kernels on $X \times Y$ and $Y \times Z$. All tensor products, pullbacks and pushforwards in this subsection are derived. These hypotheses provide the proper base-change and projection formulas and keep the resulting kernels perfect.

Write p_X, p_Y for the projections from $X \times Y$, and q_Y, q_Z for those from $Y \times Z$. The input-output rule for the first kernel is

$$\Phi_{\mathcal{K}}(F) = Rp_{Y*}(Lp_X^* F \otimes^L \mathcal{K}).$$

We want a single kernel with the same effect as applying $\mathcal{K}$ and then $\mathcal{L}$. The candidate should sum over the intermediate Y-variable while retaining the X- and Z-variables. Let

$$a = p_{12}, \quad b = p_{23}, \quad c = p_{13}$$

be the projections from $X \times Y \times Z$, and let r_X, r_Y, r_Z be its projections to the individual factors. Denote the projections from $X \times Z$ by π_X, π_Z.

Substitute the first calculation into the second. Put $A = Lp_X^* F \otimes^L \mathcal{K}$. The square with maps a, b, p_Y, q_Y is the Cartesian square expressing

$$(X \times Y) \times_Y (Y \times Z) = X \times Y \times Z.$$

Base change gives $Lq_Y^* Rp_{Y*} A \simeq Rb_* La^* A$. Substitution and the projection formula now give

$$\begin{aligned}
\Phi_{\mathcal{L}}\Phi_{\mathcal{K}}(F) &= Rq_{Z*}(Lq_Y^* Rp_{Y*} A \otimes^L \mathcal{L}) \\
&\simeq Rq_{Z*}(Rb_* La^* A \otimes^L \mathcal{L}) \\
&\simeq Rr_{Z*}(Lr_X^* F \otimes^L La^*\mathcal{K} \otimes^L Lb^*\mathcal{L}) \\
&\simeq R\pi_{Z*}(L\pi_X^* F \otimes^L Rc_*(La^*\mathcal{K} \otimes^L Lb^*\mathcal{L})).
\end{aligned}$$

The last line is the defining rule for a single transform. Thus the calculation has produced its kernel:

$$\mathcal{L} \circ \mathcal{K} = Rp_{13*}(Lp_{12}^*\mathcal{K} \otimes^L Lp_{23}^*\mathcal{L}), \qquad \Phi_{\mathcal{L}}\Phi_{\mathcal{K}} \simeq \Phi_{\mathcal{L}\circ\mathcal{K}}.$$

The middle variable has been eliminated. This is the geometric version of multiplying two matrices and summing over their shared index, with the pushforward and tensor product retaining cohomological information.

Check the unit and the order of composition. For the diagonal inclusion $\Delta : X \to X \times X$, the projection formula gives

$$Lp_1^* F \otimes^L \Delta_* \mathcal{O}_X \simeq \Delta_* F, \qquad Rp_{2*}\Delta_* F = F.$$

Hence $\mathcal{O}_\Delta$ is the identity kernel. More generally, the structure sheaf of the graph of an automorphism $f : X \to X$ sends F to $Rf_* F$. This fixes which direction the graph acts; for inversion on an elliptic curve, pushforward and pullback agree because inversion is its own inverse.

For a third perfect kernel $\mathcal{N}$ on $Z \times W$, applying the same substitutions to either parenthesization gives the kernel

$$Rp_{14*}(Lp_{12}^*\mathcal{K} \otimes^L Lp_{23}^*\mathcal{L} \otimes^L Lp_{34}^*\mathcal{N})$$

on $X \times W$; the projections are now from $X \times Y \times Z \times W$. This proves associativity through the common four-variable calculation. A map of kernels induces a map of the corresponding transforms by the same formulas, so the computation retains morphisms as well as objects.

Keep the theorem boundary visible. For the elliptic Poincaré kernel, Mukai's inversion theorem identifies its composite with the structure sheaf of the graph of inversion, shifted by $[-1]$. Given that geometric identification, the kernel calculus immediately gives $\Phi^2 \simeq (-1_E)^*[-1]$ on every input. The calculus above proves how to compose and use kernels; it does not by itself compute the Poincaré composite or prove the required base-change theorems [8].

For finite sums the same rules are ordinary sum and product. For exponential-sum sheaves, compactly supported pushforward replaces the proper coherent pushforward. These settings require their own base-change hypotheses; they are not silently identified with the proper coherent setting proved here.

