A PROOF OF THE HLRV CONJECTURE

QIYUE TANG

Abstract. We prove the Hausel–Letellier–Rodriguez-Villegas conjecture on the mixed Hodge polynomials of generic semisimple GL_n character varieties.

hexagon:2611.00441v1

Contents

1. Introduction 4
1.1. Character varieties and the mixed Hodge conjecture 4
1.2. Counting formulas and filtration comparisons 4
1.3. Stable pairs and the proof 5
2. The character-variety formula 6
3. The ordinary-curve construction 8
3.1. Notation and coefficients 9
3.2. Moduli and the two support categories 10
3.3. Established inputs 10
4. Critical models and their orientations 11
4.1. The framed critical model 11
4.2. The full deformation complex and its trace 13
4.3. Closed forms and the determinant square 15
4.4. The fixed positive quotient 16
4.5. Support gluing and the derivative 17
4.6. The comparison over the full parameter ring 20
5. Monodromic coefficients and localization 21
5.1. An explicit coefficient projector 21
5.2. The integral localization map 22
5.3. One-sided constructibility on finite-type windows 23
6. Harder–Narasimhan restrictions 23
6.1. The framed heart and its numerical bounds 24
6.2. Complete coherent flags 25
6.3. The closed-stratum counit 26
6.4. Vacuum removal and finite wall transport 28
6.5. The comparison arrow and its endpoints 29
7. The Hecke relation on the comparison arrow 32
7.1. Directed bounded presentations 32
7.2. Choices over the whole marked curve 33
7.3. The two derived words 33
7.4. A proper carrier and its supported class 35
7.5. Lifting the class to the entire arrow 36
7.6. Independence of the auxiliary presentation 38
7.7. Finite-band vanishing 39

Date: November 14, 2026.
2020 *Mathematics Subject Classification.* 14D20, 14C30, 14H60, 05E05.
Key words and phrases. Character variety, mixed Hodge polynomial, Higgs bundle, Macdonald polynomial.

8. The positive source identity 40
8.1. The unframed input and its hypotheses 40
8.2. Finite spectral addition and reduced spectral curves 42
8.3. The Hodge calculation at a nodal stratum 43
8.4. Purity and reconstruction of the source 46
9. The scalar critical primitive 46
9.1. The critical scheme on arbitrary ordinary bases 47
9.2. The normal Hessian and its determinant root 48
9.3. Pushforward and the scalar shift 48
10. Scalar packets and the global physical identity 49
10.1. Detection of the comparison error 49
10.2. The physical product over coprime packets 50
10.3. Stratified descent and induction 53
10.4. Restriction to the original nil coefficient 54
11. The fixed loci and their cohomology 55
11.1. The fixed functor 56
11.2. Closed attracting filtrations and punctual purity 57
11.3. Indices and the diagonal hook identity 58
11.4. Moving support and finite symmetric actions 60
11.5. Global purity and proof of the relative formula 61
11.6. Properness of the positive spectral map 62
11.7. The whole nil locus 63
12. Hodge recovery and the character-variety formula 64
12.1. The identity in full Hodge classes 64
12.2. Recovering perverse-graded Hodge structures 65
12.3. The hook expression 65
12.4. Proof of Theorem 3.1 66
13. Root curves and the positive comparison 67
13.1. Numerical classes and a positive length 67
13.2. A coarse positive quotient and bounded windows 68
13.3. Physical maps and their orientations on a root curve 70
13.4. The moving root Hecke correspondence 71
14. Positive root Hitchin support and canonical primitives 73
14.1. Full support in every genus 73
14.2. Active profiles and the target Euler characteristic 76
14.3. The root and coarse surface halves 76
14.4. The critical scheme and the quadratic normal bundle 78
15. The global root Fock calculation 80
15.1. Fixed objects as finite quotients 80
15.2. Local indices, coordinate changes, and purity 81
15.3. Baseline and the global series 83
16. The football calculation and the Cauchy identity 84
16.1. Vanishing of higher-rank canonical primitives 85
16.2. The exact kernel and its signs 86
17. The punctual root vertex 87
17.1. A finite character and its root-invariant index 87
17.2. Stable alphabets and the normalized series 89
17.3. The LLT expansion and the one-letter value 90
18. Supersymmetrization at fixed paired heights 91
18.1. Signed letters and the finite polynomial 91

19. The signed lattice and its intertwiners 93
19.1. The five matrices 93
19.2. The exact finite path model 94
20. Fusion with its full null submodule 96
20.1. A rank-one identity for arbitrary incoming states 96
20.2. Moving an external row past a fused pair 97
21. The capacity of an arbitrary nested block 99
21.1. The raw embedding and its kernel 100
21.2. The zero-height boundary 101
22. The right ramp and all dominance inequalities 102
22.1. Successive disappearance on the right 102
22.2. A last auxiliary color when no zero plateau is present 103
22.3. A rectangle merged into the zero-height color 104
22.4. Every prefix and dominance support 105
23. Identification with the modified Macdonald polynomial 105
24. The Chern and perverse filtrations for arbitrary flags 106
24.1. Normalizations and the statement 107
24.2. Auxiliary markings and stability 107
24.3. A universal sheaf and generators on a residue fiber 108
24.4. Residue geometry in every genus 110
24.5. The Hecke action and its reduction 111
24.6. Relative Lefschetz theory and low genus 114
24.7. The zero-residue fiber and its centered shift 115
24.8. Blockwise Gysin transfer 116
25. The weight filtration at finite-order unitary parameters 117
25.1. An algebraic finite-cover map on the entire Betti stack 118
25.2. The homogeneous weights of orbifold tautological classes 118
25.3. The coarse/root comparison is triangular 119
26. Variation of all generic eigenvalues 120
26.1. A relative full-flag family at an identity puncture 120
26.2. Uniform puncture and walk strata 121
26.3. Localization in mixed Hodge modules 122
26.4. Relative Springer sign extraction 122
27. Generic torsion parameters and the required degree 123
28. Recovery of the mixed Hodge polynomial 124
28.1. The full Hodge identity 124
28.2. Separation of the two degrees 125
28.3. The hook substitution 126
28.4. Proof of the main theorem 126
References 127

