

CORE SERIES X

Universal Residual Seeds

Deformation forests, transport categories, and conditional system completion

Visual guide

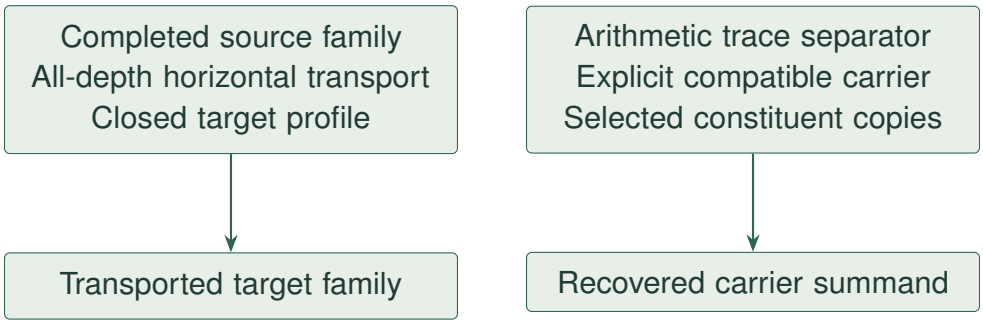
Komorebi P.

Verification edition / 20 October 2026

REALIZATION CALCULUS

The current reading contract

This guide preserves the Block 27 and scalar/Mellin worked examples. Their stated checkpoint limits belong to those source revisions. The later occupied-system endpoint is checked at its amended conditional scope: 20/20 internal claims. Actual completed source members, inverse-limit profile closure, compatible carrier presentations and constituent local comparisons remain supplied inputs.



The endpoint proves both outputs. It does not assert that their supplied local structures have been identified. Zero multiplicity is allowed; nonzero occupancy needs its own input. All-place carrier recovery needs the pointwise presentations at each place.

A checked scope and a later unpromoted scope

The final checked conditional release records two successful builds, matching proof objects, source and axiom checks, and negative controls. Its exact production source and saved records accompany this edition. A later stronger residual-system attempt records 19/20, no promoted endpoint, and an unavailable pinned Lean executable. The two records describe different input contracts.

The traditional edition gives the complete retained manuscript and the later conditional assembly. The Lean guide shows the source-exact endpoint, its proof structure and its evidence. The PDF build checks presentation. The separate October 20 production run freshly compiles the selected 122-module closure; its source selection and exact limits are recorded in the Lean guide.

READER'S GUIDE

One root is not the whole problem

Main text: §§1–3, 8–11. Current contract: Appendix D. Earlier scalar/Mellin checkpoints are historical.

Paper X separates three questions that can look deceptively similar. Does a residual object exist? Does it extend coherently through every precision? Does the resulting object occur in the required carrier? A successful answer to one question does not supply the next.

A residual root

$x^2 + 1 = 0$ over $\mathbb{F}_2$



one point mod 2

no point mod 4

A coherent branch

$x^2 + 1 = 0$ over $\mathbb{Z}_5$



reductions agree

An occupied packet

1	0
0	0

$\operatorname{rank} E = 1$

positivity needs a carrier

Three models to keep in view

The scalar equation $x^2 + 1 = 0$ has the root 1 modulo two, but no root modulo four. Over powers of five, the residues $2, 7, 57, 182$ form the beginning of a compatible branch. Neither fact constructs an automorphic carrier. In a supplied finite carrier, an idempotent such as $\operatorname{diag}(1, 0)$ detects an occupied subspace; it does not explain where that carrier came from.

Read the appendix in three passes. Pages V3–V11 explain residual digits, finite counts, and obstruction spaces. Pages V12–V20 explain directed transport, common data, and finite occupancy. Pages V21–V29 give small exact models for the newer arithmetic continuation and its remaining limits.

> **Which ledger applies?** The Block 27 and scalar/Mellin checkpoints record 17/20. The later checked assembly records **20/20 at the amended conditional scope**. Its inputs are stated on the opening pages and in Appendix D. The finite character examples below are not proofs of prescribed Frobenius occurrence.

The figures are exact-model illustrations, not pictures of an entire arithmetic moduli space. New Python checks support these examples. Saved Lean evidence is preserved separately. The October 20 production replay has its own source and compiler records in the Lean guide.

NOTATION BEFORE PICTURES

An atlas of the objects

Main text: Definitions 1.1, 5.1, 6.1, 7.1–7.2, and 8.1.

Keep the coefficient ring, its reduction, and the carrier separate. A residue class is a finite-precision address. A vector in a carrier is a different kind of object, even when both are selected by something called a packet.

Object	What it records	What it does not supply
$\mathcal{O}, \varpi, k$	Complete valuation ring, uniformizer, and residue field k of size q.	A represented deformation functor.
$\bar{x}$	One residual object in the declared moduli problem.	Other residual roots or their transports.
$X_m(\bar{x})$	Admissible objects modulo ϖ^m reducing to $\bar{x}$.	A lift of every individual member.
$(x_m)_m$	A branch satisfying $x_{m+1} \bmod \varpi^m = x_m$.	A global object without the stated reconstruction hypotheses.
w_i	A coordinate constrained to begin in $\varpi^{w_i}\mathcal{O}$.	A classification of all singular deformation rings.
$Z(T)$	The formal series $\sum_{m\geq 1} \#X_m T^m$.	The cardinality of the inverse limit.
$C^\bullet, d$	Cochains and a differential with $d^2 = 0$.	Its arithmetic identification with a Galois problem.
G	The graph of supplied admissible residual transports.	A proof that a proposed arithmetic edge exists.
$\mathscr{C}$	Common algebraic polynomials and declared local data.	A member realizing those data at every prime.
$C\boldsymbol{m} = \boldsymbol{b}$	A finite trace system for carrier multiplicities.	Integrality, positivity, or available capacity automatically.
$e_T V$	The joint eigenspace inside an existing finite carrier V.	Nonzero occupancy for every desired joint label.

Three indices that must not be conflated

The depth m counts retained digits. A coordinate index i selects a parameter or a packet. A coefficient place λ selects a different local field. Increasing m at one λ neither changes i nor constructs the member at another λ.

> **Read this as:** Specify the ambient object before applying a projector, a reduction, a rank, or a counting formula. The same-looking symbol can refer to quite different spaces.

RESIDUAL PARTITION

Higher digits cannot change the root

Main text: Theorems 2.1–2.2 and Corollary 2.3.

Work in $\mathbb{Z}/27\mathbb{Z}$. Every element has a unique form $x = 3q + r$ with $r \in \{0,1,2\}$ and $q \in \{0,\dots,8\}$. The row below is the residual root. The column contains all the remaining information at this precision.

$x = 3q + r \pmod{27}$: columns are higher-digit addresses q

	$q=0$	$q=1$	$q=2$	$q=3$	$q=4$	$q=5$	$q=6$	$q=7$	$q=8$
$r=0$	0	3	6	9	12	15	18	21	24
$r=1$	1	4	7	10	13	16	19	22	25
$r=2$	2	5	8	11	14	17	20	23	26

Adding 3 moves within one row. Adding 1 can change the row.

Read the operation, not just the arrow

An operation $x \mapsto x + 3h$ preserves r, because its correction vanishes modulo three. Repeated corrections of this form can visit every column in a row, but can never enter another row. By contrast, adding one changes the residual class. It is therefore not a positive-depth perturbation in this model.

For example, the addresses $1, 4, 7, 10, \dots, 25$ may be deeply connected by higher-digit changes. They remain disjoint from the addresses reducing to zero. In a genuine deformation problem, the same conclusion applies to operations whose reduction is the identity, not to every operation that happens to be called a deformation.

$$\operatorname{red}(x + 3h) = \operatorname{red}(x), \qquad \mathcal{M}(A) = \bigsqcup_{\bar{x}} \mathcal{M}_{\bar{x}}(A).$$

> **Read this as:** Vertical completeness can fill one entire residual fiber without reaching a single point in a different fiber.

All 27 addresses and their finite corrections are checked by integer reduction. Passing to an infinite perturbation stream additionally requires the convergence condition recorded in the formal no-crossing statement.

INVERSE-LIMIT RECONSTRUCTION

A branch fixes digits permanently

Main text: Theorem 3.2 and the compatible-digit example in §3.

The rational number $1/4$ is a five-adic integer. Its normalized residues are obtained by solving $4x_m \equiv 1 \pmod{5^m}$, not by taking decimal approximations. At the first six depths they are $4, 19, 94, 469, 2344, 11719$.

