hexagon:2611.01038v1

HECKE ACTIONS AND PURE INNER FORMS IN FAMILIES OF REAL LANGLANDS PARAMETERS

MINGHUA DOU AND JIAWEI YANG

Abstract. Let H be a connected quasi-split real reductive group. We construct, on the stack of flagged Langlands parameters, a compact object for each pure inner form of H. Its finite Hecke translates generate a summand of the periodic coherent category. These summands are mutually orthogonal and exhaust the category over the entire marked Cartan.

For each quasi-split form the corresponding summand is equivalent, as a finite Hecke module, to the periodic category of sheaves with nilpotent singular support on its enhanced real flag quotient. For a complex group regarded as real, the comparison is monoidal. For the full pure inner class we prove the equivalence on the open set where the roots taking value one span a space of rank at most one. It holds everywhere if each pure inner form is quasi-split or compact modulo its center. The comparisons preserve the marked parameter and the original boundary circle class.

For every pure inner form we also compute the maps from an arbitrary standard Hecke translate to its closed generator. The nonzero terms are lines on a finite cover of the norm torus. Their shifts and both Cartan actions agree with those on the real flag variety.

Contents

1. Introduction 1
2. The closed Borel orbit and its Hecke action 8
3. Ordinary coefficients and the quasi-split comparison 29
4. Pure inner labels, marked support and flag positions 60
5. Compact Levis and the closed real orbit 72
6. The decomposition by pure inner form 83
Appendix A. The complex-group comparison and compact-first Tate localization 98
Appendix B. Relative properness of the real parameter category 103
Appendix C. Ordinary coefficients by root system and finite verifications 107
References 127

1. Introduction

A flagged Langlands parameter for a real group specifies an odd element of the dual extension together with a Borel subgroup containing its square. The image of the square in the Cartan of this Borel is the marked boundary value. As the value varies, several real groups contribute to the same family of parameters. They are the pure inner forms of a quasi-split group, indexed by real torsors. Our first result identifies the category for the quasi-split group, including its finite Hecke action, with one summand of the spectral category. Our second result constructs a spectral summand for every pure inner label and proves that these summands are mutually orthogonal and generate the entire spectral category. In the exceptional types, the proof of the quasi-split comparison uses finite computations with Weyl groups. The programs and certificates used in the proof accompany this paper (Section C.6).

At a regular semisimple boundary value, the Borels containing a representative form a Weyl-group torsor. When a root takes value one, the variety of these Borels can have positive dimension. The table records how far

Date: November 25, 2026.

2020 *Mathematics Subject Classification.* Primary 22E45; Secondary 14F08, 17B10, 18N40.

Key words and phrases. Real reductive groups, Langlands parameters, finite Hecke categories, universal monodromy, coherent sheaves.

Minghua Dou, minghua.dou@outlook.com.

Jiawei Yang, yangjw@stu.xmu.edu.cn.

the comparison extends across these root divisors. Here the rank at a mark is the rank of the roots taking value one there.

Real forms and marked parameters	Conclusion
A quasi-split form, entire Cartan	Finite Hecke-module equivalence with its spectral summand (Theorem 1.1).
All pure labels, entire Cartan	Orthogonal decomposition and generation. Every standard return agrees on the two sides (Theorems 1.4 and 6.25).
All pure forms, marks of rank at most one	Equivalence of the full periodic categories (Theorem 6.30).
Every pure twist quasi-split or compact modulo its center	Equivalence of the full periodic categories over the entire Cartan (Theorem 1.4).
A complex group regarded as real	Monoidal equivalence over the entire Cartan (Theorem 1.3).
A pure form that is neither quasi-split nor compact modulo its center, at higher-rank root intersections	The full comparison in general remains open. It requires comparing compositions between translates of the closed generator.

1.1. **The quasi-split comparison over the marked Cartan.** Let H_0 be a connected quasi-split reductive group over $\mathbb{R}$, let $H = H_{0,\mathbb{C}}$, and choose a real Borel $B_H = T_H N_H$. Let G be the complex Langlands dual group, with pinned Borel B, Cartan T and pinned Galois involution θ. Write $R = \mathbb{C}[\pi_1(T_H)] = \mathcal{O}(T)$, let $Y = T/(1-\theta)T$, and put $A = \mathcal{O}(Y)$. The automorphic category is $\mathcal{A}(H_0) = \mathrm{Shv}_{\mathcal{N}}\big(H_0(\mathbb{R})\backslash H/N_H\big)$. Here the quotient uses the full group of real points, arbitrary stalk complexes are allowed, and $\mathcal{N}$ denotes the nilpotent singular-support condition. Right Cartan monodromy makes this category R-linear.

Let E^- be the odd component of the geometric dual extension and let $\widetilde{G} = \{(g, B') : g \in B'\}$ be the Grothendieck resolution. The square map $E^- \to G$ gives the derived stack

$$V = V_G = \mathcal{P}_E = \big[(E^- \times^R{}_G \widetilde{G})/G\big].$$

A pair (a, B') has $a^2 \in B'$ and marked boundary value $a^2 \bmod N_{B'} \in T$. Write u for the degree-two class of positive boundary rotation on the Möbius strip, and put $P = \mathbb{C}[u, u^{-1}]$. The spectral category is

$$\mathcal{T}(V) = \mathrm{Ind}\left(\mathrm{Coh}(V)^{S^1} \underset{\mathrm{Perf}(\mathbb{C}[u])}{\otimes} \mathrm{Perf}(P)\right).$$

We take circle invariants on bounded coherent objects and then invert u. On the automorphic side the subscript per denotes extension of scalars to P.

We use the actual odd square of the geometric dual extension for every isogeny type, as in (2.3). The geometric action in the original formulation is given under the adjoint-type convention [BCHN25, Section 4.2].

There is a distinguished object on each side for H_0. The unshifted extension of the universal local system on the compact enhancement of the closed real Borel orbit gives $P_C \in \mathcal{A}(H_0)$. On the parameter side let $i : F \to V$ be the locus where a itself preserves B', and set $Q = i_*\mathcal{O}_F$. Their closed coefficient ring is A. For its marking, choose the pinned odd representative τ with $\tau^2 = (2\rho^\vee)(-1)$, and put $r_\rho = (2\rho^\vee)(i)$. The point $t\tau$ has closed coordinate $[tr_\rho] \in Y$, and the boundary map in this coordinate is the norm $N([x]) = x\theta(x)$. This translation fixes the Cartan action on the coefficient ring, including along the root divisors.

Write $\mathcal{H}$ for the universal finite Hecke category of H. Its periodic monoidal comparison and the associated Steinberg stack are

$$\Phi_{\mathrm{fin}} : \mathcal{H}_{\mathrm{per}} \xrightarrow{\sim} \mathcal{T}(Z_G), \qquad Z_G = [B/B] \times^R{}_{[G/G]} [B/B].$$

The two flag projections give the input and output Cartan actions. Write $\mathcal{S}_Q$ for the localizing subcategory generated by the compact finite Hecke translates of Q.

Theorem 1.1. *For every connected quasi-split real reductive group H_0, there is an $R \otimes_{\mathbb{C}} \mathbb{C}[u, u^{-1}]$-linear equivalence*

$$E_{H_0} : \mathcal{A}(H_0)_{\mathrm{per}} \xrightarrow{\sim} \mathcal{S}_Q.$$

It sends P_C to Q, induces the prescribed identity on their closed coefficient algebra A, and intertwines the finite Hecke actions through Φ_{fin}. The action comparison is natural in the kernel and in the object and is compatible with the unit, convolution and all associativity coherences. It preserves the positive class of the original boundary circle.

Put $T_0 = T_H$ and $\Theta = H^1(\mathbb{R}, H_0)$. For a torsor class $\xi \in \Theta$ write H_ξ for its pure inner twist and put $\mathcal{A}_\xi = \mathrm{Shv}_{\mathcal{N}}(H_\xi(\mathbb{R})\backslash H/N_H)$. Inner identifications preserve the abstract Cartan. In these terms the two categories in the Families Soergel conjecture are

$$\bigoplus_{\xi\in\Theta} \mathcal{A}_\xi^{\mathrm{per}} \qquad \text{and} \qquad \mathcal{T}(V). \tag{1.1}$$

Let Θ_{qs} be the set of labels whose twist is quasi-split, put $f = \prod_{\alpha\in\Phi(G,T)}(1-e^\alpha)$ and write $U = D(f) \subset T$. Restriction to U means extension of scalars to $R[1/f]$. On this open the component group of the stabilizer of an odd element can have several characters, even though the connected centralizer of the boundary holonomy is a torus. A finite central cover of the odd component gives the associated character lines on V. Tensoring with these lines supplies the other quasi-split summands.

Theorem 1.2. *There is an $R \otimes P$-linear fully faithful finite Hecke-module functor*

$$J : \bigoplus_{\xi\in\Theta_{\mathrm{qs}}} \mathcal{A}_\xi^{\mathrm{per}} \longrightarrow \mathcal{T}(V)$$

over the whole marked Cartan. Its value on the distinguished summand is E_{H_0}, and its images for different labels are mutually orthogonal. On restriction to U it induces an equivalence between the two full categories in (1.1)*, preserving the boundary parameter u.*

There are consequently two different roles for the regular locus. For the distinguished real group, identification there gives the initial normalization of a functor that extends over every root divisor. For the full pure inner class it gives a region on which the quasi-split labels already exhaust the parameter category. A pure form of smaller real rank first appears on a proper closed subset of the Cartan. To find it one must examine flag positions over singular boundary values. We give its support and the comparison on the first root divisors below.

For a complex group viewed as real the Galois action exchanges the two factors of its complexification. The boundary data then give a Steinberg correspondence, and gluing along a flag gives convolution. This yields the following full monoidal comparison, proved in Appendix A.

Theorem 1.3. *Let L be a connected complex reductive group and $H_0 = \mathrm{Res}_{\mathbb{C}/\mathbb{R}} L$. There is a monoidal equivalence for convolution*

$$\mathcal{A}(H_0)_{\mathrm{per}} \xrightarrow{\sim} \mathcal{T}(\mathcal{P}_E)$$

over the whole Cartan in both variables. It preserves the two flag markings and the original positive circle class.

1.2. **Pure inner forms, standard kernels, and root divisors.** Fix a pure inner label ξ. Let $I = I(\xi)$ be the absolute type of a minimal real parabolic of H_ξ, and let P_ξ be the regular standard on its unique closed real flag orbit. Here regular refers to the universal local system on the enhancement of the orbit. The derived group of the minimal real Levi is compact. The corresponding spectral object in the Levi, followed by induction along the stable parabolic, gives an object $E_\xi \in \mathcal{T}(V)$. We use the Whittaker origin in the adjoint quotient of each minimal Levi to fix the labels. For the distinguished quasi-split group this construction gives Q. Let $\mathcal{S}_\xi$ be the localizing subcategory generated by the finite Hecke translates of E_ξ.

Theorem 1.4. *The finite Hecke translates of the E_ξ, over all pure labels, compactly generate $\mathcal{T}(V)$. The Hecke subcategories of distinct pure labels are mutually orthogonal over the entire marked Cartan. If every pure twist is quasi-split or compact modulo its center, the comparison in* (1.1) *is an equivalence over the entire Cartan.*

In particular $\mathcal{T}(V) \simeq \prod_{\xi\in\Theta} \mathcal{S}_\xi$, and the image of J is the product of the factors with $\xi \in \Theta_{\mathrm{qs}}$. Orthogonality holds for the full mapping complexes between distinct factors, also when their images in the Cartan intersect. For a pure form compact modulo its center the closed real flag orbit is the whole flag variety. Its factor is then its entire periodic automorphic category. The product of a complex group regarded as real with an arbitrary real torus also has a full comparison, by Corollary 4.11.

Within a single factor one can compute the effect of every standard change of flag on the closed generator. Write W for the Weyl group, W_I for the subgroup generated by I, w_I for its longest element, and $\delta_I = w_I\theta$. Before adjoining the period, the coefficient ring of the closed object is $A_I = \mathcal{O}(T/(1-\delta_I)T)$, with the norm map $[x] \mapsto x\delta_I(x)$ to T. For $w \in W$ let Δ_w be the standard kernel for relative position w. Theorem 6.25 and Proposition 5.10 identify the spectral complex $\mathrm{RHom}(\Delta_w E_\xi, E_\xi)$ and its real counterpart as bimodules. They vanish unless w commutes with δ_I. In the nonzero case there is a decomposition $w = dv$ with $v \in W_I$ and $d \in W^\theta$ preserving the positive roots of the Levi. The complex is a free rank-one $A_I \otimes P$-module in shift $[\ell(v)]$.

The input and output actions are related by the same signed Weyl transformation on both sides. Thus the part of w inside the minimal Levi gives the shift, and the other factor transports the marking.

The open set on which these results give a comparison of categories can be described in terms of the roots of the boundary element. Put $\Phi_t = \{\alpha \in \Phi(G,T) : e^\alpha(t) = 1\}$ and $U_{\le r} = \{t \in T : \operatorname{rank} \Phi_t \le r\}$. For the pure forms of a fixed minimal type I, Theorem 6.30 gives an equivalence over $U_{\le |I|}$ between the sum of their automorphic categories and the product of their spectral factors. Taking all types together gives the full equivalence in (1.1) on $U_{\le 1}$ for every connected quasi-split datum. In particular this includes each root divisor away from its intersections with independent root divisors. The support of an individual automorphic category is the union of the Weyl translates of the norm image of a maximally split Cartan of that form. Each component has dimension equal to the real rank of the form, by Lemma 4.5 and Proposition 4.6.

Two geometric constructions account for the closed objects and their labels. First write $\mathcal{C}_\xi$ for the subcategory generated under colimits by P_ξ in $\mathcal{A}_\xi^{\mathrm{per}}$. If Φ_I is the root system of the Levi, let $T[I]$ be the open where $e^\alpha(t) \neq 1$ for every $\alpha \notin \Phi_I$. Inside the Levi the roots may all take value one.

Theorem 1.5. *After choosing the open-orbit labels of the minimal Levis relative to the Whittaker origin of their adjoint quotients, there is a continuous, compact-preserving, fully faithful $R \otimes P$-linear functor*

$$\bigoplus_{\xi \in \Theta} \mathcal{C}_\xi \longrightarrow \mathcal{T}(V).$$

For a pure label of type I, its image over $T[I]$ is an orthogonal direct factor. In particular every pure form compact modulo its center contributes its entire periodic automorphic category as an orthogonal factor over the whole Cartan.

Second, the open real flag orbits have a simultaneous description for all pure inner forms. Let $\mathcal{A}_{\xi,!}^{\mathrm{open}}$ be generated under colimits by extension by zero from the open orbits in the enhanced flag variety. On the spectral side let V^{opp} be the open locus where B' and $aB'a^{-1}$ are opposite. Here the open condition is on a pair of flags. Its image under the boundary map can lie entirely over a singular Cartan value.

Theorem 1.6. *Choose an open real orbit for H_0 to fix the origin of the pure labels. For every connected quasi-split datum there is an $R \otimes P$-linear equivalence*

$$\bigoplus_{\xi \in \Theta} \mathcal{A}_{\xi,!}^{\mathrm{open}} \simeq \mathcal{T}(V^{\mathrm{opp}}).$$

The regular standard of each open orbit corresponds to the free module in its stabilizer character. Every pure inner form occurs. The right adjoint to open restriction embeds this category fully faithfully into $\mathcal{T}(V)$. If $w_0\theta = -1$ on $X^*(T)$, its marked support is the unit.

The pair of opposite Borels belongs to a single homogeneous space. Its real forms record a torsor together with an open orbit in that torsor. The stabilizer of the pair is a real torus, and its Galois cohomology is dual to the component group that labels the spectral lines. This identifies the coefficient of each open orbit, even for real forms with no support on U.

For an intermediate pure form the remaining global comparison concerns composition. On each side the standard Hecke translates of the closed object form a finite set of compact generators. To identify the two categories over the entire Cartan it suffices to identify the unital derived endomorphism algebras of their direct sums. The return formula identifies maps from a single standard translate to the closed object. The endomorphism algebra also contains maps between two translates and their compositions through a third one.

1.3. **From the real local Langlands correspondence to a family.** One starting point for the real local Langlands correspondence is Harish-Chandra's description of the discrete series. For a semisimple real group with a compact Cartan, his character formula determines the discrete series characters on the regular elements of that Cartan [Har66a]. Langlands's classification, circulated in 1973 and published in 1989, treats general irreducible admissible representations by parabolic induction. Starting with an irreducible tempered representation of a Levi subgroup, one twists it by a character whose real part lies in the positive chamber and forms the normalized induced representation. This standard representation has a unique irreducible quotient, and every irreducible admissible representation occurs in this way. On the parameter side, admissible homomorphisms from the real Weil group to the L-group organize the irreducibles into finite packets [Lan89, Bor79].

The real Weil group already explains why the nontrivial Galois element must enter a geometric parameter. It has two components, $\mathbb{C}^\times$ and $j\mathbb{C}^\times$, with $jzj^{-1} = \bar{z}$ and $j^2 = -1$. An admissible homomorphism must specify

how an odd element acts on the image of $\mathbb{C}^\times$, subject to this square relation. Passing to a conjugacy class of homomorphisms gives the parameter of a packet. To address a character question one must recover the individual representations in the packet, including the real form on which each is realized [Lan89, Bor79, ABV92].

The classification leaves a natural character problem. Given the inducing data, one knows which irreducible is the quotient. To express its character in terms of standard characters one must also determine the other composition factors of the standard representations and their multiplicities. In a block these multiplicities form a triangular matrix. Its inverse gives the coefficients expressing the irreducible characters in the basis of standard characters [Vog82].

The distinction between a packet and a block is already visible in these questions. Membership in a packet is specified by a Weil parameter. Composition factors and extensions relate representations in the same block and can involve several parameters. A geometric description of the block must therefore record both the individual parameters and how they meet. For real groups the flag variety supplied such a description on the representation side.

The flag variety gives geometric meaning to these multiplicities. Beilinson–Bernstein localization realizes Lie algebra modules with a suitable fixed infinitesimal character as twisted $\mathcal{D}$-modules on the complex flag variety [BB81]. For Harish-Chandra modules the equivariance is with respect to the complexification K of a maximal compact subgroup. A K-orbit and an equivariant local system on it determine an intersection cohomology complex on its closure. The restrictions of these complexes to smaller orbits are the subject of the calculations of Lusztig and Vogan [LV83]. Forgetting a simple step of a flag and pulling back from the partial flag variety relates the calculation on neighboring orbits. This is one origin of the finite Hecke action in the character problem. There is also a description on the orbits of the real group itself. Mirković, Uzawa, and Vilonen lift the Matsuki correspondence between the two sets of orbits to an equivalence of equivariant sheaf categories [MUV92]. We use the real group orbits below.

In this description one can see both the standard representations and their irreducible quotients. A local system on a single orbit first gives a sheaf on that orbit. Extension by zero and direct image give its standard and costandard extensions, with the usual dimension shifts. The intermediate extension gives the simple perverse sheaf. The difference between these extensions is supported on the boundary of the orbit. Thus the closure order on flags gives an induction for the character calculation, and the geometry of each boundary accounts for the entries beyond the diagonal in the multiplicity matrix [BB81, LV83].

A second use of flag geometry takes place on the dual side. Vogan's character duality relates the change of basis between standard and irreducible characters in a block to the corresponding calculation in a dual block [Vog82]. Adams, Barbasch, and Vogan construct varieties of geometric parameters for which a dual group orbit together with an equivariant local system labels a representation. Their description expresses character multiplicities through the local intersection cohomology of parameter orbit closures [ABV92]. Thus the geometry of a parameter comes with information about representations attached to neighboring orbits. One can then ask for a category of sheaves on parameters that also recovers extensions and their composition.

Koszul duality gives a model for such a statement. For the principal block of category $\mathcal{O}$, Beilinson, Ginzburg, and Soergel construct a grading in which the simple modules have linear projective resolutions. The algebra of their extensions recovers the Koszul dual block algebra [BGS96]. Soergel formulated a graded categorical version of Vogan's character duality for real groups [Soe01]. Bezrukavnikov and Vilonen proved the duality for blocks of quasi-split real semisimple groups with regular integral infinitesimal character [BV21]. Ionov and Yun gave a geometric proof using free-monodromic tilting sheaves on real group orbits. They construct a real Soergel functor by taking vanishing cycles at the closed orbit. Their comparison of completed Hom spaces is compatible with the Hecke action [IY25]. Both the closed orbit and the change of flag by a simple reflection will play a role here.

The role of the Hecke action in this development is worth making explicit. For a simple reflection, a full flag maps to a partial flag by forgetting one step. Pullback after pushforward varies the missing step through all its refinements. Composing these correspondences gives words in the simple reflections. The resulting functors come with maps between words, and one can apply such a map to an object on a real flag orbit. An equivalence of Hecke modules must carry this map, as well as each individual functor, to the corresponding map on the parameter side. This is the compatibility with convolution in the quasi-split comparison.

To pass from a block to a family, one must specify which parameter varies. In the Betti description the parameter is multiplicative. The map from the enhanced flag variety to the flag variety is a Cartan torus bundle. Given a character χ of the fundamental group of a torus fiber, one can impose monodromy χ in that fiber. One can instead take the universal local system, whose fiber is the group ring of the fundamental group and on which each loop acts by multiplication. Tensoring with the one-dimensional module of character χ then gives the local

system with that monodromy. The spectrum of the group ring is the dual Cartan torus. Completing the ring at its augmentation ideal retains the formal neighborhood of identity monodromy used in free-monodromic tilting theory. Working over the group ring itself allows us to follow a change of flag from the regular locus to a root divisor in the same family.

There is a second reason to keep the character of the torus fiber as a variable. A flag correspondence has two projections. The monodromy on the incoming flag acts on one side of a kernel and the monodromy on the outgoing flag acts on the other. On the regular locus a Weyl element transports one character to the other. When several flags lie over the same boundary value, one must still remember which projection supplies each action in order to compose two changes of flag. The abstract Cartan, before taking its Weyl quotient, supplies both markings.

There is a geometric reason for the same multiplicative parameter to occur on the dual side. In the loop description of Ben-Zvi and Nadler, real parameters arise from twisted local systems [BN13]. Picture a Möbius strip with a Borel reduction on its boundary. Going once around its core reverses orientation, so its holonomy is an element a of the odd component of the dual extension. Going once around the boundary goes twice around the core. The condition on a boundary Borel B' is therefore $a^2 \in B'$, and the image of a^2 in the Cartan quotient of B' is the boundary monodromy. Rotating this boundary gives a circle action on the parameter stack.

A Borel reduction adds an important choice to the boundary holonomy. An unmarked semisimple conjugacy class specifies a point of the Weyl quotient of the dual Cartan. Choosing a boundary flag lifts that point to the Cartan itself. For a regular semisimple holonomy, changing the lift amounts to choosing another point of a finite Weyl orbit. For a singular holonomy several boundary flags can give the same lift. The parameter stack retains the space of those flags, and a Hecke correspondence can act on the sheaves it supports. This is where the geometry of flag varieties reappears in the multiplicative family.

Ben-Zvi, Chen, Helm, and Nadler relate this stack to the geometry of Adams–Barbasch–Vogan. At a specified semisimple Cartan value their comparison identifies periodic cyclic sheaves on the completion with the periodized sheaf category on the corresponding spaces of geometric parameters [BCHN25, Theorem 4.4]. Their Families Soergel conjecture asks for the category over the multiplicative family, with the representation side summed over a pure inner class [BCHN25, Conjecture 4.5]. The indexing is significant. Two torsors can give isomorphic real groups, and each still contributes its own category. Ben-Zvi and Nadler also relate the conjecture to periodic geometric Langlands on the twistor line, whose real involution exchanges the reductions at its two marked points [BN26, Section 4.4].

To compare operations in these categories one needs a common description of their kernels. Bezrukavnikov's affine Hecke equivalence identifies constructible convolution with coherent convolution on a Steinberg space [Bez16]. Taylor develops tilting sheaves with universal monodromy for groups with connected center [Tay25]. Dhillon and Taylor establish Whittaker and affine Hecke equivalences with universal monodromy [DT26b, DT26a]. On the coherent side the Steinberg space records a group element with two Borel reductions. Composition forgets the middle Borel and retains the two outer markings. Appendix A uses the boundary circle to pass to the finite Hecke action that appears in this paper.

1.4. **The proof and the organization of the paper.** For the quasi-split comparison we start with the endomorphisms of P_C and Q and then allow a finite Hecke kernel between the two copies of the closed object. On the regular locus the two markings of a kernel are related by Weyl transport, which fixes a generic identification compatible with multiplication. Extending this identification across the singular locus depends on a statement in ordinary cohomological degrees. Write Φ for the unperiodized tame coherent functor whose periodic finite restriction is Φ_{fin}.

Theorem 1.7. *For every unshifted finite universal tilting kernel K, including finite direct sums and retracts of averaging words, the coefficient*

$$\mathcal{N}_G(K) = \mathrm{RHom}_{V_G}\big(\Phi(K)Q, Q\big)$$

is a finite projective A-module in ordinary cohomological degree zero. The A-action is the original output action and the boundary R-action is induced by the prescribed norm.

In outer type E_6, the finite calculation verifies the parabolic sources and eliminates the remaining coefficients using nonnegative equations. In types E_7 and E_8, it propagates the vanishing of positive cohomology from known sources. At length 64 in E_8 one connected component requires a geometric vanishing argument. Section C.6 gives the inputs and explains how to check these steps.

One first proves perfectness over the output Cartan and the vanishing of negative derived fibers. To control the positive degrees, one expresses the incidence for parabolic kernels as a derived intersection before applying

duality. A central cocharacter of the common Levi separates the normal terms by their weights. Its weight-zero coefficient is a calculation in a smaller group. Positivity of the decompositions of unshifted tilting kernels propagates this vanishing through the finite Weyl group. Ordinary concentration and perfectness then give projectivity. In particular the coefficient modules are reflexive. The comparison at the generic point and at the height-one root divisors identifies their multiplication and units over the whole Cartan. The circle action is retained when we pass to the periodic categories.

For the pure inner forms the proof uses the relative positions of B' and $aB'a^{-1}$. Filtering the parameter category by these positions makes it possible to ask which coefficients on a stratum are produced by changes of flag from smaller strata. Forgetting one step of a flag leads to a rank-one calculation on a projective line. The characters omitted by induction from smaller positions are precisely those supplied by the compact derived groups of the minimal real Levis. This proves generation by the E_ξ. To compute the maps for a standard kernel one must also keep its relative position, so that three flag conditions enter the correspondence. Over its open position the calculation is a derived intersection of two parabolics over the odd square. The terms normal to their common Levi have strictly negative weights in the dualizing complex. Taking invariants leaves the rank-one coefficient on the norm torus and separates the pure labels.

Section 2 constructs the finite Hecke action on the closed Borel object and compares its coefficients locally. Section 3 proves the ordinary coefficient theorem and obtains the quasi-split comparison over the entire Cartan. Section 4 describes the pure labels, their Cartan supports, and the stratification by flag position. Section 5 constructs the closed generators from compact real Levis. Section 6 proves the orthogonal decomposition, computes the standard returns, and gives the comparison along the root divisors. Appendix A treats the complex group and its finite Hecke category. Appendix B establishes relative properness and the finiteness of the coefficient complexes. Appendix C contains the proofs by root type, with the finite verification procedures and their certificates in Section C.6.

All schemes and groups on the parameter side are over $\mathbb{C}$, unless another field of characteristic zero is specified. Tensor products and fiber products in correspondences are derived. The notation RHom denotes the full cohomologically graded mapping complex. Kernels are written with their output flag first, so $K \star L$ applies L first. The automorphic simple coroot $\alpha_s^\vee \in \pi_1(T_H)$ and the spectral root β_s give the same character of T. Our root and Coxeter conventions follow [Bou02, Hum90, BB05].

Notation. The two Cartan actions record two flags even when the flags lie over the same group element. The table gives the conventions for both actions and the coefficient rings on which they act.

Notation	Meaning and definition
H_0, H, H_ξ	Quasi-split real group, its complexification, and the pure twist indexed by $\xi \in \Theta = H^1(\mathbb{R}, H_0)$ (Section 1.1).
G, B, T, θ	Dual group, pinned Borel and Cartan, and pinned Galois action (Section 1.1).
E^-, τ, r_ρ	Odd component, pinned odd lift, and the torus element $(2\rho^\vee)(i)$ used to mark $t\tau$ by $[tr_\rho]$ ((2.3)).
R, Y, A	$R = \mathcal{O}(T)$, $Y = T/(1-\theta)T$, and $A = \mathcal{O}(Y)$. The R-action on A is pullback by $[x] \mapsto x\theta(x)$ (Section 1.1).
$V, \mathcal{T}(V), P$	Flagged odd parameters, their periodic coherent category, and $P = \mathbb{C}[u, u^{-1}]$, where $\lvert u\rvert = 2$ (Section 1.1).
Z_G, C_J, Δ_w	The Steinberg stack of two flags, the kernel that forgets and refines a flag of type J, and the standard kernel for relative position w (Sections 2.1 and 2.2). For a spectral kernel the first flag is output and the second input.
$P_C, Q, \mathcal{N}_G(K)$	The quasi-split closed real generator, the sheaf of the stable-Borel locus, and the ordinary complex of maps from $\Phi(K)Q$ to Q (Theorem 1.7 and (3.48)).
$I, \delta_I, Y_{\delta_I}, D_I$	Minimal real parabolic type, $\delta_I = w_I\theta$, the norm cover $T/(1-\delta_I)T$, and its image $(T^{\delta_I})^\circ$ (Section 1.2 and Lemma 6.26).
P_ξ, E_ξ	Closed-orbit generator of the real form and its spectral object constructed using the compact derived Levi (Theorem 1.5).
$\Phi_t, U_{\le r}$	Roots taking value one at t and the locus $\operatorname{rank}\Phi_t \le r$ (Section 6.6).
$Z_G(b_0)^\circ$	Connected centralizer of a semisimple even value. It is denoted by H locally in Section 3.4 and retains the full torus T.

2. The closed Borel orbit and its Hecke action

2.1. **The closed orbit, marked coefficients, and reconstruction.** Let H_0 be a connected quasi-split reductive group over $\mathbb{R}$, put $H = H_{0,\mathbb{C}}$, and choose a Borel $B_H = T_H N_H$ defined over $\mathbb{R}$. Let G be the complex dual group, with Cartan T, Weyl group W, and pinned Galois involution θ. We write

$$R = \mathbb{C}[\pi_1(T_H)] = \mathcal{O}(T), \quad \Lambda = \pi_1(T_H) = X^*(T), \quad P = \mathbb{C}[u, u^{-1}], \quad |u| = 2.$$

The automorphic category is $\mathcal{A} = \mathrm{Shv}_{\mathrm{Nilp}}(H_0(\mathbb{R})\backslash H/N_H)$. It uses the full group of real points, arbitrary stalk complexes, and the nilpotent singular-support condition for the H-moment map, with complex covectors viewed by their real parts. Right Cartan monodromy gives its R-action. Let $\mathcal{H}$ be the universal finite Hecke category. Its periodization acts on $\mathcal{A}_{\mathrm{per}}$.

Let E be the geometric dual extension and E_{odd} its odd component. Set

$$\mathcal{V} = \mathcal{P}_E = [(E_{\mathrm{odd}} \times_G^{\mathbf{R}} \widetilde{G})/G], \qquad (a, B') \longmapsto a^2 \bmod N_{B'} \in T.$$

The marked period is the positive boundary-circle class. The Tate convention is (B.1). The finite comparison $\Phi_{\mathrm{fin}} : \mathcal{H}_{\mathrm{per}} \simeq \mathsf{T}(Z_G)$ is constructed in Theorem A.1.

Put $X = H/B_H$, $C = X(\mathbb{R})$, and $\widetilde{X} = (H/N_H)/T_H^{>0}$, where $T_H^{>0}$ is the positive real subgroup and $T_H^c = T_H/T_H^{>0}$. The unique closed real flag orbit is C. Let $i_C : \widetilde{C} \hookrightarrow \widetilde{X}$ and let L_C be its universal equivariant local system. Set $P_C = i_{C*}L_C$, without a shift. On the parameter side, $i : \mathcal{F} \to \mathcal{V}$ is the stable Borel locus $aB'a^{-1} = B'$, and $Q = i_*\mathcal{O}_{\mathcal{F}}$. Write $\mathcal{S}_Q \subset \mathsf{T}(\mathcal{V})$ for the localizing subcategory generated by all compact finite Hecke translates of Q.

The common closed coefficient ring and norm are

$$Y = T/(1-\theta)T, \quad A = \mathcal{O}(Y) = \mathbb{C}[\Lambda^\theta], \qquad N : Y \to T, \quad [t] \mapsto t\theta(t).$$

For the pinned odd lift τ put $z = \tau^2 = (2\rho_G^\vee)(-1)$ and $r_\rho = (2\rho_G^\vee)(i)$. The closed coordinate of $t\tau$ is $[tr_\rho]$. This specifies the norm marking and the signs at root walls. For later use let

$$f = \prod_{\alpha \in \Phi(G,T)} (1 - e^\alpha), \qquad U = D(f) \subset T.$$

Localization always means categorical derived extension of scalars along $R \to R[1/f]$.

2.1.1. *Coefficient categories and compact generators.* The nilpotent cotangent vectors in the fibre $\mathfrak{b}_H$ of $T^*(H/N_H)$ are exactly $\mathfrak{n}_H$. They annihilate the right Cartan directions. Consequently objects of $\mathcal{A}$ are locally constant on the enhanced inverse images of the finitely many real flag orbits. For such an orbit O, choose a real-stable Cartan in the Borel of a point, and let $D_O^c \subset T_H^c$ be the full compact image of its real stabilizer, including components. Set $\Lambda_O = \pi_1(T_H^c/D_O^c), \qquad A_O = \mathbb{C}[\Lambda_O]$. The quotient torus is connected, so A_O is a Laurent polynomial ring.

Lemma 2.1. *There is an equivalence*

$$\mathrm{Shv}_{\mathrm{loc}}(H_0(\mathbb{R})\backslash\widetilde{O}) \simeq \mathrm{Mod}(A_O),$$

whose Cartan map comes from $T_H^c \to T_H^c/D_O^c$. It retains the exact sequence

$$\pi_1((D_O^c)^0) \longrightarrow \Lambda \longrightarrow \Lambda_O \longrightarrow \pi_0(D_O^c) \longrightarrow 0. \tag{2.1}$$

Proof. The real flag stabilizer projects onto its stable real Cartan with unipotent kernel. The further kernel of projection to its compact image is an extension of a real unipotent group by a positive real torus. Both are contractible: successive vector-group quotients handle the first, and logarithm handles the second. These are the stabilizer calculations in [IY25, Lemmas 4.3, 4.5, 4.6, and 5.3]; the central torus remains in the Cartan projection for a reductive group. Removing these kernels gives the local-system category of T_H^c/D_O^c, hence modules over its fundamental group ring. The homotopy sequence of this compact torus quotient gives (2.1). □

For an involution $\delta \in W\theta$, set

$$Y_\delta = T/(1-\delta)T, \quad A_\delta = \mathcal{O}(Y_\delta), \quad N_\delta([x]) = x\delta(x), \quad V_\delta = N_\delta^{-1}(U).$$

Lemma 2.2. *If the stable Cartan of O determines δ, then $A_O = A_\delta = \mathbb{C}[\Lambda^\delta]$, with Cartan map $R \longrightarrow A_\delta, \qquad e^\lambda \longmapsto e^{(1+\delta)\lambda}$. The norm is an isogeny onto $(T^\delta)^0$. The localized stratum category is* $\mathrm{Mod}(\mathcal{O}(V_\delta))$.

Proof. Conjugation acts by $-\delta$ on the compact Cartan lattice, so D_O^c is the full kernel of the compact norm $t \mapsto t\delta(t)$. This gives $T_H^c/D_O^c \simeq ((T_H^c)^\delta)^0$, whose fundamental lattice is Λ^δ and whose quotient map on lattices is $1+\delta$. Its cokernel is the finite group in (2.1); it is not discarded. The character lattice of Y_δ is also Λ^δ. Since $(1+\delta)\Lambda$ contains $2\Lambda^\delta$, it has finite index. This proves the norm assertion and, after scalar extension, the localized description. □

Let L_O^{reg} correspond to the regular A_O-module and put $\Delta_O = j_{O!}L_O^{\mathrm{reg}}$. These standards are unshifted for the generation argument; $\Delta_C = P_C$.

Lemma 2.3. *The Δ_O are compact generators of $\mathcal{A}$. Their finite stable idempotent closure is $\mathcal{A}^c$, and restriction of a compact object to any enhanced orbit has perfect A_O-coefficients.*

Proof. The regular module is a compact generator on each stratum. Equivariant costalks commute with arbitrary direct sums. One can check this on a compact model: choose a maximal compact $K_H \subset H$ preserved by the real involution and put $K_0 = K_H \cap H_0(\mathbb{R})$. The compact enhanced flag is K_H, on which K_0 acts freely. Its quotient maps to the real equivariant quotient with contractible fibre $H_0(\mathbb{R})/K_0$. Descent identifies constructible categories, including orbit inclusions. A finite subanalytic triangulation computes each local costalk by a finite cellular complex, which commutes with direct sums even for universal Cartan coefficients. Adjunction therefore proves compactness of $j_{O!}L_O^{\mathrm{reg}}$. Open–closed localization triangles for the finite stratification prove generation. The general compact generation theorem gives its finite stable idempotent closure. Restriction of a standard is regular on its own stratum and zero on the others, so all compact restrictions are perfect. □

For the following orbit calculations use input-first right convolution. It is related to output-first left action by $F \star K = K^t \cdot F$. The universal Hecke standards are normalized by the shift $[r+\ell(w)]$, $r = \dim T_H^c$, with the compact torus orientation used for the unit. Their inverse identities are

$$\Delta_w^{\mathcal{H}} * \nabla_{w^{-1}}^{\mathcal{H}} \simeq \mathbf{1} \simeq \nabla_{w^{-1}}^{\mathcal{H}} * \Delta_w^{\mathcal{H}}. \tag{2.2}$$

The root-chart calculation in [Tay25, Propositions 2.7–2.8] retains the actual coroot and applies to every root datum. The compact Hecke category is the finite stable idempotent closure of these standards.

2.1.2. *Ascent from the closed real orbit.*

Lemma 2.4. *Let $q : Z \to Y'$ be a surjective homomorphism of compact connected tori, with lattices L, M, kernel lattice $K = \ker(L \to M)$, and $k = \operatorname{rk} K$. Then*

$$Rq_* L_L^{\mathrm{reg}} \iff \mathrm{RHom}_{\mathbb{C}[L]}(\mathbb{C}[M], \mathbb{C}[L]) \simeq \mathbb{C}[M] \otimes \det(K \otimes \mathbb{C})^\vee[-k].$$

The whole output lattice, including $M/\operatorname{im} L$, acts in this formula.

Proof. The saturated subgroup K is a direct summand of L. A basis of K gives the regular sequence $e^{k_1} - 1, \dots, e^{k_k} - 1$ in $\mathbb{C}[L]$. The dual Koszul resolution computes the derived Hom to $\mathbb{C}[L/K]$ with the displayed determinant and shift. The inclusion $L/K \subset M$ has finite index; its group-ring extension is finite étale and its trace pairing identifies the additional dual with $\mathbb{C}[M]$. Derived coinduction describes direct image of local systems, including output monodromy. An invertible character twist rescales a regular group-ring basis and does not change the conclusion. □

Lemma 2.5. *In a real simple ascent from O_x, restriction of $\Delta_{O_x} \star \Delta_s^{\mathcal{H}}$ to each upper enhanced orbit is a finite nonzero sum of shifted regular coefficient local systems, with an invertible orientation line. This holds for both upper orbits in type I and the single actual upper orbit in type II.*

Proof. The rank-one orbit picture of [IY25, Lemma 4.12] has lower fibre $\mathbb{R}P^1$ and upper disc fibres. For a fixed nonreal output flag the elliptic subgroup of $\mathrm{PSL}_2(\mathbb{R})$ acts transitively on the lower projective line. Consequently the pair correspondence is one real group orbit over each actual upper orbit.

Let $D_p \subset D_x$ be the compact image of the pair stabilizer in the input frame. Its two frame images form the graph $\{(b, s(b)) : b \in D_p\}$. One has $D_p^0 = D_x^0$: a maximal compact subgroup of the input stabilizer projects onto D_x, and its identity component maps to a compact connected subgroup of the real affine group fixing a point of $\mathbb{R}P^1$. Such a subgroup is trivial, so it fixes the nonreal output as well. Thus $T_H^c/D_p \to T_H^c/D_x$ is a finite cover of some degree $m > 0$.

After removing contractible kernels the coefficient correspondence is

$$T_H^c/D_x \xleftarrow{p} (T_H^c/D_p) \times T_H^c \xrightarrow{q} T_H^c/D_y, \quad p([t],k) = [t], \quad q([t],k) = [s(t)k].$$

Here k is output frame divided by transported input frame. The open kernel is regular in k, while the input regular module is free of rank m over the covering lattice ring. Thus their product is free of rank m over the product torus ring. The map q is surjective. Lemma 2.4 computes its direct image as m copies of the regular output module, with its determinant and shift. Geometrically the enhanced lower fibre is compact, so this is also the required compactly supported direct image. Its kernel rank is $r+1$, since $\dim D_y^0 = \dim D_x^0 + 1$, agreeing with the fibre dimension and the normalizing shifts. All compact subgroups used here include their finite components. □

Lemma 2.6. *For a complex simple ascent $O \to O'$, convolution on the open correspondence gives an equivalence of their full coefficient local-system categories, transporting the right Cartan action by s. In particular regular coefficients remain regular, up to the geometric shift and an invertible frame line.*

Proof. The rank-one picture is $\pi_s^{-1}(D) = O \sqcup O'$, with $O \simeq D$ and $O' \to D$ an affine-line bundle [IY25, Lemma 4.12(1)]. Each output in O' therefore determines its input uniquely. The two stabilizers have the same real Cartan and contractible unipotent kernels. Their full compact images in the two frames are D_x and $s(D_x)$. The coefficient diagram is the preceding product diagram with $D_p = D_x$ and output $T_H^c/s(D_x)$.

Let $j_x : \Lambda \to \Lambda_x$ be the quotient map. The fibre lattice is $\{(j_x(\gamma), -s\gamma) : \gamma \in \Lambda\}$. For a coefficient module V, choose a basis $\gamma_1, \dots, \gamma_r$, let M_i be the action of $j_x(\gamma_i)$, and let y_i correspond to $s\gamma_i$. The fibre cochains are the Koszul complex on $y_i^{-1}M_i - 1$ in $V \otimes \mathbb{C}[y_1^{\pm 1}, \dots, y_r^{\pm 1}]$. Multiplication by $y_i - M_i$ is injective and its cokernel identifies y_i with M_i. Successive elimination leaves only V in top degree r. This is natural for modules and complexes, because each quotient functor is exact in V. The Hecke-unit shift and torus orientation cancel the top degree and orientation. The output loop $s(a)$ lifts to $(a, 0)$, hence acts by the old action of a. This proves transport by the torus isomorphism s, retaining the entire lattice and the frame line. □

Theorem 2.7. *The compact category $\mathcal{A}^c$ is the finite stable idempotent closure of the compact Hecke translates of P_C. In particular*

$$\mathcal{A} = \mathrm{Loc}\{P_C \star K : K \in \mathcal{H}^c\}, \qquad Q_C = \bigoplus_{w \in W} P_C \star \Delta_w^{\mathcal{H}}$$

is a compact generator.

Proof. Let $\mathcal{D}$ be the indicated closure. It consists of compact objects by (2.2), and it contains $j_{C!}\,\mathrm{Perf}(A)$. Every nonclosed orbit has a complex descending root or a noncompact imaginary root by [IY25, Lemma 6.9]. Reversing it gives a complex or real ascent from a smaller orbit; the central torus does not change this flag geometry.

Induct on orbit dimension. For a lower orbit O_x already obtained, the object $F = \Delta_{O_x} \star \Delta_s^{\mathcal{H}}$ is extension by zero from $S = \pi_s^{-1}(\pi_s(O_x))$. This follows by composing shriek pushforwards from the open relative-position correspondence; the ambient enhanced convolution projection is proper. The triangle on S separating the lower orbit from its upper orbit or orbits is $j_! j^* F \to F \to i_* i^* F$. The last term has perfect lower coefficients by Lemma 2.3, hence lies in $\mathcal{D}$. The first term therefore lies there as well. Lemmas 2.5 and 2.6 identify its upper coefficients as finite nonzero free modules up to invertible twists. In type I, the two upper orbits are open and closed components of their union, so their projections split. A rank-one summand then gives each upper universal standard; in type II there is only one actual upper orbit. The same argument handles complex ascent. Idempotent completeness retains the summands. This proves generation by every Δ_O.

Finally the finite closure of the displayed summands of Q_C is Hecke-stable: a convolution $\Delta_w^{\mathcal{H}} * K$ of compact kernels is a finite construction from the finite Hecke standards. It contains P_C, so equals $\mathcal{A}^c$ by the first assertion. □

2.1.3. *The stable Borel stratum and its closed marking.* The geometric Galois involution is $\mathrm{Ad}(r_\rho)\theta$ [Zhu17, Lemma 5.5.5]. If σ is its order-two representative, then $\tau = r_\rho^{-1}\sigma$ acts by the pinned involution and has square $z = z^{-1}$. In the fixed-Borel presentation,

$$\mathcal{V} = [(E_{\mathrm{odd}} \times_G^{\mathbf{R}} B)/B], \quad \mathcal{F} = [B\tau/B], \quad \mathcal{T} = [T\tau/T], \quad q : \mathcal{F} \to \mathcal{T}.$$

The coordinate $y = [tr_\rho]$ satisfies

$$(t\tau)^2 = t\theta(t)z = (tr_\rho)\theta(tr_\rho) = N(y). \tag{2.3}$$

The torus-extension map $t\tau \mapsto tr_\rho\tau_0$ to an order-two representative is an actual extension isomorphism, because $r_\rho\theta(r_\rho) = z$. It therefore preserves the boundary transport and u.

Lemma 2.8. *A θ-fixed strictly dominant integral cocharacter λ_0 makes $\mathcal{F}$ a derived Θ-stratum of $\mathcal{V}$, with centre $\mathcal{T}$, where $\Theta = [\mathbb{A}^1/\mathbb{G}_m]$. The evaluations are i and q. The functor $i_*q^* : \mathrm{Coh}(\mathcal{T})_0 \longrightarrow \mathrm{Coh}(\mathcal{V})$ is fully faithful, preserves bounded coherence, and is boundary-circle equivariant. The subscript is weight zero for the central inertia cocharacter λ_0.*

Proof. Average a strictly dominant rational cocharacter under θ and clear denominators. Its parabolic is B and its centralizer is T. The attracting loci in E_{odd}, G, B are $B\tau, B, B$, and their fixed loci are $T\tau, T, T$. Mapping from Θ preserves derived fibre products. The quotient filtered-component formula [HL21, Lemma 1.6.1] therefore gives

$$[B\tau/B] \times^{\mathbf{R}}_{[B/B]} [B/B] = \mathcal{F}, \qquad [T\tau/T] \times^{\mathbf{R}}_{[T/T]} [T/T] = \mathcal{T}.$$

The second map in each product is the identity. Thus the attracting stack is classical even when $\mathcal{V}$ is derived. Evaluation at one is the closed embedding i and at zero is q. The stacks are locally almost of finite presentation with affine diagonal; hence [HL21, Theorem 2.2.2(3)] gives full faithfulness on the fixed weight-zero part of $D^-\,\mathrm{Coh}$. The map q is smooth: after the torus atlas it is $[B\tau/U] \to T\tau$, presented by the smooth projection $B\tau \to T\tau$. Thus its pullback has finite Tor amplitude, and closed pushforward retains bounded coherence. Apply the boundary action to filtered objects. It preserves the λ_0 component, and evaluation and associated graded commute with every subdivision of transport. This proves equivariance with all coherences and on the full Hom complexes. □

Lemma 2.9. *$Q = i_*q^*\mathcal{O}_{\mathcal{T}}$ is compact in $\mathsf{T}(\mathcal{V})$, and*

$$\mathrm{RHom}(Q, Q) \simeq A[u, u^{-1}]$$

as an algebra, with R-structure N^*. Consequently $\mathrm{Loc}(Q) \simeq \mathrm{Mod}(A \otimes P)$.

Proof. The structure sheaf has weight zero. By Lemma 2.8, its entire endomorphism algebra and circle action compute those of Q. The torus is linearly reductive, so $R\Gamma(\mathcal{T}, \mathcal{O}_{\mathcal{T}}) = \mathcal{O}(T)^T = \mathbb{C}[\Lambda^\theta] = A$ in degree zero, with its usual multiplication. An ordinary algebra has a discrete automorphism space; its circle action is trivial. Thus the small invariant Hom algebra is $A[u]$, and its localization is $A[u, u^{-1}]$. A small equivariant coherent object's image is compact in the prescribed Ind-completion. Compact Morita theory now gives the stated localizing category. The marking is (2.3). □

Theorem 2.10. *The periodized category supported on C is $R \otimes P$-linearly equivalent to $\mathrm{Loc}_{\mathsf{T}(\mathcal{V})}(Q)$. It sends P_C to Q and identifies the full endomorphism algebras with $A \otimes P$. On the spectral torus it uses the trivial character of the full stabilizer, followed by i_*q^*.*

Proof. For C the involution in Lemma 2.2 is θ. Closed extension is fully faithful, so $\mathrm{RHom}(P_C, P_C) = A$ with norm marking. Both categories are therefore $\mathrm{Mod}(A \otimes P)$, the second by Lemma 2.9. Their equivalence is explicitly $M \mapsto \mathrm{RHom}(P_C \otimes P, M) \otimes_{A\otimes P} Q$. It uses the computed multiplication and hence preserves full Hom and composition. □

2.1.4. *Signed Weyl transport on closed coefficients.* Let D be the full compact image of a real maximal torus in B_H. For a kernel K restricted to both closed flags, write

$$J_K(M) = (p_1)_!\big(K|_{\widetilde{C}\times\widetilde{C}} \otimes p_2^*M\big), \qquad J_w(M) = i_C^*((i_{C*}M) \star \Delta_w^{\mathcal{H}}).$$

Thus $J_w = J_{(\Delta_w^{\mathcal{H}})^t}$.

Lemma 2.11. *If $\theta(w) \neq w$, then $J_w = 0$. For $w \in W^\theta$ the real Borel pairs in relative position w form one real group orbit. Over a fixed output flag the ordinary fibre is $U_{w^{-1}}(\mathbb{R})$, of real dimension $\ell(w)$. Its enhanced coefficient map is $(T_H^c/D) \times T_H^c \longrightarrow T_H^c/D, \qquad ([t], k) \longmapsto [w^{-1}(t)k]$. The two regular coefficient systems together are regular on the product lattice.*

Proof. A real pair has Galois-fixed relative position. For a fixed w, the Tits representative of a reduced expression is independent of that expression and is fixed by the real pinning. Real Bruhat coordinates give the input cell $U_{w^{-1}}(\mathbb{R})$; conjugate pairs of roots contribute two real dimensions, so the dimension is $\ell(w)$. Choose output frame 1 and input frame $\dot{w}^{-1}$. The pair stabilizer contains the full real Cartan, with compact images $D, w^{-1}(D)$, and its other kernel is contractible. Quotient by its graph, then use input frame and frame ratio k, to obtain the product torus. Both regular systems are retained separately, proving the final assertion. □

For $w \in W^\theta$ let o_w be the top compactly supported cohomology line of the real input cell, with its stabilizer action and no chosen basis. Define

$$\epsilon_w(\mu) = (-1)^{\langle \rho_H - w^{-1}\rho_H, \mu\rangle}, \quad \mu \in \Lambda^\theta, \qquad \varphi_w(x) = \epsilon_w w^{-1} x. \tag{2.4}$$

The difference of Weyl translates of ρ_H is integral, even when ρ_H itself is not.

Proposition 2.12. *The functor J_w is transport by φ_w, with the retained line o_w. Equivalently its action on a coefficient module is $e^\mu m = \epsilon_w(\mu) e^{w\mu} m$. It is an equivalence on all local systems, and $\mathrm{RHom}(P_C \star \Delta_w^{\mathcal{H}}, P_C) \simeq \mathrm{RHom}_A(J_w(A), A)$ is free of rank one over A in degree zero. For $\theta(w) \neq w$ this Hom is zero.*

Proof. The standard kernel and proper ambient projection compute compactly supported integration over the preceding correspondence. Relative Poincaré duality retains the real-cell orientation and a shift $[-\ell(w)]$; the compact frame torus supplies its orientation and shift $[-r]$. The Hecke normalization cancels the degree, pairing the torus orientation with that of the unit. The map on fundamental groups is

$$\Lambda^\theta \oplus \Lambda \longrightarrow \Lambda^\theta, \quad (a, \gamma) \longmapsto w^{-1}a + (1+\theta)\gamma.$$

Its first summand is an isomorphism. The input coefficient complex tensored with the regular frame module is induced from that summand. Derived left Kan extension therefore leaves the input module transported by w^{-1}, naturally even for unbounded complexes.

The roots of the input cell sum to $\rho_H - w^{-1}\rho_H$. Their determinant on D is real of absolute value one and is precisely the real orientation character; the connected unipotent stabilizer preserves orientation. A loop $\mu \in \Lambda^\theta$ lifts under the compact norm to $\exp(\pi i t \mu)$, ending at $\mu(-1) \in D$. Its orientation holonomy is $\epsilon_w(\mu)$. Transport of the input along an output loop μ is by $w\mu$. This gives (2.4) with o_w retained. Closed adjunction gives the Hom formula; its input is free of rank one. The vanishing statement follows from the absence of real pairs. □

Proposition 2.13. *Conjugation by $\dot{w}^{-1}$ gives an actual circle-equivariant automorphism n_w of $\mathcal{T}$ whose map in the coordinate $x = [sr_\rho]$ is φ_w. The closed Fourier equivalence of Theorem 2.10 identifies periodized J_w with $(n_w)_*$ tensored with o_w. Moreover $\varphi_w \circ \varphi_v = \varphi_{vw}$, $\qquad N \circ \varphi_w = w^{-1} \circ N$.*

Proof. A θ-fixed Tits representative commutes with the pinned lift, so it sends $s\tau$ to $w^{-1}(s)\tau$ and conjugates the gauge group and square by the same map. It therefore respects all cyclic transport diagrams. In closed coordinates the map is $x \mapsto w^{-1}x[r_\rho/(w^{-1}r_\rho)]$. Evaluation on $\mu \in \Lambda^\theta$ gives

$$i^{\langle 2\rho_G^\vee - w^{-1}(2\rho_G^\vee), \mu\rangle} = (-1)^{\langle \rho_H - w^{-1}\rho_H, \mu\rangle}.$$

This uses the actual dual lattices, including finite characters. The ratio satisfies the cocycle identity. It is θ-fixed and its square is $z/w^{-1}z = 1$, so its norm is one. This proves both formulas and the comparison of coefficient functors, with the unchanged orientation fibre. □

2.1.5. *Reconstruction as a marked coefficient algebra.* Define $D_C \in \mathcal{H}_{\mathrm{per}}$ on compact kernels by

$$\mathrm{RHom}_{\mathcal{H}_{\mathrm{per}}}(K, D_C) = \mathrm{RHom}_{\mathcal{L}_{C,\mathrm{per}}}(J_K L_C, L_C), \tag{2.5}$$

where $\mathcal{L}_C$ is the closed-orbit local-system category and the displayed L_C is freely periodized.

Lemma 2.14. *Formula (2.5) represents a unital associative algebra object. For maps $a : J_K L_C \to L_C$ and $b : J_L L_C \to L_C$, multiplication is*

$$J_{K*L} L_C \longrightarrow J_K J_L L_C \xrightarrow{J_K b} J_K L_C \xrightarrow{a} L_C.$$

Proof. Proper base change identifies J_K with $i_C^*(K \cdot i_{C*}(-))$. The first arrow inserts the unit $\mathrm{id} \to i_{C*}i_C^*$ at the intermediate flag. For any string of kernels, successive restrictions are the restriction to their closed intersection in the common flag-tuple space. These maps provide coherent associativity and units in every arity; on the diagonal the transformation is the identity. The functor on compact kernels is exact and P-linear, hence represented under enriched Yoneda for $\mathrm{Ind}(\mathcal{H}_{\mathrm{per}}^c)$. The displayed natural pairing extends to $D_C * D_C \to D_C$ because convolution preserves colimits in each variable. The identity of L_C is its unit. This construction uses the map from the full convolution correspondence and does not assume that every intermediate flag with real endpoints is itself real. □

Lemma 2.15. *Let $\mathcal{V}'$ be a compactly generated stable presentable monoidal category with compact unit, dualizable compact objects on both sides, and separately continuous product. Suppose it acts continuously on $\mathcal{B}$ and a compact object m generates under its compact action. Then*

$$\mathcal{B} \simeq \mathrm{RMod}_{D_m}(\mathcal{V}'), \qquad \mathrm{RHom}_{\mathcal{V}'}(K, D_m) = \mathrm{RHom}_{\mathcal{B}}(K \cdot m, m).$$

The product is composition, the equivalence is $\mathcal{V}'$-linear, and m is the free module.

Proof. Let $F(V) = V \cdot m$ and let U be its right adjoint. Compact dualizable kernels have colimit-preserving right adjoint actions, so their translates of m are compact. Adjunction against these kernels makes U preserve direct sums; exactness then makes it preserve all colimits. Their generation makes U conservative. For dualizable K the projection map $K * U(b) \to U(K \cdot b)$ is an equivalence: mapping out of V identifies each side with maps from $F({}^\vee K * V)$ to b. Continuity extends it to every K, compatibly with multiplication. Hence $UF(V) = V * U(m)$. The adjunction counit and unit give the product and unit on $D_m = U(m)$; on coefficients their product is $a \circ (K \cdot b)$. Barr–Beck [Lur17, Theorem 4.7.3.5] applies because U is conservative and preserves geometric realizations. Algebras over this monad are right D_m-modules, with the stated linearity and free object. □

Theorem 2.16. *There are marked Hecke-module equivalences*

$$\mathcal{A}_{\mathrm{per}} \simeq \mathrm{RMod}_{D_C}(\mathcal{H}_{\mathrm{per}}) \simeq \mathrm{RMod}_{\Phi_{\mathrm{fin}}(D_C)}(\mathsf{T}(Z_G)),$$

carrying P_C to the free module. They preserve the boundary R-action and u and remain equivalences after every categorical derived extension of scalars $R \to S$.

Proof. Hecke standards are invertible by (2.2); their finite cones and retracts are dualizable. The hypotheses of Lemma 2.15 therefore hold after periodization, and Theorem 2.7 gives generation. Closed adjunction identifies the coefficient formula with (2.5); composition inserts exactly the intermediate restriction of Lemma 2.14. Transport the algebra and its action through Φ_{fin} for the second equivalence. The Cartan actions come from external unit endomorphisms and are preserved together with u by Theorem A.1. Finally apply $- \otimes_{\mathrm{Mod}(R)} \mathrm{Mod}(S)$ to the equivalence and its inverse. This last assertion requires no flatness of S. □

By Lemma 2.15, the spectral generator defines $D_Q \in \mathsf{T}(Z_G)$ through $\mathrm{RHom}(K, D_Q) = \mathrm{RHom}_{\mathsf{T}(\mathcal{V})}(K \cdot Q, Q)$. Its multiplication is actual composition for the square of the odd parameter.

Proposition 2.17. *The marked comparison of Theorem 1.1 is equivalent to an equivalence of unital associative algebra objects $\Phi_{\mathrm{fin}}(D_C) \simeq D_Q$ inducing the specified identification of $A[u, u^{-1}]$ on maps from the monoidal unit.*

Proof. Compactness of Q and generation of $\mathcal{S}_Q$ give its right-module description by Lemma 2.15. Extension along the marked algebra equivalence carries free module to free module and induces the prescribed functor on its localizing category, determined by its endomorphism algebra. Theorem 2.16 then gives the desired comparison. Conversely a marked Hecke-module equivalence identifies naturally all maps $K \cdot P_C \to P_C$ and $\Phi_{\mathrm{fin}}(K) \cdot Q \to Q$, with compositions for arbitrary strings. Enriched Yoneda gives the algebra equivalence and its unit marking. □

2.2. **Parabolic coefficients and the marked rank-one action.**

2.2.1. *The ordinary two-source formula.* Put $\mathcal{C}_o = [E_{\mathrm{odd}}/G]$, $X_b = [B/B]$, $Y_g = [G/G]$, and $\mathcal{V} = \mathcal{C}_o \times^{\mathbf{R}}_{Y_g} X_b$. The first map is squaring. Let $\rho : \mathcal{F} \to \mathcal{C}_o$ forget the stable Borel, and form

$$W_o = \mathcal{F} \times^{\mathbf{R}}_{\mathcal{C}_o} \mathcal{F}, \qquad h : W_o \to Z_G, \quad (a, B_1, B_2) \mapsto (a^2, B_1, B_2).$$

Lemma 2.18. *W_o is classical and a local complete intersection on a smooth atlas. Its relative-position strata are indexed by W^θ. The map h is affine. The map ρ is proper and perfect, with a circle-equivariant relative orientation $\mathcal{O}_{\mathcal{F}} \simeq \omega_\rho$ in degree zero.*

Proof. Over E_{odd} the map is $G \times^B (B\tau) \to E_{\mathrm{odd}}$, projective between smooth varieties of equal dimension $\dim G$. The two-source atlas is the pullback of the diagonal of E_{odd}. A preserved pair has $w = \theta(w)$; conversely choose a θ-fixed Tits lift and put $H_w = B \cap \dot{w} B \dot{w}^{-1}$. Its odd preserving elements are $H_w\tau$. The corresponding atlas stratum $G \times^{H_w} (H_w\tau)$ is smooth and irreducible of dimension $\dim G$, since H_w is a torus extended by a connected unipotent group. The finitely many strata give this dimension to every component. The $\dim G$ diagonal equations have the expected codimension in a smooth ambient scheme, so form a regular sequence. Their Koszul complex computes the classical fibre product. Over an atlas of Z_G, adjoining the affine variable a with its square and flag-preservation equations is a closed incidence in a product with E_{odd}; hence h is affine.

Right Maurer coordinates identify the vertical tangent of $G \times^B (B\tau)$ over G/B with $\mathfrak{b}$. Together with $\mathfrak{g}/\mathfrak{b}$, in that order, its determinant is $\det \mathfrak{g}$, also the tangent determinant of the target. Their ratio supplies the relative orientation without choosing a volume form. It is equivariant and descends. The maps of its two degree-zero sheaves have discrete mapping space, so the natural connected boundary circle preserves this trivialization. □

Proposition 2.19. *For every $K \in \mathrm{Coh}(Z_G)$,*

$$\mathrm{RHom}_{\mathcal{V}}(K \cdot Q, Q) \simeq \mathrm{RHom}_{\mathrm{QCoh}(W_o)}(h^*K, \mathcal{O}_{W_o}). \tag{2.6}$$

For a small equivariant kernel, its Tate coefficient is obtained by taking invariants of this ordinary Hom and then inverting the original u.

Proof. Set $D = \mathcal{V} \times^{\mathbf{R}}_{\mathcal{C}_o} \mathcal{F}$, with $p : D \to \mathcal{V}$ forgetting the stable input and $v : D \to Z_G$ recording its boundary square and two flags. Base change in the input gives $K \cdot Q = p_* v^* K$. The map v is the derived base change of the square map $\mathcal{F} \to X_b$, which has finite Tor dimension on smooth atlases $B\tau \to B$. Thus this action lies in bounded coherent objects. The map p is a base change of ρ, so $p^! Q \simeq p^* Q$. If $l : W_o \to D$ also imposes output stability, derived base change gives $p^* i_* \mathcal{O}_{\mathcal{F}} = l_* \mathcal{O}_{W_o}$. The two adjunctions yield

$$\mathrm{RHom}_{\mathcal{V}}(p_* v^* K, Q) = \mathrm{RHom}_D(v^* K, l_* \mathcal{O}_{W_o}) = \mathrm{RHom}_{W_o}(h^* K, \mathcal{O}_{W_o}).$$

The last adjunction is in quasicoherent complexes and requires no Tor bound for h. All maps carry the same boundary transport. Small invariants, localization, and then finite cones and retracts give the asserted Tate coefficients for all compact kernels. □

Here is the multiplication in (2.6). Let

$$D^3 = \mathcal{F} \times^{\mathbf{R}}_{\mathcal{C}_o} \mathcal{V} \times^{\mathbf{R}}_{\mathcal{C}_o} \mathcal{F}, \quad W^3 = \mathcal{F} \times^{\mathbf{R}}_{\mathcal{C}_o} \mathcal{F} \times^{\mathbf{R}}_{\mathcal{C}_o} \mathcal{F}.$$

Write $d : D^3 \to W_o$, $b : W^3 \to W_o$ for forgetting the middle flag and $e : W^3 \to D^3$ for imposing its stability. For $T_{K,L}$ the tensor of the two pulled-back kernels, $h^*(K * L) = d_* T_{K,L}$. Coefficients f, g multiply by

$$d_* T_{K,L} \longrightarrow b_* e^* T_{K,L} \xrightarrow{b_*(f_{12} \otimes g_{23})} b_* \mathcal{O}_{W^3} \xrightarrow{\mathrm{Tr}_b} \mathcal{O}_{W_o}. \tag{2.7}$$

The first arrow restricts the intermediate flag; the last is the proper duality counit using the pullback of $\omega_\rho \simeq \mathcal{O}$. Reversing the adjunctions in the preceding proof inserts the closed unit at the middle flag and then this trace, which proves the formula. For longer strings, the composition laws for these units and counits give every associativity coherence. Stability is imposed after forming the boundary convolution, also at singular holonomy.

2.2.2. *A single parabolic coefficient and its two counits.* For J a set of simple roots let $p_J : \mathcal{P}_E \to X_J = [(E_{\mathrm{odd}} \times^{\mathbf{R}}_G P_J)/P_J]$. The kernel $C_J = \mathcal{O}_{X_b \times^{\mathbf{R}}_{[P_J/P_J]} X_b}$ acts by $p_J^* p_{J*}$, with diagonal counit. The map p_J is proper and has finite Tor dimension: before the square base change its atlas is the map $P_J \times^B B \to P_J$ between smooth varieties. Thus both functors preserve bounded coherence and their adjunctions retain the original circle. Put $I = J \cap \theta J$ and

$$B_I = B \cap L_I, \quad Y_I = [L_I\tau / L_I], \quad \rho_I : [B_I\tau / B_I] \to Y_I, \quad M_I = R\rho_{I*} \mathcal{O}_{[B_I\tau/B_I]}.$$

The map ρ_I imposes preservation by the odd element, whereas p_J imposes containment of its square.

Lemma 2.20. *There is a closed derived Θ-stratum $k : S_I = [P_I\tau/P_I] \to X_J$ with centre Y_I and projection q, and $p_{J*}Q \simeq k_* q^* M_I$ in the original coherent category with its circle action.*

Proof. Choose a θ-fixed integral dominant cocharacter vanishing on I and positive elsewhere, taking zero if I is the entire diagram. Its parabolic and centralizer are P_I, L_I. The attracting loci in E_{odd}, G, P_J are $P_I\tau, P_I, P_I$, and fixed loci are $L_I\tau, L_I, L_I$. The filtered fibre-product calculation of Lemma 2.8 gives S_I and Y_I.

To see closedness, the pairs of parabolics of types $J, \theta J$ containing a common Borel form the closed orbit G/P_I, since $P_J \cap P_{\theta J} = P_I$. Pull it back by $(a, D) \mapsto (D, aDa^{-1})$. Their intersection D_I is a parabolic of type I, preserved by a because $a^2 \in D$. Conversely an a-stable parabolic of type I has a unique containing parabolic of type J. These are inverse constructions in families by the homogeneous description of the closed orbit, proving the assertion on truncations and hence that the filtered component is a Θ-stratum. Borels of P_I are inverse images of Borels of L_I. Thus $\mathcal{F} \to S_I$ is the base change of ρ_I. The projection q is smooth on its atlas $P_I\tau \to L_I\tau$. Flat base change gives the pushforward formula. All maps forget reductions or take Levi quotients and therefore commute with boundary transport and its compositions. □

Lemma 2.21. *M_I is a vector bundle concentrated in degree zero.*

Proof. On $L_I\tau$ it is the direct image of the structure sheaf by $d : L_I \times^{B_I} (B_I\tau) \to L_I\tau$. This is projective between smooth varieties of the same dimension. A θ-invariant regular dominant cocharacter supplies t with root-regular square $(t\tau)^2$. A Borel preserved by $t\tau$ contains that square, so the fibre is finite and nonempty. Properness and semicontinuity therefore make d dominant and generically finite. The tangent determinant,

$\det \mathfrak{b}_I$ times $\det(\mathfrak{l}_I/\mathfrak{b}_I)$, equals the target determinant. Hence $\omega_d \simeq \mathcal{O}$. Grauert–Riemenschneider vanishing in its projective generically finite form gives $R^j d_* \mathcal{O} = 0$ for $j > 0$; see [EL24, Remark 4.1.6]. Duality gives $\mathrm{RHom}(d_*\mathcal{O}, \mathcal{O}) = d_*\mathcal{O}$. Over a regular local target ring a minimal free resolution of positive length b would have nonzero $\mathrm{Ext}^b(-, \mathcal{O})$, since its last dual differential is zero modulo the maximal ideal. Thus $d_*\mathcal{O}$ is locally free, and descends. □

Proposition 2.22. *The ordinary coefficient of C_J is $\Gamma(Y_I, \operatorname{End} M_I)$ in degree zero; its Tate coefficient is this algebra tensored with P. Multiplication is composition using the comultiplication of $p_J^* p_{J*}$. Precomposition with its counit sends $a \in A$ to multiplication by the corresponding stable Borel function on M_I.*

Proof. The cocharacter of Lemma 2.20 is central in L_I and acts trivially on $[B_I\tau/B_I]$ and its structure sheaf, so M_I has weight zero. Weight-zero full faithfulness [HL21, Theorem 2.2.2(3)] gives $\mathrm{RHom}_{X_J}(p_{J*}Q, p_{J*}Q) = \mathrm{RHom}_{Y_I}(M_I, M_I)$. The latter is in degree zero by Lemma 2.21, affineness of $L_I\tau$, and exactness of reductive invariants. The circle acts trivially on this ordinary algebra, so invariants and localization tensor it with P. For $L = p_J^*$, $V = p_{J*}$, the adjoint of $b : LVQ \to Q$ is $Vb \circ \eta_{VQ}$. The triangle identities identify $b_1 \circ LV(b_2) \circ L\eta_{VQ}$ with composition of these adjoints. The counit becomes the identity; postcomposing it by a pushes forward multiplication by the stable Borel function. If this embedding is $\iota : A \to \Gamma \operatorname{End} M_I$, endpoints act by $e \mapsto \iota(a_1) e \iota(a_2)$. Both R-actions use N^* and agree on R^{W_J}. □

If $J \cap \theta J = \varnothing$, then ρ_I is the identity of $\mathcal{T}$. The actual counit therefore induces an isomorphism $\mathrm{RHom}(Q, Q) \to \mathrm{RHom}(C_J \cdot Q, Q)$, including on every root divisor. The opposite counit map records a different coefficient. For $a \in B\tau$ define

$$c_J(a) = \det(1 + \operatorname{Ad} a \mid \mathfrak{p}_I/\mathfrak{b}) \det(1 - \operatorname{Ad}(a^2) \mid \mathfrak{p}_J/\mathfrak{p}_I) \in A. \tag{2.8}$$

Proposition 2.23. *Let $N_I = \Gamma(Y_I, M_I \otimes M_I)$ and let $\mu_I : N_I \to A$ be multiplication. Before circle invariants, postcomposition with $\varepsilon_Q : C_J \cdot Q \to Q$ is $N_I \xrightarrow{c_J \mu_I} A$. Consequently $\mathrm{RHom}(Q, \operatorname{fib} \varepsilon_Q)$ has $H^0 = \ker \mu_I$, $H^1 = A/(c_J)$, and no other cohomology. Its Tate Hom is the two-term complex $[N_I \xrightarrow{c_J \mu_I} A]$ in degrees $0, 1$ tensored with P.*

Proof. The flag-forgetting maps have relative dualizing line canonically trivialized by the tangent determinants as in Lemma 2.18. For ρ_I denote its trace by $\delta_I : M_I \to \mathcal{O}_{Y_I}$. The perfect pairing $\delta_I(vw)$ identifies M_I with its dual and identifies $v \otimes w$ with $x \mapsto v\delta_I(wx)$ in $\mathrm{End}(M_I)$. The right adjunction for p_J thus gives the source N_I in degree zero. Moreover $\Gamma(Y_I, M_I) = \Gamma([B_I\tau/B_I], \mathcal{O}) = A$: a fixed dominant cocharacter contracts the unipotent coordinates, and twisted-invariant torus functions pull back inversely.

Compute on the dense open where $a = t\tau$ has root-regular square. Both flag maps are finite étale there. Their Jacobians in the fixed orientations are

$$J_p = \det(1 - \operatorname{Ad}(a^2) \mid \mathfrak{p}_J/\mathfrak{b}), \qquad J_\rho = \det(1 - \operatorname{Ad} a \mid \mathfrak{p}_I/\mathfrak{b}).$$

The oriented section of the relative dualizing line on a sheet is $1/J$ in canonical étale coordinates, so δ_I sums sheet values divided by J_ρ. At the output stable Borel the tensor $v \otimes w$ has diagonal matrix value $v(B)w(B)/J_\rho(B)$. The right unit inserts this sheet with factor J_p, and the left counit projects to it. The resulting map is $(J_p/J_\rho)\mu_I = c_J\mu_I$, since $1 - \operatorname{Ad}(a^2) = (1 - \operatorname{Ad} a)(1 + \operatorname{Ad} a)$ on $\mathfrak{p}_I/\mathfrak{b}$. This is linear in local tensors and hence applies to arbitrary global sections. Both maps take values in the integral ring A; equality on its dense open proves global equality. The section $v \mapsto v \otimes 1$ makes μ_I surjective. The generic value of c_J is nonzero, so it is a nonzerodivisor in A. The fibre cohomology follows. All orientations and maps retain their boundary equivariance; their degree-zero diagram takes invariants by tensoring with $\mathbb{C}[u]$ and then localizes. □

For fixed simple s, let β be the root on $\mathfrak{p}_s/\mathfrak{b}$. The pinning fixes this root vector and $e^\beta(r_\rho) = -1$. Thus $c_s = 1 - e^\beta(y)$. The boundary root divisor pulls back to $e^\beta(y)^2 = 1$, while the odd group above is supported on the scheme $e^\beta(y) = 1$.

2.2.3. *The ordinary simple kernel and central descent.*

Lemma 2.24. *Assume H is adjoint, and let Ξ be the big universal tilting object. Before taking circle invariants the ordinary tame functor sends Ξ to $\mathcal{O}_{Z_G}$ and its counit to diagonal restriction. For finite R-modules N it sends the finite standards $\Delta_w(N)$ and their finite extensions to the ordinary coherent heart.*

Proof. We work before circle invariants. Under Φ, right Whittaker averaging is the output projection $p_1 : Z_G \to X_b$ [DT26a, Theorems 1.1 and 3.2(b)]. Its left adjoint is p_1^*, of finite Tor amplitude. The averaging comonad is convolution by Ξ [DT26a, Lemma 4.3]; its value and counit at the unit give $\Phi(\Xi) = \mathcal{O}_{Z_G}$ and the diagonal restriction.

To detect the cohomology sheaves, take strictly dominant λ, μ and a finite-dimensional G-module V'. Use the test object $E = W_{-\lambda} * \Xi * Z(V') * W_{-\mu}$, $\Phi(E) = \mathcal{O}_{Z_G}(-\lambda, -\mu) \otimes V'$. The two Whittaker module actions give this identity, including the line bundles and central action. For a finite R-module N and finite Weyl element w, move the last Wakimoto factor across Hom. The identity $\ell(wt_\mu) = \ell(t_\mu) - \ell(w)$ gives

$$\mathrm{RHom}(E, \Delta_w(N)) = \mathrm{RHom}(W_{-\lambda} * \Xi * Z(V'), \nabla_{wt_\mu}(N')),$$

with the transported marking on N'. By [Tay26, Theorem 4(b)], $\Xi * Z((V')^\vee) * \nabla_{t_\lambda}$ has a costandard filtration. Over $\mathbb{C}$ every finite-dimensional representation has the required costandard filtration. Rigid duality, centrality, and the self-adjoint orientation of the big kernel give a standard filtration of the first argument. Standard–costandard adjunction then places the Hom in degree zero. The same holds for finite extensions of $\Delta_w(N)$. On the coherent side these tests compute sections after an ample twist on the two flag factors and pairing with V'. Serre vanishing and relative global generation over the affine base G detect any nonzero cohomology sheaf; reductivity lets some V' detect its sections. Thus the images of these objects lie in the ordinary coherent heart. This is the universal version of [Bez16, Proposition 39]; the stronger filtration input and its use for this heart statement are explicit in [Tay26, Theorem 4 and Remark 16]. □

Lemma 2.25. *If H is adjoint, the ordinary tame functor takes Δ_s to the fibre of $\mathcal{O}_{Z_s} \to \mathbf{1}$, where $Z_s = X_b \times^{\mathbf{R}}_{[P_s/P_s]} X_b$. After Tate localization it acts by the fibre of $p_s^* p_{s*} \to \mathrm{id}$.*

Proof. Lemma 2.24 identifies the image and the counit of the big averaging kernel. We determine the quotient selected by the two simple cells.

Restrict Ξ to the two simple Bruhat cells and push back. The map $b_s : \Xi \twoheadrightarrow \Xi_s$ is perverse surjective with standard-filtered kernel [Tay25, Proposition 3.13 and Lemma 4.6]. Its counit factors through Ξ_s and is a generator of $\mathrm{Hom}(\Xi_s, \mathbf{1}) = R$ by Whittaker adjunction. Its fibre is Δ_s by [Tay25, equation (4.4)]. The heart calculation makes $\Phi(b_s)$ a quotient of $\mathcal{O}_{Z_G}$. Apply it also after every finite target tensor N: the standards are R-flat, and the quotient remains in the heart, so it is R-flat. Generically it selects the two components $1, s$, since standard markings are their distinct Weyl graphs and standards are invertible. On either generic component, a nonzero quotient of its structure sheaf is the structure sheaf: it is a torus classifying stack over the marked generic field.

The geometric quotient $\mathcal{O}_{Z_G} \to \mathcal{O}_{Z_s}$ has the same generic kernel and is R-flat. To check flatness, reduce to the rank-one Levi. The two incidence strata are $L_s \times^{H_w} H_w$, $w = 1, s$, each of dimension $\dim L_s$, and their Cartan fibres have dimension $\dim L_s - \dim T$. The diagonal equations are a regular sequence by the expected-codimension argument of Lemma 2.18. The incidence is Cohen–Macaulay with equidimensional Cartan fibres, so miracle flatness applies. An R-flat quotient embeds in its generic localization and its kernel is the inverse image of the generic kernel. This identifies the two quotients and, because their composites to the unit agree, their diagonal counits. The maps of these heart objects have discrete mapping spaces, so their given circle structures are respected. Taking the fibre and then localizing proves the statement. □

Lemma 2.26. *For $H' = H/Z(H_{\mathrm{der}})$ with dual central cover $G' \to G$, the ordinary coefficient $\mathrm{RHom}_{\mathcal{P}_E}(\Phi(\Delta_w) \cdot Q, Q)$ is a retract of the corresponding primed coefficient, before circle invariants.*

Proof. The pinning and θ lift, and the pinned extensions map compatibly because their $(2\rho^\vee)(-1)$ agree. Let $\beta : Z_{G'} \to Z_G$ and D' be the finite central kernel. The central descent in [DT26a, proof of Theorem 1.1] and [DT26b, Sections 2.1 and 17.1–17.4] identifies spectral pullback with forgetting central descent on the constructible side. On constant loops this is direct image by the finite cover $H/N_H \to H'/N_{H'}$. Its regular local system on a w-stratum becomes the regular module of the larger lattice, with unchanged shift and closure. Therefore $\beta^* \Phi(\Delta_w) = \Phi'(\Delta'_w)$.