1. Introduction

We prove the Hausel–Letellier–Rodriguez-Villegas formula for the compactly supported mixed Hodge polynomial of a generic semisimple GL_n character variety. Theorem 2.2 treats every genus, rank, multiplicity profile, and generic determinant component. It identifies the weight and cohomological gradings with the two variables of the Macdonald generating function. The proof also gives the independence of the polynomial from the generic eigenvalues and the nonnegativity of its coefficients in the prescribed sign convention.

1.1. **Character varieties and the mixed Hodge conjecture.** The geometry underlying the problem begins with Hitchin's study of the self-duality equations and the associated algebraically completely integrable system [31, 30]. The moduli of Higgs bundles carries a proper map recording the characteristic polynomial of the Higgs field. Its general fibers are abelian varieties, and its singular fibers encode the topology of the total space. Nonabelian Hodge theory relates this space to a moduli space of local systems; Simpson developed both the correspondence and the algebraic moduli theory [62, 63]. Prescribed local monodromy at punctures corresponds to parabolic data on the Higgs side.

The two algebraic structures give different filtrations on the same underlying cohomology. The Betti moduli space has Deligne's mixed Hodge structure, whereas the Hitchin morphism gives a perverse filtration on Higgs-bundle cohomology. Even when the latter cohomology is pure, the weight filtration on the character variety contains additional information. Hausel and Thaddeus connected the Hodge theory of these moduli spaces with mirror symmetry and Langlands duality [29]. This supplied a geometric setting for comparisons between explicit invariants of different character varieties.

Hausel and Rodriguez-Villegas formulated an explicit mixed Hodge polynomial conjecture for the character variety with primitive central twist [27]. Their finite-field character calculation, together with Katz's point-counting theorem in the appendix to that paper, determined its E-polynomial. They verified the full mixed Hodge formula in rank two and formulated the curious hard Lefschetz symmetry suggested by the answer. The partition sum contains a product indexed by the boxes of a Young diagram; its factors involve the arm and leg lengths and the genus of the curve.

Hausel, Letellier, and Rodriguez-Villegas extended this formula to punctured curves with generic semisimple local monodromies [23]. Each puncture contributes an independent alphabet and a modified Macdonald polynomial. Pairing the plethystic logarithm with the complete symmetric functions extracts the prescribed eigenvalue multiplicities. Their arithmetic harmonic analysis establishes the E-polynomial formula and a related quiver-variety calculation. Conjecture 1.2.1 of [23] predicts the full mixed Hodge polynomial, its eigenvalue independence, and the resulting positivity. Section 2 fixes the genericity condition and all symmetric-function conventions used here.

The subsequent work of Hausel, Letellier, and Rodriguez-Villegas proved connectedness of every nonempty generic character variety and identified a specialization of the kernel with the Kac polynomial of a comet-shaped quiver [24]. Their cohomological interpretation of Kac polynomials and quiver Donaldson–Thomas invariants established nonnegative coefficients for those invariants [25]. The mixed Hodge conjecture connects this quiver calculation with the pure compactly supported part of the character variety.

1.2. **Counting formulas and filtration comparisons.** Several parts of the generating function have been established by different methods. Schiffmann counted indecomposable vector bundles and stable Higgs bundles over finite fields [60]. Mozgovoy and Schiffmann developed counts of semistable twisted Higgs bundles and type A quiver bundles [54]. Mellit proved the polynomiality and integrality of the Hausel–Letellier–Villegas kernels [47]. He then obtained the predicted

Poincaré polynomials in the central-twist case [46] and for generic semisimple punctured character varieties [48]. In the punctured calculation, affine Springer fibers and local factorization explain the occurrence of modified Macdonald polynomials.

The E-polynomial and Poincaré polynomial are two specializations of the mixed Hodge polynomial. The former takes an alternating sum over cohomological degrees; the latter sums over weights. The full formula requires the joint distribution of these two degrees. Mellit's survey [49, Sections 4.8–4.9] describes these specializations and records the two-variable formula as the remaining conjecture.

The $P = W$ conjecture of de Cataldo, Hausel, and Migliorini gives a geometric interpretation of the weight filtration in terms of the Hitchin map. They proved it in rank two [17]. De Cataldo, Maulik, and Shen then established it in genus two for arbitrary rank and coprime degree, using degenerations from moduli of sheaves on abelian surfaces [18]. Maulik and Shen proved it for coprime GL_n Higgs moduli in every rank on curves of genus at least two [45]. Hausel, Mellit, Minets, and Schiffmann constructed a proof using surface Hecke operators and a deformation of the algebra of polynomial Hamiltonian vector fields [26]. Their operator construction applies in every genus; the parabolic geometric application in Section 8 of that work fixes genus greater than one and complete flags. The argument below verifies the geometric inputs for all genera and passes from complete to partial flags by regular Gysin maps and flag integration. Maulik, Shen, and Yin obtained another proof of coprime GL_n $P = W$ from Fourier transforms for compactified Jacobian fibrations [43].