16.2 Reference: coherence beyond the kernel calculation

A derived category has already retained more than a list of cohomology groups. It remembers complexes and extensions, and it supplies derived tensor products and pushforwards. A stable infinity category also retains the spaces of maps and the homotopies relating their compositions.

This extra structure is relevant when a construction must vary functorially and when several ways of composing it must agree coherently.

For example, convolution is associative because both parenthesizations integrate the same external tensor product along the addition map from a triple product. For four or more factors, these identifications must themselves agree. A symmetric monoidal infinity category organizes this entire system, including the unit and symmetry constraints. The local algebra calculations remain visible: over an ordinary ring, derived tensor products still have Tor groups. With a ring spectrum as coefficients, modules can also carry stable homotopy information that an ordinary complex over a discrete ring does not encode.

This distinction matters for reconstruction. Replacing a character space by its dimension forgets its maps and multiplication. A Grothendieck class forgets a nonsplit extension. Passing to a homotopy category forgets higher information about mapping spaces. Each passage can be useful, but a reconstruction theorem must be formulated before the particular information it needs has been discarded.

16.3 Reference: spectral Cartier duality

Let R be a connective commutative ring spectrum and X a spectral abelian variety over R. Tensor product and the group law give $\mathrm{QCoh}(X)$ a Hopf-algebra structure among stable R-linear categories. Its categorical dual has convolution as its multiplication. A biextension on $X \times Y$ defines a symmetric monoidal transform from convolution on X to tensor product on Y.

Categorified Cartier duality specifies the candidate category, and Tannaka duality gives the criterion for recovering a geometric dual from it. Existence still needs geometry. Lurie constructs the functor of multiplicative line bundles Pic^m_X, proves its spectral representability using Artin's theorem, and proves that its universal biextension is perfect. Perfectness is equivalent to the associated Fourier–Mukai functor being an equivalence. The identity component of the spectral Picard space need not even be flat, so cannot simply substitute for this construction [11], Sections 4.7 and 5, especially Proposition 5.5.1 and Theorem 5.6.4.

Here the categorical formulation has become part of the construction of the parameter space itself. The circle example began with a visible space of characters, and the elliptic example began with $\mathrm{Pic}^0(E)$. The spectral construction can determine the dual through the category it must support.

16.4 Reference: integration over higher symmetries

The finite Fourier projector averages over group elements. A higher symmetry has objects, automorphisms, and potentially automorphisms between automorphisms. Even the first step changes the counting rule. A finite groupoid has cardinality

$$|\mathcal{G}| = \sum_{[x] \in \pi_0 \mathcal{G}} \frac{1}{|\operatorname{Aut}(x)|}.$$

In particular, $|BG| = 1/|G|$ for a finite group G: a single isomorphism class can still have nontrivial symmetry. For a connected space with only finitely many nonzero homotopy groups, all finite, the corresponding homotopy cardinality is

$$|X| = \prod_{i \geq 1} |\pi_i(X)|^{(-1)^i}.$$

These formulas describe rational numerical counting. A higher Fourier transform requires a version of integration inside the coefficient category itself.

The relevant condition is ambidexterity: suitably finite indexing spaces admit a compatible identification of the colimit and limit operations. For a finite set, this starts with the identification of a finite direct sum and product in an additive category. For a finite group action in a stable category, it involves the norm from homotopy orbits to homotopy fixed points. An equivalence of that norm permits a form of averaging; it is an additional property, not a formal consequence of writing down a group action. Higher semiadditivity extends this compatibility to appropriate spaces with finite homotopy data.

This is why replacing sets by spaces in a Fourier formula is a mathematical theorem rather than a change of notation. The integration must interact correctly with products, pullbacks, and composition of fibres. Once those properties hold, the kernel calculus has a higher counterpart.

16.5 Reference: a higher Kummer theorem

Barthel, Carmeli, Schlank, and Yanovski construct higher semiadditive Fourier transforms, including the classical finite transform at height zero and chromatic transforms at positive heights [12]. One application classifies certain Galois extensions of ring spectra.

Fix a prime p and $n, r \geq 1$. Let R be a $T(n)$-local commutative ring spectrum, and write $R^\times = \mathrm{gl}_1(R)$ for its connective spectrum of units. Choose a primitive higher p^r-th root of unity of height n: this is a map of spectra

$$\omega : \Sigma^n H(\mathbb{Z}/p^r) \longrightarrow R^\times$$

satisfying the primitivity condition of [12], Section 6.2. Here H denotes the Eilenberg–Mac Lane spectrum. In this chromatic setting, primitivity makes the associated Fourier transforms invertible; ω is higher coherent data, not merely an element of $\pi_0 R$.