	5^0	5^1	5^2	5^3	5^4	5^5
$x_1 = 4$	4	—	—	—	—	—
$x_2 = 19$	4	3	—	—	—	—
$x_3 = 94$	4	3	3	—	—	—
$x_4 = 469$	4	3	3	3	—	—
$x_5 = 2344$	4	3	3	3	3	—
$x_6 = 11719$	4	3	3	3	3	3

$$4x_m - 1 = 3 \cdot 5^m$$

digits already fixed never change

Compatibility is stronger than independent solvability

The residues satisfy both

$$x_{m+1} \equiv x_m \pmod{5^m}, \qquad 4x_m - 1 = 3 \cdot 5^m.$$

The latter identity shows that the residual error tends to zero in the five-adic norm. The stabilized digit sequence is

$$\frac{1}{4} = 4 + 3 \cdot 5 + 3 \cdot 5^2 + 3 \cdot 5^3 + \cdots \quad \text{in } \mathbb{Z}_5.$$

This series is not convergent in the ordinary real norm. The topology is part of the statement. To view the branch as a positive-depth coordinate above the residual root 4, use $t = x - 4 \in 5\mathbb{Z}_5$.

An inverse-limit point is the entire compatible family, not its last displayed approximation. The six rows are a finite window; the displayed exact formula supplies a branch at every depth in this particular example. In the represented deformation theorem, compatibility reconstructs maps on all ring elements, with multiplication, continuity, and locality retained.

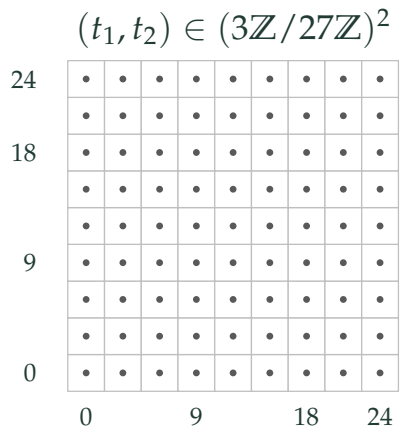
> **Read this as:** A later approximation may append a digit. It may not revise a digit already certified by reduction.

ACTUAL FINITE POINTS

A smooth chart has independent choices

Main text: Proposition 3.4 and Theorem 6.2.

Take $R = \mathbb{Z}_3[[t_1, t_2]]$ and coefficient-preserving local maps to $\mathbb{Z}/27\mathbb{Z}$. Locality forces both parameters into the maximal ideal. Each therefore has nine possible values, $0, 3, 6, \ldots, 24$.



$9 \times 9 = 81$ points

$(t_1, t_2) \mapsto (t_1, t_2) \bmod 9$

Each parent has $3^2 = 9$ children.

The chart hypothesis supplies independent digits.

Why the exponent is $d(m-1)$

At depth m, the image of the maximal ideal contains 3^{m-1} elements. Two independent parameters give $3^{2(m-1)}$ points. The grid displays all 81 points at depth three, not a random sample. Reduction to depth two maps each coordinate to one of $0, 3, 6$, and every parent has nine children.

$$\#X_m = 3^{2(m-1)}, \qquad Z(T) = \sum_{m \geq 1} \#X_m T^m = \frac{T}{1 - 9T}.$$

Counting all 27^2 parameter assignments at depth three would ignore locality. Counting only tangent directions would describe the first layer but not all later layers. The power-series presentation explains why these choices stay independent through every precision.

This is a coordinate model with a proved presentation. A general deformation ring can impose equations among the parameters, changing the counts and causing branches to die. The next pages retain the same depth convention while changing the allowed coordinates or adding a singular equation.

> **Read this as:** The smooth count is not a universal rule for deformation spaces. Its independence comes from the actual formal power-series chart.

WEIGHTED DYADIC CHARTS

Weights delay visibility

Main text: Proposition 5.3, Theorem 6.3, and §12.2.

Over $\mathbb{Z}_2$, require three coordinates to lie in $2\mathbb{Z}_2$, $2^2\mathbb{Z}_2$, and $2^4\mathbb{Z}_2$. At depth m, coordinate i has $\max(m - w_i, 0)$ free binary digits. A zero in the table means only one allowed residue: zero.

	$m=1$	$m=2$	$m=3$	$m=4$	$m=5$	$m=6$	$m=7$
$w_1 = 1$	0	1	2	3	4	5	6
$w_2 = 2$	0	0	1	2	3	4	5
$w_3 = 4$	0	0	0	0	1	2	3
count	1	2	8	32	256	2048	16384

Cell = number of free binary digits, not first separation depth.

The two different thresholds

A coordinate of weight w_i first becomes nontrivial when $m = w_i + 1$. By contrast, the first separation depth of two distinct fixed values a, b is $v_2(a - b) + 1$. Thus 1 and -1 separate at depth two, because their difference has valuation one. Neither threshold is simply the printed coordinate label.

For weights $(1, 2, 4)$,

$$N(m) = 2^{\max(m-1,0)+\max(m-2,0)+\max(m-4,0)}.$$

The first counts are $1, 2, 8, 32, 256, 2048$. Once all three coordinates are active, every additional precision digit multiplies the count by eight. Separating the initial levels from that geometric tail gives

$$Z(T) = \frac{T(1 - 6T - 8T^2 - 32T^3)}{1 - 8T}.$$

Negative coefficients in the numerator do not mean negative numbers of points; expanding the whole rational expression gives positive coefficients.

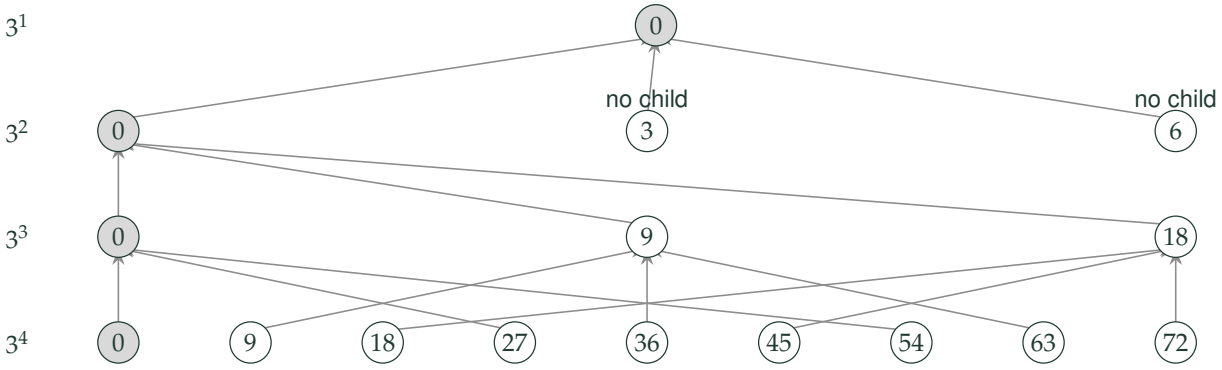
> **Read this as:** Weights change the finite-depth delay. After activation, this chart still has three independent digit streams.

SINGULAR POWER SEEDS

Many approximate roots, one limit

Main text: Theorem 6.4 and Corollary 6.6.

Now impose $x^2 = 0$ modulo 3^m. The condition is exactly $2v_3(x) \geq m$, so the roots are the multiples of $3^{\lceil m/2 \rceil}$ in that quotient. Their number is $3^{m-\lceil m/2 \rceil}$.



More finite roots, but only the zero stream persists through every precision.

Count vertices and count branches separately

There are three roots modulo nine: $0, 3, 6$. But the only roots modulo 27 are $0, 9, 18$, and all three reduce to zero modulo nine. The other two vertices at the previous level have no children. Counting them as permanent branches would be incorrect.

$$\#X_m = 1, 3, 3, 9, 9, 27, \ldots, \qquad Z(T) = \frac{T + 3T^2}{1 - 3T^2}.$$

The count-growth exponent is $\lim_m \log_3(\#X_m)/m = 1/2$. Nevertheless, a compatible branch reconstructs $x \in \mathbb{Z}_3$ with $x^2 = 0$. Since $\mathbb{Z}_3$ has no nonzero nilpotent, the only such branch is the zero branch.

More generally, $x^e = 0$ has count $q^{m-\lceil m/e \rceil}$ and count-growth exponent $1 - 1/e$. These exponents describe finite thickenings. They do not assert that the set of complete roots has a fractional geometric dimension or more than one point.