Tautological classes provide the connection between these methods. Markman's diagonal-class calculation gives generators for the cohomology of moduli spaces of sheaves on symplectic surfaces [42]. Shende determined the weights of the universal character-variety classes [61]. Their normalized Chern indices can therefore be compared with both perverse degree and Betti weight. Borri has also realized the twisted character variety as a component of the character variety of a root curve and used that realization to compute tautological weights [6]. In this paper the finite-cover argument is performed on the full orbifold local-system stack and retains the isotropy summands at every puncture.

Mellit's toric stratifications prove Hodge–Tate cohomology and curious hard Lefschetz for generic semisimple character varieties [50, Theorems 1.5–1.6]. His Grothendieck–Springer construction also permits extraction of arbitrary semisimple multiplicities from an auxiliary complete flag. We use the corresponding strata and localization maps in families over all the generic eigenvalue parameters. Su's work proves the weak geometric $P = W$ statement for very generic character varieties, where at least one local conjugacy class is regular semisimple, through their cell decompositions and dual boundary complexes [67]. These results describe complementary parts of the geometry at the Hitchin and Betti boundaries.

1.3. **Stable pairs and the proof.** The curve-counting approach has a prior formulation in the work of Chuang, Diaconescu, and Pan on BPS states and $P = W$ [14]. Chuang, Diaconescu, Donagi, and Pantev developed its parabolic version using stable pairs on local root stacks [13]. They proposed the refined BPS identification with Hitchin-theoretic invariants and the local Macdonald factors needed for the HLRV formula. Carlsson and Rodriguez-Villegas studied related symmetric-function identities through vertex operators [10]. The Hilbert-scheme calculation in [13] already proves a related K-theoretic Macdonald identity; the cohomological stable-pair comparison is the additional step used here. Monavari computed local stable-pair invariants using double nested Hilbert schemes [52] and developed the refined local Donaldson–Thomas theory of curves [53].

We carry out the cohomological comparison and local calculation needed for this program. The comparison begins with two Calabi–Yau threefolds over an ordinary curve: one has a sufficiently positive line-bundle direction, and the other has the canonical and trivial directions. A quotient of

the canonical-plus-trivial bundle produces the quadratic potential connecting them. We construct its oriented critical comparison, including families over nonreduced bases. Filtered extension correspondences give the stability-change maps, and the Hecke action is constructed on the whole comparison arrow. The support argument separates the one-packet contribution from products of smaller ranks. Its proof retains the finite permutation action needed for the symmetric-power formula.

The other calculation uses the attracting loci of stable-pair moduli. Partitions describe the transverse monomial ideals, and reverse plane partitions record modifications along the curve. The deformation character determines each shift and Tate twist. Closed attracting filtrations and purity recover the resulting cohomology objects. Section 3 gives this ordinary-curve argument in full, including the central-twist mixed Hodge formula and the separation of its two gradings.

For a root curve, the construction keeps rank, Euler characteristic, and every root-character multiplicity. The primitive contribution is obtained in the Euler characteristic corresponding to the Betti parameters. The local coefficient at a root marking is initially an independently defined symmetric function. Its two-pointed genus-zero pairing gives the Cauchy identity. A colored local calculation gives finite capacity, boundary invisibility, and dominance inequalities for every prefix. These identities and the value at the one-letter alphabet characterize the modified Macdonald polynomial.

The final comparison uses normalized universal Chern and flag classes. The Hecke algebra identifies their index filtration with centered perversity, including the zero-residue and partial-flag specializations. At finite-order unitary monodromy, an orbifold realization identifies this filtration with the Betti weight filtration. Relative toric strata and Springer sign extraction give constancy in each cohomological degree and each weight over every generic determinant component. Each such component contains a suitable finite-order unitary point.

The stable-pair identity is retained in the Grothendieck group of Hodge structures until both cohomological indices have been recovered. Purity makes the weights of the different compactly supported degrees distinct at a fixed perverse index. We can then extract those degrees separately and apply the hook substitution. This yields the full mixed Hodge formula of Theorem 2.2. The final section records the duality, signs, and shifts and derives positivity and the pure-part formula.

2. The character-variety formula

We prove the mixed Hodge formula of [23, Conjecture 1.2.1]. We begin by fixing its geometric data and symmetric-function conventions.

Fix integers $g \geq 0$, $m \geq 1$, and $n \geq 1$. For each $p \in \{1, \dots, m\}$ let $\mu^p = (\mu_1^p, \dots, \mu_{\ell_p}^p)$ be a partition of n. Choose pairwise distinct numbers $\xi_{p,1}, \dots, \xi_{p,\ell_p} \in \mathbf{C}^*$, and let $\mathcal{C}_p \subset \mathrm{GL}_n(\mathbf{C})$ be the semisimple conjugacy class with these eigenvalues and multiplicities. The tuple is *generic* if

$$\prod_{p,a} \xi_{p,a}^{\mu_a^p} = 1, \qquad \prod_{p,a} \xi_{p,a}^{\nu_a^p} \neq 1 \tag{2.1}$$

whenever $0 \leq \nu_a^p \leq \mu_a^p$ and $\sum_a \nu_a^p = r$ for every p, with $0 < r < n$. Write $\boldsymbol{\mu} = (\mu^1, \dots, \mu^m)$ and define

$$\mathcal{M}_{\boldsymbol{\mu}} = \left\{ (A_i, B_i, X_p) \in \mathrm{GL}_n(\mathbf{C})^{2g} \times \prod_{p=1}^{m} \mathcal{C}_p : \prod_{i=1}^{g} [A_i, B_i] \prod_{p=1}^{m} X_p = I_n \right\} /\!\!/ \mathrm{GL}_n(\mathbf{C}). \tag{2.2}$$

