Strong Artin

Cyclic attachment and finite realization

Komorebi P.

READER EDITION 1.0 / 8 OCTOBER 2026

A detailed account of the representation theory, exact selectors, realization criteria, and certified local lifting sectors. The global automorphic attachment remains open.

EXACT IDENTITIES / FINITE EVIDENCE / OPEN CONJECTURE

Contents

1 The conjecture and the work represented here **3**
1.1 Two strengths of attachment . . . 3
1.2 The classical range and the new boundary . . . 4
1.3 What the finite analysis proves . . . 4
1.4 How to read the formal and computational record . . . 5

2 Cyclic extension and descent torsors **6**

3 Coefficient selectors in a cyclic fibre **7**

4 Two no go theorems for invisible cyclic twists **9**

5 Local extension spectra **10**
5.1 Fourier coordinates of an extension spectrum . . . 13
5.2 Two basic extension spectrum examples . . . 14

6 Strong Artin as global spectrum alignment **14**

7 Every descent of an Artin base change is locally Artin **15**

8 Finite Satake descent **16**
8.1 Cyclotomic rigidity . . . 18

9 Finite polynomial alphabets and separating primes **18**

10 Finite trace alphabets in residual characteristic **19**

11 One prime residual attachment **22**

12 Representation theoretic collapse at the quotient prime **24**

13 The quotient prime residual congruence **24**

14 Finite Satake rigidity for l adic realizations **26**

15 Cohomological and weight raising realization gates **27**
15.1 A p adic family realization gate . . . 28

16 A strengthened recursive Strong Artin theorem **29**

17 The weight zero realization bottleneck **30**
17.1 A complex pseudocharacter realization gate . . . 31

18 A finite character separator on the Galois torsor **32**

19 Automorphic fibre separators **33**

20 Positive saturation and Artin holomorphy **33**

21 Attempted completion routes and exact failure mechanisms **34**
21.1 Total automorphic induction 35
21.2 Central and stable local invariants 35
21.3 Functional equation comparison 35
21.4 Cohomological regularization 35
21.5 Repeated weight zero and the lifting hypotheses 35
21.6 The surviving routes 36

22 Three focused theorems that would close the solvable case **36**

23 What has actually been reduced **37**

24 Verification **37**

25 Exact boundary **38**

26 From finite realization to local deformation **38**

27 Formal mechanisms and their exact hypotheses **39**
27.1 A cyclic cocycle is constrained by the norm 39
27.2 A square-zero defect gives a genuine obstruction class 40
27.3 Local conjugacy leaves a second gluing obstruction 42
27.4 The reachable equation target is determined by syzygies 42
27.5 Ordered powers have a uniform ultrametric remainder 44
27.6 An exact inverse turns the local image into one finite test 46
27.7 A noncommuting unipotent square example 47
27.8 Finite levels give a coherent branch only for a joint problem 48
27.9 What the formal receipts establish 49

28 Worked representation and selector calculations **50**
28.1 Why the prime cyclic Clifford alternative is exact 50
28.2 A rank three coefficient test that determinant misses 51
28.3 Ordered multiplicities and a common shift 52
28.4 A trace encoding, with both hypothesis failures exposed 53
28.5 A fixed alphabet and what a separating prime actually does 53

29 Statement map and the remaining arithmetic obligation **54**

1 The conjecture and the work represented here

A finite Galois representation presents a strikingly discrete object. Its matrices range over a finite group. At an unramified prime, its Frobenius eigenvalues are roots of unity. The Strong Artin conjecture asks that this discrete arithmetic information occur inside the automorphic spectrum, with the correct dimension, local parameters, and global analytic structure. Knowing a finite list of possible eigenvalues does not supply the automorphic representation that should carry them.

This piece studies that distinction at a cyclic descent step. There are exact finite choices on both sides of the proposed correspondence. The choices can be described, separated, and tested. What remains is a theorem that identifies a choice on the automorphic side with the prescribed arithmetic object. Later local deformation and matrix lifting results sharpen several algebraic parts of that problem. They do not manufacture the global automorphic carrier.

1.1 Two strengths of attachment

Let F be a number field, let $G_F = \mathrm{Gal}(\overline{F}/F)$, and let

$$\rho : G_F \longrightarrow \mathrm{GL}_n(\mathbb{C})$$

be a continuous representation with finite image. Continuity is taken for the usual complex topology. Such a representation factors through the Galois group of a finite extension. Throughout the conjecture discussion, ρ is irreducible. The trivial one dimensional representation is excluded when discussing entireness of an Artin L function, since its L function is the Dedekind zeta function.

At a finite place where ρ is unramified, fix the Frobenius convention used by the normalized local Langlands correspondence. Write

$$P_{\rho,v}(T) = \det\bigl(TI_n - \rho(\mathrm{Frob}_v)\bigr).$$

For an unramified automorphic local component π_v, write $P_{\pi,v}(T)$ for the monic polynomial of its normalized Satake conjugacy class. The convention is chosen so that an attached pair has the same polynomial. This keeps a possible geometric versus arithmetic Frobenius inversion from entering a comparison silently.

Conjecture 1.1 (Strong Artin, unramified form)**.** *Every irreducible finite image complex representation ρ of G_F has a cuspidal automorphic representation π of $\mathrm{GL}_n(\mathbb{A}_F)$ such that*

$$P_{\pi,v}(T) = P_{\rho,v}(T)$$

for all finite places outside a finite set.

We call this weak attachment, where weak refers to the comparison at almost all unramified places. It remains a powerful global assertion. It implies the expected unramified Euler factors and identifies the cuspidal automorphic representation uniquely by strong multiplicity one.

Conjecture 1.2 (Strong Artin, all local form)**.** *The representation π can be chosen so that*

$$\mathrm{rec}_{F_v}(\pi_v) \cong \rho|_{W_{F_v}}$$

for every place v, with zero monodromy at finite places.

At a finite place, the right side is viewed as a Weil Deligne representation with monodromy operator zero. At an archimedean place, the statement uses the corresponding archimedean parameter. The all local assertion includes ramified and infinity data that do not appear in the definition of weak attachment. In this piece an unramified comparison is never silently promoted to an all local theorem. A residual realization will first give weak attachment. The remaining finite set of local parameters must then be checked or covered by an additional compatibility theorem.

1.2 The classical range and the new boundary

Class field theory handles dimension one. Cyclic base change and automorphic induction provide powerful higher dimensional constructions. An induction step is available when its field tower and cuspidality conditions fall within the established induction formalism. We do not assert that every abstract monomial character has an automorphic realization merely because it is induced from a linear character. The relevant field and induction chain must be covered by a theorem [1].

The solvable two dimensional sectors treated by Langlands and Tunnell, further examples such as the solvable $\mathrm{GO}(4)$ range, and other established cases show that the conjecture is far from an empty expectation [6, 7, 5]. A recent preprint of Wang states the generalized octahedral GL_3 case with projective image $C_3^2 \rtimes \mathrm{SL}(2,3)$, together with a primitive solvable six dimensional consequence [29]. Its stated results have been checked against the submitted source. Its full proof has not been independently audited for this edition. These particular sectors must therefore not be described as untouched simply because the general problem remains open.

Arthur and Clozel's prime degree cyclic constructions and the descent and fibre statements clarified by Clozel and Rajan are the external automorphic inputs used below [1, 2]. For a cyclic extension, an invariant unitary cuspidal representation descends. When its base change is cuspidal, the cuspidal descents differ by quotient characters, and these twists are distinct. The more general solvable base change statement has extra twisting information. We use the precise prime cyclic case in the recursive argument.

The finite representation arguments have a parallel shape. If $H \triangleleft G$ and $G/H \cong C_p$, an irreducible representation either restricts to a full orbit of p subgroup types or restricts irreducibly to one invariant type. In the first case induction reconstructs the representation. In the second case an extension must be selected. The second case is where two torsors meet. Each torsor exists, and each has p elements. There is no preferred origin in the statement that they are torsors under the same character group.

1.3 What the finite analysis proves

The core argument retains the detailed residual realization analysis of the earlier project source. Stable subgroup types carry multiplicity spaces on which the cyclic quotient acts. Their ordered character multiplicities are recorded in polynomials

$$P_{\tau,V}(z) = \sum_{a \in \mathbb{Z}/p\mathbb{Z}} m_{\tau,a}(V) z^a.$$

A quotient twist rotates every such polynomial by the same amount. Moving subgroup types carry induced blocks and do not retain a quotient label. This gives a complete extension comparison, not just a determinant or total trace test.

At each place, a finite alignment set records the quotient labels that attach the local automorphic parameter to the target. The all local cyclic step is equivalent to a common label in all these sets. A nonempty exact alignment set is a coset of a subgroup of C_p, so it is a singleton or the whole label group. It may also be empty. The reduction makes the missing statement precise. It does not show that the local requirements are compatible.

A second chain starts with a descent whose base change is already of Artin type. Every descended local parameter has finite image. Every unramified Satake root belongs to one finite cyclotomic group. Consequently the Hecke polynomial alphabet is finite, its coefficients are integral, and all but finitely many coefficient primes separate its entries. A finite trace argument then bounds a semisimple residual image. A sufficiently large residual realization lifts in cross characteristic and is forced onto the exact characteristic zero polynomial alphabet. At a good separating prime, even the smaller direct torsor argument recovers the cyclic offset.

The same trace pairing gives an especially direct proof for an ordinary semisimple ℓ adic realization. Chebotarev and continuity extend a finite Frobenius trace alphabet to every image element. The semisimple image algebra then embeds the image in a finite Cartesian power of that alphabet. Thus the image is finite. The existence of the ℓ adic representation remains the input. The finite image conclusion is the output.

A third phenomenon occurs at the quotient prime p. The order p characters all reduce to the trivial character, and the unramified cyclic base change identities force the automorphic and target polynomial systems to agree modulo a coefficient prime above p. This congruence cannot choose a characteristic zero label. It is the beginning of a lifting problem. It supplies neither a residual Galois representation for an arbitrary automorphic descent nor an unrestricted automorphy lifting theorem for repeated weight zero parameters.

1.4 How to read the formal and computational record

The second part develops the later local results. Cyclic first cohomology describes infinitesimal conjugacy ambiguity under its stated norm condition. Second cohomology records a square zero extension defect. Syzygy calculations identify redundant local prescriptions after an explicit arithmetic comparison is supplied. Ordered matrix Hensel lifting gives an actual root from one finite congruence when a two sided inverse derivative certificate is available and the prescribed ball is sufficiently small. The noncommutative order of the derivative is part of the theorem.

These are useful closed sectors. They are not a proof that every arithmetic local obstruction vanishes, that every simultaneous prescription is realizable, or that a global automorphic carrier exists. The distinction is carried through the hypotheses, examples, and final statement map. A noncommuting positive matrix control and a rejected target show the local criterion working in both directions. A genuinely singular derivative exhibits the sector outside that criterion.

The accompanying sources retain historical verification reports and the precise formal statement boundaries. A historical Lean certificate verifies the theorem encoded in its source under its hypotheses. It does not verify an unstated arithmetic identification or fill a missing global premise. New computations in this edition are bounded checks of exact finite identities and figures. Their ranges are documented, and their counts are not added to overlapping historical suites as though every check were independent.

No unrestricted Strong Artin theorem and no global novelty claim is made here. The

point is a substantial, inspectable body of mathematics around the remaining gate: exact representation theory, finite selectors, rigidity, local lifting, and a clear statement of what a successful global realization theorem would still have to provide.

2 Cyclic extension and descent torsors

Let G be a finite group, let $H \triangleleft G$, and assume

$$G/H \cong C_p$$

for a prime p. Let $X = \mathrm{Hom}(G/H, \mathbb{C}^\times)$, which is cyclic of order p.

Proposition 2.1 (Galois extension torsor)**.** *Let τ be an irreducible complex representation of H that is invariant under G, and suppose τ extends to an irreducible representation of G. Then the set of isomorphism classes of extensions of τ to G is an X torsor under tensor product.*

Proof. Fix one extension ρ. Every $\chi \in X$ is trivial on H, so $\rho \otimes \chi$ restricts to τ. Conversely, let ρ' be another extension. Since both restrictions are irreducible and isomorphic, Schur's lemma shows that after one choice of an H intertwiner, the ratio between the two actions is scalar on each coset. Those scalars define a character of G/H, so $\rho' \cong \rho \otimes \chi$.

The action is free. If $\rho \cong \rho \otimes \chi$, an intertwiner restricts to a scalar on H. Commuting it through one element outside H gives $\chi = 1$. □

Classical input 2.2 (Prime cyclic cuspidal descent and fibre). Let E/F be cyclic of prime degree, and let Π be an invariant unitary cuspidal representation of $\mathrm{GL}_n(\mathbb{A}_E)$. Then Π has a cuspidal descent to F. Its cuspidal descents form a torsor under the Hecke characters corresponding by class field theory to $\mathrm{Gal}(E/F)$. In particular, distinct quotient twists of one descent remain distinct. We use the prime cyclic construction of Arthur and Clozel and the fibre statement and cyclic descent clarification of Clozel and Rajan [1, 2]. The assertion concerns a cuspidal base change. It is not a claim that quotient twisting acts freely on every cuspidal representation.

Definition 2.3 (Attached cyclic descent)**.** Let E/F be cyclic of prime degree p, and let ρ be an irreducible finite image representation of G_F whose restriction $\tau = \rho|_{G_E}$ is irreducible. Suppose τ is automorphic with cuspidal representation Π of $\mathrm{GL}_n(\mathbb{A}_E)$. An attached cyclic descent is a cuspidal descent π of Π such that

$$\mathrm{rec}_{F_v}(\pi_v) \cong \rho|_{W_{F_v}}$$

for every place v of F.

Proposition 2.4 (One attachment determines the whole fibre)**.** *If one attached cyclic descent exists, then every Galois extension in the torsor of theorem 2.1 is attached to the corresponding quotient twist of that descent.*

Proof. Local class field theory identifies a quotient character $\chi \in X$ with the corresponding Hecke character. Local Langlands is compatible with character twisting, [10]. Hence

$$\mathrm{rec}(\pi_v \otimes \chi_v) \cong \mathrm{rec}(\pi_v) \otimes \chi_v \cong \rho_v \otimes \chi_v.$$

The result holds at every place. □

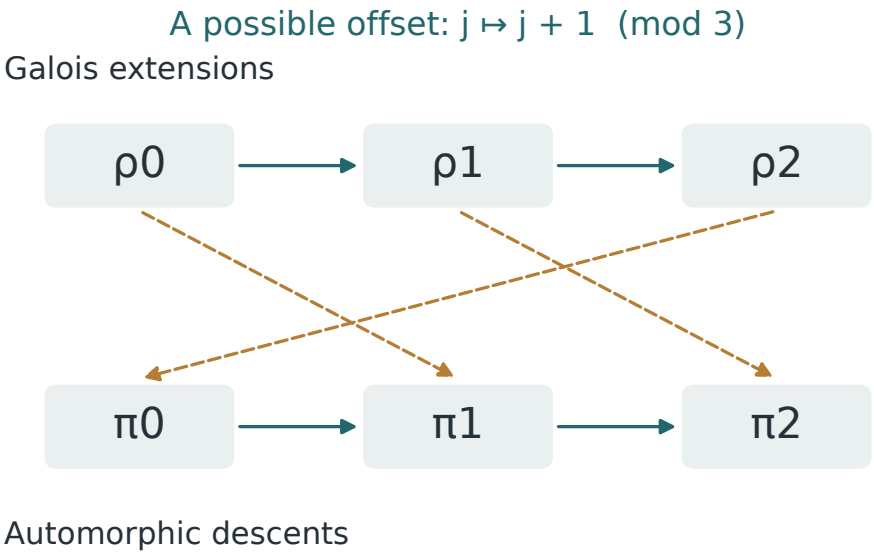


Figure 1: A quotient character rotates the two cyclic torsors. Neither torsor has a preferred origin, so the common character action does not specify their attachment.

Corollary 2.5 (Clifford recursion reduction)**.** *Suppose a finite solvable image representation is followed down a prime cyclic Clifford tower. Every induction step is handled by automorphic induction. If attached cyclic descent holds at every invariant extension step, then the original representation satisfies Strong Artin.*

Proof. Start at a one dimensional bottom constituent, which is automorphic by class field theory. Move upward through the tower. Use automorphic induction at each induction step and attached cyclic descent at each invariant extension step. The local parameters remain those of the Galois representation at every stage. □

3 Coefficient selectors in a cyclic fibre

Let ξ be a generator of X, and identify its values with pth roots of unity. For a matrix $A \in \mathrm{GL}_n(\mathbb{C})$, write

$$\det(TI - A) = \sum_{k=0}^{n} (-1)^k e_k(A) T^{n-k}.$$

Proposition 3.1 (Coefficient rotation)**.** *For $\zeta \in \mathbb{C}^\times$,*

$$\boxed{e_k(\zeta A) = \zeta^k e_k(A).}$$

Proof. The coefficient $e_k(A)$ is the trace of $\bigwedge^k A$. Scalar multiplication by ζ acts on $\bigwedge^k \mathbb{C}^n$ by ζ^k. □

Definition 3.2 (Active coefficient degrees)**.** Define

$$S(A) = \{k \in \{1, \dots, n\} : e_k(A) \neq 0\}.$$

Theorem 3.3 (One polynomial separator)**.** *Let ζ range over μ_p. The twists ζA have pairwise distinct characteristic polynomials exactly when $S(A)$ contains a degree not divisible by p.*

Proof. By theorem 3.1, the stabilizer of the characteristic polynomial is

$$\{\zeta \in \mu_p : \zeta^k = 1 \text{ for every } k \in S(A)\}.$$

Since p is prime, a nontrivial pth root satisfies $\zeta^k = 1$ exactly when $p \mid k$. □

Corollary 3.4 (Trace and determinant selectors)**.** *A nonzero trace separates all cyclic twists. The determinant separates all cyclic twists when $p \nmid n$.*

Proof. Use degree one for the trace and degree n for the determinant in theorem 3.3. □

Theorem 3.5 (Determinant selector for extension torsors)**.** *Assume $p \nmid n$. In the Galois extension torsor of an n dimensional irreducible representation, the determinant distinguishes all p extensions. In the automorphic descent torsor, the central character distinguishes all p descents. Both transformations are equivariant under the nth power map on X.*

Proof. For a quotient character χ,

$$\det(\rho \otimes \chi) = \det(\rho)\chi^n.$$

The central character satisfies the identical formula

$$\omega_{\pi\otimes\chi} = \omega_\pi \chi^n.$$

Since X has prime order p and $p \nmid n$, the map $\chi \mapsto \chi^n$ is an automorphism of X. □

Theorem 3.6 (Unique determinant matched descent)**.** *Use the cyclic Artin setup of theorem 2.3, and assume $p \nmid n$. Let π_0 be any cuspidal descent of Π. Then there is a unique quotient character $\chi \in X$ such that*

$$\boxed{\omega_{\pi_0\otimes\chi} = \det \rho}$$

under global class field theory.

Proof. The base change of ω_{π_0} is the central character of Π. Since Π is attached to $\tau = \rho|_{G_E}$, that character equals $\det \tau$. The base change of $\det \rho$ is also $\det \tau$. Hence

$$\delta = (\det \rho)\omega_{\pi_0}^{-1}$$

is a character killed by base change to E, so $\delta \in X$. We need $\chi^n = \delta$. Since X has order p and $p \nmid n$, the nth power map on X is bijective. Thus there is one and only one such χ. □

Corollary 3.7 (Rank prime sieve)**.** *At a cyclic prime step of degree p, determinant data leaves no torsor ambiguity unless $p \mid n$. The only quotient primes invisible to determinant selection are prime divisors of the rank.*

Proof. This is theorem 3.5. □

Remark 3.8. Matching determinant and central character selects a unique candidate descent when $p \nmid n$. It does not prove that the remaining local characteristic polynomials match. The selection theorem removes the finite torsor ambiguity, not the global realization problem.