> **Read this as:** The seed zeta function counts finite-level solutions. It is not a generating function for distinct infinite branches.

A SCALAR COMPANION MODEL

One equation, several local outcomes

Main text: Theorem 8.2, Corollary 8.3, and Example 8.4.

Consider the fixed scalar equation $x^2 + 1 = 0$ in $\mathbb{Z}/p^m\mathbb{Z}$. The following are exact counts for five primes and three depths. The coefficient rings and the scalar nature of the unknown are fixed throughout this page.

	$p = 2$	$p = 3$	$p = 5$	$p = 7$	$p = 13$
$m = 1$	1	0	2	0	2
$m = 2$	0	0	2	0	2
$m = 3$	0	0	2	0	2

Counts of $x \in \mathbb{Z}/p^m\mathbb{Z}$ with $x^2 + 1 = 0$.

Where failure occurs matters

At $p = 3$ and $p = 7$, no residual root exists. At $p = 2$, the residual root exists but dies immediately: every odd square is one modulo four. At $p = 5$ and $p = 13$, the two residual roots are simple, and the usual one-digit equation gives one child above each root at every depth.

At five, one branch begins $2, 7, 57, 182$; the other is its negative modulo each power of five. Arbitrarily many digits in either branch cannot supply the absent scalar root over $\mathbb{Z}_3$. Nor can one root modulo two satisfy the hypothesis that every finite level is nonempty.

The paper's system-completion theorem uses a supplied common algebraic spine, finite nonempty levels at every coefficient place, coherent reductions, and closed local conditions. The scalar example isolates why these are substantive requirements. It is not a disproof of a companion theorem whose coefficients may be enlarged or whose objects are higher-dimensional representations.

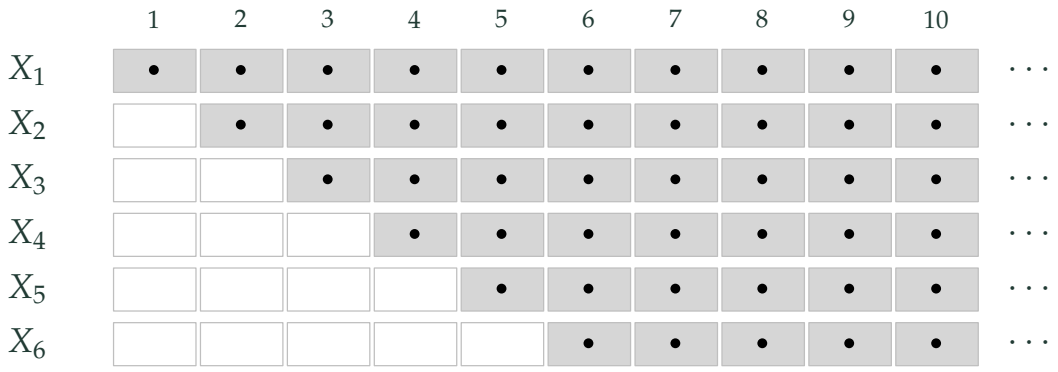
> **Read this as:** A residual root, roots at all depths, and roots at every required coefficient place are three distinct tests.

AN INFINITE-LEVEL COUNTERMODEL

Why finiteness enters compactness

Main text: Theorem 3.3 and its finite-branching hypotheses.

Let $X_m = \{n \in \mathbb{Z} : n \geq m\}$, with the inclusion $X_{m+1} \hookrightarrow X_m$ as transition map. Every level is nonempty. Every vertex has at most one child. Yet there is no compatible infinite branch.



$X_m = \{n \in \mathbb{N} : n \geq m\}$; every reduction is an inclusion.

The missing hypothesis is not surjectivity

Compatibility would force all coordinates of a branch to be one fixed integer n. Membership at every level would then require $n \geq m$ for every m, which is impossible. Each vertical column in the drawing eventually ends. A finite window always contains columns extending beyond it, but no single column survives all windows.

There is no contradiction with finite-branch compactness: these levels are infinite. Attaching a common artificial root would give that root infinitely many children. The finitely branching rooted-tree argument therefore does not apply.

In the paper, finite nonempty levels and actual total reduction maps provide compactness even when some vertices die and some transition maps are not onto. One can retain only vertices with descendants at arbitrarily large depths. Finiteness makes that surviving part nonempty and permits compatible choices.

$$X_m \neq \varnothing \text{ for all } m \quad \not\Longrightarrow \quad \varprojlim_m X_m \neq \varnothing \quad \text{without the compactness hypotheses.}$$

> **Read this as:** Arbitrarily long finite paths are not automatically one infinite path. The finite-level hypothesis prevents escape to new roots.

A TWO-PACKET COCHAIN COMPLEX

See cycles, boundaries, and classes

Main text: Theorem 4.2 and the finite-complex example.

Over $\mathbb{F}_5$, use bases (e_0, e_1), (u_0, u_1, u_2, u_3), and (f_0, f_1) in degrees zero, one, and two. The only nonzero differential values are $d_0 e_0 = u_0$ and $d_1 u_1 = f_0$. Consequently $d_1 d_0 = 0$.

$d_0 : C^0 \to C^1$

	e_0	e_1
u_0	1	0
u_1	0	0
u_2	0	0
u_3	0	0

$d_1 : C^1 \to C^2$

	u_0	u_1	u_2	u_3
f_0	0	1	0	0
f_1	0	0	0	0

$$d_1 d_0 = 0$$

$$H^0 = \langle [e_1] \rangle$$

$$H^1 = \langle [u_2], [u_3] \rangle$$

$$H^2 = \langle [f_1] \rangle$$

Packet A is exact. Packet B has zero differentials and carries all cohomology.

The quotient is the important step

In degree one, the cycles are $\operatorname{span}(u_0, u_2, u_3)$, but u_0 is a boundary. Thus the cohomology is represented by u_2, u_3, not by every vector killed by d_1. In the other degrees the surviving classes are e_1 and f_1.

$$\dim(H^0, H^1, H^2) = (1, 2, 1).$$

Packet A contains e_0, u_0, u_1, and f_0. It is exact and contributes $(0, 0, 0)$. Packet B contains e_1, u_2, u_3, and f_1, with zero differentials, and contributes $(1, 2, 1)$. Adding the packet cohomologies reconstructs the total.

The projectors must be specified in every degree and commute with the differential. Their existence is not a consequence of the printed dimensions. Identifying this finite complex with the actual continuous cohomology of an arithmetic deformation problem would also require a separate comparison.

> **Read this as:** Resolve the differential first, then pass from cycles to classes. Counting packet dimensions alone does not compute cohomology.

THE CHAIN-MAP CONDITION

An arbitrary cut need not descend

Main text: Theorem 4.2 and the scope of the cohomology splitting.

Retain the previous complex. The valid A-packet projectors are $P^0 = \mathrm{diag}(1,0)$, $P^1 = \mathrm{diag}(1,1,0,0)$, and $P^2 = \mathrm{diag}(1,0)$. They satisfy $d_r P^r = P^{r+1} d_r$.

valid P^1_A

1	0	0	0
0	1	0	0
0	0	0	0
0	0	0	0

invalid Q^1

0	0	0	0
0	1	0	0
0	0	0	0
0	0	0	0

$d_0 P^0_A - Q^1 d_0$

1	0
0	0
0	0
0	0

e_0 is retained at the source but its image u_0 is deleted at the target.

All displayed degreewise maps are idempotent. The differential test distinguishes them.

A boundary can cross an incompatible cut

Replace the degree-one projector by $Q^1 = \mathrm{diag}(0,1,0,0)$ while keeping P^0. Then

$$d_0 P^0 e_0 = u_0, \qquad Q^1 d_0 e_0 = 0, \qquad (d_0 P^0 - Q^1 d_0) e_0 = u_0 \neq 0.$$

The selected degree-zero component is sent outside the selected degree-one component. There is no subcomplex with these cuts, so the cohomology packet theorem cannot be applied to them. Both matrices remain perfectly valid idempotents; the failure is compatibility with the differential.

This example separates two kinds of obstruction. The nonzero matrix above obstructs this proposed decomposition. A nonzero class in H^2 may obstruct a lift only after the deformation problem and its obstruction map have been identified. One should not equate every failed projector commutator with an arithmetic obstruction class.