Here $[A, B] = ABA^{-1}B^{-1}$ and the quotient is for simultaneous conjugation. A common invariant subspace of dimension $0 < r < n$ would meet the eigenspaces of X_p in dimensions ν_a^p. Taking the determinant of the relation on this subspace gives $\prod_{p,a} \xi_{p,a}^{\nu_a^p} = 1$, contrary to (2.1). Every representation is therefore irreducible, and its commuting endomorphisms are scalars.

Lemma 2.1. *At an irreducible point of the relation variety, the differential of the relation map to* SL_n *is surjective. The relation variety and its geometric* PGL_n *quotient are smooth at that point.*

Proof. Write $\mathrm{Ad}_A Z = AZA^{-1}$ and use right logarithmic tangent coordinates with the trace pairing on $\mathfrak{gl}_n$. For $c = [A, B]$, put $a = \dot{A}A^{-1}$ and $b = \dot{B}B^{-1}$. Differentiation gives

$$\dot{c}c^{-1} = (1 - \mathrm{Ad}_{ABA^{-1}})a + (\mathrm{Ad}_A - \mathrm{Ad}_c)b.$$

If Z annihilates this image, set $K = ABA^{-1}$ and $U = \mathrm{Ad}_{A^{-1}} Z$. The two tangent equations give

$$\mathrm{Ad}_K Z = Z, \qquad U = \mathrm{Ad}_{c^{-1}} Z = \mathrm{Ad}_B Z.$$

The first equation also gives $\mathrm{Ad}_B U = U$. Substituting $U = \mathrm{Ad}_B Z$ and applying $\mathrm{Ad}_{B^{-1}}$ yields $\mathrm{Ad}_B Z = Z$. Thus $U = Z$, and Z commutes with both A and B.

In the product relation, the right logarithmic differential of each factor is conjugated by the preceding product. Beginning with the first commutator, the calculation just made shows successively that an annihilating covector commutes with every A_i, B_i; the preceding products then fix that covector. The tangent space to the conjugacy class of X_p has right logarithmic image $(1 - \mathrm{Ad}_{X_p})\mathfrak{gl}_n$. The same induction shows commutation with every X_p. Conversely, a common commuting endomorphism annihilates each term of the differential.

The common centralizer is $\mathbf{C}I_n$. The determinant of the relation map is identically one, and the trace pairing identifies $\mathfrak{sl}_n^*$ with $\mathfrak{gl}_n/\mathbf{C}I_n$. The differential to $\mathfrak{sl}_n$ therefore has zero annihilator and is surjective. Irreducibility makes all conjugation orbits closed and the effective PGL_n action free. The stable geometric quotient is an étale locally trivial PGL_n torsor, so smoothness descends to it. □

Subtracting the dimension $n^2 - 1$ of the relation target and then that of PGL_n shows that a nonempty $\mathcal{M}_{\boldsymbol{\mu}}$ is smooth affine of dimension

$$d_{\boldsymbol{\mu}} = 2(g-1)n^2 + 2 + \sum_{p=1}^{m} \left(n^2 - \sum_a (\mu_a^p)^2 \right). \tag{2.3}$$

This integer is even, since $n^2 - \sum_a (\mu_a^p)^2$ is even for every p. Empty varieties have zero cohomology.

Let $X_1, \ldots, X_m$ be independent alphabets. All symmetric functions below are over $\mathbf{Q}(z, w)$. For $X = (x_1, x_2, \ldots)$, put $p_j[X] = \sum_i x_i^j$ and $h_r[X] = \sum_{i_1 \le \cdots \le i_r} x_{i_1} \cdots x_{i_r}$, with $h_0 = 1$. For a partition ρ, put $p_\rho = \prod_a p_{\rho_a}$, $h_\rho = \prod_a h_{\rho_a}$, and $z_\rho = \prod_{j \ge 1} j^{m_j(\rho)} m_j(\rho)!$, where $m_j(\rho)$ counts the parts equal to j. The monomial function m_ρ is the sum of the distinct monomials whose positive exponents are the parts of ρ. Use the Hall pairing $\langle p_\rho, p_\sigma \rangle = \delta_{\rho\sigma} z_\rho$ and its product on independent alphabets. For a partition λ, write $|\lambda| = \sum_i \lambda_i$ and $\|\lambda\|^2 = \sum_i \lambda_i^2$. The arm a_b and leg l_b of a box b count the boxes to its right and below; its hook length is $h_b = a_b + l_b + 1$. Put

$$\mathcal{H}_\lambda^g(z, w) = \prod_{b \in \lambda} \frac{(z^{2a_b+1} - w^{2l_b+1})^{2g}}{(z^{2a_b+2} - w^{2l_b})(z^{2a_b} - w^{2l_b+2})}, \tag{2.4}$$

$$\Omega_{g,m}(z, w; X_1, \ldots, X_m) = \sum_\lambda \mathcal{H}_\lambda^g(z, w) \prod_{p=1}^{m} \widetilde{H}_\lambda[X_p; z^2, w^2], \tag{2.5}$$

$$\mathbb{H}_{\boldsymbol{\mu}}(z, w) = (z^2 - 1)(1 - w^2) \left\langle \operatorname{Log} \Omega_{g,m}, \prod_{p=1}^{m} h_{\mu^p}[X_p] \right\rangle. \tag{2.6}$$