4 Two no go theorems for invisible cyclic twists

The determinant selector is optimal among algebraic central characters. When the quotient prime divides the rank, applying more algebraic functors does not repair it.

Proposition 4.1 (Determinant character of an algebraic representation)**.** *Let*

$$r : \mathrm{GL}_n(\mathbb{C}) \longrightarrow \mathrm{GL}_N(\mathbb{C})$$

be an irreducible algebraic representation. Suppose the center acts with weight w, so

$$r(zI_n) = z^w I_N.$$

There is an integer a such that

$$\boxed{\det r(g) = \det(g)^a}$$

and

$$\boxed{wN = an.}$$

Proof. Every algebraic character of GL_n is an integral power of the determinant. Thus $\det \circ r = \det^a$ for one integer a. Evaluate both sides at zI_n. The left side is z^{wN}, while the right side is z^{an}. □

Theorem 4.2 (Algebraic central characters are blind when $p \mid n$)**.** *Let χ have order p, and assume $p \mid n$. For every irreducible algebraic representation r of GL_n,*

$$\boxed{\det r(\rho \otimes \chi) = \det r(\rho).}$$

The same statement holds on every irreducible summand of an arbitrary finite dimensional algebraic representation.

Proof. The scalar twist acts on $r(\rho)$ by χ^w. Hence its determinant changes by χ^{wN}. By theorem 4.1, $wN = an$, which is divisible by p. Thus $\chi^{wN} = 1$. Apply the argument to each irreducible summand. □

Corollary 4.3 (Optimality of the rank prime sieve)**.** *When $p \mid n$, no finite collection of determinant or central character tests obtained from algebraic functors on GL_n can distinguish the order p scalar extension twists.*

Proof. Every test in the collection is invariant by theorem 4.2. □

Classical input 4.4 (Stable equality of epsilon factors). Let V, V' be finite dimensional complex Weil representations with finite inertia image, of equal dimension and determinant. For every sufficiently ramified auxiliary character η,

$$\epsilon(V \otimes \eta, \Psi) = \epsilon(V' \otimes \eta, \Psi).$$

Here Ψ is a fixed nontrivial additive character and the same normalization is used on both sides. Deligne and Henniart's stability theorem applies to the virtual representation $B = V - V'$, which has dimension zero and determinant one. Their Theorem 4.6 gives the ratio as a determinant evaluation once the slope of η exceeds twice the largest relevant slope of B. That ratio is one [25]. We use this equality and avoid an unspecified choice of the auxiliary determinant evaluation element.

Theorem 4.5 (Highly ramified epsilon twists are blind)**.** *Let $V' = V \otimes \chi$, where χ has order p, and assume $p \mid \dim V$. Then V and V' have the same determinant. For every sufficiently ramified auxiliary character η,*

$$\boxed{\epsilon(V \otimes \eta, \Psi) = \epsilon(V' \otimes \eta, \Psi).}$$

Proof. The determinant changes by $\chi^{\dim V} = 1$. The two representations have the same dimension and determinant. Apply classical input 4.4. □

Remark 4.6. When the quotient prime divides the rank, neither algebraic central characters nor sufficiently ramified epsilon twists see the cyclic label. Any complete attachment theorem must use bounded depth noncentral local data, a genuine Galois realization, or an automorphy lifting theorem that controls the corresponding local component.

Theorem 4.7 (The total induced object is torsor blind)**.** *Let $H \triangleleft G$ with cyclic quotient character group X, and let V be an extension of an H representation. Then*

$$\boxed{\bigoplus_{\chi \in X} V \otimes \chi \cong \operatorname{Ind}_H^G \operatorname{Res}_H V.}$$

The left side is unchanged when V is replaced by another member of its extension torsor. Likewise, for cyclic base change,

$$\boxed{\coprod_{\chi \in X} \pi \otimes \chi \cong \mathrm{AI}_{E/F}(\mathrm{BC}_{E/F}\, \pi)}$$

inside the established automorphic induction formalism. Therefore comparison of the total induced objects cannot determine the cyclic attachment offset.

Proof. The restriction of every quotient character to H is trivial. The regular representation of G/H decomposes as the direct sum of all characters in X. Hence

$$\operatorname{Ind}_H^G \operatorname{Res}_H V \cong V \otimes \mathbb{C}[G/H] \cong \bigoplus_{\chi \in X} V \otimes \chi.$$

Replacing V by $V \otimes \chi_0$ only permutes the summands. The automorphic identity is the cyclic automorphic induction and base change analogue. It has the same permutation invariance. □

5 Local extension spectra

The determinant is only one aggregate of the local cyclic extension labels. Clifford theory gives the complete finite invariant.

Let $G/H \cong C_p$, and choose a quotient character ξ. Let G be finite and let V be a finite dimensional semisimple complex representation of G.

Definition 5.1 (Stable and moving types)**.** An irreducible representation $\tau \in \mathrm{Irr}(H)$ is stable when $\tau^g \cong \tau$ for every $g \in G$. Otherwise its orbit has size p and is called moving.

Lemma 5.2 (Extension of a stable type)**.** *Every stable irreducible complex representation of H extends to G.*

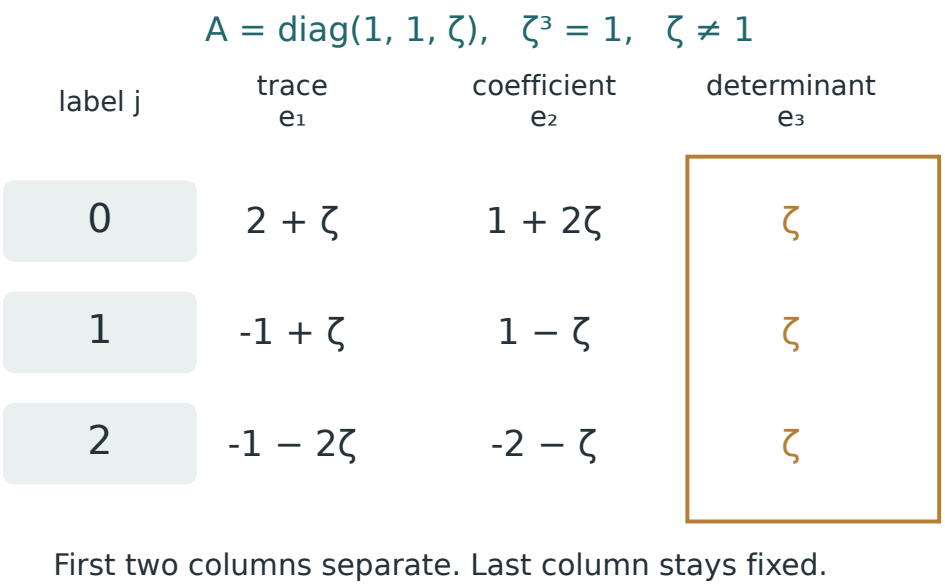


Figure 2: At quotient prime three and rank three, determinant is fixed by the scalar twist. Lower coefficient degrees can still detect the label when their coefficients are nonzero.

Proof. Choose $g \in G$ whose class generates G/H, so $g^p \in H$. Stability gives an invertible intertwiner A with

$$A\tau(h)A^{-1} = \tau(ghg^{-1}) \qquad (h \in H).$$

Iterate this identity p times. The operator $\tau(g^p)^{-1}A^p$ commutes with $\tau(H)$, so Schur's lemma makes it a nonzero scalar cI. Also A commutes with $\tau(g^p)$, by the same intertwining identity at $h = g^p$. Choose $b \in \mathbb{C}^\times$ with $b^p = c^{-1}$ and put $R = bA$. Then

$$R\tau(h)R^{-1} = \tau(ghg^{-1}), \qquad R^p = \tau(g^p).$$

These are exactly the conjugation and closing relations for the coset presentation of G. Assigning $g \mapsto R$ and $h \mapsto \tau(h)$ therefore defines a representation of G extending τ. Its restriction is irreducible, so the extension is irreducible. In cohomological language the scalar cyclic obstruction vanishes because a nonzero complex number has a pth root. This uses the coefficient field, and must not be transplanted unchanged into an arbitrary integral or residual lifting problem. □

Fix one extension $\widetilde{\tau}$ for every stable τ.

Definition 5.3 (Extension spectrum polynomial)**.** For a stable type τ, the multiplicity space

$$M_\tau(V) = \mathrm{Hom}_H(\tau, V)$$

becomes a representation of G/H after comparing the action on V with $\widetilde{\tau}$. Write

$$M_\tau(V) \cong \bigoplus_{a \in \mathbb{Z}/p\mathbb{Z}} (\xi^a)^{\oplus m_{\tau,a}(V)}.$$

Define

$$\boxed{P_{\tau,V}(z) = \sum_{a \in \mathbb{Z}/p\mathbb{Z}} m_{\tau,a}(V) z^a}$$

in $\mathbb{Z}[z]/(z^p - 1)$.

Proposition 5.4 (Effect of quotient twisting)**.** *For every $j \in \mathbb{Z}/p\mathbb{Z}$,*

$$\boxed{P_{\tau,V\otimes\xi^j}(z) = z^j P_{\tau,V}(z).}$$

Moving orbit blocks are unchanged up to isomorphism by quotient twisting.

Proof. The quotient character acts only on the multiplicity space, shifting every character label by j. For a moving type, its inertia group is H, and the corresponding irreducible block is induced from H. Tensoring an induced representation by a character trivial on H produces an isomorphic induced representation. □

Theorem 5.5 (Cyclic extension spectrum criterion)**.** *Let V and W be semisimple representations of G with*

$$\operatorname{Res}_H V \cong \operatorname{Res}_H W.$$

Then

$$\boxed{W \cong V \otimes \xi^j}$$

for some $j \in \mathbb{Z}/p\mathbb{Z}$ if and only if, for every stable irreducible H type τ,

$$\boxed{P_{\tau,W}(z) = z^j P_{\tau,V}(z).}$$

The moving orbit blocks require no additional condition.

Proof. Decompose the restriction to H into G orbits of irreducible types. For a moving orbit, Clifford theory gives one irreducible induced block, with multiplicity determined by the restriction. Quotient twisting fixes its isomorphism class.

For a stable type τ, the full τ isotypic contribution is

$$\widetilde{\tau} \otimes M_\tau(V).$$

The corresponding contribution for W is determined by $M_\tau(W)$. Tensoring by ξ^j shifts every quotient character in every stable multiplicity space by the same j. Thus the displayed polynomial identities are necessary and sufficient on each stable block. Direct sum over all blocks gives the result. □

Definition 5.6 (Alignment set)**.** For V, W with the same restriction to H, define

$$\boxed{\mathcal{J}_H(V,W) = \{j \in \mathbb{Z}/p\mathbb{Z} : P_{\tau,W} = z^j P_{\tau,V} \text{ for every stable } \tau\}.}$$

Corollary 5.7 (Structure of the alignment set)**.** *The alignment set is empty or a coset of the common rotational stabilizer*

$$\bigcap_\tau \{j : z^j P_{\tau,V} = P_{\tau,V}\}.$$

In particular, one stable type with aperiodic spectrum makes the alignment index unique whenever it exists.

Proof. If j_0 is one solution, every other solution differs from it by a rotation stabilizing every spectrum polynomial. The converse is immediate. □

Example 5.8 (Irreducible stable restriction)**.** If $\mathrm{Res}_H V$ is irreducible, its multiplicity polynomial is one monomial. Hence $\mathcal{J}_H(V, W)$ is a singleton. This recovers the ordinary extension torsor of theorem 2.1.

5.1 Fourier coordinates of an extension spectrum

Fix a primitive pth root ζ. For

$$P(z) = \sum_{a \in \mathbb{Z}/p\mathbb{Z}} m_a z^a,$$

define

$$\widehat{P}(r) = \sum_{a \in \mathbb{Z}/p\mathbb{Z}} m_a \zeta^{ar}, \qquad r \in \mathbb{Z}/p\mathbb{Z}.$$

Proposition 5.9 (Fourier trace interpretation)**.** *Let g be a lift of a generator of G/H. For a stable type τ,*

$$\boxed{\widehat{P}_{\tau,V}(r) = \mathrm{Tr}\left(g^r \mid M_\tau(V)\right).}$$

If V is twisted by ξ^j, then

$$\boxed{\widehat{P}_{\tau,V\otimes\xi^j}(r) = \zeta^{jr}\widehat{P}_{\tau,V}(r).}$$

Proof. The quotient generator acts on the ξ^a summand by ζ^a. Taking the trace of its rth power gives the Fourier sum. The twist multiplies the action of g^r by ζ^{jr}. □

Corollary 5.10 (Fourier attachment test)**.** *Suppose $\mathrm{Res}_H V \cong \mathrm{Res}_H W$. Then $W \cong V \otimes \xi^j$ if and only if*

$$\boxed{\widehat{P}_{\tau,W}(r) = \zeta^{jr}\widehat{P}_{\tau,V}(r)}$$

for every stable type τ and every $r \in \mathbb{Z}/p\mathbb{Z}$.

Proof. The discrete Fourier transform on $\mathbb{Z}/p\mathbb{Z}$ is invertible. Thus the Fourier identities are equivalent to the polynomial shift identities in theorem 5.5. □

Corollary 5.11 (One nonzero Fourier mode selects the shift)**.** *Let $r \neq 0$ modulo p. If*

$$\widehat{P}_{\tau,V}(r) \neq 0,$$

then any possible attachment shift is uniquely determined by the ratio

$$\boxed{\frac{\widehat{P}_{\tau,W}(r)}{\widehat{P}_{\tau,V}(r)} = \zeta^{jr}.}$$

Proof. Since p is prime, multiplication by the nonzero index r is invertible modulo p. Thus ζ^{jr} uniquely determines j. □

Definition 5.12 (Local spectral obstruction)**.** The local spectral obstruction vanishes when, for every stable type τ, the zero Fourier modes of the two multiplicity spectra occur at the same indices and all ratios at the remaining indices satisfy

$$\widehat{P}_{\tau,W}(r)/\widehat{P}_{\tau,V}(r) = \zeta^{jr}$$

for one common j. When $\widehat{P}_{\tau,V}(r) = 0$, the condition requires $\widehat{P}_{\tau,W}(r) = 0$, rather than omitting the comparison. Equivalently, all the Fourier identities in theorem 5.10 hold. If no such j exists, the obstruction is nonzero.

Corollary 5.13 (Finite local decision procedure)**.** *At one cyclic local step, attachment up to quotient twist is decided by finitely many traces on the stable multiplicity spaces.*

Proof. There are finitely many stable types and p Fourier modes. Apply theorem 5.10. □

5.2 Two basic extension spectrum examples

Example 5.14 (A stable cyclic type)**.** Let

$$G = C_3 \times C_2, \qquad H = C_3.$$

Fix a nontrivial character τ of H. It is stable under G. After choosing the extension with trivial C_2 coordinate, the two one dimensional extensions have spectra

$$P_{\tau,V_0}(z) = 1, \qquad P_{\tau,V_1}(z) = z.$$

The alignment set is the singleton containing the unique quotient twist carrying one extension to the other.

Proof. The quotient C_2 has the trivial and sign characters. The multiplicity space is one dimensional, so its quotient character is the complete extension label. □

Example 5.15 (A moving type and self twist)**.** Let $G = S_3$, let $H = A_3$, and let V be the standard two dimensional representation. Its restriction is

$$\operatorname{Res}_H V = \tau \oplus \tau^{-1}$$

for the two nontrivial characters of C_3. These characters form one moving orbit under $G/H \cong C_2$. Hence there is no stable extension spectrum, and

$$V \otimes \operatorname{sgn} \cong V.$$

Proof. The standard representation is induced from either nontrivial character of A_3. Quotient twisting is trivial on the inducing subgroup, so it preserves the induced representation. □

6 Strong Artin as global spectrum alignment

Let E/F be cyclic of prime degree p. Let ϵ be the corresponding idele class character. Suppose ρ is an irreducible finite image representation of G_F, its restriction τ to G_E is irreducible, and Π is an automorphic representation attached to τ. Let π_0 be one cuspidal descent of Π.

For each place v of F, put

$$R_v = \rho|_{W_{F_v}}, \qquad A_v = \operatorname{rec}_{F_v}(\pi_{0,v}).$$

Cyclic base change compatibility gives

$$\operatorname{Res}_{W_{E_w}} R_v \cong \operatorname{Res}_{W_{E_w}} A_v.$$

Definition 6.1 (Local attachment set)**.** When the local quotient has order p, define

$$\mathcal{J}_v = \{j \in \mathbb{Z}/p\mathbb{Z} : A_v \otimes \epsilon_v^j \cong R_v\}.$$

At a split place, set $\mathcal{J}_v = \mathbb{Z}/p\mathbb{Z}$.

Theorem 6.2 (Local computability of attachment)**.** *Every set $\mathcal{J}_v$ is computed by the extension spectrum criterion of theorem 5.5 applied to the finite quotient through which the two local parameters factor.*

Proof. The restrictions of A_v and R_v to W_{E_w} are isomorphic. Their common restriction has finite image because R_v does. The parameter A_v also has finite image by theorem 7.1 below, so both factor through one finite quotient after replacing it by a common finite quotient. Apply theorem 5.5. □

Theorem 6.3 (Global cyclic attachment criterion)**.** *With the preceding notation, the descent $\pi_0 \otimes \epsilon^j$ is attached to ρ if and only if*

$$\boxed{j \in \bigcap_v \mathcal{J}_v.}$$

Consequently, the cyclic Strong Artin step succeeds exactly when this intersection is nonempty.

Proof. Local Langlands is compatible with character twists. Hence

$$\mathrm{rec}(\pi_{0,v} \otimes \epsilon_v^j) \cong A_v \otimes \epsilon_v^j.$$

The global descent is attached exactly when this is isomorphic to R_v for every place v, which is the displayed condition. □

Corollary 6.4 (Determinant chosen candidate)**.** *If $p \nmid n$, determinant and central character select one candidate $j_{\det}$. The cyclic Strong Artin step is then equivalent to*

$$\boxed{j_{\det} \in \mathcal{J}_v \quad \text{for every } v.}$$

Proof. Use theorems 3.5 and 6.3. □

Remark 6.5. The local spectra are finite, but there are infinitely many places. A theorem forcing all local alignment sets to contain the same global index would prove the cyclic attachment step. The present reduction identifies that theorem precisely but does not prove it.

7 Every descent of an Artin base change is locally Artin

The automorphic descent has much stronger local structure than a general cusp form.