The square-compatible map $\gamma : W'_o \to W_o$ is on incidence atlases a finite étale pullback of $G'\tau' \to G\tau$. Changing gauge group adds exact D'-invariants. Locally its pushforward structure algebra is functions on the finite nonempty set of lifts, then functions constant on D'-orbits. It is a finite étale algebra of positive rank; normalized trace splits its unit. Adjunction consequently makes $\mathrm{RHom}_{W_o}(M, \mathcal{O}_{W_o})$ a retract of its pullback Hom for every M, even if M is unbounded. Apply (2.6) with $M = h^* \Phi(\Delta_w)$. □

Lemma 2.27. *For pairwise commuting simple roots $D = \{s_1, \dots, s_d\}$, one has $C_{s_1} * \cdots * C_{s_d} \simeq C_D$ in every order. Replacing a factor by its counit gives the corresponding parabolic restriction.*

Proof. In a parabolic of type D, its Borel variety is a product of projective lines, for every isogeny type. This follows after a central cover of its Levi, which does not change its flag variety. Given endpoints, the intermediate flag takes its already changed coordinates from the output and the rest from the input. These are algebraic parabolic-equivariant maps identifying the iterated ordinary flag correspondence with the D-correspondence. The smooth flag maps compute its derived products as well. Apply the derived loop functor to obtain the structure-sheaf kernels and their circle actions. A counit sets one coordinate equal at its endpoints, exactly the indicated restriction, compatibly with the total counit. □

2.2.4. *The real coefficient across a fixed root divisor.*

Proposition 2.28. *Let $s = s_\alpha$ be fixed by θ, let $c = \alpha^\vee \in M = \Lambda^\theta$, and let $A = \mathbb{C}[M]$ act by precomposition. Then $H^1 \operatorname{RHom}(P_C, P_C \star \Delta_s) = A/(e^c - 1)$, its H^0 is free of rank one over A, and all other cohomology vanishes. The odd coefficient survives periodization.*

Proof. Put $M_0 = \ker \alpha$, $M_p = \alpha^{-1}(2\mathbb{Z})$, $A_0 = \mathbb{C}[M_0]$, and $\overline{A} = A/(e^c - 1)$. Since $\alpha(c) = 2$, $M_p = M_0 \oplus \mathbb{Z}c$. If $\alpha(M) = 2\mathbb{Z}$, then $\overline{A} = A_0$. Otherwise choose ν with $\alpha(\nu) = 1$, write $w = e^\nu$, $b = e^{2\nu - c} \in A_0^\times$, and obtain $\overline{A} = A_0[w]/(w^2 - b)$, finite étale of rank two.

The compact pair-stabilizer image in a real rank-one fibre is the kernel of $\alpha : D_x \to \{\pm 1\}$. The norm loop μ ends at $\mu(-1)$ and interchanges the two sides exactly when $\alpha(\mu)$ is odd. Thus the chosen-side lattice is M_p. There are two upper orbits in the even case and one actual upper orbit otherwise. The upper lattice is $\Lambda^{s\theta} = M_0$, since the negative s-eigenline is that of c and θ fixes it. The lattice map from a chosen side to the upper orbit is $M_p \to M_0, \qquad \mu \mapsto (\mu + s\mu)/2 = \mu - \alpha(\mu)c/2$, surjective with kernel $\mathbb{Z}c$. Each side germ is a contractible half-interval. The link functor $L = i^* R j_*$ therefore takes (V_+, V_-) to $V_+ \oplus V_-$ in the first case, and V to $\overline{A} \otimes_{A_0} V$ in the second, with A acting through $\overline{A}$. Finite-cover induction and coinduction agree here.

Let $S = P_C \star \Delta_s$. Its closed restriction is free of rank one in degree zero by Proposition 2.12. The real-ascent integral of Lemma 2.5, with fibre dimension $r + 1$ canceled by its shift, gives upper coefficients (A_0, A_0) or A_0^2, respectively. Its gluing triple is therefore (A, V, a) with $a : A \to \overline{A}^2$. Recollement computes maps between such triples as the fibre of the difference of the closed and open Hom maps into the link Hom. Because Δ_s is invertible, $\operatorname{Ext}^1(S, S) = \operatorname{Ext}^1(P_C, P_C) = 0$. Thus $A \oplus \operatorname{End}(V) \to \overline{A}^2$, $(r, g) \mapsto a(r) - L(g)a(1)$, is surjective. For two upper orbits this makes both entries units. In the one-orbit case, after any closed-point specialization of A_0, the two sheet vectors must be independent: otherwise the common matrix action and two independent sheet scalars span at most three dimensions in a four-dimensional target. Their determinant is therefore a unit globally. Change upper bases to make $a(r) = (\bar{r}, \bar{r})$ or $(\bar{r}, w\bar{r})$, respectively.

Maps from the closed triple $(A, 0, 0)$ give $[A \xrightarrow{a} \overline{A}^2]$ in degrees $0, 1$. Its kernel is $(e^c - 1)A$, and its cokernel is $\overline{A}$, by $y - x$ or $y - wx$. Tensoring this finite complex with P proves the last claim. □

2.2.5. *Two adjunctions and the rank-one algebra.* Continue with fixed s. Put $q = e^\beta(y)$, where β is the tangent root on $\mathfrak{p}_s/\mathfrak{b}$, and let $\varphi = \varphi_s^*$ be the signed involution of A; then $\varphi(q) = q^{-1}$. Use $T = C_s$ also for its ordinary constructible preimage, $\varepsilon : T \to \mathbf{1}$ for the left counit, $\delta : T \to T * T$ for its comultiplication, and $\eta_R : \mathbf{1} \to T$ for the oriented right unit. Set $\eta = -\eta_R$.

Lemma 2.29. *There are ordinary triangles*

$$\Delta_s \longrightarrow T \xrightarrow{\varepsilon} \mathbf{1} \longrightarrow, \qquad \mathbf{1} \xrightarrow{\eta} T \longrightarrow \nabla_s \longrightarrow .$$

The composite $\varepsilon\eta$ is $e^\beta - 1 \in R$, and on P_C, Q it is $q^2 - 1$. The map $\delta\eta$ is a coevaluation for a self-adjunction of T with the correspondingly signed evaluation.

Proof. Lemma 2.25 gives the first triangle for an adjoint group, also after a central torus factor. For general isogeny type pull back its coherent fibre along the cover of Lemma 2.26. Its constructible image forgets central descent and is the primed standard. On each Bruhat stratum this is finite-cover direct image, which detects zero restriction. Hence the original fibre is extension by zero of its simple open coefficient. Finite faithful flatness makes that coefficient an ordinary rank-one projective over R, and the Laurent polynomial ring has trivial Picard group. A frame therefore identifies it with Δ_s before periodization. Write $L = p^*$, $V = p_*$ and use $p^! \simeq L$. Then $\delta = L\eta_L V$ and η_R is the unit of $V \dashv L$. The composite $\delta\eta_R$ is the unit of $LV \dashv LV$, and the right mate of ε is η_R. Taking right adjoints of the first triangle gives the second after the matched sign change. On the regular boundary, $\varepsilon\eta_R = 1 - e^\beta$ by its sheet Jacobian. Equality there extends in $\operatorname{End}(\mathbf{1}) = R$. Finally $N^* e^\beta = q^2$. □

Lemma 2.30. *The ordinary map $(\delta, \eta_R * \mathrm{id}_T) : T \oplus T \xrightarrow{\sim} T * T$ is an isomorphism.*

Proof. On an SL_2 central cover, let z be the eigenvalue on the marked output line and z_m that on the middle line. The direct image of the middle-flag structure sheaf is the locally free quadratic algebra with basis $1, z_m$ and equation $z_m^2 - \mathrm{tr}(g) z_m + 1 = 0$. Away from the two central elements a cyclic vector identifies invariant lines with roots of this equation, even at a repeated root. The determinant of the resulting map of rank-two bundles is invertible there. A divisorial zero locus cannot be confined to the excluded codimension-three locus, so it is an isomorphism everywhere. The same holds with a central torus factor. Put $k = z_m/z$, $x = z^{-2}$. Its two sheet values are $1, x$. Derived base change inserts this algebra on the two-flag incidence. The first displayed map inserts 1, the second inserts $k - x$: generically it is $1 - x$ on the output sheet and zero on the other. This equality extends by Cartan flatness from Lemma 2.25. The pair $1, k - x$ is a basis. Every rank-one Levi admits the central cover just used, with unchanged flag variety and adjunction orientation; faithful flatness detects the isomorphism downstairs. Transport through the ordinary Φ. □

Put $m = P_C$, $S = \Delta_s \cdot m$, and $E_C = \mathrm{Hom}(T \cdot m, m)$. Its product is $FG = F \circ T(G) \circ \delta_m$, its unit is $e = \varepsilon_m$, and $a \mapsto ae$ embeds A. Left and right scalars are postcomposition and precomposition by $T(a)$. Write A_φ for the usual left module with right action twisted by φ.

Lemma 2.31. *Set $h = 1 + q$. A choice of frame of* $\mathrm{Hom}(S, m)$ *gives coordinates $\ell(F) = F\eta_m/(q-1)$, $\quad b(F) = F|_S$, $\qquad E_C = \{(v, w) \in A \oplus A_\varphi : v \equiv w \pmod h\}$. The unit is $(h, 0)$.*

Proof. The first triangle and Proposition 2.12 give $0 \to Ae \to E_C \xrightarrow{b} A_\varphi \to 0$ and ordinary degree-zero concentration. Self-adjunction gives the same for $\mathrm{RHom}(m, Tm)$, with $G \mapsto \widehat{G} = T(G)\delta_m\eta_m$. The counit identity gives $\varepsilon_m\widehat{G} = G\eta_m$. Applying $\mathrm{Hom}(m, -)$ and Proposition 2.28 shows that the image of $F \mapsto F\eta_m$ is exactly $(q-1)A$. Thus ℓ is a surjective bimodule map, $\ell(e) = h$, and (ℓ, b) is injective. Initially its image has form $v \equiv \Omega w \pmod h$ for a unit $\Omega \in (A/h)^\times$.

This unit lifts to A. If $c = \alpha^\vee$ is primitive in M, extend it to a basis; $A/(1+q)$ is the Laurent ring on the remaining characters, so its units lift. Otherwise $c = 2v_0$ and $\alpha(v_0) = 1$; the root datum splits the rank-one Levi as split PGL_2 times a central real torus. Its determinant-sign character commutes with every simple correspondence, because the partial flag is fixed. On $A = A_0[z^{\pm 1}]$, $z = e^{v_0}$, it acts by $\sigma(z) = -z$ and fixes A_0. Identify its twist of the regular module by this involution. On a frame of A_φ its action is $\kappa\sigma$ with $\kappa\sigma(\kappa) = 1$. A Laurent unit has a fixed exponent under σ, so $\kappa = \pm 1$. Equivariance of ℓ, b makes the two residues of Ω at $z = i, -i$ equal or opposite. In the first case their common unit in A_0 lifts them; in the second multiply the first value by z/i. Rescale the frame by this lifted unit to get the displayed congruence. □

Lemma 2.32. *The pairing $(F, G) \mapsto \ell(FG)$ is perfect between the underlying left and right A-modules of E_C.*

Proof. Closed adjunction and Propositions 2.12, 2.28 identify opposite standard coefficients with dual free modules, one embedded by the ideal $(q-1)$. Choose generators $f : S \to m$, $g : m \to S$ with $fg = q - 1$, lift f to $\widetilde{f} : Tm \to m$, and choose $\widetilde{g} : m \to Tm$ with $\varepsilon_m\widetilde{g} = q - 1$. Then $(\varepsilon_m, \widetilde{f})$ and $(ig, \widetilde{g})$ are the two bases in the exact sequences above. Their composition matrix is

$$\begin{pmatrix} 0 & q-1 \\ q-1 & z_0 \end{pmatrix}.$$

Under $G \mapsto \widehat{G}$, its entries are $(FG)\eta_m$; hence z_0 is also divisible by $q - 1$. Dividing gives determinant -1. □

Proposition 2.33. *There is a basis e, d of E_C on either side with*

$$d^2 = d, \qquad da = \varphi(a)d + \frac{a - \varphi(a)}{1+q}, \qquad \ell(e) = 1 + q, \quad \ell(d) = 1.$$

Proof. Take $d = (1, 1)$ in the preceding coordinates. The involution is identity modulo h: $s\mu = \mu - \alpha(\mu)c$ and the sign is $(-1)^{\alpha(\mu)}$. The coordinates give both bases and the scalar relation. Write $\partial(a) = (a - \varphi(a))/h$. The identity $1/h + 1/\varphi(h) = 1$ implies $\partial^2 = \partial$ and $\partial\varphi + \varphi\partial = \varphi - 1$. Applying the scalar relation twice shows $d^2 - d$ commutes with A. Its centralizer is Ae, because φ is nontrivial and A is a domain. Thus $d^2 - d = pe$ with $\varphi(p) = p$. For $j = e - hd$ one gets $ja = \varphi(a)j$ and $j^2 = (1 + h\varphi(h)p)e$. The divided pairing in the basis e, d is $\left(\begin{smallmatrix} h & 1 \\ 1 & 1+hp \end{smallmatrix}\right)$, whose determinant is $q(1 + h\varphi(h)p)$. Perfection makes the parenthesis an invariant Laurent unit. Its exponent is fixed by s, so it belongs to $A_0^\times$, and its reduction modulo h is 1. The map $A_0 \to A/h$ is injective, since $\ker\alpha$ meets $\mathbb{Z}c$ trivially. Therefore this unit is 1 and $p = 0$. □

Proposition 2.34. *For $E_Q = \mathrm{Hom}(T \cdot Q, Q)$, there is a unital algebra isomorphism $E_C \simeq E_Q$ preserving both copies of A and $F \mapsto F\eta/(q-1)$.*

Proof. By Proposition 2.22, $E_Q = \Gamma(Y_s, \operatorname{End} M_s)$, identified by the trace pairing with $\Gamma(Y_s, M_s \otimes M_s)$. First consider ordinary multiplication of functions on the latter. On the $\mathrm{SL}_2 \times Z(L_s)^0$ central cover, θ is trivial on the derived factor. The two eigenvalues satisfy $z_i + z_i^{-1} = \operatorname{tr}(g)$; the invariant free basis is $1, z_1, z_2, z_1 z_2$. With $k = z_2/z_1$ and $x = z_1^{-2} = -q$, finite flat descent gives $\Gamma(Y_s, M_s \otimes M_s) = A[k]/((k-1)(k-x))$. The deck group changes both eigenvalues by the same sign and fixes k, x, even with its action on the central torus. The second endpoint acts by $a, \varphi(a)$ on the two sheets, and multiplication evaluates at $k = 1$.

Return to the composition product. The trace denominators are $1-x$, $1-x^{-1}$ by Proposition 2.23, so $\delta_s(1) = 1$. The identity operator has tensor values $(1-x, 0)$, hence is $e = k - x$. The tensor $d_Q = 1 \otimes 1$ is the rank-one operator $m' \mapsto \delta_s(m')1$, so $d_Q^2 = d_Q$. The two elements e, d_Q are a basis. Sheetwise, $d_Q a - \varphi(a) d_Q = ((a - \varphi(a))/h)e$. These identities extend from the dense regular open because the tensor module is free over A. Finally $c_s = 1 + x = 1 - q$ and $\eta = -\eta_R$, so the divided functional is evaluation at $k = 1$, taking e, d_Q to $h, 1$. Proposition 2.33 now gives the desired identification. □

Theorem 2.35. *For every quasi-split H_0 and every simple $s = \theta(s)$, the closed marked equivalence extends over the entire Cartan to an $\mathcal{H}_{\{s\}}$-module equivalence on the subcategories generated by P_C, Q, with all convolution coherences.*

Proof. In the ordinary simple Hecke category, let $\mathcal{B}$ be the additive idempotent closure of $\mathbf{1}, T$. Lemma 2.30 makes it monoidal, and the two triangles generate the compact stable category. Standard and costandard filtrations put its Hom complexes in degree zero, so its stable envelope is $\operatorname{Kar}(K^b\mathcal{B})$. Maps $T \to \mathbf{1}$ are generated over R by ε, maps $\mathbf{1} \to T$ by η. An endomorphism of T restricts to a scalar on Δ_s; after subtracting that outer Cartan scalar it factors through ε, hence is a multiple of $\eta\varepsilon$. These follow directly by adjunction from the triangles and $\operatorname{RHom}(\Delta_s, \mathbf{1}) = \operatorname{RHom}(\mathbf{1}, \nabla_s) = 0$.

For $m = P_C, Q$, the opposite coefficient functor has values A, E_m on these generators. On the two maps its values are $a \mapsto ae$ and $F \mapsto (q-1)\ell(F)$, with outer Cartan actions N^*. Proposition 2.34 identifies all these maps. For the pairing on $T * T$, the two components of Lemma 2.30 give FG and $(F\eta_{R,m})G$, both preserved. Pairings involving the unit are the two scalar actions. All iterated products lie in the ordinary additive category, where these coefficient diagrams have discrete mapping spaces and higher coherences are determined by ordinary composition. Extend exactly in each variable to finite cones and retracts. The small equivariant diagrams on the generators are degree zero and take invariants by tensoring with $\mathbb{C}[u]$; localization tensors them with P. The represented internal algebras are thus identified with the identity on $A \otimes P$. Lemma 2.15 proves the marked module comparison. □

2.3. **The comparison on the regular semisimple Cartan.** Let $\mathcal{I} = \{w\theta w^{-1} : w \in W\}$ and $A_\delta^U = \mathcal{O}(V_\delta)$, with the norm marking of Lemma 2.2.

2.3.1. *Clean real strata.*

Lemma 2.36. *The surviving real flag orbits after localization at f are exactly those attached to a maximally split Cartan. They are indexed by $\mathcal{I}$, and each is reached from C by complex simple ascents.*

Proof. If $(T^\delta)^0$ meets U, no root hyperplane contains its real fixed tangent space. A finite union of proper hyperplanes cannot cover that space, so it contains a regular vector. Its chamber is δ-stable, giving a positive system preserved by the real Cartan structure. The Cartan thus lies in a real Borel and is maximally split. Conversely a maximally split Cartan of a quasi-split group lies in a real Borel. Its involution preserves a chamber; averaging a vector in that chamber gives a fixed regular vector and hence a norm point in U. These Cartans are real-conjugate and their real Weyl group is W^θ [IY25, Lemmas 4.14–4.15]. Their flag orbits are $W^\theta\backslash W$, identified with $\mathcal{I}$ by $w \mapsto w^{-1}\theta w$. The closed-orbit criterion [IY25, Lemma 6.9] supplies a complex descent for every nonclosed such orbit: the alternative noncompact imaginary root cannot occur for a maximally split Cartan. Repeated descent terminates at C. For a nonsurviving orbit some root is identically one on its connected norm image; that factor of f acts by zero on the entire coefficient ring, including its finite norm cover. □

In this subsection normalize real standards by $p_O = \lfloor (\dim_{\mathbb{R}} O + \dim(T_H^c/D_O^c))/2 \rfloor$: $\Delta_O(M) = j_{O!}M[p_O]$ and $\nabla_O(M) = Rj_{O*}M[p_O]$. This is the geometric perversity of [IY25, Section 3.4]; a complex ascent increases it by one.

Lemma 2.37. *For a complex ascent $O \to O'$, the coefficient transport T_s of Lemma 2.6 gives*

$$\Delta_O(M) \star \Delta_s^{\mathcal{H}} \simeq \Delta_{O'}(T_s M), \qquad \nabla_O(M) \star \nabla_s^{\mathcal{H}} \simeq \nabla_{O'}(T_s M).$$

It preserves perfect coefficients and remains an equivalence after localization at f.

Proof. For standards, composition of shriek pushforwards and the proper enhanced projection gives extension by zero of the compact-torus integral in Lemma 2.6. Its remaining shift is exactly $p_{O'} - p_O = 1$. For costandards, locally choose a section g_x of the transitive H-action. The coordinate change $(x, z) \mapsto (x, g_x z)$ makes the Hecke kernel independent of the input coordinate. On these coordinates the external product map $Rj_*F \boxtimes_{\mathbb{C}} Rk_*E \longrightarrow R(j \times k)_*(F \boxtimes_{\mathbb{C}} E)$ is an isomorphism, computed by the cellular cross product of two finite local links. This argument works for infinite-dimensional stalks because the links have finitely many cells, and retains the two independent coefficient-ring actions. Equivariance glues these local maps. Compose star pushforwards and use equality of star and shriek on the compact torus fibre. The same coefficient transport, orientation, and shift result. The finite link models extend the formula continuously from perfect coefficients to arbitrary complexes. Finally f is Weyl-invariant. □

Lemma 2.38. *The canonical $\Delta_s^{\mathcal{H}} \to \nabla_s^{\mathcal{H}}$ is an isomorphism after inverting f.*

Proof. The punctured normal complex line to the diagonal in the simple Schubert closure has monodromy $e^{\alpha_s^\vee}$. Its cochains are $[R \xrightarrow{e^{\alpha_s^\vee}-1} R]$; they compute the cone up to the common normalization [Tay25, Proposition 4.2]. The indicated factor is inverted by f. The full coroot exponent is retained when the coroot is nonprimitive. □

Proposition 2.39. *Every surviving real standard is clean on U, and restriction gives*

$$\mathcal{A} \otimes_{\mathrm{Mod}(R)} \mathrm{Mod}(R[1/f]) \simeq \prod_{\delta \in \mathcal{I}} \mathrm{Mod}(A_\delta^U).$$

It identifies compact objects with the product of perfect module categories and retains arbitrary stalk complexes and the full norm markings.

Proof. On C, standard and costandard extensions coincide for every coefficient complex. Use the complex ascents of Lemma 2.36 and the two formulas of Lemma 2.37; their Hecke kernels coincide after localization by Lemma 2.38. Transporting any coefficient back to C proves that its standard and costandard extensions are isomorphic. The standard has zero ordinary restrictions off its stratum. The canonical standard-to-costandard map is therefore an isomorphism as well, since its cone restricts to zero on every stratum. For distinct strata, full adjunction gives $\mathrm{RHom}(j_{O!}M, Rj_{O'*}N) = \mathrm{RHom}(j_{O'}^* j_{O!}M, N) = 0$. First apply this to compact coefficients and then localize. Cleanness gives full faithfulness and orthogonality in all degrees, even if their norm images meet in U. A finite orbit filtration reconstructs every compact object; nonsurviving terms vanish and the remaining triangles split by this orthogonality. Restriction and clean extension are inverse on compacts. Ind-completion gives the asserted arbitrary-stalk equivalence. □

2.3.2. *Parameter components and their actual circle invariants.*

Lemma 2.40. *Let $\mathcal{X} = [Z/J] \to \operatorname{Spec} R$ be a noetherian derived quotient over $\mathbb{C}$, with J reductive and Z quasi-projective, locally almost of finite presentation, and admitting a linearized ample line bundle. Suppose its R-linear circle action preserves bounded coherence and the coherent t-structure. For $j : \mathcal{X}_f \to \mathcal{X}$,*

$$\mathrm{Idem}\big(\mathrm{Coh}(\mathcal{X})^{S^1}[f^{-1}]\big) \simeq \mathrm{Coh}(\mathcal{X}_f)^{S^1}, \qquad \mathsf{T}(\mathcal{X})|_{D(f)} \simeq \mathsf{T}(\mathcal{X}_f).$$

Proof. For bounded coherent M, N, ordinary Hom localizes at f. On an affine chart, truncate a bounded-above finite-cell resolution of M far enough to compute any fixed degree into N; the finite calculation commutes with flat localization. A finite affine Čech complex on the separated scheme Z, followed by exact reductive invariants, gives the same statement on the quotient. The mapping complex is bounded below. Its homotopy-invariant tower filtered by the cells of $BS^1 = \mathbb{C}P^\infty$ has successive layers shifted by $-2k$. In any fixed degree, and one degree below it, only finitely many cells contribute, uniformly after localization. The finite stages commute with flat base change, so equivariant Hom localizes as well, preserving composition.

For essential surjectivity use a finite equivariant Postnikov tower. On the ordinary coherent heart the circle action is encoded by a natural full-turn automorphism b. A heart sheaf admits an equivariant structure precisely when $b = 1$ on it, and that structure is unique, because the relevant isomorphism spaces are discrete. An equivariant coherent sheaf on the open extends as follows. Its pushforward on the classical truncation is the union of coherent J-equivariant subsheaves: twist by the ample bundle, choose finite generators, and take their finite-dimensional J-stable spans. Choose one restricting to the original sheaf. Its quotient by the image of $b - 1$ still has the same restriction and has full turn identity, by naturality. It gives an extension in the invariant heart. Full faithfulness lifts the extension maps of each finite Postnikov tower to the localized category; cones and idempotent completion prove the first assertion. The commuting central operators f, u may then be inverted in either order, followed by Ind. □

The parameter and Steinberg incidence atlases satisfy these hypotheses, with linearized ample bundles from their flag factors. Boundary rotation acts by transport and preserves their coherent hearts. The Weyl-invariant function f defines the same open at both flags of every correspondence.

Lemma 2.41. *There is an open-and-closed decomposition*

$$\mathcal{P}_E^U \simeq \coprod_{\delta\in\mathcal{I}} \{[s\tau_\delta] \in [T\tau_\delta/T] : s\delta(s)z \in U\}.$$

Here τ_δ acts by δ and squares to z. If $\delta = w\theta w^{-1}$ and $r_\delta = w(r_\rho)$, translation $s \mapsto sr_\delta$ identifies its marking with $N_\delta : V_\delta \to U$ and preserves the original circle.

Proof. The Grothendieck–Springer map is finite étale over regular semisimple holonomy: its root normal derivatives are $1-\alpha(t)$, all invertible on U. Its derived square pullback is classical there. The map $N_B \times U \to B \times_T U$, $(n,t) \mapsto ntn^{-1}$, is an isomorphism. Successive root-height quotients solve its conjugation equation linearly with coefficients $1-\alpha(t)$. Eliminating these coordinates gives $\mathcal{P}_E^U = [\{a \in E_{\mathrm{odd}} : a^2 \in U \subset T\}/T]$. Such a normalizes $Z_G(a^2)^0 = T$, even if the full centralizer is disconnected. Its component is an involution $\delta \in W\theta$. By Lemma B.3, it preserves a Borel. That Borel contains its regular square and thus the torus T, so δ preserves a chamber. These components are exactly $w\theta w^{-1}$. Conversely their invariant chambers give fixed regular vectors, so each norm image meets U. Choose $\tau_\delta = n_w \tau n_w^{-1}$; its square is the central z. The normalizer components are disjoint and T-conjugation preserves them. Finally $\delta(r_\delta) = r_\delta$ and $r_\delta^2 = z$. Translation therefore changes the square to the norm. It is induced by the actual torus-extension isomorphism $\tau_\delta \mapsto r_\delta \tau_{\delta,0}$, where the last lift has order two. Thus every cyclic transport subdivision is preserved with degree-one circle map. □

Lemma 2.42. *Let $D_\delta = T^\delta$ and $\Pi_\delta = \pi_0(D_\delta)$. Then*

$$\mathsf{T}(\mathcal{P}_E^U) \simeq \prod_{\delta\in\mathcal{I}} \prod_{\epsilon\in\Pi_\delta^\vee} \mathrm{Mod}(A_\delta^U \otimes P).$$

The norm and u are retained. The structure sheaf is the free module in the trivial finite-character factor.

Proof. The twisted conjugation map has kernel D_δ and image $(1-\delta)T$. The quotient $T \to Y_\delta$ admits a section because the corresponding quotient character lattice is free. Thus its stack is $Y_\delta \times BD_\delta$, and its open is $V_\delta \times BD_\delta$. The boundary is $N_\delta(y) \in D_\delta^0$. A coherent object has finitely many D_δ-character summands. On character χ, circle monodromy is $m_\chi(y) = \chi(N_\delta(y))$. If $\chi|_{D_\delta^0} \neq 1$, this function is nonconstant because the norm is an isogeny onto D_δ^0. Its unit fibre is smooth or empty.

For $A' = A_\delta^U$, the small invariants of a multiplier m are coherent modules over $B_m = A'[\varepsilon], \quad |\varepsilon| = -1, \quad d\varepsilon = 1-m$, with u its degree-two derived-zero-locus operator [BCHN25, Example A.13]. For nonconstant m, this is the smooth algebra $A'/(1-m)$, so its coherent modules are generated thickly by its free module. The latter has zero degree-two self-Ext, and u vanishes there. Tate localization kills the whole summand. For χ trivial on D_δ^0, $m = 1$ and the character factors through Π_δ. Finite Postnikov towers and regularity of A' show that bounded coherent B_1-modules are thickly generated by its augmentation A'. Its Koszul resolution gives $\mathrm{RHom}_{B_1}(A', A') = A'[u]$, with composition as multiplication. Localize this same u and Ind-complete. Only the finitely many characters of Π_δ remain. □

Proposition 2.43. *The localized Hecke hull is*

$$\mathcal{S}_Q|_U \simeq \prod_{\delta\in\mathcal{I}} \mathrm{Mod}(A_\delta^U \otimes P),$$

included in the previous product through the trivial finite-character sector in each component, with unchanged markings.

Proof. Fix output holonomy $t \in U$. Each input Borel contains the same torus, so $Z_G^U = \coprod_{w\in W}[U/T]$. On component w its Cartan maps are $t, w^{-1}t$, and boundary gauge is t. The preceding character calculation kills every nonzero T-character; each component's Tate category is generated by its structure kernel K_w. In the parameter action, input data and w determine a unique output flag. Its correspondence is an isomorphism graph, changing the marked involution from δ to $w\delta w^{-1}$ while retaining the odd element, its square, and transport. It takes structure sheaf to structure sheaf with trivial stabilizer character.

The stable Borel open is precisely the θ component, so $Q|_U$ is its structure sheaf. The graphs reach every member of $\mathcal{I}$, giving all the displayed free generators. Each K_w is the restriction of a global small invariant object: take the reduced schematic closure of its classical component and its structure sheaf, with canonical

boundary lift. Conversely every localized compact kernel is a finite construction from the K_w, so its translate of Q stays in these trivial sectors. Localization of a full hull generated by compact objects is the hull of their localized images. This proves the claim, with the actual maps $t, w^{-1}t$ and norm on each component. □

Theorem 2.44. *There is an $R[1/f] \otimes P$-linear equivalence $\mathcal{A}(H_0)_{\mathrm{per}}|_U \simeq \mathcal{S}_Q|_U$ extending the closed comparison and preserving u.*

Proof. Propositions 2.39 and 2.43 give the same product of module categories. On each factor the automorphic map is $1+\delta$ on lattices and the spectral map is the square in coordinate $[sr_\delta]$, both exactly N_δ^*. For $\delta = \theta$ take $w = 1$; both closed objects become the regular $A[1/N^*f] \otimes P$-module with coordinate $[tr_\rho]$. Thus the closed restriction is the prescribed one. □

2.3.3. *Return lines and their possible multiplication cocycle.* Put $B_U = R[1/f]$, $A_U = A[1/N^*f]$, and $\Gamma = W^\theta$. On graphs use $K_v * K_w = K_{vw}$, with the isomorphism defined by the unique intermediate flag. Let $C_{\mathrm{all}} = \mathcal{O}_{Z_G}$ and $T_{\mathrm{all}} = \Phi^{-1}(C_{\mathrm{all}})$. For $m = P_C, Q$, set $E(m) = H^0 \operatorname{RHom}(T_{\mathrm{all}} \cdot m, m)$, using C_{all} on the second side. Comultiplication for forgetting the entire flag gives its product and counit gives its unit e.

Lemma 2.45. *The real coefficient of T_{all} is concentrated in degree zero and free of rank $|\Gamma|$ as a left A-module. On U it is $\bigoplus_{w\in\Gamma} L_w$, where $L_w = \operatorname{Hom}(\Phi^{-1}(K_w) \cdot P_C, P_C)|_U$, $\quad ba = \varphi_w^*(a)b$. Composition identifies $L_v \otimes_{A_U} L_w$ with L_{vw}, using the graph maps just specified. The spectral return lines have canonical multiplicative generators.*

Proof. For adjoint H, $T_{\mathrm{all}} = \Xi$ by the Whittaker adjunction and has one standard at each finite Weyl element [Tay25, Proposition 3.13]. Tensor with the central torus unit, then descend by the finite cover of Lemma 2.29 for general H. Stratumwise finite-cover direct image detects support, shift, and rank-one regular coefficients. Open–closed filtration therefore supplies the same single standard per finite position downstairs. Proposition 2.12 gives a free ordinary line precisely for θ-fixed positions. The successive coefficient sequences split as left modules because their quotients are free, proving the first assertion.

On U, $C_{\mathrm{all}} = \bigoplus_w K_w$. The simple counit fibre is its simple graph before circle localization; its ordinary identification in Lemma 2.29 does not require any Galois hypothesis. Length-additive products identify all graphs with their localized standards. Transposition to output-first order and Hom into the regular module give the displayed semilinear right action. In the clean product decomposition an invertible graph permutes the indecomposable factors. Each factor has a domain as coefficient ring and no nontrivial product decomposition. A return graph sends the closed generator to its transported free line, so any frame of L_w is an isomorphism to P_C. Products of such isomorphisms are again generators of the appropriate line. The graph orbit has the full $[W : \Gamma]$ factors. On the spectral side change of flag takes the closed structure sheaf to itself; its identity section is the canonical frame and its products are graph composition. □

Choose frames $b_1 = e$, $b_w \in L_w$. The equation $b_v b_w = c(v,w) b_{vw}$ defines a class

$$[c] \in H^2(\Gamma, A_U^\times) \tag{2.9}$$

for the signed action φ^*. Associativity is its cocycle identity, and changing frames is a coboundary.

Proposition 2.46. *The closed comparison extends over U as a marked finite Hecke-module equivalence if and only if the class (2.9) vanishes.*

Proof. The localized graph category has no maps between distinct graphs and ordinary endomorphism ring B_U on each structure graph. These graphs generate its Tate category. A marked comparison must supply return frames with $b_v b_w = b_{vw}$. This is necessary even periodically, since degree-zero automorphisms of the periodic free line are exactly $A_U^\times$. Conversely choose multiplicative frames. They compare the return functors on the closed free object and hence its entire module category. Choose a graph K_t to every other factor and transport the closed equivalence through it. A graph between factors $t\Gamma, t'\Gamma$ uses the return $t'^{-1}gt \in \Gamma$. The product identity gives composition compatibility. Cartan actions follow their actual graph maps. All these diagrams are ordinary and degree zero on free generators, so their coherences are unique there; exact extension to coefficients, cones, and retracts, followed by tensoring with P, gives all required coherences. □

Lemma 2.47. *For the fixed simple roots S_0, there are integral units $j_s \in E(P_C)$ satisfying $j_s^2 = e$, $j_s a = \varphi_s^*(a) j_s$, and all Coxeter relations of W_{S_0}.*

Proof. Factoring $X_b \to X_s \to Y_g$ gives a counit-induced map of coalgebras $C_{\mathrm{all}} \to C_s$; triangle identities show that it respects comultiplication as well. Thus Proposition 2.33 embeds its rank-one algebra into $E(P_C)$ and supplies $j_s = e - (1+q_s)d_s$ with the stated square and semilinearity. Pass to the fraction field K by the central

ring $N^*(R^W)$, over which A is finite. The return graphs have distinct automorphisms of K: a fixed strictly dominant vector has trivial Weyl stabilizer, so Γ acts faithfully on the fixed subspace. Their generic centralizer is Ke. For distinct s,t, let $v=(j_sj_t)^{m_{st}}=\lambda e$ generically. Both v,v^{-1} act on the finite free left A-module, so their characteristic polynomials make λ and λ^{-1} integral. Normality gives $\lambda\in A^\times$. Conjugation by either simple unit inverts it; its Laurent exponent belongs to both distinct negative eigenlines and hence is zero. Thus $\lambda=\pm1$. At $y_0=[r_\rho]$, all signed returns fix the point and $q_s(y_0)=q_t(y_0)=-1$. The simple units are therefore identity modulo its left coefficient ideal. A telescoping expansion of $v-e$ remains in that ideal by semilinearity. The unit survives the quotient: otherwise right multiplication would force the positive-rank free module to equal its maximal-ideal multiple. Hence the sign is positive. □

2.3.4. *Ordinary universal microlocalization and normalization.* Choose a regular nilpotent covector $\xi\in i\mathfrak{h}_0^*$, whose Springer fibre is a unique real Borel $x\in C$ [IY25, Lemma 8.2]. Its moment fibre $\Xi_\xi\subset T^*\widetilde{X}$ is the compact Cartan torsor over x. Fix a base point, identifying it with T_H^c, and the corresponding point in the torsor $\Xi_{-\xi}$. For compact universal real F, define $V_{\mathbb{R}}(F)$ by microlocalizing along $\widetilde{C}$, restricting to Ξ_ξ, and taking its monodromy module. Normalize $V_{\mathbb{R}}(P_C)=A$ in degree zero. For a compact Hecke kernel define $V_\mu(K)$ by microlocalization along the enhanced inverse image of the ordinary diagonal, restriction to $\Xi_{-\xi}\times\Xi_\xi$, and shift $[-r]$. Invert the first right Cartan action to make it a left action: $r_1vr_2=v(\iota(r_1)\otimes r_2)$, $\iota(e^\lambda)=e^{-\lambda}$. Then $V_\mu(\mathbf{1})=R$.

Lemma 2.48. *For compact universal objects, each stalk, link and Morse complex used to define V_μ or $V_{\mathbb{R}}$ is computed by a finite complex of perfect coefficient modules. These models commute with derived scalar extension.*

Proof. Each compact object has finite stratification and perfect stratum coefficients. A small stratified Morse pair and its links admit finite compatible triangulations. Their cochains are bounded complexes of finitely many coefficient terms, with differentials finite sums of parallel transports and attachment maps. Universal local systems introduce no infinite product. The dimensions of these finite pairs and bounded coefficient amplitudes give one common truncation for any finite diagram of stalks, links, and Morse groups. Derived extension of scalars commutes with these models, including augmentation and its flat completion; nonflat specializations retain their Koszul terms. □

Lemma 2.49. *On compact universal objects, microlocalization sends convolution of two Hecke kernels, and the action of a Hecke kernel on a real sheaf, to derived tensor product over R. These identifications respect units and iterated convolution and give $V_{\mathbb{R}}(P_C)=A$ with the norm marking.*

Proof. We now retain these ordinary coefficients in the geometric proof of [IY25, Proposition 8.10]. For $G_F=\mathrm{pr}_1^*F\otimes K$ the two singular supports are noncharacteristic and lie in $(\Lambda_{\mathbb{R}}\times0)+\Lambda_{\mathcal{H}}$. The enhanced output projection is proper. A covector of this sum has moment values $(v-w,w)$. To lie over $0\times\Xi_\xi$ requires $v=w=\xi$. Regularity forces both ordinary flags to equal x. The remaining microlocal correspondence is precisely the second projection $\Xi_\xi\times\Xi_\xi\to\Xi_\xi$; this is [IY25, equation (8.13)]. The smooth map $\widetilde{X}^2\to\widetilde{X}\times[H\backslash\widetilde{X}^2]$ identifies the relevant sheaf with the pullback of the external product of the two microlocalizations. Microlocal proper image, smooth pullback, and external product give the claimed comparison; their compatibilities are [IY25, Proposition 8.11 and Remark 8.12], using [KS90, Propositions 4.3.4–4.3.5]. The finite cochain models of Lemma 2.48 justify them for universal stalks.

In an oriented basis of Λ, integration over the middle compact Cartan is the group-cochain Koszul complex. Its self-duality gives

$$R\Gamma(T_H^c,L_M\otimes L_N)[r]=(M\otimes_{\mathbb{C}}N)\otimes^{\mathbf{L}}_{R,\mathrm{diag}}\mathbb{C}.$$

The unit orientation fixes the top line. Inversion on the first Hecke action converts the right side to $M\otimes^{\mathbf{L}}_R N$; the Hecke normalization cancels the integration degree. This proves the module formula with the exact norm on P_C. Three enhanced flags prove the monoidal formula in the same way, with moments $(-\xi,0,\xi)$. The two parenthesizations are proper images and smooth pullbacks in one common flag-tuple diagram; base change and composition identify them with the same oriented Koszul integration. For a flag string of any finite length, the face maps compose the same correspondences and the degeneracies insert the diagonal. Integration over the ordered product of compact Cartans commutes with these maps. Hence the comparison identifies the full convolution bar diagrams, including all higher coherences. The diagonal is the regular-module evaluation and gives the unit. □

Lemma 2.50. *On compact universal objects there are natural, unital, associative identifications*

$$V_\mu(K*L)=V_\mu(K)\otimes^{\mathbf{L}}_R V_\mu(L),\qquad V_{\mathbb{R}}(F\star K)=V_{\mathbb{R}}(F)\otimes^{\mathbf{L}}_R V_\mu(K).$$

They use the original uncompleted group rings. The module $V_{\mathbb{R}}(P_C)$ has its original norm marking. Both $V_\mu(\Delta_s)$ and $V_\mu(\nabla_s)$ are ordinary free graph bimodules R_s, in degree zero.

Proof. Lemma 2.49 gives both tensor identities. It remains to compute the images of the two inverse kernels at a simple reflection.

The inverse identities for simple standards make their images inverse derived bimodules. Their two Cartan actions factor through the graph of s on the open cell and continue through extension by zero or direct image. An invertible graph bimodule over connected R is a line in one constant degree: it is perfect, and its derived endomorphism complex is R. Over any residue field the derived fibre must therefore be a one-dimensional graded vector space in one degree. Minimal free complexes lift this conclusion locally, and connectedness makes the degree constant. The line is free since $\mathrm{Pic}(R) = 0$.

To determine the degree, apply the Cartan augmentation to the model of Lemma 2.48. In the normal complex disc to the simple diagonal, the costandard is $Rj_*\mathbb{C}_{D^*}[1]$. The triangle $\mathbb{C}_D \to Rj_*\mathbb{C}_{D^*} \to \mathbb{C}_{\{0\}}[-1] \to$ shows that its microlocal group at the nonzero simple conormal is $\mathbb{C}$ in degree zero; the locally constant term has no such microstalk. The Cartan shift $[r]$ cancels the definition's $[-r]$. The other Cartan action is already fixed by the graph relation. This also agrees with the completed calculation of [IY25, Proposition 8.4 and Sections 8.3, 8.6]. The inverse simple image has degree zero as well. □

Write $V = V_\mu$ and use its bimodules on left R-modules after the input-output switch. For a boundary coalgebra T' put $E_V(T') = \mathrm{Hom}_R(V(T') \otimes_R R, R)$, with product from its actual comultiplication and the tensor constraint of V.

Lemma 2.51. *For every simple s there is a unit $j_s \in E_V(C_s)$ with $j_s^2 = e$, $j_s a = s(a) j_s$, reducing to e at the identity of the boundary torus.*

Proof. Put $h = 1 - e^\beta$ for the ordinary simple tangent root. Applying V to the two triangles of Lemma 2.29 shows that $E = E_V(C_s)$ is free of rank two, with e part of a basis. The functional $\lambda(F) = FV(\eta_R)$ is surjective because the second triangle has a free graph quotient, and $\lambda(e) = h$ by the Jacobian. Its kernel is a free rank-one module; choose a generator j.

On $D(h)$ the simple parabolic has two distinct refinements. Their identity and change-of-refinement graphs are disjoint and compose as actual changes of flags. The second graph is the counit fibre and squares to the identity. Its two Cartan actions differ by s, and the diagonal functional kills it. Thus generically $ja = s(a)j$, and torsion-freeness extends this to E. Its square is ve for $v \in R$ because the generic diagonal line is the centralizer and e is in a basis. Associativity gives $s(v) = v$. Since j is a frame on $D(h)$, v is a unit there.

For any prime component $\mathfrak{p}$ of $h = 0$, reflection acts trivially on $R/\mathfrak{p}$. The split surjection λ has kernel generated by $\bar{j}$, containing $\bar{e}$. Since e is unimodular, $\bar{e} = k\bar{j}$ for a unit k on that component. Semilinearity makes its coefficient ideal stable under multiplication by j, so $v|_{\mathfrak{p}} = k^{-2}$. No prime factor of h divides v. It is therefore a Laurent unit globally.

Write $v = ce^\gamma$ with $s\gamma = \gamma$, and let β_0 be the primitive root-line character. On the identity component of $h = 0$, v is a square, so $\gamma = 2\mu + d\beta_0$ for an integral character μ and integer d. Applying $1 + s$ gives $\gamma = \mu + s\mu$. Choose $a = \sqrt{c}\, e^\mu$; then $a\, s(a) = v$, and $a^{-1}j$ is an involution. Its residue at identity is a sign times e; choose the sign of $\sqrt{c}$ to make it e. □

Proposition 2.52. *Over U, the graph bimodules $V(\Phi^{-1}K_w)$ have frames t_w with $t_1 = 1$ and $t_v t_w = t_{vw}$ for the tensor constraint of V and the specified coherent graph maps.*

Proof. The standard filtration of T_{all}, evaluated by V, makes $E = E_V(C_{\mathrm{all}})$ a free ordinary left R-module of rank $|W|$, because every standard is a length-additive product of simple graph lines. The counit-induced coalgebra map $C_{\mathrm{all}} \to C_s$ sends the unit of Lemma 2.51 to a unit in E. Over U the algebra is the sum of the dual graph lines with the specified graph product, and their generic Cartan automorphisms are the distinct Weyl elements. Fractions are taken over the central ring R^W. For $s \neq t$, $(j_s j_t)^{m_{st}}$ is a generic scalar. The finite-free integrality and normality argument of Lemma 2.47 makes it a Laurent unit. Both reflections invert it, so its exponent lies in both distinct negative root lines and is zero. It is a sign. Every simple unit is identity modulo the augmentation ideal, preserved by every reflection, so the telescoping argument makes the sign positive. The units therefore define a group map from W. Its value j_w is a frame of the dual graph line, detected by its generic Cartan action and inverse. Let t_w be the unique frame with $j_w(t_w) = 1$. The product of the dual frames, defined by the actual coalgebra and graph maps, gives exactly $t_v t_w = t_{vw}$. □

Theorem 2.53. *For every connected quasi-split real reductive group, the closed comparison extends over the entire regular semisimple open to a marked finite Hecke-module equivalence*

$$\mathcal{A}(H_0)_{\mathrm{per}}|_U \simeq \mathcal{S}_Q|_U.$$

It preserves the boundary Cartan, the original period, and every convolution coherence.

Proof. The ordinary real microlocal functor sends P_C to A, and its endomorphism a to multiplication by a. After localization it sends that module to A_U. For $w \in \Gamma$, choose any isomorphism b_w in its return line. The module constraint of Lemma 2.50 gives $V(\Phi^{-1}K_w) \otimes_{B_U} A_U \xrightarrow{\sim} A_U$. Evaluate it on $t_w \otimes x$ to obtain an operator $D_w(x)$. The actual endpoint marking of Lemma 2.45 says $D_w(ax) = \varphi_w^*(a)D_w(x)$. Thus $D_w(1)$ is a unit. Rescale b_w by its inverse, so that $D_w(1) = 1$, taking $b_1 = e$.

The module constraint identifies composition at the intermediate flag with composition of these operators. Proposition 2.52 ensures that the operator for $b_v b_w$ is $D_v D_w$, with the very same graph product. It and D_{vw} both send 1 to 1. The two coefficients lie in one free return line, and rescaling its frame rescales this value by the same unit. Hence $b_v b_w = b_{vw}$. The cocycle (2.9) vanishes, and Proposition 2.46 extends the comparison to all factors with the stated data. □

2.4. **Singular parameters and finite extension criteria.**

2.4.1. *Full flag averaging and the Demazure algebra.* For a split real group, the coordinate $z = y/r_\rho$ turns the signed return action into the usual Weyl action. Put $S = A^{\varphi W}$ and $K = \mathrm{Frac}(A)$.

Theorem 2.54. *If the derived subgroup of the dual group G is simply connected, there are unital S-algebra isomorphisms*

$$E(P_C) \simeq \mathrm{End}_S(A) \simeq E(Q).$$

They preserve the two copies of A and take the simple operators to

$$d_s(a) = \frac{a - \varphi_s^*(a)}{1 - e^{\beta_s}(z)}.$$

Their Tate comparison tensors these ordinary algebras with the original P.

Proof. The real coefficient is finite free by Lemma 2.45. The parameter algebra is $\Gamma([E_{\mathrm{odd}}/G], \mathrm{End}\, M)$, where $M = \rho_* \mathcal{O}_{\mathcal{F}}$ is a vector bundle, by Proposition 2.22. Proposition B.10 gives finiteness over the actual Cartan, hence over output A. Trace duality identifies its left module with $\Gamma(\mathcal{F}, \rho^* M^\vee)$. The integral atlas $B\tau$ dominates $\mathrm{Spec}\, A$, so any nonzero coefficient acts injectively on this bundle and its sections. Thus it is torsion-free.

Restriction from full to simple flag correspondences puts $j_s = e - (1 + q_s)d_s$ in both algebras. The proof of Lemma 2.47 applies also to a finite torsion-free module: the determinant trick gives integrality of a scalar, and Nakayama on its full-support localization makes the unit survive at $[r_\rho]$. Thus the j_w multiply by the Weyl group law and span the generic graph lines. Both algebras embed in $K \rtimes W$ by $j_w \mapsto w$. Independence of field automorphisms and $[K : K^W] = |W|$ identify this with $\mathrm{End}_{K^W}(K)$. Since $q_s = -e^{\beta_s}(z)$, the simple elements are the displayed divided differences.

Let $\mathscr{D}$ be generated by these and multiplication by A. The simply connected derived hypothesis gives torsion-free fundamental group for a maximal compact subgroup of G. The integral Demazure theorem [HLS10, Proposition 2.2] identifies its algebra with endomorphisms of the torus representation ring over the group representation ring. The former is free of rank $|W|$ over invariants, so complexification commutes with these endomorphisms. Their isobaric operator for $-\beta_s$ is $1 + e^{\beta_s}(z)d_s$; hence $\mathscr{D} = \mathrm{End}_S(A)$.

To exclude extra operators in either finite algebra E, choose an S-basis of A. Then $\mathrm{Mat}_{|W|}(S) \subset E \subset \mathrm{Mat}_{|W|}(K^W)$. For $x \in E$, matrix units put $x_{ij}e_{ii}$ in the finite corner algebra $e_{ii}Ee_{ii}$. Thus every x_{ij} is integral over the normal ring S, and belongs to S. Equality follows. All complexes and structure maps were ordinary degree zero; their given circle lifts take invariants and localize by tensoring with P. □

2.4.2. *Commuting parabolic actions.* Let J consist of pairwise commuting simple reflections and contain at most one element from each two-element θ-orbit. Then $I = J \cap \theta J$ consists of its individually fixed roots. For $L \subset J$, put $T_L = \Phi^{-1}(C_L)$ and $E_L(m) = \mathrm{RHom}(T_L \cdot m, m)$, with the coherent action on Q. Parabolic restriction $C_M \to C_L$ induces a unital coefficient map $E_L(m) \to E_M(m)$.

Lemma 2.55. *For disjoint θ-stable types J_1, J_2 and kernels $K_i \in \mathcal{H}_{J_i}$, the natural $J_{K_1 * K_2} \to J_{K_1} J_{K_2}$ is an isomorphism. The iterated assertion for pairwise disjoint stable types is compatible with all subdivisions.*

Proof. With real endpoint Borels fixed, an intermediate Borel lies in the two real endpoint parabolics. If they contain a common Borel, use it as origin. Their intersection has type $J_1 \cap J_2 = \varnothing$ and is that Borel. It is unique and real. Restriction imposing reality of the middle flag is therefore an isomorphism on every stalk of the integrand; its entire compact frame fibre is retained. Proper base change proves the statement. Group all but one type and repeat for several factors. These are the restriction maps of Lemma 2.14, already compatible with subdivision. □

Lemma 2.56. *For $L \subset I$, $E_L(P_C)$ is ordinary degree zero and freely generated on the left by $d_D = \prod_{s\in D} d_s$, $D \subset L$, with $d_\varnothing = e$.*

Proof. Choose an ordering. The restriction of d_s to the standard fibre is a generator $f_s : \Delta_s P_C \to P_C$; on closed coefficients it is an isomorphism of transported free lines. For $s \notin L$, Lemma 2.55 and closed adjunction identify precomposition by $T_L(f_s)$ with an isomorphism $E_L(P_C) \xrightarrow{\sim} \mathrm{RHom}(T_L\Delta_s P_C, P_C)$, with the transported input mark. Apply coefficients to $T_L\Delta_s \to T_{L\cup\{s\}} \to T_L \to$. Starting from A, it gives an ordinary short exact sequence. The lifts of the quotient basis are Fd_s: in the product of flag lines, the complementary counits delete exactly the coordinates inserted by comultiplication, by triangle identities. Adjoin these lifts to the embedded basis and induct. □

Lemma 2.57. *The same d_D form a free left A-basis of $E_L(Q)$ in degree zero. On both sides these products are independent of order.*

Proof. The spectral coefficient is $\Gamma(Y_L, \mathrm{End}\, M_L)$, and its underlying module is $\Gamma(Y_L, M_L \otimes M_L)$. Pass to the central cover $Z(L_L)^0 \times \prod_{s\in L} \mathrm{SL}_2$. The involution is trivial on each derived factor. With $x_s = -q_s$, and eigenvalue ratios $k_s = z'_s/z_s$, the rank-one calculation gives $A[k_s : s \in L]/((k_s - 1)(k_s - x_s) : s \in L)$. Every basis monomial in the k_s is unchanged by central changes of lift and descends, retaining central coupling. The trace is the product trace. The identity in one factor is $k_s - x_s$ and d_s is 1 there. Thus d_D is the polynomial $\prod_{s\in L\setminus D}(k_s - x_s)$, a basis, and the operators commute. On the real side the map $E_L(P_C) \to E(P_C)$ is generically the graph inclusion, hence injective by the freeness just proved. Orthogonal units j_s commute by Lemma 2.47. Their other roots fix q_s, so the corresponding d_s commute generically and then integrally. □

For arbitrary $L \subset J$, let $L_0 = L \cap I$. The actual counit restriction induces $E_{L_0}(m) \simeq E_L(m)$ on both sides. Indeed the real fibre is filtered by length-additive standards containing a moved letter whose mate is absent; none of their positions is θ-fixed, so Proposition 2.12 kills its coefficient. Spectrally, the two closed strata of Lemma 2.20 are the same $[P_{L_0}\tau/P_{L_0}]$: intersecting the partial flag with its odd conjugate recovers it. Both direct images come from the same $q^* M_{L_0}$, and weight-zero full faithfulness makes their coefficient map the identity.

Theorem 2.58. *For J as above, the closed marked comparison extends over the full Cartan to an $\mathcal{H}_J$-module equivalence on its generated subcategories, preserving all convolution maps and coherences.*

Proof. The ordinary additive idempotent category generated by T_L, $L \subset J$, is monoidal. Distinct factors commute by flag coordinates; repeated ones split by Lemma 2.30. The two simple triangles provide standard and costandard filtrations, hence degree-zero Hom. Their stable envelope contains all W_J standards and is the compact parabolic category. Generically its return group is $W_J \cap W^\theta = W_I$. Match the products of commuting j_s on both sides to obtain one graph comparison by Proposition 2.46. It matches $d_s = (e - j_s)/(1 + q_s)$ because this denominator is invertible when the boundary root q_s^2 differs from 1. The integral basis maps $d_D \mapsto d_D$ therefore restrict to that comparison. Naturality and binary products hold generically and have torsion-free targets by the preceding lemmas, so hold integrally, including after all direct-sum decompositions. All iterated products lie in the ordinary additive category and have degree-zero coefficient diagrams. Their coherences are determined by composition. Exact extension to cones and retracts, then small invariants and inversion of u, gives the internal algebra comparison. Use Lemma 2.15. □

2.4.3. *The parabolic corner near a root divisor.* Assume θ fixes all simple roots and put $S = R^W$. Let W_t denote the full Weyl stabilizer of $t \in T$.

Lemma 2.59. *After a Weyl translate suppose $W_t \subset W_J$. There is a pointed affine étale neighborhood $V \to T/W$ with a lift $V \to T/W_J$ through the image of t in T/W_J. On V the Cartan algebra separates into factors B_i, $i \in W/W_J$. For the identity factor the corner of the finite Hecke category is $\mathcal{H}_J \otimes_{R^{W_J}} \mathcal{O}(V)$.*

Proof. Replace t by a Weyl translate and assume $W_t \subset W_J$. The finite map $T/W_J \to T/W$ is étale at its image: completions along the two orbits split into point factors, and invariants identify both local completions with

$\widehat{\mathcal{O}}^{W_t}_{T,t}$, with identity map between them. Choose an affine étale neighbourhood V there and put $S' = \mathcal{O}(V)$. On the selected Cartan branch $B' = R \otimes_{R^{W_J}} S'$, all $1 - e^\gamma$ for $\gamma \notin \Phi_J$ may be made invertible: if $e^\gamma(t) = 1$, then $s_\gamma \in W_t$ and its root belongs to Φ_J. Remove images of the other-root divisors under the finite quotient. The diagonal of the separated étale map selects an open-and-closed branch. After shrinking, its W/W_J translates are disjoint and exhaustive, as at the chosen finite orbit. Hence

$$R \otimes_S S' = \prod_{i \in W/W_J} B_i, \qquad B_e = B'.$$

Let e_i be the unit idempotents. A standard has endpoints related by its Weyl element, so the selected (e, e) corner is generated precisely by W_J standards. The same-parabolic subcategory is full and monoidal; flat central extension and the common R^{W_J}-action identify this corner with $\mathcal{H}_J \otimes_{R^{W_J}} S'$. □

Lemma 2.60. *Let e_i be the idempotent of B_i and $S' = \mathcal{O}(V)$ as in Lemma 2.59. For each i there is an invertible kernel U_i from the identity corner to the i-th corner. It carries $e_e m_{S'}$ to $e_i m_{S'}$ for either closed object $m = P_C, Q$, with its marking.*

Proof. For a minimal representative $w_i = s_k \cdots s_1$ of $w_i W_J$, all pulled-back gallery roots $s_1 \cdots s_{j-1}(\alpha_{s_j})$ lie outside Φ_J. At each step the simple correspondence splits into identity and change-of-refinement graphs. Its standard is the latter, with opposite graph as inverse. Their products give invertible kernels $U_i \in e_i \mathcal{H}_{S'} e_e$.

These graphs also transport the whole closed object. On each simple root-regular open the integral $j_s = e - (1 + q_s)d_s$ lies entirely in the return coefficient: its identity projection would be both scalar and φ_s-semilinear in a domain and must vanish. Its square is the counit, so it is an isomorphism of the translated object with the closed object. The two rank-one calculations give this on P_C, Q. Compose these maps along the outside-root gallery to obtain $U_i(e_e m_{S'}) \simeq e_i m_{S'}$. At each step the full coefficient module is transported by the signed Weyl map, which intertwines the norm with boundary reflection. Thus every branch of the closed norm cover is retained. □

Theorem 2.61. *Suppose W_t is contained in a conjugate of W_J. If the full-Cartan comparison holds for the $\mathcal{H}_J$-action on the closed objects, then it extends to the full finite Hecke action on a pointed affine étale neighbourhood of the image of t in* Spec S. *It preserves the marking on the entire base-changed closed object. In particular this holds when the simple reflections of J commute.*

Proof. Apply Lemma 2.59 after translating t. Use the kernels of Lemma 2.60 to compare the module categories. The selected module component is generated from $e_e m_{S'}$ by the selected corner, since $e_e K e_i m_{S'} = e_e K U_i (e_e m_{S'})$. The hypothesized comparison therefore compares these corner modules. Transport every (i, j) kernel by $K \mapsto U_i^{-1} K U_j$. Inverse-kernel units and counits give composition and all coherences, and the specified return maps match each closed summand and its marking. The final assertion is Theorem 2.58. □

In particular, at a general single-root point with $W_t = \langle s_\beta \rangle$, the exact central fibre at a marked lift is $\mathbb{C}[\varepsilon]/(\varepsilon^2)$: in logarithmic coordinates the reflection fixes its tangent hyperplane and negates the normal coordinate. Derived scalar extension retains that normal direction and every norm-cover branch. The full stabilizer hypothesis is necessary when additional Weyl elements fix the point.

2.4.4. *The split extension criterion.* For a word $\mathbf{s} = (s_1, \ldots, s_k)$ in a split group, write $T_{\mathbf{s}} = T_{s_1} * \cdots * T_{s_k}$, $C_{\mathbf{s}} = \Phi(T_{\mathbf{s}})$, and $N_{\mathbf{s}} = \mathrm{RHom}_{\mathcal{P}_E}(C_{\mathbf{s}} \cdot Q, Q)$. These are ordinary complexes with output A-action.

Theorem 2.62. *For a split group of arbitrary isogeny type, the conclusion of Theorem 1.1 follows if each $N_{\mathbf{s}}$ is concentrated in degree zero and finite reflexive over A.*

Proof. For each w, take a reduced Bott–Samelson tilting. Its top standard has multiplicity one, and its other standards have smaller Bruhat labels. The hypotheses apply to these words, so the general ordinary criterion, Theorem 3.38, proves the assertion. That theorem's proof is independent of the present split criterion: it uses the parabolic corner of Theorem 2.61, the marked rank-one comparisons, and ordinary tilting generation. This also specifies the common generic normalization in every central form. □

Corollary 2.63. *For every split real reductive group, the full flag coefficient algebras $E(P_C)$ and $E(Q)$ are isomorphic with both markings on the full Cartan. Their common algebra is the subalgebra of* $\mathrm{End}_{A^W}(A)$ *generated by A and $d_s(a) = (a - s(a))/(1 - e^{\beta_s}(z))$.*

Proof. Cover G by its simply connected derived group times its connected centre, with Cartan ring A'. The odd cosets are copies of the groups with conjugation; the flag varieties are unchanged. Proper flat base change pulls the stable Borel bundle and its endomorphisms to the finite étale cover. Invariant sections inject, and right Maurer orientations identify each simple trace operator with its namesake upstairs. By Theorem 2.54, the upstairs algebra is its Demazure algebra. Fixed reduced products d_w are a basis over either torus ring by [HLS10, Theorem 2.1(ii) and Lemma 2.4]. The generic injection preserves j_w, so an upstairs element in $K \rtimes W$ has its d_w coefficients in $A' \cap K = A$, using integrality of the cover and normality of A. This bounds $E(Q)$ by the downstairs Demazure algebra, while its simple counits give every generator. Hence equality holds and $E(Q)$ is free in every isogeny type. Lemma 3.36 and the reflexive-extension argument in the proof of Theorem 3.38 apply to this single full-flag kernel: both sides are already free and ordinary. It extends the generic identification. Products, unit, and simple operators agree by torsion-freeness. □

2.4.5. *Two distinct simple changes at singular holonomy.* Still in the split case, let $s \neq t$ and

$$S_i = [(P_i\tau \times^{\mathbf{R}}_{P_i} B)/B], \quad \iota_i : S_i \to \mathcal{V}, \quad Z = S_s \times^{\mathbf{R}}_{\mathcal{V}} S_t.$$

Their ordinary intersection is $\mathcal{F}$ because $P_s \cap P_t = B$. Put $R_B = \mathcal{O}(B\tau)$.

Lemma 2.64. *On a torus-stable chart around $B\tau$, the derived intersection has an augmented dg model over R_B, freely generated in degrees $-1, -2$ by strictly positive dominant weights. Its differential preserves generator number. On each augmentation quotient J^n/J^{n+1}, the S_s-algebra action factors through R_B.*

Proof. Take the big-cell chart where both $a\tau^{-1}$ and a^2 lie in U^-B. A negative-root coordinate x_γ has function weight $+\gamma$. The boundary condition has Koszul algebra $D = \mathcal{O}(\Omega)[\eta_\gamma : \gamma > 0], \quad |\eta_\gamma| = -1, \quad d\eta_\gamma = f_\gamma$. The model D_i for S_i sets $x_\gamma = \eta_\gamma = 0$ for $\gamma \neq \alpha_i$. For $\gamma \neq t$, write $f_\gamma = \sum_{\delta \neq t} A_{\gamma\delta} x_\delta$ with homogeneous coefficient weight $\gamma - \delta$. A semifree resolution of D_t over D adjoins

$$|p_\delta| = -1,\ dp_\delta = x_\delta, \qquad |u_\gamma| = -2,\ du_\gamma = \eta_\gamma - \sum_{\delta \neq t} A_{\gamma\delta} p_\delta, \quad \gamma, \delta \neq t.$$

The differential squares to zero. First quotient by the regular x_δ via their Koszul generators; the remaining (η_γ, u_γ) pairs contract, leaving D_t. Polynomial even generators in degree -2 are essential.

Tensor this resolution with D_s. Contract (p_s, x_s), then replace η_s by $\eta_s - \sum_{\delta \neq s,t} A_{s\delta}(0, b) p_\delta$ and contract it with u_s. The result is

$$D_Z = R_B[p_\delta, u_\gamma : \gamma, \delta \neq s, t], \quad dp_\delta = 0, \quad du_\gamma = -\sum_{\delta \neq s,t} A_{\gamma\delta}(0, b) p_\delta.$$

The weights are respectively $\delta, \gamma > 0$ and the differential is linear in generators. The D_s-action sends x_s to zero and η_s to the indicated linear combination of p's; it therefore acts through R_B on the generator-number associated graded. □

Lemma 2.65. *Let E_i be a vector bundle on S_i whose restriction to the atlas $B\tau$ has a torus-equivariant frame of conjugation weight zero. With right Maurer orientation the augmentation induces, naturally in the bundles,*

$$R\Gamma(\mathcal{F}, (E_t|_{\mathcal{F}})^\vee \otimes E_s|_{\mathcal{F}}) \xrightarrow{\sim} \mathrm{RHom}_{\mathcal{V}}(\iota_{t*}E_t, \iota_{s*}E_s).$$

Proof. The augmentation gives $\pi_{s*}\pi_t^* E_t \to j_{s*}(E_t|_{\mathcal{F}})$. Base change and adjunction put the required Hom on S_s. For $j_s : \mathcal{F} \to S_s$, the normal tangent of $B\tau$ inside $P_s\tau$ and the obstruction tangent of the square condition in P_s/B are the same line in right Maurer coordinates. Their determinants cancel in virtual codimension zero. Thus $j_s^! E_s = E_s|_{\mathcal{F}}$, and dualizing the augmentation gives exactly the displayed map. Its cone is supported on $B\tau$, so the chart of Lemma 2.64 computes its sections by localization.

For $n > 0$, J^n/J^{n+1} is a finite free R_B-complex in degrees $[-2n, -n]$, all of whose generators have strictly positive weight. Tensor the input bundle before dualizing; its action on this graded term factors through its restriction to $B\tau$. The relative duality just computed makes the dual term a finite complex in degrees $[n, 2n]$ with strictly negative weights. The functions of R_B have nonpositive weights, and both bundle frames have weight zero. Thus the positive graded terms contribute only strictly negative cohomology weights after duality.

This calculation converges to the full ordinary Hom. The tail J^n tensored with the bundle is in degrees at most $-n$, while the target bundle on the one-equation presentation of S_s is in degrees at least -1. Orthogonality puts Hom from that tail in degrees at least $n - 1$. Every fixed degree therefore stabilizes on a finite quotient. Inducting along its finite filtration gives strictly negative weights for the canonical cone. Finally take the full $B = T \ltimes U_B$ gauge. Unipotent Lie cochains contribute $\bigwedge^\bullet \mathfrak{u}_B^*$, again strictly negative in positive degrees. Their

finite cohomological spectral sequence has zero torus invariants on every term of the cone. Hence its derived B-invariants vanish. □

Proposition 2.66. *For every split group and distinct simple s, t, the ordinary coefficient $N_{(s,t)}$ is free of rank four over output A and concentrated in degree zero. Its marked generic comparison extends over the entire Cartan.*

Proof. Self-adjunction gives $N_{(s,t)} = \mathrm{RHom}(C_tQ, C_sQ)$. Lemma 2.20 and base change write $C_iQ = \iota_{i*}E_i$, where E_i pulls back the vector bundle M_i. With one stable Borel marked, the rank-one ratio basis $1, z'/z$ on a central cover is conjugation invariant and unchanged by central changes of lift; it descends to a weight-zero frame. The dual frame has the same property. Lemma 2.65 reduces to their ordinary sections on $\mathcal{F}$. Nonconstant unipotent functions and positive-degree unipotent cochains have strictly negative weights, so only Cartan functions remain.

The marking on $\mathcal{F}$ is the middle one. Output A acts on E_s through the additional stable Borel before its pushforward. The section algebra of $E_s|_{\mathcal{F}}$ is free of rank two for the middle marking by its ratio basis and also for the other marking: it is the stable pair algebra in the Levi, and swapping the two Borels exchanges the markings. On the cover this sends z'/z to z/z' and commutes with every central lift. Unipotent radical coordinates do not affect the zero-weight sections. The dual frame of E_t now gives two copies of this rank-two output module. It is therefore free of rank four. Lemma 3.36 and the reflexive-extension argument of Theorem 3.38 apply to this single already free coefficient and extend its specified generic comparison. □

2.4.6. *A finite tilting criterion.*

Theorem 2.67. *For a split group choose, for every $w \in W$, a compact universal tilting B_w with one standard subquotient Δ_w and all others below w. If each ordinary coefficient $\mathrm{RHom}_{\mathcal{V}}(\Phi(B_w)\cdot Q, Q)$ is concentrated in degree zero and finite reflexive over output A, then the categorical comparison holds on the entire marked Cartan.*

Proof. This is the split, top-multiplicity-one case of Theorem 3.38. The additive generation needed there is proved in Lemma C.11; in particular, no implication from a periodic comparison to ordinary concentration is used. □

3. Ordinary coefficients and the quasi-split comparison

The regular comparison fixes the identification of the two coefficient algebras over the generic point. To extend it across the root divisors we first prove that the ordinary coefficients of tilting kernels are projective. Parabolic induction supplies the initial kernels. The root-system arguments completing the induction are given in Appendix C. We then compare the two projective modules at each divisor, using the same generic identification for all kernels.

3.1. **Ordinary coefficients and parabolic sources.** The comparison problem is governed by ordinary coefficients. This subsection first establishes the finiteness and lower amplitude that hold without concentration. We then prove concentration for words whose two ends are supported in one proper stable Levi. The middle parabolic need not be stable under the involution. This flexibility is used in outer type A and in the exceptional finite calculations.

3.1.1. *Perfection and output fibers.* All coefficients in this section are taken before circle invariants or periodization. The exterior generator for a Koszul equation has degree -1 and its even resolution variable has degree -2. The dualizing complexes and quotient gauges are retained throughout.

Write

$$A = \mathcal{O}\big(T/(1-\theta)T\big), \qquad W_o = F \times^{R}_{[E^-/G]} F, \qquad N_G(K) = R\,\mathrm{Hom}_{V_G}(KQ, Q).$$

The output map $q : W_o \to \operatorname{Spec} A$ uses the original marking $[tr_{\rho,G}]$. It is not a map to a Weyl quotient. Finite groups, central tori and the actual odd square remain in every quotient.

The following statement will provide the algebraic input to concentration.

Proposition 3.1. *For every unshifted finite universal tilting kernel K,*

$$N_G(K) \in \mathrm{Perf}(A), \qquad N_G(K) \otimes^{L}_{A} k(y) \in D^{\geq 0}(k(y)) \quad \text{for every output point } y. \tag{3.1}$$

Projectivity of the degree-zero cohomology is a consequence of concentration, rather than an input to its proof. The distinction is relevant in codimension at least three, as the example below shows.

Proof. Bounded coherent endpoints and finite coefficients. Proposition 2.19 proves, for every $K \in \mathrm{Coh}(Z_G)$,

$$N_G(K) \simeq R\,\mathrm{Hom}_{\mathrm{QCoh}(W_o)}(Lh^*K, \mathcal{O}_{W_o}). \tag{3.2}$$

Its proof also establishes that KQ is bounded coherent. The relevant pullback is a derived base change of the map $B\tau \to B$ between smooth varieties, hence has finite Tor dimension, and the remaining push is proper. Thus Proposition B.10 applies to the two endpoints KQ, Q. It gives bounded, finite cohomology over the actual boundary Cartan ring R. This action on the coefficient factors through the original output action $R \to A$. A finite set of R-module generators consequently also generates each cohomology module over A. The Laurent torus algebra A is regular of finite global dimension. It follows that $N_G(K)$ is A-perfect. This proof does not use concentration.

Flatness of the actual output target. Lemma 2.18 proves that the incidence atlas $\widetilde{W_o}$ is a classical complete intersection, is Cohen–Macaulay, and is pure of dimension $\dim G$. Its finite relative-position stratification has pieces $G \times^{H_w} (H_w\tau), \qquad H_w = B \cap \dot{w}B\dot{w}^{-1}, \qquad w \in W^\theta$. On each piece the output map is induced by the smooth surjection $H_w\tau \longrightarrow T \longrightarrow T/(1-\theta)T$, followed by the fixed marking translation. The quotient by H_w is compatible with this map because twisted conjugation changes its Cartan coordinate by an element of $(1-\theta)T$. Each output fiber of each stratum is smooth and pure of dimension $\dim G - \dim \mathrm{Spec}\, A$. The fiber of the whole atlas is therefore pure of that dimension: the generic point of any irreducible component lies in one of the finitely many locally closed strata, and its local dimension is the dimension of that stratum. Miracle flatness [The26, Tag 00R4] now applies, since the source is Cohen–Macaulay and the target is regular. Thus

$$q : W_o \longrightarrow \mathrm{Spec}\, A \quad \text{is flat.} \tag{3.3}$$

Smooth descent passes flatness from the atlas to the stack.

The heart input. Lemma 2.24 shows that finite standard-filtered Hecke objects with finite Cartan coefficients are sent by Φ to the ordinary coherent heart. The proof uses ample Wakimoto tests and standard–costandard Hom vanishing. For the general central form, the pullback identity in Lemma 2.26, $\beta^*\Phi(\Delta_w) = \Phi'(\Delta'_w)$, and the finite etale incidence cover with a finite central gerbe give faithfully flat, conservative, t-exact descent of this heart assertion. Hence $Lh^*K \in D^{\le 0}$ for the actual standard-filtered kernels used here.

Tensoring the target with a perfect complex. For $L \in \mathrm{Perf}(A)$ there is a canonical identity

$$N_G(K) \otimes^L_A L \simeq R\,\mathrm{Hom}_{W_o}(Lh^*K, \mathcal{O}_{W_o} \otimes^L_A L). \tag{3.4}$$

Both sides are exact and preserve retracts as functors of L; the identity holds for $L = A$, so it holds on its thick envelope. This argument makes no compactness assertion about Lh^*K and requires no finite Tor bound on h. The A-action is indeed the target action: following an output multiplication endomorphism through the two adjunctions in Proposition 2.19 gives multiplication on $\mathcal{O}_{W_o}$.

At a closed point y the residue field has a finite Koszul resolution over A. By (3.3), $\mathcal{O}_{W_o} \otimes^L_A k(y) = \mathcal{O}_{(W_o)_y}$ is in the heart. The source in (3.4) is in $D^{\le 0}$, so ordinary t-structure orthogonality makes the Hom connective in cohomological degrees ≥ 0. Perfection extends the same assertion to every prime point: the locus where a perfect complex has no negative fiber degrees is open, and the torus is Jacobson. This proves (3.1). □

3.1.2. *Detection and concentration.*

Lemma 3.2. *Let P be perfect over a Noetherian ring, and suppose all its derived residue fibers have no negative cohomology. Then locally P admits a bounded finite free representative with terms only in nonnegative degrees. In particular it has nonnegative Tor amplitude.*

Proof. Over a local ring take a bounded finite free representative. Whenever a differential has a unit matrix entry, use invertible row and column operations and cancel the resulting contractible two-term summand. Finitely many cancellations produce a minimal representative, all of whose differentials have entries in the maximal ideal. Its reduction modulo that ideal has zero differential. The assumed fiber vanishing therefore forces every negative term to be zero. Tensoring this model with any module is visibly nonnegative. The assertion is local, so it gives the stated Tor-amplitude conclusion. □

Lemma 3.3. *Let P satisfy Lemma 3.2 and $H^j(P) = 0$ for every $j > 0$. Then $P \simeq M[0]$ for a finite projective module M.*

Proof. Use a local finite free model $F^0 \to \cdots \to F^b$. If $b > 0$, the equality $H^b(P) = 0$ makes $F^{b-1} \to F^b$ surjective. Since F^b is projective, this surjection splits; cancel its contractible summand. Repeating removes every positive term and leaves a finite free module in degree zero. These local descriptions show that $H^0(P)$ is finite locally free. On the affine base it is finite projective. This is the amplitude-zero case of the standard perfect-complex criterion [The26, Tag 0658]. □

Reflexivity. Over the regular torus ring A, (3.1) already implies that $H^0(P)$ is finite reflexive. Indeed, locally $H^0(P) = \ker(F^0 \to F^1)$ is torsion-free, and $I = \mathrm{im}(F^0 \to F^1)$ is torsion-free. At a local ring of dimension $d \geq 2$, $\operatorname{depth} I \geq 1$, and the depth lemma in $0 \longrightarrow H^0(P) \longrightarrow F^0 \longrightarrow I \longrightarrow 0$ gives depth at least two for $H^0(P)$. It is thus torsion-free and S_2 over a normal domain, hence reflexive [The26, Tag 0AY6].

For example, over $R = \mathbb{C}[x, y, z]_{(x,y,z)}$ the complex

$$[R^3 \xrightarrow{(x\ y\ z)} R] \quad \text{in degrees } 0, 1 \tag{3.5}$$

is perfect and has nonnegative derived fibers. Its degree-zero kernel is the second syzygy of the residue field, of rank two and with three minimal generators. Thus projectivity requires a further vanishing argument.

Completion and detection. For a maximal ideal $\mathfrak{m}$, $A_{\mathfrak{m}} \to \widehat{A}_{\mathfrak{m}}$ is faithfully flat [The26, Tag 00MC]. Flatness gives

$$H^j(P \otimes_A^L \widehat{A}_{\mathfrak{m}}) = H^j(P)_{\mathfrak{m}} \otimes_{A_{\mathfrak{m}}} \widehat{A}_{\mathfrak{m}}. \tag{3.6}$$

Consequently completion detects the corresponding localized vanishing. If all maximal localizations vanish, the finite module vanishes globally. Likewise, for a perfect complex over a local ring, $P \otimes^L k = 0$ implies $P = 0$: the minimal model in Lemma 3.2 then has every term zero, without any prior degree-zero-projectivity assumption. This is the derived Nakayama step needed for the final E8 coefficient.

For a noncentral norm Proposition 3.24 identifies the local coefficient with coefficients in its connected centralizer. Formula (3.6) then applies to the finite modules on these local charts.

Perfection and nonnegative derived fibers are preserved by finite direct sums, retracts and flat completed scalar extension. Thus they apply to the actual completed canonical summands once Proposition 3.16 realizes these as retracts of ordinary word kernels. No global algebraization of a completed projector is needed for this conclusion.

3.1.3. *A common Levi and an arbitrary middle parabolic.* For a parabolic type J, let C_J be the unshifted averaging kernel. Only the common end type will be required to be stable.

For one averaging kernel Proposition 2.22 expresses the ordinary coefficient as maps between copies of the vector bundle $M_{J\cap\theta J}$ on the affine odd Levi coset, followed by reductive invariants. Lemma 2.21 proves local freeness by duality for the generically finite crepant flag map and Grauert–Riemenschneider vanishing. For an inherited central square the determinant calculation is unchanged, and a stable dominant cocharacter gives a point with regular semisimple square.

Two inductions enter the extension from one kernel to a word. A central norm block lies over the image of an element of $Z(G)$ in the even adjoint quotient, with the odd markings retained.

Step	Inductive input
Central norm, end words in I	Ordinary word coefficients for connected factors of Levis with roots in Φ_I. Their ranks are at most $\lvert I\rvert$, strictly smaller than the ambient semisimple rank.
Noncentral norm b	Ordinary word coefficients for the connected factors of $Z_G(b)^\circ$, with their inherited actions. The order in Section 3.4.6 places D_n before B_n and the classical types before G_2, F_4, including proper centralizers of full rank.

The first step invokes the smaller Levi theorem after the flag correspondence has identified its coefficient. In the second step, Proposition 3.24 identifies the flag components and both markings before drawing any conclusion about cohomological degrees.

Theorem 3.4. *Fix a proper θ-stable standard parabolic type I, and let U, V be unshifted averaging words supported in I. Let J be any standard parabolic type. Assume ordinary concentration for all inherited connected cores inside I, with their actual central square, spectator torus, markings and component summands. Then the coefficient of UC_JV has no positive cohomology at every central norm block. It is therefore ordinary at that block. If the transported coefficients in all required proper norm centralizers are also ordinary, the marked centralizer formula globalizes this conclusion and the coefficient is finite projective over the output Cartan.*

Fix $b_0 \in Z(G)$ and put $C_0 = \mathcal{O}_{T/W,\bar{b}_0}$, with residue field k. For the geometric words considered here equation (3.53) is faithfully flat base change along the isogeny of marked tori. Its proof keeps the finite central inertia of the source. It therefore suffices to calculate for a simply connected derived group times a torus. The torus contributes its degree-zero Laurent ring with its norm action, as in the product calculation preceding Proposition 3.25. For this product C_0 is regular by the Steinberg invariant-ring theorem. The norm on the closed torus is finite onto its image, and $T \to T/W$ is finite. Hence Proposition 3.1 makes the localized coefficient bounded with finite C_0-cohomology. We shall prove

$$N_G(UC_JV) \otimes^L_{C_0} k \in D^{\leq 0}(k). \tag{3.7}$$

The finite Koszul resolution of k has nonpositive degrees and weight zero under conjugation. In particular negative Tor groups can occur in this fiber.