The exact checks verify both commuting diagrams for the valid family and reject the altered family. They also enumerate the finite kernels and images, so the dimension comparison is not inferred from the drawing.

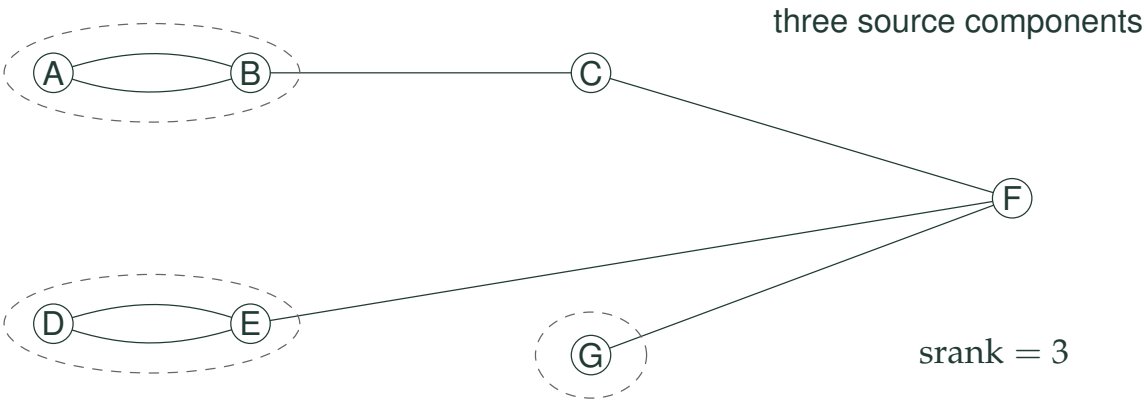
> **Read this as:** A projector on each vector space is not yet a projector on the complex. The differential must respect the chosen packets.

DIRECTED TRANSPORT

How many starting seeds are necessary?

Main text: Theorem 7.3 and the seven-root example.

The arrows below are the entire supplied transport graph. Vertices joined by paths in both directions form strongly connected components. The components $\{A, B\}$, $\{D, E\}$, and $\{G\}$ have no incoming edge from outside themselves.



A minimum set meets each dashed source component.

A lower bound and a construction

Any seed set reaching every vertex must meet all three source components. Otherwise there is no path into the omitted source. Conversely, one representative from each source reaches C and F as well. The directed seed rank is therefore three.

There are precisely four minimum seed sets:

$$\{A, D, G\}, \quad \{A, E, G\}, \quad \{B, D, G\}, \quad \{B, E, G\}.$$

The example checks this by exhaustive enumeration of all 2^7 subsets, independently of the strongly connected component calculation. Adding an admissible edge cannot increase the required seed count. In particular, adding $F \to G$ makes G reachable from the other sources, reducing the count to two.

The graph theorem is exact for the graph supplied. It does not establish that a proposed twist, induction, or arithmetic transfer preserves a determinant, local condition, or occupancy theorem. Those facts justify the arrows before graph optimization begins.

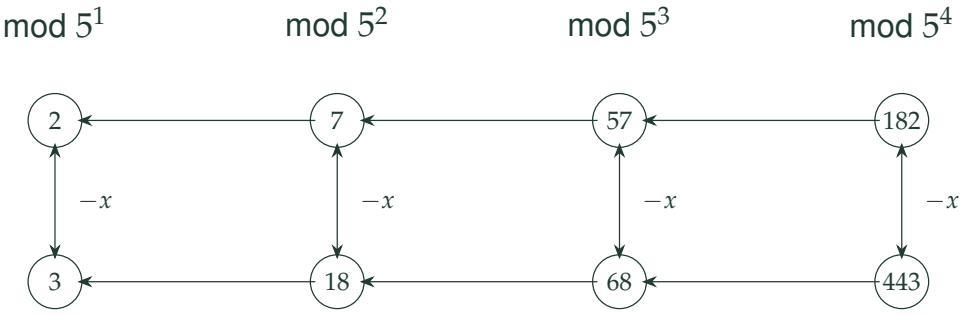
> **Read this as:** Count source components, not all components and not all vertices. But do not add an arithmetic arrow merely because it would improve the count.

ONE REVERSIBLE EXAMPLE

A transport must work at every depth

Main text: Definition 7.1 and the reduction-compatibility requirement.

For the two roots of $x^2 + 1 = 0$ modulo powers of five, negation is an actual reversible transport. It preserves the equation, is its own inverse, and commutes with reduction.



Negation commutes with reduction. Translation by 1 fails already at $2 \mapsto 3 \mapsto 4$ mod 5.

One successful input is not a natural transformation

At every depth m,

$$N_m(x) = -x \pmod{5^m}, \qquad r_{m+1,m} N_{m+1} = N_m r_{m+1,m}.$$

The two compatible branches beginning at 2 and 3 are exchanged. Because $2 \not\equiv 3 \pmod 5$, this transport changes the residual root. It is horizontal rather than a positive-depth correction.

Translation by one is a useful false control. Modulo five it takes the root 2 to the root 3, but takes the root 3 to 4, which is not a root. A single successful arrow between selected points does not define an operation on the whole deformation problem.

Even negation is admissible only for the equation and conditions stated here. An extra condition that fixes the residual root at 2 would not be preserved. Likewise, a representation-theoretic dual or twist needs its own compatibility checks for the chosen local data and carrier.

> **Read this as:** A transport is a family of maps respecting equations and reductions. A convenient pair of endpoints is not enough.

Both branches are computed by exact one-digit lifting. Negation is checked on all roots at the displayed depths; its universal reduction identity follows directly from ring homomorphism laws.

THE ARITHMETIC SPINE

Common coefficients, different residues

Main text: Definition 8.1 and Theorem 8.2.

A common spine records algebraic data before reducing at a coefficient prime. For a small model, take the one polynomial

$$P(T) = T^2 - 5T + 6 = (T-2)(T-3) \in \mathbb{Z}[T].$$

Its reductions at five and seven look different, but they come from the same coefficients.

$$P(T) = T^2 - 5T + 6 = (T-2)(T-3)$$

coefficient images under the declared maps

	T^2	T	1
$\mathbb{Z}$	1	-5	6
$\mathbb{F}_5$	1	0	1
$\mathbb{F}_7$	1	2	6

Different residue rows come from the same algebraic polynomial.

Agreement means descent from supplied data

Modulo five the polynomial is $T^2 + 1$; modulo seven it is $T^2 + 2T + 6$. The roots 2 and 3 are the reductions of the same integer roots in both places. The coefficient lists are not required to be equal as raw residues in different rings.

The paper's common spine is larger: it includes characteristic polynomials at the permitted arithmetic places and the specified determinant, weight, and local labels. Requiring every finite member to match those reductions is a compatibility condition on existing objects. It is not a construction of the objects.

Finite agreement also leaves ambiguity. Adding 35 to an integer coefficient changes neither its residue modulo five nor its residue modulo seven. Thus two finite coefficient observations do not determine a unique global polynomial without a height bound or further exact information. Conversely, one arbitrarily chosen family of local objects need not descend to an algebraic spine.

> **Read this as:** The common spine is input shared across coefficient places. It is not created by comparing unrelated residue labels after the fact.

CHARACTERS BEFORE ARITHMETIC PRIMES

A finite trace matrix can separate types

Main text: Theorem 9.1; retained scalar/Mellin continuation: arithmetic-system separation.

Use the three irreducible characters of S_3: trivial, sign, and standard. On the identity, a transposition, and a three-cycle, their values form the matrix below. Rows are samples; columns are constituent types.

C: actual S_3 character table

	trivial	sign	standard
1	1	1	2
(12)	1	−1	0
(123)	1	1	−1

$$\boldsymbol{m} = \begin{pmatrix} 2 \\ 0 \\ 1 \end{pmatrix} \quad = \quad \boldsymbol{b} = \begin{pmatrix} 4 \\ 2 \\ 1 \end{pmatrix}$$

$$\det C = 6$$

$$C^{-1}\boldsymbol{b} = (2, 0, 1)^T$$

Class representatives are finite group elements, not a proof that suitable Frobenius primes occur.

What full rank buys

Since $\det C = 6$, these samples distinguish every rational multiplicity vector. A target with two trivial constituents, no sign constituent, and one standard constituent has

$$\boldsymbol{m} = \begin{pmatrix} 2 \\ 0 \\ 1 \end{pmatrix}, \qquad \boldsymbol{b} = C\boldsymbol{m} = \begin{pmatrix} 4 \\ 2 \\ 1 \end{pmatrix}.$$

Solving $C\boldsymbol{m} = \boldsymbol{b}$ recovers the unique rational vector. The standard character is not introduced as an arbitrary list: the checks construct matrices for a transposition and a three-cycle, verify their group relations, and obtain their traces.