Theorem 7.1 (Local Artin descent)**.** *Let E/F be a finite extension, let π be an automorphic representation of $\mathrm{GL}_n(\mathbb{A}_F)$, and suppose its base change Π to E exists with local compatibility. Assume every local Weil Deligne parameter of Π has zero monodromy and finite Weil image. Then the same is true for every local parameter of π.*

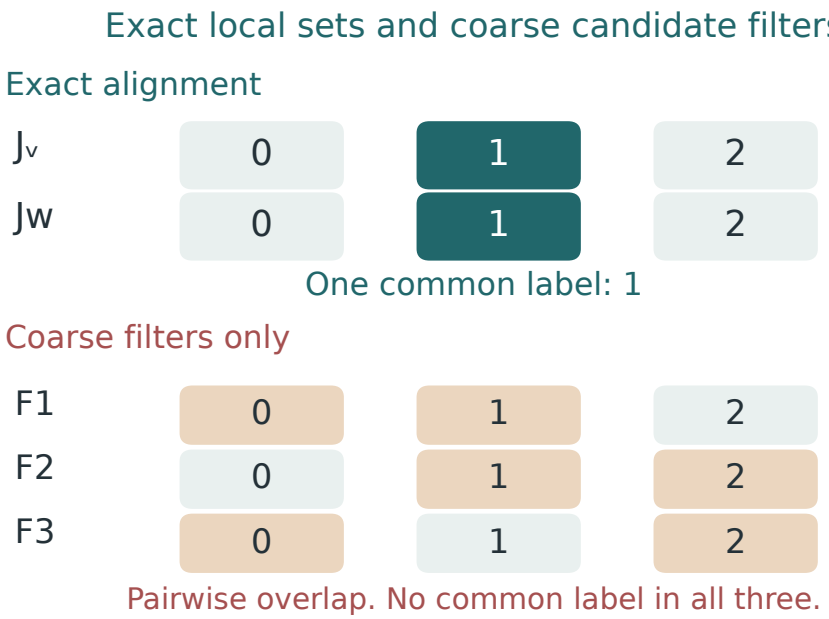


Figure 3: Exact alignment requires one common global label. The upper rows give exact local sets. The lower rows are coarse filters, which need a separate intersection test.

Proof. Fix v and $w \mid v$, and write the local Weil Deligne parameter as (ϕ_v, N_v). Base change restricts ϕ_v to W_{E_w} and keeps the monodromy operator. Since the parameter of Π_w has zero monodromy, $N_v = 0$.

The image of W_{E_w} under ϕ_v is finite. The subgroup W_{E_w} has finite index in W_{F_v}. Choose finitely many coset representatives. The full image $\phi_v(W_{F_v})$ is contained in the union of the corresponding finitely many cosets of the finite subgroup $\phi_v(W_{E_w})$. Hence it is finite. □

Corollary 7.2 (Tempered local descent)**.** *Under the hypotheses of theorem 7.1, every local parameter is bounded. Thus every local component of π is tempered in the normalized local Langlands correspondence.*

Proof. A finite subgroup of $\mathrm{GL}_n(\mathbb{C})$ is bounded. Local Langlands for general linear groups identifies bounded parameters with tempered representations [9]. □

Corollary 7.3 (Artin infinity type)**.** *If the base change is attached to a finite image global Galois representation, every archimedean local parameter of the descent has finite image.*

Proof. Apply theorem 7.1 at the archimedean Weil groups. The subgroup corresponding to the extension has finite index. □

8 Finite Satake descent

Let E/F be a finite Galois extension of degree d. Suppose π base changes to an automorphic representation Π attached to a finite image Artin representation τ. Let M be the exponent of the finite image of τ. Fix a finite set S containing the exceptional comparison places and all ramification of the extension and the representations. The unramified conclusions in this

section refer to places outside this fixed set. Weak attachment over the extension field suffices for these conclusions. If all local attachment over the extension field is known, the preceding local finiteness theorem supplies the stronger conclusion at every place.

Theorem 8.1 (Finite Satake roots)**.** *At every finite place v where E/F, π, and Π are unramified and the stated Artin polynomial comparison holds above v, every normalized Satake root of π_v belongs to*

$$\boxed{\mu_{dM}.}$$

Proof. Choose $w \mid v$ and let $f = f(w/v)$. If the normalized Satake roots of π_v are $\alpha_1, \dots, \alpha_n$, unramified base change replaces the Satake conjugacy class by its fth power. Hence the roots of Π_w are α_i^f. Since Π_w is attached to a representation of exponent dividing M,

$$(\alpha_i^f)^M = 1.$$

Thus $\alpha_i \in \mu_{fM}$. Since $f \mid d$, this group is contained in μ_{dM}. □

Corollary 8.2 (Integral Hecke polynomials)**.** *Every normalized Hecke polynomial of π at a place outside the fixed exceptional set S has coefficients in*

$$\boxed{\mathbb{Z}[\mu_{dM}].}$$

In particular, all of its coefficients are algebraic integers.

Proof. The coefficients are elementary symmetric polynomials in roots of unity from μ_{dM}. □

Corollary 8.3 (Finite Hecke field)**.** *The field generated by the normalized Hecke coefficients of π outside S is a number field contained in*

$$\boxed{\mathbb{Q}(\mu_{dM}).}$$

There are only finitely many possible normalized Hecke polynomials of degree n at these places.

Proof. Every coefficient belongs to the fixed cyclotomic field by theorem 8.2. Every root lies in the finite set μ_{dM}, so only finitely many multisets of n roots, and therefore only finitely many Hecke polynomials, can occur. □

Corollary 8.4 (Finite order central character)**.** *The central character of π has finite order.*

Proof. Outside the fixed exceptional set, every unramified determinant value lies in μ_{dM}. The Hecke character ω_π^{dM} is therefore trivial at every unramified place outside a finite set. The one dimensional strong multiplicity one statement, or equivalently uniqueness of a Hecke character from its almost everywhere unramified values, makes it the trivial global character. Thus ω_π has finite order. This argument does not require a separate local finiteness assertion at an exceptional place. □

Definition 8.5 (Finite Satake cusp form)**.** A cuspidal representation of $\mathrm{GL}_n(\mathbb{A}_F)$ is finite Satake when, outside one finite exceptional set, its normalized unramified Satake roots lie in one finite subgroup of $\mathbb{C}^\times$. The exceptional set is fixed, not allowed to change separately for each coefficient or prime.

Corollary 8.6 (Cyclic Artin descent is finite Satake)**.** *Every cuspidal cyclic descent of an automorphic Artin representation is finite Satake.*

Proof. Apply theorem 8.1. □

8.1 Cyclotomic rigidity

Proposition 8.7 (Finite number of unramified Hecke polynomials)**.** *Under the hypotheses of theorem 8.1, put $N = dM$. The number of possible normalized Hecke polynomials of degree n outside S is at most*

$$\boxed{\binom{N+n-1}{n}.}$$

Proof. A polynomial is determined by the multiset of its n Satake roots. Those roots lie in the N element set μ_N. The number of multisets of size n selected from N elements is the displayed binomial coefficient. □

Corollary 8.8 (Finite trace alphabet)**.** *The normalized Hecke traces of π outside S belong to a finite subset of $\mathbb{Z}[\mu_{dM}]$.*

Proof. Each trace is a sum of n roots from the fixed finite set μ_{dM}. □

Corollary 8.9 (Unramified Ramanujan property)**.** *Every normalized Satake root at the allowed unramified places has complex absolute value one.*

Proof. Every root is a root of unity by theorem 8.1. □

9 Finite polynomial alphabets and separating primes

Let K be a number field containing the normalized unramified Hecke coefficients of a finite Satake descent, the character values of the target Artin representation, and the values of the quotient character. Enlarge it to a splitting field in which the target and its restriction have matrix realizations. Let $\mathcal{O}_K$ be its ring of integers.

Definition 9.1 (Finite polynomial alphabet)**.** The finite polynomial alphabet of the cyclic attachment problem is the union

$$\mathcal{A} = \mathcal{A}_\pi \cup \mathcal{A}_\rho,$$

where $\mathcal{A}_\pi$ is the set of normalized Hecke polynomials occurring for π outside the fixed exceptional set, and $\mathcal{A}_\rho$ is the set of characteristic polynomials occurring for the finitely many twists

$$\rho \otimes \epsilon^j, \qquad j \in \mathbb{Z}/p\mathbb{Z}.$$

Proposition 9.2 (Finiteness of the alphabet)**.** *The set $\mathcal{A}$ is finite and is contained in $\mathcal{O}_K[T]$.*

Proof. The descent alphabet is finite by theorems 8.1 and 8.7. The Artin image and the quotient twist group are finite, so only finitely many characteristic polynomials occur on that side. Every root is a root of unity, so every coefficient is an algebraic integer. □

Definition 9.3 (Separating coefficient prime)**.** A finite place λ of K is separating for $\mathcal{A}$ when reduction modulo λ is injective on $\mathcal{A}$.

Theorem 9.4 (Existence of separating primes)**.** *All but finitely many finite places of K are separating for $\mathcal{A}$.*

Proof. For every ordered pair of distinct polynomials $P, Q \in \mathcal{A}$, choose one nonzero coefficient difference. Only finitely many finite places divide one of these finitely many nonzero algebraic integers. Outside their union, no two polynomials have the same reduction. □

Lemma 9.5 (Good reduction of an absolutely irreducible finite representation)**.** *Let τ be an absolutely irreducible representation of a finite group over a number field K. In a fixed matrix realization, exclude the finitely many coefficient places at which a group matrix has a denominator. The resulting stable local lattices have absolutely irreducible reduction at all but finitely many further places.*

Proof. Absolute irreducibility and Burnside's theorem give group elements $g_1, \dots, g_{n^2}$ whose matrices form a K basis of $M_n(K)$. Vectorize these matrices in the fixed basis. Their determinant is nonzero. Exclude places at which that determinant is zero or has a denominator, together with denominators of the finitely many group matrices. At a remaining place the coordinate lattice is stable, and the reduced matrices span the full matrix algebra over the residue field. They continue to span it after extending the residue field. Any invariant subspace is then invariant under every matrix, so it is zero or the full space. This proves absolute irreducibility of the reduction. The argument is local at the good places and does not require a globally free integral lattice over $\mathcal{O}_K$. □

Definition 9.6 (Good separating prime)**.** A finite place λ is good for the cyclic attachment problem when

1. it is separating for $\mathcal{A}$
2. its residue characteristic is different from p
3. the reduction of $\tau = \rho|_{G_E}$ is absolutely irreducible.

Corollary 9.7 (Abundance of good primes)**.** *All but finitely many finite places of K are good.*

Proof. Combine theorems 9.4 and 9.5 and exclude the finitely many places above p. □

10 Finite trace alphabets in residual characteristic

The next theorem gives an explicit bound for a finite residual image using only its trace alphabet.

Theorem 10.1 (Finite trace image bound)**.** *Let k be algebraically closed of characteristic zero or characteristic $\ell > n$. Let $\Gamma \subseteq \mathrm{GL}_n(k)$ be a subgroup, and assume its defining module k^n is semisimple. Suppose its trace set is finite, with*

$$t = \#\{\mathrm{Tr}(\gamma) : \gamma \in \Gamma\}.$$

Then Γ is finite and

$$|\Gamma| \le t^{\dim_k A} \le t^{n^2}, \qquad A = \mathrm{span}_k \Gamma.$$

In particular this bounds a finite semisimple residual matrix group in characteristic greater than its rank.

Proof. Because products of group matrices are group matrices, A is a finite dimensional subalgebra of $M_n(k)$. An A submodule is the same thing as a Γ submodule. The given module is semisimple and faithful for A. The Jacobson radical annihilates every simple constituent,

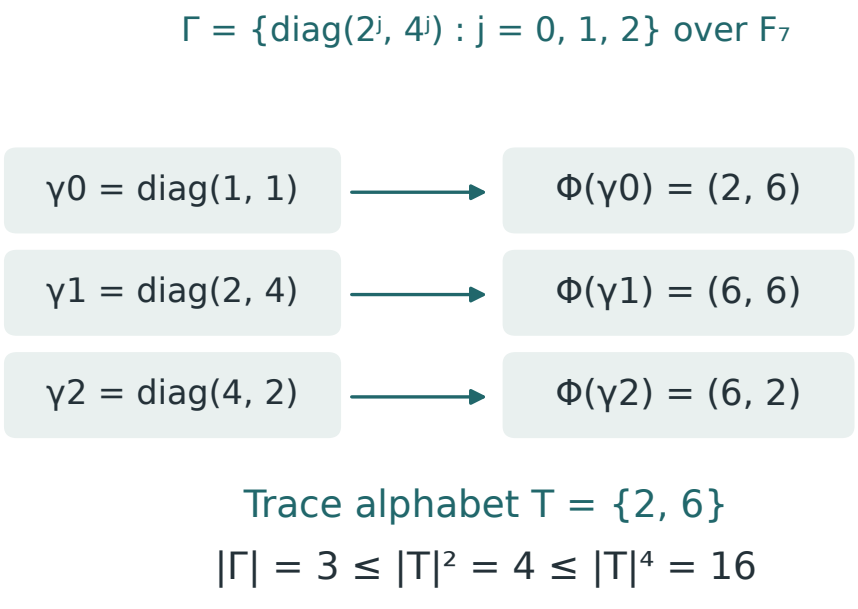


Figure 4: The nondegenerate trace pairing records each image element by finitely many traces of products with basis elements. This gives an injective encoding by the finite alphabet.

hence annihilates this faithful module. It is therefore zero. The split semisimple algebra has the form

$$A \cong \bigoplus_i M_{d_i}(k), \qquad k^n \cong \bigoplus_i (k^{d_i})^{\oplus m_i}, \qquad \sum_i d_i m_i = n.$$

Faithfulness gives $m_i \geq 1$ for every block. The trace pairing

$$B(x, y) = \mathrm{Tr}_{k^n}(xy)$$

restricts on the ith block to $m_i \,\mathrm{Tr}(XY)$. Matrix trace pairing is nondegenerate, as is seen by pairing matrix units E_{ab} with E_{ba}. Each m_i is nonzero in k: this is automatic in characteristic zero, and follows from $1 \leq m_i \leq n < \ell$ in residual characteristic. Thus B is nondegenerate on A.

Choose a basis $\gamma_1, \ldots, \gamma_r$ of A from the spanning set Γ, where $r = \dim_k A \leq n^2$. Send an image element to

$$\Phi(\gamma) = \big(\mathrm{Tr}(\gamma\gamma_1), \ldots, \mathrm{Tr}(\gamma\gamma_r)\big).$$

All coordinates lie in the trace set, because $\gamma\gamma_i \in \Gamma$. If two tuples agree, then $\gamma - \gamma'$ pairs to zero with every basis element, hence every element of A. Nondegeneracy gives $\gamma = \gamma'$. Therefore Φ injects Γ into a set of size t^r. This proves finiteness as well as the bound. Prior finiteness of Γ was not needed in this argument. □

Remark 10.2 (Historical position of the trace bound). The proof is elementary once the semisimple image algebra and its trace pairing are written down. No claim of historical priority is made. The theorem should be compared with modular Jordan type bounds and with the literature on finite trace sets in Galois representations.

Corollary 10.3 (Uniform bound for a residual Galois image)**.** *Let*

$$\bar{r} : G_F \longrightarrow \mathrm{GL}_n(\overline{\mathbb{F}}_\ell)$$

be continuous and semisimple, with $\ell > n$. If its Frobenius traces belong to a set of size at most t on a Chebotarev dense set, then

$$\boxed{|\mathrm{im}\,\bar{r}| \leq t^{n^2}.}$$

Proof. Continuity gives finite image. Chebotarev shows that every conjugacy class of the finite image occurs as an unramified Frobenius class. Apply theorem 10.1. □

Classical input 10.4 (Cross characteristic lifting). Let Γ be finite and let $\ell \nmid |\Gamma|$. Every semisimple representation of Γ over an algebraically closed field of characteristic ℓ is the reduction of a characteristic zero representation after a finite extension of coefficient fields. Irreducible representations lift to irreducible representations, and characteristic polynomials reduce compatibly. This is the nonmodular case of Brauer character theory [21, 22, 23].

Definition 10.5 (Universal bounded cyclotomic alphabet)**.** For positive integers n and B, let $\mathcal{Q}_{n,B}$ be the finite set of monic degree n polynomials whose roots are roots of unity of order at most B.

Theorem 10.6 (One large residual realization gives an Artin representation)**.** *Let K be a number field, let S be a finite set of places of F, and let $\mathcal{P} \subseteq \mathcal{O}_K[T]$ be a nonempty finite set of monic degree n polynomials. Assume all their roots are roots of unity. For $v \notin S$, choose $P_v \in \mathcal{P}$.*

There is a finite set $\Sigma(\mathcal{P})$ of rational primes, fixed from the alphabet before any realization is chosen, with the following property. Let λ be a coefficient place whose residue characteristic ℓ is outside $\Sigma(\mathcal{P})$. Suppose a continuous semisimple representation

$$\bar{r}_\lambda : G_F \longrightarrow \mathrm{GL}_n(\overline{\mathbb{F}}_\ell)$$

is unramified outside S and the places above ℓ, and has characteristic polynomial $\overline{P_v}$ at every remaining unramified place. Then there is a finite image semisimple complex representation r with polynomial P_v at those places. If $\bar{r}_\lambda$ is absolutely irreducible, then r is irreducible.

Proof. Let $t \geq 1$ be the number of trace values in the source alphabet and put $B = t^{n^2}$. First fix a finite normal extension $L/\mathbb{Q}$ containing K and the roots of unity of orders dividing

$$D_B = \mathrm{lcm}(1, 2, \ldots, B).$$

Every polynomial in $\mathcal{Q}_{n,B}$, the bounded cyclotomic alphabet, has coefficients in L. Form the union of all Galois conjugates in L of $\mathcal{P} \cup \mathcal{Q}_{n,B}$. This is a fixed finite set of integral polynomials. For each distinct pair choose one nonzero coefficient difference. Exclude the rational primes below the prime divisors of all these differences, and exclude the rational primes $\ell \leq \max(n, B)$. This defines a finite exceptional set independent of any residual representation or choice of lift. Outside it, reduction at every place of L above ℓ is injective on this enlarged polynomial set.

Continuity of $\bar{r}_\lambda$ into the discrete residual matrix group gives finite image Γ. Every conjugacy class of that image occurs at unramified Frobenius places away from the finite excluded set by Chebotarev. Its trace set has size at most t, because the Frobenius polynomials are reductions of the prescribed alphabet. Theorem 10.1 gives $|\Gamma| \leq B < \ell$, so $\ell \nmid |\Gamma|$.

Use the cross characteristic lifting statement in classical input 10.4, in a splitting modular system compatible with the chosen coefficient place. It gives an ordinary characteristic zero representation of Γ whose reduction is the prescribed residual representation. Compose it with $G_F \to \Gamma$. Each lifted eigenvalue has order dividing the order of an element of Γ, hence at most B. Every lifted Frobenius polynomial consequently belongs to $\mathcal{Q}_{n,B}$.

Only characteristic polynomials, not all matrix entries of the lift, are needed in the fixed separation field L. The lift can be realized over a further number field and reduced at a compatible place above λ. Its Frobenius polynomial and P_v have equal reduction there. On restriction to L, injectivity of the fixed enlarged alphabet makes them equal in characteristic zero. The chosen compatible embedding fixes the specified copy of K, so the equalities recover the original P_v simultaneously. The nonmodular correspondence preserves irreducibility. This proves the final assertion. □

Corollary 10.7 (Large residual realization of a finite Satake cusp form)**.** *Let π be a finite Satake cuspidal representation whose normalized Hecke polynomials lie in K. If one sufficiently large residual Galois realization with the expected unramified polynomials exists, then π has a finite image complex Galois realization. Absolute irreducibility of the residual realization gives an irreducible Artin representation.*

Proof. Apply theorem 10.6 to the finite Hecke polynomial alphabet of π. □

11 One prime residual attachment

Definition 11.1 (Residual realization of a descent)**.** Let λ be a finite place of the coefficient field. A residual realization of π at λ is a continuous semisimple representation

$$\bar{r}_{\pi,\lambda} : G_F \longrightarrow \mathrm{GL}_n(\bar{k}_\lambda)$$

whose characteristic polynomial at every unramified place outside a finite set is the reduction of the normalized Hecke polynomial of π.