The empty partition contributes 1. With μ_{Mob} the Möbius function, the plethystic logarithm is

$$\operatorname{Log} F = \sum_{r \ge 1} \frac{\mu_{\mathrm{Mob}}(r)}{r} \log \psi_r(F), \quad \psi_r(z, w, p_j[X_p]) = (z^r, w^r, p_{rj}[X_p]). \tag{2.7}$$

Every coefficient of a prescribed positive alphabet degree is a finite expression. When a rank variable is adjoined, its Adams operation is $R \mapsto R^r$.

Our convention for the modified Macdonald function is

$$\widetilde{H}_\lambda[X;q,t] = t^{n(\lambda)} J_\lambda\Big[\frac{X}{1-t^{-1}};q,t^{-1}\Big], \qquad n(\lambda)=\sum_{i\ge 1}(i-1)\lambda_i. \tag{2.8}$$

Here $J_\lambda = c_\lambda P_\lambda$, $c_\lambda(q,t) = \prod_b (1-q^{a_b}t^{l_b+1})$, and P_λ is monic triangular in the monomial symmetric functions and orthogonal for

$$\langle p_\rho, p_\sigma\rangle_{q,t} = \delta_{\rho\sigma} z_\rho \prod_j \frac{1-q^{\rho_j}}{1-t^{\rho_j}}.$$

In particular $\widetilde{H}_\lambda[1;q,t]=1$.

For a complex algebraic variety Y, let W and F denote the weight and Hodge filtrations on its cohomology and put

$$h_c^{a,b;D}(Y) = \dim_{\mathbf{C}} \operatorname{Gr}_F^a \operatorname{Gr}^W_{a+b} H_c^D(Y,\mathbf{C}).$$

Write

$$H_c(Y;x,y,u) = \sum_{a,b,D} h_c^{a,b;D}(Y) x^a y^b u^D.$$

If the cohomology is Hodge–Tate, these numbers vanish for $a\neq b$. Put $Q=xy$ in this case and write $H_c(Y;Q,u)$.

Theorem 2.2. *For every generic tuple of semisimple conjugacy classes as above, the compactly supported cohomology of $\mathcal{M}_{\boldsymbol{\mu}}$ is Hodge–Tate, and*

$$H_c(\mathcal{M}_{\boldsymbol{\mu}};Q,u) = (u\sqrt{Q})^{d_{\boldsymbol{\mu}}}\mathbb{H}_{\boldsymbol{\mu}}(-Q^{-1/2}, u\sqrt{Q}). \tag{2.9}$$

The polynomial is independent of the generic eigenvalues. Moreover, $\mathbb{H}_{\boldsymbol{\mu}}(-z,w)\in \mathbf{Z}_{\ge 0}[z,w]$. If the variety is nonempty, this polynomial has degree $d_{\boldsymbol{\mu}}$ in each variable. Its pure compactly supported part satisfies

$$\sum_k h_c^{k,k;2k}(\mathcal{M}_{\boldsymbol{\mu}})Q^k = Q^{d_{\boldsymbol{\mu}}/2}\mathbb{H}_{\boldsymbol{\mu}}(0,\sqrt{Q}). \tag{2.10}$$

The theorem includes divisible multiplicity profiles and all generic determinant components, with the full weight and cohomological grading.

3. The ordinary-curve construction

Let C be a smooth projective connected complex curve of genus $g\ge 0$, and write K_C for its canonical line bundle. For a positive integer n and a primitive nth root of unity ζ, set

$$\mathcal{M}_{g,n} = \left\{ (A_1,B_1,\dots,A_g,B_g)\in \mathrm{GL}_n(\mathbf{C})^{2g} : \prod_{j=1}^g [A_j,B_j] = \zeta I_n \right\} /\!\!/ \mathrm{GL}_n(\mathbf{C}). \tag{3.1}$$

Here $[A,B]=ABA^{-1}B^{-1}$ and $/\!\!/$ denotes the affine geometric invariant theory quotient for simultaneous conjugation. Put $N_n=(g-1)n^2+1$. For $g\ge 1$, this twisted character variety is smooth of complex dimension $2N_n$. The Hodge–Tate assertion is recovered from the weight comparison at the end of the construction. We write W for Deligne's weight filtration, use ordinary cohomology, and set

$$H(\mathcal{M}_{g,n};q,t) = \sum_{D,k} \dim \operatorname{Gr}^W_{2k} H^D(\mathcal{M}_{g,n},\mathbf{Q})\, q^k t^D. \tag{3.2}$$

We use the arm, leg, and hook lengths fixed in Section 2, and the factor $\mathcal{H}^g_\lambda(z,w)$ of (2.4). With R a formal rank variable, define

$$\Omega(R;z,w) = \sum_\lambda R^{|\lambda|}\mathcal{H}^g_\lambda(z,w), \tag{3.3}$$

$$\overline{H}_n(z,w) = (z^2-1)(1-w^2)[R^n]\operatorname{Log}\Omega(R;z,w). \tag{3.4}$$