This finite example explains the linear algebra after separating samples are available. For the arithmetic theorem, the samples must be genuine unramified Frobenius classes with the required compatibility and allowed-prime conditions. Selecting abstract group elements does not construct such primes.

The scalar/Mellin checkpoint supplies a separator assembly conditional on arithmetic coverage. The later checked source `Claim91ConstructedClosure` discharges that interface for its rational-base contract. Neither result follows from the finite character table displayed here.

> **Read this as:** Invertibility answers reconstruction from these samples. Arithmetic occurrence answers whether the required samples can be taken at actual primes.

INTEGRAL, NONNEGATIVE, WITHIN CAPACITY

A solution is not yet a summand

Main text: Theorem 9.2 and its multiplicity-recovery example.

Keep the same separating matrix C and suppose the available carrier has capacities $\boldsymbol{c} = (2,1,1)$. A recovered vector describes a genuine sub-sum of that carrier only if each entry is an integer and $0 \le m_i \le c_i$.

trace data $\boldsymbol{b}$	$C^{-1}\boldsymbol{b}$	capacity $(2,1,1)$
$(4,2,1)$	$(2,0,1)$	allowed
$(1,0,0)$	$(\frac{1}{6},\frac{1}{6},\frac{1}{3})$	nonintegral
$(2,-2,-1)$	$(-1,1,1)$	negative
$(5,3,2)$	$(3,0,1)$	over capacity

Three independent rejection tests

The trace vector $(1,0,0)$ gives $(1/6,1/6,1/3)$, so it fails integrality. The vector $(2,-2,-1)$ gives $(-1,1,1)$, which is integral but not nonnegative. The vector $(5,3,2)$ gives $(3,0,1)$, which exceeds the available capacities. In each case the linear equation itself has a unique rational solution.

The accepted example $\boldsymbol{b} = (4,2,1)$ recovers $(2,0,1)$. When compatible constituent families are already supplied, the same admissible multiplicities define the corresponding direct sum at every coefficient place. This is the finite algebraic core of summand transport.

It does not supply a compatible family for an arbitrary object, construct all local attachments, or ensure that a proposed carrier contains the desired representation. The historical Claim 9.2 checkpoint left these comparisons partial. The later checked carrier theorem retains explicit constituent-copy presentations and supplied local comparisons rather than manufacturing them.

> **Read this as:** Invert the trace matrix, then check the representation-theoretic constraints. A unique vector in a field need not be a realizable multiplicity vector.

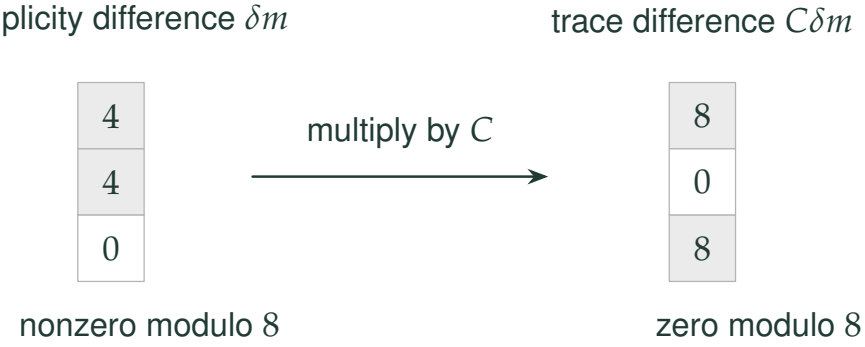
All accepted and rejected vectors are solved over exact rationals. The code rejects the wrong cases by their stated reason, not by a floating-point tolerance.

EXACT PRECISION LOSS

Separation can consume a digit

Main text: The finite trace-recovery discussion in §9.

The character matrix has determinant six, so its inverse contains a factor of two in the denominator. That arithmetic fact is visible as a loss of one binary digit in trace reconstruction.



Modulo 16, the traces separate them. $v_2(\det C) = v_2(6) = 1$.

A concrete ambiguity modulo eight

The multiplicity difference $\delta\boldsymbol{m} = (4, 4, 0)$ is nonzero modulo eight, but

$$C\,\delta\boldsymbol{m} = (8, 0, 8) \equiv 0 \pmod 8.$$

Thus traces known modulo eight do not distinguish those two residue vectors. The same trace difference is nonzero modulo sixteen, so one more trace digit resolves this ambiguity.

More generally, the kernel of C modulo 2^n has exactly two elements:

$$(0, 0, 0), \qquad (2^{n-1}, 2^{n-1}, 0).$$

To see this, the second equation gives $m_1 = m_2$; the third gives $m_3 = 2m_1$; the first then gives $6m_1 = 0$. Three is a unit modulo 2^n, leaving precisely the displayed possibilities.

Conversely, $C\delta\boldsymbol{m} \equiv 0 \pmod{2^{n+1}}$ implies $6\delta\boldsymbol{m} \equiv 0 \pmod{2^{n+1}}$ by the adjugate identity, hence $\delta\boldsymbol{m} \equiv 0 \pmod{2^n}$. Capacity bounds can remove ambiguity earlier in a particular carrier, but do not change this universal linear precision calculation.

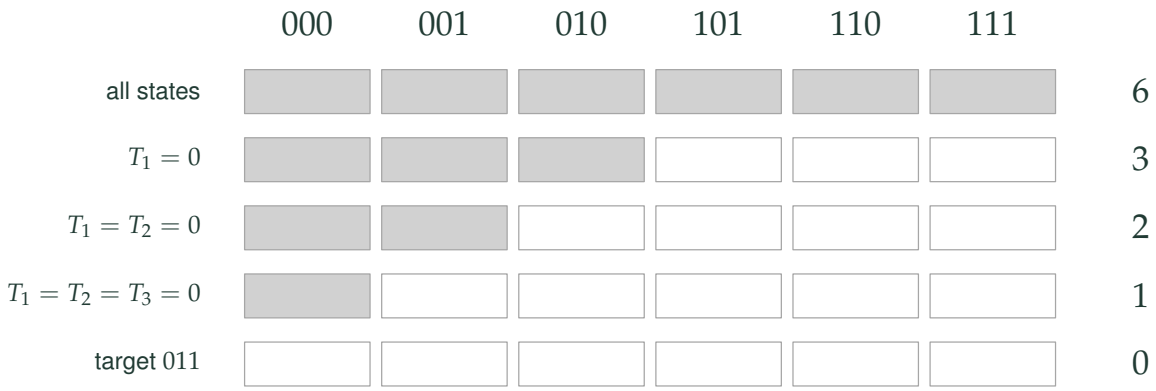
> **Read this as:** Full rank over characteristic zero is not the same as invertibility over every residue ring. Denominator valuations determine the needed extra precision.

EXACT JOINT PROJECTORS

Pairwise occupancy can still fail jointly

Main text: Theorem 10.1 and the finite occupancy models.

Let a six-dimensional rational carrier have the joint eigenvalue labels shown below for three commuting diagonal operators T_1, T_2, T_3. Each coordinate is zero or one. Selecting a zero uses $I - T_i$; selecting a one uses T_i.



Each pair of requirements from 011 is occupied. Their triple intersection is empty.

The full conjunction is a different test

For target $(0,0,0)$, successive conditions reduce the ranks from 6 to 3, then 2, then 1. The joint projector is nonzero, so that label is occupied in this carrier.

For target $(0,1,1)$, every pair of conditions is occupied: $(0,1,0)$ witnesses the first two, $(0,0,1)$ the first and third, and $(1,1,1)$ the second and third. But the triple itself is absent. Multiplying all three selectors gives the zero matrix.

$$E_{\boldsymbol{a}} = \prod_{i=1}^{3} \begin{cases} I - T_i, & a_i = 0, \\ T_i, & a_i = 1. \end{cases} \qquad \operatorname{Tr}(E_{\boldsymbol{a}}) = \operatorname{rank}(E_{\boldsymbol{a}}).$$

The checks compute all eight joint projectors, their pairwise orthogonality, and their sum. The sum is the identity even though two labels are absent: those two projectors are simply zero.