Theorem 11.2 (One prime residual attachment)**.** *Use the cyclic Artin setup of theorem 2.3. For this unramified theorem it is enough that the lower cuspidal representation Π is weakly attached to τ, with π a cuspidal cyclic descent. Let λ be a good separating prime. Suppose π has a residual realization $\bar{r}_{\pi,\lambda}$. Then there is a unique index*

$$j_\lambda \in \mathbb{Z}/p\mathbb{Z}$$

such that

$$\boxed{P_{\pi,v}(T) = P_{\rho\otimes\epsilon^{j_\lambda},v}(T)}$$

for every finite place v outside a finite set. Consequently,

$$\pi \otimes \epsilon^{-j_\lambda}$$

is weakly attached to ρ.

Proof. Restrict the residual realization to G_E. At every unramified place w of E above a good place v of F, cyclic base change identifies the characteristic polynomial of the restriction with

the reduction of the Hecke polynomial of Π_w. Since Π is attached to τ, this is the characteristic polynomial of $\overline{\tau}(\mathrm{Frob}_w)$. Chebotarev density and Brauer Nesbitt give

$$(\overline{r}_{\pi,\lambda}|_{G_E})^{\mathrm{ss}} \cong \overline{\tau}.$$

The right side is absolutely irreducible and has the full dimension n. Hence the restriction itself has one composition factor and is irreducible, so it is isomorphic to $\overline{\tau}$.

The representations $\overline{r}_{\pi,\lambda}$ and $\overline{\rho}$ are therefore two extensions of the same absolutely irreducible representation of G_E. Their ratio is a character of G_F/G_E. Since the residue characteristic is different from p, the reduced quotient character has order p, and there is a unique index j_λ with

$$\overline{r}_{\pi,\lambda} \cong \overline{\rho} \otimes \overline{\epsilon}^{j_\lambda}.$$

At every good unramified place,

$$P_{\pi,v} \equiv P_{\rho\otimes\epsilon^{j_\lambda},v} \pmod{\lambda}.$$

Both polynomials belong to the finite alphabet $\mathcal{A}$. The prime is separating, so the congruence is equality in characteristic zero. Uniqueness follows from the freeness of the extension torsor. □

Corollary 11.3 (Only finitely many local checks remain)**.** *Under theorem 11.2, all local attachment to ρ is equivalent to checking the label $-j_\lambda$ at every exceptional finite place excluded from the unramified comparison and at every archimedean place.*

Proof. At the unramified places covered by the theorem, equality of the semisimple characteristic polynomials identifies the local parameters. At each remaining finite or archimedean place, use the corresponding local attachment condition. There are finitely many such places. When the stronger all local setup over the extension field holds, the finite image local parameters at a cyclic quotient are compared by the exact extension spectrum criterion, and the split restrictions already give equality. Under only lower weak attachment, the exceptional local comparison is a separate premise and finite local Artin type is not inferred there. These checks include the infinity parameters and any exceptional unramified place, rather than just the ramified places. □

Corollary 11.4 (Infinitely many residual realizations suffice)**.** *If residual realizations of π exist at infinitely many coefficient places, then one of them is good, and π is weakly attached after one quotient twist.*

Proof. Only finitely many coefficient places are not good by theorem 9.7. Apply theorem 11.2 at any good place carrying a residual realization. □

Corollary 11.5 (A congruent cohomological eigensystem suffices)**.** *Let λ be good. Suppose the reduction of the Hecke eigensystem of π occurs in a setting that supplies a semisimple residual Galois representation with the expected unramified characteristic polynomials. Then the conclusion of theorem 11.2 holds.*

Proof. The supplied representation is a residual realization in the sense of theorem 11.1. □

12 Representation theoretic collapse at the quotient prime

Let K be a number field containing the values of the order p quotient characters, let λ_p be a place above p, and let $\mathcal{O}$ be the corresponding valuation ring.

Theorem 12.1 (Semisimple collapse of cyclic extension labels)**.** *Let G be finite, let $H \triangleleft G$ with $G/H \cong C_p$, and let V, W be semisimple representations over a number field K with isomorphic restrictions to H. Choose stable lattices at a coefficient place above p. Over an algebraic closure of the residue field their semisimplified reductions satisfy*

$$\overline{V}^{\mathrm{ss}} \cong \overline{W}^{\mathrm{ss}}.$$

The same statement applies to finite image representations of an arbitrary group after passing to a common finite quotient that retains the cyclic quotient.

Proof. For $g \in G$, write its commuting primary decomposition as $g = g_p g_s$, where g_p has p power order and g_s has order prime to p. The image of g_s in C_p is trivial, so $g_s \in H$. In residue characteristic p, the matrix of g_s is diagonalizable after extending the field, and the matrix of g_p is unipotent. They commute, so the latter preserves each eigenspace of the former. On an eigenspace with eigenvalue α, the product is α times a unipotent operator. Its characteristic polynomial has only the root α, with that eigenspace's dimension as multiplicity. Therefore the reduced characteristic polynomials of g and g_s agree.

The restrictions of V, W agree on H, so their reduced characteristic polynomials agree at g_s, hence at every g. The classical modular Brauer Nesbitt theorem identifies the semisimplifications over the algebraic closure of the residue field [21, 16]. We use its full characteristic polynomial form, valid in arbitrary residual characteristic. Chenevier's checked preprint Theorem A, Lemma 1.12(ii), and Theorem 2.12 give the determinant and uniqueness interface. No restriction $p > n$ is needed for this uniqueness assertion. This proof does not assume that the original number field is a splitting field, or choose stable type extensions over it. For finite image representations of a larger group, quotient by the intersection of their kernels with H. The resulting finite quotient has the same argument. □

Corollary 12.2 (Residual collapse in the global and local fibres)**.** *At a coefficient prime above the cyclic degree p, all global Galois extensions of a fixed irreducible restriction have the same semisimplified reduction. At any one place, local parameters with the same restricted finite image parameter have the same semisimplified reduction after passing to a common finite local quotient and compatible coefficient field.*

Proof. Apply theorem 12.1 in the global finite quotient for the first assertion, and in the common local finite quotient for the second. The comparison is always between representations of the same group. In the cyclic descent setup the required local restriction equality is supplied by full local base change compatibility and attachment over the extension field. □

13 The quotient prime residual congruence

Let λ_p be a finite place of the coefficient field above the cyclic degree p, chosen after adjoining the values of the quotient character. Then

$$\epsilon \equiv 1 \pmod{\lambda_p}.$$

Definition 13.1 (Frobenius transform of a polynomial)**.** If

$$P(T) = \prod_{i=1}^{n}(T - \alpha_i),$$

define

$$P^{[p]}(T) = \prod_{i=1}^{n}(T - \alpha_i^p).$$

This is a polynomial expression in the coefficients of P.

Lemma 13.2 (Frobenius transform after reduction)**.** *Let P have integral coefficients at a place of residue characteristic p. Then*

$$\boxed{\overline{P^{[p]}}(T) = \mathrm{Frob}\big(\overline{P}(T)\big),}$$

where Frobenius raises every coefficient to its pth power. In particular, the map

$$\overline{P} \longmapsto \overline{P^{[p]}}$$

is injective over the residue field.

Proof. Over an algebraic closure of the residue field,

$$\overline{P^{[p]}}(T) = \prod_{i}(T - \overline{\alpha}_i^p).$$

The elementary symmetric functions of the pth powers are the pth powers of the elementary symmetric functions in characteristic p. Finite residue fields are perfect, so Frobenius is injective. □

Theorem 13.3 (Canonical quotient prime congruence)**.** *Use the cyclic Artin setup of theorem 2.3. For every finite place v outside a finite set,*

$$\boxed{P_{\pi,v}(T) \equiv P_{\rho,v}(T) \pmod{\lambda_p}.}$$

Thus the semisimplified reduction of ρ is residually automorphic with the mod λ_p Hecke eigensystem of every cuspidal descent of Π.

Proof. Assume first that v splits in E. Local base change is the identity at every place above v, so the two polynomials agree before reduction.

Assume that v is unramified and nonsplit. Since the extension has prime degree, its residue degree is p. Unramified base change raises the Satake conjugacy class of π_v to its pth power. Restriction of ρ_v to W_{E_w} does the same to its Frobenius class. Since Π_w is attached to $\rho|_{W_{E_w}}$,

$$P_{\pi,v}^{[p]}(T) = P_{\rho,v}^{[p]}(T).$$

Reduce modulo λ_p and use theorem 13.2. Injectivity of residue Frobenius gives the desired congruence. The finitely many ramified and excluded places do not affect residual attachment. □

Corollary 13.4 (All cyclic twists collapse residually)**.** *For every $j \in \mathbb{Z}/p\mathbb{Z}$,*

$$\overline{\rho \otimes \epsilon^j} \cong \overline{\rho}$$

at the semisimple character level modulo λ_p. Likewise, all automorphic descents have the same mod λ_p unramified Hecke eigensystem.

Proof. The values of ϵ are pth roots of unity and reduce to one at λ_p. Apply theorem 13.3 to every descent twist. □

Definition 13.5 (Quotient prime lifting gate)**.** The quotient prime lifting gate is the assertion that the matching residual unramified eigensystems supplied by theorem 13.3, together with the local deformation conditions of ρ, lift to characteristic zero automorphy of ρ.

Theorem 13.6 (Quotient prime reduction of a cyclic step)**.** *If the quotient prime lifting gate holds for the cyclic step, then the target representation ρ is automorphic. Hence Strong Artin for all finite solvable images follows if this gate holds at every invariant extension step of every solvable Clifford tower.*

Proof. The residual polynomial input is theorem 13.3. The first conclusion is the definition of the lifting gate. The second follows by the Clifford recursion in theorem 2.5. □

Remark 13.7. The quotient prime often divides the residual image order, and the resulting residual representation can be nonadequate or nonregular. The congruence theorem supplies matching residual unramified eigensystems. It does not supply the required weight raising, local component control, or nonregular automorphy lifting theorem.

14 Finite Satake rigidity for l adic realizations

Theorem 14.1 (Finite Frobenius trace rigidity)**.** *Let $L/\mathbb{Q}_\ell$ be finite and let $r : G_F \to \mathrm{GL}_n(L)$ be continuous, semisimple, and unramified outside a finite set. Suppose its traces at a density one set of unramified places belong to one finite set T, of size t. Then*

$$|r(G_F)| \leq t^{n^2}.$$

In particular r has finite image.

Proof. The image is compact. Choose a stable integral lattice after conjugation, so reductions modulo successive powers of a uniformizer give finite quotient images. Fix $g \in G_F$ and one such quotient. The Chebotarev class of the image of g has positive density. A density one set meets it, even after removing the finite set of ramified places. Thus some selected Frobenius matrix is conjugate to the image of $r(g)$ modulo that power, and their traces are congruent.

For each precision, $\operatorname{Tr} r(g)$ is therefore arbitrarily close to an element of T. The finite set T is closed in L, so $\operatorname{Tr} r(g) \in T$. Equivalently, one value in T occurs along an unbounded sequence of precisions, and its distance from $\operatorname{Tr} r(g)$ tends to zero. This uses conjugacy classes, which suffice because trace is conjugacy invariant. It does not require choosing a preferred Frobenius element inside each class.

Extend scalars to an algebraic closure of L. Semisimplicity is preserved in characteristic zero: the finite dimensional semisimple image algebra is separable. Apply the characteristic zero case of theorem 10.1 to $\Gamma = r(G_F)$. Its faithful defining representation is semisimple and its trace set has size at most t. This gives the stated bound directly, without an additional monodromy or torus argument. □

Corollary 14.2 (Cyclotomic Frobenius rigidity)**.** *Let r be as in theorem 14.1. If every Frobenius eigenvalue on a density one set belongs to one fixed finite group μ_N, then r has finite image.*

Proof. The Frobenius traces belong to the finite set of sums of n members of μ_N. Apply theorem 14.1. □

Theorem 14.3 (One semisimple l adic realization is enough)**.** *Let π be a finite Satake cuspidal representation with algebraic unramified Hecke polynomials. Fix an embedding of their number field into $\overline{\mathbb{Q}}_\ell$. Suppose there is a continuous semisimple representation $r : G_F \to \mathrm{GL}_n(L)$, for a finite extension $L/\mathbb{Q}_\ell$, unramified outside a finite set, whose Frobenius polynomials agree with the embedded Hecke polynomials at almost all places. Then r has finite image and determines a finite image complex representation weakly attached to π.*

Proof. Finite Satake gives a finite set of embedded Frobenius traces. Theorem 14.1 makes the image finite. A characteristic zero representation of a finite group has cyclotomic character values and can be realized over an algebraic number field after choosing a splitting field. Choose the ordinary complex realization in a modular system compatible with the fixed embedding of the Hecke field. Characteristic polynomials then give the same algebraic polynomials under the two embeddings, so the complex representation is weakly attached to π.

The assertion is about a globally semisimple representation. Frobenius semisimplicity alone is not substituted for that hypothesis. If an available construction supplies a representation not stated to be semisimple, its continuous semisimplification may be used when it has the same required unramified polynomials and finite ramification. The theorem then proves finite image of that semisimplification, not automatically of the original representation with extension data. □

Corollary 14.4 (Finite Satake forms in existing realization ranges)**.** *In every automorphic setting that supplies the continuous semisimple, finitely ramified realization required by theorem 14.3, finite Satake implies an Artin realization.*

Proof. Use theorem 14.3. □

15 Cohomological and weight raising realization gates

Classical input 15.1 (A torsion cohomology realization range). Scholze's Corollary 5.4.3 supplies a unique continuous semisimple residual representation for a Hecke eigensystem actually occurring in the stated cohomology of a locally symmetric space, under the field, level, coefficient, and theorem dependencies of that source [17]. The specified polynomial has coefficients $q_v^{i(n+1)/2} T_{i,v}$, rather than an unspecified normalization of the eigenvalues. A fully stated unconditional subrange uses a CM field containing an imaginary quadratic field and the stated compatible set of rational primes. The more general totally real or CM formulation in the original article records stabilization dependencies. A determinant over a coefficient ring modulo a bounded nilpotent ideal and an actual residual representation are distinct outputs. For the gate here we require the latter and an explicit identification with our reduced normalized Hecke polynomials.

Classical input 15.2 (A holomorphic limit realization range). Goldring's theorem, with an appendix by Shin, treats a unitary similitude group from a simple unitary Kottwitz datum and a cuspidal representation whose infinity component is an X holomorphic limit of discrete series [18]. At a coefficient prime satisfying the good reduction and reflex field conditions, it

supplies a semisimple l adic representation over the CM field. Its local compatibility is stated on the density one split prime set with the twist $|\cdot|^{(1-n)/2}$ in the source normalization. This is a concrete realization range. It is not a theorem for every singular infinity type or every finite Satake representation of GL_n.

Corollary 15.3 (A normalized torsion realization closes the weak step)**.** *Use the cyclic Artin setup and a good separating prime λ. Suppose the reduced eigensystem of a descent occurs in the cohomology range of classical input 15.1, with all source hypotheses met, and an explicit normalization comparison identifies the source Frobenius polynomials with $\overline{P_{\pi,v}}$. Then the cyclic step is weakly attached after one quotient twist.*

Proof. The declared occurrence, source theorem, and normalization comparison supply exactly the residual realization required by theorem 11.2. No occurrence or normalization is inferred merely from finite Satake data. □

Corollary 15.4 (A normalized holomorphic limit realization is finite image)**.** *Suppose the representation supplied in classical input 15.2 is continuous, semisimple, and finitely ramified as in that source. Suppose further that a specified comparison identifies its density one split prime Frobenius polynomials with one fixed cyclotomic alphabet from the target normalized finite Satake data. Then that Galois representation has finite image.*

Proof. The explicit comparison gives a finite trace alphabet on the density one set covered by the source's local compatibility. Apply theorem 14.1. The normalization and density condition are essential premises. This corollary does not assert all local attachment or an arbitrary transfer from GL_n into the unitary setting. □

Classical input 15.5 (Parallel weight one examples). Gee and Kassaei's companion construction is for a totally real field unramified at $p > 2$, with an irreducible modular residual representation satisfying the stated residual image and local distinct eigenvalue hypotheses. Their Artin consequence assumes the relevant modularity statement for every totally odd residual two dimensional representation and every coefficient prime [19]. Dimitrov and Wiese construct a continuous semisimple residual representation from a parallel weight one Hilbert eigenform of prime to p level. Its trace and determinant data require the stated Hecke and diamond character normalization [20]. These theorems illustrate realization and weight one methods in their precise ranks and hypotheses. They do not supply residual representations or weight raising in arbitrary rank.

Remark 15.6. These results demonstrate two ways to close the new gates. Torsion or coherent cohomology supplies a separating residual realization. Hasse invariant methods raise the quotient prime congruence to a setting where Galois representations and lifting theorems are available. Neither construction is presently known for every rank and every finite Satake descent.

15.1 A p adic family realization gate

Definition 15.7 (Pseudocharacter family gate)**.** A pseudocharacter family gate for a finite Satake descent π consists of a reduced p adic analytic space X, a closed point x valued in a finite extension of $\mathbb{Q}_\ell$, and a continuous dimension n characteristic zero pseudorepresentation over $\mathcal{O}(X)$. It factors through the Galois group unramified outside one finite set, and its unramified characteristic polynomials specialize at x to the embedded Hecke polynomials of

π. We use a point and coefficient setting in which characteristic zero realization produces a continuous semisimple representation after a finite coefficient extension. The existence of this family and the occurrence of the specified point are part of the gate.

Theorem 15.8 (A pseudocharacter family closes finite Satake realization)**.** *If a finite Satake descent has the pseudocharacter family gate of theorem 15.7, then it has a finite image semisimple Galois realization with the prescribed almost everywhere polynomials.*

Proof. Specialize at x and use characteristic zero pseudorepresentation theory in the declared coefficient setting [16]. This gives a continuous semisimple representation over a finite extension of $\mathbb{Q}_\ell$, factoring through the stated finitely ramified quotient. Its Frobenius polynomials are the required specialization. Apply theorem 14.3. Neither specialization nor rigidity constructs the analytic space, the Galois pseudorepresentation on it, or the point corresponding to the descent. □

Corollary 15.9 (Eigenvariety interpolation criterion)**.** *Suppose a finite Satake descent is a closed finite valued point on a p adic eigenvariety or completed Hecke space, and the point and Galois pseudorepresentation satisfy every coefficient, continuity, finite ramification, and specialization condition of theorem 15.7. Then the descent has a finite image Galois realization. Interpolation from classical points may be a way to construct the pseudorepresentation, but density by itself is not a substitute for these conditions.*

Proof. The interpolated pseudorepresentation is a pseudocharacter family gate. Apply theorem 15.8. □

Remark 15.10. The formal specialization theorem is complete. The unresolved geometric problem is to place every finite Satake descent on a space carrying such a pseudorepresentation. Weight one and holomorphic limit constructions prove this in important low weight settings, but no theorem presently covers every rank and every solvable cyclic descent.

16 A strengthened recursive Strong Artin theorem

Definition 16.1 (Realization gate at an extension step)**.** At an invariant cyclic extension step, a realization gate is any one of the following inputs for one cuspidal descent:

1. a residual realization at one good separating prime
2. a continuous semisimple finitely ramified l adic realization, including one obtained by a pseudocharacter family gate
3. a quotient prime weight raising and automorphy lifting theorem for the target Artin representation.