The empty partition contributes 1. Here μ_{Mob} denotes the Möbius function and Log is the plethystic logarithm for the Adams operations $\psi_d(R,z,w)=(R^d,z^d,w^d)$; explicitly,

$$\operatorname{Log} F=\sum_{d\geq 1}\frac{\mu_{\mathrm{Mob}}(d)}{d}\log F(R^d;z^d,w^d). \tag{3.5}$$

Theorem 3.1. *For every $g\geq 2$ and every $n\geq 1$,*

$$H(\mathcal{M}_{g,n};q,t)=(t\sqrt{q})^{2N_n}\overline{H}_n\left(\sqrt{q},-\frac{1}{t\sqrt{q}}\right). \tag{3.6}$$

The expression on the right is independent of the choice of the square root and is a polynomial in q,t.

3.1. **Notation and coefficients.** For the proof of Theorem 3.1, assume $g\geq 2$ and put $h=g-1$. Choose a maximal rank for the argument; every construction below is coefficientwise in this finite rank range. Choose a line bundle L of sufficiently large even degree ℓ and sections

$$q_0=(a,b):\mathcal{O}_C\oplus K_C\twoheadrightarrow L,\qquad A=K_CL^{-1}. \tag{3.7}$$

Such sections exist without common zero for ℓ sufficiently large. The extension class of $0\to A\to \mathcal{O}_C\oplus K_C\to L\to 0$ is $\alpha_q\in H^1(C,K_CL^{-2})$. The two threefolds

$$X_0=\operatorname{Tot}_C(L\oplus A),\qquad X_1=\operatorname{Tot}_C(K_C\oplus\mathcal{O}_C) \tag{3.8}$$

have the volume forms induced by their determinant identifications with K_C. Write

$$B_r=\ell r^2,\qquad a_r=-hr-\frac{\ell r(r-1)}{2},\qquad \alpha_r=\frac{a_r}{r},\qquad c_r=\alpha_r-\frac{1}{2r},\qquad N_r=hr^2+1. \tag{3.9}$$

The symbol a_r in (3.9) is unrelated to an arm length a_b.

For a line bundle D on C, set $\mathcal{A}_r^D=\bigoplus_{j=1}^r H^0(C,D^j)$. The Hitchin map records the coefficients of the characteristic polynomial, using the same convention for all D. Multiplication of monic spectral polynomials defines the finite addition maps

$$\mathrm{add}:\prod_\nu \mathcal{A}^D_{r_\nu}\longrightarrow\mathcal{A}^D_r,\qquad \sum_\nu r_\nu=r. \tag{3.10}$$

Their finiteness will also be proved directly below. We use $\mathrm{Exp}_{\mathrm{add}}$ and $\mathrm{Sym}_{\mathrm{add}}$ for the symmetric-power operations using these maps, including their finite permutation projectors.

We write MHS for rational mixed Hodge structures, $\mathrm{MHM}(Y)$ for rational algebraic mixed Hodge modules on Y, and $\mathrm{MMHM}(Y)$ for their monodromic version; put $\mathrm{MMHS}=\mathrm{MMHM}(\mathrm{pt})$. The symbol t_0 denotes classical truncation, and IC_Y is the perverse-normalized intersection complex. All six operations are derived. The notation ω_Y denotes the Hodge dualizing complex, $\mathbb{D}$ denotes Verdier duality, and ${}^p\mathcal{H}^i$ denotes perverse cohomology. We use the perverse normalization of monodromic vanishing cycles ϕ_f. When taking vanishing cycles and their products, we use the conventions of [15]. Set

$$\mathbf{T}=\mathbf{Q}(-1),\qquad \mathbb{L}=\mathbf{T}[-2],\qquad \vartheta^2=\mathbf{T},\qquad \Lambda=\vartheta[-1]. \tag{3.11}$$

For every integer m, write $\mathbf{T}^{m/2}=\vartheta^m$ and $\mathbb{L}^{m/2}=\Lambda^m$. The tensor product in MMHM is the monodromic convolution product. The half-Tate object ϑ has even self-symmetry; Λ has odd self-symmetry. Thus $\Lambda^2=\mathbb{L}$. We prove the relevant permutation calculation in Section 11. The intrinsic weights are $\mathrm{wt}(\vartheta)=1$ and $\mathrm{wt}(\Lambda)=0$. For a pure complex K of weight w, ${}^p\mathcal{H}^i(K)$ is pure of weight $w+i$.

The ordinary embedding $\mathrm{MHM}(Y)\to\mathrm{MMHM}(Y)$ is exact, fully faithful, and monoidal. It has an exact fiber realization to ordinary mixed Hodge modules on its image. We use this to return to ordinary coefficients only after proving that the objects in question belong to that image. All preceding products are formed in MMHM. If K has bounded finite-dimensional cohomology,

$[K]$ denotes its alternating class in $K_0(\mathrm{MHS})$ or $K_0(\mathrm{MMHS})$, as appropriate. Square brackets around a complex, such as $[P_i[-i]]$, always retain its Hodge structure.

For a smooth moduli space, $\mathbf{Q}[d]$ denotes its constant Hodge module normalized in complex dimension d. For an oriented critical moduli space Y, Φ_Y is its normalized monodromic Donaldson–Thomas coefficient. On a smooth zero-potential critical chart its value is the corresponding constant object with the normalization (3.11). The construction and all orientation maps used in this paper are given in Section 4.