> **Read this as:** A formula for a projector does not force its image to be nonzero. Occupancy must be tested on the complete set of required conditions.

TRACE AND SEMISIMPLICITY

The hypotheses do real work

Main text: Theorem 10.1 and the checked false controls following it.

The finite occupancy theorem uses a split semisimple commuting family in characteristic zero. Two tiny matrices show why neither the coefficient condition nor semisimplicity should disappear from the statement.

I_5 over $\mathbb{F}_5$

1	0	0	0	0
0	1	0	0	0
0	0	1	0	0
0	0	0	1	0
0	0	0	0	1

rank 5, trace 0

$N^2 = 0$

0	1
0	0

only eigenvalue 0

$\dim \ker N = 1$

sole-eigenvalue product $= I_2$

Keep characteristic zero for trace detection, and split semisimplicity for the Lagrange formula.

Trace zero need not mean an empty image

On $\mathbb{F}_5^5$, the identity is an idempotent of rank five, but its field-valued trace is $5 = 0$ in $\mathbb{F}_5$. Rank and localization still detect its nonzero image. A characteristic-zero integer trace cannot be silently transferred to this setting.

An eigenvalue label need not select ordinary eigenvectors

For the nilpotent matrix $N = \left(\begin{smallmatrix} 0 & 1 \\ 0 & 0 \end{smallmatrix}\right)$, the only eigenvalue is zero. A Lagrange selector that multiplies over all other eigenvalues is an empty product, hence I. Its image has dimension two, while $\ker N$ has dimension one. The selector does not project onto the ordinary eigenspace.

The whole two-dimensional space is the generalized zero-eigenvalue packet. Paper XII's Artin-packet treatment is designed to retain such nilpotent structure. That is a different target from Paper X's split semisimple joint-eigenspace formula.

> **Read this as:** A correct theorem in one coefficient category is not a theorem in all categories. Changing from ordinary eigenspaces to generalized packets also changes the object being detected.

CLOSEDNESS IS NOT DECORATION

A condition must survive the limit

Main text: Theorems 8.2 and 11.1; the inverse-limit closure hypotheses.

In $\mathbb{Z}_3$, every term of the sequence $x_k = 3^k$ is nonzero. Yet $x_k \to 0$ in the three-adic topology. The property “nonzero” is not closed, so a limit of objects satisfying it can lose that property.

	3^0	3^1	3^2	3^3	3^4	3^5	3^6
$a_k = 3^0$	1	0	0	0	0	0	0
$a_k = 3^1$	0	1	0	0	0	0	0
$a_k = 3^2$	0	0	1	0	0	0	0
$a_k = 3^3$	0	0	0	1	0	0	0
$a_k = 3^4$	0	0	0	0	1	0	0
$a_k = 3^5$	0	0	0	0	0	1	0

Every a_k is nonzero. Every fixed finite digit window eventually reads only zeros.

Polynomial conditions behave differently

For a continuous polynomial f, if a convergent sequence satisfies $f(x_k) = 0$ for every k, then continuity gives $f(x) = 0$ at its limit. More generally, a compatible residue branch satisfying $f(x_m) \equiv 0 \pmod{3^m}$ reconstructs an exact root once the complete ring and its reduction maps have been identified.

The sequence 3^k here is *not* presented as a compatible branch of nonzero residues modulo 3^k: those particular residues are zero. It is a convergent sequence of nonzero elements of the complete ring. The distinction prevents an incorrect finite-level interpretation of the picture.

In an arithmetic application, local conditions may be given by quotient equations, integral structures, or other data. The completion theorem requires a proof that the chosen condition is preserved by the limit under construction. The word “admissible” does not itself provide that proof.

> **Read this as:** Compatibility reconstructs an object. Closedness is what carries the declared equations or conditions to that object.

The finite diagram displays the retreat of the first nonzero digit. The convergence assertion follows from $|3^k|_3 = 3^{-k}$, not from extrapolating a finite numerical plot.

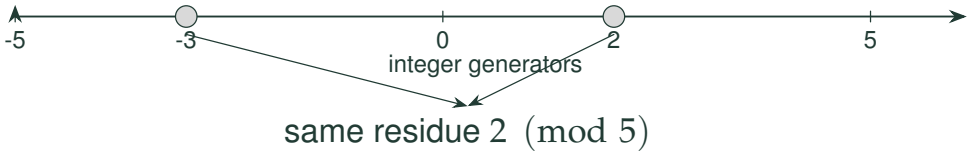
ELEMENTS VERSUS IDEALS

The sign of a norm is arithmetic data

Main text: §9, arithmetic residue-counting continuation; signed-norm comparison.

Even in the number field $\mathbb{Q}$, congruent generators need not have congruent absolute norms. The integers 2 and -3 are congruent modulo five. Their signed field norms are congruent, but their absolute values are not.

$$N_{\mathbb{Q}/\mathbb{Q}}(-3) = -3 \qquad N_{\mathbb{Q}/\mathbb{Q}}(2) = 2$$



$$N((-3)) = 3 \qquad N((2)) = 2$$

The ideal norm is positive. Passing to absolute values must retain the real sign.

The smallest counterexample is already decisive

For $K = \mathbb{Q}$,

$$N_{K/\mathbb{Q}}(2) = 2, \qquad N_{K/\mathbb{Q}}(-3) = -3 \equiv 2 \pmod 5,$$

whereas

$$N((2)) = 2, \qquad N((-3)) = 3 \not\equiv 2 \pmod 5.$$

Here $N((a))$ is the positive index of the ideal (a), equal to $|a|$. The determinant norm of an element is signed. Taking absolute values is not a ring operation and does not commute with reduction in this example.

The retained scalar/Mellin continuation counts the two real norm signs separately before reassembling its generator-to-ideal sums. Its exact comparison also includes the roots-of-unity multiplicity and the norm of the chosen inverse-class representative. Omitting those factors would change the arithmetic sum, not merely its normalization.

The finite-field norm on the following pages is a determinant in a finite algebra. It has no ordering or absolute-value sign. Its role is to control residue characters before the separate real-lattice counting argument.

> **Read this as:** Do not infer an ideal-norm residue from a generator congruence without accounting for the sign in the element-to-ideal comparison.

THE QUOTIENT ALGEBRA $\mathbb{F}_{25}$

A norm is a determinant you can compute

Main text: §9; retained scalar/Mellin continuation: actual quotient determinant norms.

Use $A = \mathbb{F}_5[u]/(u^2 - 2)$. Since two is not a square modulo five, this quotient is a field with 25 elements. Every element is uniquely $a + bu$, with $a, b \in \mathbb{F}_5$.

$$N(a + bu) = a^2 - 2b^2 \pmod 5$$

	$b = 0$	$b = 1$	$b = 2$	$b = 3$	$b = 4$
$a = 0$	0	3	2	2	3
$a = 1$	1	4	3	3	4
$a = 2$	4	2	1	1	2
$a = 3$	4	2	1	1	2
$a = 4$	1	4	3	3	4

$$u^2 = 2 \text{ in } \mathbb{F}_{25}$$

$$M_{a+bu} = \begin{pmatrix} a & 2b \\ b & a \end{pmatrix}$$

$$\#N^{-1}(0) = 1$$

$$\#N^{-1}(r) = 6 \text{ for } r \neq 0$$

All 25 elements are shown. Every nonzero element is a unit.

Coordinates, multiplication, and norm agree

In the basis $(1, u)$, multiplication by $a + bu$ has matrix

$$M_{a,b} = \begin{pmatrix} a & 2b \\ b & a \end{pmatrix}, \qquad N(a + bu) = \det M_{a,b} = a^2 - 2b^2 \pmod 5.$$

The table shows every norm value. Only the zero element has norm zero. Each nonzero norm value has exactly six preimages. This makes the later character cancellation visible without estimating any sums.

The checks enumerate all 25 matrices, verify the determinant formula, and verify norm multiplicativity on every pair of field elements. They also change to a different invertible basis and check that conjugating the multiplication matrix leaves its determinant unchanged.

This model is the reduction of the quadratic integer algebra $\mathbb{Z}[\sqrt{2}]$ modulo five. Its finite arithmetic illustrates the determinant construction used in the continuation; it is not a computation of all ideal classes or a general ideal-counting theorem.