Theorem 16.2 (Realization gates imply solvable Strong Artin)**.** *Let ρ be an irreducible finite image representation with solvable image. Follow it through a prime cyclic Clifford tower. Suppose every invariant extension step admits a realization gate in the sense of theorem 16.1. Then ρ is weakly automorphic. If the remaining exceptional finite and archimedean local parameters align at every step, then the automorphic representation is attached at every place.*

Proof. Begin with the one dimensional bottom representation and use class field theory. At an induction step, the distinct conjugate subgroup constituents give distinct automorphic

conjugates by Chebotarev and strong multiplicity one. Prime cyclic automorphic induction therefore gives a cuspidal representation with the required almost everywhere parameters. At an invariant extension step, the lower cuspidal representation is Galois invariant because its conjugates have the same almost everywhere parameters as the invariant subgroup representation. Strong multiplicity one gives the invariance needed for cyclic descent. Choose a cuspidal descent. Lower weak attachment suffices for the finite Satake root bound outside a fixed finite set. The separate all local theorem retains its stronger local compatibility premise.

A separating residual realization gives the required characteristic zero quotient twist by theorem 11.2. A sufficiently large residual realization gives a finite image complex representation by theorem 10.6. An l adic realization has finite image by theorem 14.3. Its restriction is identified with the lower Artin representation by Chebotarev, so it lies in the extension torsor and one quotient twist selects the target. The quotient prime lifting gate gives the target directly by theorem 13.6. Iteration gives weak automorphy at the top. The exceptional finite and archimedean local checks upgrade weak attachment to all place attachment by theorem 11.3. □

Corollary 16.3 (A single residual theorem target)**.** *The unramified form of Strong Artin for every finite solvable image would follow from a theorem supplying, at every cyclic extension step, one residual realization of a descent at one good separating coefficient prime.*

Proof. Use the first realization gate in theorem 16.2. □

Corollary 16.4 (A single l adic theorem target)**.** *The unramified form of Strong Artin for every finite solvable image would follow from a theorem supplying one continuous semisimple finitely ramified l adic realization for every finite Satake cuspidal descent.*

Proof. Use the second realization gate in theorem 16.2. □

17 The weight zero realization bottleneck

The local finiteness theorem shows that the descended automorphic object already has the local shape of an Artin representation. The remaining problem is global realization.

Classical input 17.1 (Conditional construction of Kim and Yamauchi). Kim and Yamauchi's construction for $\mathrm{GL}_n(\mathbb{A}_{\mathbb{Q}})$ uses five declared inputs: a finite Hecke field, integrality with a common denominator, cuspidal Galois conjugates, residual Galois realizations for all sufficiently large coefficient primes, and a Rankin Selberg coefficient bound obtainable from the required exterior power functoriality. Their theorem is a conditional construction of Artin representations, not a theorem producing the residual realizations for every weight zero cusp form [4]. The finite alphabet and cross characteristic recovery below occupy closely related representation theoretic ground. Historical novelty is not inferred from their reformulation.

Theorem 17.2 (Automatic finite Satake inputs over the rational field)**.** *Let $E/\mathbb{Q}$ be finite Galois, and let π be a cuspidal representation of $\mathrm{GL}_n(\mathbb{A}_{\mathbb{Q}})$ whose base change to E is attached to a finite image Artin representation. Then the following inputs in the Kim and Yamauchi strategy are automatic:*

1. *normalized Hecke polynomials outside the fixed exceptional set are integral*
2. *the normalized Hecke field away from that set is finite.*

The remaining hypotheses concern the global Galois realization, its conjugates, and the required exterior power functoriality.

Proof. The first two statements are theorems 8.2 and 8.3. The description of the remaining inputs is the content of classical input 17.1. □

17.1 A complex pseudocharacter realization gate

Lemma 17.3 (Complex representations of profinite groups have finite image)**.** *Let G be profinite and let $r : G \to \mathrm{GL}_n(\mathbb{C})$ be continuous for the usual complex topology. Then $r(G)$ is finite.*

Proof. The Lie group $\mathrm{GL}_n(\mathbb{C})$ has a neighborhood U of the identity that contains no nontrivial subgroup. Continuity gives a neighborhood $r^{-1}(U)$ of the identity in G. A profinite group has a basis of open normal subgroups, so one such subgroup N lies inside $r^{-1}(U)$. Its image is a subgroup contained in U, hence is trivial. Therefore r factors through the finite quotient G/N. This proof uses a group neighborhood theorem. It does not use the false general topological assertion that an arbitrary continuous image of a totally disconnected space is totally disconnected. □

Classical input 17.4 (Characteristic zero pseudocharacters). Over an algebraically closed field of characteristic zero, a semisimple pseudocharacter of dimension n is the trace of a unique semisimple n dimensional representation up to isomorphism. We use this standard pseudocharacter realization principle in the form needed below [16].

Theorem 17.5 (Complex pseudocharacter gate)**.** *Let π be a finite Satake cuspidal representation of $\mathrm{GL}_n(\mathbb{A}_F)$. Suppose there exists a continuous dimension n complex pseudocharacter T of G_F, in the characteristic zero sense of classical input 17.4, with characteristic polynomials at almost all unramified Frobenius classes equal to those of π. Then π has a finite image semisimple Galois realization with those almost everywhere polynomials.*

Proof. The algebraic pseudocharacter realization theorem supplies a semisimple representation r with trace T. Its continuity deserves a separate argument. Let A be the span of its image. The trace pairing on A is nondegenerate, by the characteristic zero semisimple algebra argument in theorem 10.1. Choose a basis $r(g_1), \dots, r(g_d)$ from the image. The functions

$$g \longmapsto \mathrm{Tr}(r(g)r(g_i)) = T(gg_i)$$

are continuous. Nondegeneracy says that these finitely many trace functions determine all linear coordinates of $r(g)$ in A. Inverting their fixed coefficient matrix proves that every matrix entry of $r(g)$ is continuous. Now theorem 17.3 gives finite image. The polynomial equality is the hypothesis. This continuity step is not silently supplied by the algebraic realization statement alone. □

Corollary 17.6 (Pseudocharacter form of the realization bottleneck)**.** *For a finite Satake descent, constructing one continuous characteristic zero pseudocharacter with the prescribed Hecke polynomials is enough to produce the missing Artin representation.*

Proof. Apply theorem 17.5. □

Definition 17.7 (Finite Satake realization principle)**.** The finite Satake realization principle asserts that a cuspidal representation of $\mathrm{GL}_n(\mathbb{A}_F)$ with finite local Langlands image at every place and finite Satake data globally is attached to a finite image complex Galois representation.

Theorem 17.8 (Finite Satake realization implies solvable Strong Artin)**.** *Assume the finite Satake realization principle over every number field and in every rank. Then Strong Artin holds for every irreducible finite image Galois representation with solvable image.*

Proof. Proceed up a prime cyclic Clifford tower. At an induction step, use automorphic induction. At an invariant extension step, the already automorphic restriction gives a Galois invariant cuspidal representation over the extension field. Cyclic base change descent supplies a cuspidal descent over the smaller field. By theorems 7.1 and 8.6, that descent satisfies the finite Satake realization principle, so it has a finite image Galois representation.

Base change of the realized representation has the same unramified Frobenius polynomials as τ. Chebotarev density and semisimplicity identify its restriction to G_E with τ. Since τ is irreducible, the realized representation is irreducible and lies in the finite extension torsor. By theorem 2.4, twisting the automorphic descent by a quotient character selects the prescribed extension at that step. Iteration reaches the original representation. □

Corollary 17.9 (A single new theorem target)**.** *Strong Artin for all finite solvable images would follow from a global theorem realizing every finite Satake cuspidal representation as an Artin representation.*

Proof. This is theorem 17.8. □

Remark 17.10. The finite Satake realization principle is not proved here. It is a focused weight zero automorphic to Galois problem. The point of theorem 17.8 is that it replaces a separate attachment theorem at every solvable cyclic extension step by one global realization theorem.

18 A finite character separator on the Galois torsor

The Galois extension torsor itself always admits a degree one separator.

Theorem 18.1 (Trace separator for irreducible extensions)**.** *Let $H \triangleleft G$ with $G/H \cong C_p$, and let ρ be an irreducible extension of an irreducible G invariant representation of H. Then there exists $g \in G \setminus H$ such that*

$$\boxed{\operatorname{Tr}\rho(g) \neq 0.}$$

Consequently, the p quotient twists of ρ are distinguished by their traces at g.

Proof. The p extensions $\rho \otimes \xi^j$ are pairwise nonisomorphic by theorem 2.1. Their characters agree on H. If $\operatorname{Tr}\rho(g) = 0$ for every $g \notin H$, then all p characters would agree on every element of G, contradicting their pairwise nonisomorphism. Thus some $g \notin H$ has nonzero trace. At that element,

$$\operatorname{Tr}(\rho \otimes \xi^j)(g) = \xi(g)^j \operatorname{Tr}\rho(g),$$

and these values are distinct. □

Corollary 18.2 (Chebotarev trace separator)**.** *Let the finite image of ρ cut out a finite Galois extension. There is an unramified finite place v whose Frobenius element lies outside H and satisfies*

$$\operatorname{Tr} \rho(\mathrm{Frob}_v) \neq 0.$$

At that place, the Galois extension torsor is separated by the first Hecke coefficient.

Proof. Choose the conjugacy class of the element from theorem 18.1 and apply Chebotarev density [8]. □

Remark 18.3. The existence of a Galois trace separator does not show that the same prime is a trace separator for an arbitrary automorphic descent. Aligning these finite separators is another expression of the attachment problem.

19 Automorphic fibre separators

Theorem 19.1 (Existence of an automorphic coefficient separator)**.** *Let E/F be cyclic of prime degree p. Let π be cuspidal and assume its base change to E is cuspidal. Let ϵ be the quotient character. Then there is an unramified finite place v and an index k with $p \nmid k$ such that the degree k coefficient of the normalized Hecke polynomial of π_v is nonzero and $\epsilon_v(\mathrm{Frob}_v) \neq 1$.*

Proof. The representations π and $\pi \otimes \epsilon$ are distinct by the fibre theorem in classical input 2.2. Strong multiplicity one implies that their unramified local components differ at some finite place v. If $\epsilon_v(\mathrm{Frob}_v) = 1$, the two unramified components are identical, so the selected place has nontrivial quotient value.

Let A_v be the normalized Satake matrix. Twisting scales it by $\zeta = \epsilon_v(\mathrm{Frob}_v)$, a nontrivial pth root. Since the characteristic polynomials differ, theorem 3.3 gives an active coefficient degree k not divisible by p. □

Corollary 19.2 (Finite separator on both torsors)**.** *Both the Galois extension torsor and the automorphic descent torsor admit a finite unramified coefficient test that separates all quotient twists.*

Proof. Use theorems 18.2 and 19.1. □

Remark 19.3. The remaining issue is not finite separability of either torsor. Each torsor is already finitely separable. The issue is proving that the finite labels chosen on the two sides have the same offset.

20 Positive saturation and Artin holomorphy

Strong Artin is stronger than the Artin holomorphy conjecture. A separate positive method can prove holomorphy without constructing an automorphic representation.

Let K/F be finite Galois with group G. For a complex character θ of a subgroup, write $L(s, \theta)$ for the corresponding Artin L function.

Definition 20.1 (Entire packet set)**.** An entire packet set $\mathcal{E}$ is a collection of characters of subgroups of G whose Artin L functions are known to be entire.

Examples include nontrivial one dimensional characters and characters already known to satisfy Strong Artin with nontrivial cuspidal automorphic image.

Theorem 20.2 (Positive saturation criterion)**.** *Let χ be a nontrivial irreducible character of G. Suppose there are a positive integer m, nonnegative integers a_i, subgroups H_i, and characters $\theta_i \in \mathcal{E}$ such that*

$$\boxed{m\chi = \sum_i a_i \operatorname{Ind}_{H_i}^G \theta_i.}$$

Then $L(s,\chi)$ is entire.

Proof. Artin formalism gives

$$L(s,\chi)^m = \prod_i L(s,\theta_i)^{a_i}.$$

The right side is entire by hypothesis. Brauer induction implies that $L(s,\chi)$ is meromorphic [11]. If it had a pole at any point, its mth power would have a pole there. Thus it has no poles and is entire. □

Definition 20.3 (Entire induction cone)**.** Fix finitely many known entire induced packets $\Theta_1, \dots, \Theta_r$ in the character lattice of G. Define

$$C_{\mathcal{E}} = \left\{ \sum_{i=1}^{r} a_i \Theta_i : a_i \in \mathbb{Z}_{\geq 0} \right\}.$$

Its saturation is

$$\operatorname{Sat}(C_{\mathcal{E}}) = \{x : \text{some } m > 0 \text{ has } mx \in C_{\mathcal{E}}\}.$$

Corollary 20.4 (Finite cone certificate)**.** *If an irreducible character χ belongs to $\operatorname{Sat}(C_{\mathcal{E}})$, then its Artin L function is entire. Membership is an integer feasibility problem once a character table and the generator decompositions are known.*

Proof. Saturation gives precisely the relation in theorem 20.2. The coefficient equality is a finite system of linear Diophantine equations in the irreducible character basis. □

Remark 20.5 (Relation with monomial generalizations). Positive multiple and separation conditions on induced characters have a substantial literature. Generalized monomial groups and almost monomial groups were introduced precisely to study Artin holomorphy [12, 13]. The cone formulation packages the same analytic principle while allowing the generator set to include any independently known entire packet. Historical novelty should therefore be assessed against that literature rather than inferred from the formulation alone.

Remark 20.6 (Why signed Brauer induction is insufficient). Brauer induction permits negative coefficients. Those coefficients place Artin L functions in denominators and yield meromorphic continuation, not holomorphy. The positivity in theorem 20.2 is the essential analytic input.

21 Attempted completion routes and exact failure mechanisms

The preceding theorems arise from several distinct attempts to close the cyclic gate. Recording the failed routes is mathematically useful because each failure isolates information that a complete theorem must add.

21.1 Total automorphic induction

The automorphic induction of the attached representation over E is already attached to

$$\operatorname{Ind}_{G_E}^{G_F} \tau \cong \bigoplus_{j=0}^{p-1} \rho \otimes \epsilon^{j}.$$

It is also the isobaric sum of all descents. By theorem 4.7, this total object is invariant under rotating either torsor. It proves the direct sum correspondence but cannot identify one constituent with the target extension.

21.2 Central and stable local invariants

Determinant and central character select the candidate when $p \nmid n$. When $p \mid n$, theorem 4.2 shows that every algebraic central character is invariant. Deligne and Henniart stability gives the analogous blindness for sufficiently ramified character twists. The missing information is therefore carried by bounded depth noncentral extension spectra.

21.3 Functional equation comparison

Brauer induction gives meromorphic continuation and functional equations for the target Artin functions. Automorphic theory gives the corresponding statements for the descents. Base change identifies the product over all cyclic twists. The product identity is a norm identity and is invariant under a cyclic permutation of the factors. Without holomorphy, a converse theorem, or a primitive factorization theorem strong enough to isolate the factors, it does not determine the offset.

21.4 Cohomological regularization

Torsion cohomology, coherent cohomology, Hasse invariants, and eigenvarieties provide residual or l adic realizations in important ranges. A general finite Satake descent has singular weight zero infinity type and need not lie in any presently available cohomological realization theorem. The missing global step is therefore not local finiteness. It is placement of the weight zero eigensystem in a geometric or p adic family carrying a pseudorepresentation.

21.5 Repeated weight zero and the lifting hypotheses

If a continuous l adic realization has finite image, then at each place of F above the coefficient prime its restriction to the local Galois group becomes trivial after a finite extension of that local field. Its Hodge Tate weights are therefore all zero. There are n copies of the same weight. For $n > 1$ this is not regular. Regularity is not a stylistic label that can be removed from a theorem statement after finding a convenient residual congruence.

For example, the regular algebraic polarized lifting theorem of Barnet–Lamb, Gee, Geraghty, and Taylor requires regularity together with potential diagonalizability at coefficient places, irreducibility and image conditions, and suitable residual automorphy [26]. Finite image weight zero data do not satisfy its distinct weight hypothesis in arbitrary rank. A quotient prime congruence therefore cannot simply be inserted into that theorem.

Special nonregular theorems also exist. Calegari and Geraghty treat a nonregular ordinary symplectic setting with weights of the form $(0, 0, j-1, j-1)$, under their numerical, residual

modularity, image, and local conditions [27]. This shows that regularity is not required by every conceivable lifting method. It does not supply the all zero, arbitrary rank theorem required by the present cyclic gate. The correct conclusion is a mismatch between the declared target and currently invoked theorem hypotheses, rather than a universal prohibition on nonregular automorphy lifting.

The quotient prime adds another distinction. Because it can divide the target image order, its residual representation need not have the image properties assumed by a particular lifting theorem. A local algebra certificate can solve a specified matrix deformation problem while these global image, automorphic occurrence, and classicality requirements remain. Keeping those requirements explicit is what makes the proposed lifting route a testable mathematical target.

21.6 The surviving routes

After these tests, three routes remain genuinely viable:

1. construct one residual representation at a good separating prime
2. construct one semisimple l adic representation with the expected Frobenius polynomials
3. raise and lift the canonical quotient prime congruence while controlling the target local component.

For the first route, the good prime cyclic comparison is theorem 11.2. A sufficiently large residual prime additionally permits the intrinsic recovery theorem theorem 10.6. The second route uses theorem 14.3. The third is the natural nonregular automorphy lifting problem.

22 Three focused theorems that would close the solvable case

The broad finite Satake realization principle is no longer the only route. The present analysis isolates three independent sufficient theorems.

Definition 22.1 (Separating residual realization principle)**.** Every finite Satake cyclic descent has a residual Galois realization at one good separating coefficient prime.

Definition 22.2 (Finite Satake l adic realization principle)**.** Every finite Satake cuspidal representation has one continuous semisimple finitely ramified l adic realization with the expected unramified characteristic polynomials.

Definition 22.3 (Quotient prime lifting principle)**.** At every invariant cyclic Artin step, the quotient prime residual congruence of theorem 13.3 can be raised and lifted to the target characteristic zero Artin representation.

Theorem 22.4 (Any one principle closes solvable Strong Artin)**.** *Each of the principles in theorems 22.1 to 22.3, if valid over every number field and in every rank, implies Strong Artin for every finite solvable image representation in the almost everywhere sense.*

Proof. Apply theorem 16.2 with the corresponding gate at every invariant extension step. □

Remark 22.5. The separating residual principle is the weakest input at one fixed cyclic step because the target Artin representation is already known. The l adic principle is more intrinsic

Table 1: The Strong Artin reduction after the maximal cyclic analysis.

Block	Status	Remaining issue
Cyclic tower induction	established	applicable tower and cuspidality hypotheses
Solvable Clifford recursion	finite and exact	invariant extension steps
Local extension spectra	exact finite criterion	exceptional finite and infinity checks after weak attachment
Finite Satake alphabet	proved finite	none
Separating coefficient primes	all but finitely many	one residual realization
One prime residual attachment	proved	supply the residual representation
Quotient prime residual congruence	automatic	weight raising and nonregular lifting
l adic cyclotomic rigidity	proved	supply one l adic realization
Torsion cohomology route	available in declared cohomological settings	occurrence of the descent eigensystem
Holomorphic limit route	available in the declared unitary setting	transfer of the target descent into that setting
Positive Artin cone	exact sufficient criterion	finding positive certificates beyond known classes
General solvable Strong Artin	three conditional reductions	residual, l adic, or lifting realization gate

and realizes a finite Satake cusp form without a preselected target. The quotient prime principle is arithmetically canonical but confronts the hardest residual image and weight zero lifting issues.

23 What has actually been reduced

At a cyclic extension step the quotient prime congruence of the unramified polynomial systems is forced by base change. The target already has its residual Galois representation by reduction. The congruence supplies an automorphic residual eigensystem matching it. It does not itself construct a Galois representation for an arbitrary descent or prove a lifting theorem. Exact characteristic zero attachment requires either a second separating residual realization, an ordinary l adic realization, or a weight zero lifting theorem from the quotient prime congruence.