Put $H = I \cap J$, let U_I denote the unipotent radical of P_I, and write $P_H^{L_I} = P_H \cap L_I$. The stacks for the middle step are

$$Y = E^- \times^R_{G,\mathrm{sq}} P_J, \qquad X_J = [Y/P_J], \qquad X_H^I = \left[(L_I\tau) \times^R_{L_I,\mathrm{sq}} P_H^{L_I}/P_H^{L_I}\right].$$

In each fiber product the odd element maps by squaring and the even parabolic by inclusion. Each group in the denominator acts by conjugation. The middle correspondence is

$$X_H^I \xleftarrow{q_H} \mathcal{S}_H = \left[(P_I\tau) \times^R_{P_I,\mathrm{sq}} P_H/P_H\right] \xrightarrow{k_H} X_J. \tag{3.8}$$

Here q_H takes the Levi quotient and k_H uses $P_H \subset P_J$ and $P_I\tau \subset E^-$. Both P_H and $\theta(P_H)$ are contained in P_I, so twisted conjugation preserves $P_I\tau$ even when J is not stable. Let $p_J : V_G \to X_J$ and $p_H^I : V_{L_I} \to X_H^I$ forget the full flags. Use $P_\varnothing = B$ in the source of (3.8) to obtain $\mathcal{S}_\varnothing$. Its maps $q_\varnothing$ to V_{L_I} and $k_\varnothing$ to V_G take the Levi quotient and remember the ambient full flag. Write $\alpha_I = k_{\varnothing *} q_\varnothing^*$.

Lemma 3.5. *Let Q_I be the closed Borel object of L_I. Put $M = V^I Q_I$ and $N = (U^I)^{\mathrm{rev}} Q_I$, where the superscript I uses the same averaging letters in L_I and* rev *reverses their order. For $D = M, N$ put $E_D = k_{H*} q_H^* Rp^I_{H*} D$. Then*

$$N_G(UC_JV) \simeq \mathrm{RHom}_{X_J}(E_M, E_N). \tag{3.9}$$

The output Cartan acts at the closed end of N. Both sides carry the action of the even adjoint quotient through the class of the square on X_J.

Proof. On the atlas Y the choices of k_H form a closed derived incidence in the proper flag variety $Z = P_J/P_H$. Thus k_H is proper. The decompositions $P_I = L_I \ltimes U_I$, $P_H = P_H^{L_I} \ltimes U_I$ and $B = B_I \ltimes U_I$ give the derived Cartesian square

$$\begin{array}{ccc} \mathcal{S}_\varnothing & \xrightarrow{q_\varnothing} & V_{L_I} \\ {\scriptstyle h}\downarrow & & \downarrow{\scriptstyle p_H^I} \\ \mathcal{S}_H & \xrightarrow{q_H} & X_H^I. \end{array}$$

Indeed, a full refinement is determined by its Levi flag. The condition that the even element lie in the refined Borel is the pullback of the Levi containment condition along the map on squares. Forgetting the refinement gives the left vertical map h, and $p_J k_\varnothing = k_H h$. Proper base change gives

$$Rp_{J*}\alpha_I D \simeq k_{H*} q_H^* Rp^I_{H*} D. \tag{3.10}$$

In the odd-stable Borel locus the analogous Levi square identifies the pulled-back structure sheaf, giving $\alpha_I Q_I = Q$. For a step inside I, the partial flag remembers the type-I parabolic. Both refinements are therefore refinements in its Levi. The same base change for the forgetful map and its pullback intertwines this step with its L_I-operator. Iteration identifies $VQ = \alpha_I M$ and $U^{\mathrm{rev}} Q = \alpha_I N$.

Each parabolic averaging operator is self-adjoint. The partial flag map is a base change of the map of smooth Grothendieck resolutions and has finite Tor dimension. On its source and target right translation identifies the determinants of group, square equation and gauge, giving trivial dualizing lines in virtual dimension zero. Thus $p^! = p^*$ for this map and the two adjunctions make $p^* Rp_*$ self-adjoint. Move U across Hom in reverse order and use $p_J^* \dashv Rp_{J*}$. Formula (3.10) on both ends now gives (3.9). The adjunctions carry scalar multiplication at the output through the same flags. The common even coordinate is the class of a^2. □

Lemma 3.6. *The two objects E_M, E_N of Lemma 3.5 are self-dual for $D = R\mathcal{H}om_{X_J}(-, \mathcal{O}_{X_J})$. Their mapping complex is the equivariant dual of the derived tensor of their two incidence sources on $Y \times Z^2$, computed with the relative dualizing complex of the flag projection. Taking the derived fiber over C_0 tensors the dualizing target with k over C_0.*

Proof. All the square-parameter stacks above have virtual dimension zero and their canonical coherent-duality target is $\mathcal{O}[0]$. On an atlas the determinant cancellation is the right-Maurer exact sequence of the group tangent, square normal and gauge tangent. For the source (3.8), if i, j, h count positive roots in I, J, H, respectively, it is the bundle identity

$$\det \mathfrak{p}_I \, \frac{\det \mathfrak{p}_J}{\det \mathfrak{p}_H} \, \frac{\det \mathfrak{p}_H}{\det \mathfrak{p}_I} = \det \mathfrak{p}_J, \qquad \dim P_I + (j-h) - (i-h) = \dim P_J. \tag{3.11}$$

The group quotient cancels the remaining determinant. The cancellation takes place in the determinant lines of the equivariant bundles.

The seed Q_I is self-dual for this duality: its smooth odd Borel source also has virtual dimension zero and the same canonical orientation, so proper duality applies. For every partial flag map p_L, finite Tor dimension and the same dimension-zero relative orientation give $p_L^! = p_L^*$. Hence $D(p_L^* R p_{L*} M) = p_L^* R p_{L*} D(M)$. Every actual word translate of Q_I is therefore self-dual. The radical pullback q_H preserves this normalization: its tangent and full U_I gauge cancel, giving $q_H^! = q_H^*$. Proper duality then makes the two actual proper pushes E_M, E_N in (3.10) self-dual. Consequently

$$R\mathcal{H}om(E_M, E_N) \simeq D(E_M \otimes^L E_N). \tag{3.12}$$

Here D on the tensor product denotes the indicated derived internal Hom, including when its argument is unbounded below. On $Y \times Z$ let $\mathscr{E}_M, \mathscr{E}_N$ denote the closed direct images defining E_M, E_N before forgetting the flag choice. With coordinates in the order (Y, Z, Z) put

$$\mathscr{K} = L\,\mathrm{pr}_{12}^* \mathscr{E}_M \otimes^L L\,\mathrm{pr}_{13}^* \mathscr{E}_N.$$

By proper base change and the projection formula the pullback of $E_M \otimes^L E_N$ to Y is $R\pi_* \mathscr{K}$, where $\pi : Y \times Z^2 \to Y$ forgets the two flag choices. For this smooth proper projection,

$$\pi^! \mathcal{O}_Y = \mathcal{O}_Y \boxtimes \omega_{Z^2}[2 \dim Z]. \tag{3.13}$$

Consequently the complex on the left of (3.7) is

$$\mathrm{RHom}_{[(Y \times Z^2)/P_J]}\big(\mathscr{K},\, (\mathcal{O}_Y \boxtimes \omega_{Z^2}[2 \dim Z]) \otimes^L_{C_0} k\big),$$

where the even adjoint quotient acts through the square on Y. Indeed proper duality gives this formula before tensoring with k, by the adjunction on quasi-coherent complexes for the smooth flag projection. Its source can be unbounded below. Lemma 3.10 gives the degreewise bounds in the local calculation. For a perfect C_0-complex the target-tensor identity follows by finite cones and retracts from the identity for C_0, as in (3.4). Apply it to the Koszul resolution of k. □

Lemma 3.7. *Every orbit of the diagonal P_J-action on Z^2 contributing to Lemma 3.6 has a θ-fixed minimal representative $w \in {}^H W_J{}^H$. The intersection of its two type-I parabolics has Levi factor L_K with $\Phi_K = \Phi_I \cap w\Phi_I$; the residual partial flag in this Levi has type $H_K = H \cap K$.*

Proof. The diagonal P_J-orbits in Z^2 have minimal representatives $w \in {}^H W_J{}^H$. Such a representative is also minimal in its W_I double coset: a simple root outside J cannot be a descent of an element of W_J, and the remaining I-descents are exactly the forbidden H-descents. An odd element preserving the two remembered type-I parabolics preserves their ordered relative position. Thus the actual source is empty unless $w = \theta(w)$, whether or not J is stable.

For a surviving w, set

$$R_w = P_I \cap wP_I, \quad S_w = R_w \cap P_J = P_H \cap wP_H, \quad \Phi_K = \Phi_I \cap w\Phi_I, \quad H_K = K \cap H.$$

Minimality makes K a standard Levi type. Since I and w are stable, so is K. The negative roots common to P_I and wP_I are the negative roots of L_K. Every root of the common unipotent radical is thus positive, so this radical lies in $B \subset P_J$. Taking its quotient gives

$$R_w / S_w = L_K / P_{H_K}^{L_K}. \tag{3.14}$$

The actual common square core is

$$\mathcal{C}_w = [(R_w \tau) \times^R_{R_w, \mathrm{sq}} S_w / S_w]. \tag{3.15}$$

This is well defined without stability of S_w. Its square condition is precisely the partial boundary condition left by the original source diagram. Its canonical dualizing target is $\mathcal{O}[0]$, because

$$\det(\mathfrak{p}_J/\mathfrak{s}_w)\frac{\det \mathfrak{r}_w}{\det(\mathfrak{r}_w/\mathfrak{s}_w)} = \det \mathfrak{p}_J. \tag{3.16}$$

□

Fix a surviving representative w. Take a positive integral multiple λ_I of the sum of the fundamental coweights for the simple roots outside I, chosen in the cocharacter lattice of the derived group. It is θ-stable. Put $\lambda = \lambda_I + w\lambda_I$. Every root of R_w is nonnegative on both summands. Their sum is zero exactly on Φ_K and strictly positive on the full common radical. Write the semisimple odd form as $a = u\ell\tau$ with u in the common radical and $\ell \in L_K$. Its central square gives $\mathrm{Ad}(\ell\tau)(u) = u^{-1}$. In characteristic zero $u^{1/2}$ is defined algebraically and $u^{-1/2}au^{1/2} = \ell\tau$. We use this representative for the semisimple form and keep the radical in the equivariant Hom.

Lemma 3.8. *For $B_w = \{\alpha > 0 : w^{-1}\alpha < 0\}$, the weights of the common forbidden odd coordinates, the two endpoint square lists outside the common Levi, the residual even generators, and the middle conormal coordinates are, respectively,*

$$\begin{aligned} P &= \Phi^+ \setminus (\Phi_I^+ \cup w\Phi_I^+ \cup B_w),\\ E_1 &= \Phi_I^+ \setminus (\Phi_H^+ \cup \Phi_K^+),\\ E_2 &= w\Phi_I^+ \setminus (w\Phi_H^+ \cup \Phi_K^+),\\ U &= \Phi^+ \setminus (\Phi_J^+ \cup \Phi_K^+),\\ M_w &= \Phi_J^+ \setminus (\Phi_H^+ \cup w\Phi_H^+ \cup B_w). \end{aligned} \tag{3.17}$$

All these weights are strictly positive on λ, and their characters satisfy

$$P \sqcup E_1 \sqcup E_2 = U \sqcup M_w. \tag{3.18}$$

Proof. A negative root direction $-\alpha$ belongs to P_I precisely when $\alpha \in \Phi_I^+$. It belongs to the conjugate parabolic wP_Iw^{-1} precisely when $w^{-1}\alpha < 0$ or $w^{-1}\alpha \in \Phi_I^+$. This gives P. In the first Levi the square cuts out $P_H^{L_I}$, with function roots $\Phi_I^+ \setminus \Phi_H^+$. Removing the common Levi gives E_1. Transport to the other end gives E_2. The ambient square has function roots $\Phi^+ \setminus \Phi_J^+$, with residual list U after removing the common Levi. Subtracting the orbit tangent $\mathfrak{p}_J/\mathfrak{s}_w$ from the two tangents of Z leaves the roots $-\alpha$ for $\alpha \in M_w$.

Minimality makes B_w disjoint from Φ_I^+ and $w\Phi_I^+$, and $B_w \subset \Phi_J^+$. The two positive systems of the Levis meet in Φ_K^+. The left side of (3.18) therefore contains once each positive root outside $\Phi_K^+ \cup \Phi_H^+ \cup w\Phi_H^+ \cup B_w$. Since $\Phi_K^+ \cap \Phi_J^+ \subset \Phi_H^+$, the right side has the same multiplicities. Each displayed list is disjoint from B_w and Φ_K^+. For a positive root outside B_w, both α and $w^{-1}\alpha$ are positive. Their pairings with λ_I are nonnegative and both are zero precisely on Φ_K^+. This proves strict positivity and verifies the weight balance for (3.11) and (3.16). □

Lemma 3.9. *On a surviving orbit the derived tensor of the incidence sources has a semifree model over the square incidence (3.15). Its remaining generators in negative degree have strictly positive λ-weights.*

Proof. Use products of complementary root groups followed by the common subgroup as charts for each pair of parabolics. Their tangent maps are isomorphisms along the common subgroup by the direct decomposition of root spaces. The charts are therefore étale there. Use the analogous even charts and pull their equations back by the square map. Let R_0 be an ordinary smooth chart, let $x_1, \dots, x_a$ be the regular source-forbidden coordinates, and let

$$D = R_0[\eta_1, \dots, \eta_b], \quad d\eta_\rho = f_\rho, \qquad S = (R_0/(x))[\epsilon_1, \dots, \epsilon_c], \quad d\epsilon_i = g_i.$$

Here the η_ρ, ϵ_i have degree -1. Lift each g_i to R_0. The restricted ambient square equation is in the ideal of the source equations, and lifting this equality to R_0 gives the homogeneous expression

$$f_\rho = \sum_i A_{\rho i} g_i + \sum_j H_{\rho j} x_j.$$

The D-algebra

$$\mathcal{P} = D[p_j, \epsilon_i, u_\rho], \quad dp_j = x_j, \quad d\epsilon_i = g_i, \quad du_\rho = \eta_\rho - \sum_i A_{\rho i}\epsilon_i - \sum_j H_{\rho j} p_j \tag{3.19}$$

is a semifree resolution of S. Here each p_j has degree -1 and u_ρ has degree -2. Taking the appropriate weight components gives homogeneous coefficients on each torus chart. Indeed the differential squares to zero. Replace η_ρ by the right side of du_ρ. Each resulting u_ρ, du_ρ pair is contractible, by a homotopy dividing only by positive integers. What remains is the regular Koszul resolution of the x_j tensored with the source square algebra. Thus its augmentation is a quasi-isomorphism with precisely the original ambient action.

Tensor the two copies of (3.19) over the actual D before dualizing. Independent forbidden ordinary coordinates are eliminated by regular Koszul contractions. On the common neutral square list $\Phi_K^+ \setminus \Phi_{H_K}^+$ the two sections are restrictions of the same norm section. Their difference cancels one duplicate against the corresponding ambient even variable with coefficient one. No derivative such as $1 + \alpha(a_0)$ is inverted.

After eliminating each independent forbidden odd coordinate, one exterior generator remains for each shared forbidden coordinate in P. The exterior square generators outside the common Levi have lists E_1, E_2, and the even generators for ambient square equations outside that Levi have list U. The common square equations and their differentials give (3.15). Lemma 3.8 gives strict positivity for every remaining generator. □

Lemma 3.10. *Suppose a semifree source has finitely many negative-degree generators, all of degrees -1 or -2, over a core of fixed finite cohomological width, and the duality target has fixed finite width. For each fixed ordinary Hom degree only finitely many negative-degree monomials are relevant. The same remains true after finite Čech and Lie-cochain totalizations.*

Proof. A monomial of augmentation length m lies in degrees between $-2m$ and $-m$, up to a fixed core shift. A map to a bounded target in a fixed Hom degree can therefore use only bounded m. There are finitely many such monomials and finitely many relevant target degrees. A finite extra totalization changes the bound only by a constant. This proves stabilization degree by degree. Degree-zero ordinary coordinate powers are not controlled by this estimate; they are controlled by the strict weight inequalities and the proper finite-dimensional choice space. Normal-pole powers are treated separately in the next lemma. □

Lemma 3.11. *On a middle orbit the dualizing target for the support filtration has nonpositive normal weights. Its weight-zero normal factor is the unshifted residue line with trivial determinant character.*

Proof. For a middle orbit closed in a successive smooth open, choose normal functions $z_1, \dots, z_c$ of positive weights β_i. The actual dualizing factor obtained from (3.13),(3.16) is $\Omega^c_{\mathrm{normal}}[c]$. Its supported local cohomology has top terms

$$\frac{dz_1 \wedge \cdots \wedge dz_c}{z_1^{n_1} \cdots z_c^{n_c}}, \qquad n_i \geq 1, \quad \mathrm{wt} = \sum_i (1 - n_i)\beta_i \leq 0. \tag{3.20}$$

The unique neutral term is the leading residue. Its degree is zero and its determinant character is trivial. All other poles have strictly negative weight. The leading residue is annihilated by the normal ideal. Thus a differential with coefficient in that ideal acts trivially on the leading residue. The other pole orders account for the normal thickenings in the Čech complex. The identification with the dual Koszul residue [Har66b], [The26, Tags 0BQZ and 0BJD] glues by the determinant transformation of forms and poles, hence is the proper duality counit with the original orientation. Tensoring the finite weight-zero Koszul target for k leaves all these weight inequalities unchanged and only adds a fixed finite cohomological width. □

Lemma 3.12. *In either endpoint gallery of a surviving orbit, the fixed components are galleries in L_K with successive parabolic steps in L_K and final projection to $L_K/P_{H_K}^{L_K}$. The neutral equations on these galleries are the infinitesimal stability equations for the nilpotent part of the holonomy.*

Proof. Write $L = L_I$. The restriction of λ to L is dominant. Its zero roots are the roots of L_K. An endpoint word is represented by its actual smooth proper gallery X, followed by the map $f : X \to L/P_H^L$, where $P_H^L = L \cap P_H$. The initial flag lies in the fixed flag scheme of the chosen semisimple odd form $a_0 \in L_K\tau$. Its square is central, so its action on flags has order dividing two. The fixed scheme is smooth in characteristic zero. Each subsequent partial-flag correspondence is smooth and proper over the preceding gallery. The equations for the variable nilpotent holonomy will be imposed on this smooth space.

On the central norm fiber, write the holonomy in its semisimple-centralizer slice as $a_0 \exp(x)$, where x commutes with a_0 and is nilpotent. The initial odd stability condition and the subsequent square stability conditions become the infinitesimal stability conditions for x on the remembered flags; the logarithm of the unipotent part of the square is $2x$, and 2 is a unit. On a moving normal direction of the a_0-fixed locus, the eigenvalue ζ of a_0 is different from 1. The operator $1 - \zeta \exp(\operatorname{ad} x)$ is invertible, since x is nilpotent and commutes

with a_0. These directions form regular Koszul pairs; the fixed directions retain their full x-stability equations. Thus their derived equations are the Koszul equations of the infinitesimal action on the gallery. In particular the exterior generator corresponding to a local function z has the same λ-weight as z.

For the scheme-theoretic identification use the Cartan-adic exponential and logarithm of Lemma 3.20 in the semisimple slice. The fixed part is calculated in the centralizer of a_0. In a faithful matrix representation, a sufficiently high power of the nilpotent matrix belongs to the Cartan ideal, so these series converge without completing the unipotent coordinates. The chain rule identifies the two actual section Koszul complexes. The norm followed by the Weyl quotient is finite onto its image, so on a selected Cartan branch the central maximal ideal and the Cartan maximal ideal have cofinal powers. Flat completion therefore preserves tensoring with the entire finite C_0-Koszul complex for the central residue, including its negative Tor terms. Each finite jet is algebraic in the unipotent coordinates. We can thus take rational invariants after passing to the algebraic derived fiber. The finite Koszul factors for the central parameters have weight zero and remain in the calculation.

The neutral galleries can be identified without a classification of the word. The λ-fixed components of L/B_L are

$$L_K \dot{d} B_L/B_L \simeq L_K/B_K, \qquad d \in {}^K W_L,$$

where the representatives are minimal for left W_K-cosets. In particular $L_K \cap {}^{\dot{d}}B_L = B_K$. For a parabolic step of type $J' \subset I$, put $P^L_{J'} = L \cap P_{J'}$, and take two minimal representatives $d, e \in {}^K W_L$. Their fixed components have the same image in $L/P^L_{J'}$ exactly when $dW_{J'} = eW_{J'}$. Indeed, their base points are B_K-fixed, and a partial flag variety of L_K has a unique B_K-fixed point. On such a branch the step is the usual partial-flag correspondence for the standard parabolic $L_K \cap {}^{\dot{d}}P^L_{J'} = L_K \cap {}^{\dot{e}}P^L_{J'}$. Consequently a fixed gallery is an actual iterated gallery for L_K, with these successive partial types. A fixed initial component must also be preserved by a_0. After choosing a compatible stable Borel, its minimal representative satisfies $\theta(d) = d$; a θ-fixed Tits lift identifies its initial condition with the actual odd fixed-flag condition for L_K. This identification transports the original monodromy marking and component projector.

The final map to the extremal component $F = L_K/P^{L_K}_{H_K} \subset L/P^L_H$ retains the actual partial boundary. For a last full-flag component it requires $dW_H = W_H$, and the final map is the projection to $L_K/P^{L_K}_{H_K}$. Thus the neutral cross-Hom has the partial push appearing in (3.24), including its middle averaging factor. The same description applies on each component of a disconnected initial fixed locus. □

Lemma 3.13. *The support filtration on either endpoint gives relative source generators of nonnegative λ-weight. Against the target of Lemma 3.11, the weight-zero Hom uses the galleries of Lemma 3.12 with their derived stability equations.*

Proof. Apply the opposite Białynicki–Birula support filtration of the smooth projective gallery X. Take the derived pullback along $F \to L/P_H$ on each term of this support filtration. A term meeting $f^{-1}(F)$ has its fixed component over F, by equivariance. The endpoint normal functions have the strictly positive weights E_1, E_2 of (3.17). Their pullbacks, their Koszul generators, and their even comparison generators from (3.19) therefore have positive weight.

On an opposite BB cell, the ordinary relative coordinate functions have positive weight. Write z_i for normal functions to that cell, of weights $-a_i < 0$. Their infinitesimal-stability exterior generators ξ_i also have weights $-a_i$. The supported normal monomials consequently have relative weights

$$\sum_i a_i(n_i - e_i) + \mathrm{wt}(m) \geq 0, \qquad n_i \geq 1, \quad e_i \in \{0, 1\}, \tag{3.21}$$

where m is a monomial in the remaining nonnegative relative generators. The unique neutral normal monomial is $\prod_i \xi_i/z_i$. Its support and exterior degrees cancel, and so do its full determinant characters. The remaining neutral equations are exactly the infinitesimal-stability equations on the fixed gallery described above. Thus its underlying gallery is smooth and its incidence algebra contains the full derived equations on that gallery. The contractions in (3.19) remove regular coordinate pairs.

Over the full radical base, a neutral relative residue need not be a cocycle. For example, $d\xi_z = v$ of weight $-a$ gives $d(\xi_z/z) = v/z$. This term is retained in the source differential. The following elementary observation explains precisely why it does not affect the neutral Hom. If a strict free-cell source has relative generator weights at least zero and the target has weights at most zero, a weight-zero map kills every positive generator. Terms from a neutral generator to a positive generator are thereby killed. A term from a positive generator to a neutral generator has a positive-weight coefficient, which also annihilates the weight-zero target. Hence

the induced weight-zero Hom differential is the differential of the neutral core, even when the neutral source generators do not span a subcomplex.

To carry this out in rational equivariant modules, fix an ordinary Hom degree. Lemma 3.10 bounds the number of negative-degree generators. For a fixed remaining weight bound, only finitely many monomials in strictly positive-weight ordinary coordinates or additional poles can contribute. Thus taking the weight-zero Hom here creates no unrestricted product of localized representations. The BB filtration is finite and its residue identifications glue with their determinant transformations. This gives a finite filtration whose source generators have nonnegative relative weights and whose neutral pieces compute marked L_K-gallery coefficients. The full radical is retained until Lemma 3.14 is applied to the entire equivariant proper push. □

An endpoint stratum need not be stable under the full common radical. Having computed the weights of the underlying λ-module, we can now take radical invariants on the whole equivariant proper push.

Lemma 3.14. *Let $R = H_0 \ltimes U$, where U is a unipotent algebraic group in characteristic zero, and let $\lambda : \mathbb{G}_m \to Z(H_0)$ act on $\mathfrak{u}$ with strictly positive weights. Let E be the actual complex of rational R-modules. Assume $H^q(E)$ has weights at most zero for every q. Then the natural edge map gives*

$$\big(R\Gamma(U,E)\big)_0 \simeq E_0, \qquad R\Gamma(R,E)_0 \simeq R\Gamma(H_0,E_0). \tag{3.22}$$

Here a subscript zero denotes the exact λ-weight-zero functor; H_0 is not required to be reductive.

Proof. One may see the rational unipotent comparison directly. The characteristic-zero exponential identifies U with an affine space. Its polynomial de Rham complex resolves the constant functions, by the Euler contracting homotopy. Tensor with a rational module V. In left Maurer coordinates each term, with the diagonal action untwisted, is a cofree rational U-module $\mathcal{O}(U) \otimes W$, hence injective. Taking invariants and evaluating invariant forms at the identity gives $V \otimes \Lambda^\bullet \mathfrak{u}^*$ with its usual Chevalley–Eilenberg differential. This construction is functorial in V and has width $\dim U$, so it also computes complexes by totalization. This is the rational group/Lie-cohomology comparison of Hochschild [Hoc61]; see also [Kas15, Section 2, Theorem 2.1]. The polynomial de Rham construction works over the characteristic-zero coefficient field used here.

Filter the resulting complex by exterior degree. Its finite-width spectral sequence has first page

$$H^q(E) \otimes \Lambda^p \mathfrak{u}^*, \qquad 0 \le p \le \dim U. \tag{3.23}$$

For $p > 0$ every weight is strictly negative. For $p = 0$ the weight-zero term is exactly $H^q(E)_0$. Since the filtration has finite width, there is no convergence problem even when the ordinary Hom is unbounded above. The weight-zero spectral sequence has only column zero, and the projection to zero cochains is a quasi-isomorphism on weight zero. This is natural for H_0 and for all endpoint actions. Taking H_0 derived invariants proves the second assertion. Finite component-group invariants may then be taken without loss in characteristic zero. □

Lemma 3.15. *Let $(C, \mathfrak{m}, k)$ be a Noetherian local ring and let P have bounded finite cohomology. If $P \otimes^L_C k \in D^{\le 0}$, then $P \in D^{\le 0}$.*

Proof. If $q > 0$ is the largest index with $H^q(P) \neq 0$, the top edge of the hyper-Tor spectral sequence gives $H^q(P \otimes^L_C k) = H^q(P) \otimes_C k$. Indeed the incoming differentials to the zero-Tor term would require a higher cohomology row, while outgoing differentials would require a positive cohomological Tor column. Equivalently this follows by tensoring the last truncation triangle and using right t-exactness of derived tensor with a module. Nakayama makes the right side nonzero, a contradiction. □

Proof of Theorem 3.4. Apply Lemmas 3.5 and 3.6 to the end words. Filter Z^2 by its finitely many diagonal P_J-orbits. Lemma 3.7 gives the common square incidence on every contributing orbit. There Lemma 3.9 computes the source and Lemma 3.11 gives the supported dualizing target. Each strictly positive source generator has strictly negative weight after duality. Functions on the common unipotent radical have nonpositive weights as well. The cohomology of the Hom on the orbit therefore has no positive weights. Its weight-zero terms come from the endpoint galleries by Lemma 3.13. The bounds of Lemma 3.10 and the pole estimate (3.20) apply in each mapping degree. Apply Lemma 3.14 to the entire equivariant push on this orbit to take the radical invariants.

For a fixed pair of endpoint galleries, let M'_K, N'_K be the two L_K-word translates of Lemma 3.12. Their coefficient is

$$R\operatorname{Hom}_{X^K_{H_K}}(Rp^K_{H_K *}M'_K, Rp^K_{H_K *}N'_K). \tag{3.24}$$

Moving the end words across Hom as in (3.9) gives a word in L_K with one C_{H_K} in the middle. Each connected factor of L_K has its roots inside I. The hypothesis of Theorem 3.4 on these factors makes the coefficient before

the C_0-fiber a module D_w in degree zero, with its inherited central square, torus, Cartan markings and finite summand. Hence in the derived fiber its contribution is

$$D_w \otimes^L_{C_0} k \in D^{\leq 0}(k). \tag{3.25}$$

The finite target-tensor identity retains every negative Tor of D_w. The support triangles for the middle orbits and endpoint galleries preserve $D^{\leq 0}$. This proves (3.7).

Apply Lemma 3.15 to (3.7). Proposition 3.1 independently excludes negative ordinary cohomology, so the original source coefficient is ordinary at the chosen central local ring. Flat scalar completion preserves that assertion. The same computation can be made at every Artin quotient $C_0/(t_1^m, \dots, t_d^m)$ of a regular parameter system: the parameter sequence still gives a Koszul resolution of width d and weight zero, so its derived tensor with D_w again lies in $D^{\leq 0}$.

At a noncentral norm b, Proposition 3.24 and (3.50) identify the localized coefficient with both markings with a sum of the stated coefficients in $Z_G(b)^\circ$. By the second hypothesis these too are ordinary. Formula (3.6) now detects positive-degree vanishing on the whole output torus. Lemma 3.3 gives finite projectivity. □

3.2. **Ordinary completed summands and marked rank-one identities.** We pass from geometric averaging words to a finite set of canonical summands, then record the rank-one identities that propagate their coefficient vanishing. All complexes are unperiodized. A grading shift of a polynomial Soergel bimodule will be forgotten; a cohomological shift will never be forgotten. This distinction is essential in using positivity to deduce vanishing in a fixed ordinary degree.

3.2.1. *A central completion on which coefficient projectors act.* Put $R = \mathcal{O}(T)$ and $C_0 = R^W$. The common even holonomy makes C_0 central in the finite boundary Hecke category. The original odd output algebra $A = \mathcal{O}(T/(1-\theta)T)$ is generally not central in that monoidal category. Its map from C_0 is the original norm, including the fixed marking translation. That norm is finite onto its image: on the appropriate invariant torus it is an isogeny onto a subtorus, followed by the finite Weyl quotient. Thus A is finite as a C_0-module; a kernel of $C_0 \to A$ causes no problem.

Let $\mathcal{T}$ be the ordinary additive category of finite universal tiltings. For a maximal ideal $\mathfrak{m} \subset C_0$ set $C = \widehat{(C_0)}_{\mathfrak{m}}$ and define the scalar extension by $\mathrm{Hom}_{\mathcal{T}_C}(X,Y) = \mathrm{Hom}_{\mathcal{T}}(X,Y) \otimes_{C_0} C$. For each j, the actual coefficient functor $F_j(X) = H^j\mathcal{N}(X)$ extends C-linearly. It then extends to $\mathrm{Kar}(\mathcal{T}_C)$ by taking images of projectors in the idempotent-complete category of finite $A \otimes_{C_0} C$-modules. Consequently a completed retract of an object with $F_j = 0$ also has $F_j = 0$. No completed projector has to descend to the Laurent category.

The finite algebra decomposes as

$$A \otimes_{C_0} C = \prod_{y \mapsto \mathfrak{m}} \widehat{A}_y. \tag{3.26}$$

This retains every actual marked factor and its full adic structure. Scalar extension is flat, so it does not change an ordinary degree. One may first prove vanishing over C and then project to a factor of (3.26); this does not require the two odd Cartan actions to be equal.

3.2.2. *Canonical summands and positive direct sums.*

Proposition 3.16. *At a completion whose even semisimple value lies in $Z(G)$, the category of unshifted tilting words admits the canonical summands of polynomial Soergel theory. Forgetting internal grading gives ordinary direct sums with nonnegative multiplicities. In particular the simple-product formula below holds without any cohomological shift.*

Proof. Work first with semisimple simply connected dual group and add the central torus by Proposition 3.24, so that the constructible group has connected center. Taylor's ordinary Soergel functor identifies the Bott–Samelson tilting word category with the category of multiplicative Bott–Samelson bimodules $R \otimes_{R^{s_1}} R \otimes_{R^{s_2}} \cdots \otimes_{R^{s_m}} R$. Its full faithfulness is an assertion about Hom^0; the derived Hom between tilting words is concentrated in degree zero. We use the ordinary full-faithfulness and generation theorems of Taylor [Tay25, Section 8.1, Propositions 8.3 and 8.6, Theorem 9.4]. Thus no periodic comparison is being used at this step.

Suppose the even value b is central in G. Its Weyl orbit is a singleton, so $R \otimes_{C_0} C = \widehat{R}_b$. The characteristic-zero formal logarithm at b is W-equivariant and gives

$$\widehat{R}_b \simeq \widehat{P}_0, \qquad P = \mathrm{Sym}(\mathfrak{t}^*), \qquad C \simeq (\widehat{P}_0)^W.$$

Finite-group invariants and finite-module Hom commute with this flat completion. In particular the actual simple bimodule becomes

$$\widehat{P}_0 \otimes_{(\widehat{P}_0)^s} \widehat{P}_0. \tag{3.27}$$

The same holds for parabolic kernels, products, morphisms and idempotents.

Let B_w be the polynomial indecomposable Soergel bimodule. Its graded character is the Kazhdan–Lusztig character in characteristic zero [KL79, Soe07] and [EW14, Theorem 1.1 and Corollary 1.2]. A finite graded word has a decomposition

$$B(\underline{s}) = \bigoplus_{w,n} B_w(n)^{\oplus m_{w,n}}, \qquad m_{w,n} \geq 0.$$

Forget the internal grading *before* completing. Every $B_w(n)$ is now the same ordinary bimodule B_w, and (3.27) realizes the decomposition as an actual unshifted direct sum in $\mathrm{Kar}(\mathcal{T}_C)$. Indecomposability after completion is not needed. The multiplicity is $\sum_n m_{w,n}$.

In particular, for a simple reflection s,

$$C_s B_w \simeq \begin{cases} B_w^{\oplus 2}, & sw < w, \\ B_{sw} \oplus \displaystyle\bigoplus_{v<w,\, sv<v} B_v^{\oplus \mu(v,w)}, & sw > w, \end{cases} \qquad \mu(v,w) \in \mathbb{Z}_{\geq 0}. \tag{3.28}$$

There is an identical right-sided formula. Every summand below the head in an ascent has length strictly less than $\ell(w)$, even when its Kazhdan–Lusztig coefficient comes from a large Bruhat gap.

□

3.2.3. *The marked rank-one identities.* Pass to the central cover with simply connected derived group. For a fixed simple node s, work in its standard Levi $L = L_s$. The stable-Levi embedding sends its odd Borel object Q_L to Q_G, intertwines C_s, and preserves the original norm and mark. Lemma 3.33 constructs the functor on objects of zero weight for the defining cocharacter. The argument applies to a fixed node also when the involution acts nontrivially on other roots. For a nontrivial involution on the Levi center, first use the central cover of (3.53) and the product formula of Section 3.4.5; the rank-one calculation is on the derived factor, with that central norm base change retained.

On the semisimple factor assume $\theta = 1$ and write $a = t\tau$, $\tau^2 = z$, and $y = tr_G$, where $r_G^2 = z$. Over a central even value b the raw branches satisfy $t^2 = z^{-1}b$. The original completed object splits as $\widehat{Q} = \bigoplus_y Q_y$, and

$$s \bullet y = s(t) r_G. \tag{3.29}$$

The root value $\alpha_s(t)$ is $+1$ or -1. A white value is $+1$; a black value is -1. All raw lifts, including spin or half-weight coordinates, remain in this definition.

Lemma 3.17. *On this cover, at a central even value, the action of the fixed simple kernel on a marked closed object satisfies* (3.30) *on the white branches and* (3.31) *on the black branches. These are isomorphisms of ordinary complexes over the completed even central ring C.*

Proof. Let p forget the rank-one flag in the square-stable parameter, and let ρ be the odd-stable Borel map to the odd sector. Then $C_s = p^* R p_*$ and Q is the actual push of the odd-stable locus. If $\alpha_s(t) = 1$, odd stability and square stability agree on the entire completed sector. One way to see the scheme and derived assertion is to use the square map on the genuine rank-one group: on its derived factor its derivative in right-translation coordinates is $1 + \mathrm{Ad}(a)$. On the sector with central semisimple part and arbitrary unipotent part it is invertible in characteristic zero. An invariant flag for a^2 is then a flag for its uniquely specified odd square root. The ordinary rank-one Grothendieck push is the rank-two algebra $R\pi_*\mathcal{O} = \mathcal{O}_L \otimes_{R^s} R$. Choose ω with $\langle \omega, \alpha_s^\vee \rangle = 1$. Set $u = e^\omega$, $D = \mathbb{C}[\ker(\langle -, \alpha_s^\vee \rangle : X^*(T) \to \mathbb{Z})]$ and $q = e^{2\omega - \alpha_s} \in D^\times$. Then $s(u) = q/u$, which gives $R = D[u^{\pm 1}]$, $R^s = D[u + q/u]$, with basis $1, u$. Proper *derived* base change, including all unipotent fibers, therefore yields an ordinary object isomorphism

$$C_s Q_y \simeq Q_y^{\oplus 2} \qquad \text{(white)}. \tag{3.30}$$

The frame need not diagonalize the odd input Cartan action.

If $\alpha_s(t) = -1$, the two marked points t and $s(t)$ are distinct: $s(t) = t\alpha_s^\vee(-1)$ and the coroot is primitive in the simply connected derived form. In the rank-one odd regular semisimple sector, choosing either eigenvalue selects a unique stable line. The eigenvalue difference is a unit, so these lines are subbundles in families. Each

completed marked Borel branch maps isomorphically to the same odd sector, with its full torus stabilizer. Their partial pushes agree; pulling back under p gives

$$C_sQ_y \simeq C_sQ_{s\bullet y} \qquad \text{(black)}. \tag{3.31}$$

The averaging functor remains on both sides. An adjoint rank-one group can have an extra Weyl component in the stabilizer; (3.31) is first proved in the simply connected form and the vanishing of coefficient cohomology descends by Lemma 2.26.

All maps keep the same even holonomy. Hence (3.30)–(3.31) are C-linear ordinary isomorphisms. Passing to the ambient group through the genuine stable-Levi functor proves the required fixed-node identities there. □

Proposition 3.18. *On the cover of Lemma 3.17, the white and black identities* (3.30)–(3.31) *hold at each mark over a central even value. They are natural under endomorphisms of the middle Hecke kernel and therefore give the separate endpoint identities* (3.32).

Proof. Set $\mathcal{N}_{y,z}(X) = \mathrm{RHom}(XQ_z, Q_y)$. Right endpoint operations act on z; ordinary self-adjunction of C_s moves a left endpoint to the target and acts on y. Thus Lemma 3.17 gives

$$\begin{aligned}
&\mathcal{N}_{y,z}(C_sX) \simeq \mathcal{N}_{y,z}(X)^{\oplus 2} && (s \text{ white at } y),\\
&\mathcal{N}_{y,z}(XC_s) \simeq \mathcal{N}_{y,z}(X)^{\oplus 2} && (s \text{ white at } z),\\
&\mathcal{N}_{y,z}(C_sX) \simeq \mathcal{N}_{s\bullet y,z}(C_sX) && (s \text{ black at } y),\\
&\mathcal{N}_{y,z}(XC_s) \simeq \mathcal{N}_{y,s\bullet z}(XC_s) && (s \text{ black at } z).
\end{aligned} \tag{3.32}$$

These are natural in the middle X. For $e \in \mathrm{End}(X)$ they intertwine the induced endomorphism C_se (or eC_s), so they extend to the image of every completed projector *of the middle* X. This does not assert compatibility with arbitrary endomorphisms of the larger word C_sX. For a descent we first put $X = B_w$ in (3.32), then use the separate object identity $C_sB_w \simeq B_w^{\oplus 2}$; an isomorphism of two-copy coefficient objects is enough to transport vanishing. A white ascent uses the genuine direct summands in (3.28). No coherent simultaneous choice of all the rank-two frames is necessary to propagate vanishing.

Equations (3.28) and (3.32) apply separately to every ordinary H^j. For a descent, the black relation transports the vanishing predicate between raw marks. For a white ascent the positive head and its lower summands are genuine direct summands. These are exactly the operations used by the finite closure arguments below. □

3.3. **Vanishing at a marked diagonal.** This section proves the geometric vanishing used to anchor the finite argument. The result holds for every root system. Its strict inequality comes from the last omitted letter of a reduced expression; neither the rank nor the length of the expression enters the argument.

3.3.1. *Statement and canonical normalization.* Work over an algebraically closed field k of characteristic zero. Fix a connected reductive group G and a Borel pair $T \subset B$. Let $K = Z_G(t)^\circ$ for a semisimple element $t \in T$, and put $B_K = B \cap K$. Write $\mathfrak{u}$ and $\mathfrak{u}_K$ for the nilpotent radicals of their Lie algebras, and set $n = \dim G/B, \qquad n_K = \dim K/B_K, \qquad r = n - n_K$. Let $\mathcal{Z}_w \subset \mathfrak{b} \times G/B$ be the closure of the Grothendieck correspondence of relative position w, with the first flag fixed. The first Cartan map is $\mathcal{Z}_w \to \mathfrak{t}$. Define the scheme-theoretic fiber and the two incidence schemes

$$\begin{aligned}
A &= \mathfrak{u} \times G/B, & S_w &= \mathcal{Z}_w \times_{\mathfrak{t}} \{0\},\\
I_G &= \{(x,F) \in A : x \in \mathfrak{b}_F\}, & A_K &= \mathfrak{u}_K \times K/B_K,\\
Y_K &= \{(x,F) \in A_K : x \in \mathfrak{b}_{K,F}\}.
\end{aligned}$$

There is a closed map $i : Y_K \to I_G$. All actions are the actual B_K actions: conjugation on x and translation on F.

The full graphs are Cohen–Macaulay and their Cartan maps are flat [BHS19, Proposition 2.3.3]. Consequently S_w is Cohen–Macaulay of dimension n. Its normalized canonical source is the ordinary sheaf

$$\mathcal{K}_w = \mathcal{E}xt_A^n(\mathcal{O}_{S_w}, \mathcal{O}_A) \otimes \omega_{G/B} = \omega_{S_w} \otimes \det \mathfrak{u}. \tag{3.33}$$

In particular this definition includes its equivariant character. The reduction of S_w is not substituted for its scheme-theoretic fiber. The whole incidence I_G may have a nonzero nilpotent second Cartan even though its first Cartan is zero.

Theorem 3.19. *If $w \notin W_K$, then*

$$R\Gamma([I_G/B_K], R\mathcal{H}om_{I_G}(\mathcal{K}_w, i_*\mathcal{O}_{Y_K})) = 0. \tag{3.34}$$

The assertion includes every ordinary cohomological degree of the incidence Hom complex.

We use two published results about full Grothendieck graphs. Besides the flatness just cited, reduced convolution of their simple structure sheaves gives the full graph structure sheaf [BR12, Theorem 2.2.1 and Remark 2.2.2(2)]. For a reductive group these results follow from the semisimple case by splitting the central Lie algebra and using a characteristic-zero central isogeny. Such an isogeny identifies the flag varieties and the semisimple Lie algebras. No assertion about an intrinsic reduced nilpotent component is needed.

3.3.2. *The gallery and its strict weight inequality.* Choose a reduced expression $w = s_1 \cdots s_m$, with simple roots α_i, and let $f : X \to G/B$ be its Bott–Samelson flag gallery. The first flag is fixed. Let $\mathfrak{b}_i$ denote the Borel bundle at the ith flag. Over an adjacent pair of flags the simple full graph is a vector bundle $V_i \subset \mathfrak{b}_{i-1}$ of codimension one, also mapping into $\mathfrak{b}_i$. For example, in the lower flag chart of SL_2 its equation is $2a + bz = 0$; the entire incidence has equation $z(2a + bz) = 0$. At a repeated coordinate flag the quotient $Q_i = \mathfrak{b}_{i-1}/V_i$ is a Cartan line of character zero. At an applied simple reflection its character is the transported simple root.

The projections of this bundle to the endpoint flag varieties are smooth. Thus intermediate flags may be matched ordinarily, while the equalities of their holonomies must be imposed derivedly. The eigenspaces of $\mathrm{Ad}(t)$ give a B_K-stable splitting $\mathfrak{u} = \mathfrak{u}_K \oplus \mathfrak{u}_\perp$. All the holonomy matches, the first Cartan equation, and the restriction of the first holonomy to $\mathfrak{u}_K$ form the linear bundle map

$$L : V = \bigoplus_{i=1}^{m} V_i \longrightarrow E = \bigoplus_{i=1}^{m-1} \mathfrak{b}_i \oplus \mathfrak{t} \oplus \mathfrak{u}_\perp. \tag{3.35}$$

Its derived zero scheme has the Koszul model $\mathrm{Sym}(V^*) \otimes \Lambda(E^*)$. Give both polynomial variables and exterior generators holonomy degree one. Each holonomy degree h is the finite bundle complex

$$\bigoplus_{p=0}^{\min(h,\mathrm{rk}\,E)} \mathrm{Sym}^{h-p}(V^*) \otimes \Lambda^p(E^*), \qquad \text{ordinary degree } -p. \tag{3.36}$$

For the regular embedding $j : K/B_K \hookrightarrow G/B$, set $D_j = R\mathcal{H}om(j_*\mathcal{O}_{K/B_K}, \mathcal{O}_{G/B}) = j_* \det N_j[-r]$. Tensor (3.36) with Lf^*D_j and call the result $\widetilde{C}_h$. The full derived convolution theorem and projection formula identify $Rf_*\widetilde{C}_h$ with the hth holonomy part of the affine holonomy pushforward of

$$C = (\mathcal{O}_{S_w} \otimes^L_A \mathcal{O}_{A_K})[-r] \otimes \det N_j. \tag{3.37}$$

In particular all higher direct images of the gallery and all negative Tor groups of the intersection remain in this model.

Fix $v \in W_K$, a strictly dominant integral cocharacter λ_0 for G, and $\lambda = v\lambda_0$. The opposite flag chart admits a T-stable product decomposition

$$\begin{aligned} vU^-B/B &= U_v \times D_v, & f^{-1}(vU^-B/B) &= U_v \times X', \\ U_v &= U_K \cap vU^-v^{-1}, & X' &= f^{-1}(D_v). \end{aligned} \tag{3.38}$$

One obtains the first product by ordered root coordinates in the unipotent group vU^-v^{-1}; the second follows from equivariance of f. In these coordinates the holonomy changes to $\mathrm{Ad}(u^{-1})x$. The variety X' is smooth. Actual U_v-equivariant sheaves descend through this free product; no action of the whole stabilizer on the slice is assumed.

Every flag tangent weight at v is negative for λ. Consequently the λ-attracting set of v in G/B is its single point, and $Z_v = f^{-1}(v)$ is a union of whole attracting cells in X. The Białynicki–Birula decomposition of the smooth projective gallery is filtrable [BB76, Theorem 3], so it gives a finite filtration of Z_v. In the smooth slice X', each successive cell is a regular closed subvariety after the previous support strata have been removed.

For a coordinate mask ϵ ending at v, put

$$a_i = s_1^{\epsilon_1} \cdots s_i^{\epsilon_i}, \qquad q_i = \langle \lambda, a_{i-1}\alpha_i \rangle, \qquad D = \sum_{\substack{\alpha \in \Phi_K^+ \\ \langle \lambda, \alpha \rangle < 0}} |\langle \lambda, \alpha \rangle|.$$

The step tangent character is $a_{i-1}\alpha_i$ at an applied letter and its negative at an omitted letter. Removing the free orbit U_v deletes precisely the negative tangent weights whose total magnitude is D. The total magnitude of the remaining negative normal weights to the attracting cell is therefore

$$s'_\epsilon = \sum_{\substack{i \ applied \\ q_i<0}} |q_i| + \sum_{\substack{i \ omitted \\ q_i>0}} q_i - D. \tag{3.39}$$

At the fixed mask, telescoping the Borel bundles in (3.35) gives the identity of virtual T-characters

$$[\operatorname{coker} L] = \sum_i [Q_i] - [\mathfrak{u}_K] + [\ker L]. \tag{3.40}$$

Here Q_i has character $a_{i-1}\alpha_i$ for an applied letter and zero for an omitted one. A common compatible holonomy in $\ker L$ belongs to $\mathfrak{u}_K \cap \operatorname{Ad}(v)\mathfrak{u}$. Its λ-weights are strictly positive. Thus all weights of $\operatorname{Sym}((\ker L)^*)$ are nonpositive, while the largest weight in the entire exterior algebra of $(\operatorname{coker} L)^*$ is

$$P_\epsilon = \sum_{\substack{i \ applied \\ q_i<0}} |q_i| - D. \tag{3.41}$$

The use of the whole exterior algebra makes the bound independent of the holonomy degree.

Every mask ending at $v \in W_K$ is proper if $w \notin W_K$. For its last omitted position j, all later letters are applied, and $v^{-1}a_{j-1}\alpha_j = s_m \cdots s_{j+1}\alpha_j \in \Phi_G^+$. Positivity is the reduced-word criterion for the reversed word. It implies $q_j > 0$, and hence the strict inequality

$$s'_\epsilon - P_\epsilon = \sum_{\substack{i \ omitted \\ q_i>0}} q_i > 0. \tag{3.42}$$

We next justify using this fixed-point calculation for supported complexes. A T-invariant affine gallery chart about a fixed mask contains its whole attracting cell. Indeed, its invariant closed complement cannot contain a point contracting to the mask. Lift eigenbases of the bundle fibers to homogeneous sections. Their failure locus is invariant and avoids the fixed point; its ideal contains a weight-zero function nonzero there. Inverting that function makes these sections frames and preserves the whole cell. The same operation permits removal of the preceding support strata. A maximal-rank minor of L at the fixed point uses matching domain and codomain characters, so its determinant has weight zero. After inverting it, cancel this invertible linear Koszul block. The residual bundle fibers are $\ker L$ and $\operatorname{coker} L$. This cancellation preserves holonomy degree and introduces neither a determinant character nor an ordinary shift; constant rank away from the fixed point is unnecessary.

The dual flag Koszul complex Lf^*D_j has nonpositive fixed-fiber weights, since its generators are tangent-normal characters at v. On an attracting cell D_ϵ, homogeneous Nakayama gives frames with the fixed-fiber weights for the residual bundles. Functions on the cell have nonpositive weights. If N_ϵ is its normal bundle in the smooth filtration open, local cohomology of a bundle E_0 along D_ϵ has an exhaustive pole filtration with successive pieces $E_0|_{D_\epsilon} \otimes \det N_\epsilon \otimes \operatorname{Sym}^a N_\epsilon$, $\quad a \geq 0$, in its normal codimension. The determinant weight is $-s'_\epsilon$, and the additional symmetric powers have nonpositive weights. Combining (3.41) and (3.42), every such supported term has weight at most $P_\epsilon - s'_\epsilon < 0$. The finite cell filtration and the exhaustive pole filtrations give

$$\operatorname{wt}_\lambda H^a R\Gamma_{Z_v}(X', \widetilde{C}_h|_{X'}) < 0 \qquad (a \in \mathbf{Z}, h \geq 0). \tag{3.43}$$

All filtrations here are in rational torus modules. Proper pushforward with support transfers the result to (3.37). Affine holonomy pushforward is exact, and algebraic direct sums in h commute with the finite Cech support complexes. The argument also permits tensoring with any finite bundle complex whose fixed-fiber weights are nonpositive: the same inequality holds on each pole-filtered cochain term.

3.3.3. *Incidence Ext and the full stabilizer.* The incidence $I_G \subset A$ is the zero scheme of the flag vector field, with equation bundle $\Omega_{G/B}$. Above the relative position cell v, the holonomy fiber has dimension $n - \ell(v)$. Thus every stratum has dimension n, and the n equations have the expected codimension in the smooth $2n$-dimensional ambient. They are a regular sequence. The same argument applies to $Y_K \subset A_K$.

Ambient perfect biduality and (3.33) give

$$\begin{aligned} R\mathcal{H}om_A(\mathcal{K}_w, \mathcal{O}_{Y_K}) &= (\mathcal{O}_{S_w} \otimes_A^L \mathcal{O}_{Y_K})[-n] \otimes \omega_{G/B}^{-1} \\ &= C \otimes_{A_K}^L \mathcal{O}_{Y_K}[-n_K] \otimes \omega_{K/B_K}^{-1}. \end{aligned} \tag{3.44}$$

The last finite factor is the dual incidence Koszul complex for K/B_K; its terms are $\Lambda^d T_{K/B_K}$, $0 \le d \le n_K$. At v all its generators have negative λ-weight. Consequently (3.43) applies to this ambient Hom after support on the slice.

To recover the Hom over I_G, take a T-invariant affine chart and write $R = Q/(f_1, \dots, f_n)$, with Q regular and f_i the actual regular incidence sequence. Its equation characters are $\chi_i = v\beta_i$, $\beta_i \in \Phi_G^+$, and are all positive for λ. Write $M = \mathcal{K}_w$ and $N = i_*\mathcal{O}_{Y_K}$ on the chart. Both are ordinary finite R-modules. Choose a finite equivariant projective Q-resolution $P_Q \to M$ in nonpositive degrees. Such a resolution follows from homogeneous finite generation and regularity: its last ordinary-projective syzygy is equivariantly projective by taking the weight-zero part of an ordinary splitting.

Homogeneous higher homotopies for the f_i give the Eisenbud–Shamash resolution. After Hom into N, its underlying module and differential have the form

$$\mathrm{Hom}_Q(P_Q, N) \otimes k[\theta_1, \dots, \theta_n], \qquad |\theta_i| = 2, \quad \mathrm{wt}(\theta_i) = -\chi_i, \quad d = d_0 + \sum_{a \neq 0} \theta^a d_a. \tag{3.45}$$

See [AG02, Theorem 4.1, Remark 4.4, and Appendix Remark A.2]. The projective version is obtained by the same higher-homotopy construction; homogeneous choices are made by exact weight projection. The decreasing filtration by operator degree has associated graded the ambient Hom tensored with the corresponding polynomial operator space. All higher-homotopy differentials in (3.45) are retained.

Both $\mathrm{Hom}_Q(P_Q, N)$ and a finite Cech support complex have nonnegative ordinary degrees. Therefore only operator degrees $2|a| \le j$ occur in total degree j. The filtration is degreewise finite after support, so its spectral sequence converges without completion or an infinite product. Every operator has negative λ-weight. We conclude that the true supported incidence Hom on the slice has strictly negative weights in every cohomological degree.

We record the elementary rational-cohomology fact used for descent. For a connected unipotent group U in characteristic zero and an arbitrary rational U-module V, rational group cohomology is computed by $C^\bullet(\mathfrak{u}, V)$; for $H = T \ltimes U$ one takes its T-invariants. This is the characteristic-zero comparison of Hochschild [Hoc61]. One may verify it directly as follows. Lie cohomology has degree-zero functor $V \mapsto V^U$. Embed V into a cofree rational module $k[U] \otimes V$ by its coaction. The Lie-cochain complex of a cofree module is the algebraic de Rham complex of U tensored with the underlying vector space. Since $U \simeq \mathbf{A}^{\dim U}$ in characteristic zero, the polynomial Poincare homotopy makes its positive cohomology zero. Thus Lie cohomology is an effaceable delta functor and computes the derived invariants. Exactness of torus invariants gives the assertion for H. The construction commutes with filtered colimits and has cohomological dimension at most $\dim U$.

Let $C_v = B_K v B_K/B_K$ and $H_v = B_K \cap v B_K v^{-1}$. Work in a B_K-stable flag open where C_v is closed, and let J be its ideal. After affine holonomy pushforward, any cohomology sheaf $\mathcal{M}$ of the coefficient supported on C_v is rational, quasicoherent and J-torsion. The filtration $F_a\mathcal{M} = \mathrm{ann}_{\mathcal{M}}(J^a), \qquad a \ge 1$, is exhaustive and its successive quotients are actual B_K-equivariant sheaves on C_v. On (3.38) the supported complex is pulled back from the slice, and its cohomology is $\mathcal{O}_{U_v} \boxtimes M'$. This assertion holds also for the unbounded incidence Hom: resolve its ordinary slice source by finite-rank projectives in each nonpositive degree, pull back flatly, and take Hom into its ordinary target term by term. The finite support complex introduces only finitely many terms in each total degree.

The ideal J is pulled back from the slice ideal. Restriction to the identity section of U_v is exact on these pulled-back objects, so the fiber of F_a/F_{a-1} at v is a T-subquotient of M' and has strictly negative weights. Equivariant descent on $C_v = B_K/H_v$ identifies this fiber with its actual rational H_v-representation. The unipotent radical of H_v has roots $\Phi_K^+ \cap v\Phi_K^+$, all positive for λ. Its dual Lie-cochain factors have negative weights. Hence the Lie-cochain complex tensored with the fiber has no weight zero, and $R\Gamma([C_v/B_K], F_a/F_{a-1}) = 0$. Pass through the filtration and its filtered colimit. Lie-cochain width and finite affine-cover width bound the equivariant cohomological dimension, so the cohomology-sheaf spectral sequence also converges for the whole bounded-below Hom complex. Its contribution supported on C_v is zero. Finally the finitely many B_K Bruhat cells filter K/B_K. For every contributing cell the endpoint lies in W_K and its mask is proper. The support triangles prove Theorem 3.19.

3.3.4. *The multiplicative marked fiber.* The following formal comparison is used only to identify the schematic source and its canonical specialization. The weight arguments above take place on algebraic fibers, never on an arbitrary completed coordinate ring.

Lemma 3.20. *Exponential and logarithm identify the completion of $\mathfrak{b} \to \mathfrak{t}$ along $\mathfrak{u}$ with the completion of $B \to T$ along U. This identification is conjugation equivariant, preserves flag-stability incidences over nonreduced bases,*

and identifies completed full graph closures. The normalized canonical module of a multiplicative full graph has derived first-Cartan fiber (3.33).

Proof. Choose a faithful representation with B upper triangular. For the additive and multiplicative Cartan ideals I and J, the universal matrices satisfy $X^N \in I\,\mathrm{End}(V)$ and $(g-1)^N \in J\,\mathrm{End}(V)$ for a fixed N. Consequently

$$\exp X = \sum_{a\geq 0} X^a/a!, \qquad \log g = \sum_{a\geq 1} (-1)^{a+1}(g-1)^a/a$$

converge in these Cartan-adic topologies, with no completion in unipotent coordinates. Membership of $\exp X$ in B can be checked on the polynomial defining relations of B in matrices, after clearing powers of the determinant. For a relation of degree d, each expansion term contains at most d matrix-power factors. If their total power tends to infinity, one factor does also, and $X^N \in I$ makes these terms tend to zero I-adically. The identities of the ordinary formal exponential therefore remain valid under this substitution. This justifies evaluation without asserting continuity for arbitrary formal series in all matrix entries.

The resulting map induces the formal exponential on the Cartan and the polynomial exponential $\mathfrak{u} \simeq U$ on the closed fiber. Its left-translated differential is $(1-\exp(-\operatorname{ad} X))/\operatorname{ad} X$. On the closed fiber $\operatorname{ad} X$ is nilpotent, so this is invertible. Formal etaleness and the closed-fiber isomorphism identify the Cartan completions; the convergent logarithm is the inverse. Conjugation equivariance is immediate from the series. For a flag stabilizer, use a representation in which it stabilizes a flag or a line: polynomials in a preserving matrix and their convergent sums preserve it. The inverse has the same property. Thus the incidence functors agree over nonreduced bases, including their nilpotent second-Cartan coordinates.

The flag coordinate is unchanged. Schematic closure of the top Bruhat incidence commutes with flat completion: locally its ideal is saturation by the finitely generated boundary ideal; the colon chains stabilize and flat base change preserves their kernels. The completed full graph ideals therefore agree. Additive Cohen–Macaulayness and Cartan flatness transfer to the formal graph, and faithful flatness of local completion transfers the regularity of the Cartan sequence back along the unipotent fiber. The sequence is regular on both the graph and its maximal Cohen–Macaulay canonical module. Compute the latter by finite Ext in a regular ambient scheme and use canonical adjunction for the regular sequence. The Cartan cotangent determinant is a trivial conjugation character. The full correspondence is the associated bundle $G \times^B \mathcal{Z}_w$ over the first flag variety. Restricting its canonical sheaf to the first-flag slice contributes the cotangent determinant of G/B at B, whose B-character is $\det \mathfrak{u}$. Thus its normalized fiber is $\omega_{S_w} \otimes \det \mathfrak{u}$, giving (3.33) in ordinary degree zero. □

We now assume that G_{der} is simply connected and that the automorphism on G itself is the identity. Thus the odd component is $G\tau$ with τ central and $\tau^2 = z \in Z(G)$. Let $t \in T$ satisfy $b_0 = zt^2 \in Z(G)$, and use the original shifted mark $y = tr_G$. The centralizer $K = Z_G(t)$ is connected [Ste65, Remark 2.10]. For reductive G with simply connected derived subgroup, this follows by passing through $Z(G)^\circ \times G_{\mathrm{der}} \to G$: the commutator of two lifts has first coordinate 1, and hence cannot be a nontrivial element of the diagonal kernel. Its centralizer is the image of the connected product centralizer.

Let

$$F_w = [C_{s_1} \longrightarrow \mathbf{1}] \star \cdots \star [C_{s_m} \longrightarrow \mathbf{1}], \qquad C_{s_i} \text{ in ordinary degree } 0,$$

be the positive reduced Rouquier kernel, and write $P_w = \mathcal{N}_G(F_w)_{y,y}$ for the coefficient completed at the original output point. The definition of F_w includes the actual restriction maps in the two-term factors.

Proposition 3.21. *With these hypotheses and conventions,*

$$P_w \otimes^L_{A_{\mathrm{out}}} k \simeq R\Gamma([I_G/B_K], R\mathcal{H}om_{I_G}(\mathcal{K}_{w^{-1}}, i_*\mathcal{O}_{Y_K}))\,. \tag{3.46}$$

In particular $w(t) \neq t$ implies $P_w = 0$.

Proof. The exact sequence [BR12, (1.10.2)] identifies a two-term simple Rouquier kernel with its canonical simple graph sheaf in ordinary degree zero; its character is given in [BR12, Remark 1.10.5]. The canonical-sheaf assertion of [BR12, Theorem 2.2.1] reverses the simple factors. Hence, with our convolution convention, F_w is the normalized canonical sheaf of the full graph of label w^{-1}. Translate by b_0^{-1} and apply Lemma 3.20. It identifies the completed multiplicative diagrams, their restriction maps, and their canonical sources. Fixing the first flag leaves precisely the normalization $\mathcal{K}_{w^{-1}}$ of (3.33).

The ordinary norm formula gives $\mathcal{N}_G(F_w) = R\,\mathrm{Hom}_{W_o}(Lh^*F_w, \mathcal{O}_{W_o})$, with the output Cartan acting on the target. Its residue field has a finite Koszul resolution. Tensoring with this finite complex commutes with Hom by finite cones and retracts, without requiring the source Lh^*F_w to be perfect over W_o. The odd incidence is classical Cohen–Macaulay and flat over the output Cartan. In the present untwisted case, its relative-position

stratum has a holonomy fiber $U \cap vUv^{-1}$ of dimension $n - \ell(v)$ over a cell of dimension $\ell(v)$ at every Cartan point; constant fiber dimension and the regular base give flatness. Thus target tensoring produces the ordinary structure sheaf of the actual output fiber and keeps the source pullback derived.

After fixing the first flag, that fiber has the full B gauge and atlas $\{(a, F) : a \in B\tau,\ \mathrm{Cartan}(a) = t,\ aF = F\}$. Conjugation gives a scheme isomorphism

$$B \times^{B_K} (tU_K) \xrightarrow{\sim} B|_t, \qquad [b, tu] \longmapsto btub^{-1}. \tag{3.47}$$

In ordered root coordinates, eliminate the noncentralizing directions successively by root height. Their linear coefficients $1 - \alpha(t)$ are units, while commutators involve only larger heights. This constructs the inverse modulo the residual B_K action. Equivalently, the semisimple part of an upper triangular matrix with fixed diagonal t is obtained by polynomial spectral projectors; distinct eigenvalues have invertible differences. The universal triangular matrix is annihilated by fixed powers of the eigenvalue polynomials, so the construction commutes with arbitrary base change, including nonreduced bases.

The same projectors preserve every invariant flag. A flag fixed by tu is therefore fixed by both t and u. The fixed-flag scheme $(G/B)^t$ is the disjoint union of K-flag components indexed by $W_K\backslash W$; on the component indexed by v the semisimple second Cartan is $v^{-1}t$. Connectedness of $Z_G(t)$ identifies the Weyl stabilizer of t with W_K. The diagonal input idempotent thus selects exactly K/B_K. It imposes no second-Cartan equation. Quotienting (3.47) by B retains the whole B_K gauge. Polynomial logarithm on U_K identifies the selected fiber, as a quotient stack, with $[Y_K/B_K]$.

For $a = tu\tau$, its actual square is $a^2 = b_0u^2$. After translation and logarithm the norm is $x \mapsto 2x$. This scaling preserves the conic graph and its canonical sheaf. The first even Cartan is zero, so the norm map factors through I_G. Lemma 3.20 identifies its genuine canonical specialization with $\mathcal{K}_{w^{-1}}$. The subsequent pullback to Y_K stays derived. Derived closed-map adjunction transfers it to $i_*\mathcal{O}_{Y_K}$, which proves (3.46).

The order of operations is essential: tensor the original coefficient with the finite output residue complex, identify the algebraic fiber $[Y_K/B_K]$, and then compute rational invariants. No rationality of arbitrary Cartan-completed functions is used. If $w(t) \neq t$, then $w^{-1} \notin W_K$, so Theorem 3.19 annihilates the derived fiber. The coefficient is perfect over the completed output Cartan by Proposition 3.1, applied termwise to the finite Rouquier complex. Derived Nakayama gives $P_w = 0$. □

The identity assumption on the automorphism in this proposition concerns the group, including its central torus. An automorphism trivial on the roots need not be the identity there. The product and central-isogeny reductions in the global argument separate that torus contribution before applying this proposition.

Corollary 3.22. *Suppose the coefficients of all canonical labels of length less than L are ordinary, and let $\ell(w) = L$. At a diagonal raw mark t with $w(t) \neq t$, the coefficient of B_w has no positive ordinary cohomology.*

Proof. The degree-zero term of F_w is the reduced Bott–Samelson word BS_w. Every term in positive degree is a shorter word, whose canonical summands have labels of length less than L. If T_w is its stupid truncation in degrees at least one, contravariance gives a triangle $\mathcal{N}(\mathrm{BS}_w) \longrightarrow \mathcal{N}(F_w) \longrightarrow \mathcal{N}(T_w)$. By hypothesis the last term lies in $D^{\leq -1}$. The head B_w occurs once in BS_w, and its other canonical summands have shorter labels. Therefore

$$H^j\mathcal{N}(B_w)_{y,y} \simeq H^j\mathcal{N}(F_w)_{y,y} = 0 \qquad (j > 0)$$

by Proposition 3.21. □

Remark 3.23. The incidence tail is present even in rank one. At a noncentral mark for SL_2, the local rings and canonical source are

$$R = k[u, z]/(uz^2), \qquad S = R/(uz), \qquad \mathcal{K} = S\langle\alpha\rangle, \qquad \mathrm{wt}(u) = -\alpha, \quad \mathrm{wt}(z) = \alpha.$$

The alternating resolution by uz and z gives

$$\mathrm{Ext}^j_R(\mathcal{K}, k) = k\langle-(\lfloor j/2\rfloor + 1)\alpha\rangle \qquad (j \geq 0).$$

Every ordinary Ext group is nonzero, while all torus invariants vanish. Dropping the canonical character would already leave an invariant in degree zero. This illustrates why both the canonical normalization and the full incidence resolution are required.

3.4. **Semisimple localization and marked descent.**

3.4.1. *Conventions and the ordinary coefficient.* Let G be a connected complex reductive group with a pinned involution θ preserving (B,T). In the prescribed odd extension write $\tau^2 = c \in Z(G)^\theta$. Fix $r \in T^\theta$ with $r^2 = c$ when using the canonical marking; during an inherited-datum calculation the marking is retained as part of the datum and need not be re-chosen. Set $Y = T/(1-\theta)T, \qquad A = \mathcal{O}(Y), \qquad q(t\tau) = [tr], \qquad N([tr]) = t\theta(t)c.$ Here $(1-\theta)T$ is the image subtorus. The coordinate $[tr]$ is the original closed marking, and t is its raw odd coordinate. Write $\mathcal{C}_o = [G\tau/G]$, let $F_G = [B\tau/B]$ be the odd stable Borel stack, and put $W_{o,G} = F_G \times^{\mathrm{R}}_{\mathcal{C}_o} F_G, \qquad h_G(a, B_0, B_1) = (a^2, B_0, B_1)$. The ordinary source formula of Proposition 2.19 is

$$\mathcal{N}_G(K) = \mathrm{RHom}_{\mathrm{QCoh}(W_{o,G})}(Lh_G^* K, \mathcal{O}_{W_{o,G}}). \tag{3.48}$$

Every assertion in this section is made before taking the original small circle invariants and inverting its positive parameter u. In particular, no periodic vanishing is used to infer ordinary vanishing.

3.4.2. *The full semisimple-centralizer formula.* Fix a semisimple $b_0 \in T^\theta$ and let $H = Z_G(b_0)^\circ, \quad B_H = B \cap H, \quad \Phi_H = \{\alpha \in \Phi_G : \alpha(b_0) = 1\}$. The full torus T lies in H; the inherited odd coset is $H\tau$, with the same square c and the same output map $[tr]$. The restriction of θ can have signs on the simple root spaces of H. We retain these signs until the explicit transport argument below. Define

$$H^G = \{h \in H : \det(1 - \mathrm{Ad}(h)|\mathfrak{g}/\mathfrak{h}) \neq 0\}, \qquad U_H = D(\delta_H) \subset Y,$$

where

$$\delta_H(y) = \prod_{\alpha \in \Phi_G \setminus \Phi_H} (1 - \alpha(N(y))).$$

Choose the positive representatives $D = \{d \in W_G : d^{-1}\Phi_H^+ \subset \Phi_G^+\} \subset W_H \backslash W_G$. There is one such representative in each left coset: the intersection of a positive system with Φ_H is a positive system of Φ_H, and W_H acts simply transitively on its chambers. The set D is θ-stable. This argument does not require H to be a standard Levi.

Proposition 3.24. *Restriction to the representable etale chart $[H^G/H] \to [G/G]$ identifies the Grothendieck stack with a disjoint union of H-Grothendieck stacks indexed by D, and its Steinberg stack with the full $D \times D$ matrix. Convolution is matrix convolution. If $K_{d,e}$ denotes the resulting matrix entry, then*

$$\mathcal{N}_G(K)|_{U_H} \simeq \bigoplus_{d \in D^\theta} \mathcal{N}_H(K_{1,d})|_{U_H}. \tag{3.49}$$

The output algebra on both sides is the original A. Both odd endpoints, all matrix indices at intermediate flags, and all stabilizer groups are retained.

Proof. The conjugation map $G \times^H H^G \to G$ has differential $(X \bmod \mathfrak{h}, Y) \longmapsto (1 - \mathrm{Ad}(h))X + Y$, an isomorphism on this open. It is therefore etale. If $h \in H^G$ and h_s is its semisimple part, the absence of eigenvalue 1 on the normal representation gives $Z_G(h_s)^\circ \subset H$. Every h-stable Borel contains a maximal torus containing h_s. That torus lies in H and contains the central element $b_0 \in H$. Thus every such Borel is fixed by b_0.

The scheme-theoretic fixed locus of b_0 on G/B is the disjoint union of H/B_H indexed by D. To check exhaustion, put a maximal torus of a fixed Borel into T using H and then use the Weyl flags. The tangent spaces are exactly the root spaces in Φ_H; the fixed locus of a diagonalizable group on this smooth flag variety is smooth. On normal directions the equations for h-stable flags have invertible differential $1 - \mathrm{Ad}(h)$. This identifies the derived Grothendieck diagrams, and derived fiber products give the entire matrix. Proper base change on three or more flags gives its convolution and associativity.

It remains to identify the specified odd output, rather than merely an etale chart of the odd quotient. Order positive root coordinates by ambient height and let N_{tr} be the ordered product of root subgroups outside Φ_H. This is a transversal variety, not in general a subgroup. Triangular root multiplication gives $N_{\mathrm{tr}} \times U_{B_H} \simeq U_B$. For $a_H \in B_H\tau$, consider na_Hn^{-1}. At each root height the linear coefficient on the outside coordinates is $1 - L_t, \qquad L_t = \mathrm{Ad}(t)\theta$. Since $L_t^2 = \mathrm{Ad}(N([tr]))$, its inverse on U_H is $(1 + L_t)(1 - L_t^2)^{-1}$. All other terms have smaller height. Solving successively proves the algebraic isomorphism $B \times^{B_H} (B_H\tau)|_{U_H} \simeq (B\tau)|_{U_H}$. Taking the full Borel quotient identifies the original output source with $F_H|_{U_H}$. Uniqueness in the height recursion also identifies automorphism groups: a gauge between two normalized points has trivial outside transversal and belongs to B_H.

The odd chart $[(H\tau)^G/H] \to [G\tau/G]$ is etale by the same differential calculation, because $1 - \mathrm{Ad}(a)$ has inverse obtained from $1 - \mathrm{Ad}(a^2)$. Every a-stable input Borel is a^2-stable and hence lies in one of the fixed-flag components D. The inner factor in H preserves a component and θ permutes the components; exactly the indices D^θ can support stable input Borels. A θ-fixed pinned Tits representative of d identifies that component

with the inherited H-source. Normal differentials are again invertible. Consequently the pullback of the input source is the disjoint union over D^θ of F_H.

Combining this input calculation with the specified output identification gives

$$W_{o,G}|_{U_H} \simeq \coprod_{d\in D^\theta} W_{o,H}|_{U_H}.$$

On its d-th component h_G is precisely the $(1,d)$-matrix boundary map. Formula (3.48) proves (3.49). Scalar localization is justified by the bounded one-endpoint proper-push model in Proposition 2.19; compactness of Lh_G^*K is not required. □

Simple entries are actual unshifted kernels. For an ambient simple root α_s and $d \in D$, if $\beta = d\alpha_s \in \Phi_H$, then β is a positive simple root of H. Otherwise expressing it as a sum of positive H-roots and applying d^{-1} would express α_s as a sum of positive ambient roots. The corresponding diagonal entry is C_β^H. If $d\alpha_s \notin \Phi_H$, both d and ds are distinguished, and the partial-flag component has two full-flag refinements. Its block is

$$\begin{pmatrix} \mathbf{1}_H & \mathbf{1}_H \\ \mathbf{1}_H & \mathbf{1}_H \end{pmatrix}.$$

These are isomorphisms of partial-flag diagrams before forming kernels. There are no cohomological shifts or omitted normal factors. Hence a simple word restricts to finite direct sums of actual H-words.

The marked input projector. An additional formula is required if a coefficient is split by both closed markings. On the d-component the raw H input coordinate is sent to $d^{-1}t$. Define

$$\chi_d(y) = [d^{-1}(yr^{-1})r], \qquad d \in D^\theta. \tag{3.50}$$

Then $q_{\mathrm{in},G} = \chi_d \circ q_{\mathrm{in},H}$, while the output mark is unchanged. After any common central completion that splits the marked points, the ambient projector at y_{in} pulls back to the H input projector at $\chi_d^{-1}(y_{\mathrm{in}})$. Explicitly,

$$e_{y_{\mathrm{out}}}\mathcal{N}_G(K)e_{y_{\mathrm{in}}} \simeq \bigoplus_{d\in D^\theta} e_{y_{\mathrm{out}}}\mathcal{N}_H(K_{1,d})e_{\chi_d^{-1}(y_{\mathrm{in}})}$$

on the localized completed source. A nonexistent branch contributes zero. This follows from the maps of actual Cartan functions and therefore also holds for nonreduced formal neighborhoods. It is not an equality of untransported input labels. For concentration, each projector is simply an exact retract; no extra vanishing theorem is needed.

3.4.3. *Signed actions and arbitrary inherited central squares.* Suppose σ preserves (B_H, T) but has root-space signs. Let θ_p be its pinned automorphism on the complete based root datum. Choose $s \in T$ with $\sigma = \mathrm{Ad}(s)\theta_p$. Such an s exists for every central form: the prescribed simple-root scalars define a character of the root lattice, and it extends to $X^*(T)$ because $\mathbb{C}^*$ is divisible. Put $z_s = s\theta_p(s) \in Z(H)$, $\eta = s^{-1}\tau$, $\eta^2 = c_p = z_s^{-1}c$. The actual odd-extension isomorphism is $g\tau \mapsto gs\eta$ and $(gs)\theta_p(gs)c_p = g\sigma(g)c$. It preserves the odd element, its square, both stable endpoint conditions, and their full stabilizers. If the old marking is $[tr_{\mathrm{old}}]$ and a pinned theorem uses $[t_pr_p]$, where $t_p = ts$, the required translation is

$$\phi(y) = y[sr_pr_{\mathrm{old}}^{-1}]. \tag{3.51}$$

Changing c_p to a canonical central square c_{can} is a separate operation. The odd twisted-conjugation stack, its stable flags, and W_o do not depend on the central square. Only their even boundary map changes, by simultaneous central translation m_z with $z = c_{\mathrm{can}}c_p^{-1}$. The multi-flag diagram of every geometric word is invariant under this translation, so $Lm_z^*C_{\mathbf{s}} \simeq C_{\mathbf{s}}$. Thus

$$\mathcal{N}_{\sigma,c,r_{\mathrm{old}}}(C_{\mathbf{s}}) \simeq \phi^*\mathcal{N}_{\theta_p,c_{\mathrm{can}},r_p}(C_{\mathbf{s}}). \tag{3.52}$$

This proves transfer of global ordinary-word concentration. For a local theorem its open set must also be pulled back by ϕ^{-1}. Central translation transports the boundary functions and circle holonomy; it does not prove an unchanged full marked categorical equivalence for an arbitrary coherent kernel. This restricted scope is sufficient for the centralizer induction.

3.4.4. *Finite central isogenies.* Let $\pi : G' \to G$ be a compatible pinned central isogeny with kernel D. Write Y', A' for its marked closed torus and ring. The map $Y' \to Y$ is finite etale and surjective. For an ordinary geometric word, the precise base-change statement is

$$\mathcal{N}_{G'}(C'_{\mathbf{s}}) \simeq \mathcal{N}_G(C_{\mathbf{s}}) \otimes_A A'. \tag{3.53}$$

An unrestricted version for arbitrary character sectors would be false.

Here is the character correction. On the output-Borel stable-pair atlas $W_G = [Z/B]$, the lifted incidence is $Z' = Z \times_T T'$. Set

$$D_- = (1-\theta)D, \quad \overline{D} = D/D_-, \quad K_{\mathrm{inertia}} = \frac{D \cap (1-\theta)T'}{(1-\theta)D}, \quad D_Y = \frac{D}{D \cap (1-\theta)T'}.$$

The stable-pair map factors as a D^θ-gerbe followed by a finite etale $\overline{D}$-torsor. Since diagonalizable groups in characteristic zero have exact invariants, its pushforward structure algebra is the finite algebra $\mathcal{E}$ of that torsor. In general the coefficient upstairs is $\mathrm{RHom}(M, \mathcal{E})$, with $M = Lh_G^* K$, rather than immediate coarse-torus base change.

For $t \in T^\theta$, choose a lift t' and define $\delta(t) = (1-\theta)t' \bmod D_-$. This gives $T'^\theta \longrightarrow T^\theta \xrightarrow{\delta} K_{\mathrm{inertia}} \longrightarrow 1$. The fixed subgroup T^θ_{ad} is a connected torus: its character coinvariants are free, since θ permutes the basis of simple roots. The identity component of T'^θ maps onto it by the differential criterion. Therefore $T^\theta = \pi((T'^\theta)^\circ) Z(G)^\theta, \qquad \delta(Z(G)^\theta) = K_{\mathrm{inertia}}$. The fixed center is central inertia on W_G. A geometric word has its trivial character: this is true for structure-sheaf simple kernels and is preserved by all word pullbacks, tensors, and proper pushforwards. Hence only the trivial fixed-center-character part of $\mathcal{E}$ contributes. It is $\mathcal{E}^{K_{\mathrm{inertia}}} = \mathcal{O}_{W_G} \otimes_A A'$, because $T'/(D \cap (1-\theta)T') \simeq T \times_Y Y'$. This proves (3.53) by (3.48). Tensoring with finite projective A' commutes with this Hom even if $Lh_G^* K$ is unbounded, since it uses only finite direct sums and retracts. The same multi-flag diagrams prove naturality, products, and the unit. Compatible pinnings send $r' \mapsto r$ and $c' \mapsto c$, so these are the original markings.