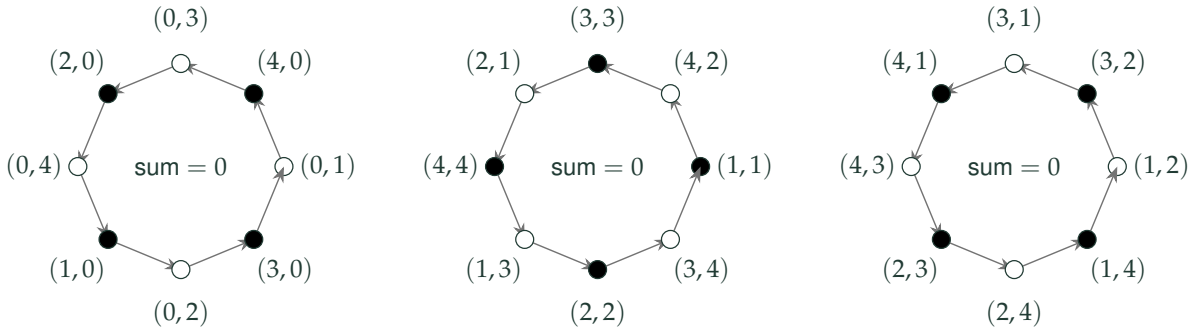
> **Read this as:** A quotient norm is built from the actual multiplication operator. Its basis independence is a determinant identity, not a choice of convenient coordinates.

A NONSCALAR UNIT WITNESS

Permutation forces a zero mean

Main text: §9; retained scalar/Mellin continuation: norm-residue lattice mean cancellation.

Let χ be the quadratic character modulo five, extended by $\chi(0) = 0$. In the field on the preceding page, choose $v = 2u$. Its norm is $N(v) = 2$, so $\chi(N(v)) = -1$.



Multiply by $v = 2u$. Its norm is 2, so $\chi(N(vz)) = -\chi(N(z))$.

Filled node: $+1$. Open node: -1. The omitted zero element has weight 0.

The whole sum is paired with its negative

Multiplication by v permutes the 25 elements. Norm multiplicativity gives

$$\sum_{z \in A} \chi(N(z)) = \sum_{z \in A} \chi(N(vz)) = -\sum_{z \in A} \chi(N(z)).$$

The sum is an integer, so it is zero. The origin contributes zero. The other 24 elements fall into three cycles of length eight under multiplication by v; the character signs alternate around each cycle.

The same reasoning works with a complex root of unity $\chi(N(v)) \neq 1$: subtracting the two equal sums gives $(1 - \chi(N(v)))S = 0$. The essential inputs are an actual permutation, a multiplicative norm, and a nontrivial character value. They are not consequences of drawing an alternating cycle.

In the arithmetic continuation the quotient is the actual additive module J/mJ associated with an ideal representative. This finite-field example isolates the cancellation mechanism while keeping that broader module and lattice comparison separate.

> **Read this as:** The witness is one unit that changes the character value. Once it exists, the permutation cancels the entire finite mean.

THE EXACT SCALAR TEST

A nontrivial character may still be exceptional

Main text: §9; scalar/Mellin checkpoint: scalar-power and nonscalar-witness boundaries.

A scalar $a \in \mathbb{F}_5^\times$ acts on a degree-d algebra with determinant a^d. Therefore the scalar cancellation route tests $\chi(a^d)$, not merely $\chi(a)$.

	$d=1$	$d=2$	$d=3$	$d=4$	$d=5$	$d=6$
trivial	no	no	no	no	no	no
quartic	yes	yes	yes	no	yes	yes
quadratic	yes	no	yes	no	yes	no
conjugate quartic	yes	yes	yes	no	yes	yes

For $\chi_k(2) = i^k$, a witness exists exactly when $kd \not\equiv 0 \pmod 4$.

A failed sufficient test is not an impossibility theorem

Let $\chi_4(2) = i$ generate the character group of $\mathbb{F}_5^\times$, and write $\chi_k = \chi_4^k$. Since 2 generates the unit group,

$$\exists a : \chi_k(a^d) \neq 1 \quad \Longleftrightarrow \quad kd \not\equiv 0 \pmod 4.$$

For the quadratic character $k = 2$ and degree $d = 2$, every scalar test is trivial. Yet page V23 supplies the nonscalar witness $2u$, proving cancellation in that specific algebra. Scalar exceptional does not mean exceptional for every possible quotient unit.

Conversely, the quartic character $k = 1$ works in degree two even though $\gcd(2,4) \neq 1$. Thus the simpler degree–totient coprimality criterion is sufficient, not necessary. With degree three, every nontrivial character passes because exponentiation by three permutes the units.

The retained scalar/Mellin continuation states its public hypothesis in the weaker exact scalar-power form. It does not claim to construct a nonscalar witness for every scalar-exceptional arithmetic character.

> **Read this as:** Test the character after the degree-power map. Nontriviality before that map can disappear after it.

A CRT MODEL FOR REPRESENTATIVE CHOICE

Avoid finitely many primes constructively

Main text: §9; retained scalar/Mellin continuation: finite-prime avoidance and automatic representatives.

A representative needed by an arithmetic argument often has to avoid a prescribed finite collection of primes. A small integer version asks for

$$x \equiv 1 \pmod 2, \qquad x \equiv 2 \pmod 3, \qquad x \equiv 4 \pmod 5.$$

There is exactly one solution modulo 30: $x = 29$.

0	1	2	3	4	5
6	7	8	9	10	11
12	13	14	15	16	17
18	19	20	21	22	23
24	25	26	27	28	29

Choose one residue in each quotient

$$a \equiv 1 \pmod 2$$

$$a \equiv 2 \pmod 3$$

$$a \equiv 4 \pmod 5$$

$$a = 29, \quad \gcd(29, 30) = 1$$

Shaded values avoid 2, 3, and 5 simultaneously.

Avoidance and ideal class are different constraints

The selected residues are nonzero modulo each forbidden prime, so $\gcd(29, 30) = 1$. The eight units modulo 30 are displayed separately. They all avoid the primes, but only one also satisfies the three specific residue prescriptions.

The continuation proves an ideal-theoretic analogue with more structure: in every prescribed inverse ideal class it constructs a nonzero integral representative whose norm is coprime to the nonzero modulus. That removes a representative-choice assumption from the public norm-character bound. It does not remove the remaining character, lattice, or Frobenius-occurrence obligations.

The integer CRT calculation is an illustration, not a proof of the ideal-class representative theorem. In particular, enumerating integers cannot impose a nontrivial ideal class. The full construction uses prime-ideal avoidance and its source comparison to the class group.

> **Read this as:** A finite avoidance argument supplies an admissible representative. It does not by itself make the arithmetic object being counted, or the desired Frobenius class, occur.

All 30 residue classes are tested exactly. The unique CRT result and the unit count are independent finite outputs of the check script.

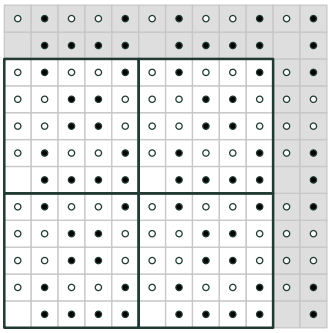
A FULL LATTICE-BOX CALCULATION

Cancellation leaves a boundary problem

Main text: §9; retained scalar/Mellin continuation: additive residue counts and norm-character bounds.

Extend the function $f(a,b) = \chi(a^2 - 2b^2 \bmod 5)$ periodically to $\mathbb{Z}^2$. Every complete 5×5 residue tile has sum zero by the finite-field calculation. Only incomplete tiles can contribute to a box sum.

$[0,12)^2 \cap \mathbb{Z}^2$



four full 5×5 cells

100 bulk points: sum 0

44 boundary points

actual total $S_{12} = 1$

$|S_L| \le L^2 - (L - (L \bmod 5))^2$ for this square-lattice model.

Separate bulk cancellation from boundary size

For $0 \le a, b < 12$, four complete tiles fill a 10×10 square and contribute zero. The remaining 44 points lie in the two boundary strips. Their exact sum is one. Since $|f| \le 1$, a bound of 44 follows without needing that final cancellation.

For a general positive integer L, write $L = 5q + r$, $0 \le r < 5$. The complete tiled square has side $L - r$, so

$$|S_L| \le L^2 - (L - r)^2 = 2Lr - r^2 \le 8L.$$

The area is $X = L^2$, making this a square-root-in-area bound for this square-box model. The proof is the exact tile decomposition; checking many L values is only a regression check of the implementation.