24 Verification

The accompanying verification programs test the finite claims used in the theorem chain, including the trace alphabet bound, cross characteristic toy lifts, quotient prime collapse, and the no go identities. They check cyclic torsor rotations, coefficient selectors, extension spectrum shifts, determinant selection, finite Satake root bounds, local alignment intersections,

and positive cone certificates on declared small examples. The release checks also cover labels, citations, proof counts, and source integrity.

Finite verification cannot prove the global automorphic theorems imported from the literature and does not establish the cyclic attachment conjecture. Its purpose is to detect algebraic and transcription errors in the finite reductions.

25 Exact boundary

The article does not prove Strong Artin for every solvable image. The full force audit ruled out several natural shortcuts and strengthened the residual and l adic realization gates. It proves a canonical unramified residual congruence at every cyclic step, a one prime weak attachment theorem under a residual realization input, and semisimple l adic finite image rigidity.

The following remain open in full generality:

1. producing one residual realization of every finite Satake descent at a separating coefficient prime
2. producing one continuous semisimple finitely ramified l adic realization of every finite Satake descent
3. raising the quotient prime congruence to a regular or otherwise liftable automorphic setting in arbitrary rank
4. proving the required nonregular automorphy lifting and classicality theorems
5. aligning all remaining exceptional finite and archimedean local parameters when the realization theorem supplies only almost everywhere compatibility
6. extending the argument from solvable finite images to arbitrary finite images
7. proving Artin holomorphy for arbitrary irreducible finite image representations.

Scholze's torsion cohomology theorem, Goldring and Shin's holomorphic limit theorem, and parallel weight one companion methods show that each realization gate is achievable in substantial arithmetic settings [17, 18, 19, 20]. The missing result is a theorem covering every finite Satake descent created by solvable Clifford recursion.

26 From finite realization to local deformation

The global gate is now sharply stated, but it has not disappeared. A separating coefficient prime converts a supplied residual Galois realization into an exact cyclic label and weak characteristic zero attachment. The quotient prime forces a different statement: equality of the unramified residual eigensystems, with all order p labels collapsed. A continuous semisimple l adic realization becomes finite image by the trace pairing theorem. Each route begins with its declared existence or lifting input.

The later formal work studies the local algebra surrounding that input. It distinguishes a full cyclic cohomology group from a convenient local line, a genuine second cohomology defect from a claim that the defect vanishes, and a certified local matrix inverse from the harder singular derivative sector. The next sections retain those distinctions and give detailed

proofs and worked controls. They are part of the same conjecture program because they tell us which local obligations have an exact certificate and which still require an arithmetic theorem.

These sufficient routes are not asserted to be mutually equivalent constructions. In particular, an abstract pseudocharacter does not automatically place the eigensystem in cohomology or on an eigenvariety. Nor does an exact Hensel root automatically come from an arithmetic local Galois representation. Such identifications require additional data and proofs. The final statement map returns to those remaining obligations after the local calculations.

27 Formal mechanisms and their exact hypotheses

The formal work is most useful when its hypotheses remain visible. A theorem that identifies an obstruction is different from a theorem that makes that obstruction vanish. A theorem that reconstructs a matrix from compatible residues is different from a theorem that attaches a Galois representation to an automorphic object. This section develops the algebra behind the checked interfaces, including examples where an attractive stronger statement fails. The purpose is to make the output of each mechanism clear enough to use in a future arithmetic argument.

The coefficient rings, groups, and local restrictions below are part of the input. Their connection to a particular number field or automorphic packet requires additional theorems. Nothing here constructs a global automorphic carrier. The matrix results do give actual integral matrices, rather than a formal symbol for a proposed limit. Their scope is therefore substantial, but local and explicit.

27.1 A cyclic cocycle is constrained by the norm

Let $C_e = \langle \gamma \rangle$ be a finite cyclic group of order e, acting linearly on a module M. Write T for the action of γ and set

$$N = 1 + T + \cdots + T^{e-1}.$$

The relation $T^e = 1$ implies $N(T-1) = (T-1)N = 0$. These identities explain why the image of $T-1$ lies in the kernel of N, but they do not identify the two spaces.

A one-cocycle is a function $a : C_e \to M$ satisfying

$$a(gh) = a(g) + g \cdot a(h).$$

Putting $g = h = 1$ gives $a(1) = 0$. If $x = a(\gamma)$, repeated application of the cocycle equation gives

$$a(\gamma^r) = \sum_{j=0}^{r-1} T^j x \qquad (1 \le r \le e).$$

The group relation $\gamma^e = 1$ therefore forces $Nx = 0$. Conversely, if $Nx = 0$, this formula defines a cocycle. Whenever exponents wrap around e, the discrepancy is a translate of Nx, which is zero. Thus evaluating at the chosen generator identifies cocycles with $\ker N$.

For a zero-cochain $y \in M$, its coboundary at γ is $(T-1)y$. Taking the quotient by these changes gives the precise presentation

$$H^1(C_e, M) \simeq \ker N / \operatorname{im}(T-1). \tag{27.1}$$

If the norm operator is identically zero, this reduces to $M/\operatorname{im}(T-1)$. That reduction has a hypothesis. Dropping the kernel from the general formula introduces classes that cannot satisfy the cyclic relation.

Suppose now that e is invertible in the coefficient ring. Define $S = \sum_{j=0}^{e-1} jT^j$. A direct coefficient calculation gives

$$(T-1)S = e - N.$$

For $x \in \ker N$, take $y = e^{-1}Sx$. Then $(T-1)y = x$, so every cocycle is a coboundary and $H^1(C_e, M) = 0$. The averaging argument uses division by e at its decisive step. In residue characteristic dividing e, the same display remains an identity but no longer constructs y.

Example 27.1 (The same generator action gives different cyclic cohomology)**.** Over $\mathbb{F}_3$, let $M = \mathrm{Mat}_2(\mathbb{F}_3)$ and let T be conjugation by $F = \operatorname{diag}(1, -1)$. The diagonal entries of a matrix form the $+1$ eigenspace of T. The off-diagonal entries form its -1 eigenspace. Both spaces have dimension two.

For C_2, the norm $1 + T$ is multiplication by 2 on the diagonal space and zero on the off-diagonal space. Hence its kernel is the off-diagonal space. The operator $T - 1$ is invertible on that space and zero on the diagonal space. Formula (27.1) gives $H^1(C_2, M) = 0$.

Let instead the generator of C_6 act by the same T. On the diagonal space, its norm is multiplication by $6 = 0$ in $\mathbb{F}_3$. On the off-diagonal space, the alternating six-term sum is also zero. Now $N = 0$ on all of M, while the image of $T - 1$ is still only the off-diagonal space. Therefore $H^1(C_6, M)$ has dimension two. The generator matrix alone does not determine this cohomology without the order and coefficient characteristic.

In Frobenius coordinates the relevant action is often $T(X) = FXF^{-1}$. The images of $T - 1$ and $1 - T$ coincide, since multiplication of a source vector by -1 exchanges the signs. This is the bridge to the concrete quotient by Frobenius coboundaries. The recovered v0.18 source proves that bridge under the explicit zero-norm hypothesis. It does not claim that every Frobenius quotient is automatically an arithmetic local cohomology group. One must first identify the acting group and the coefficient module.

27.2 A square-zero defect gives a genuine obstruction class

Let $q : R' \to R$ be a ring map with kernel I, and assume $I^2 = 0$. Let $\rho : G \to R^\times$ be a group representation. Choose unit-valued lifts $r_g \in (R')^\times$ with $q(r_g) = \rho(g)$ and $r_1 = 1$. Such a choice is an input to this calculation. It is not already a multiplicative lift. Define its normalized defect by

$$c(g, h) = (r_g r_h - r_{gh}) r_{gh}^{-1} \in I. \tag{27.2}$$

Equivalently, $r_g r_h = (1 + c(g, h)) r_{gh}$. The coefficients are elements of the actual ring kernel, with its additive group, rather than an unrelated vector space of the same dimension.

The action of G on I is $g \cdot x = r_g x r_g^{-1}$. It is independent of the chosen unit lift. Indeed, any other unit lift has the form $(1 + b)r_g$ with $b \in I$. Since $(1 + b)^{-1} = 1 - b$ and products of two elements of I vanish, conjugating $x \in I$ by $1 + b$ leaves x unchanged. The same observation shows that the defect in the multiplication of the r_g does not prevent these conjugations from forming a genuine group action on I.

Compute $(r_g r_h) r_\ell$ and $r_g (r_h r_\ell)$. After expressing both in terms of $r_{gh\ell}$, the first is

$$(1 + c(g, h) + c(gh, \ell)) r_{gh\ell},$$

while the second is

$$(1 + g \cdot c(h, \ell) + c(g, h\ell))r_{gh\ell}.$$

All omitted cross products belong to I^2. Associativity therefore yields

$$g \cdot c(h, \ell) - c(gh, \ell) + c(g, h\ell) - c(g, h) = 0. \tag{27.3}$$

This is the usual inhomogeneous two-cocycle identity. Normalization at the identity follows from $r_1 = 1$.

Change the chosen lifts to $r'_g = (1 + b(g))r_g$, where $b : G \to I$ and $b(1) = 0$. Expanding without commuting the factors gives

$$r'_g r'_h = (1 + b(g) + g \cdot b(h) + c(g, h))r_{gh}.$$

To compare with r'_{gh}, multiply by $1 - b(gh)$. The new defect is

$$c'(g, h) = c(g, h) + g \cdot b(h) - b(gh) + b(g). \tag{27.4}$$

Thus changing the section adds a two-coboundary. The class $[c] \in H^2(G, I)$ does not depend on the section.

Proposition 27.2 (Exact lifting criterion)**.** *With the hypotheses above, a unit-valued representation lifting ρ exists if and only if $[c] = 0$.*

Proof. If $c = \delta b$, replace r_g by $(1 - b(g))r_g$. Equation (27.4) shows that the corrected defect is zero. The corrected units are therefore multiplicative and reduce to ρ. The equality $c(1, 1) = 0$ forces $b(1) = 0$ for a primitive in the displayed normalization. Conversely, a multiplicative unit lift has zero defect. Its difference from the original section is a kernel-valued one-cochain, so the same change formula proves that the original class is zero. □

The cocycle identity proves compatibility of the defect. It does not prove that its cohomology class vanishes. This distinction is indispensable when an obstruction is used as the target of a resolution argument.

Example 27.3 (An obstruction that really does not vanish)**.** Reduce $\mathbb{Z}/27\mathbb{Z}$ to $\mathbb{Z}/9\mathbb{Z}$. Its kernel is $9\mathbb{Z}/27\mathbb{Z}$, whose square is zero. Map a generator of C_3 to the unit 4 modulo 9. This defines a representation because $4^3 = 1$ modulo 9.

A lift of the generator must be $4 + 9t$ modulo 27, for $t = 0, 1, 2$. Direct expansion gives

$$(4 + 9t)^3 \equiv 4^3 \equiv 10 \pmod{27}.$$

None has cube 1. Hence no representation lift exists and the obstruction class is nonzero. Conjugation is trivial in this scalar example. The kernel is additively $\mathbb{F}_3$, and the standard cyclic resolution gives $H^2(C_3, \mathbb{F}_3) = \mathbb{F}_3$ for the trivial action. This is a failure of lifting, rather than a failure to find a convenient section.

The defect can also be seen without an existence argument. Choose the section $r_1 = 1$, $r_\gamma = 4$, and $r_{\gamma^2} = 16$ modulo 27. These reduce to the required units modulo 9. The products of consecutive powers have no defect until the exponent passes through three. At that point $4^3 = 10 = 1 + 9$ modulo 27. Consequently the normalized defect is 9 when the exponents wrap around and zero otherwise. Under the identification $I \simeq \mathbb{F}_3$, this is the ordinary carry cocycle.

For a proposed correction b, sum the equations $c(\gamma^i, \gamma) = b(\gamma) - b(\gamma^{i+1}) + b(\gamma^i)$ for $i = 0, 1, 2$. The right side telescopes to $3b(\gamma) = 0$, using $b(1) = 0$ and the trivial action. The left side is $9 \neq 0$ in I. There is therefore no primitive. This explicit calculation recovers the nonzero obstruction class directly from its normalized representative.

For topological groups the cochains and correcting primitives must belong to the required continuous complex. An abstract primitive need not be continuous. The recovered continuous modules impose continuous sections and local restrictions, with a discrete coefficient ring. They preserve those hypotheses through their lifting criteria. Replacing continuous cohomology by abstract group cohomology would require a separate comparison.

27.3 Local conjugacy leaves a second gluing obstruction

Suppose an absolute two-cocycle has been cancelled by a continuous one-cochain b_0. Prescribed local lifts need not agree with the resulting global lift. Their remaining differences are local one-cocycles. Changing b_0 to another absolute primitive changes these differences by the restriction of a global one-cocycle. Conjugating the local lift at index i by a unit $1 + x_i$, with $x_i \in I$, changes the local cocycle by a local coboundary. The conjugators x_i are independently chosen at each local index. Requiring one common conjugator is a stronger problem.

After taking the local coboundary quotients, the remaining class belongs to

$$\operatorname{coker}\left(H^1_{\mathrm{cont}}(G, I) \longrightarrow \prod_i H^1_{\mathrm{cont}}(G_i, I)\right). \tag{27.5}$$

Here the local groups, their maps into G, and the precise continuous coefficient action must be supplied. The recovered source constructs the corresponding quotients directly in its square-zero data.

To see the criterion, let a_i be the local mismatch after b_0. A match up to the allowed local conjugations exists exactly when

$$a_i = \operatorname{res}_i(z) + \delta x_i$$

for a global continuous one-cocycle z and kernel elements x_i. This is exactly the assertion that the image of $(a_i)_i$ in (27.5) is zero. If b_1 is another absolute primitive, then $b_1 - b_0$ is a global cocycle. Its restriction disappears in the cokernel, so the remaining class does not depend on which primitive was used.

The v0.29 exact sequence identifies this secondary cokernel with the kernel of the map that forgets local gauge data from the relative obstruction space. Exactness at that position does not make the forgetting map surjective. It also does not make the localization map in (27.5) surjective. These are separate arithmetic questions. In particular, solving the absolute H^2 obstruction does not by itself solve a family of prescribed local conditions.

27.4 The reachable equation target is determined by syzygies

Let $J : V \to W$ be a linear map over a field k. The correction equation $Ju + e = 0$ has a solution precisely when $e \in \operatorname{im} J$. A useful description of this image uses all the dual relations among the equations. Let

$$S = \ker(J^* : W^* \to V^*), \qquad C = \{w \in W : \phi(w) = 0 \text{ for every } \phi \in S\}.$$

Every Ju belongs to C. Conversely, if w does not belong to $\operatorname{im} J$, its nonzero class in $W/\operatorname{im} J$ is detected by an algebraic linear functional. Pulling that functional back to W gives an element of S that does not annihilate w. Hence

$$\operatorname{im} J = C. \tag{27.6}$$

This algebraic dual argument works without a finite-dimensional hypothesis. It makes no assertion about continuous duals in a different topology.

Corestricting J to C makes it surjective onto its actual reachable target. The equations have not changed. By contrast, adjoining variables that absorb arbitrary errors changes the system unless their removal is proved. A solution of an augmented equation is not automatically a solution of the original equation. The original target and the quotient that records unreachable errors must be kept separate.

Example 27.4 (A reachable target with a visible failed correction)**.** Let $V = k^2$, $W = k^3$, and $J(u, v) = (u, v, u + v)$. Its image is the plane $w_3 = w_1 + w_2$. The dual kernel is generated by the functional $\phi(w_1, w_2, w_3) = w_3 - w_1 - w_2$. The error $(1, 2, 3)$ is compatible in every characteristic and is cancelled by $(u, v) = (-1, -2)$. The error $(0, 0, 1)$ is incompatible in every field, since its value under ϕ is 1. No choice of the two correction variables can cancel it.

If a third variable t is added so that the derivative becomes $J'(u, v, t) = (u, v, u + v + t)$, the enlarged map is surjective. It cancels the second error with $(u, v, t) = (0, 0, -1)$. But the nonzero value of t records exactly the failed original relation. This enlarged solution cannot be declared a solution of the original equations by forgetting t. One must show that the additional variable is removable, or that the target error already lies in the original compatible plane.

For a nonlinear system $F(x) = 0$, a relation among the residual Jacobian rows is only a first-order identity. Consider

$$F(x) = (x, x^2), \qquad J_x(h) = (h, 2xh).$$

At the origin the derivative has horizontal image. The functional selecting the second coordinate is a residual syzygy, but it does not annihilate $F(x)$ for general x. The exact relation is the moving relation

$$(-x, 1)F(x) = 0.$$

Furthermore,

$$F(x + h) - F(x) - J_x(h) = (0, h^2).$$

Modulo 8, taking $x = 0$ and $h = 2$ gives remainder $(0, 4)$, outside the horizontal target. Keeping only the syzygy at the origin would therefore lose a genuine nonlinear error even within the same residual fibre.

The graded correction argument retains an exact moving relation $B_x F(x) = 0$ and a residual exactness statement

$$\operatorname{im} \overline{J} = \ker \overline{B}. \tag{27.7}$$

If $F(x) = \pi^r e$, the moving identity gives $\pi^r B_x e = 0$. A faithful leading layer allows this equality to imply $\overline{B}\,\overline{e} = 0$. Equation (27.7) then supplies a residual correction $\overline{u}$ with $\overline{J}\,\overline{u} = -\overline{e}$. Lift $\overline{u}$ to u and put $x' = x + \pi^r u$. The linear part of $F(x')$ is divisible by π^{r+1}. For polynomial systems the remainder contains at least two copies of π^r, so it is divisible by π^{2r} and hence by π^{r+1} when $r \geq 1$. This produces the next layer.

Over a suitable truncated principal coefficient ring, iteration terminates at $\pi^N = 0$ and produces a root of the original system. The recovered v0.25 modules derive the necessary faithful-layer property under their stated principal residue and nilpotence hypotheses. They still require the moving identity and exactness. Equality of selected dimensions alone does not prove either one. A correct dimension count can help prove exactness after an inclusion is established, but cannot replace that inclusion.

The normalization by π^r deserves care. Writing $F(x) = \pi^r e$ need not determine e uniquely in a ring with torsion. If e' is another choice, then $\pi^r(e - e') = 0$. The faithful leading layer used by the theorem ensures that $e - e'$ has zero residue below the terminal cutoff. Thus the residue correction problem is independent of the chosen normalized error. The same hypothesis converts $\pi^r B_x e = 0$ into the residue relation needed for exactness. These are two concrete reasons why the coefficient-ring condition is part of the theorem. Formal cancellation of π^r in a ring where it is a zero divisor would be invalid.

The moving relation is evaluated at the current approximate root, rather than at the initial point alone. The residual class of that root is kept fixed during every correction, so its residual derivative and residual relation operator remain the same. The exact relation changes over the coefficient ring as the root changes. This is how the iteration retains the full nonlinear equations while continuing to solve one fixed residual linear problem. The distinction between a fixed residual operator and a fixed integral operator is the mechanism that permits redundant equations without an assumed integral right inverse.

There is also an integral warning. In $R = \mathbb{Z}/4\mathbb{Z}$, the submodule $C = 2R$ contains 2, and 2 is divisible by 2 in R. Yet 2 is not divisible by 2 inside C, since $2C = 0$. Thus an ambient divisibility statement need not be an internal divisibility statement in a compatible target. An integral splitting cannot be inferred from the residual vector-space picture. The graded argument avoids silently assuming such a splitting by working with the normalized residue of each actual error.