For the categorical descent use the specific real quotient $q : H_0 \longrightarrow H_0/Z(H_{0,\mathrm{der}})$. Its target is an adjoint derived group times a real torus. The central descent in the proofs of Lemmas 2.26 and 2.29, from [DT26a, proof of Theorem 1.1], gives $\beta_Z^* \Phi \simeq \Phi'$ oblv on ordinary compact affine kernels. Here oblv is direct image along the corresponding finite cover on each enhanced Bruhat stratum. Restriction gives the ordinary monoidal square for finite kernels.

The real coefficient base change has the same marked ring map. Indeed, if $c : H_1 \to H_2$ is the real isogeny, write X_i for compact enhanced flags and $\widehat{H}_1 = c^{-1}(H_2(\mathbb{R})) \subset H_1(\mathbb{C})$. The map of real enhanced quotient stacks is the finite cover with deck group $\widehat{H}_1/H_1(\mathbb{R})$. Retaining the input enhancement, the mixed action stack $[H_1(\mathbb{R})\backslash(X_1 \times X_2)]$ is the base change of the H_2 action stack along the input cover. Its base change along the complex pair cover is the H_1 action stack. These two Cartesian identities give the module identity $E(K * M) = F(K) * E(M)$ for the finite direct images E, F. On the closed orbit, regular-local-system induction gives $E(P_1) = P_2$ and $e^* P_2 = P_1 \otimes_{A_1} A_2$. Adjunction gives

$$\mathrm{RHom}(F(K)P_2, P_2) \simeq \mathrm{RHom}(KP_1, P_1) \otimes_{A_1} A_2. \tag{3.54}$$

The normalization is coefficient projection to the subgroup lattice, not the unnormalized algebra trace. Thus the unit vector is 1. Higher action diagrams give products and coherences. This argument uses the full groups of real points; their identity components cannot be substituted.

Concentration and finite projectivity descend faithfully through $A \to A'$. For comparison maps, use the generic microlocal normalization of Theorem 2.53. Deck translations are represented by the central kernel and commute with the affine W^θ action. The microlocal return operator is normalized to send 1 to 1; this condition is preserved by the finite-cover induction of regular local systems. Consequently the specified generic comparison is deck-equivariant and descends. The two pullbacks of its integral extension to $A' \otimes_A A'$ agree generically and hence everywhere by torsion-freeness; faithful flat descent gives the comparison, inverse, products, and unit over A.

3.4.5. *Products, central tori, exchanged factors, and globalization.* On a direct product, the ordinary multi-flag sources are products. A word splits into its commuting factor words, and its coefficient is the tensor product of their coefficients. A torus contributes its marked Laurent ring in degree zero. Apply the exact quotient $H_0 \to H_0/Z(H_{0,\mathrm{der}})$ on the real side. Its dual is a finite central cover by a product of simply connected simple groups and a torus. The pinned involution permutes the simple factors and preserves the torus factor; on a fixed semisimple factor with trivial diagram action it is the identity. The preceding descent restores the coupled

central lattice and all stabilizers. In particular a nontrivial involution on the central torus is never discarded when treating a root-trivial action.

For exchanged simple factors, Appendix A gives an ordinary two-flag model. The spectral parameter stack is the group Steinberg stack and Q is its diagonal unit; the real category is the finite Hecke category with its regular unit. The ordinary tame equivalence identifies the left and right actions and their coefficient Homs. Standard–costandard adjunction makes word coefficients ordinary before circle invariants. The orientation-cover calculation in Lemma A.3 identifies the original positive circle with the Steinberg circle, with no extra factor of two. Central descent then covers exchanged factors with finite central or torus couplings.

Proposition 3.25. *Suppose an ordinary central-norm theorem is known for a connected pinned root action, at every original output mark and for every finite geometric word. Suppose the ordinary global word theorem is known for every connected factor, with its inherited action, of every proper semisimple norm centralizer. Then all global ordinary word coefficients of the ambient datum are concentrated in degree zero. Using Proposition 3.1 and Lemma 3.3, their degree-zero modules are finite projective over the original A.*

Proof. Fix a closed output point y and take $b_0 = N(y)$. If b_0 is central, use the stated central theorem. Otherwise Proposition 3.24 restricts every word to finite sums of ordinary words in $H = Z_G(b_0)^\circ$ on an open neighborhood of y. Signed transport, product decomposition and exact central descent put each entry into its known actual inherited theorem. Input marked summands are the retracts specified in (3.50). Thus every completed output coefficient is ordinary. For finite cohomology modules, vanishing at all maximal localizations, or at their faithfully flat completions, gives global vanishing. Proposition 3.1 and Lemma 3.3 then force the degree-zero module to be projective: a module of finite projective dimension with no positive Tor at any residue field has projective dimension zero locally. Finiteness globalizes this to finite projectivity. □

3.4.6. *The order of the centralizer induction.* The following order makes the centralizer induction acyclic. The connected proper-centralizer dependencies are as follows; torus and component data are handled by the preceding proofs, not by suppressing them in this list.

Ambient action	Connected proper factors needed
Root-trivial A_n	Root-trivial A of smaller rank.
Outer A_n	Smaller internally outer A blocks and exchanged A blocks.
Root-trivial C_n	Root-trivial C_r, $r < n$, and type A.
Root-trivial D_n	Root-trivial D_r, $r < n$, and type A.
Root-trivial B_n	Root-trivial B_r, $r < n$, and D_s, $s \le n$, and type A.
Outer D_{n+1}	Outer D_r, $r \le n$, root-trivial D_s, $s \le n$, and type A.
G_2, F_4	Classical blocks, including a same-rank classical pseudo-Levi when present.
Root-trivial E_6	Type A, D of rank at most 5.
Outer E_6	Smaller A, D, including outer D_5; root-trivial connected factors have rank at most 4.
Root-trivial E_7	Type A, D, E_6.
Root-trivial E_8	Type A, D, E_6, E_7.

For completeness, the classical bounds follow directly from eigenspace decomposition of the actual vector norm. In orthogonal type decompose into V_+, V_- and reciprocal pairs $E_\lambda \oplus E_\lambda^*$. An orthogonal odd matrix whose square is -1 on V_- has equally many i and $-i$ eigenvalues there and determinant 1. Reciprocal pairs also have determinant 1. In outer D the determinant -1 is therefore carried by nonzero V_+; this gives the outer D_r factor and leaves the D_s factor root-trivial. A noncentral norm excludes the whole vector space from a single central eigenspace. In odd orthogonal type the possible $+1$ eigenline and $2n$-dimensional -1 eigenspace give the genuine same-rank D_n dependency in B_n. Symplectic eigenspaces give C_r, C_s with $r, s < n$ and reciprocal type-A blocks. The central scalar relating the vector square to the prescribed norm only interchanges the sign labels and does not change the bound.

In outer E_6, the fixed character space has rank 4, so a root-trivial connected subsystem has rank at most 4. A proper full-rank connected A_6 or D_6 subsystem cannot occur: its root-lattice determinant ratio to E_6 would be $7/3$ or $4/3$, whereas a full-rank sublattice has integral squared index. The only connected exceptional type of rank at most 6 is E_6 itself; a subsystem of that type already has all 72 roots. Thus the new largest outer proper factor is D_5. The same determinant/root-count argument excludes connected full-rank proper exceptional factors in the other simply laced exceptional rows. In E_8, in particular, every proper connected component is already classical, E_6, or E_7.

3.5. **The adjacent rank-two orbit.** We establish the low-rank comparison for the involution exchanging the two simple roots of type A_2. By central descent it is enough to use $G = SL_3$; its real dual is the quasi-split adjoint group $PU(2,1)$. Write $R = \mathbb{C}[x^{\pm1}, y^{\pm1}]$, $R^W = \mathbb{C}[u,v]$, $u = x + y/x + y^{-1}$, $v = y + x/y + x^{-1}$. The actual norm map sends x, y to q. Thus the boundary torus element is $\mathrm{diag}(q, 1, q^{-1})$, and $A = \mathbb{C}[q^{\pm1}]$, $S = \mathbb{C}[c]$, $c = q + q^{-1}$, $u, v \longmapsto c + 1$. For this pinned SL_3 lift, $r_G = \mathrm{diag}(-1, 1, -1)$, so translation by r_G leaves $q = a/f$ unchanged. We therefore keep the original marking in these coordinates.

Theorem 3.26. *The marked comparison holds for the adjacent A_2 involution. Before circle invariants, the full-averaging coefficient algebras on both sides identify, under the prescribed generic comparison, with $\mathrm{End}_S(A)$. Both Cartan actions are composition with multiplication on A.*

3.5.1. *The spectral full-averaging algebra.*

Lemma 3.27. *The spectral full-averaging algebra is $E_Q = \mathrm{End}_S(A)$, with the indicated Cartan actions and trace (3.55).*

Proof. Set

$$J = \begin{pmatrix} 0 & 0 & 1 \\ 0 & -1 & 0 \\ 1 & 0 & 0 \end{pmatrix}, \qquad \theta(g) = Jg^{-T}J, \qquad M = gJ, \qquad k = MM^{-T}.$$

Twisted conjugation becomes congruence $M \mapsto hMh^T$, and $c(M) = \mathrm{tr}(k) - 1$ restricts to $q + q^{-1}$. The identity $kM = Mk^{-T}$ and odd-dimensional skew symmetry show that k has eigenvalues $1, q, q^{-1}$ generically. On that dense open, congruence diagonalization shows that the invariant functions restrict injectively to $A^{q \mapsto q^{-1}} = S$; the displayed global function c proves $\mathcal{O}(G\tau)^G = S$. For a tangent vector $\dot{M} = XM$, $\mathrm{tr}(X) = 0$, $dc_M(X) = \mathrm{tr}((k - k^{-1})X)$. Thus $dc = 0$ exactly when $k^2 = 1$: any scalar value of $k - k^{-1}$ must vanish on its 1-eigenvector. The two possibilities for k are 1 and a conjugate of $\mathrm{diag}(1, -1, -1)$. In the first case M is a nondegenerate symmetric form, with stabilizer SO_3; in the second it is a symmetric one-dimensional block plus a skew-symmetric two-dimensional block, with connected stabilizer Sp_2. Both loci have dimension 5 in the smooth 8-dimensional SL_3.

Consequently the cameral cover $D = \{(M, q) : q^2 - c(M)q + 1 = 0\}$ is Cohen–Macaulay and regular in codimension one: its singular locus is contained in $q = \pm1$, $dc = 0$, of codimension at least three. The twisted Grothendieck resolution $\widetilde{G\tau}$ maps properly and birationally to D, since its two generic stable flags have distinct coordinates q, q^{-1}. Irreducibility of $\widetilde{G\tau}$ gives the quadratic function-field extension, so D is integral and hence normal. Normal proper-birational pushforward, together with Lemma 2.21, gives an equivariant algebra identity $R\rho_*\mathcal{O}_{\widetilde{G\tau}} = \mathcal{O}_{G\tau} \otimes_S A$. Proposition 2.22, followed by the endomorphism bundle and exact G-invariants, therefore gives the actual spectral full-averaging coefficient algebra $E_Q = \mathrm{End}_S(A)$. Its two A actions are composition with multiplication on A, and its unit is the identity. The odd stable-flag differential has factors $1 - q^{-1}$ on the exchanged negative simple-root pair and $1 + q^{-1}$ on the highest root line. The original right-Maurer orientation gives Jacobian $1 - q^{-2}$ and the trace

$$\delta(f) = \frac{f(q)}{1 - q^{-2}} + \frac{f(q^{-1})}{1 - q^2}. \tag{3.55}$$

Its pairing in the basis $(1, q)$ is

$$\begin{pmatrix} 1 & c \\ c & c^2 - 1 \end{pmatrix}, \qquad \det = -1,$$

so it remains perfect at both ramification points. □

3.5.2. *A normalized ordinary operator map.* Let P_C be the original real closed object and put $E_R = \mathrm{Hom}(C_{\mathrm{all}}P_C, P_C)$, with the coalgebra coefficient product. Its real standard filtration makes it ordinary and finite free of rank two over A, by Lemma 2.45. We construct the specific ordinary map

$$E_R \longrightarrow \mathrm{Hom}_R(R \otimes_{R^W} A, A) = \mathrm{End}_S(A). \tag{3.56}$$

The underlying microlocal realization of C_{all} is $R \otimes_{R^W} R$, with its Sweedler coalgebra $\epsilon(a \otimes b) = ab$, $\Delta(a \otimes b) = a \otimes 1 \otimes b$. Lemma 3.32 identifies this coalgebra with the actual Whittaker adjunction coalgebra, using its unit at the Whittaker unit and the prescribed oriented tensor constraint. Applying the ordinary real microlocal functor therefore gives (3.56) with its coefficient product, unit and both Cartan actions. Under the single choice of reflection frames in that lemma, its regular restriction is exactly the generic comparison of Theorem 2.53. It is generically injective and hence injective, since E_R is A-torsion-free.

The full derived microlocal target is $R\mathrm{Hom}_{R^W}(A,A)$, whose normal Ext^1 is contributed by the divisor $u=v$. Equation (3.56) uses its degree-zero Hom; the ordinary operator algebra is not the entire derived target.

3.5.3. *The completion at $q=1$.*

Lemma 3.28. *The injection (3.56) is an isomorphism after completion at $q=1$.*

Proof. Put $\widehat{A}_+=\mathbb{C}[[q-1]]$, $\widehat{S}_+=\mathbb{C}[[c-2]]$ and complete R at $x=y=1$. The closed-orbit stabilizer is connected. Indeed, in a Witt basis its preimage in $SU(2,1)$ is a semidirect product of a real unipotent group isomorphic as a variety to $\mathbb{R}^3$ with $\{\mathrm{diag}(z,\bar z/z,\bar z^{-1}) : z\in\mathbb{C}^*\}$. Both factors are connected, and the quotient by the central μ_3 remains connected. Thus the closed orbit has a single unipotent local-system label. The closed cross-action stabilizer is $W^\theta=\{1,w_0\}$, acting by $q\mapsto q^{-1}$. The closed block algebra B_0 of Ionov–Yun is consequently $\widehat{S}_+$. Theorem 9.21 of [IY25] identifies this action by evaluation on the closed object, and Corollary 9.22 gives full faithfulness on ordinary completed free-monodromic tiltings for this adjoint datum, which is nice by [IY25, Lemma 9.3]. The closed object is tilting, and its full-averaging translate is tilting by [IY25, Proposition 6.4]. The enhanced block algebra and the two modules are

$$\widehat{\mathcal{B}}=\widehat{S}_+\otimes_{\widehat{R^W}}\widehat{R},\qquad V_R(P_C)=\widehat{A}_+,\qquad V_R(C_{\mathrm{all}}P_C)=\widehat{\mathcal{B}}\otimes_{\widehat{S}_+}\widehat{A}_+.$$

The indicated $\widehat{S}_+$ action is the original action on P_C, not a change of formal coordinate. Full faithfulness gives

$$\widehat{E}_{R,+}=\mathrm{Hom}_{\widehat{\mathcal{B}}}(\widehat{\mathcal{B}}\otimes_{\widehat{S}_+}\widehat{A}_+,\widehat{A}_+)=\mathrm{End}_{\widehat{S}_+}(\widehat{A}_+).$$

The completion of the universal finite Morse-cochain functor of Lemma 2.50 is the Ionov–Yun functor; the tensor normalization above identifies this map with the completion of (3.56). Ordinary Hom commutes with completion here by the finite standard/costandard filtrations of these tiltings. Thus faithful flatness proves equality of the specified ordinary lattices at $q=1$. This is an ordinary completed theorem, used only at its unipotent monodromy, with no periodized inference. □

3.5.4. *The completion at $q=-1$.*

Lemma 3.29. *The injection (3.56) is an isomorphism after completion at $q=-1$.*

Proof. The boundary point is $t_0=\mathrm{diag}(-1,1,-1)$. Its Weyl stabilizer is the highest-root reflection. The clean s-graph, since $\alpha_s(t_0)=-1\neq 1$, transports it to $s(t_0)=\mathrm{diag}(1,-1,-1)$, whose stabilizer is the standard t-Levi. Etale-locally the characteristic polynomial has its relatively prime Hensel factors of degrees one and two. The spectral projectors give $V=V_+\oplus V_-$, and the selected flag branch is $F_1=V_+$, $F_2=V_+\oplus\ell$ with a stable line $\ell\subset V_-$. Consequently its full-averaging pair and triple flag diagrams are precisely the t-parabolic pair and triple diagrams. Conjugation from the centralizer Levi is etale, as $1-\mathrm{Ad}(t_0)$ is invertible off its roots, so this identification includes the quotient gauges and all derived fiber products. Changing each Borel by the same clean s-graph commutes with forgetting and repeating flags. Hence it identifies the counit and comultiplication as well as the kernel objects. Only the selected middle idempotent contributes to a coefficient at the selected seed. With $m_1=\Delta_sP_C$ this proves $\widehat{E}_{R,-}=\mathrm{Hom}(C_tm_1,m_1)^\wedge_{q=-1}$. m_1 is the transported seed; the original seed itself has only its original branch projector.

This is an actual real rank-one coefficient. A closed $PU(2,1)$ flag consists of an isotropic line in its orthogonal plane. The s-ascent changes it to a positive line in a degenerate plane. Forgetting that plane has base the open orbit of positive lines, whose stabilizer is $(U(1)\times U(1,1))/U(1)_{\mathrm{diag}}\simeq U(1,1)$. The fiber is the projective line in the Hermitian quotient of signature $(1,1)$: the closed orbit is $\mathbb{RP}^1$, and the two open orbits are the two disks. The transported middle orbit is exactly that closed rank-one orbit. The associated-bundle description identifies the equivariant sheaves and their full enhanced Cartan. Extension from the inverse image of the open base is fully faithful and intertwines the t-Hecke correspondence, which keeps the positive line fixed. Propositions 2.33 and 2.34 therefore apply to this $U(1,1)$ coefficient.

Its dual Levi is GL_2 with involution $M\mapsto M/\det M$, trivial on SL_2 and inverse on its center. The full boundary norm is $\mathrm{diag}(q_L,q_L^{-1})$, while the transported original block is $\mathrm{diag}(q,q^{-1})$. Hence $q_L=q$ with its original sign, rather than the insufficient statement $q_L^2=q^2$. The pinned odd action on the original highest-root line has the sign -1, so its raw rank-one root coordinate is $-q$ up to inversion, consistently with this marking.

The rank-one order of Proposition 2.33 has ordinary basis $1,d$ and return $j=1-(1+q)d$, with $j^2=1$ and $ja=\iota(a)j$. In the generic microlocal embedding write $j=b(q)\iota$. Since $j\equiv 1 \pmod{q+1}$ as an integral operator,

$b = j(1) \equiv 1 \pmod{q+1}$ is a local unit. Then

$$\frac{1-b\iota}{1+q} = b\frac{1-\iota}{1+q} + \frac{1-b}{1+q}$$

shows that its local lattice is the span of 1 and $(1-\iota)/(1+q)$. As $q - q^{-1} = (q+1)(q-1)/q$ and $(q-1)/q$ is a local unit, this is exactly the ordinary divided-difference basis of $\mathrm{End}_S(A)$ at $q = -1$. Thus the specified injection is an isomorphism at both omitted points and on the regular locus; its finite cokernel is zero globally. □

3.5.5. *Integral normalization and extension.*

Proof of Theorem 3.26. The preceding lemmas identify both ordinary full-averaging algebras with $\mathrm{End}_S(A)$ under the same generic map. If $D = (1-\iota)/(q-q^{-1})$, the integral matrices in the S-basis $(1, q)$ are

$$m_q = \begin{pmatrix} 0 & -1 \\ 1 & c \end{pmatrix}, \quad D = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}, \quad \delta = 1 + q^{-1}D = \begin{pmatrix} 1 & c \\ 0 & 0 \end{pmatrix}.$$

In particular $D = m_q(\delta - 1)$, $m_q D = e_{22}$, $1 - m_q D = e_{11}$ and $m_q e_{11} = e_{21}$ exhibit every matrix unit. The normalized return is $\iota(q) = q^{-1}$. If one uses the folded root $2(\alpha_s + \alpha_t)$, its divided-difference generator is $d = 1 - \delta = (1-\iota)/(1-q^2)$; hence $1 - (1-q^2)d = \iota$. The trace itself is not this generator.

The other five coefficients of the six-element A_2 triangular tilting family of Proposition C.1 are identified with A by their actual counits: the nonidentity standards in their counit fibers are not θ-fixed, so their closed real coefficients vanish, and the surviving identity has residue 1. Proposition C.1 gives the same residue on the spectral side. Naturality and products agree generically and therefore integrally in their torsion-free ordinary targets. Finite tilting extension and the discrete-mapping-space argument give the marked unital comparison. Only then are the original circle invariants and original u used. The already proved exact central descent transfers this result from the $SL_3/PU(2,1)$ model to its coupled central forms and spectator tori. □

3.6. **The integral comparison and the main theorems.** The regular comparison of Theorem 2.53 holds for every quasi-split datum. To extend it across the discriminant, we compare the ordinary coefficient lattices at height one. Only three local root systems occur: a fixed A_1, two orthogonal A_1 factors exchanged by the involution, and an A_2 whose simple roots are exchanged. Their marked comparisons are respectively Proposition 2.34, Appendix A, and Theorem 3.26. The essential normalization issue is that these local comparisons must extend one common generic map. We address it before the extension argument.

3.6.1. *Whittaker normalization and stable Levi subgroups.*

Lemma 3.30. *For the real adjoint-derived-times-torus quotient, the ordinary microlocal functor on compact finite Hecke kernels agrees with the universal Whittaker functor with its two Cartan actions. The identification intertwines the convolution tensor constraints.*

Proof. We identify the two tensor constraints on the opposite cell. Write $Y = H/N_H$, let N^- be opposite to N_H, and put $r = \dim T_H$. Use relative position coordinates $g_1^{-1}g_2$ for an ordered pair of enhanced flags. After fixing the ordinary first flag, N^- is a slice to the diagonal, with base point 1. The restriction of $(-\xi, \xi)$ to this slice is ξ. Choose the opposite unipotent so that this restriction is df_1 for a homomorphism $f : N^- \to \mathbb{C}$ nonzero on each negative simple root space. Indeed in $\mathfrak{n}_H$ a regular nilpotent is conjugate under N_H to the sum of its nonzero simple-root components: the open B_H-orbit is the regular orbit, and the Cartan factor of a conjugating element fixes the higher components at zero. Changing the opposite by this unipotent does not change the abstract Cartan or either frame. Use the real part of f for the positive Morse test.

On N^- the graph of df meets the possible singular support only at 1. To see this, express a cotangent vector using right translation of N^-. A lift from the conormal to a B_H-orbit of flags has moment in $\mathfrak{n}_H$. Pairing with $\mathrm{Lie}\, N^-$ determines that moment uniquely, and df is translation invariant, so the moment must be ξ. The flag at the other end then contains ξ in its nilpotent radical. The Springer fiber of ξ is the single first flag. The slice meets this flag only at 1. At that point the relevant conormal is smooth: off the diagonal two distinct flags cannot contain the same regular nilpotent in their nilpotent radicals, and regularity is open. Thus the transverse test computes the stalk of the microlocalization at the stated covector. Keep the frame coordinates in this calculation. In Taylor's notation $i : N^- \hookrightarrow B^- = N^- T_H$ fixes the relative frame. Local constancy in the complex torus directions and their complex orientation give $i^! = i^*[-2r]$ on the restricted kernels. Consequently Taylor's formula $\phi_{f,1}(i^!K|_{B^-})[r]$ after Definition 3.3 of [Tay25] is precisely this normal Morse group shifted by $[-r]$. This gives a natural identification on kernels. Use the absolute normal coordinate $n \mapsto nt_2$ for the second flag when the frames are t_1, t_2. In this coordinate the test remains $f(n)$. Moving the two frames replaces $g_1^{-1}g_2$

by $t_1^{-1}g_1^{-1}g_2t_2$, with the covector transported at the same time. It gives the inverse first right action and the positive second right action, including on loops. These are Taylor's left and right monodromy.

For a pair of composable changes write n_1, n_2 for the successive slice coordinates with the middle frame fixed. Their composite is n_1n_2 and $f(n_1n_2) = f(n_1) + f(n_2)$. In the local-cohomology description of the Morse test, the external-product map is induced by the inclusion of supports

$$\{\operatorname{Re} f(n_1) \geq 0,\ \operatorname{Re} f(n_2) \geq 0\} \ \subset\ \{\operatorname{Re}(f(n_1) + f(n_2)) \geq 0\}.$$

The microlocal exterior-product map uses this inclusion. In the corepresenting description of vanishing cycles it is the Thom–Sebastiani map in [Tay25, (3.1)]. The middle change of frame acts by $(b_1, b_2) \mapsto (b_1t^{-1}, tb_2)$. Descent along this torus applies the same diagonal group cochains on both sides. In an oriented basis their Koszul complex, shifted by r, gives the derived tensor product over R with the indicated inversion. This is the averaging map in [Tay25, (3.2)]; there the outer shift by r uses the same complex orientation. Finally push along multiplication to the relative output chart. The microlocal proper-image map is obtained by applying pushforward to the supported sections and the counit. Its restriction to this chart is the change-of-support and adjunction map in [Tay25, (3.3)]. In making the restriction, Lemma 3.6 of loc. cit. eliminates the complementary opposite Bruhat strata: a negative simple root subgroup acts trivially there, and its nonzero image under f makes the resulting direct image on $\mathbb{C}$ translation invariant. Equivalently the proper microlocal correspondence at $(-\xi, \xi)$ has only the first ordinary flag as intermediate flag. Thus the identification intertwines the two binary maps on their defining operations, and it sends the unit to the regular R-module with its unit vector. For three successive changes the two support inclusions factor through the same condition on $\operatorname{Re}(f(n_1) + f(n_2) + f(n_3))$. The two integrations over the middle frames are pushforward from the same ordered torus product. Composition of pushforwards gives its oriented integration, with the usual tensor-product signs. Longer strings use the identical support and torus operations in each arity. All these are finite local Morse or link computations and integrations over compact frame tori, so they retain the universal coefficient rings. □

Lemma 3.31. *For the same quotient, the ordinary microlocal image of parabolic averaging of type J is the R-bimodule $R \otimes_{R^{W_J}} R$. Its counit is multiplication, and its comultiplication sends $a \otimes b$ to $a \otimes 1 \otimes b$. The construction commutes with restriction to the standard Levi, retaining the entire Cartan.*

Proof. Use the tensor identification of Lemma 3.30. Let $L \dashv U$ be Whittaker induction and restriction, let ω be the Whittaker unit, and let $\Xi = L\omega$. The definitions and adjunction of [DT26b, §12.1] identify the averaging comonad with LU. The universal Whittaker functor is represented by Ξ, and its bimonodromy isomorphism identifies $\operatorname{End}(\Xi) = R \otimes_{R^W} R$, with $1 \otimes 1$ corresponding to id_Ξ [Tay25, Theorem 1.2]. Write η for the unit of this adjunction. After forgetting the left R-action the right adjoint is the Whittaker functor, and the left adjoint takes M to $\Xi \otimes_R M$, with its Hecke-module structure [DT26b, Definitions 12.1.1–12.1.3 and Lemmas 12.1.4–12.1.5]. The element corresponding to id_Ξ is $\eta_R(1)$. The comultiplication is $L(\eta_R)$. Its adjoint is $UL(\eta_R) \circ \eta_R$, which the tensor identification above sends to $\eta_R(1) \otimes \eta_R(1)$. Thus the bimodule comultiplication takes $1 \otimes 1$ to $1 \otimes 1 \otimes 1$; bimodule linearity determines all its values. The counit corresponds to the identity map of R, so the induced bimodule counit is multiplication. The prescribed tame equivalence identifies this averaging adjunction with the proper flag adjunction $p^* \dashv p_*$, by [DT26a, Theorem 1.1 and the proof of Lemma 4.3]. Thus these are the coalgebra maps of the actual parabolic kernel, not coalgebra maps selected after an abstract bimodule isomorphism.

On the constructible side let $\check{G}$ be the complexification of the real quotient and use a check accent for its corresponding subgroups. For a standard Levi, use the closed embedding $\check{P}_J/\check{U} \simeq \check{L}_J/\check{U}_{L_J} \longrightarrow \check{G}/\check{U}$. Its intersection with the opposite big cell is $\check{N}^-_{L_J}\check{T}$. The Whittaker function restricts to the Levi Whittaker function; the Cartan shift is unchanged because the full $\check{T}$ is retained. Closed adjunction therefore identifies the Whittaker realization of the extended Levi kernel with its Levi realization. Restricting opposite-cell multiplication and the same oriented Cartan integration proves tensor compatibility. Applying the preceding adjunction argument within L_J gives $R \otimes_{R^{W_J}} R$ and its displayed coalgebra. The ordinary universal Struktursatz [Tay25, Theorem 9.4] makes this compatibility faithful on kernel morphisms. Its connected-center hypothesis is satisfied by the real quotient under consideration and by its standard Levis. □

Lemma 3.32. *After passage to the real adjoint-derived-times-torus quotient, the ordinary microlocal functor of Lemma 2.50 identifies the averaging coalgebra for a standard parabolic type J with $R \otimes_{R^{W_J}} R$, with counit $\epsilon(a \otimes b) = ab$ and comultiplication $\Delta(a \otimes b) = a \otimes 1 \otimes b$. These identifications are compatible with restriction to a standard Levi, with its full torus retained. The simple units in Lemma 2.51 may consequently be chosen to*

act as the actual reflections on R. With this choice, the generic comparison of Theorem 2.53 agrees with the ordinary operator maps obtained from the same microlocal functor, including the map (3.56).

Proof. Lemma 3.31 gives the first assertion. To fix the generic normalization, apply the calculation to a simple reflection.

For $J = \{s\}$, the ordinary coefficient algebra is now the actual operator algebra $\mathrm{End}_{R^s}(R)$. The operator s is integral, satisfies $s^2 = 1$ and $sa = s(a)s$, and reduces to the identity modulo the maximal ideal at the boundary identity. The last assertion also follows from

$$s = 1 - (1 - e^{\alpha_s})\frac{1 - s}{1 - e^{\alpha_s}};$$

the divided difference is an integral operator. Hence s is one of the permitted units of Lemma 2.51. Fix these units once and for all in Proposition 2.52. Their products are the actual Weyl operators, so the dual graph frames are exactly the frames used by the ordinary operator realization. The normalization $D_w(1) = 1$ in Theorem 2.53 then identifies each return with the corresponding semilinear Weyl operator on the original A. This proves the last assertion without an independent change of frames in a local chart. □

Lemma 3.33. *Let θ preserve a pinned Borel pair (B, T), let J be θ-stable, and write $P = LU$ for the corresponding standard parabolic, with $B_L = B \cap L$. Choose a θ-fixed cocharacter λ which is zero on J and strictly positive on the roots of U. Retain the full torus, the given odd element τ, and its original square. Put*

$$\mathcal{V}_G = [(G\tau \times^{\mathbf{R}}_{G} B)/B], \qquad \mathcal{V}_L = [(L\tau \times^{\mathbf{R}}_{L} B_L)/B_L], \qquad \mathcal{S} = [(P\tau \times^{\mathbf{R}}_{P} B)/B].$$

The maps $k : \mathcal{S} \to \mathcal{V}_G$ and $q : \mathcal{S} \to \mathcal{V}_L$ give a fully faithful functor

$$k_* q^* : \mathrm{Coh}(\mathcal{V}_L)_{\lambda=0} \longrightarrow \mathrm{Coh}(\mathcal{V}_G).$$

It preserves bounded coherence and sends Q_L to Q_G. For $D \subset J$, it intertwines the actual D-parabolic pull–push kernels, their units and counits, and their compositions. All these identifications preserve the given boundary-circle action. No freeness of the action of U is asserted.

Proof. On the atlas with the output Borel fixed, the Levi projection identifies $P\tau \times^{\mathbf{R}}_{P} B$ with $U \times (L\tau \times^{\mathbf{R}}_{L} B_L)$. The gauge group is still $B = B_L \ltimes U$. In particular, q has the quotient of this entire smooth U-factor as its fiber; one does not delete its stabilizer. The pullback q^* is flat and preserves bounded coherence.

The functions on U have nonpositive λ-weights, with weight-zero part precisely the constants, while $\mathfrak{u}$ has strictly positive weights. We retain its full Lie-cochain complex below. These are the radical weights in weight-zero full faithfulness [HL21, Theorem 2.2.2(3)].

We also check the normal directions of k, rather than replacing the derived stratum by its classical support. Work in the open $U^- LU \subset G$ containing P. The coordinates x_i on U^- have strictly positive function weights. The square-stability equations in the negative root spaces outside L have the same property; denote their Koszul variables by η_j, in degree -1. Setting all these x_i, η_j to zero gives the actual algebra of $\mathcal{S}$. Write ζ_a for the remaining, weight-zero Levi incidence variables. On a homogeneous affine chart the relevant dg algebra has the form

$$\mathscr{R} = \mathscr{R}_0[x_i, \eta_j, \zeta_a], \qquad d\eta_j = \sum_i a_{ji} x_i, \qquad d\zeta_a = g_a + \sum_i b_{ai} x_i,$$

Here $\mathscr{R}_0$ consists of the smooth degree-zero coordinates on the Levi atlas and on U, and g_a is the corresponding Levi equation. Homogeneous affine localizations at weight-zero functions are allowed; their restrictions to the stratum depend only on the Levi coordinates. All a_{ji}, b_{ai} are ordinary degree-zero functions, and $\mathscr{S} = \mathscr{R}/(x_i, \eta_j)$. First adjoin p_i with $|p_i| = -1$ and $dp_i = x_i$. Only in this enlarged dg algebra make the triangular changes

$$\eta'_j = \eta_j - \sum_i a_{ji} p_i, \qquad \zeta'_a = \zeta_a - \sum_i b_{ai} p_i.$$

They satisfy $d\eta'_j = 0$ and $d\zeta'_a = g_a$. An $\mathscr{R}$-semifree resolution of $\mathscr{S}$ is obtained by additionally adjoining $|u_j| = -2$ and $du_j = \eta'_j$. Its augmentation is a quasi-isomorphism: cancel the x_i, p_i pairs, and then the η'_j, u_j pairs. To apply this calculation to arbitrary bounded coherent E, F, write D_L for the Levi dg algebra and take a bounded-above, degreewise finite semifree resolution $P_E \to E$ over D_L, with all its free generators in weight zero. After the changes of variables above, the enlarged algebra $\widetilde{\mathscr{R}} = \mathscr{R}[p_i, u_j]$ is, up to these localizations, $D_L \otimes \mathcal{O}(U) \otimes \mathbb{C}[x_i, p_i, \eta'_j, u_j]$. The last factor contracts to its constants; localization preserves this quasi-isomorphism. Thus $\widetilde{\mathscr{R}} \otimes_{D_L} P_E \to q^* E$ is an $\mathscr{R}$-semifree resolution. Apart from the generators of P_E, its free generators have strictly positive λ-weights. The target $q^* F$ has only nonpositive weights, so the

weight-zero part of its Hom complex keeps exactly the normal constants and the constants in $\mathcal{O}(U)$. It is the full complex $\mathrm{RHom}_{D_L}(E,F)$; arbitrarily high Levi Ext groups have not been truncated. This computation may be made after forgetting the U-action. One then applies rational Chevalley–Eilenberg cochains to the actual equivariant Hom object [Kas15, Section 2, Theorem 2.1], using their finite-width cohomology spectral sequence; no artificial U-action on an individual normal chart or resolution is asserted. Its positive exterior degrees have negative weight, so the weight-zero Hom is exactly the Levi Hom. For a fixed ordinary degree only finitely many monomials in the negative-degree resolution generators can contribute to Hom into a bounded coherent target. There is no completion or infinite product of weight spaces in this argument. Taking exact λ-weight zero and then the full B_L-descent, and gluing over the charts, gives full faithfulness of k_*q^*.

The map k is a closed derived stratum, so its pushforward preserves bounded coherence. The odd stable-Borel locus belongs to this stratum; its square with the Levi stable-Borel locus is Cartesian over q. Flat base change consequently gives $k_*q^*Q_L = Q_G$. If $D \subset J$, two flags in one D-partial flag have the same containing J-parabolic. The full D-flag correspondence thus restricts to $\mathcal{S}$, and its Levi quotient is the corresponding Levi correspondence. These are derived Cartesian squares with the same U-gauge retained. Proper base change and the projection formula intertwine the pull–push functors. The comparison is constructed from these squares, so it also intertwines their adjunction units and counits. Applying the same construction to every tuple of flags gives the composition and higher coherence diagrams. The cyclic diagrams for the boundary action only repeat and compose transport by a^2; Levi projection preserves a, its square and those compositions. Thus all the maps above are equivariant for the original circle. □

Lemma 3.34. *For a θ-stable standard type J, the spectral and real stable-Levi embeddings preserve the closed object, the original marked Cartan, and every J-parabolic kernel with its adjunction maps. They intertwine the prescribed finite boundary equivalences.*

Proof. Apply Lemma 3.33 to the inherited odd extension of L_J, keeping its full torus. Its weight-zero functor sends Q_{L_J} to Q_G and intertwines every J-parabolic correspondence with its adjunction maps. In particular, this comparison retains the unipotent stabilizers of the filtered component. To compare intrinsic central squares, put $\tau_{L_J} = (r_G r_{L_J}^{-1})^{-1}\tau_G$. The marking is unchanged, since $[t(r_G r_{L_J}^{-1})r_{L_J}] = [tr_G]$.

Let $P_{H,J} = L_{H,J}U_{H,J}$ be the corresponding real parabolic. The stabilizer of the standard real partial flag is $P_{H,J}(\mathbb{R})$. Over its real orbit the enhanced full-flag space is the associated bundle with fiber $P_{H,J}(\mathbb{C})/N_H \simeq L_{H,J}(\mathbb{C})/(L_{H,J} \cap N_H)$. Taking the quotient by $H_0(\mathbb{R})$ gives the quotient of this fiber by $P_{H,J}(\mathbb{R})$. The real unipotent radical acts trivially on it and is contractible by the exponential map. The homotopy quotient therefore agrees with the quotient by $L_{H,J}(\mathbb{R})$ for equivariant Betti sheaves. The Levi contains the full enhanced torus. Every J-change of full flag keeps the partial flag fixed, so applying this description to a pair or a triple of flags identifies the kernel, the two adjunctions, and their composites. Extension along the partial-flag orbit sends the closed Levi object to the ambient closed object, with the same Cartan action.

For compatibility with the specified boundary equivalence, the same-parabolic Steinberg correspondence is Cohen–Macaulay and has equidimensional Cartan fibers, as is seen from its relative-position strata. It is Cartan-flat by miracle flatness. The flat-quotient argument of Lemma 2.25 therefore identifies the coherent image of parabolic averaging with its actual structure-sheaf kernel C_J; the quotient and the three-flag correspondence fix its coalgebra maps. Lemma 3.32 now gives the same ordinary Soergel realization for the two functors on Levi Bott–Samelson tiltings. Taylor's full faithfulness identifies their maps and tensor constraints. Extension through bounded complexes and retracts gives the asserted compatibility for all finite kernels. The exact central descent in Section 3.4 restores arbitrary central forms. □

3.6.2. *The height-one local models.*

Lemma 3.35. *In each of the three stable orbit blocks A_1, $A_1 \times A_1$ and A_2, the marked comparison can be chosen to extend the generic operator map fixed in Lemma 3.32. Its return can be chosen integrally to act as the actual folded reflection on the raw odd torus.*

Proof. For a fixed simple root, put $h = 1 - e^{\beta}(t)$, where $t = yr^{-1}$ is the raw odd coordinate. Proposition 2.33 gives a basis $1, d$ and a return $j = 1 - hd$, with $j^2 = 1$. In the ordinary microlocal operator realization its generic action is $b\varphi_s$. Here $b = j(1)$ belongs to A, and $j^2 = 1$ gives $b\varphi_s(b) = 1$, so b is a Laurent unit. Since d is an integral operator, $b \equiv 1 \pmod{h}$. Consequently

$$\frac{1 - b\varphi_s}{h} = b\frac{1 - \varphi_s}{h} + \frac{1 - b}{h}$$

shows that the real coefficient order is precisely $A\,1 + A(1-\varphi_s)/h$. The spectral rank-one calculation of Proposition 2.34 gives this same order in the same generic operator space. Thus $b^{-1}j$ is an integral return acting as φ_s. This does not assume that the original rank-one choice already had $b = 1$. The comparison also preserves the adjunction maps: the counit gives 1, and the opposite rank-one functional kills j and takes 1 to h, so it is unchanged by this normalization.

For an exchanged orthogonal pair, the ordinary two-flag equivalence of Appendix A identifies the coefficient calculation with that of a single complex rank-one Hecke category. The full torus, the proper flag adjunction and the Whittaker unit are retained; hence Lemma 3.32 identifies its return with the actual reflection. The original circle orientation is the one of Lemma A.3. For an adjacent pair, Theorem 3.26 identifies the actual order with $\mathrm{End}_{\mathbb{C}[q+q^{-1}]}\mathbb{C}[q^{\pm 1}]$ under that same generic map, so the integral operator $q \mapsto q^{-1}$ is the required return. Stable-Levi transport and central descent preserve these assertions and their original markings. □

Lemma 3.36. *Suppose the dual derived group is simply connected. At every height-one prime of the marked output ring A, the fixed generic comparison of Theorem 2.53, normalized as in Lemma 3.32, extends over a faithfully flat local etale chart to a marked comparison of the local finite Hecke actions.*

Proof. Only primes on the pulled-back root discriminant require consideration. The character of a boundary root α under the norm is $\alpha + \theta\alpha$ on the raw odd torus; its central-square factor is trivial. Two independent such characters cannot vanish on the same prime divisor. Thus the roots vanishing at its generic point lie on a single ray of theta-average roots. These rays form the folded root system, whose Weyl group is $\Gamma = W^\theta$ [Hai18, Definition 3.3 and Remark 3.4]. Move the ray by Γ to the ray of a simple orbit O. For a positive absolute root α on that ray, every simple coefficient of $\alpha + \theta\alpha$ is nonnegative. Its support is contained in O, so the support of α is contained in O. The absolute block is consequently A_1, $A_1 \times A_1$, or A_2; the last case is the adjacent simple orbit described in [Hai18, Lemma 3.2]. In the A_2 block the norm characters are q, q, q^2. The branch $q = -1$ is retained inside that full stable A_2 Levi, even though only its highest root vanishes there.

Semisimple centralizers are connected when the derived group is simply connected [Ste68]. Thus the full Weyl stabilizer at the generic divisor is generated by the vanishing roots and is contained in W_J, where $J = O$ is the stable type just selected. The purely complex corner construction in the proof of Theorem 2.61 applies without a splitness assumption: $T/W_J \to T/W$ is etale at this point; after shrinking an etale neighborhood, the Cartan algebra splits into the selected J-factor and its disjoint translates. The selected Hecke corner is the scalar extension of the J-Hecke category. The other corners are transported by invertible outside-root graph kernels. This description retains the derived central fiber and all its nilpotent normal directions.

We check the closed seeds and the generic normalization in this matrix description. A closed norm branch in a coset wW_J requires $\theta(w)W_J = wW_J$, since both its norm and the selected norm are θ-fixed and the stabilizer lies in W_J. Conversely a stable coset has a unique minimal representative fixed by θ, and this representative transports an actual closed branch. A reduced folded gallery to it crosses only outside-J roots: its absolute word is reduced, and minimality excludes inversions in Φ_J. For each gallery step use the integral return normalized in Lemma 3.35 for the relevant fixed A_1, exchanged orthogonal pair, or adjacent A_2 comparison. These are available on the whole low-rank Cartan, so other ambient roots may remain singular. Lemma 3.34 transports them to the ambient chart. Under Lemma 3.32 each return acts by its actual folded reflection. Its product along any chosen gallery is therefore the already fixed generic Weyl operator. No additional folded Coxeter-normalization argument is needed.

In the selected corner apply the marked comparison for J. Transport every other corner by the chosen invertible graph kernels. Units and counits of these inverse kernels preserve composition and all higher convolution diagrams. The preceding operator calculation shows that the resulting comparison restricts to the same generic map on every corner. Formula (3.50) records the original input mark; the output mark is unchanged. Thus all closed summands, including their full completed factors, are retained.

Finally, this etale chart is valid even when the norm image has positive codimension in T/W. Let $S_0 = \mathrm{im}(\mathcal{O}(T)^W \to A)$. The norm is finite onto its image: the image of its character map contains twice the invariant character lattice. Hence A is finite over S_0, and localization at $\mathrm{Frac}(S_0)$ gives $\mathrm{Frac}(A)$. If S' is the chosen etale chart over $\mathcal{O}(T)^W$ and $\mathfrak{q}$ is a lift of the height-one prime $\mathfrak{p}$, the map $A_{\mathfrak{p}} \longrightarrow (A \otimes_{\mathcal{O}(T)^W} S')_{\mathfrak{q}}$ is flat and local, hence faithfully flat. Choosing the chart first at the geometric generic point and then spreading it out supplies such a lift. This proves the asserted local form. □

3.6.3. *Circle actions and the finite ordinary criterion. The finite geometric circle input.* The trivial action on finite kernels used below is fixed before any comparison of ordinary coefficient modules. Let $I^u = \mathrm{ev}_0^{-1}(N_H)$.

Rotation $g(z) \mapsto g(az)$ on the positive-loop substack $I^u\backslash H[[z]]/I^u$ commutes with evaluation at zero. The constant-loop inclusion

$$N_H\backslash H/N_H \longrightarrow I^u\backslash H[[z]]/I^u$$

is equivariant for the *strictly trivial* action on its source: rotation fixes the constant group H, both constant copies of N_H, and both constant Cartan groups pointwise. Evaluation removes only the contractible pro-unipotent kernel contained in I^u and is an equivalence for the specified weakly constructible sheaf theory. Thus the equivalence is equivariant on the full quotient-stack nerves, not merely on their orbit spaces. In particular it keeps the universal Cartan local systems and both actions of the uncompleted group ring. Evaluation and inclusion identify the entire convolution bar diagrams, since evaluation is a group homomorphism and rotation commutes with it. The induced trivial circle action is therefore monoidal with its specified unit and all coherences. The circle enhancement of Proposition A.11 transports precisely this action to the images of finite kernels on the spectral side.

Lemma 3.37. *Let a $\mathbb{C}$-linear stable category $\mathcal{D}$ carry its given circle action, and let Q be an equivariant object. Suppose its* entire derived *endomorphism algebra is the ordinary algebra* $\mathrm{RHom}_{\mathcal{D}}(Q,Q) \simeq A[0]$. *Let an ordinary additive monoidal category of tilting kernels $\mathcal{U}$ act on $\mathcal{D}$, compatibly with the given circle action, and fix the trivial circle action on $\mathcal{U}$ supplied by its finite geometric model. Suppose $M_K = \mathrm{RHom}_{\mathcal{D}}(KQ,Q)$ lies in degree zero for every $K \in \mathcal{U}$. Then the full coefficient diagram, including the two A-actions, its unit, kernel naturality and all coefficient products, has a unique trivialization of its circle action with the underlying diagram fixed. This is a coherent A-linear statement, not only a statement about the action on cohomology. Its invariant and Tate coefficients are consequently $M_K \otimes_{\mathbb{C}} \mathbb{C}[u]$ and $M_K \otimes_{\mathbb{C}} \mathbb{C}[u,u^{-1}]$ for the original positive class u.*

Proof. The equivariance of Q gives its derived endomorphism algebra an actual action in E_1-algebras. Composition gives each M_K its equivariant output action by that algebra, and the input action is obtained by applying K to endomorphisms of Q. These actions and their commuting data are defined before taking cohomology.

The full subcategory of ordinary algebras inside connective derived algebras has discrete mapping spaces [Lur17]. Hence $\mathrm{Aut}_{E_1}(A[0])$ is a discrete group. An action on this specified algebra is a pointed map $BS^1 = K(\mathbb{Z},2) \longrightarrow B\,\mathrm{Aut}_{E_1}(A[0])$. The target is a 1-type, while the source is simply connected; the space of such pointed maps is contractible. Thus the actual algebra action is coherently trivialized, with the given identification of A retained. The module actions just constructed now give coherent circle actions on M_K as A-modules. This supplies the A-linearity that boundary R-linearity alone would not supply.

For completeness, the same argument applies simultaneously to the whole diagram, rather than separately to its values. Mapping spaces between modules in degree zero are discrete. Moreover, for any finite list of such modules, their derived tensor product is in $D^{\leq 0}$; its mapping space to a module in degree zero is again discrete, since negative Ext vanishes by the t-structure. Therefore all mapping spaces used for the A-actions, the coefficient products and the unit are discrete. Natural transformations and compatibility data form limits of these spaces and remain discrete. In particular the automorphism space of this entire ordinary coefficient diagram is a discrete group. Applying the preceding pointed-map argument to this group gives a unique simultaneous trivialization. It preserves both A-actions, not just the boundary subring.

An isomorphism between two such ordinary diagrams is automatically compatible with these trivializations: the corresponding space of equivariant lifts with that underlying isomorphism is contractible. This proves the assertion needed for the integral coefficient comparison. We have not applied this argument to the automorphism space of an entire stable category, which need not be discrete; that separate circle enhancement requires its monodromy argument.

Finally, invariants of the trivial action on $M_K[0]$ are $\mathrm{RHom}_{\mathbb{C}}(C_*(BS^1), M_K)$. Since M_K is in degree zero, only one cochain degree contributes to each total degree, giving $M_K \otimes \mathbb{C}[u]$ without a power-series completion. The identification is multiplicative on the whole diagram. Localizing this small invariant diagram at the actual class u gives the claimed Tate diagram; no circle or generator has been rescaled. □

Theorem 3.38. *For each $w \in W$, choose an unshifted finite tilting T_w supported on the Bruhat interval $[1,w]$, whose top-standard multiplicity is a nonzero finite free module. Suppose* $\mathrm{RHom}_{\mathrm{PE}_G}(\Phi(T_w)Q, Q)$ *is concentrated in degree zero and finite reflexive over the original marked output algebra A. Then the categorical comparison holds for the quasi-split datum. The resulting equivalence preserves the closed object, both Cartan actions, the original positive period, and all convolution coherences.*

Proof. Lemma C.11 shows that every finite tilting is an additive retract of a finite direct sum of the T_w. Consequently every averaging word has an ordinary finite reflexive coefficient. On the real side, Proposition 2.12

and a standard filtration make every word coefficient finite free and ordinary: a theta-fixed standard contributes one A-line and a nonfixed standard contributes zero. Hom complexes between word kernels are ordinary by standard–costandard adjunction.

Use the exact central quotient of Section 3.4; its dual has simply connected derived group. Theorem 2.53, with the choice fixed in Lemma 3.32, identifies the two ordinary coefficient spaces over $\mathrm{Frac}(A)$. At each height-one prime, Lemma 3.36 gives an integral local comparison extending precisely this map. Taking degree zero recovers the ordinary lattices from the local comparison because both coefficients are ordinary; their circle invariants and periodization are the tensor product with the original $\mathbb{C}[u, u^{-1}]$. Faithfully flat descent gives equality of the lattices over $A_{\mathfrak{p}}$.

The Laurent ring A is normal. A finite reflexive module inside its fraction vector space is the intersection of its height-one localizations. Indeed its double dual consists exactly of vectors on which all integral linear forms have values in A, and membership in A is detected at height one. The common generic comparison therefore identifies the coefficient modules integrally for every word. Naturality for kernel morphisms, binary coefficient products and the unit hold generically and take values in torsion-free targets; they hold integrally. The unit is still the identity of the original A.

Here is the coherence argument before periodization. The additive category of word tiltings and its coefficient diagrams have ordinary mapping complexes. Their mapping spaces are discrete, so the verified associativity and unit diagrams specify all higher coherences. Separately exact extension in each variable through bounded complexes and retracts preserves them, as in Proposition A.10. The prescribed boundary equivalence is equivariant for the original positive circle by Proposition A.11. The ordinary calculation in the proof of Lemma 2.9 gives the full derived algebra $\mathrm{RHom}(Q, Q) = A[0]$ on both sides, before invariants. Lemma 3.37 therefore supplies the coherent A-linear trivialization of the entire ordinary coefficient diagram. The integral comparison respects it, including both Cartan actions, products, and the unit. Taking their invariants and inverting the original u yields the marked unital coefficient-algebra equivalence required by Proposition 2.17. This gives the finite Hecke-module equivalence. The deck-equivariant faithful descent proved in Section 3.4 returns it to the original central form. □

3.7. **Proof of the main theorems.** We now assemble the ordinary and categorical arguments. The two inductions are distinct. Parabolic source arguments use connected cores contained in a proper stable parabolic. Semisimple localization uses the connected factors of a proper norm centralizer; such a centralizer can have full semisimple rank, so rank alone would not order that induction. The list in Section 3.4.6 provides the required order.

Proof of Theorem 1.7. Use the central descent of Section 3.4 to pass to the simply connected derived cover, retaining its full torus and finite inertia. Torus factors have no root incidence. Exchanged simple factors are covered by the two-flag comparison of Theorem A.1. Thus it remains to treat a connected root system with the trivial diagram action, or one of the order-two outer actions in types A, D, E_6.

The common-Levi source theorem, Theorem 3.4, allows an arbitrary middle parabolic. Each of its connected cores is contained in a proper stable end and has smaller derived rank. The different-end theorem, Theorem C.10, supplies the additional sources with four parabolic averaging factors; its neutral terms are cross-Hom spaces between vector bundles and require no induction on arbitrary words. The outer A_2, A_3, A_4 bases in Section C.1, the uniform type-A argument, and the classical maximum principle in Section C.2 therefore prove all classical central blocks. Their proper-centralizer dependencies are classical and precede them in the order of Section 3.4.6. Proposition 3.25 gives the statements over their whole Cartans.

For the exceptional types with trivial diagram action, Section C.3 proves G_2, F_4, E_6 by the source and endpoint rules with geometric diagonal anchors. For the outer E_6 action, Theorem C.14 uses the proper sources and the exact nonnegative elimination of the complete endpoint system. Every connected factor needed by those sources or by their proper centralizers belongs to an earlier row.

Section C.5 proves the two remaining central calculations. In type E_7, the conservative support guards for higher Kazhdan–Lusztig terms license only implications whose other terms are already zero. Their exhaustive replay leaves no marked coefficient unresolved. In type E_8, the equality forests reduce the length induction to one component at length 64. The chosen label is an involution and moves its chosen raw diagonal mark. Proposition 3.21 therefore annihilates its standard coefficient in every ordinary degree. By the shorter-length hypothesis, Corollary 3.22 annihilates the positive cohomology of its canonical tilting coefficient. The verified tree propagates this zero to all 35 419 states in the component. Induction now supplies the shorter-length premise for

each later forest layer. Thus every marked central block in E_7 and E_8 is ordinary. Their proper norm centralizers have connected factors already treated above, so Proposition 3.25 applies again.

These local statements concern summands of the ordinary completed tilting category, by Proposition 3.16. Every finite averaging word decomposes into them at the corresponding central completion. The marked centralizer formula of Proposition 3.24 retains the output algebra and transports the input idempotent by (3.50). Hence all positive cohomology of each original word coefficient vanishes at every output point. Its cohomology modules are finite by Proposition 3.1; local detection and faithful flatness of completion give global positive vanishing.

The independently proved nonnegative output-fiber statement in Proposition 3.1 now applies. Lemma 3.3 shows that the coefficient is finite projective in degree zero over the original output ring A. Ordinary tilting generation and idempotent completion extend this conclusion from finite words to all kernels in the statement. Exact central descent restores every central form, coupled character lattice, and torus factor. □

Proof of Theorem 1.1. Let D_C and D_Q denote the internal endomorphism objects of the closed real generator and the stable Borel object for the two finite Hecke actions. We compare them as algebras. On a compact kernel K, Yoneda evaluates their underlying objects as the coefficient complexes for (KP_C, P_C) and $(\Phi(K)Q, Q)$. In each case the product for two kernels sends a pair of maps $a : KX \to X$, $b : LX \to X$ to $a \circ K(b) : (K * L)X \to X$, using the action constraint. Lemma 2.14 realizes this product by restriction of the correspondence of three flags. The identity of X gives the unit. Theorem 2.10 fixes the normalization at the closed generator on the original output ring A.

Take an unshifted averaging word. The real coefficient is ordinary and finite free by its standard filtration and Proposition 2.12; the spectral coefficient is ordinary and finite projective by Theorem 1.7. The comparison on the regular locus respects kernel morphisms and composition. Fix its generic value with the Whittaker normalization of Lemma 3.32. Write M_C and M_Q for the two ordinary coefficient modules inside their common fraction vector space. If $\mathfrak{p}$ has height one, Lemma 3.36 identifies the localized lattices. On the discriminant the calculation uses the fixed-root comparison of Proposition 2.34, the product comparison of Appendix A, or the adjacent-root comparison of Theorem 3.26. The last calculation treats each of the two ramification points with its own mark. All three maps restrict to the fixed generic comparison, so the height-one identifications give equalities of submodules in the same vector space.

The Laurent ring A is normal and both coefficient modules are reflexive. Consequently

$$M_C = \bigcap_{\operatorname{ht} \mathfrak{p}=1} (M_C)_{\mathfrak{p}} = \bigcap_{\operatorname{ht} \mathfrak{p}=1} (M_Q)_{\mathfrak{p}} = M_Q$$

in positive dimension, and in dimension zero the generic identification already suffices. Naturality for a morphism of words, preservation of a binary product, and preservation of the unit are equalities in torsion-free output modules. They follow from their generic values. On unshifted tiltings the mapping spaces in the coefficient diagram are discrete. Thus the equalities for the ordinary products give associativity in all arities. Exact extension to bounded complexes and passage to retracts preserve these coherences. This is the algebra comparison of Theorem 3.38, including descent of the original output and input markings.

Lemma 3.37 identifies the circle trivializations of the ordinary coefficient diagrams compatibly with both actions and products. Taking small invariants and inverting the original u gives an equivalence of the represented unital algebras $D_C \simeq D_Q$. The compact translates generate on the real side by Theorem 2.7 and on the spectral side by the definition of $\mathcal{S}_Q$. Proposition 2.17 identifies the two module categories from their internal algebras, sending the free module to the free module. Under this identification the distinguished objects correspond, and multiplication gives the prescribed naturality and convolution compatibilities for their translates. □

Proof of Theorem 1.3. For a restriction of scalars the two factors on the dual side are exchanged. Lemma A.3 identifies the odd parameter with an even element and two boundary flags, and identifies the boundary rotation with the canonical rotation on the resulting Steinberg stack. On the real side Lemma A.5 identifies the enhanced-flag quotient with the finite Hecke model. Under the monoidal equivalence of Proposition A.11 the two rotations agree before periodization. Proposition A.16 then identifies the entire Tate category of kernels with the periodized finite Hecke category. Both flag variables give their original Cartan actions by Lemma A.12. The source has one pure inner label by Lemma A.4. These identifications compose to the asserted monoidal equivalence. □

4. Pure inner labels, marked support and flag positions

Let $B_0 = B_H$ and $\Gamma = \mathrm{Gal}(\mathbb{C}/\mathbb{R})$. We prove Theorem 1.2 using the pure torsor labels and the original odd-component marking of (2.3). The global functor uses a choice of extension of a finite character and an inner identification of real forms. Its source labels, its restrictions on the fixed center, and its essential image on the regular semisimple open are intrinsic.

Write $Z = T \setminus U$. For an R-linear category $\mathcal{C}$ we use $\mathcal{C}_Z$ for the full subcategory of objects that become zero upon inverting f. Corollary 6.31 identifies these supported categories after the Hecke subcategories for the different minimal types have been separated.

4.1. Pure inner labels and lines on the parameter stack.

Lemma 4.1. *The inclusion $T_0 \to H_0$ identifies $H^1(\mathbb{R}, T_0)$ with Θ_{qs}. Under torus duality this labels Θ_{qs} by the characters of $\Pi = \pi_0(T^\theta)$.*

Proof. The unipotent radical of B_0, also after twisting by a torus cocycle, has trivial first Galois cohomology. Indeed it has a filtration by real vector groups, for which additive Hilbert 90 applies. The splitting of the Borel by T_0 then gives $H^1(\mathbb{R}, T_0) = H^1(\mathbb{R}, B_0)$. A lift of a torsor of class ξ to B_0 is a reduction to a real Borel. Up to isomorphism these lifts are the orbits of $H_\xi(\mathbb{R})$ on the real points of the twisted flag variety. They exist precisely when the twist is quasi-split. Real Borels of a quasi-split group are conjugate by its real points, giving exactly one lift in this case.

For the second assertion let $\Lambda = X_*(T_{0,\mathbb{C}}) = X^*(T)$. The exponential sequence of $T_0(\mathbb{C})$ has kernel $2\pi i\Lambda$, with Galois action $-\theta$ under the displayed lattice identification. The additive Lie algebra has zero higher cohomology by averaging. Cyclic group cohomology therefore gives

$$H^1(\mathbb{R}, T_0) = H^2(\Gamma, 2\pi i\Lambda) = \ker(1+\theta : \Lambda \to \Lambda)/(1-\theta)\Lambda.$$

This is the torsion subgroup of $\Lambda/(1-\theta)\Lambda$: over $\mathbb{Q}$ the image of $1-\theta$ is the negative eigenspace, and twice each integral anti-invariant belongs to that image integrally. The quotient lattice with its torsion is the character group of T^θ, whose torsion subgroup consists exactly of characters of Π. This is the real-torus pairing of [Kot86, (1.1.1)]. □

To construct the lines globally, let G_{sc} be the simply connected cover of the derived group and $C = Z(G)^\circ$. Form the central isogeny

$$p : G' = G_{\mathrm{sc}} \times C \longrightarrow G, \qquad p(s,c) = \bar{s}\, c^2.$$

Let $D = \ker p$, $T' = p^{-1}(T)$, and let θ act on both factors of G' by the lifted pinned action and its action on C. Adjoin a pinned odd element τ' with square $((2\rho_{\mathrm{sc}}^\vee)(-1), 1)$, using the coroots of G_{sc}. This square is fixed and central, and its image is τ^2. The assignment $\tau' \mapsto \tau$ thus gives a map of extensions. The odd components form a finite étale D-torsor $E'_- \to E_-$.

Lemma 4.2. *For $t \in T^\theta$ and a lift $t' \in T'$ the rule*

$$d(t) = [t'\theta(t')^{-1}] \in D_\theta = D/(1-\theta)D$$

induces an injection $\Pi \hookrightarrow D_\theta$. The map $Z(G)^\theta \to \Pi$ is surjective.

Proof. Changing the lift multiplies the representative by an element of $(1-\theta)D$. The kernel is $p((T')^\theta)$, since a lift with trivial class can be adjusted by D to be fixed. The torus of G_{sc} has character lattice with basis the fundamental weights, permuted by θ. Its fixed subgroup is connected. For $c \in C^\theta$ its image $c^2 = c\theta(c)$ lies in the connected image of the norm from C. Thus the entire kernel lies in $(T^\theta)^\circ$. Conversely, the isogeny induces an isomorphism of Lie algebras with involution, so the connected fixed torus upstairs maps onto $(T^\theta)^\circ$.

For the adjoint torus $T_{\mathrm{ad}} = T/Z(G)$ the simple roots give a permuted basis of its character lattice, so T_{ad}^θ is connected. The map $(T^\theta)^\circ \to T_{\mathrm{ad}}^\theta$ is surjective by exactness of invariants on the complex Lie algebras. Given $t \in T^\theta$, divide by an element of $(T^\theta)^\circ$ with the same adjoint image. The result lies in $Z(G)^\theta$ and represents the component of t. □

Lemma 4.3. *For each $\varepsilon \in \Pi^\vee = \mathrm{Hom}(\Pi, \mathbb{C}^\times)$ there is an invertible line L_ε on V with an action of the boundary circle. Its fixed central inertia character is the pullback of ε along $Z(G)^\theta \to \Pi$. Tensoring by it commutes with finite Hecke convolution. One may take $L_1 = \mathcal{O}_V$.*

Proof. By Lemma 4.2 extend ε to a character χ of D_θ. Such an extension exists because $\mathbb{C}^\times$ is divisible. On E_- take the line associated to $E'_- \to E_-$ with the convention $[da', v] = [a', \chi(d)v]$. To act by $g \in G$, choose a lift $g' \in G'$ and conjugate a' by g'. Changing g' by $d \in D$ changes the result by $d\theta(d)^{-1}$, on which χ is trivial. This gives an algebraic G-linearization by descent. If $c \in Z(G)^\theta$, any lift of c is central in G': its commutator with the connected group G' takes values in the finite kernel. Its action on the line is therefore the scalar $\chi(d(c))$. Pull the line on $[E_-/G]$ back to V.

For the action of the boundary circle the full turn at a is a^2. Its lift $(a')^2$ centralizes a', so this turn acts identically on the associated line. This computation takes place on the finite étale torsor and is an identity of descent maps. Equivalently, on $[E_-/G]$ the boundary gauge a^2 defines the circle action after forgetting the flag. A line on this smooth quotient lies in the coherent heart and its isomorphisms form a discrete space. The full turn being the identity is precisely the condition for the equivariant lift, as in the heart calculation in Lemma 2.40. Equivariant pullback gives the lift on V; the dual line gives its inverse.

In a change of boundary flag both projections remember the same odd element and give the same map to $[E_-/G]$. The two pullbacks of the line are thus canonically identified. The projection formula commutes tensoring by the line with the kernel action. For a string of flags these identifications all come from the same line on the base, so they respect the unit and convolution. They hold before taking circle invariants, hence preserve the class u after localization. □

4.2. **The global embedding and its restriction.**

Lemma 4.4. *The objects of $\mathcal{S}_Q$ have trivial $Z(G)^\theta$-character. The subcategories $L_\varepsilon \otimes \mathcal{S}_Q$ are therefore pairwise orthogonal in both directions.*

Proof. The fixed center acts trivially on the incidence atlas since it fixes both the odd element and the flag. Its action on the stable Borel structure sheaf Q is trivial. We verify that all finite kernels have trivial diagonal $Z(G)$-inertia on Z_G. The universal standard Δ_w is convolution-invertible by the standard inverse formula (2.2). Its endomorphism complex is that of the unit, hence R in degree zero: the regular local system is the free rank-one module on the torus and the unipotent stabilizers are contractible. After periodization its graded endomorphism ring is $R \otimes P$. The degree-zero part R has no nontrivial idempotents, so the image of Δ_w in the periodic finite Hecke category is homogeneous for diagonal central inertia. It remains nonzero upon inverting f: right Cartan multiplication identifies the unit endomorphisms with those of this invertible kernel, and its localized endomorphism ring is $R[1/f] \otimes P$.

Over a regular semisimple holonomy the two ordered flags give Weyl graphs, each with stabilizer T. On each graph the compact-first Tate category of $U \times BT$ has only trivial T-weight. In weight $\lambda \neq 0$, the positive turn is the function $t \mapsto \lambda(t)$ on U and its derived unit fiber is smooth or empty. Thus its coherent category vanishes after inverting the cohomological operator, by the argument of Lemma 2.42. Here the rotation is the Steinberg rotation of Lemma A.3. The center embeds in this torus. It follows that the central character of the global image of each Δ_w is trivial. The Bruhat filtration generates all compact kernels from the universal standards by finite cones and retracts, using the free module generator on each torus stratum. Pullback, tensor product and the proper projection from the flag correspondence therefore preserve the fixed-center character of an object of $\mathcal{T}(V)$. In particular all compact Hecke translates of Q are neutral.

On a quotient with a diagonalizable central group acting trivially on the base, equivariant complexes split by characters and all maps preserve them. The character projections commute with the boundary action and hence pass to circle invariants, inversion of u, and Ind-completion. Lemmas 4.2 and 4.3 distinguish the characters of the indicated subcategories, proving the assertion for their full mapping complexes. □

For each $\xi \in \Theta_{\mathrm{qs}}$ choose an inner isomorphism from H_ξ to H_0 over $\mathbb{R}$. One way to make the choice is to represent the torsor by a T_0-cocycle. Its image in the adjoint torus is a coboundary: the cocharacter lattice of the adjoint torus has a basis of fundamental coweights permuted by Galois, so the calculation in Lemma 4.1 gives trivial H^1. Conjugation by a representative of the coboundary identifies the two real structures, preserving the abstract Cartan marking. Equivalently, right multiplication on the enhanced flag is unchanged by the corresponding left translation. Let $\varepsilon(\xi)$ be the character in Lemma 4.1. Use the single-form functor $\mathcal{E}_0$ of Theorem 1.1 to set

$$J_\xi : \mathcal{A}_\xi^{\mathrm{per}} \simeq \mathcal{A}_0^{\mathrm{per}} \xrightarrow{\mathcal{E}_0} \mathcal{S}_Q \xrightarrow{L_{\varepsilon(\xi)} \otimes -} \mathcal{T}(V).$$

Here 0 denotes the distinguished torsor, for which we take the identity choices. Each functor is fully faithful. Lemma 4.4 gives full faithfulness on the finite direct sum. The comparison with the Hecke action and the markings follows from Lemma 4.3. This proves the global assertion of Theorem 1.2.

4.3. **The support of an individual pure inner form.** Fix a label ξ, and let O run over the orbits of $H_\xi(\mathbb{R})$ on the complex flag variety. Choose a stable real Cartan in a Borel of O, and let δ_O denote its involution in the marked abstract Cartan. Set

$$D_O = (T^{\delta_O})^\circ, \qquad Y_O = T/(1-\delta_O)T, \qquad N_O : Y_O \longrightarrow D_O, \quad [y] \longmapsto y\delta_O(y).$$

Write I_O for the ideal of D_O in R. The support of an R-linear category in this subsection consists of the scheme points t such that extension of scalars to the residue field $\kappa(t)$ is nonzero. This definition applies to the periodic category and keeps the marking on T. Ionov and Yun identify the union of the Weyl translates of the maximally split support in the completed Cartan, and compute their category over the generic field of the invariant base [IY25, Section 7.1 and Proposition 7.11]. We work over the multiplicative Cartan, keeping the finite covers N_O and their behaviour under base change.

Lemma 4.5. *Choose a maximally split Cartan of H_ξ and one of its Borel markings, with marked involution δ_*. Set $D_* = (T^{\delta_*})^\circ$. Then*

$$\bigcup_O D_O = \bigcup_{w\in W} wD_* =: \Sigma_\xi.$$

Each distinct translate on the right occurs for precisely one flag orbit. All remaining D_O have smaller dimension. Each irreducible component of Σ_ξ has dimension $\operatorname{rank}_{\mathbb{R}} H_\xi$.

Proof. Let J be the real Cartan of O, and $S \subset J$ its split part. In the cocharacter space of J, the positive eigenspace of its Galois involution is $X_*(S)_{\mathbb{Q}}$. Choose a maximal split torus $S' \supset S$, and a real maximal torus $J' \supset S'$. Over $\mathbb{C}$, the tori J and J' are maximal tori of the connected reductive centralizer of S. Conjugation in that centralizer identifies their cocharacter spaces by a map which fixes $X_*(S)$. It follows that, after suitable markings, the positive eigenspace for J is contained in that for J'. Markings of complex maximal tori by Borels differ by W. On the abstract cocharacter space of the original group choose a rational positive definite form invariant under the finite group generated by W and θ. This identifies the cocharacter space with its dual, compatibly with every marked involution. Thus the same inclusion holds in the cocharacter space of the dual torus T. Since the two subtori are connected, the inclusion of rational cocharacter spaces gives $D_O \subset wD_*$ for some w. The dimension of D_O is exactly $\dim S$, and it equals the real rank precisely for maximally split J.

The maximally split real Cartans are conjugate under $H_\xi(\mathbb{R})$. Fix one. The Borels containing its complexification form a Weyl torsor. Two of them give the same real orbit precisely when their markings differ by the real Weyl group: stable Cartans in a fixed Borel are conjugate under its real stabilizer. For this Cartan the real Weyl group is W^{δ_*} [IY25, Lemmas 4.3, 4.14 and 4.15]. The stabilizer in W of D_* is also W^{δ_*}. Indeed, a Weyl element preserving the positive eigenspace preserves its orthogonal complement, which is the negative eigenspace, and hence commutes with δ_*. The converse follows directly. This gives the asserted bijection from the maximal orbits to the distinct translates of D_*. □

Proposition 4.6. *The support of $\mathcal{A}_\xi^{\mathrm{per}}$ is exactly Σ_ξ. Its ideal $I(\Sigma_\xi)$ is the radical of the kernel of the central action*

$$R \longrightarrow H^0 \operatorname{End}\bigl(\mathrm{Id}_{\mathcal{A}_\xi^{\mathrm{per}}}\bigr).$$

If the nonempty flag orbit dimensions take m distinct values, then $I(\Sigma_\xi)^m$ annihilates the identity functor.

Proof. Lemmas 2.1–2.3 apply to H_ξ: their proofs use the stable Cartan of each individual orbit, its real stabilizer and finite orbit recollement, for a connected reductive group and its full group of real points. The stratum algebra is $A_O = \mathcal{O}(Y_O)$, with R-action induced by N_O. This map is a finite surjective isogeny onto D_O. Consequently A_O is finite locally free of positive rank over R/I_O. Its ordinary fibre at a point of D_O is nonzero, so its derived fibre there is nonzero as well. Off D_O its derived fibre vanishes. Tensoring with the nonzero flat periodic algebra P preserves these statements.

The continuous R-linear stratum extension and restriction satisfy $j_O^* j_{O!} = \mathrm{Id}$. They retain this identity after extension of scalars. Conversely, regular stratum standards generate before and after scalar extension. It follows that the support of the category is the union of the stratum supports, which is Σ_ξ by Lemma 4.5.

Filter by the closed unions of real flag orbits of dimension at most d. Successive open–closed triangles give a functorial tower for the identity with one graded functor for each of the m nonempty dimensions. Each such functor factors through the stratum restriction and extension functors for the disjoint orbits of that dimension. A scalar in $I(\Sigma_\xi)$ acts as zero on every graded functor since it maps to zero in every A_O. In a triangle $F' \to F \to F''$ of functors, if a scalar kills F'', its action on F factors through F' in the homotopy category of functors. Hence multiplying by a scalar that kills F' kills F. Induction on the tower gives the power bound. Conversely, a

scalar annihilating the identity annihilates the regular coefficient module of every stratum, by evaluation on its standard and restriction back to the stratum. It lies in every I_O, whose intersection is $I(\Sigma_\xi)$. □

Let $\mathscr{D}_\xi$ be the set of distinct maximal subtori wD_* in Lemma 4.5. Write O_D for the unique orbit marked by $D \in \mathscr{D}_\xi$. Define a closed subset and its ambient open complement by

$$B_\xi = \bigcup_{\dim D_O < \operatorname{rank}_{\mathbb{R}} H_\xi} D_O \;\cup\; \bigcup_{\substack{D,D' \in \mathscr{D}_\xi \\ D \neq D'}} (D \cap D'), \qquad \Omega_\xi = T \setminus B_\xi.$$

All intersections here are removed in full, including disconnected intersections of subtori. The open meets every component of Σ_ξ in a dense open. For a quasi-compact open $U' \subset T$ write $\mathcal{C}|_{U'} = \mathcal{C} \otimes_{\mathrm{Mod}(R)} \mathrm{QCoh}(U')$. This agrees with ordinary localization on principal opens.

Theorem 4.7. *Restrictions to the orbits O_D give an equivalence of $\mathrm{QCoh}(\Omega_\xi) \otimes \mathrm{Mod}(P)$-linear categories*

$$\mathcal{A}_\xi^{\mathrm{per}}|_{\Omega_\xi} \;\simeq\; \prod_{D \in \mathscr{D}_\xi} \big(\mathrm{QCoh}(N_{O_D}^{-1}(D \cap \Omega_\xi)) \otimes_{\mathbb{C}} \mathrm{Mod}(P)\big).$$

The two extensions of every displayed stratum, standard and costandard, are canonically isomorphic on Ω_ξ. The equivalence restricts to the corresponding compact objects.

Proof. Over the open all standards belonging to smaller norm images vanish. Indeed the localized coefficient ring of each such stratum is zero. For a maximal component D, first work on the open $U_D = \Omega_\xi \setminus \bigcup_{D' \neq D} D'$. There is only one surviving stratum there. The finite orbit filtration, applied to an object whose restriction to that stratum vanishes, shows that the object vanishes on U_D. Here and below restriction to the Cartan open preserves the functorial filtration: all the stratum functors are continuous and R-linear, as follows from the finite link models in Lemma 2.3. The ordinary restriction to the unique surviving stratum is an equivalence since its composite with extension by zero is the identity and its kernel vanishes by the filtration. More explicitly, an arbitrary object is reconstructed by the one remaining term of that filtration, and the extension by zero functor is fully faithful by adjunction. The canonical $j_{O_D!} \to Rj_{O_D*}$ is an isomorphism: its cone restricts to zero on O_D, hence vanishes on U_D. For locally closed strata this map is the one obtained by factoring the inclusion into an open and a closed immersion.

The U_D cover Ω_ξ. Distinct such opens overlap only outside Σ_ξ, where all the categories in question are zero. Equivalently, one can use the direct sum of the standard extension functors, globally defined on Ω_ξ, as the candidate inverse to the displayed restriction. Its two composites are equivalences locally on the covering, hence globally: for a finite affine open cover, the augmented flat Čech resolution of the structure sheaf detects zero after tensoring with any linear category. The stratum coefficient description gives the asserted QCoh factors since N_{O_D} is finite. An equivalence with a continuous inverse preserves compact objects. □

For a closed point $t \in D \cap \Omega_\xi$, put $L = \ker(X^*(T) \to X^*(D))$ and $c = \operatorname{rank} L$. This lattice is saturated since D is a torus. Choosing a basis of L, the derived fibre in Theorem 4.7 takes the form

$$\mathcal{A}_\xi^{\mathrm{per}} \otimes_{\mathrm{Mod}(R)} \mathrm{Mod}(\mathbb{C}_t) \simeq \prod_{y \in N_{O_D}^{-1}(t)} \mathrm{Mod}\Big(\bigwedge\nolimits^{\bullet}(\mathbb{C}^c[1]) \otimes_{\mathbb{C}} P\Big).$$

Indeed the basis characters minus one form a regular sequence cutting out D. Splitting its character lattice sequence identifies T with a product of D and a complementary torus. The fibre algebra in those normal directions is the Koszul algebra with zero differential and c generators in degree -1. Along D the norm is finite étale, giving one copy for each point of the ordinary norm fibre at t. This describes the derived specialization of the open family. In particular, a component of positive codimension can have odd endomorphisms in its derived fibre even when the coefficient algebra of the family is concentrated in even degrees.

Lemma 4.8. *If H_ξ is not quasi-split, then $\mathcal{A}_\xi|_U = 0$.*

Proof. By Proposition 4.6 it suffices to consider a connected norm image $(T^\delta)^\circ$ in Σ_ξ. If that image meets U, no root vanishes identically on its real cocharacter space. A finite union of proper hyperplanes cannot cover this space. A fixed regular vector then lies in the interior of a Weyl chamber, which must be δ-stable. The corresponding positive system on the dual root datum gives a Galois-stable Borel of H_ξ containing the chosen real Cartan. This would make H_ξ quasi-split. Thus for every orbit at least one factor $1 - e^\alpha$ acts by zero. Localizing kills the sheaf category on that orbit. Induction through the finite orbit stratification by the open-closed triangles gives the claim. □

Proof of the regular assertion of Theorem 1.2. Let $\mathscr{I} = \{w\theta w^{-1} \mid w \in W\}$. For $\delta \in \mathscr{I}$ put $\Pi_\delta = \pi_0(T^\delta)$, $Y_\delta = T/(1-\delta)T$, $N_\delta([t]) = t\delta(t)$, and $A_\delta^U = \mathcal{O}(N_\delta^{-1}(U))$. The full parameter category has the decomposition

$$\mathcal{T}(V)|_U \simeq \prod_{\delta\in\mathscr{I}} \prod_{\nu\in\Pi_\delta^\vee} \mathrm{Mod}(A_\delta^U \otimes P). \tag{4.1}$$

We recall the inputs to this formula from Lemmas 2.40, 2.41 and 2.42. A regular boundary element determines its connected centralizing torus. The odd element normalizes this torus and preserves a chamber for some choice of flag, giving $\delta \in \mathscr{I}$. The corresponding component is the open in a twisted torus quotient with stabilizer T^δ. Translation by the pinned marking identifies its boundary map with N_δ. In stabilizer weight λ the full turn acts by the multiplier $\lambda \circ N_\delta$. If the weight is nontrivial on $(T^\delta)^\circ$, its unit fiber is smooth or empty. The cohomological operator of this fiber kills its perfect objects after inversion of u, as in [BCHN25, Example A.13]. For a component-group character the multiplier is identically one, yielding the periodic module category in (4.1). This calculation uses the single boundary circle. By Proposition 2.43, $\mathcal{S}_Q|_U$ is exactly the product of the $\nu = 1$ factors.