Let M be a connected n-finite $H(\mathbb{Z}/p^r)$-module spectrum: $\pi_i M$ is finite for $1 \leq i \leq n$ and zero otherwise. Its Pontryagin dual is the spectrum

$$M^* = \mathrm{RHom}_{\mathbb{Z}}(M, \mathbb{Q}/\mathbb{Z}), \qquad \pi_i M^* \simeq \mathrm{Hom}_{\mathbb{Z}}(\pi_{-i} M, \mathbb{Q}/\mathbb{Z}).$$

The notation uses the underlying complex of abelian groups. Put $G = \Omega\Omega^\infty M$, the loop space of the underlying infinite loop space. Their higher Kummer theorem identifies the space of extensions as

$$\{G\text{-Galois extensions of } R\} \simeq \mathrm{Map}_{\mathrm{Sp}}(\Sigma^n M^*, R^\times). \tag{70}$$

The right-hand side is a mapping space of spectra, and the left-hand side is the space of G-Galois extensions in the $T(n)$-local category. The finiteness, localization, and higher-root hypotheses specify where the classification holds. The source is Theorem C of [12], a connected case of Theorem 7.34.

Compare the role of the inputs. For a cyclic field extension, roots of unity make the character lines available, and their multiplication gives a radical class. Here the symmetry is G, the coefficients form a ring spectrum, and a mapping space into units records the extension together with its higher automorphisms. Retaining the mapping space is part of the strength of the classification: its connected components record equivalence classes, while its higher homotopy retains the symmetries of the reconstruction data.

The same work also constructs categorical Fourier equivalences for suitable local systems of chromatic modules. These results enlarge both sides of the elementary equivalence between representations and character-graded vector spaces. They provide an algebraic-topological setting in which duality, descent, and Galois reconstruction can continue beyond ordinary fields and varieties.

16.6 What the examples establish together

The reusable calculations have explicit inputs and outputs: a permutation determines weighted root coordinates; average and half-difference expose a square; a coefficient recurrence constructs every chain in a truncated-polynomial tensor product; remainders choose a terminating rank-and-degree reduction; and an intermediate-variable elimination composes kernels. Each calculation includes its verification and return map. Applications to elliptic bundles, cohomological bounds and higher coefficients retain separate theorem dependencies. A reference theorem extends the range of an operation only when its hypotheses have been supplied.

A useful comparison with *What a Boundary Map Can Prove* is to ask which operation has become elementary. There an incidence map turns a cycle condition into linear algebra and a determinant into a tree count. Here a selected response to a permutation turns unknown roots into invariant powers and compatible products. The comparison concerns these constructed operations, not a claim that one final framework explains every example.

Further reading

For Lagrange resolvents and the classical formulas, see Beukers [1], Sections 9.2–9.3. Milne [2], Chapters 4–5, develops the quartic pairing resolvent, cyclic extensions, Kummer theory, and solvability; Chapter 8 relates finite étale algebras to covers. Bluher [3] treats Dickson and Chebyshev identities and their finite-field permutation properties. Baez [4], Section 6.1, gives the categorical character decomposition. The Stacks Project [5, 6] supplies the torsor and Kummer formulations.

Shaska [13], Sections 3.1 and 5, studies genus-two elliptic subcovers and split Jacobians. Liu, Maxim, and Wang [15] develop Mellin transformations using universal local systems. The circle calculation here is a small model of that construction. Schröer [7] relates unipotent bundles to formal-group convolution; Chen [8] explains elliptic Fourier–Mukai calculations. Hernández Ruipérez and Tejero Prieto [14], Section 2, supply the transform and stability statements used in the rank-and-degree reduction.

Elkies [16], especially Exercise 12, connects Jacobi sums to Fermat curves. Otsubo and Yamazaki [17] construct the motivic lifts of Gauss–Jacobi identities; the tensor isomorphism is Theorem 1.1, and the Frobenius and period comparisons are discussed in Theorem 1.2 and Remarks 7.4–7.5. For the quantitative application, Michel [18], Section 4, introduces the cohomological estimates, Katz [9] develops Kloosterman sheaves and monodromy, and Laumon [10] supplies geometric Fourier theory.