27.5 Ordered powers have a uniform ultrametric remainder

Let E be a complete normed ring with $\|1\| = 1$, a submultiplicative norm, and the ultrametric inequality

$$\|x + y\| \le \max(\|x\|, \|y\|).$$

Assume $\|A\| \le 1$. For a nonnegative integer d, define

$$D_{A,d}(h) = \sum_{i=0}^{d-1} A^i h A^{d-1-i}, \tag{27.8}$$

$$R_{A,d}(h) = (A + h)^d - A^d - D_{A,d}(h). \tag{27.9}$$

For $d = 0$ the sum is empty and both maps are zero. The order of the factors in (27.8) matters. Replacing this operator by $dA^{d-1}h$ requires A and h to commute. The replacement is invalid for general matrices.

The recursive identities are

$$D_{A,d+1}(h) = D_{A,d}(h)A + A^d h,$$
$$R_{A,d+1}(h) = R_{A,d}(h)(A + h) + D_{A,d}(h)h.$$

They follow by multiplying $(A + h)^d$ on the right by $A + h$ and collecting the terms with zero,

one, and at least two occurrences of h. This derivation does not rearrange any noncommuting factors.

For degree three, the linear term is

$$D_{A,3}(h) = A^2h + AhA + hA^2,$$

and the nonlinear remainder is

$$R_{A,3}(h) = Ah^2 + hAh + h^2A + h^3.$$

The middle word hAh cannot be absorbed into Ah^2 in a matrix ring. The remainder estimate keeps all four words and bounds each by its norm, so it respects the ordered expansion. This also explains why the estimate is not an assertion of commutative binomial calculus in disguise. The scalar bound remains uniform because the norm of a finite ultrametric sum is controlled by its largest summand, rather than by the number of summands.

Lemma 27.5 (Uniform remainder estimate)**.** *For $0 \le r \le 1$ and $\|h\|, \|k\| \le r$,*

$$\|D_{A,d}(h)\| \le \|h\|, \qquad \|R_{A,d}(h) - R_{A,d}(k)\| \le r\,\|h-k\|.$$

In particular $\|R_{A,d}(h)\| \le r\,\|h\|$.

Proof. The recursive formula for D shows its additivity. Each term of its ordered sum has norm at most $\|h\|$, because every power of A has norm at most one. A finite ultrametric sum has norm at most the largest term, proving the first estimate without a factor d.

Induct on d for the second estimate. The case $d = 0$ is immediate. Write $D_d = D_{A,d}$ and $R_d = R_{A,d}$. Subtracting the remainder recursions gives

$$\begin{aligned} R_{d+1}(h) - R_{d+1}(k) = {} & (R_d(h) - R_d(k))(A+h) \\ & + R_d(k)(h-k) \\ & + (D_d(h) - D_d(k))h + D_d(k)(h-k). \end{aligned}$$

The first term has norm at most $r\,\|h-k\|$ by induction and $\|A+h\| \le 1$. Taking $h = k$ and the other argument zero in the inductive estimate gives $\|R_d(k)\| \le r\,\|k\| \le r$, so the second term has the same bound. Additivity gives $\|D_d(h) - D_d(k)\| \le \|h-k\|$, which bounds the third term. Finally $\|D_d(k)\| \le \|k\| \le r$ bounds the fourth term. Ultrametricity bounds their sum by the same maximum. Since $R_{A,d}(0) = 0$, taking $k = 0$ gives the final claim. □

For $E = \mathrm{Mat}_n(\mathbb{Q}_p)$ use the entrywise supremum norm and normalize $|p|_p = p^{-1}$. Matrix multiplication satisfies

$$\left|(XY)_{ij}\right|_p = \left|\sum_t X_{it}Y_{tj}\right|_p \le \max_t |X_{it}|_p \left|Y_{tj}\right|_p \le \|X\|\,\|Y\|.$$

Taking the maximum over the finite set of entries proves submultiplicativity. The same ultrametric calculation removes a matrix dimension factor. Completeness is inherited from the finite product of copies of $\mathbb{Q}_p$. These are the hypotheses required by the actual matrix application, rather than properties borrowed from a complex matrix norm.

27.6 An exact inverse turns the local image into one finite test

Assume that $D = D_{A,d}$ has an additive two-sided inverse Q, and that $\|Qz\| \le K\|z\|$, with $K \ge 0$. Choose $0 \le r \le 1$ with $Kr < 1$. For a target U, set $\Delta = U - A^d$. The proposed correction map is

$$\Phi(h) = Q(\Delta - R_{A,d}(h)). \tag{27.10}$$

Both inverse identities will be used. $DQ = 1$ turns a fixed point into a root. $QD = 1$ proves the necessary condition from an existing root.

Theorem 27.6 (Exact image of a small power ball)**.** *Under these hypotheses, there exists a unique h with $\|h\| \le r$ and $(A+h)^d = U$ if and only if $\|Q\Delta\| \le r$.*

Proof. Suppose $\|Q\Delta\| \le r$. On the closed r-ball, the remainder estimate and ultrametricity give

$$\|\Phi(h)\| \le \max\big(\|Q\Delta\|, K\|R_{A,d}(h)\|\big) \le \max(r, Kr\|h\|) \le r.$$

For h, k in that ball,

$$\|\Phi(h) - \Phi(k)\| \le Kr\|h - k\|.$$

The ball is complete and $Kr < 1$. The contraction theorem gives a unique fixed point h. Apply D to its equation to obtain $Dh = \Delta - R_{A,d}(h)$, which is exactly $(A+h)^d = U$.

Conversely, a root in the ball satisfies $\Delta = Dh + R_{A,d}(h)$. Applying Q gives

$$Q\Delta = h + QR_{A,d}(h).$$

Its norm is at most r, since the second summand has norm at most $Kr\|h\| \le r$. Any such root is a fixed point of (27.10), so the contraction also proves uniqueness among all roots in this ball. □

The proof gives a controlled approximation process. Start at $h_0 = 0$ and iterate $h_{j+1} = \Phi(h_j)$. Then $h_1 = Q\Delta$, and

$$\|h_{j+1} - h_j\| \le (Kr)^j \|h_1\|.$$

For $t > j$, ultrametricity bounds $\|h_t - h_j\|$ by the largest of the intervening increments, which is at most $(Kr)^j\|h_1\|$. The limit therefore satisfies the same bound for its distance from h_j. With $r = p^{-m}$ and $K = p^s$, every iteration gains at least $m - s$ digits in this estimate. This is a quantitative consequence after the inverse and target test have been certified. It is not a search over unrelated roots at successively deeper levels.

Uniqueness is restricted to the chosen ball. Two different balls can each contain a root of the same target. A failed target congruence rules out the specified residue prefix, while an accepted congruence determines one root with that prefix. Covering all possible prefixes requires a separate argument, especially near centers where the derivative is singular or where the available precision is too shallow.

For integral matrices suppose the inverse has the certified form

$$Q = p^{-s}B, \qquad s \in \mathbb{N},$$

where B is an integral linear operator on the matrix entries. In a fixed ordering of the n^2 entries it is an n^2 by n^2 matrix over $\mathbb{Z}_p$. This is an operator on the matrix space, not ordinary left multiplication by an n by n matrix. Integral coefficients imply $\|Bz\| \le \|z\|$, so $K = p^s$ is a valid bound.

There is a finite exact certificate for the inverse. For every matrix unit E_{ij} check

$$QD(E_{ij}) = E_{ij}, \qquad DQ(E_{ij}) = E_{ij}.$$

Linearity and the matrix-unit basis extend both identities to every matrix over $\mathbb{Q}_p$. These equalities are exact in the coefficient field. They are not merely congruences modulo p. In particular, an operator singular over $\mathbb{Q}_p$ cannot satisfy them.

Take $r = p^{-m}$ with $m > s$. Then $Kr = p^{s-m} < 1$, and

$$\begin{aligned} \left\| Q(U - A^d) \right\| \le p^{-m} &\iff \left\| B(U - A^d) \right\| \le p^{-(m+s)} \\ &\iff B \operatorname{vec}(U - A^d) = 0 \pmod{p^{m+s}}. \end{aligned} \tag{27.11}$$

The last congruence is entrywise. It depends only on the reductions of the target, center, and integral operator modulo p^{m+s}. The theorem therefore decides the existence and uniqueness of a root in $A + p^m \operatorname{Mat}_n(\mathbb{Z}_p)$ by one finite target test, once the exact inverse certificate is supplied.

The root constructed over $\mathbb{Q}_p$ is integral. Its center is integral, and every correction entry has norm at most $p^{-m} \le 1$. If $d > 0$ and U is a unit, the root F is also a unit. Indeed, F commutes with $U = F^d$ and hence with U^{-1}, and $F^{d-1}U^{-1}$ is a two-sided inverse of F. Thus the unit version has an actual output in $\mathrm{GL}_n(\mathbb{Z}_p)$.

The value of s measures denominators of the whole inverse derivative. It is not automatically $v_p(d)$. The proof makes no statement about shallow balls with $m \le s$, where its contraction estimate is insufficient. It also does not provide automatic certificate discovery for arbitrary p-adic input. Exact rational examples can supply certificates by ordinary linear algebra, but a general certified implementation is an additional task.

27.7 A noncommuting unipotent square example

Work over $\mathbb{Z}_2$ and take

$$A = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}, \qquad d = 2.$$

For $X = \begin{pmatrix} a & b \\ c & e \end{pmatrix}$, the ordered derivative is

$$D_{A,2}(X) = AX + XA = \begin{pmatrix} 2a + c & a + 2b + e \\ 2c & c + 2e \end{pmatrix}.$$

Using the entry order (a, b, c, e), its matrix and a scaled inverse are

$$J = \begin{pmatrix} 2 & 0 & 1 & 0 \\ 1 & 2 & 0 & 1 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & 1 & 2 \end{pmatrix}, \qquad B = \begin{pmatrix} 2 & 0 & -1 & 0 \\ -1 & 2 & 1 & -1 \\ 0 & 0 & 2 & 0 \\ 0 & 0 & -1 & 2 \end{pmatrix}.$$

Direct multiplication gives $BJ = JB = 4I_4$. Therefore $Q = \frac{1}{4}B$ is the exact inverse and $s = 2$. The entry $\frac{1}{4}$ in this inverse prevents clearing its denominators with only one factor of 2.

The residual derivative is singular. In characteristic two, $D_{A,2}(I_2) = 2A = 0$, while $I_2 \neq 0$. Nevertheless J is invertible over $\mathbb{Q}_2$, with determinant 16. Residual singularity and singularity over the fraction field are different tests.

Choose $m = 3$, so the local root ball is $A + 8\,\mathrm{Mat}_2(\mathbb{Z}_2)$ and the finite target test is modulo $2^{m+s} = 32$. Consider

$$F = \begin{pmatrix} 1 & 1 \\ 8 & 1 \end{pmatrix}, \qquad U = F^2 = \begin{pmatrix} 9 & 2 \\ 16 & 9 \end{pmatrix}.$$

Here

$$AF = \begin{pmatrix} 9 & 2 \\ 8 & 1 \end{pmatrix}, \qquad FA = \begin{pmatrix} 1 & 2 \\ 8 & 9 \end{pmatrix},$$

so the root does not commute with the center. Since $A^2 = \begin{pmatrix} 1 & 2 \\ 0 & 1 \end{pmatrix}$, the target difference has entry vector $(8, 0, 16, 8)$. Applying B gives

$$B\,\mathrm{vec}(U - A^2) = (0, 0, 32, 0).$$

The congruence modulo 32 passes. The theorem supplies uniqueness in the specified ball, beyond the direct multiplication that exhibits this root.

For a negative control, take

$$U_{\mathrm{bad}} = \begin{pmatrix} 1 & 10 \\ 0 & 1 \end{pmatrix}.$$

The difference from A^2 is $(0, 8, 0, 0)$ in entry coordinates, and B sends it to $(0, 16, 0, 0)$. This is nonzero modulo 32, so no square root of U_{bad} lies in that ball. This is a failure for the specified prefix. It is not a claim that the target has no square root anywhere in a larger coefficient field.

The recovered exact checker compares the finite image with the finite criterion. Its unipotent square case tests 256 corrections at the stated horizon and 4096 corrections one level deeper. Both give the same 16 target residues as the criterion. These are bounded corroborating calculations. Infinite existence and uniqueness follow from the general contraction proof and its exact hypotheses.

There is a genuinely singular contrast. Let $A_0 = \mathrm{diag}(1, -1)$ over $\mathbb{Q}_p$. Then $D_{A_0,2}(E_{12}) = 0$, although $E_{12} \neq 0$. Thus this derivative has no inverse over the fraction field, including when p is odd and 2 is a unit. The certificate theorem cannot cover this center. A universal claim that invertible centers always have invertible power derivatives would be false.

27.8 Finite levels give a coherent branch only for a joint problem

Let X_r be finite sets with genuine reduction maps $\pi_{r,t} : X_t \to X_r$ for $r \leq t$. Assume identity and composition laws for these maps. Let $P_r \subseteq X_r$ be solution sets, and assume reduction preserves solutions. Then

$$(\forall r,\ P_r \neq \varnothing) \iff (\exists (x_r)_r,\ x_r \in P_r \text{ and } \pi_{r,t}(x_t) = x_r \text{ for } r \leq t)\,. \tag{27.12}$$

The nontrivial implication is finite branch compactness. For each depth, a solution at that depth gives a finite path of solutions at all earlier depths. There are arbitrarily long paths. The resulting tree has finite branching because every level is finite. An infinite branch follows from the finite-branch form of König's lemma. The reduction maps on the solution sets need not be surjective. Compactness selects persistent choices instead of extending a previously chosen node at every step.

This matters for simultaneous matrix prescriptions. At level r, a node can contain a shared matrix F_r together with one local conjugator $z_{i,r}$ for each member of a finite family. The equations are

$$z_{i,r}F_r^{d_i}z_{i,r}^{-1} = U_i \pmod{p^r},$$

with the prescribed prefix of F_r and prescribed residual identity for each conjugator. These are a single joint predicate on a finite tuple. The arithmetic equations and prefixes are preserved by reduction. Compactness therefore produces compatible towers for the shared matrix and for all local conjugators, provided this joint problem is nonempty at every depth.

An individual solution for each prescription is weaker. Even the scalar requirements $F = 1$ and $F = -1$ modulo an odd prime have solutions separately, but have no common solution. Likewise, solving every member of a family at each depth with a different shared matrix does not establish the premise of (27.12). The compatible branch must come from one family of jointly satisfiable nodes.

Nor does a successful shallow node necessarily extend. For example, the scalar prefix $F = 4$ modulo 9 satisfies $F^3 = 1$ there. Its lifts modulo 27 are $4, 13, 22$, all with cube 10. That branch stops at the next level. Merely checking a finite number of successful levels gives no all-depth conclusion without another theorem.

Conversely, for a downward-closed tower of finite solution sets indexed by all natural-number levels, absence of an infinite branch implies that some finite level is empty. This gives a finite obstruction witness in principle. It supplies no uniform computable bound on the depth of that witness. The matrix inverse theorem above provides an explicit horizon only in its certified nonsingular sector with $m > s$. It removes the all-depth premise for one local power problem in that sector. It does not remove the joint all-depth premise for arbitrary simultaneous prescriptions.

Finally, compatible integral matrix residues determine an actual matrix over $\mathbb{Z}_p$ entry by entry. Compatible inverse residues determine its inverse when the nodes are units. Equality of all reductions then implies equality of the integral matrices. These reconstruction steps turn the selected branch into a genuine integral solution. The remaining arithmetic attachment is not part of that reconstruction.

27.9 What the formal receipts establish

The recovered archives preserve the source statements and historical verification receipts. The v0.38 receipt records a successful build of its seven new modules against the pinned Mathlib environment. The earlier interfaces described here are drawn from their recovered source files and their recorded histories. Archive authentication verifies the recovered bytes. It is not a new compilation of the mathematical statements.

A fresh Lean replay was attempted during preparation of this edition, but the portable kit could not be recovered because its download returned HTTP 403. No fresh compiler invocation was therefore completed for this edition. The accompanying coverage report records that failure and distinguishes it from the historical successful receipt. The new exact finite controls also have their own tested domains. They do not substitute for a kernel replay or extend a conditional theorem to a missing arithmetic hypothesis.

The formal program supplies explicit cocycle quotients, exact obstruction criteria, careful nonlinear correction interfaces, compatible tower reconstruction, and an effective finite criterion for a substantial local matrix sector. To reach Strong Artin, these mechanisms still

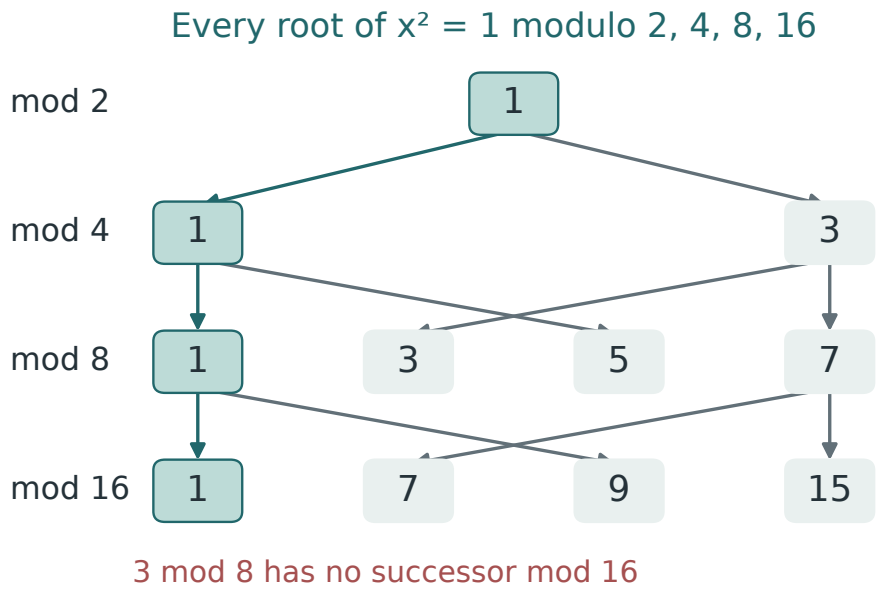


Figure 5: Scalar roots of one modulo successive powers of two form a finite prefix tree. A prefix can fail to continue even though other branches persist. An actual inverse limit needs a compatible branch at every depth. The separate matrix Hensel theorem supplies such a branch in its certified sector from one finite test.

need the arithmetic identifications and global existence arguments stated in the research boundary. Their mathematical value comes from making those remaining premises visible, while proving the local implications in full.

28 Worked representation and selector calculations

The preceding arguments use a small number of mechanisms repeatedly. This section expands those mechanisms in examples that can be checked without automorphic input. The examples are finite representation models. An incompatible local model is not an arithmetic counterexample to Strong Artin.

28.1 Why the prime cyclic Clifford alternative is exact

Proposition 28.1 (The two prime cyclic cases)**.** *Let G be finite, let $H \triangleleft G$ with $G/H \cong C_p$, and let V be irreducible over $\mathbb{C}$. Exactly one of the following occurs:*

1. *$V|_H$ is irreducible and invariant under G, and V is one of its p quotient character extensions.*
2. *$V|_H$ is a sum of p pairwise nonisomorphic conjugate irreducible representations, and V is induced from any one of them.*

Proof. Complete reducibility for the finite subgroup gives an isotypic decomposition of $V|_H$. The group G permutes these isotypic spaces by conjugating their types. More than one orbit would give a proper nonzero G stable direct sum, contrary to irreducibility. Thus there is one orbit. Since conjugation by H fixes isomorphism classes of H representations, the permutation

action factors through C_p. Its transitive orbit has size one or p.