Consider the component for $\delta = w\theta w^{-1}$. Conjugating an odd lift by a lift of a representative of w identifies the stabilizer of a point there with the transport of T^θ. On T^δ the character of L_ε is consequently $t \mapsto \varepsilon([w^{-1}(t)])$. In fact the defect of a lift commuting with the odd lift is carried under this conjugation to the very same element of the central group D as in Lemma 4.2. Multiplying the odd representative by a point of T has no effect on this computation since the lifts of torus elements commute. The character also admits a description independent of w: it is trivial on $(T^\delta)^\circ$, and its restriction to $Z(G)^\theta$ is fixed. This restriction determines it since the fixed center surjects onto Π_δ, by conjugating Lemma 4.2. For each fixed δ the characters obtained as ε varies thus exhaust $\Pi_\delta^\vee$ without repetition. Tensoring by the line and by its dual gives inverse equivalences on the indicated factors. The essential image of the sum is all of (4.1). By Lemma 4.8 the other source summands restrict to zero. This proves the equivalence for the entire pure inner class, with the action already constructed on the global functors. □

4.4. **Real torus factors.**

Corollary 4.9. *For every real torus the full equivalence in (1.1) holds over its entire marked Cartan, with all pure torsor labels and the original boundary period.*

Proof. The root system is empty, so $f = 1$. The only involution in $\mathscr{I}$ is θ. Each pure label contributes a copy of $\mathrm{Mod}(\mathcal{O}(T/(1-\theta)T) \otimes P)$, with the boundary action through the norm. Lemma 4.1 identifies these copies with all the factors in (4.1). □

4.5. **Products of real groups.**

We use one periodic variable for all factors. Write V_i for the parameter stack of a connected quasi-split real reductive group $H_{0,i}$, and let $\mathcal{A}_i^{\mathrm{all}}$ be the direct sum of the automorphic categories of its pure inner twists before periodicization.

Proposition 4.10. *For the direct product of two such groups, external product induces equivalences*

$$\mathcal{T}(V_1) \otimes_{\mathrm{Mod}(P)} \mathcal{T}(V_2) \simeq \mathcal{T}(V_1 \times V_2), \qquad \mathcal{A}_1^{\mathrm{all}} \otimes_{\mathbb{C}} \mathcal{A}_2^{\mathrm{all}} \simeq \mathcal{A}_{1\times 2}^{\mathrm{all}}.$$

They preserve the two Cartan markings and commute with external products of finite Hecke kernels. Thus the full family comparison is preserved by finite direct products.

Proof. The odd extension for the product is the fiber product of the extensions over the common parity group. The parameter stack is $V_1 \times V_2$, and rotating the boundary rotates both factors by the same angle. For bounded below complexes M, N over $\mathbb{C}$ with circle action, the natural map

$$M^{hS^1} \otimes^L_{\mathbb{C}[u]} N^{hS^1} \longrightarrow (M \otimes_{\mathbb{C}} N)^{hS^1}$$

is an isomorphism for this diagonal action. Here $\mathbb{C}[u]$ represents the cochains of BS^1 with the positive class in degree two. First take M to be a vector space in one degree. Its circle action is trivial, and in each cohomological degree only a finite skeleton of BS^1 contributes for bounded below coefficients. Taking invariants therefore commutes with tensoring by this vector space. The map is then the scalar unit isomorphism. Finite equivariant Postnikov towers give the assertion for bounded M. For general bounded below M take the filtered colimit of its truncations from above. Their tensor products with N have a common lower bound. The same finite-skeleton argument commutes this colimit with invariants on both sides. This proves the map is an isomorphism before localizing in u.

For equivariant bounded coherent objects in the two factors, ordinary Hom between their external products is the tensor product of the ordinary Hom complexes. On affine charts choose resolutions by finite free modules in each degree, bounded above. Into bounded coherent targets only finitely many terms of the tensor resolution contribute to any fixed degree. Finite affine covers with affine intersections compute sections on the incidence atlases. Taking invariants for the product of two reductive groups is exact and takes external tensors to the tensor of the invariant spaces. This proves the ordinary Hom identity, which is natural for both rotations. Applying the preceding calculation and inverting u gives full faithfulness on external products of compact objects. Finite constructions and retracts are allowed on either side, and Ind-completion gives full faithfulness on the indicated tensor product. For essential surjectivity use Theorem 4.17, whose proof applies directly to each root datum. In the product the reduced closures of positions are precisely $X_{w_1} \times X_{w_2}$, the finite characters are pairs, and their homogeneous lines can be chosen to be external products. Their coherent closed direct images and their equivariant structures are the external products of those in the factors. These objects generate the target.

Real torsors for a product are pairs of torsors, and on each pair of real forms the enhanced flag, its strata, and the quotient tori of the strata are products of the corresponding spaces. In particular the regular local system on a product stratum is the external product of the two regular systems. Extension by zero commutes with this external product. To check maps out of such a standard use adjunction with the extraordinary restriction to its stratum. In the compact model of Lemma 2.3, local cohomology at a stratum is computed by relative cellular cochains of a finite triangulated neighborhood pair. For a product stratum the relative complexes are the tensor product: the complement of the product support is the union of the two complements, and product cells give the tensor differential. Open restriction handles locally closed strata. Thus the costalk of an external product is the external product of the costalks, including their monodromy. Maps from the regular local system then evaluate the resulting module over the tensor product of the two group algebras. This proves full faithfulness on the compact standards, and these standards generate on both sides by Lemma 2.3. The flag correspondences for convolution are products and their proper pushforwards and pullbacks respect external products. Both equivalences retain the stated markings. Extending the second one to P proves the final assertion. □

Corollary 4.11. *The full comparison holds for the product of a complex group regarded as real with any real torus. More generally it holds for a product of any number of groups satisfying the hypothesis of Corollary 6.7. An anisotropic torus factor duplicates the category of the other factors once for each of its pure torsor labels, with its Cartan acting at one.*

Proof. Apply the product formula to Theorem 1.3, Corollary 4.9 and Corollary 6.7. For an anisotropic torus the dual action is inversion, so Y_θ is a point and its norm is one. Each label contributes $\mathrm{Mod}(P)$. □

4.6. **The open stratum of opposite flags.** Put $\delta = w_0\theta$, where w_0 is the longest Weyl element. Write $D = T^\delta$, $D^0 = (T^\delta)^\circ$, and $Y_\delta = T/(1-\delta)T$. The norm $N_\delta : Y_\delta \to D^0$ is an isogeny. The stratum considered here is open in the parameter stack. Its image in the marked Cartan can nevertheless lie entirely inside the root divisors.

Lemma 4.12. *Let $j : V^{\mathrm{opp}} \hookrightarrow V$ be the locus where B' and $aB'a^{-1}$ are opposite. It is classical and isomorphic to $[Ta_0/T]$, where $a_0 = n_{w_0}\tau$ and $a_0^2 = 1$. Its marking on Y_δ is N_δ. With the original boundary circle,*

$$\mathcal{T}(V^{\mathrm{opp}}) \simeq \prod_{\chi \in X^*(\pi_0 D)} \mathrm{Mod}(\mathcal{O}(Y_\delta) \otimes P).$$

Proof. Write $\mathcal{B}$ for the flag variety. The space of ordered opposite pairs is the homogeneous space G/T. An odd element acts on the pair by conjugation and interchange. In the derived fixed-pair incidence the equation $B_2 = aB_1a^{-1}$ is a graph. Eliminating it leaves exactly the derived condition that a^2 preserves B_1, with the opposite condition imposed. This incidence is classical: as one varies the odd element in its G-torsor, the image of the given pair moves submersively on G/T. Thus the graph meets the diagonal transversely. Fixing (B, B^-) gives the quotient $[Ta_0/T]$. The pinned simple representatives obey the braid relations. Applying θ to a reduced word for w_0 therefore fixes n_{w_0}. Inverting that word and moving its factors $\alpha_i^\vee(-1)$ to one side gives the product $\prod_{\alpha>0} \alpha^\vee(-1) = (2\rho^\vee)(-1) = z$, since the inversion set is the set of positive roots. Hence $n_{w_0}^2 = z$ and $(n_{w_0}\tau)^2 = z^2 = 1$. For $h, t \in T$ we have

$$h(ta_0)h^{-1} = th\delta(h)^{-1}a_0, \qquad (ta_0)^2 = t\delta(t) = N_\delta([t]).$$

Splitting the quotient of tori $T \to Y_\delta$ identifies the stack with $Y_\delta \times BD$. On stabilizer weight λ the full turn is the function $\lambda \circ N_\delta$. If $\lambda|_{D^0} \neq 1$, its unit fiber is smooth or empty, since it is the unit fiber of a nonconstant character on a torus in characteristic zero. The calculation in Lemma 2.42 then kills this weight after inverting

u. If $\lambda|_{D^0} = 1$, the turn is identically one and the free weight line generates a copy of $\mathrm{Mod}(\mathcal{O}(Y_\delta) \otimes P)$, with the stated norm marking. □

Write $C_0 = Z(G)^\theta \cap (T^\theta)^\circ$ and denote its image in $\pi_0(D)$ by C_{opp}. This map exists because W acts trivially on the center. The objects in the image of J have trivial C_0-action by Lemmas 4.3 and 4.4.

Proposition 4.13. *Restriction to V^{opp} has a continuous fully faithful right adjoint. Its image is supported on $D^0 \subset T$. For any $\chi \in X^*(\pi_0 D)$ nontrivial on C_{opp}, the corresponding copy of $\mathrm{Mod}(\mathcal{O}(Y_\delta) \otimes P)$ is orthogonal in both directions to J. Furthermore the free generator of each character factor lifts under restriction to a compact object supported on D^0 with the same C_0-character.*

Proof. Let $i : Y \to V$ be the classical closed substack of the parameter incidence where the marking belongs to D^0. By the lemma, j factors through the open immersion $h : V^{\mathrm{opp}} \to Y$. Lift χ from D to a character λ of T. This is possible because restriction of characters to a closed diagonalizable subgroup is surjective. The homogeneous line on the flag variety with fiber of weight λ at B pulls back to a line A_λ on Y. At (a, B') its boundary turn is $\lambda(b)$, where $b \in D^0$ is the mark. This function is identically one. Its restriction to the opposite stratum is a generator in weight χ: it is a line in that weight and its factor on the affine torus Y_δ is trivializable.

We describe the open immersion by one section. In the irreducible representation of highest weight 2ρ, use the highest line for the first flag and the highest line in the dual representation for the second flag. The canonical pairing is nonzero exactly on opposite flags. Indeed at relative position w its two weights can pair only if $w(2\rho) = -2\rho$. Strict dominance forces $w = w_0$, where the pairing is nonzero. The weight belongs to the root lattice, so this representation exists for every isogeny type. The pairing gives an invariant section of the dual of the tensor product of the two lines. Pull it back along $(a, B') \mapsto (B', aB'a^{-1})$ to obtain a line L with section s on Y. The fiber weights of this line are $(-2\rho, -2\rho)$. The second mark is transported from the first by θ, which fixes 2ρ, so the boundary turn on L is $(-4\rho)(b)$. Since $\delta(2\rho) = -2\rho$, this character is trivial on D^0. In semisimple rank zero the same construction uses the constant section one. Thus h is the nonvanishing locus of s, and is affine by local trivialization of L.

For $n \geq 0$ put $M_n = i_*(A_\lambda \otimes L^{\otimes n})$. These are bounded coherent heart objects. Their full turns are the identity and multiplication by s is equivariant, by the heart calculation in Lemma 2.40. The object M_0 gives the asserted compact lift after small Tate localization. Every M_n has the same C_0-character: the section trivializes L on the opposite stratum and the central fiber character of this homogeneous bundle is constant. Each M_n is annihilated by the ideal of D^0.

Let F be a small equivariant bounded coherent object of V. There is a natural identification before taking circle invariants,

$$\operatorname*{colim}_n \mathrm{RHom}_V(F, M_n) \simeq \mathrm{RHom}_{V^{\mathrm{opp}}}(j^*F, h^*A_\lambda).$$

Indeed the direct limit of the target sheaves is $i_*h_*h^*A_\lambda$: on a trivializing open this is ordinary localization by one element. Affineness of h makes this the derived direct image as well. To pass the limit through Hom, use an almost perfect resolution of F on the incidence atlas. Into targets in the heart only finitely many finite free terms affect any fixed cohomological degree. Compute global sections there by a finite affine cover with affine intersections, then take exact rational invariants of the reductive group G. This proves the identity degree by degree. It is natural for all the boundary transports. The mapping complexes have a common lower bound depending only on the upper cohomological bound for F. In the cell filtration of BS^1 only finitely many cells can affect any fixed total degree. Thus the direct limit also commutes with circle invariants. Localizing at u preserves the identification.

The functor j^* preserves compact objects: it restricts bounded coherent invariant objects before localization. Its right adjoint $j_*^{\mathcal{T}}$ exists on the Ind-categories and preserves colimits, by compact generation. Testing against the images of the small equivariant objects now identifies $j_*^{\mathcal{T}}(h^*A_\lambda)$ with the direct limit of the M_n in $\mathcal{T}(V)$. After restriction, every transition is multiplication by the invertible s. The counit is therefore an isomorphism on the free generator of every character factor. Since both functors preserve colimits, it is an isomorphism throughout $\mathcal{T}(V^{\mathrm{opp}})$. This proves full faithfulness. The same description proves the support assertion and keeps the original R- and P-actions. Distinct characters of C_0 give both directions of orthogonality by exact central weight projection. □

If $w_0\theta = -1$ on $X^*(T)$, then $D = T[2]$, Y_δ is a point, and all the objects in the proposition are supported at the unit parameter. Each nontrivial restriction to C_{opp} gives a fully faithful global copy of $\mathrm{Mod}(P)$ orthogonal to the quasi-split summands. The compact lift and the value of the right adjoint play different roles: the first extends a coherent line and the second identifies its full mapping complexes with those on the open stratum.

4.7. **The open real orbits in the entire pure inner class.** Let $\mathcal{B}_H$ denote the variety of complex Borels of H. For each $\xi \in \Theta$, write $\mathcal{A}^{\mathrm{open}}_{\xi,!}$ for the full localizing subcategory of $\mathcal{A}^{\mathrm{per}}_{\xi}$ generated by the regular standards of its open real flag orbits. The standard functor here is extension by zero from the enhanced inverse image of an open orbit. Retaining the actual orbits within each pure inner form is important, since a single form can have several.

Lemma 4.14. *Fix an open $H_0(\mathbb{R})$-orbit and a Borel B_f in it. With σ denoting its real structure on H, the torus $J_f = B_f \cap \sigma B_f$ is defined over $\mathbb{R}$. The marking by B_f gives Galois involution $\delta = w_0\theta$ on its cocharacter lattice. There are bijections*

$$\coprod_{\xi\in\Theta} \{\textit{open } H_\xi(\mathbb{R})\textit{-orbits on } \mathcal{B}_H\} \;\simeq\; H^1(\mathbb{R}, J_f) \;\simeq\; X^*(\pi_0(T^\delta)).$$

The map from the middle set to Θ is induced by $J_f \hookrightarrow H_0$. In particular every pure label occurs, with multiplicity the size of its fibre.

Proof. Give $\mathcal{B}_H \times \mathcal{B}_H$ the real structure $(B_1, B_2) \mapsto (\sigma B_2, \sigma B_1)$. Its real points have the form $(B, \sigma B)$. For a cocycle z replace σ by $\sigma_z = \mathrm{Ad}(z)\sigma$. At such a real pair the stabilizer of the diagonal complex group is $B \cap \sigma_z B$, defined over $\mathbb{R}$ for σ_z; its real points are exactly the stabilizer of B in $H_z(\mathbb{R})$. Smoothness in characteristic zero gives equality between the real dimension of the real orbit and the complex dimension of the diagonal orbit of the pair. The latter dimension is maximal precisely when the pair is opposite. Thus real open orbits are exactly the real orbits on opposite pairs. They occur in every twist. Indeed the real variety of all pairs is smooth and geometrically irreducible, and it has a real point $(B, \sigma_z B)$ for any B. Its real points are Zariski dense: real analytic coordinates at a smooth real point contain an open real ball of the full complex dimension, on which a regular function can vanish only if its holomorphic germ vanishes. The nonempty Zariski open of opposite pairs therefore has real points.

The complex group acts transitively on opposite pairs. The stabilizer of the chosen neutral pair $p_f = (B_f, \sigma B_f)$ is J_f. On its marked root datum, Galois takes the positive system for B_f to its negative. Its based action is θ, so the actual involution is $w_0\theta$; this also records its given action on the central torus.

For a pair p fixed under the twisted real structure, choose $g \in H(\mathbb{C})$ with $p = gp_f$. Then $c = g^{-1}z\sigma(g)$ fixes p_f and satisfies $c\sigma(c) = 1$. Changing g to gj with $j \in J_f$ replaces c by $j^{-1}c\sigma(j)$. Translation of p by an element of $H_z(\mathbb{R})$ preserves the class. Conversely a J_f-cocycle makes p_f itself a real pair in its pure twist. These operations are inverse on classes, and the underlying pure label of c is its image in $H^1(\mathbb{R}, H_0)$. The stabilizer used in this argument is the torus itself, so the construction retains every orbit in each fibre.

Finally, for $\Lambda = X_*(J_{f,\mathbb{C}}) = X^*(T)$ the exponential sequence has kernel $2\pi i\Lambda$, with involution $-\delta$. The calculation in Lemma 4.1 gives

$$H^1(\mathbb{R}, J_f) = \ker(1+\delta : \Lambda \to \Lambda)/(1-\delta)\Lambda = (\Lambda/(1-\delta)\Lambda)_{\mathrm{tors}} = X^*(\pi_0(T^\delta)).$$

This proves the second bijection with the indicated marking. □

Proof of Theorem 1.6. Every open orbit has marked involution δ. Indeed in the cocycle description of Lemma 4.14 the twist is by an element of the abelian torus J_f, so it has no effect on conjugation on that torus. Lemmas 2.1 and 2.2 identify the enhanced stratum category with $\mathrm{Mod}(\mathcal{O}(Y_\delta) \otimes P)$, marked by $e^\lambda \mapsto e^{(1+\delta)\lambda}$. Extension by zero along an open inclusion is fully faithful. Extensions from two distinct open orbits are orthogonal in both directions: ordinary restriction to either orbit annihilates extension by zero from the other, and ordinary and extraordinary restriction agree for an open inclusion. Different pure labels are already direct summands. Taking the finite product over the orbit labels thus computes the entire displayed source, with one factor for each $\chi \in X^*(\pi_0(T^\delta))$.

Lemma 4.12 gives the identical decomposition of the opposite spectral category, with the same norm on each factor. Identifying the two free coefficient modules defines the equivalence and preserves the original period. The global fully faithful functor follows from Proposition 4.13. If $\delta = -1$, then Y_δ is a point, $T^\delta = T[2]$, and the norm is one. □

The choice of a neutral open orbit fixes the origin of the cohomological labelling; a change of origin relabels the character factors. The construction uses open-orbit extensions on the real side and the right adjoint to open restriction on the spectral side. Comparing it with the finite Hecke functor from the closed real orbit involves the attachment of the other flag strata.

4.8. **Generators indexed by all flag positions.** For a point (a, F) of the parameter stack, consider the relative position of (F, aFa^{-1}). Write S_w for the reduced locally closed substack of position w, and X_w for its reduced closure in V. We use the original boundary circle on these stacks: its closing automorphism is a^2. Equivalently one uses the groupoid of an odd transport interchanging the two flags and restricts the even boundary transport to the indicated invariant locus. All inclusions and the projections of transport groupoids below respect this circle.

Lemma 4.15. *The positions that occur are precisely the elements with $w\theta(w) = 1$. Set $\delta = w\theta$, $J = B \cap {}^wB = T \ltimes U_w$, and $Y_\delta = T/(1-\delta)T$. There is a lift $a_w \in N_G(T)\tau$ of $w\theta$ with $a_w^2 = 1$. Then*

$$S_w = [Ja_w/J], \qquad p_w : S_w \longrightarrow [Ta_w/T].$$

Here p_w is induced by the Levi quotient of J, and the marking is $N_\delta([t]) = t\delta(t)$ in its torus coordinate.

Proof. Applying a takes the position to $\theta(w)$ and reverses the pair, so $\theta(w) = w^{-1}$. For the converse let n_w be the pinned Tits lift and put $r = (2\rho^\vee)(i)$. Multiplying a reduced word by its reversal and using $n_s^2 = \alpha_s^\vee(-1)$ gives

$$n_w n_{w^{-1}} = \prod_{\alpha>0,\, w^{-1}\alpha<0} \alpha^\vee(-1) = (\rho^\vee - w\rho^\vee)(-1) = r/w(r).$$

The pinning gives $\theta(n_w) = n_{w^{-1}}$. Since $r^2 = \tau^2$ has order dividing two and $r/w(r)$ has order dividing two, $(n_w\tau)^2 = r\,w(r)$. Thus $a_w = r^{-1}n_w\tau$ has square one. It interchanges B and ${}^wB$. All the odd elements interchanging this fixed pair form its stabilizer coset Ja_w. The variety of such ordered pairs is homogeneous under G, with stabilizer J. This proves the quotient description, including the reduced structure since the coset is smooth. Conjugation by a_w preserves both T and U_w. Projection of the square to the Cartan therefore gives the asserted norm. □

Lemma 4.16. *Pullback by p_w induces an equivalence on the small coherent Tate categories and their Ind-completions. In particular,*

$$\mathcal{T}(S_w) \simeq \prod_{\chi \in X^*(\pi_0(T^\delta))} \mathrm{Mod}\big(\mathcal{O}(Y_\delta) \otimes P\big).$$

Proof. Choose an integral strictly dominant cocharacter v of T and put $\nu = v + \delta(v)$. A root of U_w is positive for both flags. The involution δ permutes these roots, so each pairs strictly positively with ν. Also $\delta(\nu) = \nu$, whence ν acts trivially on the torus atlas Ta_w. On that atlas the category splits by the integer weight of ν on stabilizers. A coherent object uses finitely many such weights. For each integer k consider the image by p_w^* of the weight-k category before taking circle invariants.

Here is the calculation of maps between the images. It suffices to take character lines lifted from T: the torus quotient is $Y_\delta \times B(T^\delta)$ after a choice of splitting, each stabilizer character extends to T, and the coordinate ring of Y_δ is regular. For source and target lines of weights k, l, respectively, maps are computed by the rational T-invariants of the Lie algebra cochains of $\mathfrak{u}_w$ with coefficients in $\mathcal{O}(Ja_w)$ tensored with the difference of the two characters. This is the unipotent cohomology calculation in characteristic zero [Hoc61]. In the coordinates $U_w \times T$, the nonconstant monomials in the radical coordinates have strictly negative ν-weights: use the equivariant exponential map on U_w. Every factor of $\mathfrak{u}_w^*$ has strictly negative weight as well. Thus maps vanish if $l < k$. If $l = k$, the weight-zero subcomplex has only the functions of t, in cochain degree zero. Its remaining torus invariants compute precisely the maps on $[Ta_w/T]$. These identifications are induced by pullback and are compatible with composition. They extend to finite constructions and retracts, giving full faithfulness within each weight.

These images generate the bounded coherent category before taking invariants. For a coherent equivariant module on the smooth affine atlas Ja_w, choose finitely many generators over its coordinate ring. Their rational J-spans are contained in a finite-dimensional representation. This gives an equivariant surjection from a free bundle with representation coefficients. Iterate on its coherent kernels. Ext between coherent heart objects on the quotient vanishes above $\dim J + \dim U_w$: ordinary sheaf Ext on the smooth affine atlas vanishes above $\dim J$, and derived invariants are computed by at most $\dim U_w$ unipotent cochain degrees followed by exact torus invariants. For a truncation of the free resolution of length exceeding this bound, the connecting extension from the original module to its shifted last syzygy therefore vanishes. The module is a retract of the truncated complex of free bundles. The group J is connected solvable, so each representation coefficient has a finite filtration by characters. Postnikov towers prove the generation assertion for bounded complexes.

It follows that every object lies in the thick subcategory generated by a finite interval of weights. Within this interval the fully faithful weight categories give a semiorthogonal decomposition in order of increasing weight. Indeed for two successive parts the vanishing of maps from the higher part to the lower in every degree makes the triangle with higher subobject and lower quotient functorial. Such triangles are closed under cones and retracts and contain both parts. Induction gives the successive exact projections. Rotation preserves each part, since p_w preserves the boundary action. Uniqueness of the projections makes the triangles equivariant, so every small invariant object has the same finite filtration by invariant weight pieces. Full faithfulness on each piece persists after taking homotopy invariants of its mapping complexes.

On the torus the calculation of Lemma 2.42 applies over the full Y_δ. A character nontrivial on the connected norm image has a smooth unit fibre for its multiplier and vanishes after inverting u. Every surviving character is trivial on ν, whose image is connected and fixed by δ. Thus all surviving pieces belong to the single fully faithful weight-zero category. This proves the equivalence, first for small objects after idempotent completion and then for their Ind-completions. □

For each surviving χ choose a character $\lambda \in X^*(T)$ extending it from T^δ. On X_w pull back the homogeneous line of weight λ from the first flag and denote it by $A_{w,\chi}$. Its full turn is the value of λ on the mark. The mark takes values in $(T^\delta)^\circ$ on S_w, hence on its reduced closure; this value is one. The line consequently has an equivariant structure by the coherent heart calculation of Lemma 2.40. Its restriction to S_w generates the indicated character factor, since any line on Y_δ is trivializable. Let $C_{w,\chi}$ be its coherent direct image along $X_w \to V$, with this circle action, followed by periodic localization.

Theorem 4.17. *The finitely many $C_{w,\chi}$ compactly generate the entire spectral category $\mathcal{T}(V)$.*

Proof. We first realize the open $S_w \subset X_w$ by a single section of a line bundle. For the dominant regular weight $\gamma = 2\rho$, the highest line in the representation of highest weight γ defines a line on each flag. With first flag B, the lines over the Schubert closure of the cell w lie in the span of the B-translates of a vector of weight $w\gamma$. All other weights in that span are obtained by adding nonzero sums of positive roots. Projection to the extremal weight is therefore a B-equivariant map to the line of character $w\gamma$. It is nonzero on the cell w. On a cell $v < w$ its value is zero: along a Bruhat chain the weights $v\gamma - w\gamma$ are nonzero sums of positive roots, by the reflection formula and regular dominance. Thus inducing the map from the first flag gives an invariant section on the orbit closure of pairs, of weights $(w\gamma, -\gamma)$. Pull it back to X_w, with line L and section s. The second flag mark is obtained by θ from the first one, so the full-turn character is $w\gamma - \theta\gamma$. Here $\theta\gamma = \gamma$ and $w\gamma = \delta\gamma$, giving value one on $(T^\delta)^\circ$. Hence L and s are circle-equivariant on X_w. The nonvanishing locus is exactly S_w. In semisimple rank zero take the constant section one.

For bounded equivariant coherent F, F' on X_w, write $j : S_w \to X_w$. Then the maps after restriction are computed by the telescope

$$\operatorname*{colim}_{n\geq 0} \mathrm{RHom}_{X_w}(F, F' \otimes L^n)^{hS^1} \simeq \mathrm{RHom}_{S_w}(j^*F, j^*F')^{hS^1}.$$

On a trivialization of L this is localization by one function, and j is affine. On the incidence atlas an almost perfect resolution of F uses only finitely many terms in any fixed degree into these uniformly bounded targets. A finite affine cover with affine intersections and exact rational G-invariants give the identity before circle invariants. The mapping complexes have a common lower bound, so finite skeleta of BS^1 suffice in any fixed degree before inverting u. Localization in u preserves the resulting identity. In particular the Verdier quotient by the cones of multiplication by s, with all twists, embeds fully faithfully in the small open category after idempotent completion. To see the quotient formula, formally invert the commuting maps $F' \to F' \otimes L$; powers of s give fractions with the displayed colimit as mapping complex. Its kernel is generated thickly by their cones.

The open category is generated by the restrictions of the $A_{w,\chi}$, by Lemma 4.16. Hence on X_w these lines together with objects supported on the boundary generate. More explicitly every cone of multiplication by s on a small invariant representative has bounded cohomology supported on $X_w \setminus S_w$. Twists, retracts and cones after localization give the same thick generation statement. Each boundary cohomology sheaf has a finite filtration by modules on the reduced lower closures: take powers of the ideal of the reduced boundary, then successively multiply by the ideals of its component closures. If a lower closure meets X_w only in a closed subset, the module so obtained is pushed from that subset into the lower closure. The filtration is equivariant. On the heart the closing automorphism of the original sheaf is the identity and restricts to the identity on each subquotient. Bounded Postnikov towers lift the resulting generation statement to equivariant coherent complexes. Pushing into V and inducting in Bruhat order now treats every X_w. Finally every equivariant coherent heart object of V is a module on the classical truncation. The same ideal filtration for the union of all the X_w, including a power

of its nilradical, expresses it in terms of their direct images. Taking bounded Postnikov towers, inverting u and forming the Ind-category proves the theorem. Compactness follows because each specified generator is the image of a small coherent invariant object. □

Proposition 4.18. *For every $w\theta(w) = 1$, the enhanced real flag strata of position w in the full pure inner class have, after a choice of origin, the same index set $X^*(\pi_0(T^\delta))$ and the same marked coefficient categories as in Lemma 4.16.*

Proof. Write the real conjugation in a quasi-split pinning of H as $\theta_H c$, where c is split conjugation. Put $r_H = (2\rho_H^\vee)(i)$, and take the pinned representatives in H. Then $z_w = r_H^{-1} n_w$ is a Galois cocycle: $c(r_H) = r_H^{-1}$, and the product of z_w with its conjugate is $r_H^{-1} w(r_H) n_w n_{w^{-1}} = 1$ by the inversion formula in Lemma 4.15. In this pure twist the pair $(B_H, {}^w B_H)$ is a real point for the action exchanging the two flags. Its stabilizer J_H is the intersection of the two Borels with this real structure. Projection onto its stable Cartan with involution δ induces a bijection on first Galois cohomology: there is a stable torus section, and the unipotent radical of every torus twist has trivial first cohomology by its filtration by real vector groups.

Using this pair as origin, an orbit in any pure twist corresponds to a cocycle of J_H. Explicitly, if the twisting cocycle is z and the pair is translated from the origin by g, the relative cocycle is $g^{-1} z(\theta_H c)(g) z_w^{-1}$ for the real structure at the origin. Changing representatives gives the usual coboundaries, and twisting by such a cocycle reverses the construction. The exponential-sequence calculation of Lemma 4.1 gives precisely the stated character group. The compact enhanced orbit calculation in Lemmas 2.1 and 2.2 gives the norm N_δ and the module category of its finite cover, including all components of the real stabilizer. □

4.9. **The successive quotients for the flag filtration.** Order the twisted involutions by a linear extension of Bruhat order, starting with the identity. Filter the spectral category by the localizing subcategories generated by the $C_{w,\chi}$ up to each position. On the real side filter by the closed union of the corresponding positions, in each pure twist. We calculate the successive quotients with their original Cartan actions. This requires an ambient calculation: pushing coherent sheaves from a closed stratum can introduce extension classes from its normal directions.

Lemma 4.19. *Let M be an invariant open set of ordered flag pairs in which the orbit O of position w is closed. Assume that M is preserved by the pinned involution composed with interchange. Let V' be the open part of the derived parameter stack over M. The direct image $i_* : \mathcal{T}(S_w) \to \mathcal{T}(V')$ is fully faithful. Its image is generated by the small equivariant coherent objects supported on S_w after periodic localization.*

Proof. We first work with ordinary coherent complexes. On the incidence atlas for V', let Z be the derived equality of x and $a \cdot x$ in M, for $x \in M$ and a in the odd component. Here $a \cdot x$ includes the interchange of the two flags. Eliminating the second flag recovers the original derived incidence, as in Lemma 4.12. Form $Z' = Z \times_M^R O$ using the coordinate x, and write S for the intrinsic equality on O. There are equivariant maps $S \xrightarrow{b} Z' \xrightarrow{c} Z$. The equality on O is transverse because the action of G on O is transitive. Thus S is smooth, it is the classical truncation of Z', and $[S/G] = S_w$.

Let N be the pullback of the normal bundle of O in M. At the pair $(B, {}^w B)$ it is the J-representation

$$N = \mathfrak{g}/(\mathfrak{b} + {}^w\mathfrak{b}).$$

Indeed the orbit tangent is the diagonal image of $\mathfrak{g}$ in the sum of the two flag tangents; subtracting representatives gives the quotient. Its torus weights are the negatives of the roots of U_w. In particular all its ν-weights are strictly negative for the cocharacter chosen in Lemma 4.16.

For vector bundles A, B' on S we compute the normal terms in maps from $c_* b_* A$ to $c_* b_* B'$. By derived base change for the definition of Z', the underlying complex of $Lc^* c_* b_* A$ is obtained by tensoring $i_* A$ over M with $\mathcal{O}_O$. A Koszul resolution for the regular embedding $O \subset M$ has zero differentials after this tensor product. Consequently its cohomology in degree $-p$, as a module on $t_0 Z' = S$, is $A \otimes \bigwedge^p N^\vee$. These identifications are intrinsic by the conormal description of Tor. The resulting finite Postnikov filtration on Z' is G-equivariant and its degree-zero quotient is the adjunction map to $b_* A$.

The sheaf Ext on Z' from $b_* D$ to $b_* B'$, for a vector bundle D on S, is zero in odd degrees. In degree $2h$ it is $D^\vee \otimes B' \otimes \mathrm{Sym}^h N$. To see this, choose local smooth coordinates adapted to O near a point of the diagonal of $M \times M$. On the atlas with $x \in O$ both arguments of the equality land in O. The differences of the normal coordinates are identically zero. The remaining differences impose the transverse equality on O. The local algebra is therefore the exterior algebra over $\mathcal{O}_S$ on the normal covectors in degree -1, with zero differential. Resolving its augmentation gives the symmetric algebra on the dual generators in cohomological degree two. On

overlaps the degree-two identification is dual to the intrinsic conormal generator in degree -1; multiplication gives the asserted identification in each degree. This calculation requires only the sheaf Ext groups.

Take for A, B' the lines pulled back from the torus whose stabilizer characters are trivial on $(T^{\delta})^{\circ}$. Both lines have ν-weight zero at the torus. By adjunction for c, the two calculations above express the contributions to the ambient Hom by the equivariant cohomology on S of

$$A^{\vee} \otimes B' \otimes \bigwedge^{p} N \otimes \operatorname{Sym}^{h} N, \qquad p, h \geq 0.$$

For $p + h > 0$ every fiber weight of the last two factors is strictly negative. On the atlas Ja_w the functions and the unipotent cochains have nonpositive weights by Lemma 4.16. Exact torus invariants thus annihilate the entire complex for $p + h > 0$. The spectral sequences used here converge degree by degree: the first filtration is finite, the second starts in degree zero, and the quotient atlas has bounded cohomological dimension. Only $p = h = 0$ remains. The induced isomorphism from the intrinsic Hom to the ambient Hom is the map of direct image, as is seen from the two augmentations. For boundary rotation use the composite i itself. It is the map of derived twisted loop stacks induced by the inclusion of the orbit in the open space of pairs. Rotation through the full boundary commutes with this map by functoriality of precomposition on the orientation double cover. The two-step factorization above is used only for ordinary coherent Hom. Hence the Hom isomorphism passes to circle invariants and inversion of u. The specified lines compactly generate by Lemma 4.16, and direct image sends them to compact objects. This proves full faithfulness on the Ind-categories.

For the image statement, apply Postnikov truncations to an invariant bounded coherent complex with support on S. Each cohomology module is a module over $t_0 Z$ killed by a power of the ideal of S. The ideal filtration and the circle structure on the heart exhibit it as iterated extensions of direct images from S. After localization this gives the required image. □

Theorem 4.20. *On both sides of* (1.1)*, the successive quotient at position w is equivalent, over the marked Cartan, to*

$$\prod_{\chi \in X^{*}(\pi_0(T^{w\theta}))} \operatorname{Mod}\big(\mathcal{O}(Y_{w\theta}) \otimes P\big).$$

The marking on each factor is $N_{w\theta}$. On the spectral side the image of $C_{w,\chi}$ gives a free generator of its factor.

Proof. Remove from the space of flag pairs the orbit closures of all earlier positions and the boundary of the orbit O of position w. The resulting open M contains O as a closed orbit. Interchange composed with the pinned action preserves the removed sets, because it fixes each of their indexing twisted involutions. Write $k : V' \to V$ for the resulting open immersion. Let $\mathcal{D}$ be the spectral filtration term for w and $\mathcal{D}_{-}$ the preceding term.

Write $l : X_w \to V$. For a chosen line $A = A_{w,\chi}$, use the section s of the line L constructed in Theorem 4.17 and form the telescope of $M_n = l_{*}(A \otimes L^{n})$ under multiplication by s. Each transition has cone supported on the boundary. The ideal filtration and induction in that theorem place this cone in $\mathcal{D}_{-}$, so each M_n and its telescope belong to $\mathcal{D}$. For a small equivariant coherent F in the ambient category we have

$$\operatorname*{colim}_{n} \operatorname{RHom}_{V}(F, M_n)^{hS^1} \simeq \operatorname{RHom}_{V'}(k^{*}F, i_{*}(A|_{S_w}))^{hS^1}.$$

Indeed ordinary localization of the target on X_w gives the direct image from its affine open S_w. The composite $S_w \to X_w \to V$ also factors as $S_w \to V' \to V$, giving the formula by adjunction. The finite-degree resolution and common lower bound of Theorem 4.17 apply on the ambient incidence atlas since every M_n is in the heart. Thus filtered colimits commute with the indicated Hom and with circle invariants, and the formula persists after inverting u.

Earlier generators restrict to zero under k^{*}, so the telescope is right orthogonal to $\mathcal{D}_{-}$. The map to it from M_0 has cone in $\mathcal{D}_{-}$. It therefore represents the localization of $C_{w,\chi}$ in $\mathcal{D}/\mathcal{D}_{-}$. The formula computes maps between every pair of localized generators as maps between their direct images in V'. By Lemma 4.19, these are exactly their intrinsic maps on S_w, with composition induced by restriction. The quotient is compactly generated by these images since the killed subcategory is generated by compact objects. Their restrictions generate $\mathcal{T}(S_w)$. Lemma 4.16 proves the spectral assertion.

For a real form the flag-pair map sends a flag to the pair formed with its real conjugate, hence inverse images of the closed position unions are closed. Within a position its real orbits are open and closed: the real points of the smooth orbit map at a real pair are submersive onto the real locus, and this holds at each orbit. The same partition applies after passing to the enhanced flag. Extension by zero and restriction for this finite orbit filtration identify the successive quotient with the product of its orbit categories, as in Proposition 4.18. More

explicitly on a locally closed orbit the standard extension of its free monodromic system has nilpotent singular support by the orbit construction of Lemma 2.2; the localization triangles recover all objects of the filtration from these standards. Periodic scalar extension preserves these quotients. Applying Proposition 4.18 gives the specified indices and marking. □

5. Compact Levis and the closed real orbit

To an individual pure inner form we now attach a compact object on the parameter stack. When the derived real group is compact, its entire flag variety is a single closed orbit. In general the same construction applies inside the Levi of a minimal real parabolic. The stable-parabolic correspondence carries its object to the ambient group and identifies the maps between closed generators.

5.1. **Orthogonal factors from compact real groups.** For the formal comparison in this subsection assume H_0 is semisimple of adjoint type. Let $\xi \in H^1(\mathbb{R}, H_0)$ be a pure inner label for which $H_\xi(\mathbb{R})$ is compact. We retain the geometric Langlands extension and the original boundary circle. Let V be its parameter stack, let $b : V \to T$ be the boundary map, and let $\mathfrak{m} \subset R = \mathcal{O}(T)$ be the ideal of $1 \in T$. The proof below uses the ABV parameter attached to this actual pure inner label. These are the root-data conventions in [BCHN25, Section 4.2].

Let $\widehat{V}_1$ denote the formal completion along $b^{-1}(1)$. We use the ordinary supported category

$$\mathrm{QC}^!(\widehat{V}_1) = \mathrm{QC}^!_{V_1}(V), \qquad \mathrm{Coh}(\widehat{V}_1) = \mathrm{Coh}_{V_1}(V).$$

These equalities are the ordinary convention of [BCHN25, Appendix A.2.1]. In particular the displayed left side of [BCHN25, Theorem 4.4] is $\mathrm{QC}^!(\mathbb{L}_1^\eta)^{\mathrm{Tate}}$, without the hat used there for the larger continuous coherent category. Small invariants and Tate localization below are formed from $\mathrm{Coh}_{V_1}(V)$. Denote the resulting Ind-category by $\mathcal{T}_1$.

Lemma 5.1. *For a compact pure inner form one has an R-linear equivalence*

$$\mathrm{Shv}_{\mathcal{N}}(H_\xi(\mathbb{R})\backslash H/N_H) \simeq \mathrm{Mod}(\mathbb{C}),$$

where R acts on the right by evaluation at one. Its periodization is $\mathrm{Mod}(P)$*, for $P = \mathbb{C}[u, u^{-1}]$.*

Proof. Iwasawa decomposition for the compact real form gives $H = H_\xi(\mathbb{R})A_{\mathrm{pos}}N_H$. The intersection of the compact group with any conjugate of N_H is trivial. Consequently the quotient is the contractible positive real Cartan A_{pos}. Residual Cartan local constancy, which is the nilpotent singular-support condition on the enhanced flag variety, is ordinary local constancy on this quotient. Its category is $\mathrm{Mod}(\mathbb{C})$. The compact Cartan acts through a contractible quotient, so every monodromy operator acts as the identity. This gives the stated R-marking. □

Lemma 5.2. *In the unit ABV parameter space, the label of the trivial representation of $H_\xi(\mathbb{R})$ determines a character local system ε_ξ on an open orbit with finite stabilizer. Its intersection complex L_ξ is clean,*

$$j_!\varepsilon_\xi[\dim O] \simeq L_\xi \simeq j_*\varepsilon_\xi[\dim O], \qquad \mathrm{RHom}(L_\xi, L_\xi) = \mathbb{C}.$$

It is orthogonal in both directions to every other irreducible equivariant intersection complex. The corresponding statement holds simultaneously for all the compact pure labels.

Proof. This is the compact-form calculation in [Soe01, Section 3.5]. Its pure-inner scope follows from [Soe01, Section 1.3]: the choice $F = 1$ in that paper gives exactly $H^1(\mathbb{R}, H_0)$. The calculation uses the ABV character-duality formula to obtain orthogonality to the other parameters and identifies the compact parameter with an open orbit for the symmetric subgroup coming from a split dual real form. Its stabilizer is finite. There is no principal-block assumption in this calculation.

Here are the consequences for the sheaf itself. The restrictions i^*L_ξ and $i^!L_\xi$ to the complementary closed union of orbits are bounded constructible objects. Its irreducible equivariant intersection complexes generate that bounded category under finite cones and shifts. Orthogonality, in the two adjunctions for i, therefore forces both restrictions to vanish. Recollement gives the two displayed clean-extension isomorphisms. The endomorphism calculation can now be made on the open orbit: for its finite stabilizer D, characteristic zero gives $\mathrm{RHom}_{BD}(\varepsilon_\xi, \varepsilon_\xi) = \mathbb{C}$. Verdier duality gives the opposite direction of orthogonality as well; the dual parameter still belongs to the same compact pure label.

The finite-stabilizer local system is compact on the open orbit, and $j_!$ preserves compactness because its right adjoint j^* preserves colimits. On the finite stabilizer orbit the orthogonal projection to the chosen local system is the character projector. Its two extensions agree by cleanness. The direct images along the open commute

with direct sums on this character: their stalks are computed by the finite links of a triangulation adapted to the finite orbit stratification. Hence the two extensions of the projector give the same factor for arbitrary coefficients. After periodization it is $\mathrm{Mod}(P)$. □

Lemma 5.3. *The compact parameter in Lemma 5.2 belongs to the unit ABV space for the geometric extension.*

Proof. In [Soe01, Section 1.3] the dual extension is the pinned semidirect product. Write its odd lift as τ_0, so that $\tau_0^2 = 1$. The parameter at infinitesimal character λ has equation $y^2 = \exp(2\pi i\lambda)$. The trivial compact representation has Harish–Chandra parameter ρ_H. Under root duality,

$$\exp(2\pi i\rho_H) = (2\rho_G^\vee)(-1) = z.$$

Let τ be the pinned odd lift of the geometric extension, with $\tau^2 = z$. Both lifts act on G by the same pinned involution θ. The map $g\tau_0 \mapsto g\tau$ is an isomorphism of odd G-spaces for twisted conjugation and satisfies

$$(g\tau)^2 = (g\tau_0)^2 z.$$

Since $z^2 = 1$, it identifies the arithmetic ABV equation $y^2 = z$ with the geometric equation $a^2 = 1$, preserving flags and stabilizers. This is the classical-base identification used in the unit formal comparison. The original boundary circle remains the geometric one; changing the odd square in the arithmetic presentation requires the central twist explained in [BCHN25, Appendix A.2.3].

The open split-type parameter is represented arithmetically by $n_{w_0}\tau_0$, whose square is z, and geometrically by $n_{w_0}\tau$, whose square is one. Its action on T is inversion, so the open stabilizer is $T[2]$. Thus the square-one display in the compact calculation is interpreted after this geometric normalization when $z \neq 1$. □

Lemma 5.4. *The inclusion of supported coherent objects induces a fully faithful functor $\mathcal{T}_1 \to \mathcal{T}(V)$. Its essential image is the subcategory of $\mathfrak{m}$-torsion objects for the central R-action. Compact objects of $\mathcal{T}_1$ are compact in $\mathcal{T}(V)$, and each is killed by a power of every element of $\mathfrak{m}$ in the homotopy category.*

Proof. Let $t_1, \dots, t_r$ be torus coordinates, let $f_i = t_i - 1$, and let $\mathsf{K} = \bigotimes_i \mathrm{cofib}(f_i : R \to R)$. Supported coherent objects form a full stable subcategory of $\mathrm{Coh}(V)$. The circle action preserves it, and its mapping complexes agree with the global mapping complexes, including the action of u. Full faithfulness therefore persists after taking small invariants, localizing in u, idempotent completion, and Ind-completion.

The localizing subcategory of $\mathfrak{m}$-torsion objects in any compactly generated R-linear category is generated by $\mathsf{K} \otimes_R C$, where C ranges over compact generators. Indeed its colocalization is the tensor product of the fibers of $R \to R[f_i^{-1}]$. Each fiber is a filtered colimit of shifted Koszul complexes for powers of f_i; these complexes are finite extensions of the Koszul complex for f_i. Taking the product over all i gives the asserted generators.

Each $\mathsf{K} \otimes_R C$ has a small representative which is bounded coherent and supported over one. Conversely a bounded coherent supported object has a finite Postnikov filtration whose heart layers are annihilated by powers of the f_i. The geometric circle acts exactly on this filtration. Its heart layers have discrete mapping spaces, so the same annihilation holds equivariantly. The finite filtration gives nilpotence on the whole object. Its localizations at the f_i vanish, so it belongs to the torsion subcategory. This proves equality of essential images.

The displayed torsion generators are compact globally. Every compact object of their localizing subcategory belongs to their finite thick hull. It is therefore globally compact. On the Koszul generators each f_i is nullhomotopic, and finite cones and retracts preserve annihilation by some power, proving the last assertion. □

Proposition 5.5. *There is a compact object $M_\xi \in \mathcal{T}(V)$, supported at the unit boundary parameter, which generates an orthogonal direct factor*

$$\mathcal{T}(V) \simeq \mathrm{Mod}(P) \times \mathcal{T}(V)'_\xi.$$

The free generator of the first factor maps to M_ξ, and the coefficient map on that factor is evaluation $R \to \mathbb{C}$ at one. Distinct compact pure labels give mutually orthogonal factors.

Proof. Apply the unit comparison of [BCHN25, Theorem 4.4], in the ordinary supported convention recorded above, to L_ξ of Lemma 5.2. Lemma 5.3 identifies its parameter with the geometric unit. The result is a compact object $E_\xi \in \mathcal{T}_1$ generating an orthogonal factor $\mathrm{Mod}(P)$. It has endomorphism algebra P. Let M_ξ be its fully faithful global image under Lemma 5.4.

Each f_i acts nilpotently on M_ξ, while $\mathrm{End}^0(M_\xi) = \mathbb{C}$. Hence f_i acts by zero. This identifies the action on the degree-zero endomorphisms. Equivalently, one can identify the marking on the factor by open restriction. The clean ABV object restricts to its character on the finite-stabilizer orbit. The unit comparison commutes with this restriction: the unipotent loop description, restriction to formal loops and the local Koszul equivalence for filtered differential operators restrict to the same open orbit, as does Riemann–Hilbert. Here $w_0\theta = -1$ on the

full torus. The orbit is $BT[2]$. Its unipotent loop stack is itself, since a finite group in characteristic zero has no nonidentity unipotent elements and its Lie algebra is zero. On this opposite stratum the boundary mark is identically one. Restriction is fully faithful on the factor, since it takes the generator to the free character line and induces the identity on its scalar algebra P. Thus the factor carries the original action through evaluation at one. In particular the argument does not require a comparison between a formal parameter and its chosen logarithm.

We verify the orthogonal splitting globally. Let $\Gamma_{\mathfrak{m}}$ be the torsion colocalization in the R-linear category $\mathcal{T}(V)$. The cone of $\Gamma_{\mathfrak{m}}X \to X$ is the finite Čech totalization of the localizations on $D(f_i)$. Each term has at least one f_i invertible. For every $W \in \mathrm{Mod}(P)$, every f_i acts by zero on $M_\xi \otimes_P W$. Centrality consequently gives

$$\mathrm{RHom}(M_\xi \otimes_P W, X) \simeq \mathrm{RHom}(M_\xi \otimes_P W, \Gamma_{\mathfrak{m}}X),$$
$$\mathrm{RHom}(X, M_\xi \otimes_P W) \simeq \mathrm{RHom}(\Gamma_{\mathfrak{m}}X, M_\xi \otimes_P W).$$

For the orthogonality assertion used here, multiplication by f_i on either argument induces the same endomorphism of the mapping complex. If that endomorphism is simultaneously invertible and zero, the complex vanishes. Finite Čech totalization preserves this vanishing.

Inside $\mathcal{T}_1$, the projection of $\Gamma_{\mathfrak{m}}X$ onto the factor generated by E_ξ is $E_\xi \otimes_P W_X$, where $W_X = \mathrm{RHom}(M_\xi, X)$. Thus the cone of the global evaluation $M_\xi \otimes_P W_X \to X$ is orthogonal to M_ξ in both directions, by the two displayed identities. The evaluation triangle splits, and its orthogonal decomposition is functorial. This gives the claimed product of presentable categories. The same argument handles the finite collection of compact labels simultaneously, using their mutual orthogonality from Lemma 5.2. □

Suppose now that the semisimple rank is positive. Let $\mathcal{J}$ be generated by the objects $L_\eta \otimes \Delta_w Q$ arising from the quasi-split pure labels. Each standard Δ_w is convolution-invertible. By the marked endomorphism calculation of Lemma 2.9, invertibility of the character line, and regular localization in Theorem 2.53, we have

$$\mathrm{End}^0(L_\eta \otimes \Delta_w Q) = A, \qquad (L_\eta \otimes \Delta_w Q)|_U \neq 0,$$

where A is the integral coordinate ring of the closed-seed torus and U is the root-regular open. The projection of each such generator to $\mathrm{Mod}(P)M_\xi$ is zero. Indeed an orthogonal decomposition would give an idempotent in A. This idempotent is zero or one; the latter would make the whole generator supported at one, contradicting its nonzero restriction to U. Hence M_ξ is orthogonal to $\mathcal{J}$, and the compact pure factor extends the quasi-split comparison as an orthogonal summand. Together these factors and J identify, over the entire Cartan, the sum of the quasi-split and compact pure summands with a full spectral subcategory. Their compatibility here is as marked periodic linear categories.

5.2. **Closed real orbits and the Levi of a minimal parabolic.** For a pure label ξ, let $I(\xi)$ be the absolute type of a minimal real parabolic of H_ξ. It is independent of the parabolic, since minimal real parabolics are conjugate under the full group of real points. The type is θ-stable. Fix a type I that occurs. In H_0 choose the corresponding standard parabolic $P_{0,I}$, with its quasi-split Levi M_0, and write L_I for the dual standard Levi in G. Put $\delta_I = w_I\theta$, where w_I is the longest element of W_I, and $A_I = \mathcal{O}(T/(1-\delta_I)T)$. Use $\mathcal{C}_\xi$ for the full localizing category generated by the regular coefficient standard of the closed real flag orbit in $\mathcal{A}^{\mathrm{per}}_\xi$.

Lemma 5.6. *For each ξ of type I there is a unique lift $\eta(\xi)$ to $H^1(\mathbb{R}, M_0)$ associated with its minimal parabolic. The derived group of the Levi $M_{\eta(\xi)}$ is anisotropic, and the Levi acts transitively on its complex flags through its full real group of points. The closed real flag category $\mathcal{C}_\xi$ is equivalent to the entire periodic category of this Levi form, which is* $\mathrm{Mod}(A_I \otimes P)$ *marked by N_{δ_I}.*

Proof. Reductions of the given torsor to $P_{0,I}$, up to isomorphism, are the $H_\xi(\mathbb{R})$-orbits on its real parabolics of type I. Since a parabolic of that type is minimal for this form, they constitute one orbit. The cohomology of $P_{0,I}$ is that of its Levi: its unipotent radical, also after twisting, has a filtration by real vector groups with trivial positive Galois cohomology. This gives the stated lift. A proper real parabolic of the derived Levi would yield a proper real parabolic inside the chosen minimal one. Hence the derived Levi is anisotropic, and its real points form a compact real form of the complex semisimple derived group. The compact group acts transitively on the complex flag variety, for instance by Iwasawa decomposition. Passing to the reductive Levi preserves transitivity.

In the projection of the full complex flag variety to the parabolics of type I, the real minimal parabolic orbit is closed: it is the real locus of a projective real variety. Its inverse image is a single real orbit, by the transitivity just proved on each fibre. This is the closed orbit by the uniqueness of the closed real flag orbit [IY25, Lemma 6.9]. On the enhanced flags over a fixed minimal parabolic $P = MU$, the fibre is $P/N_H \simeq M/N_M$ after

choosing a complex flag in that parabolic. Here the complex radical U is contained in N_H; its real points in the stabilizer are contractible. Equivariant descent over the real homogeneous base therefore identifies the enhanced closed-orbit coefficient category with the entire enhanced category of $M_{\eta(\xi)}$. Extension along the closed real orbit is fully faithful. The single orbit of the Levi is open. Its coefficient description and marking follow from Theorem 1.6 for M_0. □

The choice of a neutral open orbit for M_0 in Lemma 4.14 associates to this unique Levi orbit a character $\chi_\xi \in X^*(\pi_0(T^{\delta_I}))$. Different pure labels of type I give different characters: the attached Levi orbit determines its Levi torsor and hence its image in $H^1(\mathbb{R}, H_0)$.

Write P_ξ for the regular standard of the closed enhanced orbit in the form H_ξ. On the real side these objects suffice even for forms of positive real rank smaller than the quasi-split rank.

Proposition 5.7. *For each pure inner form the finite family $\{P_\xi \star \Delta_w^{\mathcal{H}} : w \in W\}$ compactly generates $\mathcal{A}_\xi$. Its periodic images compactly generate $\mathcal{A}_\xi^{\mathrm{per}}$.*

Proof. Fix the form. Its unique closed orbit supplies the initial regular standard by Lemma 5.6. For any other orbit the real rank-one orbit classification and [IY25, Lemma 6.9] give an ascent from a smaller orbit. Use induction on the dimension and let Δ_x be the regular standard on that smaller orbit. The lower restriction of $\Delta_x \star \Delta_s^{\mathcal{H}}$ has perfect lower coefficients by Lemma 2.3. Indeed the compact enhanced flag model and the finite stratified link calculation there apply to the given group of real points. In a real ascent the input-frame map on the compact torus is a finite cover and the output is the torus quotient of Lemma 2.5. Its derived direct image gives a nonzero finite free regular module on each upper orbit, up to shift and orientation. In a complex ascent the unique input flag gives the regular module on the unique upper orbit by Lemma 2.6. Both calculations use the actual real stabilizers of the two orbits. These stabilizers and the rank-one classification apply to every real form.

The open–closed triangle on the inverse image of the partial-flag orbit now isolates the extension by zero from the union of upper orbits. Each is a component of that union, so taking summands gives its individual regular standard. This proves by induction that every stratum standard belongs to the thick subcategory of compact Hecke translates of P_ξ. The standards generate all compact objects by the finite-stratum argument in Lemma 2.3. Finally the compact Hecke kernels themselves are generated by the finitely many $\Delta_w^{\mathcal{H}}$, so their displayed translates suffice. Scalar extension preserves these compact generators. □

Write H_ξ for the chosen real group, with its full group of real points, and let C be its closed orbit on the complex flag variety of H. Choose an output flag $B \in C$, a real minimal parabolic $P \supset B$, and a maximally split real Cartan $T_0 \subset B$. Let I be the absolute type of P, let δ be the marked Galois involution of its root datum, and let $D \subset T^c$ be the full compact image of $T_0(\mathbb{R})$. Set $M = \pi_1(T^c/D)$, $A = \mathbb{C}[M]$, and $r = \dim T^c$. Write i_C for the enhanced closed-orbit inclusion and $P_C = i_{C*}L_C^{\mathrm{reg}}$.

Lemma 5.8. *The centralizer W^δ of the marked involution in the Weyl group normalizes W_I. Each $\omega \in W^\delta$ has a unique decomposition $\omega = vu$ in which $u \in W_I$ and $v\Phi_I^+ = \Phi_I^+$. The intersection of C with $B\omega B/B$ is $\mathbb{R}^{\ell(v)} \times \mathbb{C}^{\ell(u)}$ in root coordinates. For $\omega \notin W^\delta$ the intersection is empty.*

Proof. Let S be the maximal split subtorus of T_0. The roots vanishing on S are precisely Φ_I, and the derived group of the Levi of P is compact. Thus δ acts by -1 on Φ_I. The real Weyl group of T_0 is W^δ: this is [IY25, arXiv version 2, Lemma 4.14(2)], whose hypothesis is maximal splitness and which uses the full group of real points. It normalizes W_I. Multiplying any of its elements by a unique element of W_I makes it preserve Φ_I^+. These representatives form a subgroup; indeed each sends the positive system of Φ_I to itself. Their elements are the minimal representatives of their W_I double cosets, and $\ell(vu) = \ell(v) + \ell(u)$.

The ordinary closed orbit is the inverse image of the real minimal parabolics in H/P. On that real flag space the rational Bruhat decomposition [Spr98, Chapter 15] has cells $U_v(\mathbb{R})\dot{v}P/P$, indexed by the preceding representatives. Here U_v is the product of the root groups in the inversion set of v. It contains no root of Φ_I, is defined over $\mathbb{R}$, and has complex dimension $\ell(v)$. Exponential coordinates identify $U_v(\mathbb{R})$ with $\mathbb{R}^{\ell(v)}$. The representatives can be chosen in $N_{H_\xi(\mathbb{R})}(T_0)$ by the real Weyl-group equality.

Fix the output B. Over the complex partial-flag cell $U_v\dot{v}P/P$, the Schubert cell with index vu is the product of U_v with the Levi Schubert cell of index u. This follows directly by factoring the positive root groups of its inversion set into those outside Φ_I and those in $v\Phi_I^+ = \Phi_I^+$. The Levi cell is $\mathbb{C}^{\ell(u)}$. Every flag in a real minimal parabolic belongs to C, since its compact derived Levi acts transitively on the complex Levi flag variety.

Conversely every flag in C lies over a real minimal parabolic. Consequently

$$C \cap BvuB/B \simeq \mathbb{R}^{\ell(v)} \times \mathbb{C}^{\ell(u)},$$

and the intersection is empty for all other Schubert indices. □

Lemma 5.9. *With the output flag fixed, put $\omega = w^{-1} \in W^\delta$. Over the cell of Lemma 5.8, integration over the compact torus of the enhancement transports a derived coefficient module by ω_*, tensors with the torus orientation line and shifts by $[-r]$. Thus an output loop $b \in M$ acts by the original action of $\omega^{-1}b$. For the regular local system in degree zero the resulting cohomology is concentrated in degree r.*

Proof. We retain the enhancement in this calculation. The root coordinates of Lemma 5.8 give a frame over the real unipotent cell. On the Levi cell use compact Iwasawa decomposition: its positive Cartan factor disappears in the compact enhancement, and the resulting frame is equivariant for the compact Cartan. Choose an actual real normalizer representative of $\omega = w^{-1}$. After quotienting the contractible positive-Cartan and unipotent stabilizers, the two enhanced coordinates have the form $t_1, t_2 \in T^c$, with the complete stabilizer D acting through $(d, \omega(d))$. Thus the torus part of the correspondence is

$$Z_w = (T^c \times T^c)/\operatorname{graph}(\omega|_D) \simeq (T^c/D) \times T^c, \qquad k = t_2\omega(t_1)^{-1}.$$

The output map is $([t_1], k) \mapsto [\omega(t_1)k]$. Constants introduced by the normalizer frame are retained as fibers of the universal local system; choosing their bases is unnecessary. The frame construction is on the cell, so no transitivity of its compact torus action is being asserted.

Let $\Lambda = \pi_1(T^c)$ and let $j : \Lambda \to M$ be the actual map. It can have a kernel or a nontrivial finite cokernel. The output lattice map is

$$M \oplus \Lambda \longrightarrow M, \qquad (a, \gamma) \longmapsto \omega(a) + j(\gamma).$$

It is surjective, with kernel $\gamma \mapsto (-\omega^{-1}j(\gamma), \gamma)$ and section $a \mapsto (\omega^{-1}a, 0)$. Hence its compact torus fiber has dimension r. The input regular system tensored with the regular Hecke system is the regular system of $\mathbb{C}[M \oplus \Lambda]$. The Koszul complex for a basis of the displayed kernel computes its direct image as the regular $\mathbb{C}[M]$-module in degree r, with its torus orientation line. The same graph Koszul resolution is free on both coefficient sides, so it gives the stated transport equivalence for every derived coefficient module. This calculation retains all components of D, rather than replacing D by its identity component. □

Proposition 5.10. *Use input-first right convolution as in Section 2.1, so the open support of Δ_w has $\operatorname{pos}(B_{\mathrm{in}}, B_{\mathrm{out}}) = w$, and normalize Δ_w by $[r + \ell(w)]$. Then*

$$i_C^*(P_C \star \Delta_w) = 0 \qquad \text{if } w \notin W^\delta.$$

For $w \in W^\delta$, write $w^{-1} = vu$, where $u \in W_I$ and v is the unique representative of its coset which preserves Φ_I^+. Then

$$i_C^*(P_C \star \Delta_w) \simeq L_w^{\mathrm{reg}}[-\ell(u)].$$

Here L_w^{reg} is a regular A-line. More precisely, the closed restriction on arbitrary coefficient modules is the equivalence induced by $w^{-1} : T^c/D \to T^c/D$, tensored with the orientation local system of the real cell and with the line determined by the chosen frame for the normalizer representative, followed by $[-\ell(u)]$. In particular

$$\mathrm{RHom}(P_C \star \Delta_w, P_C) \simeq \mathrm{RHom}_A(L_w^{\mathrm{reg}}, A)[\ell(u)]$$

is a regular Laurent line in one degree. The assertions hold for every pure inner form separately and for the original, uncompleted coefficient rings.

Proof. With the output fixed, the input-first convention tests the cell of index w^{-1}. Lemma 5.8 gives vanishing for the empty intersection. Otherwise Lemma 5.9 computes the integral over the enhanced torus. It remains to integrate along the cell and include the shift of the standard kernel.

The real cell contributes compactly supported cohomology in degree $\ell(v) + 2\ell(u)$. Its orientation character on D is trivial on D^0; the complex Levi cell has its canonical complex orientation. The resulting finite orientation local system on T^c/D is pulled back through $\pi_1(T^c/D) \to \pi_0(D)$. Tensoring it with the regular local system preserves a regular Laurent line. The use of compact supports follows from the standard extension $j_!$ and the proper projection on compact enhanced flags. Combining the cell degree, the torus degree and the standard shift gives

$$(r + \ell(w)) - r - \ell(v) - 2\ell(u) = -\ell(u).$$

The torus orientation is paired with the same orientation used for the convolution unit. Finally C is closed, so adjunction with i_{C*} gives the asserted mapping complex. None of these steps uses completion or localization of the Cartan coefficients. □

Proposition 5.11. *For each absolute type I there is a continuous fully faithful $R \otimes P$-linear functor*

$$\bigoplus_{\xi:\ I(\xi)=I} \mathcal{C}_\xi \longrightarrow \mathcal{T}(V).$$

Write E_ξ for the image of the free closed-orbit standard. Then $\mathrm{RHom}(E_\xi, E_\xi) \simeq A_I \otimes P$ as an algebra over the marked base. The images for distinct labels of type I are orthogonal in both directions in all degrees. For every label there is also an ambient compact object obtained from a compact Levi lift of the same opposite-stratum character.

Proof. Work first with the intrinsic geometric extension of L_I. Write j_I for the opposite open in its parameter stack. By Proposition 4.13 its restriction functor has a continuous fully faithful right adjoint. Apply this right adjoint to the factor of character χ_ξ. The factors for distinct characters are orthogonal, and the compact coefficient generator of the open has endomorphism algebra $A_I \otimes P$ there.

Choose a θ-fixed integral cocharacter λ_I which vanishes on the simple roots in I and is positive on the unipotent radical of the dual standard parabolic. Its image is central in L_I and is contained in $(T^{\delta_I})^\circ$: both w_I and θ fix the cocharacter. Consequently its action on every surviving opposite character factor is trivial. On the whole Levi stack this cocharacter acts trivially on the atlas, so its inertia weights split the category. Restriction to the open preserves the weights. Its right adjoint preserves each weight as well, by adjunction with all the other weight categories. The same decomposition holds for the small circle invariants and after inversion of u; before Ind-completion each coherent object has finite support for this central weight decomposition.

We now apply the fully faithful weight-zero functor $\Phi_I^G = k_{I*} q_I^*$ of Lemma 3.33. For its inherited extension the intrinsic odd representative is $\tau_{L_I} = (r_G r_{L_I}^{-1})^{-1} \tau_G$, as in Lemma 3.34. The correcting element is central in L_I: the difference of the two positive coroot sums pairs to zero with every simple root in I. This identification retains the actual odd square and marked boundary. The functor is fully faithful after small circle invariants since mapping complexes and their actions agree; localization and Ind-completion preserve full faithfulness. We may thus apply the functor to the right-adjoint images before establishing their compactness. Composition with the opposite comparison and Lemma 5.6 proves the functor and the Hom assertions.

For the compact lift take the bounded coherent equivariant lift of the free weight line from Proposition 4.13 in the Levi. One may project to λ_I-weight zero without changing its open restriction. The projection is a direct summand, and Φ_I^G preserves compact objects of weight zero. This produces the claimed compact object. The self-maps computed above concern the right-adjoint image E_ξ. □

Proposition 5.12. *If $I(\xi') \not\subseteq I(\xi)$, then $\mathrm{RHom}_{\mathcal{T}(V)}(E_\xi, E_{\xi'}) = 0$. In particular the seeds are orthogonal in both directions for incomparable types.*

Proof. Write $I = I(\xi)$, $J = I(\xi')$, and denote the closed parabolic incidence used by Φ_J^G by $\mathcal{S}_J$. The proof of Proposition 4.13 writes the Levi right adjoint of the free character line as a sequential colimit of bounded coherent equivariant heart objects $M_{J,n}$. The transition is multiplication by the section cutting out the complement of the Levi opposite open on the classical supporting stack. All these objects have weight zero for λ_J, as in the proof of Proposition 5.11. Hence $E_{\xi'}$ is the colimit of the compact Tate images of $k_{J*} q_J^* M_{J,n}$. The underlying objects are still in the heart since q_J is flat and k_J is an affine closed immersion.

On underlying quasi-coherent complexes their colimit is the direct image from the locus in $\mathcal{S}_J$ where the induced Levi flags are opposite on that classical support. Indeed the Levi colimit is the direct image of the open line by the affine localization described in Proposition 4.13. Flat base change along q_J gives its pullback as the direct image from this inverse-image locus, and the closed pushforward along k_J preserves the colimit. The two ambient flags on that locus have relative position w_J: they lie in the same J-parabolic and their Levi positions are opposite. Every pair on $\mathcal{S}_I$ has relative position in W_I. Each simple reflection in J occurs in a reduced expression for w_J, so w_J belongs to W_I only if $J \subseteq I$. The direct-image locus is consequently disjoint from $\mathcal{S}_I$.

Take a single compact telescope term of E_ξ, represented by its bounded equivariant coherent heart object C supported on $\mathcal{S}_I$. Its derived quasi-coherent restriction to the preceding locus is zero. Adjunction shows that maps from C to the underlying direct image vanish. Into the indicated heart targets, in each fixed degree a bounded-above, degreewise finite resolution of C uses only finitely many terms. A finite affine cover of the incidence atlas followed by reductive invariants therefore commutes with this filtered colimit in each degree. The

mapping complexes have a uniform lower bound, so taking circle invariants commutes with the same colimit degree by degree by the finite-cell truncations of BS^1. Inverting u and using compactness of the Tate image gives $\mathrm{RHom}_{\mathcal{T}(V)}(C, E_{\xi'}) = 0$. Finally E_ξ is the colimit of these terms, so mapping out of it gives a limit of zero complexes. This proves the assertion. □

For each ξ of type I, let c_ξ be the pullback of χ_ξ along $Z(G)^\theta \to \pi_0(T^{w_I\theta})$. This map is defined because every Weyl element fixes the center pointwise.

Proposition 5.13. *If $c_\xi \neq c_\eta$, then for all compact finite Hecke kernels K, K' one has $\mathrm{RHom}(K\cdot E_\xi, K'\cdot E_\eta) = 0$. The localizing subcategories generated by their respective translates are mutually orthogonal.*

Proof. On the Levi opposite orbit the regular weight line has central character c_ξ. The telescope in Proposition 4.13 has this same character termwise: the twisting line is defined by the pairing of two root-lattice representations and has trivial central character. The pullback and closed direct image defining Φ_I^G intertwine the action of $Z(G)^\theta$. Thus E_ξ has character c_ξ in the ambient category. In the convolution the kernel has trivial diagonal central character by Lemma 4.4. Pullback to the flag correspondence, tensor product and equivariant direct image therefore retain c_ξ. The fixed center is diagonalizable over $\mathbb{C}$ and acts trivially on the atlas. Its distinct character categories are orthogonal and closed under colimits. This proves both assertions, including convolution through flags outside the Levi. □

There is also a flag-position restriction within each central character. Use $\mathrm{pos}(B_1, B_2) = w$ when, after translating B_1 to B, the second flag has a representative in BwB/B. For $x, y \in W$, write $x * y$ for the Demazure product. Multiplication on the right by a simple reflection s replaces x by the larger of x, xs in Bruhat order. Iteration along a reduced word gives the general product.

Proposition 5.14. *Let K be a bounded coherent circle-equivariant kernel on the Steinberg stack supported on pairs with input-to-output relative position at most v. Write $I = I(\xi)$ and $J = I(\eta)$. If*

$$w_J \not\leq v^{-1} * w_I * \theta(v),$$

then $\mathrm{RHom}_{\mathcal{T}(V)}(K \cdot E_\xi, E_\eta) = 0$. The vanishing extends to the thick subcategory of periodic kernels generated by such kernels.

Proof. On the support of one compact telescope term of E_ξ, the input flag B_0 and its conjugate aB_0a^{-1} are refinements of a common I-parabolic. Their position is at most w_I. In the convolution, a point of the support of the relative tensor product with K supplies a second flag B_1 with $\mathrm{pos}(B_0, B_1) \leq v$. The sequence

$$B_1,\ B_0,\ aB_0a^{-1},\ aB_1a^{-1}$$

has successive positions bounded by $v^{-1}, w_I, \theta(v)$. For such a sequence the position of the two ends is at most the Demazure product of the bounds. Indeed refine each bound into a reduced word, permitting repeated flags at each simple step. The images of these galleries are the corresponding Schubert closures. On appending a simple step the maximal endpoint changes from x to $\max(x, xs)$, by the two Bruhat cells in a minimal parabolic. Induction gives the asserted bound, also for specializations. The convolution projection is proper, so its direct image has support in that same closed bound on the position of B_1, aB_1a^{-1}.

By Appendix B, convolution with the chosen kernel takes the compact telescope term to an equivariant bounded coherent complex. Its restriction to the direct-image locus of the target telescope in Proposition 5.12 is zero, because every pair there has position w_J. The degreewise colimit and adjunction calculation of that proposition applies to this bounded coherent source, using its upper cohomological bound for the circle-invariant calculation. Maps from the convolved term to E_η therefore vanish after inverting u. Convolution commutes with the colimit forming E_ξ, and maps out of this colimit are a limit of zero complexes. Exactness and passage to summands give the last assertion. □

For any θ-stable set of simple roots I, write $P_I = L_IU_I$ in the dual group, and let Φ_I be its Levi root system. In the marked Cartan put

$$T[I] = \{t \in T : e^\alpha(t) \neq 1 \text{ for all } \alpha \in \Phi(G,T) \setminus \Phi_I\}.$$

All the roots of L_I may be singular on this open. The stable parabolic determines a whole component there.

Theorem 5.15. *Over $T[I]$, the parameter stack V_{L_I} with its inherited odd component is isomorphic, with its original circle and marking, to the open and closed substack of V on which the odd element preserves the I-parabolic of the flag. On small Tate categories Φ_I^G on this open extends to the inclusion of a direct factor on all weights. If $I(\xi) = I$ and $I(\eta) \not\subseteq I$, the restrictions of E_ξ, E_η to $T[I]$ are orthogonal in both directions.*

Proof. Use the parabolic incidence $\mathcal{S}_I$, with maps k, q of Lemma 3.33. Let F be the given flag, $P = P_I(F)$, and $h = a^2$. Two parabolics over our point contain h, namely P and aPa^{-1}. They have the same type since I is θ-stable. In the partial Grothendieck resolution $\{(h, P') : h \in P'\} \to G$, the relative tangent complex at the first point is the two-term complex with differential $1 - \mathrm{Ad}(h)$ on $\mathfrak{g}/\mathfrak{p}$, in compatible translations. Its eigenvalue test can be made in the output Borel: in a weight filtration the diagonal entries of $\mathrm{Ad}(h)$ are the negative radical-root evaluations on the mark. Thus the partial resolution is étale at this lift. Conjugation by a shows the same for the second lift, because a commutes with h. The diagonal of the separated étale locus over G is open and closed. Pulling it back along the two lifts makes their equality an open and closed condition on the underlying parameter stack.

We verify its derived structure. In the atlas with fixed output Borel, the square sends the pair $(G\tau, P_I\tau)$ to (G, P_I). At a point with $a \in P_I\tau$, its normal differential in right translations is $1 + \mathrm{Ad}(a)$ on $\mathfrak{g}/\mathfrak{p}_I$. Its product with $1 - \mathrm{Ad}(a)$ is $1 - \mathrm{Ad}(h)$, already proved invertible. The local defining equations of P_I thus pull back to independent normal equations of $P_I\tau$. More explicitly they vanish on that smooth coset and generate its conormal space at the point. Nakayama's lemma gives its ideal locally, and the Koszul complex resolves the quotient since the codimension is unchanged. Factoring the further derived base change by $B \to P_I$ now identifies k with the inclusion of the entire indicated open and closed derived substack.

On its atlas the Levi projection leaves the smooth factor U_I and the U_I gauge action. For a Levi element a_L in the odd component, conjugation of a_L by $v \in U_I$ gives radical coordinate $v\,\mathrm{Ad}(a_L)(v)^{-1}$. On $\mathfrak{u}_I$, $1 - \mathrm{Ad}(a_L)^2$ is invertible by the positive radical-root evaluations on the same mark. Thus $1 - \mathrm{Ad}(a_L)$ is invertible there and on the successive quotients of the lower central series. The map in radical coordinates is an isomorphism: successively solve on those vector quotients with the indicated invertible linear operators. This gives a regular inverse in families, since only their invertible determinants occur in the denominators. The calculation also holds over the derived Levi incidence by pullback from the smooth open of odd Levi elements defined by this determinant condition. The radical factor is therefore a free transitive gauge factor, leaving precisely the Levi parameter stack. Both maps k, q preserve transport by the actual square, which proves the statement about the circle. Circle invariants, inversion of u, and Ind-completion respect a finite disjoint union.

The restricted E_ξ lies in the Levi component by construction. For $J = I(\eta) \not\subseteq I$, the right-adjoint telescope defining E_η comes from flags of relative position w_J, as in the proof of Proposition 5.12. This locus misses the Levi component, since flags in a common I-parabolic have position in W_I. Restricting the underlying direct image to the open and closed component gives zero. The same degreewise calculation against compact equivariant objects as in that proposition applies on this open. Thus E_η has zero component in the Levi factor. The direct product decomposition gives both vanishings. □

5.3. **Compact closed-orbit generators for all pure inner forms.** For each minimal Levi choose the neutral open orbit so that its image in the adjoint group is the Whittaker origin for the discrete series packet of the trivial compact representation. Such a choice exists: the central quotient induces an isomorphism of the complex flag varieties, and any point of the prescribed open orbit defines a neutral open orbit upstairs by the opposite-pair criterion. For the neutral torus of semisimple rank zero the adjoint group is trivial.

Lemma 5.16. *Let M_0 be a quasi-split real group with a pure inner form compact modulo its center. Write $L = \widehat{M_0}$, $L_s = \widehat{(M_0)_{\mathrm{ad}}}$ and $C = Z(L)^\circ$. The central isogeny $p : L_s \times C \to L$ induces a map π of parameter stacks commuting with the original boundary circle. It factors as a gerbe for a finite diagonalizable group followed by a representable finite étale map. Its bounded coherent pullback and direct image are adjoint in both orders. These adjunctions persist in the compact-first Tate categories and commute with restriction to the opposite open.*

Proof. Set $K = \ker p$ and $\delta = w_I\theta$, with w_I the longest element for M_0. The projection from K to $Z(L_s)$ is injective since $C \to L$ is an inclusion. On the Cartan of L_s the involution δ is inversion: the derived group of the given real form is compact and acts transitively on its flags, whose marked real Cartan therefore has no split part on the semisimple roots. Since w_I acts trivially on the center, $\theta(k) = k^{-1}$ for every $k \in K$.

The two geometric extensions are related by sending the pinned odd element upstairs to the pinned odd element downstairs. The cocharacters which are sums of positive coroots are carried to one another by p; the torus has no roots. Hence this map respects the actual central squares. The two flag varieties are the same. On the incidence atlases the map of odd elements is a K-torsor, and the flag equation is pulled back, including its derived structure: it asks that the square fix the chosen flag, an identical action on the common flag variety. Write $\widetilde{X} \to X$ for this torsor of incidence atlases. Translation by k sends $\tilde{a}$ to $k\tilde{a}$, while gauge conjugation by k sends it to $k^2\tilde{a}$. Both modified lifts have square $\tilde{a}^2$ since $k\theta(k) = 1$. Thus $K[2]$ acts trivially as gauge and

$K/K[2]$ acts freely by translation through K^2. Rigidifying this trivial gauge subgroup gives

$$[\widetilde{X}/(L_s \times C)] \longrightarrow [\widetilde{X}/((L_s \times C)/K[2])] \simeq [(\widetilde{X}/K^2)/L] \longrightarrow [X/L].$$

The first map is a $BK[2]$-gerbe. The second is the finite étale torsor for K/K^2, as may be checked on the atlas X. This distinguishes the ineffective gauge group from the deck transformations of the odd lift. The original map comes from applying a homomorphism of geometric extensions to framed local systems. It commutes with rotation of the boundary. The factorization respects rotation as well: the central translations preserve the transported flags and their actual closing square.

For the gerbe, pullback identifies sheaves downstairs with the summand of trivial $K[2]$-weight, and direct image projects to this summand. For the finite étale torsor, the trace pairing identifies the left and right adjoints to restriction of scalars. These operations preserve bounded coherence and the adjunctions are natural under base change. Their units and counits, given by the finite weight projection and the finite trace, are invariant under rotation. Passing to small invariants, inverting the same u, and taking Ind-completions gives the stated functors and adjunctions. Open base change holds on the above two presentations, and the inverse image of the opposite open is the opposite open upstairs. □

We first work on the product cover of Lemma 5.16. Its factors have dual groups L_s (semisimple simply connected) and C (a torus). On the real side the semisimple factor corresponds to the adjoint group of the minimal Levi. We use the opposite open in each of these parameter stacks.