3.2. **Moduli and the two support categories.** A compactly supported stable pair on X_i is a map $\mathcal{O}_{X_i} \to F$ with F pure of dimension one and zero-dimensional cokernel. Write $P_{r,\chi}(X_i)$ for its moduli space, where the pushdown of F to C has rank r and Euler characteristic χ. The framed stack used to compare stability chambers is the derived stack of surface pairs on $S = \mathrm{Tot}_C(L)$, with its two additional dual fields in the critical model. Its ordinary surface data are

$$(E, \Theta, s), \qquad \Theta : E \longrightarrow E \otimes L, \quad s : \mathcal{O}_C \longrightarrow E. \tag{3.12}$$

Here E is a vector bundle. The additional fields, the relation variable, and the coherent realization of its framing are described in Section 4. We use Euler slope $\mu(E) = \chi(E)/\operatorname{rk}(E)$ and the Harder–Narasimhan filtration of the L-Higgs pair (E, Θ). A cutoff c means $\mu_{\min}(E, \Theta) > c$. The symbol $\mathscr{P}_{r,\chi}(c)$ denotes the corresponding nonzero-section surface window. The potential is

$$g_q(E, \Theta) = \frac{1}{2}\langle \alpha_q, \mathrm{Tr}(\Theta^2)\rangle. \tag{3.13}$$

It factors through a function $g_{\mathcal{A}}$ on $\mathcal{A}_r^L$.

For a rank partition $\nu = (r_1, \ldots, r_s)$ of r, let $\Sigma_\nu \subset \mathcal{A}_r^L$ be the image of spectral addition. There are two different Serre subcategories of the perverse heart $\mathrm{MMHM}(\mathcal{A}_r^L)$ in the argument. The category $\mathscr{S}_r$ consists of objects whose simple constituents have strict supports among the Σ_ν. It includes the single rank-r semistable unframed coefficient. The category $\mathscr{I}_r$ is the Serre subcategory generated by the perverse constituents of the finite-addition convolutions of the framed and unframed coefficients occurring in the proper flag correspondences, with at least two positive-rank factors, all of rank less than r. In taking these generators we first perform the indicated shifts, Tate twists, and source vanishing-cycle operations, and then take perverse cohomology and its simple constituents. In particular $\mathscr{I}_1 = 0$. For either category $\mathscr{J}$, let $D_{\mathscr{J}}$ denote the complexes whose perverse cohomology objects all belong to $\mathscr{J}$.

Direct images from the intermediate stacks can be unbounded on one side. Statements about them are always made in a fixed perverse degree. Grothendieck classes are taken at the bounded proper endpoints. A finite-type smooth presentation gives a common lower standard bound on the ordinary pullbacks to its Čech nerve. The mth totalization fiber has this bound increased by m; therefore only finitely many levels contribute in any fixed degree. On the finite-dimensional Hitchin base, the same observation gives a finite Postnikov range for each fixed perverse degree of a finite convolution. All Grothendieck classes used later belong to bounded proper stable-pair or Hitchin endpoints.

3.3. **Established inputs.** For a sufficiently positive effective L, the positive Hitchin support theorem and cohomological χ-independence hold in the derived category of mixed Hodge modules [44, Theorems 0.1 and 0.4]. The smooth-stack integrality theorem of [9, Theorem 10.2.7 and Lemma 10.2.9], with its commutative cotangent orientation, applies to the fixed-slope positive Higgs stack. Section 8 verifies its hypotheses and fixes the scalar stabilizer parameter.

The integral localization theorem of [35, Corollary 7.19 and Remark 7.20] supplies the morphism for filtered oriented 3-Calabi–Yau objects. We use its chart-level monodromic Hodge realization together with the specialization exchange of [20]. The smooth descent and the comparison of

orientations are included below. The conductor transversality criterion of [51, Theorem 4.16(3)] supplies smoothness of the relative Hilbert schemes on the reduced positive spectral locus.

Finally, smooth semiprojective Higgs moduli have pure cohomology [28, Corollary 1.3.2]. The $P = W$ theorem for coprime GL_n Higgs moduli is [45, Theorem 0.2]. The rational tautological classes have the Hodge types stated in [61]; together with tautological generation, they identify the resulting weight pieces with Tate structures. These facts are used only after the two stable-pair calculations and their comparison have been proved.

4. Critical models and their orientations

We first construct the comparison of coefficients used below. All determinants in this section are graded determinants. Products of determinants therefore include their Koszul commutativity maps. The derived category of perfect complexes is used with its stable enhancement, so that a fiber, a map of triangles, and their compatibilities refer to the corresponding enhanced constructions.

4.1. **The framed critical model.** Let N be a rank-two bundle on C, with a fixed identification $\det N = K_C$, and choose an exact sequence

$$0 \longrightarrow K_C L^{-1} \longrightarrow N \xrightarrow{q} L \longrightarrow 0. \tag{4.1}$$

Its signed extension class is denoted by $\alpha_q \in H^1(C, K_C L^{-2})$. The sign is fixed by writing a transition between local lifts of the quotient coordinate as $X_j = X_i - c_{ij}\Theta$, with c_{ij} representing α_q. Put $S = \mathrm{Tot}_C L$, with projection $p : S \to C$. The map $\mathrm{Tot}_C N \to S$ is the K_S-torsor with class $\alpha_q y$, where y is the tautological section of the pullback of L.