In the size p case, choose one isotypic space U. Its translates are pairwise distinct isotypic spaces, with direct sum V. The stabilizer of its type is exactly H. If U had a proper nonzero H submodule, the direct sum of its p translates would be a proper nonzero G submodule of V. Hence U is irreducible. The natural map from $\mathrm{Ind}_H^G U$ to V, sending each coset copy to its translated space, is an isomorphism. Its restriction is precisely the stated orbit sum.

In the size one case, write $V|_H \cong \tau \otimes M$, where H acts trivially on the multiplicity space M. The stable type τ extends by theorem 5.2. Remove that extension action from the action on V. The remaining action on M factors through C_p. If M had a proper quotient stable subspace, its tensor product with τ would give a proper G submodule of V. Thus M is irreducible as a complex representation of C_p. Every such irreducible is one dimensional. Therefore $V|_H \cong \tau$, with multiplicity one, and the remaining quotient character is its extension label. The torsor theorem gives the p possible extensions. This proves both alternatives and excludes an intermediate multiplicity case. □

For a solvable finite group, a subnormal composition series has prime cyclic factors. Applied recursively to the constituents encountered along the series, this proposition separates induction steps from invariant extension steps. It does not say that every irreducible representation is induced all the way from a linear character. An invariant restriction can survive at positive rank. That surviving case is the attachment gate used in the main text.

For S_3 and $H = A_3$, the standard two dimensional representation lies in the moving case. A three cycle has the two nontrivial cube roots as eigenvalues. Restriction is the sum of the two nontrivial C_3 characters. Induction from either one gives the standard representation. Its sign twist is isomorphic to itself. This self twist is allowed precisely because the restricted representation is reducible. It does not contradict freeness of the torsor whose restriction is irreducible.

28.2 A rank three coefficient test that determinant misses

Let ω be a primitive cube root of unity and put

$$A = \mathrm{diag}(1, 1, \omega).$$

Expanding directly gives

$$\det(TI - A) = T^3 - (2 + \omega)T^2 + (1 + 2\omega)T - \omega.$$

The three elementary symmetric coefficients are

$$e_1 = 2 + \omega, \qquad e_2 = 1 + 2\omega, \qquad e_3 = \omega.$$

For scalar twist $\omega^j A$, they become

$$\omega^j(2 + \omega), \qquad \omega^{2j}(1 + 2\omega), \qquad \omega^{3j}\omega = \omega.$$

The determinant is constant for $j = 0, 1, 2$, but the first coefficient distinguishes the three labels because $2 + \omega \neq 0$. Thus rank divisibility causes blindness of the determinant, rather than blindness of every coefficient.

Now compare the balanced matrix

$$A_{\mathrm{bal}} = \mathrm{diag}(1, \omega, \omega^2).$$

Its polynomial is $T^3 - 1$. Multiplying all eigenvalues by ω merely permutes their multiset. Every lower coefficient vanishes, and the determinant is fixed. Its full characteristic polynomial is therefore blind to the label. This is an actual local rotational stabilizer. A claim that one characteristic polynomial always detects a quotient twist would fail on this example.

At quotient prime three and rank two, the determinant changes by ω^{2j}. If a required determinant ratio is ω^b, the unique candidate satisfies $2j = b$ modulo three, so $j = 2b$. This solves a finite selection problem. It still does not show that every other coefficient of the candidate matches the target.

The all local comparison includes infinity. For a real place with complex quadratic extension, the quotient is C_2. Let the target finite parameter have sign multiplicities $(3, 0)$, and let a second parameter have multiplicities $(1, 2)$. Both have determinant sign positive. Twisting the second by sign gives $(2, 1)$, which still does not equal $(3, 0)$. Its exact attachment set is empty. This is a finite local model showing why determinant selection, even with quotient prime coprime to rank, does not discharge the remaining infinity condition.

28.3 Ordered multiplicities and a common shift

Take $G = C_2 \times C_3$, let $H = C_2$, and let ξ be the quotient character taking a generator of C_3 to ω. Let τ be the nontrivial character of H, extended with trivial quotient action. Consider a semisimple representation with spectrum

$$P_{\tau,V}(z) = 2 + z.$$

Its quotient multiplicity vector is $(2, 1, 0)$. Twisting by ξ gives

$$zP_{\tau,V}(z) = 2z + z^2,$$

whose vector is $(0, 2, 1)$. Twisting by ξ^2 gives $(1, 0, 2)$. These are ordered vectors. Knowing that each has total multiplicity three does not retain the extension label.

Suppose a second stable H type σ occurs with spectrum one. If the target has spectra $2z + z^2$ for τ and z^2 for σ, the first block demands shift one and the second demands shift two. The two restrictions to H have the same multiplicities. Each block separately is a quotient twist of the matching source block. The full representation is not one quotient twist, because the required shifts disagree. This is the reason the spectrum criterion quantifies over one common j, rather than an independent j_τ for each type.

For a balanced multiplicity spectrum $1 + z + z^2$, both nonzero Fourier modes vanish:

$$1 + \omega + \omega^2 = 0, \qquad 1 + \omega^2 + \omega^4 = 0.$$

The spectrum 3 has the same total multiplicity, but both nonzero modes equal three. Comparing only ratios for which the source mode is nonzero would omit both decisive comparisons. The zero support clause in theorem 5.12 is therefore necessary. A zero source mode requires a zero target mode. It is not an instruction to skip that mode.

Changing the reference extension of τ relabels both spectra by the same rotation. The attachment set does not change. Likewise changing the chosen generator of the quotient rescales all labels by an invertible element of $\mathbb{Z}/p\mathbb{Z}$. These coordinate choices affect the displayed polynomials, but not the existence of an isomorphism after one common quotient twist.

For prime cyclic quotients a nonempty exact alignment set is either a singleton or all labels. Hence two nonempty exact singleton conditions $\{1\}$ and $\{2\}$ are incompatible, while $\{1\}$ together with all labels is compatible. Coarse coefficient filters may produce other finite subsets, but those are not automatically exact extension spectrum sets. Keeping that distinction prevents a visual or computational filter from being assigned a stronger meaning than its construction supports.

28.4 A trace encoding, with both hypothesis failures exposed

Let

$$\Gamma = \{\operatorname{diag}(\omega^a, \omega^{-a}) : a = 0, 1, 2\}.$$

The trace alphabet is $\{2, -1\}$. The image algebra has dimension two, with basis the identity and $\gamma_1 = \operatorname{diag}(\omega, \omega^2)$. The trace coordinate map from theorem 10.1 gives

$$\Phi(\gamma_0) = (2, -1), \qquad \Phi(\gamma_1) = (-1, -1), \qquad \Phi(\gamma_2) = (-1, 2).$$

The tuples are distinct. The sharper algebra bound is $3 \leq 2^2 = 4$, while the rank only bound is $3 \leq 2^{2^2} = 16$. The latter is deliberately rough. Its role is a uniform bound independent of the coefficient prime, not an efficient classification of the group.

Semisimplicity cannot be omitted. Over $\mathbb{C}$, the matrices

$$\begin{pmatrix} 1 & a \\ 0 & 1 \end{pmatrix}, \qquad a \in \mathbb{Z},$$

form an infinite group with the single trace value two. The invariant first coordinate line has no invariant complementary line. The trace pairing does not see the upper right nilpotent direction. This is exactly the step that fails without semisimplicity.

The characteristic restriction also has a visible purpose. In characteristic two, choose a primitive cube root ω and the scalar group

$$\{I_2, \omega I_2, \omega^2 I_2\}.$$

It is a faithful semisimple representation, as a direct sum of two identical one dimensional characters. Every trace is zero because its multiplicity two vanishes in the coefficient field. The alphabet has size one, but the group has size three. The would be bound $3 \leq 1^4$ is false. In the proof, the block pairing multiplier $m_i = 2$ is zero. Requiring $\ell > n$ prevents precisely this multiplicity cancellation.

In characteristic zero the finite trace theorem itself is enough to prove finiteness once global semisimplicity and a finite image trace alphabet are known. The l adic part of the work is therefore not another mysterious finite image principle. It is the arithmetic passage from density one Frobenius data to the trace set of the full compact image, followed by this explicit algebraic encoding.

28.5 A fixed alphabet and what a separating prime actually does

For an elementary example, compare the integral polynomials

$$P(T) = T^2 - 2T + 1, \qquad Q(T) = T^2 + 2T + 1.$$

Their middle coefficients differ by four. Reduction modulo two makes them equal. Reduction modulo any odd rational prime keeps them distinct. A separating prime is a coefficient prime with this latter injectivity property. It does not say which polynomial is realized by a global Galois representation. It says that, once a realization gives one of the alphabet polynomials modulo that prime, a different alphabet entry cannot imitate it.

For a cyclotomic alphabet over a number field, a nonzero algebraic integer difference has only finitely many prime ideal divisors. Choosing one nonzero coefficient difference for every pair therefore produces the finite bad set. Distinct roots can reduce to the same root at some bad prime, and repeated root multiplicities can become inseparable. The argument avoids those accidents by testing the polynomial alphabet, rather than presuming that each individual eigenvalue is already distinguished.

The large residual theorem uses two alphabets. One is the supplied automorphic alphabet. The other contains every degree n polynomial whose roots have orders at most $B = t^{n^2}$. A field containing $\mu_{\operatorname{lcm}(1,\dots,B)}$ contains all their coefficients. Both the field and the exceptional coefficient primes are fixed before choosing a residual realization. This order matters. A proof that chose a separation field from the coefficients of an as yet unknown lift and then called the excluded primes a predetermined set would leave a circular quantifier.

The resulting bound is not claimed to be practical. Even t and n of moderate size make B and the universal cyclotomic field enormous. The theorem establishes an existence implication with a finite exceptional set. A useful computational implementation for a particular finite image may exploit its actual exponent and much smaller image algebra, but that is an additional optimization and must retain the same separation logic.

29 Statement map and the remaining arithmetic obligation

The reader edition has three layers of evidence. The finite algebra proofs in the main text and local appendix are written out. Classical automorphic, local Langlands, and representation realization results are imported with their source hypotheses. The formal archive records exactly encoded algebraic theorems and historical build evidence. These layers support one another without being interchangeable.

Statement group	What the proof supplies	What remains an input
Cyclic extension torsor	All extensions of one invariant irreducible type differ by a unique quotient character	An automorphic attachment to select the matching extension
Coefficient and multiplicity selectors	Exact rotation, stabilizer, and common shift tests	Actual local data and compatibility of the selected global label
Finite Satake descent	A bounded root of unity alphabet outside a finite exceptional set	The known lower automorphic correspondence and base change compatibility
Trace encoding	Finite image and the bound under semisimplicity and characteristic hypotheses	A representation whose trace alphabet has the required form

Statement group	What the proof supplies	What remains an input
Separating residual recovery	Exact weak characteristic zero attachment from a supplied good residual realization	Existence of that residual realization and remaining local checks
Large residual recovery	A finite ordinary lift and exact alphabet recovery at sufficiently large good primes	The residual representation and classical cross characteristic lifting theorem
Quotient prime congruence	Equal reduced unramified polynomial systems with collapsed labels	A suitable weight zero lifting theorem with target local control
Semisimple l adic rigidity	Finite image from continuous finitely ramified realization and finite Frobenius traces	Existence and exact polynomial normalization of that realization
Cyclic cohomology and syzygies	Algebraic spaces, redundancy relations, and conditional local obstruction criteria	Arithmetic Selmer comparison and the required terminal obstruction vanishing
Certified matrix Hensel sector	Unique actual local root from one finite test when $m > s$ and the two sided inverse certificate holds	A certificate, singular sector analysis, and arithmetic or global realization
Solvable recursion	A complete implication from gates at every invariant step	Uniformly supplying those gates over the encountered fields and ranks

The historical paper contains additional theorem statements, corollaries, and small example verifiers. Its source and audit are retained as evidence, with the current statement map recording the narrowed claims used in this edition. The older global semisimplicity, Fourier zero support, coefficient field, and exceptional place issues have been corrected in the reader text. The later formal checkpoints are not retroactive certificates of every classical theorem or every prose argument in that historical paper.

The remaining global question can now be stated without a vague final step. At an invariant cyclic restriction, the automorphic descent fibre is known to exist. Can one prove that a descent has a continuous Galois realization with the exact normalized polynomial data, or lift the canonical quotient prime congruence to the prescribed target with its local conditions? If the answer supplies only almost everywhere polynomials, the exceptional finite and archimedean conditions must still be addressed for the all local version.

This is an existence and compatibility obligation. A finite trace alphabet controls a representation after it exists. A syzygy certificate removes redundant equations after the deformation comparison is supplied. A Hensel certificate solves a specified local matrix equation in its certified ball. None of those conclusions creates the global automorphic to Galois arrow by itself. Their value is that the resulting arrow, when constructed, will arrive at a tightly constrained and independently checkable destination.

References

[1] J. Arthur and L. Clozel, *Simple Algebras, Base Change, and the Advanced Theory of the Trace Formula*, Annals of Mathematics Studies 120, Princeton University Press, 1989. Prime cyclic base change and induction are used with their cuspidality hypotheses.

[2] L. Clozel and C. S. Rajan, Solvable base change, *Journal für die reine und angewandte Mathematik* 772 (2021), 147–174. DOI: 10.1515/crelle-2020-0023. Checked preprint arXiv:1806.02513v1, 7 June 2018. Theorem 3(2) and the ensuing cyclic descent remark are the fibre inputs.

[3] C. S. Rajan, On the image and fibres of solvable base change, *Mathematical Research Letters* 9 (2002), 499–508.

[4] H. H. Kim and T. Yamauchi, A uniform structure on subgroups of $GL_n(\mathbb{F}_q)$ and its application to a conditional construction of Artin representations of GL_n, *Journal of the Ramanujan Mathematical Society* 32(1) (2017), 75–99. Checked preprint arXiv:1403.5535v3, 7 April 2016. Theorem 1.1 is conditional on its five inputs. The inherited citation key records the preprint year.

[5] D. Ramakrishnan, Modularity of solvable Artin representations of $GO(4)$ type, *International Mathematics Research Notices* 2002(1), 1–54. DOI: 10.1155/S1073792802000016.

[6] R. P. Langlands, *Base Change for $GL(2)$*, Annals of Mathematics Studies 96, Princeton University Press, 1980.

[7] J. Tunnell, Artin's conjecture for representations of octahedral type, *Bulletin of the American Mathematical Society* 5(2) (1981), 173–175.

[8] J. Neukirch, *Algebraic Number Theory*, Springer, 1999.

[9] G. Henniart, Une preuve simple des conjectures de Langlands pour $GL(n)$ sur un corps p-adique, *Inventiones Mathematicae* 139(2) (2000), 439–455.

[10] M. Harris and R. Taylor, *The Geometry and Cohomology of Some Simple Shimura Varieties*, Annals of Mathematics Studies 151, Princeton University Press, 2001. Checked author preprint dated 30 August 1999, Lemma 12.7(4–6), for twisting, cyclic prime base change, and induction compatibility. Its numbering differs from the published book.

[11] R. Brauer, On Artin's L-series with general group characters, *Annals of Mathematics* 48 (1947), 502–514.

[12] J. König, Solvability of generalized monomial groups, *Journal of Group Theory* 13(2) (2010).

[13] M. Cimpoeaş and F. Nicolae, Artin L-functions of almost monomial Galois groups, *Forum Mathematicum* 32(4) (2020).

[14] K. S. Brown, *Cohomology of Groups*, Graduate Texts in Mathematics 87, Springer, 1982.

[15] I. M. Isaacs, *Character Theory of Finite Groups*, Academic Press, 1976.

[16] G. Chenevier, The p-adic analytic space of pseudocharacters of a profinite group and pseudorepresentations over arbitrary rings, in *Automorphic Forms and Galois Representations*, Vol. 1, London Mathematical Society Lecture Note Series 414, Cambridge University Press, 2014, 221–285. DOI: 10.1017/CBO9781107446335.008. Checked arXiv:0809.0415v2, 18 July 2013: Proposition 1.27, Theorem 2.12, Corollary 2.13, and Section 4.1. These are preprint numbering. The algebraic realization theorem alone does not assert continuity for the usual complex or l-adic topology.

[17] P. Scholze, On torsion in the cohomology of locally symmetric varieties, *Annals of Mathematics* 182(3) (2015), 945–1066. DOI: 10.4007/annals.2015.182.3.3. Corollary 5.4.3, p. 1060, is the residual eigensystem input. Field conditions, cohomological occurrence, normalization, and the source's historical dependencies are retained.

[18] W. Goldring, with an appendix by S. W. Shin, Galois representations associated to holomorphic limits of discrete series, *Compositio Mathematica* 150(2) (2014), 191–228. DOI: 10.1112/S0010437X13007355. Theorem 1.2.1 is used only with its unitary datum, coefficient prime, and split prime normalization conditions.

[19] T. Gee and P. Kassaei, Companion forms in parallel weight one, *Compositio Mathematica* 149(6) (2013), 903–913. DOI: 10.1112/S0010437X12000875. Checked arXiv:1206.6631v1, 28 June 2012. Its companion and Artin consequences retain their respective modularity hypotheses.

[20] M. Dimitrov and G. Wiese, Unramifiedness of Galois representations attached to Hilbert modular forms mod p of weight one, *Journal of the Institute of Mathematics of Jussieu* 19(2) (2020), 281–306. DOI: 10.1017/S1474748018000026. Checked preprint version 4, 28 November 2017. The full Frobenius polynomial comparison also requires the diamond character determinant normalization.

[21] R. Brauer and C. Nesbitt, On the modular characters of groups, *Annals of Mathematics* 42(2) (1941), 556–590.

[22] C. W. Curtis and I. Reiner, *Representation Theory of Finite Groups and Associative Algebras*, John Wiley and Sons, 1962.

[23] J.-P. Serre, *Linear Representations of Finite Groups*, Graduate Texts in Mathematics 42, Springer, 1977.

[24] P. Deligne, Les constantes des équations fonctionnelles des fonctions L, in *Modular Functions of One Variable II*, Lecture Notes in Mathematics 349, Springer, 1973, 501–597.

[25] P. Deligne and G. Henniart, Sur la variation, par torsion, des constantes locales d'équations fonctionnelles de fonctions L, *Inventiones Mathematicae* 64(1) (1981), 89–118. Theorem 4.6, p. 110, supplies the stable epsilon ratio used here.

[26] T. Barnet-Lamb, T. Gee, D. Geraghty, and R. Taylor, Potential automorphy and change of weight, *Annals of Mathematics* 179(2) (2014), 501–609. DOI: 10.4007/annals.2014.179.2.3. Checked arXiv:1010.2561v4, 9 December 2013. Theorem 4.2.1 uses regular algebraic and further lifting hypotheses.

[27] F. Calegari and D. Geraghty, Minimal modularity lifting for nonregular symplectic representations, *Duke Mathematical Journal* 169(5) (2020), 801–896. DOI: 10.1215/00127094-2019-0044. This is a special nonregular range, not the arbitrary all zero weight gate.

[28] P. Deligne and J.-P. Serre, Formes modulaires de poids 1, *Annales scientifiques de l'École Normale Supérieure* 7(4) (1974), 507–530. DOI: 10.24033/asens.1277. Lemma 3.2, p. 513, is a primary instance of the Frobenius determination argument.

[29] J. Wang, Strong Artin conjecture for generalized octahedral representations in GL_3, arXiv:2609.38231v1, 28 September 2026. Theorem 1.1 and Corollary 5.4 state the special three and six dimensional cases mentioned here. Statements checked. The full proof has not been independently audited for this reader edition.