Lemma 5.17. *For the compact pure form of the adjoint real Levi, let F_s be its line on the opposite open and j_s the inclusion. The object $E_s = j_{s*}F_s$ is compact. Maps to or from E_s are computed by restriction to the character factor of F_s.*

Proof. First consider the adjoint real group $(M_0)_{\mathrm{ad}}$. The pure form induced by the chosen label is compact. If this adjoint group is trivial, its parameter category is $\mathrm{Mod}(P)$ and the next identification is the identity. In positive semisimple rank, Proposition 5.5 supplies an orthogonal factor $\mathrm{Mod}(P)$ with compact generator E_s in the global parameter category $\mathcal{T}(V_s)$. After a shift it restricts to its character line F_s on the open $B(T_s[2])$, where T_s is the Cartan of L_s. To check this identification, use the open immersion of the finite-stabilizer orbit into the unit ABV stack in the construction of [BCHN25, Theorem 4.4]: the identification with unipotent loops restricts to this open, restriction from unipotent to formal loops commutes with open restriction, and the Koszul equivalence commutes with schematic pullback [Che23, Theorem 3.4.9]. Riemann–Hilbert likewise commutes with the open immersion. On a finite stabilizer the tangent complex is zero, and the character line retains its weight. The norm and central-square normalization here are those of Lemma 5.3.

Let j_s be this spectral opposite open. By adjunction $\mathrm{RHom}(E_s, j_{s*}F_s) = P$. The object $j_{s*}F_s$ is indecomposable since its degree-zero endomorphism ring is $\mathbb{C}$. Its projection to the orthogonal factor generated by E_s is nonzero, so it lies in that factor and is E_s itself. The complement has zero F_s-component under restriction: maps from that component to the free module vanish by the right adjunction, and a free module over the graded field P detects every nonzero module by maps into it. Thus restriction computes maps to and from E_s. □

Lemma 5.18. *On the product cover the pullback of the opposite character line of the given pure label has a compact extension. Restriction computes maps both to and from this extension in its character factor.*

Proof. Take E_s and F_s from Lemma 5.17. We describe the character after pullback along the central isogeny. Write J_f for the marked real Cartan of the neutral open orbit. Its map to the adjoint Cartan sends the cocycle of an open orbit in any pure twist to the cocycle of the image orbit, by the formula $g^{-1}z\sigma(g)$ of Lemma 4.14. For a regular discrete series parameter, marking the fundamental Cartan by its positive chamber identifies these orbit labels with the admissible embeddings of its torus. In the adjoint group the pure enhancement is the Tate–Nakayama character of this cocycle measured from the Whittaker-generic embedding [AK26, Lemmas 2.3.1, 4.1.1 and Section 4.3]. The chosen real Levi is transitive on its flags, and its image is the unique flag orbit of the compact adjoint form. Hence the pullback of its open character χ_ξ to $T_s[2]$ is exactly the character of F_s. This step uses the pure packet comparison only for the semisimple adjoint quotient.

For the remaining torus factor put $D_C = C^\theta$, $Y_C = C/(1-\theta)C$. The product-cover parameter stack is $V_s \times [C/C]_\theta$. Splitting the quotient of tori identifies its second factor with $Y_C \times BD_C$. The full pullback of the open line has weight (F_s, λ) on $T_s[2] \times D_C$. Here λ is trivial on D_C°, since the original character factors through $\pi_0(T^\delta)$. Any line factor on the affine torus Y_C is trivializable. On weight λ the second positive turn is the scalar $\lambda(c\theta(c)) = 1$. Thus this weight category is obtained from the category on $V_s \times Y_C$ with the circle on V_s, by tensoring with the weight line. This computation uses the diagonal boundary rotation with its original period.

Over Y_C take $E_s \boxtimes \mathcal{O}_{Y_C}$, and tensor by the weight line. It is compact since product pullback preserves bounded coherent representatives and their retracts. To verify its two clean adjunctions, compare its maps to a compact X, and maps from X, with the corresponding maps on the open. Other stabilizer weights give zero in both directions. Within its weight, specialize these comparisons at any closed $y \in Y_C$. The finite-Tor inclusion of y has coherent pullback and direct image. The projection formula identifies tensoring with $\mathbb{C}_y$ with their composite. Adjunction therefore gives exactly the two comparisons for E_s, hence both are isomorphisms at every y.

These closed points detect the comparison cones. Indeed their two periodic cohomology groups are finite over $R_s \otimes \mathcal{O}(Y_C)$, where $R_s = \mathcal{O}(T_s)$. For the source of either comparison this follows from Theorem B.11, applied to the product with a split torus whose dual is Y_C and restricted to neutral torus inertia: its added marked coordinates are the squares of the coordinates on Y_C, and finite generation persists over the larger coordinate ring. On the open the claim is the corresponding calculation for perfect modules over $\mathcal{O}(Y_C) \otimes P$. Kernels and cokernels preserve finiteness. For a nonzero such cone choose a maximal ideal where some cohomology is nonzero. Its restriction to $\mathcal{O}(Y_C)$ is a closed point y. In the local ring, tensoring successively with the Koszul factors for its torus coordinates minus their values at y preserves nonvanishing by Nakayama and preserves finite cohomology. This contradicts the vanishing of the derived fibre. The two global comparisons are therefore isomorphisms. They extend to arbitrary objects by compactness and colimit preservation. We have obtained a two-sided compact extension on the product cover. □

Proposition 5.19. *For every pure label of minimal Levi type I, the right-adjoint extension of its opposite character line in $\mathcal{T}(V_{L_I})$ is compact. On that character factor this extension is also left adjoint to restriction followed by the character projection.*

Proof. Write j for the opposite open downstairs and $\tilde{j}$ for its inverse image. Let F_0 be the downstairs free character line and π_o the restriction of π to the opens. By Lemma 5.18, $\pi_o^* F_0$ has a two-sided compact extension $\widetilde{E}$. For any object X downstairs, Lemma 5.16 identifies the maps in both directions between X and $\pi_* \widetilde{E}$ with those between $j^* X$ and $\pi_{o*}\pi_o^* F_0$; for example

$$\mathrm{RHom}(X, \pi_* \widetilde{E}) = \mathrm{RHom}(\pi^* X, \widetilde{E}) = \mathrm{RHom}(\pi_o^* j^* X, \pi_o^* F_0) = \mathrm{RHom}(j^* X, \pi_{o*}\pi_o^* F_0).$$

The reverse direction uses the other adjunction for π and π_o. The line F_0 is a summand of $\pi_{o*}\pi_o^* F_0$: the weight-zero gerbe contributes the identity, and the finite torsor unit splits by the trace divided by its degree. The restriction isomorphisms identify endomorphisms of the extension with endomorphisms on the open. The projector to F_0 therefore lifts and splits on the compact object $\pi_* \widetilde{E}$. Its image is a compact two-sided extension restricting to F_0, hence equals $j_* F_0$. This argument works with the actual deck correspondence and requires no choice of descent of an individual lifted character. The adjunctions extend from the free line to its full factor by the module description on the open and by continuity. □

Corollary 5.20. *There is a continuous, compact-preserving, fully faithful $R \otimes P$-linear functor*

$$\bigoplus_{\xi \in \Theta} \mathcal{C}_\xi \longrightarrow \mathcal{T}(V)$$

taking the closed real regular standards to the seeds E_ξ.

Proof. Proposition 5.19 and the stable-Levi functor give compactness. Equal types are treated by Proposition 5.11, and incomparable types by Proposition 5.12. For $I \subsetneq J$ consider the images inside V_{L_J}. The stable-Levi embeddings through nested parabolics compose: the inverse image of the I-parabolic of L_J in the ambient J-parabolic is the I-parabolic, and the atlas of the derived incidence is the product with the J-radical. Flat base change followed by composition of the closed pushforwards proves the assertion with full gauge groups and the original rotation. Both objects are in the weight-zero range for $L_J \to G$, since its defining cocharacter is central in L_J and fixed by $w_I \theta$. The object from I restricts to zero on the opposite open of L_J: its two flags share a proper parabolic. The two adjunctions of the proposition now give vanishing in both directions within L_J, hence also in G by full faithfulness. This establishes the simultaneous embedding for all labels. □

Proof of Theorem 1.5. The embedding and compactness follow from Corollary 5.20. Within the minimal Levi, let p be restriction to the opposite open followed by projection to the given character. Its fully faithful section e is both left and right adjoint by Proposition 5.19. The counit $epX \to X$ has cone in $\ker p$, which is orthogonal to the image of e in both directions by the two adjunctions. Thus this image is an orthogonal direct factor. Over $T[I]$ the entire Levi category is itself an orthogonal factor by Theorem 5.15. Restriction of a direct product to a base open preserves its two projections. This proves the factor assertion. If H_ξ is compact modulo its center,

its minimal parabolic is the entire group, so $T[I] = T$ and all its flags form one orbit. The full automorphic summand therefore occurs as an orthogonal factor over T. □

Take the semisimple adjoint real quotient, with dual group G_s. Keep the compact parameter a_0, put $K = Z_{G_s}(a_0)$, and write $\mathcal{B} \to \mathcal{P}_s$ for the smooth proper projection of complex flag varieties. The compact parameter determines the clean extension of an equivariant local system $L = j_!F[\dim O] = Rj_*F[\dim O]$ on the open K-orbit O of $\mathcal{B}$; its stabilizer is finite and F has rank one (Lemma 5.2). Let p_s denote this flag projection and the induced map on parameter stacks.

Lemma 5.21. *For each simple root s the proper direct image of L on $[\mathcal{P}_s/K]$ is zero.*

Proof. The image of O is a single open K-orbit Q in $\mathcal{P}_s$. Over $p \in Q$ the fibre $\mathcal{B}_p$ is $\mathbb{P}^1$, and $O \cap \mathcal{B}_p$ is a nonempty open orbit of the stabilizer K_p. The identity component of the stabilizer, K_p°, has dimension one, since the flag stabilizer is finite. It is affine and has an open orbit of the same dimension. The intersection $O \cap \mathcal{B}_p$ is irreducible, so this connected group acts transitively there. Its orbit is $\mathbb{A}^1$ or $\mathbb{G}_m$, by the classification of connected affine algebraic groups of dimension one. Locally analytically on Q the fibration and open suborbit are products: pull back along local sections of $K \to Q$. Thus cleanness restricts to cleanness on the fibre. A rank-one local system on $\mathbb{A}^1$ is constant and has a nonzero direct-image stalk at the missing point. This excludes that case. On $\mathbb{G}_m$ the same punctured-disc calculation forces the monodromy $m \neq 1$. Its cochain complex is $[\mathbb{C} \xrightarrow{m-1} \mathbb{C}]$, so its cohomology vanishes. Over a point outside Q the fibre consists entirely of boundary flags, where L has zero stalks. Proper base change gives $R(\mathcal{B}/K \to \mathcal{P}_s/K)_*L = 0$. □

Lemma 5.22. *On the semisimple adjoint real quotient, the corresponding compact spectral object E_s has zero direct image under the partial flag projection.*

Proof. To transport Lemma 5.21 to the parameter category, first identify the relevant formal neighborhoods. For either flag variety $X = \mathcal{B}, \mathcal{P}_s$, the formal neighborhood of $[X/K]$ in the geometric parameter stack is the formal loop space of $[X/K]$ at its constant loops. Indeed $[X/K]$ is the closed locus over the semisimple odd orbit of a_0, with $a_0^2 = 1$. Conjugation on $G_s \times^K K$, $[g, h] \mapsto ga_0hg^{-1}$, induces an isomorphism of completions at $h = 1$ and at that orbit. Its tangent map is an isomorphism by the decomposition of the Lie algebra into the two eigenspaces of $\mathrm{Ad}(a_0)$; the map is étale there and identifies the closed orbits. In this slice the square is $b = h^2$. Squaring is a conjugation-equivariant formal isomorphism at 1, and the remaining derived flag equation says precisely that b fixes $x \in X$. This is the derived loop incidence for $[X/K]$. These identifications commute with the flag projection and the given boundary rotation. In the boundary-local-system description the restricted even local system has holonomy b; in the slice its structure group is reduced to K. Changing the base flag along a boundary arrow transports the odd holonomy by conjugation and transports b by the same arrow. Conversely, the formal odd holonomy is $a_0\sqrt{b}$; the formal square root commutes with conjugation by K. The identifications respect subdivision and composition of the arrows, and the positive closing arrow is $b = a^2$. This is the unit case of the loop construction in [BN13, Theorem 4.12 and Corollary 4.13], expressed in boundary coordinates.

The regular holonomic D-module corresponding to L is compact on $[\mathcal{B}/K]$: its local system on the finite-stabilizer open is compact, and extension by zero is left adjoint there to continuous open restriction. Apply the ordinary, unrenormalized graded Koszul equivalence of [Che23, Theorem 3.4.12]. Its inverse sends this D-module to a compact graded Tate object on formal constant loops; this theorem commutes with proper direct image. By Lemma 5.21 the graded loop direct image of this object for $\mathcal{B}/K \to \mathcal{P}_s/K$ vanishes. Forget the auxiliary $\mathbb{G}_m$-grading and restrict along the affinization $S^1 \to B\mathbb{G}_a$. This operation takes small equivariant coherent representatives to small circle-equivariant coherent representatives and sends the inverted degree-two generator to the inverted original boundary class with the positive orientation. It commutes with proper direct image. Coherent objects on the formal constant-loop neighborhood are exactly coherent objects set-theoretically supported on its zero section in the ambient stack, by the ordinary convention of [BCHN25, Appendix A.2.1]. Their inclusion into the ambient coherent category commutes with the proper flag projection, by composition on supported objects. These functors extend through idempotent completion. We have produced a compact global object H killed by p_{s*}.

On the finite-stabilizer opposite open its restriction is the periodic character line F, up to a shift: schematic open pullback commutes with the same graded theorem, and on the finite quotient the tangent complex is zero. The adjoint compact seed $E_s = j_*F$ belongs to an orthogonal $\mathrm{Mod}(P)$-factor by Proposition 5.5 and Lemma 5.17. Adjunction gives $\mathrm{RHom}(H, E_s) = P$ up to shift. The projection of the compact H to this factor is therefore E_s up to shift, since duality detects this assertion for perfect modules over the graded field P. Consequently E_s is a

summand of H up to shift, and $p_{s*}E_s = 0$. This argument uses no ungraded equivalence for the partial flag category. □

Proposition 5.23. *In the Levi parameter category let E be the compact character extension of Proposition 5.19. For every simple root s of the Levi, let p_s forget the full flag down to its s-partial flag, retaining the odd element and the derived condition that its square fixes the flag. Then $p_{s*}E = 0$. In particular the finite kernel $C_s = p_s^* p_{s*}$ kills the ambient closed seed for every $s \in I(\xi)$.*

Proof. Lemma 5.22 proves the assertion on the adjoint real quotient. On the product cover of Lemma 5.16, partial flag projection is the product of the proper map of Lemma 5.22 with the identity on the twisted central torus factor. Proper base change therefore kills the external compact extension of Lemma 5.18, with its original torus weight. The finite cover maps commute with the proper flag projections, so composition of direct images kills its finite direct image as well. The downstairs Levi seed is a retract of that direct image and is therefore killed. Compatibility of the stable-Levi functor with the s-kernel gives the ambient assertion. □

Corollary 5.24. *Let $I = I(\xi)$ and $J = I(\eta)$. For $w \in W$, take the shortest element d of $W_J w W_I$ and a length-additive factorization $w = vdu$ with $v \in W_J$, $u \in W_I$. For the input object on the right of the kernel, there is an isomorphism of underlying periodic complexes*

$$\mathrm{RHom}(\Delta_w \cdot E_\xi, E_\eta) \simeq \mathrm{RHom}(\Delta_d \cdot E_\xi, E_\eta)[\ell(u) + \ell(v)].$$

In particular if $w \in W_J W_I$, this coefficient is zero for distinct labels. For equal labels and $w \in W_I$ it is $(A_I \otimes P)[\ell(w)]$. If I is the full set of simple roots, $\mathrm{Loc}(E_\xi)$ is stable under the entire finite Hecke category and is mutually orthogonal to all Hecke translates of the closed generators of the other pure labels.

Proof. For a simple Levi reflection the standard filtration of the tilting kernel gives the triangle $\Delta_s \to \Xi_s \to \mathbf{1} \to \Delta_s[1]$ [Tay25, (4.4)]. The finite comparison takes Ξ_s to the parabolic averaging kernel, so Proposition 5.23 makes the middle term zero on the seed. Thus $\Delta_s \cdot E_\xi \simeq E_\xi[-1]$ for $s \in I$, and the inverse kernel acts by the opposite shift. Iterating along reduced Levi words gives the shifts for u and v. The length-additive factorization is obtained by successively removing a left descent in J or a right descent in I until the double-coset minimum is reached. The standard product for this factorization is $\Delta_v * \Delta_d * \Delta_u$. Apply the shift for the rightmost factor to the source and move the leftmost factor to the target by its inverse. If $d = 1$, Corollary 5.20 computes the remaining coefficient. For equal labels the resulting degree agrees with the periodization of Proposition 5.10.

For the last assertion all standard kernels and their inverses preserve $\mathrm{Loc}(E_\xi)$, by the shift computation. The standards generate every compact kernel by cones and retracts, so all compact kernels and their two adjoints preserve this same subcategory. Moving a compact kernel across Hom by adjunction and using Corollary 5.20 gives orthogonality for the translates. These translates are compact, so the orthogonality passes to the localizing categories they generate. □

6. The decomposition by pure inner form

The compact objects just constructed give the factors of the parameter category. We first prove generation by finite Hecke translation. To prove orthogonality we compute the maps from a single standard translate to a closed generator. Finally, we compare the full factors over the open set determined by the rank of the roots with value one.

6.1. Hecke generation for the entire pure inner class.

Write q_s for the map on the spectral parameter spaces forgetting the simple flag at s, and $D_s = q_s^* q_{s*}$ for the action of the averaging kernel. If the two full flags have position x, their partial flags have position the double coset $W_s x W_{\theta(s)}$. Here the group on the right is generated by $\theta(s)$, since the second flag is the odd translate of the first. In the following calculation we use the quotient at a position from Theorem 4.20.

For the following three lemmas fix $w\theta(w) = 1$ and a simple left descent s of w. Write v for the shortest element of $W_s w W_{\theta(s)}$, and put $\delta = w\theta$ and $\epsilon = v\theta$. The maps r_v and r_w forget the chosen step of the flag on the intrinsic twisted loops of the two corresponding orbits. Both land in the twisted loops of the same partial orbit.

Lemma 6.1. *The support of $D_s C_{v,\lambda}$ is contained in the closure of position w. Its coefficient at w is $r_w^* r_{v*}(A_{v,\lambda}|_{S_v})$.*

Proof. A left descent implies that $\delta^{-1}\alpha_s = \delta\alpha_s$ is negative. The identity $w\theta(w) = 1$ makes $\theta(s)$ a right descent. If $sw \neq w\theta(s)$, the double coset has four elements with maximum w and minimum $v = sw\theta(s)$. Otherwise it has two elements and $v = sw = w\theta(s)$. This second case is exactly $\delta\alpha_s = -\alpha_s$. In either case $v\theta(v) = 1$. The

preimage of the partial orbit closure is the full orbit closure of its maximum. Indeed over each partial pair the two fibers are projective lines, and the inverse image of the partial cell has a dense full cell at its maximum. The projection is smooth, so the closure of this preimage is the preimage of the closure. This gives the support bound on pull–push.

Let O_p be the orbit of these partial pairs. Over it the minimum full orbit O_v is closed and the maximum O_w is open in the inverse image. Remove the boundary of O_p from the partial pair space. There the input object over O_p is the direct image from its intrinsic twisted loops S_v of its character line: the support X_v restricts to the smooth reduced S_v. Denote by S_p the intrinsic twisted loops of $[O_p/G]$, and by r_v, r_w the maps to it. The coefficient of the output at the maximum is the object on S_w

$$r_w^* r_{v*}(A_{v,\lambda}|_{S_v}).$$

Here are the base changes in this formula. The direct image from the minimum factors through S_p by composition. Pullback to the full parameter stack of a direct image from S_p is computed by the derived Cartesian square of twisted loop stacks. On pairs before taking twisted loops this square is the pullback of O_p along the smooth map forgetting the two flags. In its inverse image the maximum is open. Removing the other positions in the fiber, then taking twisted loops, identifies the pullback there with the closed direct image from S_w of the displayed object. Twisted loops preserve derived fiber products by their equalizer description. The closed immersion on the partial side is affine, so the base-change identity for the direct image follows by extension of scalars on the atlas. Lemma 4.19 and Theorem 4.20 now identify the coefficient as asserted. All these maps preserve the original even rotation. □

Lemma 6.2. *If $sw \neq w\theta(s)$, the coefficient of Lemma 6.1, as the input character varies, gives a free line in every character factor at w.*

Proof. In the four-element case, fix a partial pair lifted at its minimum to $(B, {}^vB)$. Its stabilizer is $B \cap {}^vB$. In fact adding $-\alpha_s$ to the first Borel and $-v\theta(\alpha_s)$ to the second adds no common root: each was negative for the other Borel by minimality, and the two added roots are distinct. The intersections are connected, as is seen from the common torus and root groups. Thus $O_v \to O_p$ is an isomorphism. Simultaneously changing the first flag by s and the second by $\theta(s)$ gives the torus marking at the maximum, with $\delta = s\epsilon s$. The map r_w on the torus quotients is induced by this change of marking, up to a translation coming from the odd lifts. It is an isomorphism there: on the torus stabilizers it is conjugation by a representative of s, and on the odd torus cosets it is the same conjugation followed by the change of base lift. It permutes their finite characters. Pullback therefore gives a free line in every surviving factor by Lemma 4.16. □

Lemma 6.3. *If $sw = w\theta(s)$, the coefficients of Lemma 6.1 generate exactly the character factors at w on which $\alpha_s^\vee(-1)$ acts trivially.*

Proof. In the two-element case the common Levi of the partial stabilizer has semisimple rank one. More precisely its root groups outside the common torus consist of the two groups for $\pm\alpha_s$ and its unipotent radical: this follows from $v\theta(\alpha_s) = \alpha_s$ and the same intersection of the two parabolic root sets. Its action on the two projective lines of flags is through this rank-one Levi. The pairs of flags of minimum position are the graph of the equivariant identification of those two lines. Indeed their stabilizer in the common Levi is a Borel, projecting in each factor to the chosen Borel. The maximum consists of the complement of that graph. Odd transport exchanges the two factors and preserves the graph. It thus defines one projective transformation of the common line, followed by interchange. Over a point of S_w it exchanges the two distinct chosen points. Such a transformation is a nonidentity involution: in coordinates placing the points at zero and infinity it is $x \mapsto c/x$ with $c \neq 0$. Its two fixed lines are distinct. These lines parametrize the flags of minimum position lying over this partial parameter.

Over the partial open where the fixed lines are simple, r_v is finite étale of degree two. In a frame its equation is the fixed-point equation on the projective line; its derivative is invertible at each fixed point. This also proves the assertion for the derived equality. The open contains the entire image of S_w. It follows that the coefficient of Lemma 6.1 pulls back to a vector bundle on the atlas Ta_w whose fiber is the sum of the two input lines over these fixed flags. To compute its characters write $D = T^\delta$, $D' = T^\epsilon$ and $K = \ker(\alpha_s : D \to \mathbb{G}_m)$. The finite action of D on the two fixed flags has kernel K: on the projective coordinate it acts by $x \mapsto \alpha_s(d)x$, with $\alpha_s(d) \in \{1, -1\}$. In the common Levi, K is central, and its inclusion in the stabilizer torus at either minimum flag is the same. To see this, diagonalize the projective transformation by conjugation in the rank-one derived group; this conjugation fixes K pointwise. The transported involution on root data at that flag is ϵ, which fixes α_s. Thus K acts on each input line by $\lambda|_K$.

The tori satisfy

$$(D')^\circ = D^\circ\, \alpha_s^\vee(\mathbb{G}_m), \qquad K \cap (D')^\circ = D^\circ\, \langle \alpha_s^\vee(-1)\rangle.$$

For the first equality, $\epsilon = s\delta$ changes the negative coroot line to a fixed line and agrees with δ on the root hyperplane. This gives equality of Lie algebras, hence of connected subtori. The second equality uses $\alpha_s \circ \alpha_s^\vee(t) = t^2$ and $D^\circ \subset K$. If $\chi(\alpha_s^\vee(-1)) = 1$, its restriction to K is therefore trivial on $K \cap (D')^\circ$ and extends to a character λ of $\pi_0 D'$. Indeed characters of a subgroup of a finite abelian group extend over $\mathbb{C}$. If $D = K$, the two lines have this same character. Otherwise they form one orbit and the representation on their sum is induced from $\lambda|_K$, so contains χ. Conversely the input character is trivial on $\alpha_s^\vee(-1)$, so a character taking the value -1 on that element cannot occur. This applies at every torus parameter. For a fixed input λ, each character extending $\lambda|_K$ occurs with rank $2/|\alpha_s(D)|$, also when the double cover has monodromy over Y_δ. After character projection the periodic coefficient is thus a vector bundle of this positive rank on Y_δ. Here one may restrict to the torus section of p_w; it is an inverse on the Tate category by Lemma 4.16. The bundle in this character is a compact generator on the affine torus: its dual is finite projective, and evaluation split by coevaluation divided by the rank exhibits the free module as a retract of its tensor with the dual. Tensoring by a finite projective module is itself a summand of a finite direct sum. This proves the required generation for the indicated characters. □

Lemma 6.4. *Suppose $w\theta(w) = 1$ and s is a simple left descent of w. Set $\delta = w\theta$, let v be the minimum of $W_s w W_{\theta(s)}$, and put $\epsilon = v\theta$. All the objects $D_s C_{v,\lambda}$ are supported on positions at most w. If $\delta\alpha_s \neq -\alpha_s$, their images generate the entire quotient at w. If $\delta\alpha_s = -\alpha_s$, their images generate all factors of the quotient whose character satisfies $\chi(\alpha_s^\vee(-1)) = 1$.*

Proof. Lemma 6.1 gives the support bound. In its four-element case Lemma 6.2 gives the regular generators of every factor. In its two-element case $\delta\alpha_s = -\alpha_s$, and Lemma 6.3 gives precisely the indicated factors. □

Proposition 6.5. *Let I be θ-stable and suppose $\delta = w_I\theta$ is minus the identity on Φ_I. For the Whittaker origins of Theorem 1.5, the map $\xi \mapsto \chi_\xi$ identifies the pure real forms of minimal-parabolic type I with the characters χ of $\pi_0(T^\delta)$ satisfying*

$$\chi(\alpha_s^\vee(-1)) = -1 \qquad (s \in I).$$

The closed generator E_ξ is supported in positions $v \le w_I$, and in the quotient of Theorem 4.20 at w_I it is the free module in the single character χ_ξ, with its original norm marking.

Proof. Let M_0 be the quasi-split real Levi of the standard parabolic $P_{0,I}$. Choose its reference flag above a Whittaker-generic reference flag of $M_{0,\mathrm{ad}}$. Their complex flag varieties are canonically identical. The map on real points contains the identity component of the adjoint real points, by surjectivity on Lie algebras, so the chosen flag still has open real orbit. Write J_f for its real Cartan. Lemma 4.14 sends any χ to a torus cocycle $c \in H^1(\mathbb{R}, J_f)$, hence a pure label η of M_0 together with its actual open orbit. This is a bijection on pointed pairs. Indeed in the notation of that lemma the cocycle of a flag represented by g is $g^{-1} z\sigma(g)$. Passing to the adjoint quotient sends it to the projected cocycle. By functoriality of the torus pairing the adjoint orbit therefore has character $\chi|_{T_{\mathrm{sc}}[2]}$, where T_{sc} is the Cartan of the simply connected dual of $M_{0,\mathrm{ad}}$. The character is measured from the chosen Whittaker origin, as in the proof of Proposition 5.19.

On the adjoint semisimple root datum, $w_I\theta = -1$ says that the based outer action is opposition, the outer action of the compact real form. That compact form therefore exists in this inner class, and adjointness makes it a pure inner form. It is unique as an inner class of real structures because compact real structures are inner conjugate. Its real group is transitive on complex flags. Its parameter is the clean character of Lemma 5.2. For each simple root of the Levi the simply connected root SL_2 has its two root spaces interchanged by the opposite odd element. The fixed torus in this SL_2 acts on the open $\mathbb{G}_m$ of its flag line with stabilizer $\{1, -1\}$. A meridian therefore lifts to a path ending at -1, whose image in $T_{\mathrm{sc}}[2]$ is $\alpha_s^\vee(-1)$. Cleanness forces its monodromy to be nontrivial, by the punctured-disc calculation. These simple coroot elements form a basis of $T_{\mathrm{sc}}[2]$. Thus the compact adjoint character is exactly the unique character taking all these values to -1. The condition in the statement means precisely that the projected pointed orbit is this compact adjoint orbit. A finite central isogeny preserves compactness of the semisimple real Lie algebra, so this is equivalent to compactness of the derived group of M_η.

Map c through $M_0 \subset P_{0,I} \subset H_0$, obtaining ξ and a real parabolic of type I with Levi M_η. Any proper real parabolic inside this parabolic would induce a proper real parabolic in the Levi. Its derived group is compact, so none exists. Hence $I(\xi) = I$. Conversely minimal real parabolics in the given twist are conjugate under its full real group, and $H^1(\mathbb{R}, P_{0,I}) = H^1(\mathbb{R}, M_0)$ by unipotent descent. Their Levi label is thus unique, and its compact derived group has one orbit on complex flags. This proves both surjectivity and injectivity, including

the empty-root case. Any other reference above the same adjoint Whittaker orbit changes the origin by a torus class trivial after adjoint projection; it permutes the displayed set. The same chosen reference is used in E_ξ and in the displayed bijection.

Finally $E_\xi = k_* q^* e_{\chi_\xi}$ is supported where the standard partial flag of type I is preserved by the odd element. There the two full flags have relative position in W_I, hence at most w_I. On the top open of this stable-parabolic locus the two Levi flags are opposite. Its pullback atlas is U times the opposite Levi incidence atlas, by Lemma 3.33; in particular it is reduced there. Its stabilizer presentation is precisely $[Ja_{w_I}/J] = S_{w_I}$ with $J = T \ltimes U$, since the two opposite Levi Borels meet in T. The Levi generator restricts to the single line of character χ_ξ. The two torus coordinates agree: $c_0 = r_G r_L^{-1}$ is central in L, $\tau_L = c_0^{-1}\tau_G$, and $r_L^{-1} n_{w_I} \tau_L = r_G^{-1} n_{w_I} \tau_G$. Using $n_{w_I}\tau_L$ in the Levi instead changes this representative by a torus coboundary, since r_L belongs to the derived Levi torus on which $\delta = -1$. Conjugation by its torus square root fixes the stabilizer character. After removing the lower positions the direct image is therefore $i_* p_{w_I}^* \mathcal{O}_{Y_\delta, \chi_\xi}$. Lemma 4.19 and Theorem 4.20 read its intrinsic coefficient as the stated single free module. All maps retain the original odd square and norm. □

Theorem 6.6. *The full spectral category $\mathcal{T}(V)$ is compactly generated by the finite family $\Delta_w \cdot E_\xi$, with $w \in W$ and $\xi \in H^1(\mathbb{R}, H_0)$, over the entire marked Cartan.*

Proof. Let $\mathcal{L}$ be the localizing subcategory generated by the compact Hecke translates of the E_ξ. The universal standards Δ_w thickly generate the compact finite Hecke kernels by the Bruhat filtration. Their translates therefore suffice, and all these objects are compact.

Proceed through the filtration of Theorem 4.20. Assume that the preceding term $\mathcal{D}_-$ belongs to $\mathcal{L}$, and let $\mathcal{D}$ denote the term at w. The subcategory $(\mathcal{D} \cap \mathcal{L})/\mathcal{D}_-$ is localizing and full in $\mathcal{D}/\mathcal{D}_-$: localization removes a term in $\mathcal{D}_-$ by the right-orthogonal projection, so it preserves $\mathcal{D} \cap \mathcal{L}$. For any left descent s of w, the minimum v in Lemma 6.4 precedes w. Hence all its character lines have compact lifts in $\mathcal{L}$ by the induction hypothesis. Applying D_s to the specified lifts places the factors computed in that lemma in the image of $\mathcal{D} \cap \mathcal{L}$.

It remains to treat a character missed by every such descent. Let I be the set of left descents of w, and $\delta = w\theta$. For each $s \in I$ the panel lemma implies both $\delta\alpha_s = -\alpha_s$ and $\chi(\alpha_s^\vee(-1)) = -1$. For $s \notin I$ the root $\delta\alpha_s$ is positive and outside the span of I, since δ is minus the identity on that span. It stays positive under w_I. On the span of I, $w_I\delta$ carries each simple root to a positive root as well. Thus $w_I\delta$ preserves the positive system. It has the same outer action as θ, so $w = w_I$. The equality $w\theta(w) = 1$ then gives $\theta I = I$. Proposition 6.5 supplies a compact E_ξ supported in this filtration term whose image in the quotient is the free module in the remaining character. This includes the starting position $I = \varnothing$.

All factors of the quotient thus belong to the image. Since $\mathcal{D}_- \subset \mathcal{L}$, the localization triangle implies $\mathcal{D} \subset \mathcal{L}$. Induction and Theorem 4.17 prove the assertion. □

Corollary 6.7. *Suppose every pure inner twist of H_0 is either quasi-split or compact modulo its center. Then the two full categories in (1.1) are equivalent over the entire marked Cartan, with the prescribed period. More generally fix a character of $Z(G)^\theta$. If every pure label of that character is quasi-split or compact modulo its center, the same comparison holds between the sum over those labels and the entire spectral category of that character. Without this hypothesis it still holds over the complement of the finite union*

$$\bigcup_{\substack{\xi \text{ of the given character} \\ H_\xi \text{ neither quasi-split nor compact modulo center}}} W(T^{w_{I(\xi)}\theta})^\circ.$$

Proof. For the quasi-split twists take J from Theorem 1.2. Its closed generators agree with the E_ξ for $I = \varnothing$. Indeed the parabolic embedding for the torus Levi sends the trivial coefficient line to Q. On the torus the global twist L_ε has weight ε on $\pi_0(T^\theta)$: its restriction to the fixed center determines this character by Lemma 4.2, and its construction gives trivial weight on the connected fixed torus. The remaining line on Y_θ is trivializable. The projection formula identifies the two closed objects. For a compact-modulo-center twist not already counted the entire automorphic category is its closed-orbit category. Theorem 1.5 realizes it as an orthogonal factor generated by E_ξ. Corollary 5.24 gives orthogonality also with all Hecke translates of the other forms and shows that its own factor contains its Hecke translates. The direct sum of these functors is fully faithful. Its image contains the generators of Theorem 6.6, which proves essential surjectivity. For a single character apply the fixed-center weight projection to that generation theorem. Proposition 5.13 retains exactly the generators of the specified labels. The same functors on these summands then give the whole character category, including the zero category if the set of labels is empty.

For the open assertion, Proposition 4.6 places each excluded automorphic summand on its indicated Weyl union. On the spectral side the ideal of $(T^{w_I\theta})^\circ$ acts by zero on E_ξ through its norm coefficient ring. Every function on the Weyl quotient vanishing on this support therefore kills every Hecke translate as well, since the even semisimple class is unchanged on the two projections of the Hecke correspondence. The image of the support in the finite Weyl quotient is closed, and its inverse image is precisely the displayed Weyl union for this label. On the claimed open all these translates thus vanish. Localizing the generators of Theorem 6.6 and the orthogonal embeddings proves the comparison there. □

6.2. **Duality of the closed generators.** Write p_D for the forgetful map to partial flags of type D. The second partial flag in the pair description has type θD. We use $A_I = \mathcal{O}(Y_{w_I\theta})$, with its specified norm map, and the original periodic coefficient ring P.

Lemma 6.8. *Coherent duality takes each compact closed generator E_ξ to an isomorphic object. It also fixes $p_{D*}E_\xi$ up to isomorphism.*

Proof. Take duality with the dualizing complex on each stack. The ambient derived stacks are derived incidence loci of smooth groups and flags, hence are quasi-smooth and Gorenstein; bounded coherent duality is involutive. It preserves the small circle structures: duality commutes with pullback by isomorphisms, with the inverse transpose on descent arrows. On a full turn this is the action of the same gauge on the dual representation. Duality exchanges source and target in the invariant mapping complexes and descends through inversion of the boundary class and passage to compact retracts. We require this duality only for compact objects. Proper direct image commutes with bounded coherent duality, so it suffices to treat the Levi object and its pullback by the stable-parabolic map q.

Let j be the opposite Levi open and $e = e_\chi$ the compact object in Proposition 5.19. For a compact test X, duality, its two extension adjunctions and open restriction identify

$$\mathrm{RHom}(X, \mathbb{D}e) \simeq \mathrm{RHom}(e, \mathbb{D}X) \simeq \mathrm{RHom}(j^*X, \mathbb{D}(j^*e)).$$

On the smooth torus quotient $[Ta_{w_I}/T]$ the dualizing line is trivial and its virtual dimension is zero. Its torus acts trivially on the two tangent determinants. The open line is a character of the component group, which has exponent two, so its dual is itself. The right extension adjunction gives $\mathbb{D}e \simeq e$. On the atlas of Lemma 3.33, q is smooth in the Artin sense, with smooth coordinate group U and gauge kernel U. Its relative tangent complex has these two Lie algebra bundles in successive degrees. Under B_L the determinants both have character $\det(\mathrm{Ad}\,|_{\mathfrak{u}})$; unipotent conjugation and translations preserve the invariant volume form. Thus the relative dualizing line is trivial and the virtual dimension is zero. The boundary gauge acts on the two determinants by this same character, so cancellation respects the original turn. Consequently duality commutes with q^* here. The proper maps k and p_D now give the claims. This uses the two extension adjunctions only in the minimal Levi. □

6.3. **Orthogonality for distinct minimal real parabolic types.** We now compare arbitrary kernels between closed generators of different minimal types. Three relative positions enter the convolution: those of an input and an output flag with their odd translates, and that of the two flags with each other.

Lemma 6.9. *In the spectral finite Hecke category the subcategory supported on any lower Bruhat ideal of flag pairs is generated by the standards indexed by that ideal. At its maximal position d the quotient by the remaining positions has a compact free generator given by the structure sheaf of the reduced position closure. The telescope of its twists by the line cutting out the open position represents its image localized away from those lower positions.*

Proof. Use the ordinary Steinberg loop presentation of Lemma A.2. It is the exchange-group case of Theorems 4.17 and 4.20. In the gauge slice of that lemma the two positions are d, d^{-1}, determined by one another. The torus stabilizer after passing to the Levi is the connected diagonal torus, so only its trivial character survives. Its coefficient ring is the full Cartan ring. We may choose the trivial homogeneous line in Theorem 4.17. The localization telescope and its boundary cones are those in Theorem 4.20.

To check the assertion specifically for standards, start with the unit on the diagonal. Suppose that multiplying by a simple reflection on the output end increases the length, and forget that simple output flag while retaining the input flag. Over the resulting partial pair, the lower full position has a unique lift: at its shortest representative the extra negative simple root of the partial parabolic is absent from the input Borel, so the intersection stabilizer is unchanged. The lower pair orbit therefore maps isomorphically to the partial pair orbit. Its ordinary loop stack has the same property. The inverse image of this partial orbit on pairs is smooth over it and has the higher position as its open orbit. The derived Cartesian square of loop stacks, as in Lemma 6.4, computes the restriction of simple averaging to the new position as pullback of the previous free line. On the torus coefficient

this pullback is an isomorphism of markings, since both flags share the same semisimple element. The identity term of the simple counit triangle is zero there by support. Induction along a reduced word thus puts the standard on the closed Bruhat bound of its index and gives a free line in its highest quotient. If the input-first convention transposes the two flags, take the inverse index throughout. Induction on the lower ideal now proves the generation assertion. □

We record the support calculation for the three open positions. Write Z_G for the Steinberg stack, and $\mathcal{W}$ for its correspondence with the parameter stack as in (B.4). Its projections to the kernel, input and output are denoted by p, r, q, respectively. Coherent duality on a parameter or kernel stack is denoted by $\mathbb{D}$.

Lemma 6.10. *In the kernel, input and output variables take closed direct images of bounded coherent equivariant sheaves from closed substacks. In each variable twist by powers of a line and form a telescope by a section of that line. Assume these are the coherent heart telescopes of Lemma 6.9 and Proposition 5.12 for the kernel and the two closed generators. If no classical point of $\mathcal{W}$ projects to the three simultaneous nonvanishing loci, the three telescopes pair to zero under the colimit-preserving extension of*

$$(K, X, Y) \longmapsto \mathrm{RHom}\big((\mathbb{D}K)\cdot(\mathbb{D}X), Y\big)$$

from compact periodic objects.

Proof. Duality in each source makes this functor covariant and exact in each of the three compact variables. Extend it in each variable by filtered colimits. Write Z_c, V_i, V_o for the classical closed supports of the three telescopes, with their respective coefficient lines F_c, F_i, F_o and section lines L_c, L_i, L_o. The first two duals are closed direct images of the coherent duals on their supports, twisted by the inverse powers of the section lines. In the notation of (B.4) form

$$C = Z_c \times^R_{[B/B]} V_i, \qquad T' = C \times^R_V V_o.$$

The map $g : T' \to V_o$ is proper since all the support inclusions are closed and forgetting a flag is proper. Write W' for the pullback to T' of the exterior product of the untwisted duals on Z_c, V_i. The diagonal of the three index sets is cofinal. At index n, before circle invariants, the complex is

$$\mathrm{RHom}_{V_o}\big(g_*\big(W' \otimes (L_c \otimes L_i \otimes L_o)^{-n}\big), F_o\big)\,,$$

where the lines in the direct image denote their pullbacks to T'. Indeed composition with the two closed inclusions first computes the convolution by direct image from C. Adjunction with the output closed inclusion then pulls the source back to V_o, giving T' by proper derived base change. The projection formula moves the output twist to the source. These identities hold in quasi-coherent complexes and retain the full derived intersection. The diagonal transition in Hom is precomposition with multiplication by $h = s_c s_i s_o$ on the source, from its $-(n+1)$-st twist to its $-n$-th twist.

The open where h is invertible in T' is empty. Put $L = L_c \otimes L_i \otimes L_o$ on T'. The direct limit of $\mathcal{O}_{T'} \to L \to L^2 \to \cdots$ is zero: on a derived affine chart it is localization at a function nilpotent in the classical ring. The unit $\mathcal{O}_{T'}$ is compact in quasi-coherent complexes on T'. Indeed this stack has a quasi-compact separated derived scheme atlas modulo the reductive gauge group. A finite affine cover with affine intersections computes sections on the scheme, and rational invariants are exact and commute with filtered colimits. It follows by mapping the unit to the displayed direct system that some $h^N : \mathcal{O}_{T'} \to L^N$ is zero in the derived category. Its map after tensoring with W' and any line and taking direct image is zero. The underlying diagonal system of Hom complexes thus has zero colimit.

For the circle calculation, the two untwisted duals on the first supports are bounded, also after any line twist. Their external tensor over the smooth stack $[B/B]$ has a uniform upper bound independent of the index by Proposition B.12. Proper projection along the flag has bounded cohomological dimension. The original output target is in the heart, so the original Hom complexes have a common lower bound. Finite skeleta of BS^1 therefore commute this filtered colimit with invariants degree by degree, and inverting u preserves zero. This proves the claim. □

Lemma 6.11. *For a position d, the coherent dual of the standard finite kernel belongs to the localizing subcategory generated by the telescope of the free line on its exact position. Thus vanishing of the pairing in Lemma 6.10 on this telescope implies vanishing of $\mathrm{RHom}(\Delta_d E_\xi, E_\eta)$ for the specified pair of labels, without any hypothesis on returns for shorter kernels.*

Proof. Let $\mathcal{H}_{\le d}$ and $\mathcal{H}_{<d}$ be the full supported categories for the indicated Bruhat ideals. On the enhanced flag space Δ_d is extension by zero from the open cell of its support. Hence $\mathrm{RHom}(\Delta_d, C) = 0$ for every C

supported strictly below d, by open restriction and adjunction inside the closed bound. Under the complex-group equivalence of Appendix A this applies to all compact spectral kernels supported there: by Lemma 6.9 they are generated by the images of the shorter standards. Coherent duality preserves the classical support, is involutive on compact periodic kernels, and reverses the two entries of Hom. Thus $\mathbb{D}\Delta_d$ belongs to $\mathcal{H}_{\le d} \cap (\mathcal{H}_{<d})^{\perp}$. Here compact tests suffice because mapping out of a colimit gives a limit.

This intersection is the image of the right adjoint to $\mathcal{H}_{\le d} \to \mathcal{H}_{\le d}/\mathcal{H}_{<d}$. The right adjoint is continuous: the quotient sends compact generators to compacts, so maps from those generators test its compatibility with sums. The quotient has one compact free generator by Lemma 6.9. Its image under the right adjoint is the section telescope by Theorem 4.20, applied to the exchange group. Consequently the telescope generates the entire intersection. With the other two arguments fixed, the pairing of Lemma 6.10 preserves colimits in the kernel variable. It therefore vanishes on $\mathbb{D}\Delta_d$ whenever it vanishes on the telescope. Finally $\mathbb{D}E_\xi \simeq E_\xi$ by Lemma 6.8. □

Theorem 6.12. *Let $I(\xi) \neq I(\eta)$. For any two compact finite Hecke kernels K, K',*

$$\mathrm{RHom}(K \cdot E_\xi, K' \cdot E_\eta) = 0.$$

Consequently the entire spectral category is the orthogonal product of its Hecke subcategories generated by the closed objects of each fixed minimal real parabolic type. For any two labels of the same type I, including equal labels, one has

$$\mathrm{RHom}(\Delta_w E_\xi, E_\eta) = 0 \qquad (w \notin W^{w_I\theta}).$$

The factor for the empty type is equivalent to the sum of the full automorphic categories of all quasi-split twists over the entire marked Cartan.

Proof. First take $K = \Delta_w$ and $K' = \mathbf{1}$. Let d be shortest in $W_J w W_I$. Corollary 5.24 reduces its coefficient to the one for d, up to shift. We determine what the existence of a triple of exact positions implies.

In input-first notation write the kernel position from output to input as d, with output type J and input type I. Equivalently reverse the flags and use d^{-1}. Suppose a triple exists. The output and input parabolics P_J, P_I are preserved by the odd element a, and inside each of their Levis the chosen flag and its odd translate are opposite. The partial position is therefore unchanged by θ. Since I, J are stable and d is shortest, uniqueness of the shortest representative gives $\theta(d) = d$. Fixing the pair of full flags as $(B, {}^dB)$ identifies the parabolic pair as $(P_J, {}^dP_I)$. The image of the intersection of these parabolics in the output Levi contains its chosen Borel: every positive root of that Levi belongs to ${}^dB$, by left minimality. This image is a parabolic subgroup since it contains a Borel. It is preserved by the action of a on the Levi quotient. It therefore contains both opposite Borel subgroups and equals the entire Levi. A negative root from Φ_J can then occur in the intersection only if its image under d^{-1} belongs to Φ_I: the corresponding positive root was positive for ${}^dB$. Thus $d^{-1}\Phi_J \subseteq \Phi_I$. The same argument at the input, by right minimality, gives the reverse inclusion. Minimality also preserves positive systems, so d sends the simple system I to J.

Apply this root-datum conclusion on the real side to the quasi-split maximally split Cartan. Since $d \in W^\theta$, it has a representative in its real normalizer by the real Weyl group calculation used in Proposition 5.10. Conjugation by this representative carries the real standard Levi of type I to that of type J. Transport the compact-derived Levi cocycle of ξ by this real conjugation. Its image in $H^1(\mathbb{R}, H_0)$ is still ξ. The transported Levi is again compact modulo its center, so twisting the standard parabolic of type J yields a minimal real parabolic in the same twist. All its minimal real parabolics are conjugate by real points and therefore have the same absolute type. Thus an exact triple requires $I = J$. For equal types the preceding root argument requires that d preserve Φ_I^+ and commute with θ. Such a d commutes with $w_I\theta$. Every element of W_I also commutes with $w_I\theta$, since this involution acts by minus the identity on Φ_I. Hence $w \in W^{w_I\theta}$ is necessary.

For distinct types, or for equal types and w outside this centralizer, the three exact positions have empty intersection. Lemma 6.10 gives zero on their telescopes, and Lemma 6.11 gives zero for the actual standard return. This proves the assertion for all standard kernels, including the stated self-returns.

For distinct types it remains to treat the second kernel. Standards generate the compact kernels. Their inverses and their composition law allow the second kernel in the asserted Hom to be moved to the first variable, first for standards and then by finite cones and retracts for any compact kernels. This proves both directions of orthogonality. Each family consists of compact objects, so orthogonality passes to the localizing subcategories. Their finite sum is the entire category by Theorem 6.6. Finally the empty type consists of the quasi-split labels, and its closed generators are those of the fully faithful functor J, as proved in Corollary 6.7. This identifies its factor. □

Put $N_I^+ = \{d \in W^\theta : d\Phi_I^+ = \Phi_I^+\}$. This group is the quotient of $W^{w_I\theta}$ by W_I. Indeed an element commuting with the involution preserves its fixed torus and hence the roots Φ_I vanishing identically there. Its unique adjustment by W_I preserving Φ_I^+ commutes with w_I and with θ. Conversely these positive-system-preserving elements and W_I commute with $w_I\theta$.

Lemma 6.13. *Let d be shortest in dW_I. Over a point of D_I at which the roots outside Φ_I have value different from one, the standard kernel for d acts on the stable Levi chart by replacing the flag B with ${}^{d^{-1}}B$. We use the convention taking the input marking to its d-translate. If two such representatives d, d_0 give the same translate of D_I, the two identifications differ by an automorphism of the Levi data which permutes the signed factors.*

Proof. For a reduced word of d applied from right to left, the roots tested at the initial mark by the successive steps are its positive inversion roots. None belongs to Φ_I, by minimality of d. Thus each step tests a root outside the Levi at the initial marking.

On this root-regular locus the projection forgetting the simple flag is a cover with two flags. In the rank-one Levi of the partial flag the even element has two distinct eigenlines, and their fixed-point equations are étale. In the fiber product of the cover with itself the diagonal and the other graph are open and closed. The counit of $q_s^* q_{s*}$ projects to the diagonal, so the standard kernel acts by the other graph. This argument works over the original open locus, including derived base change from the remaining flag and square conditions.

On the stable Levi chart of Theorem 5.15 fix the Levi Borel $B_L = B \cap L_I$, so the odd element belongs to $L_I\tau$ and its square to B_L. The graph corresponding to a shortest representative d replaces the ambient Borel B by ${}^{d^{-1}}B$ while keeping this same Levi Borel. Indeed at every suffix v of the reduced word, $v\Phi_I^+$ is positive. The Borel ${}^{v^{-1}}B$ therefore contains B_L, so contains the even element. Two successive such flags differ in precisely the indicated simple panel. The other étale graph just described identifies their derived families. This description is equivariant under B_L and hence descends from the chart. Moving the new flag back to the standard one by a representative of d gives the translated Cartan marking.

The roots trivial identically on D_I are exactly Φ_I: the involution is minus the identity on the Levi roots, while the fixed cocharacter positive on the radical in Proposition 5.11 separates all other roots. The images of Φ_I^+ under d, d_0 are therefore the same positive system of this common root subsystem. Thus $n = d_0^{-1}d$ preserves Φ_I^+. It also commutes with δ_I, since it preserves its fixed rational subspace and its orthogonal complement under an invariant positive form. It fixes w_I by conjugation, so commutes with θ. A representative of n in the torus normalizer consequently preserves B_L, L_I and the odd coset $L_I\tau$.

On the two Levi charts the transport by d and by d_0 thus differ by conjugation by this representative on the Levi data. This follows also by comparing their final flags ${}^{d^{-1}}B$ and ${}^{d_0^{-1}}B$ in the preceding formula. The conjugation preserves the Levi opposite open and permutes its simple roots. On the torus stabilizers it therefore preserves the set of characters taking value -1 at every simple coroot at -1. Its possible change of odd base lift translates the norm cover and preserves the free module of a character factor. The right extensions of these factors, characterized by adjunction, are permuted as a collection. By Proposition 6.5 they are precisely the factors of the type-I seeds. This proves both assertions. □

6.4. **The pure labels of a fixed minimal type.** Let $\Lambda = X_*(T)$, let $Q^\vee \subset \Lambda$ be the coroot lattice, and put $F = \Lambda/Q^\vee$. We use additive notation for these lattices and their quotients. For an abelian group M with involution γ, write

$$M_\gamma = M/(1-\gamma)M, \qquad \widehat{H}_\gamma^{-1}(M) = \ker(1+\gamma : M \to M)/(1-\gamma)M, \qquad \widehat{H}_\gamma^0(M) = M^\gamma/(1+\gamma)M.$$

The first Tate group embeds in M_γ. For a minimal type I set $\delta = w_I\theta$, and let $Q_I^\vee$ be the lattice of the simple coroots in I. Write $\mathcal{X}_I$ for the set of pure labels of this type, using the Whittaker origins of Proposition 6.5.

Proposition 6.14. *Suppose I occurs. The set $\mathcal{X}_I$ is a torsor for the character group of*

$$C_I = \frac{\ker(1+\delta : \Lambda \to \Lambda)}{(1-\delta)\Lambda + Q_I^\vee} = \ker\left(\widehat{H}_\theta^{-1}(F) \xrightarrow{\partial_I} \widehat{H}_\delta^0(Q^\vee)\right).$$

For $[f] \in \widehat{H}_\theta^{-1}(F)$, the arrow sends a lift $x \in \Lambda$ to the class of $x + \delta x$. In particular, for semisimple adjoint H_0 there is at most one pure label of any given minimal type.

Proof. The map $x \mapsto x(-1)$ identifies $\widehat{H}_\delta^{-1}(\Lambda)$ with $\pi_0(T^\delta)$. Indeed on the maximal compact torus write a fixed point as $\exp(2\pi i v)$. The class of $x = (1-\delta)v \in \Lambda$ is independent of the logarithm. Dividing the point by $x(-1)$ leaves the exponential of the fixed vector $(v+\delta v)/2$, in the identity component. Changing x by $(1-\delta)\Lambda$ leaves its component unchanged. This also identifies the kernel, and the positive real part of the complex torus contributes

no components. Proposition 6.5 now identifies the quotient of any two label characters with a character of C_I, and multiplication by each such character preserves the prescribed signs.

The involution on F is θ, since the Weyl group acts trivially there. Modulo $Q_I^\vee$, the involution δ permutes the basis of simple coroots outside I. An anti-invariant in this quotient is a sum of differences of exchanged basis vectors. Subtracting $(1-\delta)q$ for a lift $q \in Q^\vee$ thus carries each anti-invariant of $Q^\vee$ into $Q_I^\vee$, and every element of the latter is anti-invariant. If the image of an anti-invariant $x \in \Lambda$ in F is a difference, subtract a difference in Λ and apply this argument in $Q^\vee$. This proves the asserted injection of C_I. The class of $(1+\delta)x$ for any lift of an anti-invariant in F is well defined modulo $(1+\delta)Q^\vee$, including changing the representative by a difference. Its vanishing says exactly that the lift can be modified by a coroot lattice element to become anti-invariant. This gives the kernel formula.

For semisimple adjoint H_0 its dual G is simply connected, so $F = 0$. □

6.5. **The return of an arbitrary standard kernel.** Fix two pure labels ξ, η of the same minimal type I. We first compute the return for a representative $d \in W^\theta$ satisfying $d\Phi_I^+ = \Phi_I^+$. Write the two flags, in order, as B and $B' = {}^dB$. Their parabolics $P = P_I$ and $P' = {}^dP_I$ have the same Levi $L = L_I$. The two Borels induce the same Borel B_L of L. Put $Q = P \cap P' = L \ltimes U$ and $J = B \cap B' = B_L \ltimes U$. Superscript $-$ on a stable parabolic or on Q denotes its coset in the odd component E^-, using the pinned lift τ. We shall use w_I for the longest element of W_I and $\delta = w_I\theta$.

There are three open conditions in the return calculation. The pair of flags must have relative position d. At each endpoint the odd element must preserve the prescribed parabolic and take its Levi flag to an opposite flag. We begin by identifying the derived space on which all three conditions hold. We then calculate which complex on that space gives the return.

Lemma 6.15. *Let T^* denote the simultaneous open in the derived intersection of Lemma 6.10. Before quotienting by the stabilizer of (B, B'), form the derived Cartesian square*

$$\begin{array}{ccc} X & \longrightarrow & J \\ \downarrow & & \downarrow{\scriptstyle \iota} \\ P^- \times^R_{E^-} (P')^- & \xrightarrow{\text{sq}} & P \times^R_G P'. \end{array} \tag{6.1}$$

The bottom arrow squares the odd element at each endpoint. The right arrow sends $b \in J$ to its two inclusions into P and P' and uses their common image in G. On the classical truncation of X let X^{op} denote the condition that the projection a_L of the common odd element to the odd Levi takes B_L to an opposite Borel. We use the same superscript for the corresponding derived open. Then

$$T^* \simeq [X^{\mathrm{op}}/J] = \left[\left(P^- \times^R_{E^-} (P')^- \right) \times^R_{P \times^R_G P'} J \Big/ J \right]^{\mathrm{op}}. \tag{6.2}$$

The J in the fiber product parametrizes the common even element. The J in the quotient is the stabilizer of both flags in G and acts by simultaneous conjugation.

Proof. On the orbit of the pair (B, B') a loop is an even element of its stabilizer J. Thus the kernel chart is $[J/J]$. At the B endpoint the supporting stack, before gauge quotient, is $P^- \times^R_P B$, with the open condition on the Levi flag. Gluing the even element of the kernel to this endpoint gives $P^- \times^R_P J$. At the B' endpoint the ambient chart is $E^- \times^R_G B'$, and its supporting chart is $(P')^- \times^R_{P'} B'$. Taking their prescribed fiber product identifies the two odd transports in E^-, and identifies their squares with the same even element of J. Pasting these derived pullbacks gives (6.1). The argument is symmetric in the two flags.

Both Levi projections of the classical common transport agree on Q^-, and the two flags project to B_L. Thus the two opposition conditions reduce to the one in the statement. To see why using the reduced supporting stacks in Lemma 6.10 gives these endpoint charts on their open strata, consider the equation $a_L^2 \in B_L$. In right-translation coordinates the derivative of squaring is $X \mapsto X + \mathrm{Ad}(a_L)X$. For $X \in \mathfrak{b}_L$ its image modulo $\mathfrak{b}_L$ spans $\mathfrak{l}/\mathfrak{b}_L$: the Lie algebra $\mathrm{Ad}(a_L)\mathfrak{b}_L$ is opposite to $\mathfrak{b}_L$. The equation is therefore transverse on this open. The unipotent radical of an endpoint parabolic imposes no additional equation, since it is contained in its Borel. The endpoint chart there is smooth, as is $[J/J]$. The remaining intersections between these three smooth charts are exactly those retained in (6.2). □

Lemma 6.16. *Let $A^\bullet$ be a connective commutative differential graded algebra with noetherian H^0 and finite cohomology over H^0 in each degree. Suppose that M is a bounded-above pseudo-coherent module, N a bounded-below module, and $h \in H^0(A^\bullet)$. There is a natural equivalence*

$$\operatorname*{colim}_h \mathrm{RHom}_{A^\bullet}(M, N) \simeq \mathrm{RHom}_{A^\bullet[h^{-1}]}(M[h^{-1}], N[h^{-1}]).$$

On the left the transition acts on N by h. The same calculation commutes localization with the right adjoint to an affine closed direct image when its structure algebra is pseudo-coherent as a module over the ambient algebra and the target is bounded below.

Proof. Choose b such that $H^i(N) = 0$ for $i < b$. For any given degree j, choose a finite approximation $F \to M$ by finite free cells whose cone has cohomology in degrees at most $b - j - 2$. Such approximations exist by pseudo-coherence. Maps from this cone to $N[j]$ and the adjacent shifts needed for H^j vanish by the cohomological bounds. Localization at h preserves these bounds, so the same approximation computes the indicated degree on the open. The assertion for F follows by finite cones and direct sums from the case $F = A^\bullet$. Filtered colimits of vector spaces are exact, proving the formula in every degree. For an affine closed immersion with algebra B over $A^\bullet$, restriction of scalars has right adjoint $N \mapsto \mathrm{RHom}_{A^\bullet}(B, N)$ with its natural B-module action. Apply the first assertion with $M = B$. □

Lemma 6.17. *Let S_o be the closed output support in Lemma 6.10, S_o^* its smooth top, and $g : T' \to S_o$ the proper projection from the derived triple intersection there. On the three smooth open strata denote the kernel, input and output lines by F_c, F_i, F_o. Denote by W the tensor product on T^* of the pullbacks of the coherent duals of F_c and F_i. Put $D = W^\vee \otimes g^*(F_o \otimes \omega_{S_o^*}^{-1})$, using g also for the restriction to T^*. Before circle invariants and periodic localization, the coefficient of the three telescopes is*

$$R\Gamma(T^*, \omega_{T^*} \otimes D). \tag{6.3}$$

Proof. Denote by W' the tensor product of the pullbacks of the untwisted duals on the closed kernel and input supports. Write F_o' for the closed output term. On T' let $\mathscr{L}$ be the tensor product of the three section lines, and h the product of their sections. At diagonal index n, Lemma 6.10 gives the two expressions

$$\mathrm{RHom}_{S_o}\big(g_*(W' \otimes \mathscr{L}^{-n}), F_o'\big) \simeq \mathrm{RHom}_{T'}\big(W', g^! F_o' \otimes \mathscr{L}^n\big). \tag{6.4}$$

Here duality is applied on the closed support, where g is proper. To identify the right adjoint, factor g through the flag bundle over S_o parametrizing the additional flag. The first map is a derived closed immersion: on truncations it imposes the closed flag incidence and support conditions. The flag-bundle projection is smooth and proper. Its right adjoint on quasi-coherent complexes is pullback tensored with its relative canonical bundle and shifted by its relative dimension. For the closed immersion use the coinduction in Lemma 6.16. This describes $g^! F_o'$ on smooth charts. It is bounded below, since F_o' is in the coherent heart, the flag shift is fixed and the structure algebra for the closed immersion is connective. The algebras on these charts have finite cohomology in each degree over their classical rings, by their presentations as iterated derived intersections of finite-type schemes. In particular the algebra of the derived closed immersion is pseudo-coherent over the ambient chart: its truncation is a closed immersion and each cohomology module is finite. The duals used to define W' are bounded coherent on their original supports. Their pullbacks and tensor product are bounded above and pseudo-coherent.

Trivialize $\mathscr{L}$ on an affine chart. The transition on the right in (6.4) is multiplication by h, so Lemma 6.16 computes its colimit by restriction to $h \neq 0$. The calculation glues on overlaps. Indeed the charts come from quasi-compact separated schemes with a reductive gauge group; a finite affine cover with affine intersections computes scheme cohomology, and rational reductive invariants are exact and commute with filtered colimits. On S_o^* the target is the line F_o. The local description of $g^!$ commutes with open restriction by Lemma 6.16. On T^* it is therefore $\omega_{T^*} \otimes g^*(F_o \otimes \omega_{S_o^*}^{-1})$. The source restricts to W, which is a line since the three open supports have smooth charts of stack dimension zero. This proves (6.3). Here the dualizing object on the possibly unbounded derived intersection means the smooth local dualizing complex computed by coinduction, with the relative dualizing line and shift removed on passage from an atlas to its stack.

The sections and the adjunction maps in (6.4) respect boundary rotation. The complexes on the left have a common lower bound by Lemma 6.10. To compute circle invariants in a fixed degree it suffices to use a finite skeleton of BS^1: the successive skeletal contributions shift the common lower bound upward by the cell dimension. Thus the colimit computation is valid before passing to these invariants and subsequently inverting u. □

Choose a θ-fixed defining cocharacter λ for P, zero on the roots of L and strictly positive on the roots of its unipotent radical. Put $\nu = \lambda + d\lambda$. Both summands are central in L and fixed by δ, so ν takes values in $(T^\delta)^\circ$. Let $M = \mathfrak{g}/(\mathfrak{p} + \mathfrak{p}')$. The roots of M pair negatively with both λ and $d\lambda$; those of U pair positively with both.

Lemma 6.18. *Set $J' = T \ltimes U$ and $a_I = r_G^{-1} n_{w_I} \tau$, where $r_G = (2\rho_G^\vee)(i)$. The classical truncation T_0^* of T^* is $[X_0/J']$ with $X_0 = J' a_I$. The action is conjugation. The dualizing complex of this classical quotient is the trivial line in degree zero.*

Proof. Before the quotient in (6.2), the classical odd element lies in Q^-. Its square must lie in J, which says that its Levi square lies in B_L. By the transversality in Lemma 6.15, the ordinary base change $Q^- \times_Q^R J$ is classical on the opposite open. The unipotent radical of B_L acts simply transitively on the Borels opposite to B_L. Use this action to fix $a_L B_L a_L^{-1}$ to the standard opposite. The condition $a_L^2 \in B_L$ then says that a_L interchanges the two opposite Borels. Their common stabilizer is T, so the remaining odd elements in the Levi form Ta_I, by Lemma 4.15. Restoring the freely varying U coordinate gives X_0, and the remaining gauge group is J'. The slice is available on the derived chart too: projection from the first parabolic to its Levi records the opposite Borel, and taking its specified value removes the same free action. Finally right translations identify the tangent bundle of the coset $J'a_I$ with the trivial bundle with fiber $\mathrm{Lie}(J')$. Conjugation acts on that fiber by the adjoint action, exactly its action on the infinitesimal gauge group. The two determinants and their shifts cancel in the quotient dualizing complex. □

Lemma 6.19. *On the slice of Lemma 6.18, the structure algebra of the derived triple intersection has a ν-equivariant model over X_0 of the form*

$$A^\bullet = \mathcal{O}_{X_0} \otimes \mathrm{Sym}\left(M_o^*[1] \oplus M_e^*[2]\right). \tag{6.5}$$

Here M_o, M_e are the odd and even copies of M. The degree -1 generators are closed; the differential on the degree -2 generators is linear in the degree -1 generators. Its linear map is induced by the odd square in (6.1).

Proof. For a root outside Φ_I, record the signs of its pairings with λ and $d\lambda$. Roots of each of the four resulting types generate unipotent subgroups $U_{++}, U_{+-}, U_{-+}, U_{--}$. Indeed each set is closed under addition of roots and lies in a positive system. Here $U_{++} = U$ and $\mathrm{Lie}(U_{--})$ projects isomorphically onto M. Multiplication gives a chart

$$U_{--} \times U_{+-} \times U_{-+} \times Q \longrightarrow G.$$

At $(1,1,1,q)$ its differential is an isomorphism by the root-space decomposition and right translation by q. The chart is therefore étale along Q. The U_{+-} axis times Q is the chart for P, and the U_{-+} axis times Q is the one for P'. Right multiplication by τ gives the corresponding chart for the odd cosets. All these charts respect the action of ν. Thus they identify the formal neighborhoods of the indicated axes along their intersection. Equivalently their Koszul resolutions may be computed étale locally there, since the full classical support of the intersection is Q in the even case and Q^- in the odd case.

In exponential coordinates the equations cutting out the two axes have the U_{--} equations in common. The other equations eliminate the U_{+-} and U_{-+} coordinates. Each common equation gives one closed exterior generator in degree -1, so the even and odd intersections have algebras $\mathcal{O}_Q \otimes \bigwedge(M_e^*[1])$ and $\mathcal{O}_{Q^-} \otimes \bigwedge(M_o^*[1])$. These are simultaneous models along the whole classical intersection, since the product chart and its equations are defined along $\{1\}^3 \times Q$.

On these algebras the square map is ordinary squaring in degree zero. Its map on the degree -1 generators is linear, by degree. To realize this description by a map of models, lift the map from the smooth affine algebra $\mathcal{O}_Q$ along the Postnikov tower. The obstructions lie in positive Ext from its projective cotangent module and vanish. The exterior generators can then be sent to cycles representing their prescribed classes. These choices can be made ν-equivariantly by exactness of torus invariants. In right-translation coordinates the map on the excess conormals is dual to $1 + \mathrm{Ad}(a)$ on the common normal quotient at $a \in Q^-$. Indeed this is the derivative of squaring and a preserves both parabolics.

For the arrow from J the degree -1 exterior generators go to zero. Resolve this augmentation of the even exterior algebra by adjoining polynomial generators in degree -2 with differentials the exterior generators. For one pair this resolution has differential $d(z^n) = nz^{n-1}dz$ and has only the coefficient ring in cohomology, since the base field has characteristic zero. Tensoring the resolutions gives the assertion for all generators. Base change to the odd algebra replaces each differential by its image under the square map. The ordinary base change gives $Q^- \times_Q^R J$, already classical on the opposite open by Lemma 6.18. Restricting to the slice gives (6.5). This construction uses equivariance of the model only for ν. □

Lemma 6.20. *Suppose that the top kernel line is the structure sheaf and the two endpoint lines have finite stabilizer characters trivial on $(T^\delta)^\circ$. For the coefficient line D of Lemma 6.17, the trace of $j : T_0^* \to T^*$ gives an equivalence*

$$R\Gamma(T_0^*, \omega_{T_0^*} \otimes j^* D) \longrightarrow R\Gamma(T^*, \omega_{T^*} \otimes D). \tag{6.6}$$

It respects the boundary circle.

Proof. On the smooth affine chart X_0 take the algebra $A^\bullet$ of (6.5). Forgetting its module structure has right adjoint coinduction, so the local dualizing complex of the derived chart is

$$R\mathcal{H}om_{X_0}(A^\bullet, \omega_{X_0}).$$

Indeed the chart projection to X_0 is affine, with identity classical truncation, and its composition with the smooth structure map of X_0 gives the dualizing complex computed in Lemma 6.17. Equivalently one may compute against bounded Postnikov approximations degree by degree. Each term of $A^\bullet$ is a finite free $\mathcal{O}_{X_0}$-module: only finitely many exterior and polynomial degrees have any given total degree. Thus its dual is computed by the ordinary dual complex. For descent from X_0 one removes the relative dualizing line and shift of the J'-torsor, as in Lemma 6.18.

Under coinduction the trace is the dual of the augmentation $A^\bullet \to \mathcal{O}_{X_0}$. In the dual complex its cone is given by terms containing at least one factor from M_o or M_e. All such monomials have strictly negative ν-weight. Indeed every root of M pairs negatively with ν. The coefficient functions on X_0 add only nonpositive weights: the T coordinate is fixed by ν, and the coordinates on U have strictly negative weights. The endpoint characters and the kernel line have weight zero by hypothesis. The determinant cancellation on each smooth top, proved by the same tangent computation as in Lemma 6.18, accounts for the dualizing lines in D. Consequently all cohomology groups of the trace cone on the atlas have strictly negative weights.

To take J'-invariants, first compute rational U-cohomology and then take exact T-invariants. The Lie algebra cochains tensor the coefficients with exterior powers of $\mathfrak{u}^*$, all of nonpositive ν-weight. Their spectral sequence for the trace cone is bounded in cochain degree and the coefficient complex is bounded below. It therefore converges, and no T-invariants survive. If $M = 0$, augmentation is already an equivalence on the atlas. Finally, the trace is induced by the canonical truncation and the counit for its proper direct image. Both truncation and this counit are natural under the rotation of the transport groupoid. This makes (6.6) circle-equivariant; the coordinates used to prove its underlying equivalence need only respect ν. □

Lemma 6.21. *Under the assumptions of Lemma 6.20, the unshifted coefficient in (6.3) is zero if the endpoint characters disagree after transport by d. Otherwise it is a free line over $A_I = \mathcal{O}(Y_\delta)$ in degree zero, where $Y_\delta = T/(1-\delta)T$. Passing to small circle invariants and then inverting u gives the same answer over $A_I \otimes \mathbb{C}[u, u^{-1}]$.*

Proof. By Lemmas 6.18 and 6.20 we must compute the coherent global sections of the ratio of endpoint lines on $[J' a_I / J']$. Pull back from $[T a_I / T]$ by the Levi projection. Every nonconstant radical function and every cochain of positive unipotent degree has strictly negative ν-weight, while the ratio of lines has weight zero. The pullback therefore gives an equivalence on global sections. It preserves the square of the odd element and hence the original boundary action.

The torus acts on the odd coset by $s : (t a_I) \mapsto s\delta(s)^{-1} t a_I$. Its orbits are the fibers of the projection to Y_δ, with stabilizer T^δ. Torus invariants are exact. A line with nontrivial stabilizer character has zero invariant sections; with trivial character it descends to a line on Y_δ. Every line on this torus is free, since its Laurent polynomial ring is factorial. The stabilizer character here is the ratio of the two endpoint characters in their common Levi frame. This proves the stated alternatives in degree zero. A vector space in this single degree has discrete automorphism space, so its circle invariants tensor it with $\mathbb{C}[u]$. Inverting u gives the asserted periodic coefficient. □

Lemma 6.22. *Every $d \in N_I^+$ acts trivially on $\pi_0(T^\delta)$, where $\delta = w_I \theta$.*

Proof. Use a Weyl- and δ-invariant positive definite inner product on the real root space. Let Φ_+, Φ_- be the roots in the positive and negative eigenspaces of δ. Here $\Phi_- = \Phi_I$. Choose positive systems in both, using Φ_I^+ on the negative side, and denote their half sums by ρ_+, ρ_-. A reflection in a root of Φ_+ fixes the negative eigenspace pointwise. Multiplying d by an element of $W(\Phi_+)$ we can make it preserve both positive systems, since d already preserves Φ_I^+. The resulting Weyl element fixes both half sums. The roots orthogonal to $\rho_- + c\rho_+$ for a generic real c are exactly those orthogonal to both. The stabilizer of this vector is generated by their reflections: move the vector to the closed fundamental chamber and take the simple walls through it. Call this root subsystem Ψ.

It is stable under δ and has no root in either eigenspace, since a root of Φ_+ or Φ_- pairs nontrivially with the half sum of its own positive system.

The irreducible components of Ψ are exchanged in pairs. Indeed on a stable component a generic vector in its fixed subspace is regular for that component: a root perpendicular to the whole fixed subspace would be anti-invariant. It defines a stable positive system. Its unique highest root would then be fixed, which is impossible. Thus $W(\Psi)^\delta$ is generated by products $s_\beta s_{\delta\beta}$ of commuting reflections from two exchanged components. For any anti-invariant cocharacter x,

$$(s_\beta s_{\delta\beta} - 1)x = (1-\delta)\big(-\langle\beta, x\rangle\beta^\vee\big).$$

Reflections in Φ_+ fix x. The component map $x \mapsto x(-1)$ identifies the anti-invariant cocharacters modulo $(1-\delta)X_*(T)$ with $\pi_0(T^\delta)$ by Proposition 6.14. This proves the assertion. □

Lemma 6.23. *On the exact kernel position of d, the top coefficient for the coherent dual of Δ_d is the structure sheaf of $[J/J]$ in degree zero. The section telescope of that line represents the dual kernel.*

Proof. Start with the unshifted unit on the diagonal. Along a reduced word, at each length-increasing step the identity term in the triangle $\Delta_s \to q_s^* q_{s*} \to \mathbf{1}$ restricts to zero on the new open position. Over that position the averaging correspondence pulls back the previous line along the panel, by the unique lower lift in Lemma 6.9. Thus the descended line on $[J/J]$ is unshifted and has trivial stabilizer character. The dualizing line of $[J/J]$ is trivial in degree zero, by the right-translation calculation of Lemma 6.18 for the group J acting on itself. Duality commutes with open restriction and with closed direct image after removing the lower positions. The coherent dual therefore has the same top line. It is right orthogonal to those lower positions, so the right extension from the quotient in Lemma 6.11 recovers it as the indicated telescope. □

Lemma 6.24. *In the first flag frame let $y = [t] \in Y_\delta$ be the coordinate of ta_I. In the frame of $B' = {}^dB$ its coordinate is*

$$y \longmapsto [d^{-1}(t)c_d], \qquad c_d = r_G/(d^{-1}r_G). \tag{6.7}$$

On each $\mu \in X^*(Y_\delta) = X^*(T)^\delta$, evaluation on c_d gives the orientation holonomy of the non-Levi real Bruhat cell in the corresponding frame. In particular the norm of (6.7) is d^{-1} applied to the first norm.

Proof. The pinned Tits lift n_d is fixed by θ. Both dw_I and $w_I d$ are length-additive products and are equal, since d preserves the Levi positive system. Their Tits lifts coincide, so n_d commutes with $n_{w_I\tau}$. Conjugating ta_I by n_d^{-1} now gives $d^{-1}(t)r_G(d^{-1}r_G)^{-1}a_I$.

Let $\mathcal{I}(d) = \{\alpha \in \Phi^+(G,T) : d\alpha < 0\}$. It contains no Levi root. For a non-Levi root, positivity is detected by pairing it with λ. Since δ fixes λ and commutes with d, it permutes $\mathcal{I}(d)$. The Weyl inversion identity gives the integral cocharacter

$$h_d = \rho_G^\vee - d^{-1}\rho_G^\vee = \sum_{\alpha\in\mathcal{I}(d)} \alpha^\vee.$$

Consequently $c_d = h_d(-1)$ is fixed by δ and has square one. This proves the norm claim. Its value on μ is

$$\mu(c_d) = (-1)^{\langle\mu, h_d\rangle}. \tag{6.8}$$

Under duality μ is a cocharacter of the marked automorphic Cartan. In the compact norm chart its loop lifts to $s \mapsto \exp(\pi i s\mu)$ for $0 \le s \le 1$, and the endpoint is $\mu(-1)$ in the compact image of the real Cartan, as in Proposition 2.12. For a character μ expressed in the second frame, the cocharacter in the first frame is $d\mu$. With the first flag fixed, the non-Levi cell of flags in position d has complexified real tangent weights $\beta^\vee$ for $\beta > 0$ with $d^{-1}\beta < 0$, by the Bruhat cell in Proposition 5.10 and root duality. These roots are precisely $\beta = -d\alpha$ for $\alpha \in \mathcal{I}(d)$. Evaluating their determinant on $(d\mu)(-1)$ gives the exponent

$$\sum_{\alpha\in\mathcal{I}(d)} \langle d\mu, -d\alpha^\vee\rangle = -\langle\mu, h_d\rangle.$$

Its parity is that of (6.8). A real linear map and its complexification have the same determinant. At this element of order at most two, the determinant is the real orientation sign. The Levi Schubert cell of Proposition 5.10 is complex and contributes its complex orientation. This proves (6.8) for the real return. Interchanging the two flags inverts the coordinate map on both sides, as in Proposition 2.13. □

Theorem 6.25. *For any two distinct pure inner labels and any two compact finite Hecke kernels K, K', one has* $\mathrm{RHom}(KE_\xi, K'E_\eta) = 0$. *The full spectral category therefore decomposes as the orthogonal product of the individual pure-label Hecke subcategories. For a single label of type I the spectral return of Δ_w vanishes unless $w \in W^{w_I\theta}$. In the latter case write $w = dv$ with $d \in N_I^+$ and $v \in W_I$. The return is a free rank-one module over $A_I \otimes \mathbb{C}[u, u^{-1}]$ in shift $[\ell(v)]$. The two coefficient actions are related by the signed Weyl transformation of the real closed-orbit return.*

Proof. Theorem 6.12 gives vanishing for distinct types and, at equal type, for $w \notin W^\delta$. Take $w \in W^\delta$ and write $w = dv$ as in the statement. Corollary 5.24 reduces the computation to d and gives the shift $[\ell(v)]$. The kernel at d has the unshifted top line by Lemma 6.23. Apply Lemma 6.17 to the two closed generators and that kernel. By Lemma 6.21 its torus coefficient compares χ_ξ with the transport of χ_η by d. Lemma 6.22 identifies the transported character with χ_η. Proposition 6.5 says these characters distinguish the pure labels. We obtain zero for distinct labels and a free Laurent line for the same label. Lemma 6.24 identifies its two Cartan actions with those of the real return. In the input-first convention the real calculation uses w^{-1}; its internal Levi factor is conjugate to v^{-1} by a Levi diagram automorphism, so it has the same length.

For completeness, fix distinct labels. The compact finite Hecke kernels K for which $\mathrm{RHom}(KE_\xi, E_\eta) = 0$ form a thick subcategory. We have just shown that this subcategory contains every standard, hence it contains every compact kernel. Moving K' to the first variable by its adjoint gives the stated vanishing between arbitrary translates. The argument applies in both directions. The translates are compact and generate by Theorem 6.6; therefore their orthogonal factors exhaust the category. □

6.6. **The full comparison along the root divisors.** For $t \in T$ put $\Phi_t = \{\alpha \in \Phi : e^\alpha(t) = 1\}$, and let $U_{\le r}$ be the open on which the rational span of Φ_t has dimension at most r. Its complement is the finite union of simultaneous unit loci of $r+1$ linearly independent roots. In particular this is a Weyl-invariant open, and each component of its complement has codimension at least $r+1$. We retain the connected norm image $D_I = (T^{w_I\theta})^\circ$ and the original norm cover $Y_{w_I\theta} \to D_I$.