For the higher-coefficient questions, Lurie [11], Sections 3–5, develops categorical Cartier duality, biextensions, Tannaka reconstruction, and spectral Fourier–Mukai duality. Barthel, Carmeli, Schlank, and Yanovski [12], especially Sections 3–5 and 7, develop higher semiadditive Fourier theory, categorical extensions, and the higher Kummer classification. These references provide the additional hypotheses and constructions used in the final section.

References

[1] Frits Beukers, *Notes on Galois Theory*, Sections 9.2 and 9.3.
https://www.staff.science.uu.nl/~beuke106/ringengalois/dicgal.pdf.

[2] J. S. Milne, *Fields and Galois Theory*, version 5.10, 2022.
https://www.jmilne.org/math/CourseNotes/FT.pdf.

[3] Antonia W. Bluher, *Permutation properties of Dickson and Chebyshev polynomials and connections to number theory*, Finite Fields and Their Applications 76 (2021), 101899. https://arxiv.org/abs/1707.06877.

[4] John C. Baez, *Higher-Dimensional Algebra II: 2-Hilbert Spaces*, Advances in Mathematics 127 (1997), 125–189, Section 6.1. https://arxiv.org/abs/q-alg/9609018.

[5] The Stacks Project Authors, *The Stacks Project*, Section 59.28, "Kummer theory," Tag 03PK. https://stacks.math.columbia.edu/tag/03PK.

[6] The Stacks Project Authors, *The Stacks Project*, Section 21.4, "First cohomology and torsors," Tag 03AG. https://stacks.math.columbia.edu/tag/03AG.

[7] Stefan Schröer, *On the ring of unipotent vector bundles on elliptic curves in positive characteristics*, 2009. https://arxiv.org/abs/0908.3743.

[8] Kaiyi Chen, *Fourier–Mukai Transforms and Vector Bundles on Elliptic Curves*, lecture slides, 2025. https://www.math.pku.edu.cn/puremath/docs/2025-05/20250530155446773636.pdf.

[9] Nicholas M. Katz, *Gauss Sums, Kloosterman Sums, and Monodromy Groups*, Annals of Mathematics Studies 116, Princeton University Press, 1988. https://web.math.princeton.edu/~nmk/Katz-GKM.pdf.

[10] Gérard Laumon, *Transformation de Fourier, constantes d'équations fonctionnelles et conjecture de Weil*, Publications Mathématiques de l'IHÉS 65 (1987), 131–210. https://numdam.org/item/PMIHES_1987__65__131_0/.

[11] Jacob Lurie, *Elliptic Cohomology I: Spectral Abelian Varieties*, 2018. https://people.math.harvard.edu/~lurie/papers/Elliptic-I.pdf.

[12] Tobias Barthel, Shachar Carmeli, Tomer M. Schlank, and Lior Yanovski, *The Chromatic Fourier Transform*, 2022. https://arxiv.org/abs/2210.12822.

[13] Tony Shaska, *Genus two curves covering elliptic curves: a computational approach*, 2012, especially Sections 3.1 and 5. https://arxiv.org/abs/1209.3187.

[14] Daniel Hernández Ruipérez and Carlos Tejero Prieto, *Fourier–Mukai transforms for coherent systems on elliptic curves*, 2006, revised 2007, especially Section 2. https://arxiv.org/abs/math/0603249.

[15] Yongqiang Liu, Laurenţiu Maxim, and Botong Wang, *Non-abelian Mellin transformations and applications*, Forum of Mathematics, Sigma 10 (2022), e100. https://doi.org/10.1017/fms.2022.91.

[16] Noam D. Elkies, *$L(s,\chi)$ as an entire function; Gauss sums*, Math 229: Introduction to Analytic Number Theory, Harvard University, 2019, especially Exercise 12. https://people.math.harvard.edu/~elkies/M229.19/lsx.pdf.

[17] Noriyuki Otsubo and Takao Yamazaki, *Motivic Gauss and Jacobi sums*, Annales de la Faculté des Sciences de Toulouse, Série 6, **35** (2026), no. 1, 95–135; especially Theorems 1.1–1.2 and Remarks 7.4–7.5. https://doi.org/10.5802/afst.1842. https://arxiv.org/abs/2402.06072.

[18] Philippe Michel, *Lectures on applied ℓ-adic cohomology*, notes from the 2016 Arizona Winter School, 2017, Section 4. https://arxiv.org/abs/1712.03173. https://swc-math.github.io/aws/2016/2016MichelNotes.pdf.