Number-field ideal sums involve different regions, lattice indices, signs, and generator multiplicities. Their $X^{1-1/d}$ exponent cannot be justified simply by relabeling this square picture as an ideal-norm ball.

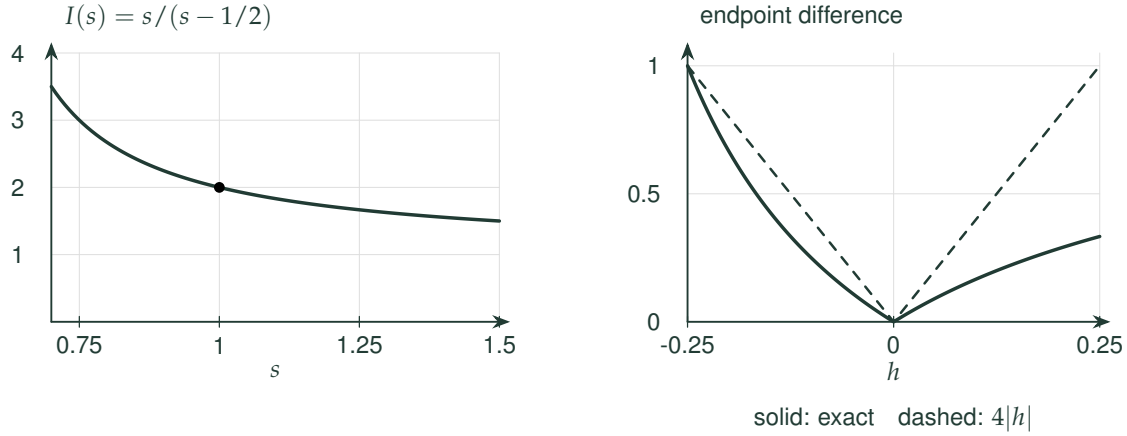
> **Read this as:** A zero finite mean removes the bulk. The remaining estimate depends on the actual geometry of the region and its boundary.

A MELLIN ENDPOINT MODEL

What a power-saving bound buys analytically

Main text: §9; retained scalar/Mellin continuation: summatory measurability and first-order Mellin control.

Use the explicit envelope $A(x) = x^{1/2}$ for $x \geq 1$. It is not an ideal-counting function. It is a model showing why a power saving can make a Mellin expression regular at $s = 1$.



Compute the endpoint rather than just naming it

For real $s > 1/2$,

$$I(s) = s \int_1^\infty A(x) x^{-s-1}\,\mathrm{d}x = \frac{s}{s-1/2}.$$

At $s = 1 + h$ this gives

$$I(1+h) - I(1) = -\frac{h}{h+1/2}, \qquad |I(1+h) - I(1)| \leq 4|h| \quad (|h| \leq 1/4).$$

The derivative at one is -2. The logarithmic moment is finite as well: $\int_1^\infty x^{-3/2} \log x\,\mathrm{d}x = 4$. These are explicit calculations; the drawn curve does not establish differentiability or convergence.

The continuation performs the corresponding analytic steps for its actual summatory functions, with their explicit measurability and growth hypotheses. The arithmetic Euler-product identity, nonzero pole residues, and the required prescribed-prime divergence remain separate. Regularity of one Mellin transform does not construct all of those inputs.

> **Read this as:** An analytic estimate controls a function already tied to arithmetic data. It is not a substitute for proving that tie.

THE REMAINING FROBENIUS STEP

Finite splitting data are not a density theorem

Main text: Theorem 9.1. Latest continuation: GeneralArithmeticCoverage and ArithmeticSystemSeparation.

For the quadratic polynomial X^2+1, reduction modulo an odd prime either gives two roots or no roots. This finite scan displays both behaviors. It is a useful model for the distinction between splitting and nonsplitting classes.

odd primes p

3	5	7	11	13	17	19	23
0	2	0	0	2	2	0	0

roots of x^2+1 in $\mathbb{F}_p$

two factorization types
already appear in this finite list

repeat one class

1	1	2
1	1	2
1	1	2

three rows, rank one

Finite character separation needs enough different evaluations, not just more evaluations.

Existence must survive arbitrary exclusions

The continuation's missing general coverage proposition asks for an actual unramified rational prime in a prescribed conjugacy class, outside any prescribed finite forbidden set, together with the corresponding absolute Frobenius witness. A list of small primes does not establish that quantifier.

The finite character argument on page V15 also needs samples that jointly separate the constituents. Repeating the identity-class row $(1,1,2)$ of the S_3 table gives rank one no matter how often it is repeated. Many samples are not necessarily informative samples.

The later checked rational-base separator composes its constructed arithmetic coverage theorem with finite linear algebra. The earlier scalar/Mellin route retained a coverage premise. The displayed scan certifies neither route: it is a finite example, not an occurrence theorem.

> **Read this as:** A separating group element, a witnessed Frobenius at one prime, and a new prime after every finite exclusion are distinct levels of evidence.

The scan exhaustively checks scalar roots for the shown primes. No density, infinitude, or general prescribed-conjugacy statement is inferred from it.

SCALAR ABSENCE VERSUS REPRESENTATION EXISTENCE

Changing coefficients changes the problem

Main text: Example 8.4 and the declared coefficient categories in §§8–9.

Over $\mathbb{F}_3$, no scalar solves $x^2 + 1 = 0$. But the matrix

$$A = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} \in M_2(\mathbb{F}_3) \quad \text{satisfies} \quad A^2 + I = 0.$$

And over $\mathbb{F}_9 = \mathbb{F}_3[u]/(u^2 + 1)$, the scalars u and $-u$ solve the same polynomial.

Scalar over $\mathbb{F}_3$	Matrix over $\mathbb{F}_3$	Scalar over $\mathbb{F}_9$
$x = 0, 1, 2$	$\begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$	$u^2 = -1$
$x^2 + 1 = 1, 2, 2$		roots $u, -u$
no solution	$A^2 + I = 0$	coefficient extension

Fix the coefficient field, dimension, and point category before declaring an occupancy failure.

Three outputs with different types

The failed scalar problem over $\mathbb{F}_3$, the successful matrix problem over $\mathbb{F}_3$, and the successful scalar problem over $\mathbb{F}_9$ are not contradictory answers. They have different allowed objects. The matrix is multiplication by u in the $\mathbb{F}_3$-basis $(1, u)$ of the extension field.

This also has a characteristic-zero analogue. The matrix with the same entries gives a rational two-dimensional representation of a cyclic group of order four. Its scalar eigenvalues require a field containing a square root of -1. No one-dimensional rational eigencharacter is created by writing the two-dimensional matrix.

Accordingly, the scalar companion failure on page V8 is a counterexample only to the fixed scalar problem. It is not a counterexample to a representation problem that permits coefficient extension or a different dimension. Those permissions must be declared before comparing fibers across primes.

> **Read this as:** An empty fiber is meaningful only after its coefficient ring and object type are fixed. Do not repair an empty fiber by silently changing either.

FINAL READER'S CHECKLIST

Read the evidence at its actual scope

Main text: Appendix D. Formal guide: statement register, dependency sources, and replay boundary.

The reader carries three distinct contracts. A larger date or a larger numerator does not make their hypotheses interchangeable.

Record	What it establishes
Historical 17/20	Block 27 and the scalar/Mellin continuation. Claims 9.1, 9.2, and 11.1 were partial in those revisions.
Checked 20/20	The amended conditional occupied-system assembly. Completed source members, closed target profiles, compatible carrier presentations, and local comparisons remain explicit inputs.
Later 19/20	A stronger residual-system attempt with no promoted terminal endpoint and an unavailable pinned Lean executable. It is a different contract, not a failed recompilation of the checked source.

Two conclusions, not an unstated identification

The terminal theorem constructs a transported target family and independently recovers a summand of the supplied carrier. To identify them requires an explicit comparison. The carrier result is pointwise at an arbitrary coefficient place. Recovering carriers at every place needs the relevant presentations there. Zero multiplicity is allowed. Nonzero occupancy needs a separate nonzero premise.

What the examples verify

The appendix supplies reproducible finite examples and exact Python checks. No plotted pixel decides a congruence, rank, determinant, root, or character sum. The original Lean sources and recorded audits are preserved. The October 10 editorial release performed no fresh Lean compilation. The separate October 20 production replay is reported in the Lean guide; it introduces no new formal theorem.

> **Read this as:** A completed conditional theorem is a theorem with its hypotheses, not a claim that every arithmetic application supplies them.