Lemma 6.26. *In a real form of minimal parabolic type I, put $r = |I|$. Every marked point in the support of a flag orbit satisfies* $\operatorname{rank} \Phi_t \ge r$. *For an orbit whose real Cartan is not maximally split this inequality is strict. If $t \in dD_I \cap U_{\le r}$, then $\Phi_t = d\Phi_I$.*

Proof. Let J be the real Cartan of the orbit, S its full split subtorus, and $S' \supset S$ a maximal real split torus of the group. Take a real maximal torus J' containing S'. Set $L = Z_H(S)$ and $L' = Z_H(S')$ inside the complexified group. Both maximal tori J, J' belong to the connected reductive group L. Its roots in J are exactly those which vanish on S. Under root duality their coroots, regarded as roots of the spectral group, vanish on the connected norm image of that orbit: a root vanishes on the fixed torus precisely when the involution takes it to its negative, and the same holds for its coroot. Hence at every marked point in the orbit support the rank of Φ_t is at least the semisimple rank of L.

Using J' as common torus, the roots of L' form a subsystem of the roots of L, and the first system has rank r since L' is a minimal real Levi. If their ranks agree, every root of L lies in the rational span of roots trivial on S', and is itself trivial on S'. The two root systems and centralizers then coincide. Hence $S' \subset Z(L)^\circ \subset J$, which forces $S' = S$ by maximality of the split subtorus of J. This proves strictness when the Cartan is not maximally split. Finally Φ_t contains $d\Phi_I$ on dD_I. A standard Levi root subsystem contains all ambient roots in its rational span. Thus equality of ranks here gives equality of root systems. □

Lemma 6.27. *On $U_{\le |I|}$ the automorphic category of a pure form of type I is the product of the coefficient categories of its maximally split flag orbits there. Each of their standard extensions is clean on this open, even where two distinct connected norm images meet.*

Proof. Only these orbits survive by Lemma 6.26. They correspond to the distinct dD_I by Lemma 4.5. Every such orbit is reached from the closed orbit by complex simple ascents. Indeed the descent criterion of [IY25, Lemma 6.9] gives on a nonclosed orbit a complex descent or a noncompact imaginary simple root. The latter would increase the split dimension by a Cayley transform and is excluded for a maximally split Cartan. A complex descent changes only its marking, so ends again on a maximally split Cartan. Dimension strictly decreases, and the closed orbit is unique by Lemma 5.6.

On the closed orbit both extensions agree. At each complex ascent let β be the dual root tested at the input flag. It is a complex root for its marked involution, whereas that involution is minus the identity on the transported Levi subsystem. Thus β lies outside that subsystem. Lemma 6.26 shows that $1 - e^\beta$ is invertible on

the part of this orbit's norm cover over $U_{\le r}$, with $r = |I|$. Here and below an empty such part contributes the zero category. By Lemma 2.38 the cone from the standard simple kernel to the costandard kernel is killed by localization at this particular root function. Its action on the input orbit standard therefore vanishes on the stated open. Lemma 2.37 identifies convolution with the standard kernel with standard extension on the upper orbit, and convolution with the costandard with costandard extension there. Its proof uses only the complex panel and holds in the given real form. The coefficient transport is an equivalence, so induction gives cleanness on every surviving orbit. For distinct such orbits ordinary restriction of extension by zero to the other orbit is zero. Adjunction with a costandard extension therefore gives orthogonality in all degrees. The finite stratum filtration now gives the product, with the full coefficient covers of Lemma 2.2. □

Fix a minimal type I, set $r = |I|$, and restrict to $U_{\le r}$. For a label ξ of type I, each standard translate $\Delta_d E_\xi$ has marked support in a Weyl translate of the connected norm image D_I. We first separate distinct such images, which can meet even on this open.

Lemma 6.28. *Take shortest representatives d_1, d_2 for their right cosets modulo W_I, with $d_1 D_I \neq d_2 D_I$. For any labels ξ, η of type I the objects $\Delta_{d_1} E_\xi$ and $\Delta_{d_2} E_\eta$ are mutually orthogonal on $U_{\le r}$.*

Proof. For compact generators their periodic mapping complexes are perfect over $R \otimes P$ by Theorem B.11. It suffices to check their vanishing after derived specialization at closed points of the open: localization at a maximal ideal and the residue-field Koszul complex detect a nonzero perfect periodic module by Nakayama's lemma applied successively to generators of the maximal ideal. For compact X, Y this specialized complex is $\mathrm{RHom}(X, Y \otimes^L_R \mathbb{C}_t)$, by finite cones and retracts. It is enough to consider points common to the two supports, since a function annihilating either generator also annihilates their maps.

Consider now $t \in d_1 D_I \cap d_2 D_I \cap U_{\le r}$. By Lemma 6.26, $d_1 \Phi_I = d_2 \Phi_I = \Phi_t$, so $n = d_1^{-1} d_2$ preserves Φ_I^+ by minimality. All the root equations outside this subsystem are invertible at the corresponding intermediate markings in the two galleries of Lemma 6.13. We may compute maps in a neighborhood of the point where they are still invertible. Equivalently apply the inverse d_1-transport to both objects, which takes the residue-field Koszul complex to that at $d_1^{-1} t$. On the stable Levi chart the first object is now on the open and closed component of Theorem 5.15 where the odd element preserves the indicated parabolic of type I.

For the second object the composite of the two graph transports starts at a chart with $a \in L_I \tau$, with full flag B and $a^2 \in B \cap L_I$. Its final full flag on this chart is ${}^{n^{-1}}B$. To verify the formula through the inverse gallery, write each intermediate marking change in the forward d_1-gallery as a suffix v of d_1. In the composite reverse gallery the corresponding Weyl element is vn; it takes Φ_I^+ to positive roots because n preserves that positive system. Thus all the proposed flags contain $B \cap L_I$, and at each step the other flag in the simple fiber is unique on the root-regular locus. This is the graph calculation in Lemma 6.13 with its full derived base change. The parabolic of the final flag is ${}^{n^{-1}}P_I$, and its Levi is L_I. Since $a \in L_I \tau$, this parabolic is preserved by a exactly when $\theta(n^{-1}) P_I \theta(n^{-1})^{-1} = n^{-1} P_I n$, in Weyl notation. A standard parabolic is self-normalizing, so this requires $n\theta(n^{-1}) \in W_I$. Both n and $\theta(n^{-1})$ preserve Φ_I^+, and the identity is the only element of W_I preserving this positive system. Therefore $\theta(n) = n$. Such an n commutes with both w_I and θ, and preserves D_I. For distinct $d_1 D_I, d_2 D_I$ this is impossible. Thus the second object lies on the open and closed complement of the first near the marked point; their maps vanish in both directions. This test used only the indicated parabolic and applies also to a disconnected centralizer of the even semisimple element. □

Lemma 6.29. *Choose one shortest representative d for each distinct dD_I. The translates $\Delta_d E_\xi$ for these representatives and the labels of type I generate $\mathcal{T}_I|_{U_{\le r}}$. For each fixed d their factors are the module categories of the transported norm covers, one for each label.*

Proof. For each distinct $D = dD_I$ choose one d shortest modulo W_I. For each label of type I take the translate $\Delta_d E_\xi$, with the convention of Lemma 6.13 for transporting the marking to D. For this single d the generators give mutually orthogonal copies of the module category of the transported norm cover, because the kernel is invertible and the original seeds have these Hom complexes. Scalar restriction retains these copies and their marked actions.

By Theorem 6.6 it suffices for generation to consider an arbitrary standard translate. The W_I-factor at the right of any such translate contributes only a shift. For a shortest representative every tested root is outside the transported Levi, so it is invertible in the required root equation on the entire norm support over $U_{\le r}$. Hence the graph calculation of Lemma 6.13 applies there. Representatives giving the same connected norm image differ by an element preserving Φ_I^+ and commuting with θ. The induced conjugation on the Levi permutes the clean

signed factors, exactly as in that lemma. The new translate belongs to the factors already chosen for that norm image. □

Theorem 6.30. *For every minimal type I occurring in the pure inner class,*

$$\bigoplus_{\xi:I(\xi)=I} \mathcal{A}_{\xi}^{\mathrm{per}}\Big|_{U_{\le|I|}} \simeq \mathcal{T}_I|_{U_{\le|I|}}$$

as categories over the marked Cartan with the original period. In particular, the full equivalence in (1.1) holds on $U_{\le 1}$ for every connected quasi-split real reductive group.

Proof. By Lemmas 6.28 and 6.29, the spectral category of type I on this open is the product of the indicated coefficient categories. On the real side the factors are those of Lemma 6.27, where the orbit for dD_I has the transported norm cover and the corresponding free module. This gives the claimed linear equivalence.

For the last assertion every nonempty type has $r \ge 1$, so restriction of its equivalence to $U_{\le 1}$ applies. The empty type is fully compared everywhere by Theorem 6.12. Taking the finite direct sum over the types finishes the proof. □

There is a common open which can include points of larger root rank. For every intermediate minimal type I, remove from T the closed subset $WD_I \cap (T \setminus U_{\le|I|})$. On the remaining open the two categories in (1.1) are equivalent. Indeed each intermediate factor vanishes off WD_I, by its generation and norm support. Its restriction there and its comparison on $U_{\le|I|}$ thus glue, by localization of marked modules, to an equivalence on the complement of the indicated closed subset. The extreme types have their comparisons everywhere.

6.7. **The category supported on the root divisors.** Write $\mathcal{T}_I$ for the factor of Theorem 6.12 of minimal type I, and $Z = T \setminus U$ for the root divisors. A subscript Z denotes the kernel of localization to U.

Corollary 6.31. *The full category supported on the root divisors has the orthogonal decomposition*

$$\mathcal{T}(V)_Z \simeq J\left(\bigoplus_{\xi\in\Theta_{\mathrm{qs}}} (\mathcal{A}_{\xi}^{\mathrm{per}})_Z\right) \oplus \bigoplus_{I\neq\varnothing} \mathcal{T}_I. \tag{6.9}$$

Each nonempty type factor retains the entire ambient mapping complexes, including extensions in directions normal to the divisors.

Proof. Localization carries the empty-type equivalence of Theorem 6.12 to the full equivalence on U in Theorem 1.2. The other factors thus become zero there. Taking the kernel on the finite direct product gives the formula. Its inclusions are fully faithful in the original parameter category. □

Together with Theorem 6.30, this describes the full category over the locus where at most one independent root takes value one. At an intersection of independent root divisors, Theorem 6.25 still separates the pure labels and computes the maps from each standard translate back to its generator. For a noncompact, non-quasi-split label the corresponding comparison of full categories at all such intersections remains open.

Appendix A. The complex-group comparison and compact-first Tate localization

Let H be a connected complex reductive group with pinned Borel $B_H = T_H N_H$, and let G be its Langlands dual. On G use the negative Borel $B = TU$ of [DT26b]. Set $R = \mathbb{C}[\pi_1(T_H)] = \mathcal{O}(T)$ and $H_{\mathbb{R}} = \mathrm{Res}_{\mathbb{C}/\mathbb{R}} H$. Pinned split conjugation c gives the real embedding $h \mapsto (h, c(h))$ in $H \times H$. Thus $\mathcal{A}_H = \mathrm{Shv}_{\mathrm{Nilp}}\big(H(\mathbb{C})\backslash(H/N_H \times H/N_H)\big)$, where h acts by $(g_1N_H, g_2N_H) \mapsto (hg_1N_H, c(h)g_2N_H)$. The quotient is a real analytic quotient stack and arbitrary stalk complexes are allowed.

The dual extension has identity component $G \times G$ and its odd coset exchanges the factors. Its parameter stack is the derived incidence quotient $\mathcal{P}_E = \big[\{(a, C_1, C_2) : a \in E_{\mathrm{odd}},\ a^2 \in C_1 \times C_2\}/(G \times G)\big]$. The incidence is pulled back from the two group Grothendieck–Springer maps along the square map. It carries the boundary circle with transport a^2.

Write $\mathsf{T}(\mathcal{P}_E)$ for the category obtained by taking circle invariants of the bounded coherent objects, inverting the positive generator u of degree two, and then taking the Ind-completion, as in (A.4).

Theorem A.1. *For every connected complex reductive group H, there is a monoidal equivalence*

$$\mathcal{A}_H \otimes_{\mathbb{C}} \mathrm{Mod}(\mathbb{C}[u,u^{-1}]) \xrightarrow{\sim} \mathsf{T}(\mathcal{P}_E)$$

over the entire Cartan in both variables. It identifies the original two right monodromies with the two dual flag markings and preserves the positive boundary class.

There are three steps. We identify $\mathcal{P}_E$ with the Steinberg stack and check rotation and the two flag markings. We then equip the ordinary tame Hecke equivalence with its circle action by computing on tilting kernels. Finally we localize at u on compact objects. This last step removes the nonfinite affine Bruhat positions and gives the finite convolution category.

A.1. **The two flag descriptions.** Write

$$\mathcal{H}_{\mathrm{fin}} = \mathrm{Shv}_{\mathrm{Nilp}}(N_H\backslash H/N_H), \quad X=[B/B], \quad Y=[G/G], \quad Z = X \times^{\mathbf{R}}_{Y} X.$$

Both bracket quotients are conjugation quotients. There are canonical identifications

$$Z \simeq [(\widetilde{G} \times^{\mathbf{R}}_{G} \widetilde{G})/G] \simeq \mathcal{L}(B\backslash G/B). \tag{A.1}$$

Its loop description supplies canonical rotation. Let $p_1,p_2 : Z \to X$ be the two projections, $a : X \to T$ the Cartan map, and $\Delta : X \to Z$ the diagonal. The first flag is output and the second input, so $K * L = p_{13*}(p_{12}^*K \otimes p_{23}^*L)$, $\quad K(N) = p_{1*}(K \otimes p_2^*N)$. All fibre products and tensor products are derived.

Lemma A.2. *There is an equivalence of derived stacks $\mathcal{P}_E \simeq Z$ preserving the ordered Cartan markings.*

Proof. Exchange coordinates can be chosen in the actual extension. If an odd element a has square (z_1,z_2) and induces $f : G_1 \to G_2$, then $f(z_1) = z_2$ because a commutes with a^2. Thus $\eta = (z_1^{-1},1)a$ satisfies $\eta^2 = (z_1^{-1},1)(1,f(z_1)^{-1})(z_1,z_2) = 1$. Its two factor maps are inverse. Identify the factors by f. For restriction of scalars the based root data are exchanged; the remaining modifications are inner and preserve the abstract Cartan lattices.

A point is now $((g_1,g_2)\eta, B_1, B_2)$, with $g_1g_2 \in B_1$ and $g_2g_1 \in B_2$. Gauge by $(1,g_2^{-1})$ to obtain $((h,1)\eta, C_1, C_2)$, $\quad h = g_1g_2$, $\quad C_1 = B_1$, $\quad C_2 = g_2^{-1}B_2g_2$. Both conditions become $h \in C_1 \cap C_2$ and the remaining gauge group is diagonal G. This is an invertible change of ambient coordinates followed by elimination of a free gauge factor, so it identifies the derived fibre products as well as their classical points. □

Lemma A.3. *Lemma A.2 identifies boundary rotation with canonical loop rotation. It identifies the two-flag and three-flag convolution diagrams and their positive circle classes.*

Proof. The orientation double cover of a Möbius strip is a cylinder. Projection to the first exchanged factor identifies exchange-group local systems on the strip with G-local systems on the cylinder; the inverse adds the deck pullback. This is the cylinder descent description of [BN13, Theorem 4.6 and Remarks 4.7–4.8]. It retains the gauge group $G \times G$, and commutes with rotations, their isotopies, and boundary gluing. It therefore identifies the full circle-action diagrams and the convolution correspondences, not only their holonomy on points.

Orient each lifted boundary by the pullback of the positive boundary orientation on the strip. Each maps to that boundary with degree one. The circle of [BCHN25, Section 4.5] rotates the covering cylinder once, hence has the same positive generator. On the gauge slice the automorphism is $((h,1)\eta)^2 = (h,h)$, which is h in the diagonal stabilizer. The additional odd-loop quotient, with its half-degree circle in [BCHN25, footnote 10], is not taken here. □

Lemma A.4. *$H^1(\mathbb{R}, H_{\mathbb{R}})$ is a singleton.*

Proof. In the two conjugate factors a Galois cocycle has form $(a,\overline{a}^{-1})$. It is $t^{-1}\gamma(t)$ for $t=(a^{-1},1)$. Thus the pure inner class has only the given restriction-of-scalars group. □

Lemma A.5. *The map $q(g_1N_H, g_2N_H) = N_Hc(g_1)^{-1}g_2N_H$ induces an equivalence $Q : \mathcal{A}_H \xrightarrow{\sim} \mathcal{H}_{\mathrm{fin}}$.*

Proof. The map is invariant under the displayed real action. Gauge by g_1^{-1}; the remaining stabilizer is N_H, since $c(N_H) = N_H$. This gives an equivalence of real analytic quotient stacks with $N_H\backslash H/N_H$.

To compare singular supports, use a conjugation-compatible invariant pairing on $\mathfrak{h}$. The real diagonal moment map is $\mu_1 + c^*\mu_2$. Its zero equation determines the first covector from the second by negation and conjugation, with adjoint coordinate changes. Each of these operations preserves the nilpotent cone. Nilpotence of the original pair is therefore equivalent to nilpotence of the remaining covector. On H/N_H the cotangent bundle is $H \times^{N_H} \mathfrak{b}_H$. The right Cartan moment map is induced by $\mathfrak{b}_H \to \mathfrak{t}_H$, and an element of $\mathfrak{b}_H$ is nilpotent exactly

when this projection is zero. Left N_H-equivariance gives the corresponding condition on the other side. Thus one obtains exactly the universal monodromic category of [DT26b, Section 2.3]. The stack equivalence respects arbitrary stalks and the universal local systems; quotienting positive real Cartan factors only removes contractible directions. □

Transport the usual finite Hecke convolution through Q:

$$F \star_q F' = Q^{-1}(Q(F) * Q(F')). \tag{A.2}$$

The normalization is that of [Tay25, Section 2.6], and its unit and associativity maps are transported with the functor.

Lemma A.6. *The two original right Cartan monodromies become, in order, the positive left and right Cartan monodromies of* $\mathcal{H}_{\mathrm{fin}}$.

Proof. Translation acts on the quotient coordinate by $x \mapsto c(t_1)^{-1} x t_2$. For a cocharacter λ, $c(\lambda(e^{2\pi i s}))^{-1} = \lambda(e^{2\pi i s})$. The first positive right loop is therefore the positive left loop, and the second is the positive right loop. For the spectral identification one may use the equivalent gauge slice $((1, g_2 g_1)\eta, g_1^{-1} B_1 g_1, B_2)$. Its remaining holonomy is conjugate to the previous one, rather than its inverse, and its flag order is unchanged. □

A.2. **Enhancing the tame equivalence to the actual circle.** Let $I^u = \mathrm{ev}_0^{-1}(N_H) \subset H[[z]]$ and $\mathcal{H}_{\mathrm{aff}} = \mathrm{Shv}_{\mathrm{Nilp}}(I^u \backslash H((z))/I^u)$. Fix the ordinary monoidal equivalence

$$\Phi : \mathcal{H}_{\mathrm{aff}} \xrightarrow{\sim} \mathrm{IndCoh}(Z) \tag{A.3}$$

of [DT26a, Theorem 1.1]. Its compact spectral objects are bounded coherent complexes. Positive rotation is $a \cdot z^\lambda = \lambda(a) z^\lambda$ as in [DT26b, Section 3.3]. Denote its monodromy transformation by m, and the canonical loop transformations on X, Z by γ_X, γ_Z.

For the extended affine Weyl group write Δ_w, ∇_w for universal standard and costandard objects in their perverse normalization. Let T_w be a universal tilting with top stratum w and $\Xi \in \mathcal{H}_{\mathrm{fin}}$ the big tilting. The finite Whittaker module is $\mathrm{QCoh}(T)$, with action represented by Ξ, and the induced module is $\mathcal{M} = \mathcal{H}_{\mathrm{aff}} \otimes_{\mathcal{H}_{\mathrm{fin}}} \mathrm{QCoh}(T)$. This is the universal Iwahori–Whittaker category. If $\lambda = \mu - \nu$ with μ, ν dominant, its Wakimoto object is $W_\lambda = \nabla_{e^\mu} * \Delta_{e^{-\nu}}$ with the canonical translation lifts of [DT26b, Section 4.2].

Lemma A.7. *The equivalence* $E : \mathrm{IndPerf}(X) \xrightarrow{\sim} \mathcal{M}$ *of* [DT26a, Theorem 3.2(b)] *intertwines* γ_X *and* $m_{\mathcal{M}}$. *On the weight-*λ *line this automorphism is* e^λ *with the stated positive sign.*

Proof. We spell out the transformation retained by the construction. The initial free-bundle functor of [DT26b, Section 8] uses $\mathcal{X} = \{(g, x, t) \in G \times (G/U)^{\mathrm{aff}} \times T : gx = xt\}$, where (a, s) acts by (aga^{-1}, axs^{-1}, t). The pair (g, t) stabilizes this point. On a free bundle associated to a representation A of $G \times T$, its action is the coaction matrix M_A of loc. cit., Section 8.2. Propositions 8.3.1(a) and 8.4.1(a) identify $F(M_A) = m_{F(A)}$. Proposition 5.2.1(c,d) identifies this monodromy with actual geometric loop rotation, rather than a separately chosen nearby-cycle operator.

The matrices commute with every free-bundle morphism. Their images are perverse; mapping spaces between perverse objects are discrete because negative Ext vanishes. Thus this is an enhanced natural transformation with its unique higher naturality coherences. Its coefficients lie in the degree-zero image of $\mathcal{O}(G \times T)$ and remain in the graded subring used in loc. cit., Section 10.5. Extend termwise through bounded complexes of free tilting bundles and their retracts. The transformation preserves the irrelevant-boundary subcategory and descends through the Verdier quotient of Proposition 10.5.2. We obtain $F\gamma_X = mF$ on $\mathrm{Perf}(B/B)$.

On the relevant open set write $x = hU$, $g = hbh^{-1}$ and $t = \pi_T(b)$. Quotient by T and gauge h to 1. The residual stabilizer is $b \in B$; its action on the weight-λ line is $\lambda(\pi_T(b)) = e^\lambda$. The negative Borel changes the ampleness convention but does not change this holonomy. It agrees with $m_{W_\lambda} = e^\lambda$ in Proposition 4.2.1(a).

Rotation is trivial on the constant-loop finite Hecke category. The positive-loop presentation is equivariantly equivalent to it after contractible pro-unipotent directions are removed. Give its Whittaker module the trivial action. The equivariant bar diagram for the relative tensor product then equips $\mathcal{M}$ with its induced action and makes averaging equivariant. The identity just proved passes to averaging and Ind-extension.

For a central cover, replacing g by zg multiplies M_A by its central-character factor. The rotation calculation in [DT26b, Lemma 17.2.2, diagram (82), and Proposition 17.3.1] gives the identical factor and retains it under descent. Hence the argument applies to arbitrary connected reductive groups and yields intrinsic holonomy on B/B. □

Let $q : \mathcal{H}_{\mathrm{aff}} \to \mathcal{M}$ be averaging, $q(A) = A \otimes \mathcal{O}_T$. It satisfies $A * q(B) = q(A * B)$.

Lemma A.8. *A morphism $f : A \to A'$ in the perverse heart of $\mathcal{H}_{\mathrm{aff}}$ is zero if its action on every object of $\mathcal{M}$ is zero.*

Proof. Let $C = \operatorname{im} f$ in the heart. Apply the hypothesis to $q(T_w)$. The averaging adjunction gives $q^L q(-) = (-) * \Xi$, so $f * T_w * \Xi = 0$. Convolution by universal tiltings is perverse exact, as used in [DT26a, Proposition 4.1]; therefore $C * T_w * \Xi = 0$. The objects D with $C * D * \Xi = 0$ form a localizing subcategory containing all tiltings, hence the whole affine Hecke category. By [DT26a, Lemma 4.3 and Proposition 4.1], C is infinitely connective. It also belongs to the heart, so t-structure orthogonality gives $C = 0$, and then $f = 0$. □

Lemma A.9. *For every compact universal tilting T, $\Phi(m_T) = \gamma_Z(\Phi(T))$.*

Proof. The construction of (A.3) through the Whittaker module gives, naturally in A, N, $E^{-1}(A * E(N)) \simeq \Phi(A) * N, \qquad N \in \mathrm{IndPerf}(X)$. Rotation is compatible with convolution, so $m_{A*M} = m_A * m_M$. On the spectral action, $\gamma_X(K * N) = \gamma_Z(K) * \gamma_X(N)$: every flag in the correspondence carries the same automorphism, and proper pushforward respects it. Lemma A.7 identifies the total transformation and the transformation on N. Canceling the latter shows that $\Phi(m_A)$ and $\gamma_Z(\Phi(A))$ act identically on every N. When $A = T$, pull their difference back through Φ. This is an endomorphism of a perverse tilting whose Whittaker action is zero. Lemma A.8 kills it. □

Proposition A.10. *Let $\mathcal{T}$ be a small ordinary additive linear monoidal category, with tensor product additive in each variable. Let $\mathcal{C} = \mathrm{Kar}(K^b\mathcal{T})$ be its enhanced stable envelope, with the totalized tensor product. Suppose $\mathcal{T}$ is replete and full in $\mathcal{C}$, with mapping complexes concentrated in degree zero. Restriction gives $\mathrm{Aut}^{\otimes}(\mathcal{C};\mathcal{T}) \simeq \mathrm{Aut}^{\otimes}(\mathcal{T})$. This space is a 1-type. A linear monoidal circle action preserving $\mathcal{T}$ is determined by its monodromy on $\mathcal{T}$.*

Proof. The stable-envelope universal property identifies exact linear functors out of $\mathcal{C}$ with additive enriched functors out of $\mathcal{T}$, including transformations and their higher homotopies. Explicitly one extends by finite totalization and splitting idempotents. Applying this property in each of n variables makes restriction fully faithful on separately exact functors $\mathcal{C}^n \to \mathcal{D}$. For $n = 2$ it identifies the tensor constraint of a functor with its restriction; all higher arities identify associativity, units, and every coherence. Conversely, termwise extension and totalization extend an additive monoidal functor and all its data. Extending an inverse proves the claimed equivalence of automorphism spaces.

Autoequivalences of the ordinary category $\mathcal{T}$ and their monoidal natural isomorphisms form an ordinary groupoid. Thus $G = \mathrm{Aut}^{\otimes}(\mathcal{C};\mathcal{T})$ is a 1-type, and BG is a connected 2-type with $\pi_1(BG) = \pi_0(G), \qquad \pi_2(BG) = \mathrm{Aut}^{\otimes}(\mathrm{id}_{\mathcal{T}})$. A circle action is a pointed map $BS^1 = K(\mathbb{Z}, 2) \to BG$. It lifts to the simply connected cover, which is $K(\mathrm{Aut}^{\otimes}(\mathrm{id}_{\mathcal{T}}), 2)$. Such a pointed map is determined by the image of $1 \in \mathbb{Z}$, namely positive monodromy. Equal monodromy therefore gives equivalent actions with underlying category fixed. □

Proposition A.11. *The tame equivalence (A.3) admits a linear monoidal enhancement identifying positive loop rotation and canonical rotation, with the same period.*

Proof. Let $\mathcal{T}$ be the ordinary category of compact universal tiltings. Standard–costandard adjunction gives

$$\mathrm{RHom}(\Delta_v, \nabla_w) = \begin{cases} R & v = w, \\ 0 & v \neq w, \end{cases}$$

with R in degree zero. Induction through standard and costandard filtrations puts all derived Hom between tiltings in degree zero. The nonpositive truncations of their dg mapping complexes are closed under composition and map by quasi-isomorphisms both to the original complexes and to their degree-zero cohomology. Thus this is an ordinary enhanced algebroid model.

The standard filtration of T_w contains one copy of Δ_w and only lower standards otherwise. Inducting on a finite Bruhat interval recovers Δ_w from T_w and lower standards: remove the filtration pieces above the unique copy, then use the triangle with the pieces below it. Standards compactly generate, since their restrictions are regular modules over the residual Cartan rings. It follows that the stable envelope of $\mathcal{T}$ is the compact affine Hecke category. The monoidal equivalence with $K^b(\mathcal{T})$ is also the one used in [DT26b, proof of Lemma 16.3.1].

Rotation on each side is monoidal: the constructible convolution correspondence is preserved by loop rotation, and the spectral flag diagrams are loop-stack diagrams. Each action functor is isomorphic to the identity by connectedness of S^1, so the replete tilting subcategories are preserved. Lemma A.9 identifies their monodromy

components. On an ordinary category this is equality of the monoidal natural transformations. Transport the spectral action through Φ and apply Proposition A.10. Ind-extension then retains the tensor and circle coherences, as continuous functors and their transformations are determined on compacts. The sign is the one fixed in Lemma A.7. □

Lemma A.12. *The functor Φ is the identity on the unit endomorphism ring R. It identifies positive left and right Cartan monodromy with multiplication through ap_1 and ap_2, respectively, including their commuting coherences and compatibility with the circle enhancement.*

Proof. The constructible unit is the regular local system on the identity orbit and has endomorphism ring R. The spectral unit is $\Delta_*\mathcal{O}_X$, with degree-zero endomorphisms $\Gamma(X,\mathcal{O}_X) = \mathcal{O}(B)^B = \mathcal{O}(T) = R$. For the middle equality, restriction to T is injective because the dense regular semisimple locus of B is conjugate into T; pullback by $B \to T$ is the inverse. Let $\rho : R \to R$ be the map induced by Φ. The monoidal functor $F : \mathrm{Perf}(X) \to \mathcal{H}_{\mathrm{aff}}$ is R-linear and sends $\mathcal{O}_X$ to the unit by [DT26b, Propositions 8.3.1(b) and 10.5.2]. Since $E = qF$, one has $q(r\mathrm{id}_{\mathbf{1}}) = E(r\mathrm{id}_{\mathcal{O}_X})$. The action comparison of [DT26a, Section 3, Proposition 3.3] identifies its other description with multiplication by $\rho(r)$ on $\mathcal{O}_X$. Multiplication by R there is faithful, hence $\rho(r) = r$.

Left and right convolution by a unit endomorphism are the two positive Cartan actions; their group-ring signs are those of [Tay25, Lemmas 2.13 and 2.17]. For our kernel order, these multiply by $(ap_1)^*r$ and $(ap_2)^*r$. Monoidality and $\rho = \mathrm{id}$ prove the statement with commuting coherences. The monoidal circle enhancement preserves the unit and both operations. □

A.3. **Finite Bruhat localization before Ind-completion.** For a compactly generated category with a circle action, we use

$$\mathcal{C}^{\mathrm{Tate}} = \mathrm{Ind}\left((\mathcal{C}^c)^{S^1} \otimes_{\mathrm{Perf}(\mathbb{C}[u])} \mathrm{Perf}(\mathbb{C}[u,u^{-1}])\right). \tag{A.4}$$

Equivariant objects here have compact underlying objects. We work by finite triangles in this small category.

Lemma A.13. *If S^1 acts on a complex torus T' through translation by a nonzero cocharacter λ, then the pullback of u to $[T'/S^1]$ is zero over $\mathbb{C}$. Its action on all equivariant local systems is naturally nullhomotopic, so their compact-first Tate category vanishes.*

Proof. Write $\lambda = d\lambda_0$ with $d > 0$ and λ_0 primitive, and extend λ_0 to a lattice basis. Then $T' = \mathbb{C}^\times \times T''$ with action $a \cdot (t,t') = (a^d t, t')$. As a homotopy quotient, $[T'/S^1] \simeq T'' \times \mathbb{R}_{>0} \times B\mu_d$. The nontrivial part of its map to BS^1 factors through $B\mu_d$, whose positive complex cohomology vanishes. A cochain primitive for the pulled-back class supplies a nullhomotopy natural in every local system and every coefficient map. It restricts to compact underlying objects. Inverting that nullhomotopic map kills the identity of each object, both before and after Ind. □

Lemma A.14. *Every object of $(\mathcal{H}^c_{\mathrm{aff}})^{S^1}$ admits a finite equivariant filtration whose layers are extensions by zero from individual affine Bruhat strata and all of whose terms have compact underlying objects.*

Proof. The compactness criterion [DT26b, Lemma 12.3.1] is finite Bruhat support with perfect R-stalks on residual tori, equivalently finite generation by universal standards. Finite lower Bruhat intervals place the support inside a finite closed union. Rotation preserves its strata. For a maximal stratum $j : V \hookrightarrow X'$ and complementary closed $i : Z' \hookrightarrow X'$, use the equivariant triangle $j_!j^*F \longrightarrow F \longrightarrow i_*i^*F$. Restrictions preserve perfect residual-torus stalks, and closed extension adds zero stalks. Extension by zero of a perfect R-complex on V is a finite construction from its universal standard. Every term is therefore still compact. Induct on the number of strata. The resulting finite triangles remain exact after inverting u. □

Lemma A.15. *For the trivial geometric circle action,*

$$\mathcal{H}^{\mathrm{Tate}}_{\mathrm{fin}} \simeq \mathcal{H}_{\mathrm{fin}} \otimes_{\mathbb{C}} \mathrm{Mod}(\mathbb{C}[u,u^{-1}]),$$

monoidally and with the marked u.

Proof. Let $\mathcal{C} = \mathcal{H}^c_{\mathrm{fin}}$. Its objects are bounded for the usual cohomological t-structure, and their truncations remain compact. Indeed, restriction is t-exact on strata, and finite cohomology modules of a perfect R-complex remain perfect because R is regular. Thus the compactness criterion still applies to each truncation.

For the trivial circle action the induced t-structure on $\mathcal{C}^{S^1}$ is objectwise. A heart object's automorphism space is discrete; a pointed map from $K(\mathbb{Z},2)$ to its classifying 1-type is trivial. Finite equivariant Postnikov towers, retaining their actual extension maps, therefore show that trivial lifts generate $\mathcal{C}^{S^1}$.

For compact A, B, their ordinary Hom complex V is bounded below. Constant cochains on BS^1 give

$$\mathrm{RHom}_{\mathcal{C}^{S^1}}(A_{\mathrm{triv}}, B_{\mathrm{triv}}) \simeq V \otimes_{\mathbb{C}} \mathbb{C}[u].$$

There is no formal-series completion in this equality: if $V^j = 0$ for $j < a$, then total degree d receives a contribution from u^n only for $0 \le n \le (d-a)/2$. The equality thus holds degree by degree and respects composition. It gives full faithfulness of the trivial-lift functor after tensoring with $\mathrm{Perf}(\mathbb{C}[u])$; the finite Postnikov argument gives generation after idempotent completion. Localize and Ind-complete. The trivial lift and cochain action preserve convolution and its unit, yielding the monoidal assertion. □

Proposition A.16. *For*

$$i : I^u\backslash H[[z]]/I^u \hookrightarrow I^u\backslash H((z))/I^u, \qquad p : I^u\backslash H[[z]]/I^u \to N_H\backslash H/N_H,$$

the functor i_*p^* induces a monoidal equivalence

$$\mathcal{H}_{\mathrm{fin}} \otimes \mathrm{Mod}(\mathbb{C}[u, u^{-1}]) \xrightarrow{\sim} \mathcal{H}_{\mathrm{aff}}^{\mathrm{Tate}}$$

with unchanged period.

Proof. Write an extended affine Weyl element as xe^λ. After contractible pro-unipotent directions are removed, its enhanced Bruhat stratum is T_H, and rotation is translation by λ; see [DT26b, Section 3.3 and Lemma 3.3.1]. This removal is compatible with equivariant descent. For $\lambda \neq 0$, Lemma A.13 kills its small equivariant category after Tate localization.

The elements with $\lambda = 0$ are the finite Weyl group. They form a Bruhat lower ideal, with union the closed positive-loop substack. Evaluation at zero has contractible pro-unipotent kernel contained in I^u. It therefore gives an equivalence of the weakly constructible categories, equivariant for rotation and the trivial action on the finite Hecke model. It preserves intrinsic compact objects in both directions.

For an invariant compact object use Lemma A.14, successively removing maximal strata outside this lower ideal. Each removed layer is zero after localization. After finitely many steps the object comes from an invariant compact object on the positive-loop substack. This proves essential surjectivity on the small localized category. Closed extension preserves full mapping complexes with their circle actions, hence remains fully faithful after invariants, localization, and Ind-completion. Lemma A.15 identifies its source with the stated periodization.

Finally, positive loops form a subgroup and evaluation is a homomorphism. The convolution diagrams agree after removal of the same contractible directions. Their units have the regular-torus normalization of [DT26b, Section 12.1]. All maps are equivariant for the same rotation circle, proving monoidality and the identity on u. □

Proof of Theorem A.1. Compose the quotient equivalence of Lemma A.5, periodized, with Proposition A.16, the Tate enhancement of Proposition A.11, and inverse pullback along $\mathcal{P}_E \simeq Z$ from Lemma A.3. Every arrow is monoidal and preserves the positive circle. The coherent compacts used in the Tate step are precisely the compact images under the ordinary tame equivalence, so the localization is the small one prescribed in (A.4). Lemma A.12 identifies both Cartan actions; Lemmas A.2 and A.6 identify them with the original markings. Lemma A.4 shows that there are no other pure inner summands. The equivalence preserves full Hom complexes, composition, and convolution; in particular it also preserves evaluation, coevaluation, and adjunctions of compact kernels. □

Appendix B. Relative properness of the real parameter category

Let L be a connected reductive group over $\mathbb{C}$, and fix an algebraic extension $1 \longrightarrow L \longrightarrow E \longrightarrow \mathbb{Z}/2 \longrightarrow 1$. Write $E_{\mathrm{odd}} = E \setminus L$, let $\mathcal{B}_L$ be the variety of Borel subgroups, and let T_L, W_L, and $R_L = \mathcal{O}(T_L)$ be the abstract Cartan, Weyl group, and Cartan ring. The square map uses the multiplication of this particular extension. Set

$$\widetilde{L} = \{(h, B) : h \in B\}, \qquad U = E_{\mathrm{odd}} \overset{\mathbf{R}}{\underset{L}{\times}} \widetilde{L}, \qquad \mathcal{P}_E = [U/L], \qquad D_0 = \mathcal{O}(E_{\mathrm{odd}})^L.$$

The derived scheme U is quasi-projective and quasi-smooth: both factors and the base in its defining derived fibre product are smooth. We always use the boundary circle, whose transport at (a, B) is a^2, and its positive class $u \in H^2(BS^1, \mathbb{C})$. For clarity, the Tate category in this section is

$$\mathsf{T}(\mathcal{P}_E) = \mathrm{Ind}\Big(\mathrm{Coh}(\mathcal{P}_E)^{S^1} \otimes_{\mathrm{Perf}(\mathbb{C}[u])} \mathrm{Perf}(\mathbb{C}[u, u^{-1}])\Big). \tag{B.1}$$

In particular, equivariance is imposed on bounded coherent objects before either localization or Ind-completion.

Lemma B.1. *The boundary Cartan map $\alpha : \mathcal{P}_E \longrightarrow T_L, \qquad (a, B) \longmapsto a^2 \bmod N_B$ is equivariant for the boundary circle and the trivial action on T_L. The resulting R_L-linear action on the coherent category restricts on $R_L^{W_L}$ to evaluation of invariant functions at a^2.*

Proof. Boundary restriction with its Borel reduction is the map

$$\mathcal{P}_E \longrightarrow [\widetilde{L}/L] \simeq \mathcal{L}(BB) \longrightarrow \mathcal{L}(BT_L) = T_L \times BT_L.$$

Here $\mathcal{L}$ denotes the derived loop stack. The boundary construction is compatible with rotation, including the identity and all compositions of the transport by a^2; see [BCHN25, Section 4.5]. Since T_L is abelian, every cyclic structure map fixes its holonomy coordinate. Projection to that coordinate proves equivariance. A Weyl-invariant function of the Cartan coordinate is the corresponding conjugation-invariant function of the even element a^2. □

Theorem B.2. *For every extension as above and every $M, N \in \mathrm{Coh}(\mathcal{P}_E)$, $\mathrm{RHom}_{\mathrm{Coh}(\mathcal{P}_E)}(M, N) \in \mathrm{Perf}(R_L)$. For every pair of compact objects of $\mathsf{T}(\mathcal{P}_E)$, $\mathrm{RHom}_{\mathsf{T}(\mathcal{P}_E)}(M, N) \in \mathrm{Perf}(R_L[u, u^{-1}])$. Both assertions use the action of Lemma B.1 and the original positive boundary class.*

B.1. **Invariant functions and contraction of nilpotent directions.**

Lemma B.3. *Every algebraic automorphism of a connected complex reductive group preserves a Borel subgroup.*

Proof. Choose a pinning. An automorphism is the product of an inner automorphism and an automorphism preserving the pinning. The former acts trivially on the cohomology of the flag variety: it is induced by an element of the connected group. The latter permutes the Schubert basis, all of whose elements have even degree, and fixes the class of a point. Its Lefschetz number is therefore a positive integer. The topological Lefschetz fixed-point theorem [Hat02, Theorem 2C.3] gives a fixed point on the projective flag variety, which is the required Borel. □

Lemma B.4. *Through evaluation at the square, D_0 is finite over $R_L^{W_L}$.*

Proof. Choose an odd element a preserving a Borel B by Lemma B.3, and fix a maximal torus $T_L \subset B$. The maximal tori aT_La^{-1} and T_L of B are conjugate by some $b \in B$ [Bor91]. Replacing a by the odd element $\tau = ba$ therefore makes τ normalize both B and T_L within the same odd normalizer coset. Write $\theta = \mathrm{Ad}(\tau)|_{T_L}$ and $z = \tau^2 \in T_L$. The involution θ preserves the positive roots. An integral strictly dominant θ-fixed cocharacter is obtained by summing a strictly dominant cocharacter with its θ-translate. It contracts N_B and fixes $T_L\tau$. Consequently restriction gives

$$\mathcal{O}([B\tau/B]) = \mathcal{O}\big(T_L/(1-\theta)T_L\big). \tag{B.2}$$

Indeed, conjugation by $v \in T_L$ changes $t\tau$ to $vt\theta(v)^{-1}\tau$, and every invariant function on the torus quotient pulls back to a B-invariant function on $B\tau$.

The square is $t\tau \mapsto t\theta(t)z$. The norm from $T_L/(1-\theta)T_L$ is an isogeny onto its image subtorus, followed by translation by z: on Lie algebras it is multiplication by two from the quotient by the anti-invariant subspace onto the invariant subspace. Thus the ring in (B.2) is finite over R_L, and hence over $R_L^{W_L}$. Every odd element preserves a Borel, so $[B\tau/B] \to [E_{\mathrm{odd}}/L]$ is surjective on geometric points. The odd coset is reduced; therefore pullback embeds D_0 in the ring in (B.2). This embedding respects evaluation at the square. Since $R_L^{W_L}$ is noetherian, its submodule D_0 is finite. □

Let $\mathcal{N}_L \subset \mathfrak{l}$ be the nilpotent cone, and put

$$I = \{(a, \xi) \in E_{\mathrm{odd}} \times \mathcal{N}_L : \mathrm{Ad}(a)\xi = -\xi\}, \qquad D = \mathcal{O}(I)^L.$$

Give linear functions in ξ degree two. Then D_0 is the degree-zero part of D.

Lemma B.5. *For every $(a, \xi) \in I$, the closure of its L-orbit contains $(a_0, 0)$ for an odd element a_0.*

Proof. There is nothing to prove when $\xi = 0$. Otherwise use the Jacobson–Morozov theorem [CM93, §3.3] in $\mathfrak{l}_{\mathrm{der}}$ to choose $\lambda : \mathbb{G}_m \to L_{\mathrm{der}}$ with $\mathrm{Ad}\,\lambda(t)\xi = t^2\xi$. If $\mathfrak{l} = \bigoplus_j \mathfrak{l}(j)$ is its weight decomposition, then

$$F^{\geq j}\mathfrak{l} = \bigoplus_{r \geq j} \mathfrak{l}(r)$$

is the canonical weight filtration of $\mathrm{ad}\,\xi$. It is characterized by

$$(\mathrm{ad}\,\xi)F^{\geq j} \subset F^{\geq j+2}, \qquad (\mathrm{ad}\,\xi)^j : \mathrm{gr}_F^{-j}\mathfrak{l} \xrightarrow{\sim} \mathrm{gr}_F^{j}\mathfrak{l} \quad (j \geq 0).$$

The centre lies in degree zero. Scaling the nilpotent operator does not change this characterization. The relation $\mathrm{Ad}(a)(\mathrm{ad}\,\xi)\,\mathrm{Ad}(a)^{-1} = -\,\mathrm{ad}\,\xi$ therefore shows that a preserves the filtration and normalizes its stabilizer $P(\lambda) = M_\lambda \ltimes R_u P(\lambda)$, $\quad M_\lambda = Z_L(\lambda)$. Conjugacy of Levi factors in a parabolic gives $v \in R_u P(\lambda)$ such that $a_0 = v^{-1}a$ normalizes M_λ. It still preserves the filtration, as does its inverse.

Every $Z(M_\lambda)$-weight space has a single λ-degree. The permutation of these spaces induced by a_0 preserves this degree, because both a_0 and its inverse preserve the filtration. Consequently λ and $a_0\lambda a_0^{-1}$ have the same adjoint action. They are central cocharacters of M_λ lying in L_{der}. Their quotient is a cocharacter of the finite group $Z(L) \cap L_{\mathrm{der}}$, and hence is trivial. Thus a_0 centralizes λ.

The exponential is a polynomial isomorphism from the Lie algebra of the unipotent radical to the radical in characteristic zero. Writing $v = \exp(\sum_{j>0} v_j)$ gives

$$\lambda(t)v\lambda(t)^{-1} = \exp\left(\sum_{j>0} t^j v_j\right) \longrightarrow 1.$$

It follows that $\big(\lambda(t)a\lambda(t)^{-1}, \mathrm{Ad}\,\lambda(t)\xi\big) \longrightarrow (a_0, 0)$. The displayed orbit map extends algebraically to $\mathbb{A}^1$; its limit remains odd since $v \in L$. □

Lemma B.6. *The ideal $D_{>0}$ is nilpotent. In particular D is a finite D_0-module and has a bounded range of nonzero graded degrees.*

Proof. A positive-degree homogeneous invariant has the same value at (a, ξ) and at the limit $(a_0, 0)$ of Lemma B.5; this value is zero. It vanishes at every geometric point and is therefore nilpotent. Hilbert finite generation for the reductive group L [MFK94, Chapter 1] makes D a finitely generated graded algebra. Choose positive-degree generators $f_1, \dots, f_r$ over D_0 and exponents $e_i \geq 1$ with $f_i^{e_i} = 0$. Then

$$(D_{>0})^{1+\sum_i (e_i-1)} = 0.$$

The monomials $\prod_i f_i^{a_i}$ with $0 \leq a_i < e_i$ span D over D_0 and have bounded degrees. □

B.2. **Singularity operators and bounded ordinary Hom.** The obstruction sheaf of U gives its scheme of singularities

$$\mathrm{Sing}(U) = \mathrm{Spec}_{U^{\mathrm{cl}}} \mathrm{Sym}_{\mathcal{O}_{U^{\mathrm{cl}}}} H^1(T_U).$$

We use the grading in which its fibrewise linear functions are operators of cohomological degree two. Choose an E-invariant nondegenerate symmetric form on $\mathfrak{l}$. On the semisimple part use the Killing form; on the centre choose nondegenerate forms separately on the two eigenspaces of the odd involution. Then $\mathfrak{b}^\perp = \mathfrak{n}_B$.

Lemma B.7. *Scheme-theoretically,* $\mathrm{Sing}(U) = \{(a, B, \xi) : a^2 \in B, \xi \in \mathfrak{n}_B, \mathrm{Ad}(a)\xi = -\xi\}$. *Forgetting B defines a proper L-equivariant map $\phi : \mathrm{Sing}(U) \to I$.*

Proof. At (a, B), put $h = a^2$ and $A = \mathrm{Ad}(a)$. The tangent complex is

$$\Big[T_a E_{\mathrm{odd}} \oplus T_{(h,B)}\widetilde{L} \xrightarrow{dn - d\pi} T_h L\Big]$$

in degrees zero and one. In right Maurer coordinates, $dn = 1 + A$, while $\mathrm{im}(d\pi) = \mathfrak{b} + (1 - \mathrm{Ad}(h))\mathfrak{l}$. A covector annihilating $\mathfrak{b}$ lies in $\mathfrak{n}_B$. It annihilates $(1 + A)\mathfrak{l}$ exactly when $A\xi = -\xi$. This equality implies $\mathrm{Ad}(h)\xi = A^2\xi = \xi$, so the remaining annihilator condition is automatic. These are equations for the kernel of maps of vector bundles, proving the assertion with its scheme structure. Since $\mathfrak{n}_B$ consists of nilpotent elements, the resulting map has target I. It is the projection of a closed incidence in $I \times \mathcal{B}_L$, hence proper and equivariant. □

Lemma B.8. *For $M, N \in \mathrm{Coh}(\mathcal{P}_E)$, regard their pullbacks as equivariant coherent objects on U. The graded sheaf*

$$\mathcal{V}(M, N) = \bigoplus_q \mathcal{E}xt^q_{\mathcal{O}_U}(M, N)$$

is finite over $\mathrm{Sym}\, H^1(T_U)$*, with generators of that operator algebra in degree two. The corresponding sheaf $\widetilde{\mathcal{V}}(M, N)$ on* $\mathrm{Sing}(U)$ *is coherent, graded, and algebraically L-equivariant.*

Proof. On an affine open derived subscheme U', the intrinsic Lie algebra $T_{U'}[-1]$ acts through Hochschild cochains. On cohomology this provides the commuting degree-two operators

$$\Gamma(U'^{\mathrm{cl}}, H^1(T_{U'})) \longrightarrow HH^2(U') \longrightarrow \mathrm{Nat}(\mathrm{id}_{\mathrm{Coh}(U')}, \mathrm{id}_{\mathrm{Coh}(U')}[2]).$$

These are the canonical operators of [AG15, Section 4.1.2, equation (4.2), and Corollary G.2.7]. Their construction by the diagonal groupoid commutes with restriction to affine opens. They therefore glue. The local finite-generation theorem [AG15, Theorem 4.1.8 and Appendix D] says that Ext is finite over the resulting algebra of functions on $\mathrm{Sing}(U')$. It follows that $\widetilde{\mathcal{V}}$ is coherent.

Algebraic equivariance must also be checked, since later we use exactness of invariants of rational representations. On $W = L \times U$, let p and s be its two projections. The relative diagonal over L is the base change of the diagonal of U. The relative version of the operator construction [AG15, Corollary G.2.8 and Section F.4.8] identifies its operators with the pullback operators. Since $T_{W/L} = p^*T_U, \qquad T_W = p^*T_U \oplus s^*T_L$ and L is smooth, the relative-to-absolute map induces an isomorphism on H^1. Flat base change for sheaf Ext gives $p^*\mathcal{V}(M,N) = \mathcal{V}(p^*M, p^*N)$ with these operators. The action map is $p \circ \varphi$, where $\varphi(l,x) = (l, lx)$ is an isomorphism over L. Equivariance of M, N identifies the two pullbacks as operator modules. Applying the same relative diagonal construction on $L \times L \times U$ establishes its cocycle identity. Thus the equivariance is algebraic, not merely an action on the geometric fibres. □

Lemma B.9. *The graded modules*

$$V_p = \bigoplus_q H^p(U^{\mathrm{cl}}, \mathcal{E}xt^q(M,N))$$

are finite rational L-equivariant $\mathcal{O}(I)$-modules and vanish for all but finitely many p. Their invariants V_p^L are finite over D_0 and have a bounded range of graded degrees.

Proof. The projection $\mathrm{Sing}(U) \to U^{\mathrm{cl}}$ is affine. Quasicoherent cohomology on the quasi-compact separated scheme U^{cl} commutes with direct sums. Properness of ϕ and affineness of I consequently give $V_p = H^p(\mathrm{Sing}(U), \widetilde{\mathcal{V}}) = \Gamma(I, R^p\phi_*\widetilde{\mathcal{V}})$. The higher direct images are coherent and vanish above a fixed bound, and they retain their grading and algebraic equivariance.

For completeness, finiteness of covariants follows directly from Hilbert finite generation applied to the square-zero algebra $\mathcal{O}(I) \oplus V_p$. Its invariant algebra is $D \oplus V_p^L$, which is noetherian. The square-zero ideal V_p^L is therefore finite over this algebra, and its action factors through D. Lemma B.6 now makes it finite over D_0. Homogeneous generators over the bounded graded ring D also prove the asserted bound on degrees. □

Proposition B.10. *For coherent M, N, the full ordinary complex $\mathrm{RHom}_{\mathcal{P}_E}(M,N)$ is bounded, has finite D_0-cohomology, and is perfect for its actual R_L-action.*

Proof. The equivariant local-to-global Ext spectral sequence is

$$H^p(U^{\mathrm{cl}}, \mathcal{E}xt^q(M,N))^L \Longrightarrow \mathrm{Ext}^{p+q}_{[U/L]}(M,N). \tag{B.3}$$

The group L is reductive in characteristic zero, so invariants of rational representations are exact. Equivariant descent therefore identifies the abutment with the full mapping complex on the quotient, without higher group cohomology. The finite p-range and the lower bound on q for bounded coherent objects give convergence. Lemma B.9 says that its E_2 page has finitely many nonzero terms, all finite over D_0.

The resulting cohomology is bounded and finite over D_0, hence finite over $R_L^{W_L}$ by Lemma B.4. Lemma B.1 supplies its full R_L-action. The same finite set of $R_L^{W_L}$-generators also generates over R_L. The Laurent polynomial ring R_L is regular of finite global dimension. A bounded complex with finite cohomology over it is perfect.

This proof uses the nilpotence of invariant singularity operators before taking the spectral sequence. Properness of the flag variety alone would not imply bounded Ext on a quasi-smooth derived scheme. □

B.3. **Finite circle invariants and the coherent Hecke action.**

Theorem B.11. *The second assertion of Theorem B.2 holds for the Tate category (B.1), with the marked class u.*

Proof. First let M, N be small coherent circle-equivariant objects. Their ordinary mapping complex V is an object of $\mathrm{Fun}(BS^1, \mathrm{Mod}_{R_L})$ by Lemma B.1, and its underlying cohomology is bounded and finite by Proposition B.10. Apply the functorial standard truncations in Mod_{R_L} objectwise to this diagram. This gives a finite equivariant Postnikov tower with successive layers $H^i(V)[-i]$. This use of truncation concerns an R_L-module with an action; it does not assume that geometric rotation preserves a t-structure on the entire coherent parameter category.

An R_L-module in the heart has a discrete automorphism space: its negative self-Ext groups vanish. An action on it is a pointed map from $BS^1 = K(\mathbb{Z}, 2)$ to the classifying space of that discrete group and is therefore trivial. Consequently $(H^i(V)[-i])^{hS^1} \simeq H^i(V)[-i] \otimes_{\mathbb{C}} \mathbb{C}[u]$. The cohomology module is perfect over R_L, so each such

layer is perfect over $R_L[u]$. Homotopy invariants preserve finite exact triangles, with their actual extension maps. Induction along the finite tower proves $V^{hS^1} \in \mathrm{Perf}(R_L[u])$.

Inverting the original u gives the full Hom in the localized small category and yields a perfect $R_L[u, u^{-1}]$-complex. Every compact object of its Ind-completion is a retract of a finite construction from small images. Finite cones and retracts in either argument preserve the asserted perfection. Together with Proposition B.10, this proves Theorem B.2. □

Proposition B.12. *Put $X_L = [\widetilde{L}/L] \simeq [B/B]$, $Y_L = [L/L]$, and $Z_L = X_L \times^{\mathbf{R}}_{Y_L} X_L$. The Tate category $\mathsf{T}(\mathcal{P}_E)$ is a continuous module over $\mathsf{T}(Z_L)$. The action preserves compact objects and the positive period. Under the complex-group comparison of Theorem A.1, the acting category is the periodized universal finite Hecke category of the complex group dual to L.*

Proof. Use the first flag of a kernel as output and the second as input. The correspondence is

$$\mathcal{W} = \mathcal{P}_E \times^{\mathbf{R}}_{X_L} Z_L, \quad \begin{aligned} p(a, B_1, B_2) &= (a^2, B_1, B_2), \\ r(a, B_1, B_2) &= (a, B_2), \\ q(a, B_1, B_2) &= (a, B_1). \end{aligned}$$

The map q is representable and proper, being a base change of $X_L \to Y_L$. Define on bounded coherent objects

$$K \cdot M = q_*(p^* K \otimes r^* M). \tag{B.4}$$

The relative external product is the pullback of $K \boxtimes M$ by a base change of the diagonal of the smooth stack X_L. It has finite Tor dimension. Indeed, on smooth presentations the tensor product is taken over a regular ring of uniformly bounded dimension. Thus the product is bounded coherent, and its proper direct image is bounded coherent as well. This argument applies to the relative product and does not assert separate finite Tor dimension for both p and r.

The diagonal kernel acts as the identity by derived base change and the projection formula: pulling back its diagonal simply repeats a flag. Associativity is proper pushforward of one relative external product on tuples (a, B_0, B_1, B_2) with a^2 in every flag. The analogous diagrams for arbitrary finite tuples provide all higher coherences; forgetting a flag is composition, and repeating a flag is the unit.

Transport by a^2 on every flag commutes with all these maps. It is simultaneously canonical rotation on the Steinberg kernel and the original boundary rotation on $\mathcal{P}_E$. The action is thus equivariant for a common circle, with the same positive class. Pass to small equivariant objects, form the balanced action over $\mathrm{Perf}(\mathbb{C}[u])$, invert u, and then Ind-complete. This gives

$$\mathsf{T}(Z_L) \otimes_{\mathrm{Mod}(\mathbb{C}[u,u^{-1}])} \mathsf{T}(\mathcal{P}_E) \longrightarrow \mathsf{T}(\mathcal{P}_E).$$

Small equivariant inputs have small equivariant output; finite constructions and retracts prove compact preservation. The final identification follows from the complex-group comparison, whose proof in Section A uses the ordinary tame equivalence and compact-first Tate localization, independently of this proposition. □

Appendix C. Ordinary coefficients by root system and finite verifications

The proof of Theorem 1.7 reduces to ordinary cohomology at marks with central even value. We give the verification for each irreducible root datum and pinned involution here. The induction also uses smaller centralizers, ordered as in Section 3.4.6. Section C.6 gives the inputs, independent verifiers and sizes of the exceptional calculations.

C.1. **The small outer type-A cases.** The induction in outer type A starts with ranks at most four. We prove ordinary concentration in those ranks directly. The rank-two calculation also supplies the ordinary input for the marked comparison in Section 3.5. In ranks three and four, finitely many parabolic sources, together with endpoint transport, generate all tiltings. The finite calculation uses only group-algebra multiplication and Bruhat order.

C.1.1. *The mixed coefficient in outer type A_2.* The rank-one case is Proposition 2.34. For the adjacent exchanged A_2 case, choose the pinned involution

$$\theta(g) = J g^{-\mathrm{t}} J, \qquad J = \begin{pmatrix} 0 & 0 & 1 \\ 0 & -1 & 0 \\ 1 & 0 & 0 \end{pmatrix}$$

on SL_3. In a root chart near its upper Borel write

$$g = \begin{pmatrix} a & b & c \\ xa & xb+d & xc+e \\ za & zb+yd & zc+ye+f \end{pmatrix}, \qquad adf = 1,$$

and put $k = g\theta(g)$. The actual square-Borel atlas is the Koszul dg algebra $D = \mathcal{O}(\Omega)[\eta_{21}, \eta_{31}, \eta_{32}], \quad |\eta| = -1, \quad d\eta_{ij} = k_{ij}$. Define $F = y(d+ex) + fx, \quad Z = z - xy, \quad H = ax + bxy + dy, \quad h = H - (e+cx)Z$. Direct multiplication gives $fk_{21} = h, \quad fk_{31} = yh + (a - f + by - cZ)Z, \quad dk_{32} = ek_{31} + F + (b+cx)z$. The actual two simple-source objects are $D_s = \mathcal{O}(\Omega)/(z, F)[\eta_{21}], \qquad D_t = \mathcal{O}(\Omega)/(Z, h)[\eta_{32}]$, with the inherited differentials and zero actions of the other square generators. These expressions follow by recovering the odd stable refinement of the retained plane or line. Their supports intersect only at the odd stable Borel: the two inputs can have relative position of length at most two, whereas the only theta-fixed A_2 relative positions are 1 and w_0.

A semifree D-resolution of D_t adjoins odd p_h, p_Z and even degree -2 variables u_h, u_{31} with $dp_h = h, \quad dp_Z = Z, \quad du_h = \eta_{21} - p_h/f$, $du_{31} = \eta_{31} - yp_h/f - (a - f + by - cZ)p_Z/f$. The displayed multiplication identities prove $d^2 = 0$. Tensor with D_s, contract the genuine $(u_h, \eta_{21} - p_h/f)$ pair, and rescale by the invertible functions $f, d+ex, d+bx$. The actual tensor becomes $R[x, (d+ex)^{-1}, (d+bx)^{-1}][p, q, u], \quad dp = xT, \quad dq = x^2, \quad du = xp - Tq$, where $R = \mathcal{O}(B)$ and $T = a - f - fx(b+cx)/(d+ex)$. Removing the regular Koszul pair (q, x^2), using the exhaustive finite u-degree filtration, gives the quasi-isomorphic model

$$\overline{D} = R[x, p, u]/(x^2, p^2), \quad |x| = 0, \ |p| = -1, \ |u| = -2, \quad dp = (a-f)x, \quad du = xp. \tag{C.1}$$

The polynomial even variable u remains present. For $I = (x, p, u)$ one has $I^n \in D^{\leq 1-n}$; the nilpotence of x is essential to this estimate.

For the invariant cocharacter $\mathrm{diag}(z, 1, z^{-1})$, the function weights of x, p, u are $1, 1, 2$, whereas all coefficients of R have nonpositive weight. The stable-core immersion into the simple source has $j^! D_s = R[0]$: its one ordinary normal coordinate and its one square equation have the same root character, so virtual dimension and determinant shift cancel. Dualizing a positive augmentation layer therefore gives strictly negative weights. The target has fixed width $[-1, 0]$, and the tail bound above makes the Hom of I^n start in degree at least $n-2$. Hence every fixed ordinary degree is computed by a finite augmentation quotient. The full unipotent Chevalley–Eilenberg complex only contributes further negative weights. Exact torus invariants kill every positive layer. The sole residue is the stable-Borel structure sheaf, with its full torus stabilizer. Consequently $\mathcal{N}_{A_2}(C_s C_t) = A[0] = \mathcal{N}_{A_2}(C_t C_s)$. The residue of the composite counit is 1, so these are the original marked output identifications.

Proposition C.1. *Every finite tilting coefficient for outer type A_2 is ordinary and finite projective over its original marked output ring.*

Proof. The six triangular tiltings $\mathbf{1}$, C_s, C_t, $C_s C_t$, $C_t C_s$, $C_{\{s,t\}}$ have tops $1, s, t, st, ts, w_0$, each once. The first five coefficients are A by their counits; the last is ordinary by Lemma 2.21. Finiteness and the ordinary nonnegative output-fiber theorem give projectivity. Lemma C.11 gives every tilting and every word. Central descent and torus products retain this statement for the actual inherited A_2 root data. The marked comparison, including the two ramification points of the norm map, is proved in Theorem 3.26. □

C.1.2. *Neutrality of the odd Grothendieck bundle.*

Lemma C.2. *Let L be connected reductive, let σ preserve a Borel pair, and let $L\tau$ have this action with τ^2 central. Put $d : L \times^{B_L} (B_L\tau) \longrightarrow L\tau, \qquad \mathcal{M}_L = Rd_*\mathcal{O}$. For every σ-fixed cocharacter ν of the compatible torus, $\mathcal{M}_L|_{Z_L(\nu)\tau}$ has only weight zero under ν.*

Proof. Lemma 2.21 makes $\mathcal{M}_L$ a vector bundle in degree zero; its proof by crepant duality does not depend on a choice of pinning scalars or a particular central square. The base $Z_L(\nu)\tau$ is connected and fixed by conjugation by ν. The restricted equivariant bundle therefore splits into weight subbundles of constant rank.

This fixed base contains every $t\tau$ with $t \in T$. A regular-square point exists there: for a positive root α, both α and $\sigma\alpha$ are positive, so $\alpha + \sigma\alpha \neq 0$; none of the equations $\alpha(t\sigma(t)\tau^2) = 1$ is identically true. Choose t off their finite union. At this point all odd stable Borels belong to the finite set of Borels fixed by its regular semisimple square. The odd stability differential is invertible in every normal root direction, since its square has no root eigenvalue 1. Thus the derived fiber is exactly a finite reduced scheme. A connected torus acts trivially on it, and proper derived base change identifies the bundle fiber with its functions. All nonzero weight subbundles have rank zero at that point, hence rank zero everywhere on the connected base. □

Lemma C.3. *Let $B = \bigoplus_{i\le 0} B_i$ be a finitely generated graded algebra, let $J = B_{<0}$, and let E be a finite graded projective B-module. If E/JE is concentrated in weight zero, then the natural map $B \otimes_{B_0} E_0 \longrightarrow E$ is an isomorphism. It preserves any commuting reductive-group action.*

Proof. Finite homogeneous generation bounds weights above. If a highest positive weight occurred in E, it could not be supplied by JE, whose contribution at that weight would require larger weights in E. It would survive in E/JE, a contradiction. Thus E has no positive weights and $E_0 = (E/JE)_0 = E/JE$. The latter is projective over B_0, being the base change of E. The displayed map is surjective by graded Nakayama. Its kernel is a finite graded module and is zero modulo J, because E is projective; a second graded Nakayama argument kills it. The map is canonical, so any commuting action is retained. □

C.1.3. *A source theorem with two invariant ends.* Throughout this argument G is connected reductive over $\mathbb{C}$ with simply connected derived group. Fix the odd coset $G\tau$ and the marked output ring A. Take two θ-stable end types K_s, K_t and arbitrary middle types J_s, J_t. We calculate the coefficient for $C_{K_t}C_{J_t}C_{J_s}C_{K_s}$ by cutting its flag correspondence at the common middle type.

Move the two end factors by ordinary self-adjunction, and use the ordinary parabolic base-change diagrams for the two middle factors. Set

$$D = J_s \cap J_t, \qquad H = P_D, \qquad Z_a = P_{J_a}/P_{J_a\cap K_a} \quad (a = s, t).$$

The stability of K_a is used here: otherwise the source would retain the additional $K_a \cap \theta K_a$ partial flag, and Z_a would not be this variety. The resulting coefficient is the cross-Hom of the two proper source pushes on the common D boundary, with their actual forgotten-flag bundle factors. The further full-D-flag factor is the ordinary even Grothendieck bundle. Thus every source, bundle, and projector in this cross-Hom comes from the original flag diagram and its original output action.

Lemma C.4. *The diagonal action of H on $Z_s \times Z_t$ has finitely many orbits. Representatives can be chosen in $W_{J_s} \times W_{J_t}$.*

Proof. The diagonal $H = P_D$ action on $Z_s \times Z_t$ has finitely many orbits. Put $U_s = U_{J_s}$ and $U_t = U_{J_t}$. Their positive root sets give $U_D = U_sU_t$. Both U_s and U_t are normal in H, with $[U_s, U_t] \subset U_s \cap U_t$, and

$$U_D/(U_s \cap U_t) \simeq U_s/(U_s \cap U_t) \times U_t/(U_s \cap U_t).$$

The group U_s acts trivially on Z_s, and U_t acts trivially on Z_t, since these groups are normal in the respective P_{J_a} and contained in the denominator of Z_a. Hence the two factors of the displayed quotient act independently on the two choices.

There are finitely many H-orbits on each Z_a, by the parabolic double-coset decomposition in the J_a Levi. On an orbit choose a minimal double-coset representative w_a. Projection to the D Levi identifies its quotient by U_D with L_D/Q_a^D, where $Q_a^D = L_D \cap {}^{\dot w_a}P_{J_a\cap K_a}$ is a parabolic containing the chosen B_D. The independence of the two unipotent factors shows that, on a product of two such orbits, the remaining orbit problem is exactly the diagonal L_D action on $L_D/Q_s^D \times L_D/Q_t^D$. Its orbits are the finite set of parabolic double cosets in W_D. This proves the assertion, with representatives (w_s, w_t) in $W_{J_s} \times W_{J_t}$ after adjusting by W_D. Choose a finite ordering in which each orbit is closed in a successive smooth open of $Z_s \times Z_t$. Supported duality is applied to this smooth orbit immersion, rather than to an orbit closure. □

For such a representative (w_s, w_t) put $R_a = {}^{\dot w_a}P_{K_a}$, $Q_a = R_a \cap P_{J_a}$, $R = R_s \cap R_t$, and $L = R \cap H$. Let v be shortest in $W_{K_s}w_s^{-1}w_tW_{K_t}$.

Lemma C.5. *There exists an odd element preserving both chosen parabolics if and only if $\theta(v) = v$. In this case there is an odd lift τ_0 with $\tau_0^2 = \tau^2$, the stabilizer of the pair of choices is L, and the two sources Y_a and their common incidence $W_{\mathcal O}$ are the derived fiber products in (C.2).*

Proof. Choose invariant dominant cocharacters λ_{K_a} vanishing exactly on K_a, and put $R_a = {}^{\dot w_a}P_{K_a} = P(\lambda_a)$, $\lambda_a = w_a\lambda_{K_a}$. The pair admits an odd element preserving both parabolics if and only if the unique minimal representative v of $W_{K_s}w_s^{-1}w_tW_{K_t}$ is fixed by θ. Indeed θ sends the orbit of this pair to that indexed by $\theta(v)$, and an odd element identifies these two pair orbits. Unsupported orbits therefore contribute zero, including their derived tensor: two coherent source supports with empty intersection have zero tensor.

On a supported orbit write $w_s^{-1}w_t = u_svu_t$ with $u_a \in W_{K_a}$, take a normalizer lift g of w_su_s, and define $\tau_0 = g\tau g^{-1}$. The θ-fixed element v has a θ-fixed Tits representative. It follows that

$$\tau_0^2 = \tau^2, \qquad \mathrm{Ad}(\tau_0)\lambda_a = \lambda_a, \qquad R = R_s \cap R_t, \quad L = R \cap H$$

and that the full common odd-preservation locus is $R\tau_0$. All these intersection groups are smooth and connected. The stabilizer of the choice pair is the entire group L.

Put $Q_a = R_a \cap P_{J_a}$. Derived associativity of the original incidence gives the unweighted source and its common core as

$$Y_a = (R_a\tau_0) \overset{\mathrm{R}}{\underset{R_a,\,\mathrm{sq}}{\times}} Q_a \overset{\mathrm{R}}{\underset{P_{J_a}}{\times}} H, \qquad W_{\mathcal{O}} = H \times^L \big((R\tau_0) \overset{\mathrm{R}}{\underset{R,\,\mathrm{sq}}{\times}} L\big). \tag{C.2}$$

Here $\mathrm{sq}(a) = a^2$ is the actual group-valued square map. It lands in R_a or R because τ_0 normalizes these groups and has the original central square. There is no assumption that τ_0 normalizes H or L. The actual twisted L conjugation preserves $R\tau_0$ and its square condition. The maps from the core to the two sources use the same square path. They define the canonical source augmentation used below. □

Fix a supported orbit and the lift τ_0 of Lemma C.5. The defining cocharacters of R_s, R_t are λ_s, λ_t, obtained by transporting the invariant dominant cocharacters of the two end types. Write $\lambda = \lambda_s + \lambda_t$ and $M = Z_G(\lambda_s) \cap Z_G(\lambda_t)$.

Lemma C.6. *On this orbit coherent duality identifies the coefficient with the supported dual of the tensor of the two source sheaves. Relative to the common incidence, the tensor has an augmentation filtration. Its piece of length e is locally finite free over the incidence algebra, in relative degrees $[-2e, -e]$.*

Proof. Right Maurer coordinates give the full determinant identity

$$\frac{\det \mathfrak{q}_a}{\det \mathfrak{p}_{J_a}} \frac{1}{\det \mathfrak{r}_a} \det(\mathfrak{r}_a/\mathfrak{q}_a) \det(\mathfrak{p}_{J_a}/\mathfrak{h}) = \det \mathfrak{h}^{-1}. \tag{C.3}$$

The source has virtual dimension $\dim H$. This is exactly the boundary dualizing line and degree. Consequently proper coherent duality converts the cross-Hom into the supported dual of the tensor of the two source sheaves with the relative dualizing target of $Z_s \times Z_t$. This operation requires neither classicality of Y_a nor perfectness of its push over the boundary. Translation to the odd coset does not alter any determinant in (C.3).

To describe this tensor, we list the roots of the two resolutions. Write $\Phi(P)$ for the root set of a subgroup containing T, and set

$$\mathcal{F}_a = \Phi_G \setminus \Phi(R_a), \qquad \mathcal{B}_J = \Phi_G \setminus \Phi(P_J), \qquad \mathcal{U}_a = \mathcal{B}_{J_a} \cap \mathcal{F}_a.$$

Thus $\mathcal{B}_J$ consists of the negative roots outside J. Simultaneous equivariant root charts near $R\tau_0$ and near L give coordinates for these actual subgroup embeddings. Such charts are obtained by multiplying the complementary root groups in a fixed order and then the common subgroup. Their differential is the root-space direct sum, so the distinguished charts are etale along the whole intersection. The Jacobian denominators have conjugation weight zero. Pull the even chart back by the actual square map.

For a source-forbidden coordinate x_f adjoin p_f in degree -1. An ambient square equation redundant on the source has an actual homogeneous ideal expression $f_\gamma = \sum_f A_{\gamma f} x_f$; adjoin u_γ in degree -2 and impose

$$dp_f = x_f, \qquad du_\gamma = \eta_\gamma - \sum_f A_{\gamma f} p_f.$$

Genuine source square generators are retained. More generally the right side also includes their actual linear combination, as in the ordinary semifree construction (3.19). This is a resolution of the original source algebra: eliminate each pair (u_γ, du_γ) and then the regular Koszul coordinates x_f. The first contraction divides only by positive integers. No root-character sum or derivative of the square map is inverted.

Tensor the two resolutions over the actual ambient square algebra. Regular-coordinate elimination and cancellation of duplicate square generators with coefficient one leave the shared lists

$$\mathcal{P} = \mathcal{F}_s \cap \mathcal{F}_t, \qquad \mathcal{U} = \mathcal{U}_s \cap \mathcal{U}_t, \qquad \mathcal{E}_0 = \mathcal{B}_D \setminus (\mathcal{U}_s \cup \mathcal{U}_t).$$

The roots in $\mathcal{E}_0 \cap \Phi(R)$ are the finite Koszul equations of the full core in (C.2). The remaining roots form $\mathcal{E}_{\mathrm{extra}}$ and retain their degree -1 square generators. Over the core, the remaining model has shared p of degree -1, shared u of degree -2, and these extra η of degree -1, with

$$dp = 0, \qquad d\eta_{\mathrm{extra}} = 0, \qquad du = \sum A(a)p.$$

The original ambient η still acts by its core or extra value plus its linear p correction. Thus $I = (p, u, \eta_{\mathrm{extra}})$ is a strict augmentation filtration for the original ambient action on the orbit model. Its length-e graded piece is finite free over the finite core in degrees $[-2e, -e]$. Local resolutions need not glue: the augmentation and its

cone are canonical, and these equivariant models calculate that same cone. In the normal directions one uses the joint supported residue filtration below; a positive generator is not declared a subcomplex before applying normal support. □

Lemma C.7. *The additional generators of the tensor in Lemma C.6 have strictly positive function weights for λ. The normal tangent directions to the orbit have strictly negative weights. The degree shift and determinant character in the leading supported costalk both vanish.*

Proof. The cocharacter λ fixes τ_0, so conjugation on $R\tau_0$ has exactly the root weights of R. A root allowed in R_a has nonnegative λ_a value, and a forbidden root has strictly negative value. A negative root outside J_a stays negative outside J_a under W_{J_a}, hence has nonpositive λ_a value. Function weights are the negatives of these root values.

For a root δ write

$$p_a = 1_{\delta\in\Phi(P_{J_a})}, \quad r_a = 1_{\delta\in\Phi(R_a)}, \quad f_a = 1 - r_a, \quad h = p_s p_t, \quad u_a = (1 - p_a)f_a.$$

The multiplicities in the three free-generator lists are $P_\delta = f_s f_t$, $U_\delta = u_s u_t$, $E_\delta = (1 - h)(1 - u_s)(1 - u_t)(1 - r_s r_t)$. The actual normal tangent to the choice orbit is the cokernel of

$$\mathfrak{h} \longrightarrow \mathfrak{p}_{J_s}/(\mathfrak{p}_{J_s} \cap \mathfrak{r}_s) \oplus \mathfrak{p}_{J_t}/(\mathfrak{p}_{J_t} \cap \mathfrak{r}_t),$$

so its root multiplicity is $n_\delta = p_s f_s + p_t f_t - h(1 - r_s r_t)$. Direct expansion gives

$$P_\delta + E_\delta - U_\delta = n_\delta. \tag{C.4}$$

For $(p_s, p_t) = (0,0), (1,0), (0,1), (1,1)$ the two sides are, respectively, $0, f_s, f_t, f_s f_t$. Thus this is the full character multiset identity $\mathcal{P} \sqcup \mathcal{E}_{\text{extra}} = \mathcal{U} \sqcup \mathcal{N}$, not merely a dimension identity.

A shared p or u root is forbidden at both ends, giving strictly positive function weight. If an extra root is source-forbidden, nonredundancy forces $p_s = 1$. It cannot be forbidden at the target also: nonredundancy would then force $p_t = 1$, contradicting $h = 0$. Thus $r_t = 1, p_t = 0$. Target allowance and its negative outside-J_t root condition force its target weight to be zero; its source weight is strictly negative. Extra η therefore has strictly positive function weight as well. The target-forbidden case is symmetric. Core coordinates and genuine core square generators have nonpositive function weight, since their roots are allowed at both ends.

Finally the normal tangent has only three possible nonzero cases. For $(p_s, p_t) = (1,1)$ both flags forbid the root. For $(1,0)$ the source forbids it and it is negative outside J_t; $(0,1)$ is symmetric. Its total λ weight is strictly negative in every case. Equations (C.3) and (C.4) therefore identify the leading supported costalk with the derived structure sheaf of the core, with zero degree shift and zero full determinant character. □

Lemma C.8. *After taking the gauge invariants of the supported dual, only the leading residue of augmentation length zero contributes. Its group of descent is $P_0 = M \cap H$, and its coefficients are the restrictions of the forgotten-flag bundles to the common Levi.*

Proof. If $z_1, \dots, z_c$ are normal functions, their weights β_i are positive. The normal dualizing target has the local cohomology representatives

$$\frac{dz_1 \wedge \cdots \wedge dz_c}{z_1^{n_1} \cdots z_c^{n_c}}, \qquad n_i \geq 1, \qquad \text{wt} = \sum_i (1 - n_i)\beta_i \leq 0.$$

Only the leading residue is neutral. It is annihilated by the normal ideal, so positive normal coefficients in a differential cannot create a missing differential into this neutral piece. Every higher pole has negative weight. Dualizing a positive augmentation generator contributes a further negative weight. Lemma C.7 and right Maurer duality identify the term with simple poles and augmentation length zero with the core, with its actual proper-duality counit.

The ordinary calculation converges degree by degree. The core and the dualizing target have fixed finite Koszul width; a length-e augmentation monomial has degree between $-2e$ and $-e$. Thus a fixed Hom degree sees only bounded e and finitely many such cells. The finite Cech complex for normal support and the finite Lie-cochain complex for the unipotent gauge change the bound by a fixed constant. Normal localization is its algebraic pole-order colimit, with all pole targets of the same finite width. Lemma 3.10 applies to this filtration, including every thickening.

We also need the actual forgotten-flag factors. Even refinements have finite projective Grothendieck presentations with conjugation-invariant matrix entries, hence weight-zero frames. For the odd refinements let $M_a = Z_G(\lambda_a)$ be the Levi of R_a. Their factors are the pullbacks of the actual odd Grothendieck vector bundles

on $M_a\tau_0$. The involution has a compatible Borel pair on each M_a. Use the lifts g and $g\dot{v}$ from the proof of Lemma C.5 at the source and target, respectively. Both λ_a and λ are fixed. Lemma C.2 consequently makes these bundles neutral on $M\tau_0$, where $M = Z_G(\lambda_s) \cap Z_G(\lambda_t)$. The intersection R is $M \ltimes U$, and every root of U has strictly positive λ value. On the affine contraction $R\tau_0 \to M\tau_0$, Lemma C.3 identifies each restricted odd bundle with the extension of its neutral vector bundle. This is canonical and M-equivariant; it applies to the dual bundle too. Thus the actual framed source contributes no positive coefficient weight that could cancel a negative residue weight. Pulling back these vector bundles before the derived tensor and then filtering gives the same assertion for the strict ambient source action.

All coordinate rings, finite projective bundles, localizations and Lie-cochains just used are algebraic rational representations. No Cartan completion or ring of arbitrary formal series is used in this proof. The weight-zero etale denominators preserve this property. The radical of L outside $P_0 = M \cap H$ has positive root weights, so its full rational Chevalley–Eilenberg cochains have nonpositive weights and their positive exterior part is strictly negative. Exact torus invariants remove every negative residue and augmentation layer. The same fixed-width argument justifies the order of these operations. The surviving contribution is the weight-zero core with its full P_0 quotient, not invariants under a reduced or omitted gauge group. □

Lemma C.9. *The contribution of the common Levi in Lemma C.8 is concentrated in degree zero.*

Proof. The group M is a connected Levi with simply connected derived subgroup, and $P_0 = M \cap H$ is parabolic in M. It contains the Borel $M \cap B$, even though it need not be preserved by $\mathrm{Ad}(\tau_0)$. The zero core is exactly $[(M\tau_0) \times^{\mathrm{R}}_{M,\,\mathrm{sq}} P_0/P_0]$ with an M-equivariant vector bundle coming from the odd frames and the even projectors. The ordinary partial Grothendieck map $\pi : M \times^{P_0} P_0 \to M$ satisfies $R\pi_*\mathcal{O} = \mathcal{O}_M \otimes_{\mathcal{O}(T)^{W_M}} \mathcal{O}(T)^{W_{P_0}}$. Here is the partial-push argument, including its ordinary vanishing. Put $A_M = \mathcal{O}(T)$ and $S_M = A_M^{W_M}$. The representation-ring freeness theorem [Ste75, Theorem 1.1] makes A_M finite free over S_M. Reynolds averaging makes $A_M^{W_{P_0}}$ a finite projective S_M-module. Both invariant rings are regular, being the character rings of reductive groups with simply connected derived subgroup. Hence $V_{P_0} = M \times_{\operatorname{Spec} S_M} \operatorname{Spec} A_M^{W_{P_0}}$ is finite flat over M and a local complete intersection; in particular it is Cohen–Macaulay. The adjoint quotient of M is smooth on its regular locus, and its nonregular locus has codimension at least three, by the regular-element theorem [Ste65]. Thus V_{P_0} is smooth over the regular locus and regular in codimension one. It is normal by the Cohen–Macaulay and codimension-one criterion. Its regular-semisimple open is the irreducible cover parametrizing the parabolic flags of type P_0, and every component meets that open by finite flatness. The natural map $M \times^{P_0} P_0 \to V_{P_0}$ is consequently proper birational, so its ordinary structure push is $\mathcal{O}_{V_{P_0}}$.

The source $M \times^{P_0} P_0$ is smooth. Right Maurer determinants give $\omega_{M\times^{P_0}P_0} = \pi^*\omega_M$: the factors $\det \mathfrak{p}_0$ and $\det(\mathfrak{m}/\mathfrak{p}_0)$ multiply to $\det \mathfrak{m}$. Grauert–Riemenschneider vanishing [GR70] for its proper birational map to the normal V_{P_0}, followed by the finite map $V_{P_0} \to M$, gives $R^i\pi_*\omega_{M\times^{P_0}P_0} = 0$ for $i > 0$. The projection formula therefore gives $R^i\pi_*\mathcal{O} = 0$ for $i > 0$. This proves the displayed formula, including its actual map from flag Cartan functions and its finite projectivity over $\mathcal{O}_M$. In particular its derived pullback under the actual square $M\tau_0 \to M$ stays a vector bundle. Proper derived base change identifies that pullback with the push of the exact core above. Tensor its M-equivariant frame bundle and take exact M-invariants on the affine coset $M\tau_0$. The result is concentrated in degree zero. This argument uses neither an identification of the twisted adjoint invariant ring nor any unproved smaller categorical comparison. □

Theorem C.10. *Let G be connected reductive over $\mathbb{C}$, with simply connected derived subgroup, and let θ preserve the chosen Borel pair. Retain the original odd coset $G\tau$, its central square, and its marked output ring A. If K_s, K_t are θ-stable standard parabolic types and J_s, J_t are arbitrary, then $\mathcal{N}_G(C_{K_t}C_{J_t}C_{J_s}C_{K_s})$ is concentrated in degree zero and finite projective over A. Empty types are allowed.*

Proof. Lemma C.4 gives a finite filtration by orbits. Lemma C.5 discards the empty incidences. On each remaining orbit Lemmas C.6–C.8 reduce the supported term to the common Levi. Lemma C.9 puts that term in degree zero. The cohomology exact sequences of this filtration show that the original framed cross-Hom, hence the stated four-parabolic coefficient, is ordinary. Its original output action has not been replaced: all augmentations, residues, adjunctions and projectors were constructed on the original marked flag diagram. Finiteness and the ordinary nonnegative output-fiber statement of Proposition 3.1, together with Lemma 3.3, now imply finite projectivity over A. Flat completion and projection to marked summands may be performed afterward; they are not ingredients in the rational weight calculation. □

The small-rank argument uses the following four instances of Theorem C.10. In each case the left outer end is empty. Therefore λ_t is regular and the common zero Levi M is a torus. No unknown lower connected concentration theorem enters their neutral cores.

Type	Weyl top	Actual ordinary kernel	Top multiplicity
Outer A_3	3421	$C_{12}C_3C_2$	1
Outer A_3	4312	$C_{23}C_1C_2$	1
Outer A_4	45321	$C_{123}C_4C_{23}$	2
Outer A_4	54312	$C_{234}C_1C_{23}$	2

The right end types are respectively $\{2\}$ and $\{2,3\}$, which are stable for the indicated diagram involutions. The two displayed middle types are disjoint, so their common middle gauge parabolic is the full Borel, not a quotient by its torus or unipotent radical.

C.1.4. *Endpoint transport and additive generation.* If $D = \{1, n\}$ is an orthogonal exchanged pair in outer A_n, the actual Levi flag variety is a product of two rank-one flag varieties. Write its two odd maps as a_1, a_n. A square-stable partial flag x_n has the unique odd-stable full refinement $(a_1(x_n), x_n)$. The equation defining the first coordinate is a graph isomorphism, so elimination identifies the derived incidence, including all square directions. It is equivariant for the full Levi, including its finite central coupling. Consequently $C_1Q \simeq C_nQ, \qquad C_DQ \simeq (C_1Q)^{\oplus 2}$. The ordinary stable-Levi embedding transports these identities to the ambient datum. For any middle kernel this gives equality of the ordinary coefficients after replacing a right endpoint 1 by n or conversely. Self-adjunction gives the corresponding left endpoint replacement. Concentration is preserved; projectivity for the actual output marking is deduced by Lemma 3.3 after concentration. The adjacent exchanged pair is never used in this operation.

Lemma C.11. *Let $\mathcal{I}$ be a finite lower Bruhat ideal. Suppose every $w \in \mathcal{I}$ has an actual finite tilting K_w supported at most w, with a positive finite free top-standard multiplicity and an ordinary coefficient. Then every finite tilting supported in $\mathcal{I}$ has an ordinary coefficient.*

Proof. Bruhat induction and top-cell recollement recover a finite positive sum of Δ_w from K_w and the already generated lower standards. A genuine matrix idempotent on that top-cell sum recovers Δ_w. Thus $B_{\mathcal{I}} = \bigoplus_{w\in\mathcal{I}} K_w$ generates this finite-support category. Standard–costandard adjunction makes $E = \mathrm{End}(B_{\mathcal{I}})$ ordinary. A further tilting T corresponds to a bounded projective E-complex P, and both P and its projective dual have cohomology only in degree zero. Split cancellation of the highest positive terms first in P and then in its nonnegative dual leaves one projective in degree zero. Hence T is an additive retract of a finite sum of $B_{\mathcal{I}}$. The coefficient functor preserves these finite sums and retracts. □

C.1.5. *Finite verification in ranks three and four.* In outer A_3, the proper stable endpoint sets are $\varnothing, \{2\}, \{1,3\}$; in outer A_4 they are $\varnothing, \{1,4\}, \{2,3\}$. For each such I, each arbitrary parabolic J, and each $u, v \in W_I$, take fixed reduced endpoint Bott–Samelson words U_u, V_v and the actual source $U_uC_JV_v$. Theorem 3.4 supplies its ordinary coefficient. The only connected nontrivial neutral inputs are the already solved fixed A_1, exchanged orthogonal pair, and outer A_2.

At $v = 1$ its standard character is computed by multiplication in the integral group algebra: a simple factor contributes $1 + s$ and C_J contributes $\sum_{w\in W_J} w$. Its Demazure maximum is its unique highest standard and has positive multiplicity. This is an ordinary standard-filtration statement, not an inference from an Euler characteristic. Enumeration gives

Type	Source words	Source heads	Additional anchors	Padding steps
Outer A_3	168	22	2	0
Outer A_4	848	92	2	26

The two anchors are precisely the respective rows of the preceding four-witness table. In outer A_4, the 26 further steps choose a reduced word of the new top and replace an endpoint 1 or 4 by its mate. The target word has its whole lower Bruhat interval already covered at that point. Lemma C.11 makes that target coefficient ordinary, and endpoint transport proves the new coefficient ordinary. Steps are executed in their recorded order; no unproved top is used in its own lower-interval premise.

The calculation is implemented in `check_outer_small_sources.py`. The independent verifier, `replay_outer_small_sources.py`, imports no producer code or KL table. It enumerates parabolic subgroups, recomputes every standard character and top multiplicity, and checks every lower interval by the subword criterion instead of the producer's permutation rank-matrix criterion. The certificates contain the actual endpoint words, parabolic types, anchor characters and all intermediate lower intervals. They verify exactly 24/24 and 120/120 tops.

Theorem C.12. *All ordinary finite-tilting coefficients in outer types A_3 and A_4 are concentrated in degree zero and finite projective over their original output ring.*

Proof. Apply Lemma C.11 to the whole Weyl group. This proves ordinary concentration for all finite tiltings in both outer types. Lemma 3.3 gives finite projectivity over the original output ring. Central descent, products with tori, and Theorem 3.38 give their marked categorical comparisons. □

C.2. **The classical groups.** The initial outer types are treated in Proposition C.1 and Theorem C.12. The arguments here are uniform in rank. The Coxeter combinatorics use the parabolic decompositions in [BB05, Chapter 2]. The argument uses actual tilting direct sums, so it does not require a braid relation for averaging kernels.

C.2.1. *Root-trivial type A.* Write $W = S_N$, $n = N - 1$, and let $I = \{2, \dots, N-1\}$. Then W_I fixes 1 and $W = W_I \sqcup W_I s_1 W_I$. This is a length-additive source statement. If $w(1) = q > 1$, set $x_q = s_{q-1} \cdots s_1$. Its entries in positions $2, \dots, N$ are increasing, so it is minimal in its right W_I coset. Therefore $w = (s_{q-1} \cdots s_2) s_1 v, \quad v = x_q^{-1} w \in W_I, \quad \ell(w) = q - 1 + \ell(v)$. The actual same-$I$ source UC_1V contains B_w as its leading summand. Its neutral cores are proper and have type A_{n-1} or A_{n-2}. Starting with tori and rank one, Theorem 3.4 and passage to retracts prove concentration in every central block. Proper norm centralizers are products of smaller A blocks with their full torus, so the centralizer formula (3.49) proves concentration at the remaining marks.

C.2.2. *Outer type A: the four source double cosets.* Now let $\theta(s_i) = s_{N-i}$, $N \geq 6$, and put $I = \{2, \dots, N-2\}$. The arbitrary-middle source theorem applies to all four families

$$UC_JV, \qquad J = \varnothing, \{1\}, \{N-1\}, \{1, N-1\}. \tag{C.5}$$

Only I must be stable; the two singleton middle types are legitimate. The neutral connected cores have ranks at most $n-2$; in the open two-panel position the core has rank $n-4$. An empty type means a torus. This remains valid at $N = 6$, so the uniform induction starts after the explicit outer ranks $n \leq 4$.

The four double cosets used below have minimal representatives $1, s_1, s_{N-1}, s_1 s_{N-1}$. Their double cosets are exactly

$$\mathcal{F}_N = \{w : w(1) \neq N, \ w(N) \neq 1\}. \tag{C.6}$$

For a direct check, remove left or right I-descents. The endpoint category is unchanged, and when no such descent remains the interior positions and the positions of interior values are increasing. In the four endpoint categories this forces precisely the four displayed representatives. Reversing the reductions gives a reduced concatenation udv, $u, v \in W_I$. Hence every label in (C.6) is an ordinary summand of an actual source from (C.5).

For an orthogonal exchanged pair $i, j = N - i$, the genuine pair-Levi graph elimination gives ordinary natural endpoint identities

$$C_iQ \simeq C_jQ, \qquad C_{ij}Q \simeq (C_iQ)^{\oplus 2}. \tag{C.7}$$

They are used over the full common central scalar completion. No unproved assertion that they fix each separate marked summand is needed. If w has both i, j as right descents, write $w = us_is_j$ with $\ell(u) = \ell(w) - 2$. The coefficient of a reduced word XC_{ij} for this expression is two copies of that of XC_i, whose every canonical label is shorter. By induction the leading summand B_w is ordinary. The left-sided argument uses adjunction.

It remains to treat $w \notin \mathcal{F}_N$. Diagram transport lets us assume $w(1) = N$; write $w = [N, a, \dots, c, b]$. If $c > b$ there are two outer right descents, so the preceding argument applies. Otherwise set $v = ws_1s_{N-1} = [a, N, \dots, b, c], \qquad \ell(v) = \ell(w)$. For a reduced word X of ws_1, the identity $\mathcal{N}(XC_1) \simeq \mathcal{N}(XC_{N-1})$ transfers ordinary vanishing between these two leading labels once the other head is independently known. The following exhaustive cases supply that reason:

Condition	Independent ordinary conclusion
$c \neq 1$	$v \in \mathcal{F}_N$.
$c = 1$, $b \neq 2$	w has the two outer left descents.
$c = 1$, $b = 2$, $a \neq N-1$	v has the two outer left descents.
$c = 1$, $b = 2$, $a = N-1$	$w = [N, N-1, \sigma, 1, 2]$ has right descents $2, N-2$, an orthogonal exchanged pair for $N \geq 6$.

All paired-descent conclusions use strictly shorter labels. The same-length exchange therefore never closes an unanchored cycle. This proves every central outer-A block. A proper semisimple norm centralizer has smaller fixed A blocks or exchanged blocks, so the centralizer formula (3.49) gives the result at every mark.

C.2.3. *A maximum principle for B, C, D and outer D.* Use $\alpha_1 = e_1 - e_2$, $\alpha_2 = e_2 - e_3$, and write $s = s_1$, $t = s_2$. The argument applies to root-trivial B_n, C_n for $n \geq 3$, root-trivial D_n for $n \geq 4$, and the fork involution of D_n for $n \geq 4$. In the last case s, t are fixed nodes. Put $I = \Delta \setminus \{1\}$, $\mathcal{B} = \{w : w(e_1) = -e_1\}$. W_I has three orbits on signed coordinate vectors, namely $\{e_1\}$, $\{-e_1\}$ and $\{\pm e_2, \dots, \pm e_n\}$. The first and third positions have minimal double-coset representatives $1, s$. Thus every $w \notin \mathcal{B}$ is in W_I or has a length-additive expression usv with $u, v \in W_I$. The source theorem makes B_w ordinary. Its possible neutral types are I and $\Delta \setminus \{1, 2\}$, both smaller.

If $w \in \mathcal{B}$, then $w^{-1}(e_1) = -e_1$. Writing $w^{-1}(e_2) = \epsilon e_a$, $a \geq 2$, shows $w^{-1}(\alpha_1) = -e_1 - \epsilon e_a < 0$. Hence s is a left descent and $sw \notin \mathcal{B}$. Moreover $tw \in \mathcal{B}$, and if t is an ascent of w, the label tw has both s, t as left descents.

At each actual central norm, the first three raw vector coordinates have the form $x_1, x_2, x_3 \in \{a, -a\}$ for a nonzero a. All spin or half-weight lifts remain part of the raw mark. The unsigned permutations s, t preserve those lifts and the central square. The colors are

$$s \text{ white} \iff x_1 = x_2, \qquad t \text{ white} \iff x_2 = x_3. \tag{C.8}$$

These are conditions on closed points indexing full completed factors, not equations imposed on their adic coordinates. Pigeonhole implies that an unsigned S_3 permutation makes s white. If s is black, one of y, sy makes t white.

Fix any $j > 0$ and the complete common even local ring C, with residue field k. Use the finite additive zero detector $m_w(y, z) = \dim_k \big(H^j \mathcal{N}(B_w)_{y,z} \otimes_C k\big)$. The tensor here is the ordinary tensor of an already formed finite cohomology module. Nakayama detects its zero; no derived fiber is being substituted for the coefficient. Write $C_r B_w = \bigoplus_v B_v^{\oplus a^r_{v,w}}$ and set

$$d_w = \sum_{x \leq w} P_{x,w}(1) > 0.$$

Augmentation of the canonical character gives $\sum_v a^r_{v,w} d_v = 2d_w$. A white mark gives $2m_w = \sum_v a^r_{v,w} m_v$, whereas a black descent transports $m_w(y, z) = m_w(ry, z)$.

If $w \in \mathcal{B}$ and s is white, the ordinary predecessor sw and its positive ascent head give $m_w = 0$. If $w \in \mathcal{B}$ has both s, t as left descents, descent transport through $\langle s, t \rangle = S_3$ makes s white, so again $m_w = 0$ for all marks. Suppose a positive maximum $M = \max_{w,y,z} m_w(y, z)/d_w$ exists. Its label lies in $\mathcal{B}$, its node s is black, and t is an ascent. Black descent transport at s makes t white without changing the maximum. The white equation now says

$$M = \sum_v \frac{a^t_{v,w} d_v}{2d_w} \frac{m_v}{d_v}.$$

The coefficients are nonnegative and sum to one. The head tw has positive coefficient and has the two descents s, t, hence its value is zero. Every other value is at most M. This contradicts $M > 0$. Thus every positive row vanishes.

The low bases are tori, rank one, $B_2 = C_2$ with all eight labels already sources, and the type-A identifications $D_3 = A_3$ (including its outer action). In D_4 every nontrivial diagram involution is conjugate, by diagram relabeling, to the fork exchange used here. Order-three triality is outside the pinned-involution setting of the conjecture. Globalization has an acyclic order: D_n, C_n before B_n, outer D_{n+1}. The same-rank proper D_n centralizer in B_n is therefore already known; the other required factors have smaller appropriate rank or type A. The full inherited squares, central forms and spectator tori are supplied by the marked descent proof. After global positive vanishing, the nonnegative-output-fiber lemma supplies projectivity.

C.3. **Exceptional central blocks.**

C.3.1. *Finite forcing and geometric anchors.* The finite part of the proof is a calculation in a finite Weyl group with its *full* raw marked set. It uses the actual direct sums and natural transformations of Section 3.2. In particular, it does not infer concentration from a Grothendieck-group identity, from a generic rank, or from periodized cohomology.

Fix a central even value and a raw marked Weyl orbit Ω. The full odd semisimple sectors associated with different such orbits are disjoint open-and-closed sectors after completion. The boundary action keeps the odd

element, so coefficients between different orbits vanish. We may therefore work on $\Omega \times \Omega$, while keeping every point of Ω and its complete local ring.

For a fixed ordinary $j > 0$ write $M_w(y,z) = H^j\mathcal{N}(B_w)_{y,z}$. The finite proof propagates the predicate $M_w(y,z) = 0$; it works simultaneously for every $j > 0$ and needs no bound on j.

Same-I source labels and their certificate. For each maximal proper stable I, seed the family with the leading labels w_J of the actual parabolic kernels C_J. Propagate only by left or right simple ascents whose node belongs to that *same* I. Every resulting word has the form UC_JV, with both end words supported in I. The source theorem makes its coefficient ordinary. Its positive leading canonical summand is B_w, so B_w is an ordinary source label at every marked sector.

One may store these families as bits: a bit records the omitted node, or omitted involution orbit. A step across a node removes every bit whose I excludes that node. It must not merge the end types and then close the union under arbitrary convolution. Predecessor-gather and upward-scatter recurrences compute the same finite family array; the independent source audits compare them pointwise.

A minimal-length equality graph. Suppose all labels shorter than L are ordinary. Write $D = D_L(w)$ and $E = D_R(w)$ for a label of length L. The descent part of (3.28),(3.32) makes vanishing constant on each product of actual descent orbits $(A,B) \in (W_D\backslash\Omega) \times (W_E\backslash\Omega)$. A state (w,A,B) is a *sink* if neither descent orbit contains a point having a white descent. Every nonsink is ordinary: transport to a white descent s, apply the white identity to the shorter predecessor sw or ws, and take the positive ascent head.

If s is white at the left mark and $sw > w$, formula (3.28) has only labels shorter than w below its head. Thus

$$M_{sw}(y,z) \simeq M_w(y,z)^{\oplus 2}. \tag{C.9}$$

An identical statement holds on the right. Its use is an equivalence of *vanishing predicates*; cancellation of arbitrary objects is not asserted. Two length-L sinks are joined if they are white predecessors of the same length-$(L+1)$ descent-orbit state. The exact left incidence condition is

$$A \cap A' \cap \mathrm{White}_s \neq \varnothing, \qquad B \cap B' \neq \varnothing, \tag{C.10}$$

with the white condition on the second factor for a right incidence. These intersections prevent the use of incompatible raw witnesses.

Lemma C.13. *If each equality component of the length-L nonsource sinks contains an ordinary state, or meets a head which is an ordinary source or has an ordinary source/nonsink predecessor of length L, then every length-L label is ordinary.*

Proof. Nonsinks are ordinary by the preceding shorter-length argument. Each head incidence gives (C.9). Descent transport at the common head makes vanishing equivalent along every edge. A specified zero anchor therefore kills its finite connected component. A cycle with no anchor proves nothing and is recorded as unresolved. □

Diagonal geometric anchors. The uniform local theorem gives an additional actual anchor. In a semisimple simply connected datum with $\theta = 1$, let p be a full raw mark with central square, and $K_p = Z_G(p)$. In the subscripts below, p denotes the original marked idempotent at $y = pr_G$. If $w(p) \neq p$, then

$$\mathcal{N}(\Delta_w)_{p,p} = 0. \tag{C.11}$$

The statement includes the full derived output fiber, the original marked projector, and derived Nakayama. It is proved in Section 3.3.

For clarity, the transfer from (C.11) to the tilting label is elementary. The positive Rouquier complex [Rou04] for a reduced word of w has the Bott–Samelson word in degree zero. All its positive-degree terms are words obtained by deleting at least one letter, and hence have canonical labels shorter than w. Under the shorter-length induction their coefficients are ordinary. Contravariance places these terms in strictly negative total degrees. Consequently

$$H^j\mathcal{N}(\Delta_w)_{p,p} \simeq H^j\mathcal{N}(B_w)_{p,p} \quad (j > 0), \tag{C.12}$$

since the other summands of the degree-zero Bott–Samelson word are also shorter. No assertion about equality in degree zero is used.

Thus a length-L state (w,A,B) is an ordinary geometric anchor whenever

$$\text{there is } p \in A \cap B \text{ with } w(p) \neq p. \tag{C.13}$$

This anchor is applied only at the current induction length, not at an unproved longer head. The ordinary implication is therefore acyclic. For E_8 it removes the single previously unanchored core. It also gives a shorter fresh proof of the small exceptional cases.

The small exceptional types. The program `small_marked_closure.py` starts only with the Cartan matrices. It enumerates the Weyl group as exact integral root matrices, checks every length by counting positive roots sent negative, and checks the left/right actions and their commutation. It constructs all parabolic longest elements and compares the two independent source recurrences. It imports no former source table or closure output.

For G_2, F_4 and simply connected E_6, the normalized raw central branches are the actual coroot-lattice points modulo 2. For E_6 the center has order three, so multiplication by a central square root transports every central norm to this same normalized finite mark calculation. No order-two central lift is lost. The action is $p \longmapsto p - \alpha_s(p)\alpha_s^\vee \pmod 2$. The code computes all descent orbits from this raw action, constructs every sink target, checks every white-head incidence and zero boundary, and adds exactly the geometric anchors (C.13).

Type	$\lvert W\rvert$	Same-I source labels	Raw orbit sizes	Final unresolved
$B_2 = C_2$	8	8	not needed	0
G_2	12	8	$1, 3$	0
F_4	1152	660	$1, 3, 12$	0
Root-trivial E_6	51840	43201	$1, 36, 27$	0

For $B_2 = C_2$, every label is already a source; the assertion is therefore independent of the central branch. For G_2 the nontrivial block has four sink targets across all lengths. For F_4 the two nontrivial blocks have 6 and 556 targets; for E_6 they have 4707 and 475. Each target lies in an anchored component in the final replay. The all-white blocks have no nonsource sink.

The inputs, source recurrences, root actions and all anchor witnesses are specified in Section C.6. In particular, the finite calculation checks the combinatorial hypotheses of Lemma C.13; the geometric hypotheses are supplied by Theorems 3.4 and 3.19.

C.4. **Outer type E_6.** We use the Bourbaki numbering with edges $13, 34, 45, 56, 24$ and involution $1 \leftrightarrow 6$, $3 \leftrightarrow 5$, fixing 2 and 4. The Weyl group has 51 840 elements. The geometric source theorems provide many ordinary canonical coefficients; marked endpoint identities then give an exact finite system whose only nonnegative solution in each positive cohomological degree is zero. All decompositions are the unshifted ordinary decompositions of Proposition 3.16.

C.4.1. *Averaging on exchanged components.* Call a subset K admissible when it is a union of entire connected components of $D = K \cup \theta K$. Thus every θ-fixed connected component of D is contained in K, and for an exchanged pair of components $C, \theta C$, the subset K contains one or both components. This definition is stronger than merely requiring D to be stable.

We first establish the ordinary endpoint identity behind this definition. Consider a genuine stable Levi whose derived root system consists of two orthogonal components exchanged by the odd involution. Its flag variety is intrinsically a product $\mathcal{F}_C \times \mathcal{F}_{\theta C}$, even when its full torus has a finite central coupling. If an odd element a has factor maps a_1, a_2, a square-stable second flag F_2 determines the unique odd-stable refinement $(F_1, F_2) = (a_1 F_2, F_2)$. In the actual derived incidence, the first odd-stability equation is a graph equation eliminating F_1 with coefficient one. The remaining equation is precisely the original square-stability equation for F_2. Consequently the forgetful map from the odd-stable flag source is an isomorphism onto the partial source, as a derived correspondence, and

$$C_C Q \simeq \mathcal{O}_V \simeq C_{\theta C} Q. \tag{C.14}$$

The maps preserve the odd element, its inherited central square, all spectator torus variables, and the complete gauge action.

For simply connected derived groups the ordinary even Grothendieck push has the form

$$R\pi_* \mathcal{O} = \mathcal{O}_{L_C} \otimes_{\mathcal{O}(T)^{W_C}} \mathcal{O}(T), \tag{C.15}$$

where the last torus algebra is free of rank $\lvert W_C\rvert$ over its invariant subring. A basis gives an actual equivariant vector bundle isomorphism; the adjoint action on this basis is trivial. Proper derived base change, including singular square fibers, therefore gives

$$C_C^2 \simeq C_C^{\oplus |W_C|}, \qquad C_{C \cup \theta C} Q \simeq (C_C Q)^{\oplus |W_C|}. \tag{C.16}$$

These assertions do not use the cohomology of a reduced flag fiber in place of the actual Grothendieck push. They transport to the ambient group through the genuine stable-Levi correspondence. The central-cover descent needed for general central forms preserves concentration; it need not descend a particular noncanonical choice of free basis.

For a general admissible K, choose one component in each exchanged pair of D. Apply (C.14)–(C.16) separately to these mutually orthogonal components, retaining every fixed component unchanged. There is an actual endpoint object E_D and positive integers

$$m_K = \prod_{\substack{\{C,\theta C\}\subset D\\ C\cup\theta C\subset K}} |W_C|, \qquad m_D = \prod_{\{C,\theta C\}\subset D} |W_C|$$

such that

$$C_K Q \simeq E_D^{\oplus m_K}, \qquad C_D Q \simeq E_D^{\oplus m_D}. \tag{C.17}$$

In particular, $C_D Q$ is a positive finite number of copies of $C_K Q$. For two admissible subsets K, L with the same union D one obtains

$$(C_K Q)^{\oplus m_L} \simeq (C_L Q)^{\oplus m_K}. \tag{C.18}$$

This is a natural identity over the even central boundary. Applying an arbitrary genuine middle kernel, or using ordinary self-adjunction at the left endpoint, gives the corresponding identity of coefficient functors. Its naturality commutes with idempotents of a completed canonical middle object. No assertion of full marked-torus linearity of a chosen basis is made here.

It follows that for admissible K_t, K_s and arbitrary J_t, J_s, the coefficient of

$$C_{K_t} C_{J_t} C_{J_s} C_{K_s} \tag{C.19}$$

is ordinary. Indeed, padding both endpoints with (C.17) identifies a positive number of copies of this coefficient with the coefficient having stable endpoints D_t, D_s. The latter is ordinary by Theorem C.10. Finite direct sums detect the vanishing of every positive cohomology group. Because the whole completed Q was retained throughout this argument, concentration then holds for every marked matrix summand.

C.4.2. *Sources and endpoint transfer.* There are exactly 32 admissible subsets for this diagram. Encoding a subset by $\sum_{i\in K} 2^{i-1}$, they are

$$\begin{gathered} 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 16, 18, 20, 22, \\ 28, 30, 32, 33, 34, 35, 40, 41, 42, 43, 48, 50, 53, 55, 61, 63. \end{gathered} \tag{C.20}$$

The exact source construction uses the following operations.

Four-parabolic source heads. Take the Demazure product of the four longest parabolic elements in (C.19). The actual parabolic word has this canonical head with positive multiplicity. The ordinary canonical decomposition of Proposition 3.16 makes that head an actual retract of an ordinary object. The resulting number of distinct heads is $8,285$. The intermediate counts after successive factors are $32, 412, 3106, 8285$.

Same-I source heads. For each proper stable I, start from w_0^J and perform arbitrary left or right simple ascents from I. A sequence of such ascents is a literal word UC_JV with endpoint letters in I. Its canonical head is an ordinary retract by Theorem 3.4 and Proposition 3.16. For this initial family restrict to J with $J \cap \theta J \subset I$. This step enlarges $8,285$ to $9,605$ distinct known labels. Its use of proper I ensures strict root-count decrease in every neutral source core; there is no outer-E_6 concentration assumption in this step.

Safe endpoint transfer. Let K, L be admissible with $K \cup \theta K = L \cup \theta L$. Suppose K lies in the right descent set of w, and put $u = ww_0^K$. Then

$$\ell(u) + \ell(w_0^K) = \ell(w), \qquad v = u \star w_0^L. \tag{C.21}$$

Assume that *every* canonical label $z \le v$ is already ordinary. Choose a reduced word $\mathbf{u}$ for u. All canonical summands of $C_{\mathbf{u}} C_L$ are then ordinary, since their support is contained in $[e, v]$. Equation (C.18) gives positive finite copies of $\mathcal{N}(C_{\mathbf{u}} C_K)$ and $\mathcal{N}(C_{\mathbf{u}} C_L)$ that are isomorphic. Thus the former is ordinary, and its canonical head B_w is ordinary. The left version uses $u = w_0^K w$ and the left Demazure product.

The full lower-ideal hypothesis is essential. Knowing only that the head v is ordinary would not justify this transfer. The calculation checks all $z \le v$ at each transfer, not merely a head flag. Iterating these safe transfers first enlarges $9,605$ to $14,867$.

Stable-orbit omission sources. Finally take

$$J = \{1,6\}, \quad \{3,5\}, \quad \{2\}, \quad \{4\}, \qquad I = \Delta \setminus J. \tag{C.22}$$

Each I is stable and proper. All left/right I-ascent families from w_0^J are again supplied by Theorem 3.4. They enlarge the set to $22,429$; another safe endpoint closure gives exactly $26,242$.

The verification uses the 72-root permutation representation, with lengths checked by inversion counts and all $459,588$ strong Bruhat covers constructed by deletion of letters in reduced words. It constructs the source set from literal words, without an initial membership table. Its complete source ancestry is

direct four-parabolic witnesses	$8,285$
additional sources with a fixed stable end type	$1,320$
additional omission-source witnesses	$7,562$
safe endpoint transfers	$9,075$
total initial labels	$26,242$.

(C.23)

Every direct witness is replayed as an actual word, and every safe transfer verifies its entire Bruhat lower ideal. The ancillary program `e6_source_replay.cpp` records a literal source or an ordered endpoint-transfer witness for every resulting label.

C.4.3. *The four raw marks and the actual pair identities.* Write the invariant character coordinates of the raw torus element t as $a = e^{\omega_1+\omega_6}$, $b = e^{\omega_3+\omega_5}$, $c = e^{\omega_2}$, $d = e^{\omega_4}$. The corresponding marked object is indexed by $y = tr_G$. The norm has fundamental-weight coordinates (a, c^2, b, d^2, b, a). The center is μ_3, on which θ acts by inversion. A central square of an odd element must be fixed by this action, so its only possible value is 1. The actual raw central marking fiber is therefore

$$a = b = 1, \qquad c, d \in \{1, -1\}. \tag{C.24}$$

No central lift has been discarded: $1 - \theta$ acts by squaring on μ_3, which is surjective. Label the four raw marks by $0 = (+,+), 1 = (-,+), 2 = (+,-), 3 = (-,-)$ in the (c,d) coordinates.

The fixed-node root evaluations and actions are

$$\alpha_2(t) = d, \quad \alpha_4(t) = c, \qquad s_2(c,d) = (cd, d), \quad s_4(c,d) = (c, cd). \tag{C.25}$$

These actions have orbits $\{0\}$ and $\{1,2,3\}$; the two exchanged pair reflections below fix every mark, so these are also the full W^θ orbits. An odd element whose square is 1 is semisimple. Its stable Borels have raw torus coordinates in a single twisted Weyl orbit: conjugate a stable Borel pair to the chosen pair and then use the normalizer of the torus. Hence two distinct displayed orbits cannot occur at the two ends of the same ordinary source. Their formal completions remain disjoint, so the corresponding derived source is zero. The possible ordered sectors are therefore $(0,0)$ and the nine pairs in $\{1,2,3\}^2$. The panel identities of Proposition 3.18 supply the two fixed-node row types of the finite certificate, with precisely the actions in (C.25).

For the exchanged pairs, one cannot simply project the whole-Q identity (C.14) by an arbitrary marked-torus idempotent. The required stronger fact has a separate geometric proof. For an orthogonal pair $\{s,t\}$ exchanged by θ, retain its actual Levi L and full torus. An odd element a exchanges the two projective-line factors by maps A_s, A_t. A square-stable pair (x_s, x_t) determines two odd-stable pairs $(x_s, A_t x_s)$, $(A_s x_t, x_t)$. The inverse takes the first coordinate of the first pair and the second coordinate of the second. This identifies the actual incidences

$$V_L \simeq \mathcal{F}_L \times^R_{[L\tau/L]} \mathcal{F}_L. \tag{C.26}$$

For derived verification on an equivariant central cover, write $a = (g_1, g_2)\tau$, set $h = g_1 g_2$, and transport the second flag by g_2^{-1}. The two square-stability sections become the two ordinary h-stability sections by an invertible change of coordinates. Hence their Koszul diagrams, including all derived intersections, agree. The intrinsic factor-map description shows descent through the central coupling and preservation of spectators and the actual square.

Under (C.26), Q is the diagonal pushforward and the partial maps are the two projections. Their composites with the diagonal are identities, so proper push followed by derived pull gives

$$C_s Q \simeq \mathcal{O}_{V_L} \simeq C_t Q \tag{C.27}$$

over the odd adjoint-character quotient. This quotient is $Y_{\mathrm{raw}}/\langle st \rangle$ for the pair Levi. After completion, its factors correspond to the finite st-orbits of raw marks. If $st(y) = y$, its factor contains exactly the one marked summand

Q_y. Restricting (C.27) to that entire factor therefore gives

$$C_sQ_y \simeq C_tQ_y. \tag{C.28}$$

The whole completed factor, including its nilpotents, is retained; ramification does not create extra raw marks. The rank-two ordinary even Grothendieck push additionally gives $C_{st}Q_y \simeq (C_sQ_y)^{\oplus 2}$.

On the full invariant character lattice of outer E_6,

$$s_1s_6 : a \longmapsto b/a, \qquad s_3s_5 : b \longmapsto ad^2/b, \tag{C.29}$$

with the other displayed coordinates unchanged. Both transformations fix all four marks in (C.24). Thus (C.28) gives exactly the marked identities $C_1Q_y \simeq C_6Q_y$ and $C_3Q_y \simeq C_5Q_y$ used in the third row type of the certificate. Genuine stable-Levi transport and ordinary self-adjunction give their ambient and left-end versions. Applying a completed canonical middle preserves naturality and its Karoubi projector. These are ordinary identities on the same completed even central ring as the fixed-node rows.

C.4.4. *Additional sources with arbitrary middle parabolics.* Next use every arbitrary middle parabolic J in the four maximal proper stable types I, obtained by omitting respectively $\{1,6\}$, $\{3,5\}$, $\{2\}$ or $\{4\}$. Starting from each longest w_J, extend on either side only by length-increasing simple letters in I. The actual source is $C_{\mathbf{u}}C_JC_{\mathbf{v}}$; its top occurs with positive unshifted multiplicity. The union gives 19 892 source labels. Its union with the previous family has 26 760 labels. Only the latter 19 892 family is used to supply additional zeros in the final positive elimination below.

One useful source in this family is $T = C_4C_{1356}C_{2345}C_{1356}C_4$. It is a source of Theorem 3.4 with $I = \{1,3,4,5,6\}$, $J = \{2,3,4,5\}$ and endpoints supported in that outer-A_5 Levi. Independent standard-character multiplication gives mass 995 328, support 17 104, and unique top $v = 1423154234565423143546$, $\ell(v) = 22$, top multiplicity 16. The deterministic Weyl enumeration labels this top 39447. Every standard support label lies below it. Thus its ordinary concentration implies that this canonical top is ordinary in every marked sector.

C.4.5. *All simple endpoint equations with exact lower terms.* For each positive cohomological degree j, localize and complete at the common central norm point. Denote the ordinary coefficient module in a sector (z,y) by $H^j_{w,z,y} = H^j\big(e_z\mathcal{N}(B_w)e_y\big)$. It is finite over the completed common central ring. Choose any residue field k there and set $m_{w,z,y} = \dim_k(H^j_{w,z,y} \otimes k)$. These are finite nonnegative integers, additive under the actual unshifted direct sums. Every ordinary source has $m = 0$. Vanishing of all these numbers detects the finite modules by Nakayama. No derived-fiber shift, generic-rank test or lifting of a numerical ray to an indecomposable object is being used.

For a simple ascent $ws > w$, the exact unshifted canonical product is

$$B_wC_s = B_{ws} \oplus \bigoplus_{x<w,\ xs<x} B_x^{\oplus\mu(x,w)}.$$

For a descent it is $B_w^{\oplus 2}$. The left formulas are analogous. On a white fixed endpoint the coefficient of this product equals two copies of the original coefficient. On a black fixed endpoint its coefficient equals that at the reflected raw mark. For the exchanged orthogonal pairs the actual marked identities are $C_1Q_y \simeq C_6Q_y$ and $C_3Q_y \simeq C_5Q_y$. Self-adjunction supplies their left versions. All equations preserve the common central action. They therefore give exact homogeneous linear equations in the nonnegative m's; identical terms may be subtracted as integers.

We explain why the coefficient list is exhaustive. If $\ell(w) - \ell(x) > 1$ is odd and some right descent of w is an ascent of x, then $P_{x,w} = P_{xs,w}$ and the degree bound kills the highest possible coefficient. The same is true on the left. Thus a nonzero higher-gap μ requires both descent containments; all comparable gap-one pairs have coefficient 1. An independent left-q KL recurrence, with no inverse or involution cache quotient, checked all 481 490 possible higher-gap pairs whose lower label is not among the 26 242 ordinary labels. Exactly 84 837 are positive, and their maximum is 3. The largest polynomial coefficient computed was 159, and the largest subtraction product was 40. Integer arithmetic was used throughout.

Using these independently checked coefficients, the root-permutation engine rebuilt every nonzero endpoint equation, 1 374 050 in total. It then performed 240 317 successive deductions of vanishing. At each step the targeted variable still occurred and all still-unknown terms in its actual equation had one sign; positivity then forced them to vanish. This left exactly 15 663 states and rebuilt all their 90 782 residual rows, coefficient by coefficient.

C.4.6. *The final positive elimination and its implication.* The additional arbitrary-middle sources immediately set 256 of those residual states to zero. Starting with the remaining 15 407 variables, we rebuild a rational graph from all residual rows at every step. A two-term row $ax + by = 0$ with $ab < 0$ identifies $y = (-a/b)x$ with positive rational ratio. A cycle with inconsistent ratios forces its whole component to zero. A row whose surviving terms have only one sign also forces every occurring component to zero. After contraction the full original row system is rebuilt, so no row is dropped on the basis of a numerical-rank heuristic.

Active before	Two-term components	Active after	One-sign rows	Inconsistent cycles
15407	138	115	752	3998
115	46	33	864	120
33	17	1	1884	0
1	1	0	4	0

All ratios are exact rational numbers, and every equation is retained when passing to the contracted variables. Every operation is valid for nonnegative real numbers, and hence for the actual m's. Its conclusion is that all 518 400 retained marked coefficients have zero positive cohomology; the other sectors were already zero by support. As the degree $j > 0$ was arbitrary, ordinary concentration follows. Proposition 3.1 gives the absence of negative cohomology, and Lemma 3.3 gives finite projectivity. Proper semisimple centralizers of outer E_6 have only the already-solved classical factors, including outer D_5; Proposition 3.25 globalizes the result. Theorem 3.38 then gives the marked categorical comparison.

The ancillary implementation separates the root model, the three independent checks and the final rational elimination:

Program	Role
`e6_root_permutation.hpp`	Roots, lengths and Bruhat order
`e6_source_replay.cpp`	Geometric source labels
`e6_mu_left_audit.cpp`	Kazhdan–Lusztig coefficients
`e6_event_independent_audit.cpp`	Endpoint equations
`e6_augmented_positive_replay.py`	Nonnegative elimination

All exact input tables are supplied with these programs.

Theorem C.14. *For outer type E_6, every ordinary finite-tilting coefficient is concentrated in degree zero and finite projective over the original marked output algebra. Consequently the categorical comparison holds for this type, including all central forms and torus factors.*

Proof. The finite argument proves the central-norm assertion at every raw mark. The dependencies of every proper norm centralizer are classical by Section 3.4.6. Apply Proposition 3.25, Theorem 3.38, and the central descent of Section 3.4. □

C.5. **Types E_7 and E_8.** The remaining exceptional calculations are finite but large. We specify the inference rules and the complete counts; independent verifiers implement these rules without using the search order. The algorithms and their input formats are given in Section C.6. Geometric anchors are used only at the current induction length.

C.5.1. *E_7: raw marks, source labels, and safe guards.* Use the chain $1-3-4-5-6-7$ with node 2 attached to 4. The exact regular action has 2 903 040 Weyl elements and longest length 63. The proper same-I families, for all seven maximal proper I and all middle parabolics, supply 1 689 999 source heads, leaving 1 213 041 labels. Both endpoint words of each family keep the same I. An independent source replay compares the predecessor-gather and upward-scatter constructions at every label; the geometric interpretation is Theorem 3.4 and the positive canonical-head theorem, not closure under arbitrary convolution.

In fundamental-character coordinates a raw mark is $b \in (\mathbb{Z}/4)^7$, subject to $A_{E_7} b = 0 \pmod 2$. There are 256 raw marks, with Weyl orbits of sizes 1, 1, 63, 63, 56, 72. The sum of their squares is 16 260. These are the supported ordered sectors. Both central norm branches are included. The two singleton and two 63-point orbits have raw square 1 and actual norm z_G. The 56- and 72-point orbits have actual norm 1. No square-root translation on the order-two center is used. The all-white singleton orbits are settled from the identity source by increasing-length white-head induction.

There are two distinct finite compressions. For a fixed label w, actual black-descent transport makes its zero predicate constant on the $W_{D_L(w)} \times W_{D_R(w)}$ classes of its endpoints. Separately, adding the central phase $2p$, where $p = (0, 1, 0, 0, 1, 0, 1)$, preserves root colors and intertwines every simple action. It exchanges the two

63-point proof graphs and acts freely within the 56- and 72-point orbits. Thus quotient-mark graphs of sizes $63, 28, 36$ can be used, with ordered-sector lift multiplicities $2, 4, 4$. Each zero assertion means zero at every supported raw lift. Only the proof graphs are identified by this central symmetry; isomorphisms between the corresponding coefficient modules are unnecessary.

The following test bounds the support of the terms whose Kazhdan–Lusztig coefficients have not been evaluated.

Lemma C.15. *For a left ascent C_sB_u, a term B_v with nonzero coefficient and $\ell(u) - \ell(v) > 1$ satisfies*

$$\ell(u) - \ell(v) \geq 3 \text{ odd}, \qquad D_L(u) \subset D_L(v), \quad D_R(u) \subset D_R(v), \quad s \in D_L(v). \tag{C.30}$$

The right version interchanges the endpoints.

Proof. The simple product formula in Section 3.2 requires $v < u$, $s \in D_L(v)$ and a nonzero top coefficient $\mu(v, u)$ in the Kazhdan–Lusztig polynomial $P_{v,u}(q)$. Thus $\ell(u) - \ell(v) = 2k + 1$ for some $k \geq 1$, and $\mu(v, u)$ is the coefficient of q^k. Suppose $t \in D_R(u)$ is a right ascent of v. The descent identity gives $P_{v,u} = P_{vt,u}$. Since $\ell(u) - \ell(vt) = 2k$, the degree of $P_{vt,u}$ is at most $k - 1$ (and the polynomial vanishes if $vt \not\leq u$). This contradicts $\mu(v, u) \neq 0$, so $D_R(u) \subseteq D_R(v)$. Applying the same argument to u^{-1} and v^{-1} proves the left descent inclusion. For a comparable pair at gap one, $P_{v,u} = 1$, so that the coefficient is one. □

The engine uses the head and appropriate strong covers as guaranteed positive outputs. All higher-gap terms are treated only by the superset guard (C.30), even omitting Bruhat comparability. Thus a possibly spurious guard candidate may prevent an inverse deduction but cannot justify one. A white inverse is accepted only when the head, all strong-cover outputs and every guarded higher-gap candidate are already zero. Then the actual product has no positive cohomology, so the white identity makes the middle zero. Forward head and cover deductions and descent transport use only genuine positive outputs.

The strong coatoms are computed by the Coxeter lifting identity

$$\mathrm{Low}(us) = \{u\} \cup \{xs : x \in \mathrm{Low}(u),\ \ell(xs) = \ell(x) + 1\} \quad \text{when } \ell(us) = \ell(u) + 1. \tag{C.31}$$

Indeed, delete the final letter to obtain u; deleting a preceding letter yields xs, and it has the required length exactly under the displayed condition. Conversely every length-one Bruhat predecessor is a deletion in the chosen reduced word. There are 36 011 108 strong coatoms in this enumeration. An independent implementation checked (C.31) against all 459 588 reduced-word deletion covers of E_6; the E_7 implementation applies the same proved recurrence to its independently audited action table.

For efficiency, the first implementation buckets unknown lower-label sectors by fourteen left/right descent bits and uses a superset zeta transform. At middle length L, only lengths at most $L - 3$ and parity $L - 3$ are active. A left query uses $(D_L(u) \cup \{s\}, D_R(u))$; the right query is analogous. A stale unknown snapshot can only enlarge the guard. After every inverse pass, all genuine forward implications and descent transports are saturated.

C.5.2. *E_7: exhaustive closure and its independent replay.* The exact histories below start after the first forward closure:

$$\begin{aligned}
28:&\quad 46\,478\,035 \to 11\,294\,003 \to 1\,167\,970 \to 30\,575 \to 80 \to 0,\\
63:&\quad 919\,204\,701 \to 483\,936\,187 \to 150\,711\,464\\
&\quad \to 24\,409\,345 \to 1\,288\,091 \to 8\,797 \to 0,\\
36:&\quad 597\,996\,740 \to 460\,046\,750 \to 343\,090\,606\\
&\quad \to 271\,961\,463 \to 217\,098\,726 \to 154\,905\,330\\
&\quad \to 85\,546\,982 \to 29\,163\,609 \to 4\,637\,949\\
&\quad \to 289\,356 \to 2\,232 \to 0.
\end{aligned}$$

The numbers are quotient-mark ordered sectors. The productive inverse row totals for $28, 63, 36$ are respectively 806 931, 3 690 417, 5 160 850.

The independent engine stores dense raw quotient-mark matrices rather than descent-class matrices. It applies the simple transports explicitly and builds the high-gap guards by direct submask enumeration rather than a zeta transform. It reconstructs the initial zero matrices, compares them pointwise, and obtains every displayed pass count and every inverse-row total exactly. The pointwise comparisons cover 951 024 144, 4 814 559 729 and 1 572 101 136 bits for quotient sizes $28, 63, 36$. The group action and source tables are shared, separately audited inputs; the state propagation and guard algorithms are independent.

The initial number of unknown supported raw coordinates is $1\,213\,041 \cdot 16\,260 = 19\,724\,046\,660$. It decreases to zero. Over all Weyl labels the proof covers 47 203 430 400 supported label-sector coordinates.

Proposition C.16. *For each canonical tilting label in a central marked block of type E_7, the coefficient has no positive ordinary cohomology.*

Proof. Fix the cohomological degree $j > 0$. For each raw sector the parabolic sources give initial vanishing in degree j. Lemma C.15 justifies every use of the inverse white endpoint identity in the computation: the calculated set contains every possible additional summand. Thus vanishing passes through each of the forward and inverse steps above. The independent replay reaches all three nonconstant quotient orbits, including their full raw lifts. The singleton orbits were settled by white-head induction. This covers all central sectors for the chosen j, and the same computation applies to every $j > 0$. □

Finite projectivity follows in the global argument from Lemma 3.3, after the independent amplitude and finiteness hypotheses there have been checked.

C.5.3. *E_8: exact enumeration and anchored head forests.* Use the same diagram, extended by node 8 at the end of the chain. Write each element uniquely as $w = ur_q$ with $u \in W(E_7)$ and $0 \le q < 240$, where r_q are minimal representatives for the left parabolic cosets. This enumerates all 696 729 600 elements. The small transition audit checks all 1 920 right Schreier transitions on all 240 roots, the reduced parents, minimality and distinctness of the cosets. The remaining left action is generated by commuting left and right multiplication; all involution and length-change tests are checked. The all-label same-I source gather/scatter comparison supplies 171 947 976 heads, leaving 524 781 624 labels.

The full 256 raw marks have orbits of sizes $1, 120, 135$. The singleton is all white, and different-orbit endpoint sectors vanish by the actual semisimple support. The finite targets for each nontrivial orbit are every nonsource sink $(w, A, B), \quad A \in W_{D_L(w)}\backslash\Omega, \quad B \in W_{D_R(w)}\backslash\Omega$. At each length L, all shorter coefficients are the induction hypothesis. A nonsink at length L has a white descent after actual descent transport and is zero from its shorter white predecessor. For a white ascent, all terms other than the length-$L+1$ head have length below L. Hence vanishing at the middle and head is equivalent. The same-length graph theorem of the preceding finite-forcing section applies.

For clarity, a valid left head incidence has the exact conditions

$$h = sw, \quad \ell(h) = L + 1, \qquad A \cap A_h \cap \mathrm{White}_s \neq \varnothing, \quad B \cap B_h \neq \varnothing. \tag{C.32}$$

The right incidence interchanges endpoints. All intersections use full raw-index sets. A color mask merely serializes a class; it does not choose unrelated raw representatives at successive steps. An equality edge consists of two such incidences to the same marked head.

The search stores sparse forests leading to a source head or to a source/nonsink predecessor. If a component exhausts without a boundary it remains unresolved. A capped search is also unresolved. The independent verifier reconstructs every sink target from the Weyl action. It verifies each equality and boundary incidence, forms the equality components, and records each component without an ordinary boundary. A cycle alone supplies no zero.

For each component, an *ordinary boundary* means either a source head or a common head leading to a source or nonsink predecessor as above. An equality edge transports vanishing within the component by Lemma C.13.

Lemma C.17. *For the two nontrivial raw orbits, the numbers of sink targets and recorded relations, summed over the lengths* 0 *through* 120*, are as follows.*

Raw orbit	*Sink targets*	*Verified forest records*	*Unanchored component*
120	88 916 761	88 916 761	*none*
135	461 736 429	461 736 428	*one, at length* 64

In the exceptional component the verifier finds 35 419 *states. All other components have an ordinary boundary in the sense just defined.*

Proof. At length L the verifier lists the labels with $\ell(w) = L$, excludes the source labels, and takes the product of the two sets of sink descent classes for each remaining label. This reconstructs the targets before the forest records are read. A record joins two targets through a common head, or identifies an ordinary boundary for a target. The tests in (C.32), on both ends of each edge, decide which records are valid. Joining the verified edges and marking the components that contain a boundary gives the table. For the 135-orbit at length 64, the verifier also compares the entire set of targets in components without a boundary to the supplied list of residual states. Section C.6 identifies the verification programs and their per-length outputs.

Both enumerations and their independent replays cover all 121 lengths from 0 to 120. In the 135-orbit, lengths 0–63 have a boundary in every component, so induction from the identity applies successively. At length 64 every component except the specified one has a boundary. Each layer 65–120 is a proved implication conditional on all shorter lengths. These later implications were not used to assert the missing length-64 vanishing. □

C.5.4. *The diagonal coefficient at length 64.* The single unanchored component contains 35 419 sink states and 10 313 length-64 Weyl labels. We identify one explicit state in the component, then compute its coefficient geometrically. The digits in the word below are the indices of the simple reflections.

Lemma C.18. *The residual list in Lemma C.17 contains a diagonal state with Weyl element and raw phase given by*

$$w = 5423143565423143546765423143542654787654231435426543176542348765, \tag{C.33}$$
$$\ell(w) = 64, \qquad D_L(w) = D_R(w) = \{5\}, \qquad w^2 = 1.$$

The phase at both endpoints is

$$b = (0, 0, 0, 0, 0, 2, 0, 2) \pmod 4. \tag{C.34}$$

Only simple root 5 *is black at* b*, and*

$$w(b) = (0, 0, 2, 2, 0, 0, 2, 2) \neq b. \tag{C.35}$$

If t *is the raw torus element with phase* b*, its connected centralizer* $K = Z_G(t)$ *has* 56 *positive roots, Cartan determinant* 4*, and type* D_8*.*

Proof. The two endpoint marks have index 5 and black-color mask 16. The residual list and the compact word dictionary both give the word (C.33) for factor-code 359034028. To verify the root data, generate all roots from the eight simple roots by simple reflections. There are 240 roots, of which 120 are positive. Multiplying the eight-dimensional integral reflection matrices in the order (C.33) gives an involution that sends 64 positive roots to negative roots. Applying the matrix and its inverse to the simple roots gives the two descent sets. Acting on the phase modulo four gives (C.35). Among the positive roots, those with trivial value at t are the 56 positive roots of K. Removing those expressible as sums of two of these positive roots gives the simple system. Its Cartan matrix has determinant 4 and diagram D_8. The compact core verifier checks membership of the selected state in the tree. □

Lemma C.19. *Assume that the coefficients of labels of length less than* 64 *in this orbit are ordinary. The coefficient of* B_w *in the diagonal sector of Lemma C.18 has no positive ordinary cohomology.*

Proof. At this central norm, the actual diagonal marked fiber is the one-sided nilpotent incidence for K, with the second-Cartan nilpotents retained. The canonical full-graph source for the positive Rouquier standard is the w^{-1} source in the uniform local theorem. Equation (C.35) says that $w^{-1} \notin W_K$. Theorem 3.19 therefore gives zero in every ordinary degree of the actual derived output fiber of this standard coefficient. The coefficient is perfect by Proposition 3.1. Derived Nakayama makes the completed coefficient itself zero.

All canonical labels of length less than 64 in this orbit are already ordinary. In the positive Rouquier complex for (C.33), the degree-zero Bott–Samelson term contains B_w once; its other canonical summands and every term in positive Rouquier degree have shorter labels. The contravariant coefficient of the latter subcomplex lies in degrees at most -1. Its triangle with the zero standard coefficient identifies the positive cohomology of the Bott–Samelson term with zero. Thus the selected B_w coefficient has no positive cohomology. This is the application of Corollary 3.22 at the chosen raw mark. □

Proposition C.20. *For each canonical tilting label in a central marked block of type* E_8*, the coefficient has no positive ordinary cohomology.*

Proof. Induct on the Weyl length. The identity label is a parabolic source. For a given length, Lemma C.13 propagates vanishing from every ordinary boundary in Lemma C.17. This proves the assertion for all lengths up to 63. At length 64 there is one component in the 135-orbit still to treat. Lemma C.19 applies since its shorter-length hypothesis has now been proved. It gives vanishing at the state in Lemma C.18.

The compact core certificate has 35 418 equality edges. Its independent replay generates 20 717 explicit root matrices, rechecks all lengths and descent sets, builds the full raw descent-class partitions, and verifies the exact incidences (C.32). The independent root-matrix verification checks every edge and confirms that (C.33)–(C.34) belongs to the tree. Hence the anchor propagates to every state of the component.

This completes length 64. Increasing-length induction now discharges each premise of the independently checked layers 65–120. Thus both nontrivial raw orbits, as well as the singleton, are ordinary. This proves positive-degree vanishing in every sector. □

The all-length forest contains 550 653 189 records for 550 653 190 targets; the one missing zero premise is precisely supplied by the preceding geometric anchor. The full forests are regenerable from the algorithms in Section C.6. Together with the marked centralizer formula of Proposition 3.24, Proposition C.20 supplies the E_8 step of the assembly in Section 3.7.

C.6. **Verification of the finite calculations.** The finite calculations use the parabolic source theorems, the marked identities of Proposition 3.18, the positivity of ordinary canonical decompositions, and the diagonal anchor of Corollary 3.22. The programs enumerate their finite inputs and check their consequences. All arithmetic that affects a conclusion is integral or exact rational arithmetic.

C.6.1. *Data and implications.* A Weyl element is represented by its action on the roots or by a factorization whose transitions have been checked on all roots. Its length is the number of positive roots sent to negative roots. The two descent sets are computed independently by left and right simple multiplication. A raw mark records the entire torus phase in the prescribed central fiber; a binary color is only a derived description of its white simple roots. Descent classes are generated from the actual simple actions on these raw marks.

Fix a central norm and let C be the common completed even central ring at that norm, with residue field k. For an ordinary degree $j > 0$, write $m_{w,a,b} = \dim_k\big(H^j\mathcal{N}(B_w)_{a,b} \otimes_C k\big)$. The endpoint identities are C-linear, even when they transport an odd Cartan action. The coefficient modules are finite over C, so these dimensions are finite and Nakayama detects their vanishing. They are nonnegative integers. All finite arguments apply uniformly to them, so they make no assumption about an upper bound on j. There are four permitted operations.

(1) A proved parabolic source gives zero for each of its positive canonical summands. Source generation retains the end parabolic; an unrestricted convolution closure would not be a valid replacement.
(2) A white endpoint identity and a positive canonical product give a linear equality with nonnegative coefficients. If one side is zero, every positive term on the other side is zero.
(3) A black endpoint identity transports the two marked rows with the retained averaging factor. The operator C_s is not cancelled.
(4) At a minimal unresolved length, the shorter terms disappear. Two white predecessors of a common head then have equal scalar detectors. An entire equality component vanishes once one vertex has a proved zero boundary, including a geometric diagonal anchor.

These operations are consequences of isomorphisms of ordinary complexes before taking dimensions. Taking a single residue field or a single positive degree cannot introduce a negative coefficient. Finite generation then turns the resulting pointwise assertions into vanishing of cohomology modules.

C.6.2. *Classical groups and small exceptional groups.* The uniform classical arguments are proofs in arbitrary rank; finite tests are used only to check their implementation conventions. The nonstable-middle root check covers 3 541 minimal positions in outer A_2, A_3, A_4, A_5, D_4. Among them, 921 positions are fixed by the involution, including 442 with nonstable middle parabolic. The checks compare the full root multisets and the strict weight signs in the common-Levi proof.

The outer A_3 and A_4 certificates enumerate every Weyl element. Their initial source supports contain 22 of 24 labels and 92 of 120 labels. The explicit different-end source words supply the missing anchors; safe endpoint padding completes the A_4 case. An independent implementation uses the subword criterion for Bruhat intervals to check those padding steps.

The root-trivial G_2, F_4, E_6 calculations enumerate the full groups, the central marks, and the marked endpoint classes. They use the geometric anchor only under the shorter-length hypothesis. For G_2 and F_4, a second implementation keeps the individual raw states throughout, so the result does not depend on the primary descent-class compression. The data and numerical totals are given in Section C.3.

C.6.3. *Outer E_6: exact endpoint rows.* The root-permutation model reconstructs the 72 roots and all 51 840 Weyl elements. Deleting letters in reduced words and checking the length drop gives all 459 588 strong Bruhat covers. The initial source set of 26 242 labels is reconstructed from 17 167 direct source words and 9 075 safe transfers whose entire lower ideals have already been accepted.

The endpoint calculation requires 481 490 Kazhdan–Lusztig coefficients, including the zero values. The independent recurrence checks all of them; 84 837 are positive. The event verifier then reconstructs the coefficient vectors of 240 317 ordered endpoint events. After this elimination, it writes the 90 782 residual rows on 15 663 variables.

The enlarged common-Levi source family has 19 892 labels. It is constructed independently of the candidate residual solution. Adding its zero conditions reduces the active residual system by exact rational elimination. A one-sign row forces every variable in its support to zero. A two-term row with opposite signs gives a positive ratio. Ratios propagate in a graph; an inconsistent ratio around a cycle forces that component to zero. After merging these components, all rows are reduced again. The successive active counts are $15\,407 \longrightarrow 115 \longrightarrow 33 \longrightarrow 1 \longrightarrow 0$. Every one of the 518 400 marked coefficient positions is therefore ordinary. No primitive-object interpretation of an intermediate positive vector is needed.

C.6.4. *E_7: bounds for the lower terms.* The E_7 model has 2 903 040 Weyl elements. Its independently checked source set contains 1 689 999 labels, leaving 1 213 041 labels for endpoint propagation. All 256 raw marks are retained. Their Weyl-orbit sizes are 1, 1, 63, 63, 56, 72. The constant-color cases are immediate. The three nonconstant quotient calculations have mark sizes $63, 28, 36$ and lift multiplicities $2, 4, 4$, respectively.

Higher-gap terms in the canonical product are controlled by the support bound in Lemma C.15. The guard retains every label with the necessary descent inclusions and odd length gap at least three. It deliberately need not decide the remaining Bruhat relation or the numerical value of the coefficient: this makes the set larger, so absence of an unresolved guarded term is sufficient. Strong-cover terms are enumerated exactly and have coefficient one. Thus a white inverse step is accepted only when every possible other term is already zero.

The primary implementation uses descent-class bitsets and a superset zeta transform. The independent implementation expands to raw marks and uses direct submask guards. It reconstructs the initial zero array and every subsequent guard pass. Both implementations reach zero unresolved positions in each of the three cases and reproduce the pass histories in the main text. In the full raw system, 19 724 046 660 initially unresolved coordinate positions become zero. Neither implementation treats exclusion from a possible maximum as an established zero.

C.6.5. *E_8: the full forest and the portable core.* Use the parabolic factorization $|W(E_8)| = |W(E_7)| \cdot 240 = 696\,729\,600$. Each of the 240 distinguished cosets is specified by a reduced word. The 1 920 simple right transitions are checked on all 240 roots. The retained E_7 action table supplies multiplication inside each coset. The missing simple left action is generated from these words and checked to be an involution with length difference one on every Weyl element.

The source array is a byte for each element, with its eight bits recording the eight proper maximal end parabolics separately. The primary generator gathers permitted predecessors. An independent generator scatters permitted successors and compares every family byte. The resulting counts are $\#\text{source} = 171\,947\,976$, $\#\text{remaining} = 524\,781\,624$. The 256 raw marks have Weyl orbits of sizes $1, 120, 135$. The all-white orbit is immediate. After the descent reduction, the two other orbits have 88 916 761 and 461 736 429 sink targets.

At each length the verifier reconstructs every target and every descent class from the raw phase actions. A forest record consists of four unsigned 64-bit integers: the record kind, first sink, second sink and head. Kind 0 asserts equality of two sinks through a common white head. Kind 1 gives a source head. Kind 2 gives an ordinary predecessor boundary. Adjacency, raw-class intersections, and the shorter-length premises are checked for every record. A union-find calculation then tests whether each component meets a zero boundary. The order of the records has no mathematical significance.

All layers of the 120-mark orbit close. In the 135-mark orbit, the layers below 64 close, and length 64 has exactly one unanchored component of 35 419 states. All later layers close conditionally on the shorter-length assertion. These conditions are discharged in the length induction of Section C.5, using the geometric anchor at length 64.

The whole finite forest has 550 653 189 records for 550 653 190 targets: the difference is the one missing zero anchor. Its compact core consists of 35 418 edges. The ancillary package includes this tree, its complete state list, the raw mark actions, and a word dictionary. A separate verifier reconstructs 20 717 integer root matrices from the dictionary and checks all core lengths, left and right descents, mark intersections, and connectivity. It reads no full E_7 or E_8 action array. It also checks that the selected state used by the geometric proof is a vertex of the verified tree.

For that selected state, the word in the main text has length 64, is an involution, and has both descent sets equal to $\{5\}$. The raw phase is $(0, 0, 0, 0, 0, 2, 0, 2)$ and its image is $(0, 0, 2, 2, 0, 0, 2, 2)$. Independent constructions

using simple-root integer matrices and the Euclidean D_8 plus half-spin model of the E_8 roots give the same result. The root centralizer has 56 positive roots and type D_8. Thus the phase is moved, exactly as required in Proposition 3.21.

C.6.6. *The accompanying source package.* The source archive contains the verification programs, input data, and certificates in `anc/`. Its `README.txt` gives the commands in dependency order, the input files, and the expected outputs. There are three separate checks a reader can perform. First, `run_checks.py -outer-e6` recomputes the outer E_6 result of Theorem C.14 and the tree and root computations used in Lemmas C.18 and C.19. Second, `check_e7_outputs.py` compares the pass histories used in Proposition C.16, and `check_e8_outputs.py` compares the generator and verifier records for all 242 orbit-length pairs in Lemma C.17. This also identifies the full list of residual keys with the input to the compact tree verifier. Third, the full regeneration commands construct the large binary tables and check each guard or forest incidence against the Weyl action. Only this third procedure replays all the individual relations in the two large calculations.

The portable checks require Python 3 and a C++17 compiler. Assertions must remain enabled. The full exceptional runs should be executed serially when memory is limited. Rebuilding the full E_8 forest requires approximately 17.6 GB for forest records and 3.5 GB for the action and source caches, besides working memory. These large regenerable files are omitted from the source package; their generators, input data, verification programs, and recorded hashes are included. The compact core can be checked independently without generating them.

For this manuscript we replayed the portable core, the selected root and mark calculation, the small Schreier model, the small outer sources, and all the outer-E_6 source, coefficient, event, and positive-elimination verifiers. We also checked the split-E_6 closure and the raw-E_7 mark orbits. The core replay reconstructed all 35 419 vertices and checked all 35 418 edges. Starting from the distributed source, we regenerated the full E_7 group table and the initial zero arrays for the 28-, 63- and 36-point mark orbits. The independent dense audit compared every initial entry. On each orbit the two guard programs reproduced the recorded numbers of unknowns at every pass and ended with zero. For E_8 we regenerated the complete left-action and source tables and the count of sink targets. At length 64 of the 135-orbit the generator produced 14 812 036 records. The independent verifier reconstructed 14 812 037 targets, found the one residual component, and compared its 35 419 states to the supplied list. Re-extracting the compact core gave the supplied tree. The three large tables and the binary forest for this layer match the recorded SHA-256 digests. For the other E_8 layers the package contains the records of the earlier full runs. The README and the dated verification record give the measured running times, memory use and disk requirements. A matching digest certifies file identity alone. The mathematical conclusions require the incidence checks, the geometric vanishing at the chosen state and the length induction of Section C.5.

References

[ABV92] Jeffrey Adams, Dan Barbasch, and David A. Vogan, *The Langlands classification and irreducible characters for real reductive groups*, Progress in Mathematics, vol. 104, Birkhäuser, Boston, 1992, doi:10.1007/978-1-4612-0383-4.

[AG02] Luchezar L. Avramov and Daniel R. Grayson, *Resolutions and cohomology over complete intersections*, Computations in Algebraic Geometry with Macaulay 2 (David Eisenbud, Daniel R. Grayson, Michael Stillman, and Bernd Sturmfels, eds.), Algorithms and Computation in Mathematics, vol. 8, Springer, Berlin, 2002, https://macaulay2.com/Book/ComputationsBook/chapters/completeIntersections/ci-wrapper.pdf, pp. 131–178.

[AG15] Dima Arinkin and Dennis Gaitsgory, *Singular support of coherent sheaves, and the geometric Langlands conjecture*, Selecta Math. (N.S.) **21** (2015), no. 1, 1–199, Theorem numbering follows the October 31, 2014 author version; doi:10.1007/s00029-014-0167-5.

[AK26] Jeffrey Adams and Tasho Kaletha, *Discrete series L-packets for real reductive groups*, Essential Number Theory **5** (2026), no. 1, 141–219, doi:10.2140/ent.2026.5.141.

[BB76] Andrzej Białynicki-Birula, *Some properties of the decompositions of algebraic varieties determined by actions of a torus*, Bull. Acad. Polon. Sci. Sér. Sci. Math. Astronom. Phys. **24** (1976), no. 9, 667–674.

[BB81] Alexander Beilinson and Joseph Bernstein, *Localisation de $\mathfrak{g}$-modules*, C. R. Acad. Sci. Paris Sér. I Math. **292** (1981), no. 1, 15–18.

[BB05] Anders Björner and Francesco Brenti, *Combinatorics of Coxeter groups*, Graduate Texts in Mathematics, vol. 231, Springer, Berlin, Heidelberg, 2005, doi:10.1007/3-540-27596-7.

[BCHN25] David Ben-Zvi, Harrison Chen, David Helm, and David Nadler, *Between coherent and constructible local Langlands correspondences*, The Langlands Program, Proc. Sympos. Pure Math., vol. 112.2, Amer. Math. Soc., 2025, Numbering follows arXiv:2302.00039v1; doi:10.1090/pspum/112.2/02067, pp. 421–470.

[Bez16] Roman Bezrukavnikov, *On two geometric realizations of an affine Hecke algebra*, Publ. Math. Inst. Hautes Ètudes Sci. **123** (2016), 1–67, doi:10.1007/s10240-015-0077-x; arXiv:1209.0403.

[BGS96] Alexander Beilinson, Victor Ginzburg, and Wolfgang Soergel, *Koszul duality patterns in representation theory*, J. Amer. Math. Soc. **9** (1996), no. 2, 473–527, doi:10.1090/S0894-0347-96-00192-0.

[BHS19] Christophe Breuil, Eugen Hellmann, and Benjamin Schraen, *A local model for the trianguline variety and applications*, Publ. Math. Inst. Hautes Ètudes Sci. **130** (2019), 299–412, doi:10.1007/s10240-019-00111-y.

[BN13] David Ben-Zvi and David Nadler, *Loop spaces and representations*, Duke Math. J. **162** (2013), no. 9, 1587–1619, doi:10.1215/00127094-2266130; arXiv:1004.5120v3.

[BN26] ———, *Potent categorical representations*, Proceedings of the International Congress of Mathematicians 2026, vol. 4, SIAM, 2026, doi:10.1137/25M1806934, pp. 471–491.

[Bor79] Armand Borel, *Automorphic L-functions*, Automorphic Forms, Representations and L-Functions, Part 2 (Armand Borel and William Casselman, eds.), Proc. Sympos. Pure Math., vol. 33, American Mathematical Society, Providence, RI, 1979, doi:10.1090/pspum/033.2/546608, pp. 27–61.

[Bor91] ———, *Linear algebraic groups*, second ed., Graduate Texts in Mathematics, vol. 126, Springer, New York, 1991.

[Bou02] Nicolas Bourbaki, *Lie groups and Lie algebras, chapters 4–6*, Elements of Mathematics, Springer, Berlin, 2002, doi:10.1007/978-3-540-89394-3.

[BR12] Roman Bezrukavnikov and Simon Riche, *Affine braid group actions on derived categories of Springer resolutions*, Ann. Sci. Èc. Norm. Supér. (4) **45** (2012), no. 4, 535–599, https://www.numdam.org/item/ASENS_2012_4_45_4_535_0/.

[BV21] Roman Bezrukavnikov and Kari Vilonen, *Koszul duality for quasi-split real groups*, Invent. Math. **226** (2021), no. 1, 139–193, doi:10.1007/s00222-021-01050-4; arXiv:1510.08343v7.

[Che23] Harrison Chen, *Categorical cyclic homology and filtered $\mathcal{D}$-modules on stacks: Koszul duality*, 2023, Theorem numbering in arXiv:2301.06949v3.

[CM93] David H. Collingwood and William M. McGovern, *Nilpotent orbits in semisimple Lie algebras*, Van Nostrand Reinhold Mathematics Series, Van Nostrand Reinhold, New York, 1993.

[DT26a] Gurbir Dhillon and Jeremy Taylor, *Tame local Betti geometric Langlands*, 2026, arXiv:2501.14157v2, 5 October 2026; numbering follows this version.

[DT26b] ———, *The universal monodromic Arkhipov–Bezrukavnikov equivalence*, 2026, arXiv:2501.14156v2, 5 October 2026; numbering follows this version.

[EL24] Lawrence Ein and Robert Lazarsfeld, *Lectures on the syzygies and geometry of algebraic varieties*, 2024, Authors' preliminary manuscript, 4 May 2024.

[EW14] Ben Elias and Geordie Williamson, *The Hodge theory of Soergel bimodules*, Ann. of Math. (2) **180** (2014), no. 3, 1089–1136, doi:10.4007/annals.2014.180.3.6; arXiv:1212.0791.

[GR70] Hans Grauert and Oswald Riemenschneider, *Verschwindungssätze für analytische Kohomologiegruppen auf komplexen Räumen*, Invent. Math. **11** (1970), 263–292, doi:10.1007/BF01403182.

[Hai18] Thomas J. Haines, *Dualities for root systems with automorphisms and applications to non-split groups*, Represent. Theory **22** (2018), 1–26, doi:10.1090/ert/512.

[Har66a] Harish-Chandra, *Discrete series for semisimple Lie groups. II. Explicit determination of the characters*, Acta Math. **116** (1966), 1–111, doi:10.1007/BF02392813.

[Har66b] Robin Hartshorne, *Residues and duality*, Lecture Notes in Mathematics, vol. 20, Springer, Berlin, 1966, Lecture notes of a seminar on the work of A. Grothendieck.

[Hat02] Allen Hatcher, *Algebraic topology*, Cambridge University Press, Cambridge, 2002, https://pi.math.cornell.edu/~hatcher/AT/AT.pdf.

[HL21] Daniel Halpern-Leistner, *Derived Θ-stratifications and the D-equivalence conjecture*, 2021, arXiv:2010.01127v2; numbering follows this version.

[HLS10] Megumi Harada, Gregory D. Landweber, and Reyer Sjamaar, *Divided differences and the Weyl character formula in equivariant K-theory*, Math. Res. Lett. **17** (2010), no. 3, 507–527, doi:10.4310/MRL.2010.v17.n3.a10.

[Hoc61] Gerhard Hochschild, *Cohomology of algebraic linear groups*, Illinois J. Math. **5** (1961), no. 3, 492–519, doi:10.1215/ijm/1255630894.

[Hum90] James E. Humphreys, *Reflection groups and Coxeter groups*, Cambridge Studies in Advanced Mathematics, vol. 29, Cambridge University Press, Cambridge, 1990, doi:10.1017/CBO9780511623646.

[IY25] Andrei Ionov and Zhiwei Yun, *Tilting sheaves for real groups and Koszul duality*, Compos. Math. **161** (2025), no. 7, 1698–1777, Numbering follows arXiv:2301.05409v2; doi:10.1112/S0010437X25007134.

[Kas15] Hisashi Kasuya, *Extended simplicial rational Nomizu's theorem*, 2015, arXiv:1410.3176v4.

[KL79] David Kazhdan and George Lusztig, *Representations of Coxeter groups and Hecke algebras*, Invent. Math. **53** (1979), 165–184, doi:10.1007/BF01390031.

[Kot86] Robert E. Kottwitz, *Stable trace formula: elliptic singular terms*, Math. Ann. **275** (1986), 365–399.

[KS90] Masaki Kashiwara and Pierre Schapira, *Sheaves on manifolds*, Grundlehren der mathematischen Wissenschaften, vol. 292, Springer, Berlin, 1990.

[Lan89] Robert P. Langlands, *On the classification of irreducible representations of real algebraic groups*, Representation Theory and Harmonic Analysis on Semisimple Lie Groups (Paul J. Sally and David A. Vogan, eds.), Math. Surveys Monogr., vol. 31, American Mathematical Society, Providence, RI, 1989, Originally circulated in 1973; doi:10.1090/surv/031/03, pp. 101–170.

[Lur17] Jacob Lurie, *Higher algebra*, 2017, Author's version, September 2017; https://www.math.ias.edu/~lurie/papers/HA.pdf.

[LV83] George Lusztig and David A. Vogan, *Singularities of closures of K-orbits on flag manifolds*, Invent. Math. **71** (1983), 365–379, doi:10.1007/BF01389103.

[MFK94] David Mumford, John Fogarty, and Frances Kirwan, *Geometric invariant theory*, third ed., Ergebnisse der Mathematik und ihrer Grenzgebiete (2), vol. 34, Springer, Berlin, 1994.

[MUV92] Ivan Mirković, Tohru Uzawa, and Kari Vilonen, *Matsuki correspondence for sheaves*, Invent. Math. **109** (1992), 231–245, doi:10.1007/BF01232026.

[Rou04] Raphaël Rouquier, *Categorification of the braid groups*, 2004, arXiv:math/0409593.

[Soe01] Wolfgang Soergel, *Langlands' philosophy and Koszul duality*, Algebra—Representation Theory (Constanţa, 2000), NATO Sci. Ser. II Math. Phys. Chem., vol. 28, Kluwer, Dordrecht, 2001, doi:10.1007/978-94-010-0814-3_17; see also the author's correction to p. 405, pp. 379–414.

[Soe07] ———, *Kazhdan–Lusztig-polynome und unzerlegbare Bimoduln über Polynomringen*, J. Inst. Math. Jussieu **6** (2007), no. 3, 501–525, doi:10.1017/S1474748007000023.

[Spr98] Tonny A. Springer, *Linear algebraic groups*, second ed., Progress in Mathematics, vol. 9, Birkhäuser, Boston, 1998.

[Ste65] Robert Steinberg, *Regular elements of semi-simple algebraic groups*, Publ. Math. Inst. Hautes Ètudes Sci. **25** (1965), 49–80, https://www.numdam.org/item/PMIHES_1965__25__49_0/.

[Ste68] ———, *Endomorphisms of linear algebraic groups*, Memoirs of the American Mathematical Society, vol. 80, American Mathematical Society, Providence, RI, 1968, doi:10.1090/memo/0080.

[Ste75] ———, *On a theorem of Pittie*, Topology **14** (1975), no. 2, 173–177.

[Tay25] Jeremy Taylor, *Universal monodromic tilting sheaves*, Represent. Theory **29** (2025), 871–892, Numbering follows arXiv:2305.03033v7; doi:10.1090/ert/705.

[Tay26] ———, *The tilting property of Whittaker averaged central sheaves*, 2026, arXiv:2606.25387v1.

[The26] The Stacks Project Authors, *The Stacks Project*, 2026, Individual tags are specified at the point of use; https://stacks.math.columbia.edu.

[Vog82] David A. Vogan, *Irreducible characters of semisimple Lie groups IV. Character-multiplicity duality*, Duke Math. J. **49** (1982), no. 4, 943–1073.

[Zhu17] Xinwen Zhu, *An introduction to affine Grassmannians and the geometric Satake equivalence*, Geometry of Moduli Spaces and Representation Theory, IAS/Park City Math. Ser., vol. 24, American Mathematical Society, Providence, RI, 2017, arXiv:1603.05593v2, pp. 59–154.