Let $\mathscr{P}_L$ be the derived stack of surface pairs $\mathcal{O}_S \xrightarrow{s} G$, with G perfect and properly supported, and write

$$J_G = \mathrm{fib}(\mathcal{O}_S \to G).$$

We work on the locus on which the ordinary sheaf G is finite over C. It can then be described by its curve pushdown E, the endomorphism $\Theta : E \to E \otimes L$, and its section. The regular function used on this stack is

$$g_q(E, \Theta, s) = \frac{1}{2}\,\mathrm{Tr}_C\Big(\alpha_q \,\mathrm{Tr}(\Theta^2)\Big). \tag{4.2}$$

Here $\mathrm{Tr}_C : H^1(C, K_C) \to \mathbf{C}$ is Serre trace. Formula (4.2) defines a function on the derived stack as well: use a Čech or Thom–Whitney representative of α_q and the universal trace. Its derivative in a bundle and field direction (b, v) is $\mathrm{Tr}_C(\alpha_q \,\mathrm{Tr}(\Theta v))$, since the additional term $\mathrm{Tr}(\Theta[b, \Theta])$ vanishes.

The physical enhancement of a stable pair on $X = \mathrm{Tot}_C N$ is the moduli of perfect complexes

$$I = \mathrm{fib}(\mathcal{O}_Y \to F), \qquad I|_D \simeq \mathcal{O}_D, \quad Y = \mathbf{P}_C(\mathcal{O}_C \oplus N), \quad D = \mathbf{P}_C N.$$

The boundary is framed by the indicated isomorphism. Whenever a determinant trivialization is also used, its restriction agrees with the determinant of this boundary framing, with the specified homotopy of agreement. On a formal neighborhood of this locus, boundary framing agrees with fixing the determinant: near $\mathcal{O}_D$, the perfect-object stack is the derived Picard stack, and $R\Gamma(Y, \mathcal{O}_Y) \to R\Gamma(D, \mathcal{O}_D)$ is an equivalence. Thus restriction to X, followed by extension with the fixed determinant, also identifies the formal problems for the two compactifications used below. The stable-pair locus is an open condition on the classical truncation. Pure one-dimensional sheaves in this locus have locally free curve pushdown: a curve-torsion submodule would generate a zero-dimensional submodule under the finite support algebra.

Proposition 4.1. *The boundary-framed physical enhancement is naturally equivalent, on the indicated locus, to* $\mathrm{Crit}(\mathscr{P}_L, g_q)$. *The equivalence is defined over connective Artin differential graded algebras, and respects maps of objects, their homotopies, direct sums, and finite exact diagrams.*

Proof. Compactify the torsor over S as $\pi : Z = \mathbf{P}_S(V) \to S$, using the lines convention, where

$$0 \longrightarrow K_S \longrightarrow V \longrightarrow \mathcal{O}_S \longrightarrow 0, \qquad H = \mathcal{O}_Z(D_\infty).$$

The projective bundle formulas give

$$R\pi_* H = V^\vee, \qquad R\pi_* H^{-1} = 0, \qquad R\pi_* H^{-2} = K_S[-1]. \tag{4.3}$$

The relative diagonal has resolution

$$0 \longrightarrow H_1^{-1} H_2^{-1} \det(V)^\vee \longrightarrow \mathcal{O}_{Z \times_S Z} \longrightarrow \mathcal{O}_{\Delta_Z} \longrightarrow 0.$$

For a framed perfect object I, define

$$G = R\pi_*(IH^{-1})[1], \qquad U = R\pi_*(IH^{-2}K_S^{-1})[1], \qquad Q = GK_S^{-1}.$$

Applying this diagonal resolution and then tensoring by H recovers

$$I \simeq \mathrm{fib}(\pi^* U \longrightarrow \pi^* G\, H). \tag{4.4}$$

Adjunction identifies the arrow with $U \to GV^\vee$. Since $H|_{D_\infty} = K_S^{-1}$, the framing identifies its boundary restriction with $U \to Q$. Thus the framed object is equivalently a map of triangles

$$\begin{array}{ccccccc}
\mathcal{O}_S & \longrightarrow & U & \longrightarrow & Q & \longrightarrow & \mathcal{O}_S[1] \\
{\scriptstyle s}\downarrow & & \downarrow & & {\scriptstyle \mathrm{id}}\downarrow & & {\scriptstyle s[1]}\downarrow \\
G & \longrightarrow & GV^\vee & \longrightarrow & Q & \xrightarrow{c_G} & G[1].
\end{array}$$

If $e : Q \to \mathcal{O}_S[1]$ is the upper connecting map, the middle arrow includes the specified homotopy $s[1]e \simeq c_G$. In cone coordinates, write φ for a degree-zero cochain in $\mathrm{RHom}_S(Q, G)$, and view e as a degree-one cochain in $\mathrm{RHom}_S(Q, \mathcal{O}_S)$. The equations are

$$d\varphi + se = c_G, \qquad de = 0. \tag{4.5}$$

Conversely these equations construct the upper triangle, its map to the lower triangle, and the fiber (4.4). Both composites are the unit and counit of the displayed diagonal kernel. This also proves the assertion on mapping spaces and higher homotopies, after every derived base change.

The tangent of $\mathscr{P}_L$ at a pair is $\mathrm{RHom}_S(J_G, G)$. Surface duality with proper support identifies its shifted cotangent with $\mathrm{RHom}_S(GK_S^{-1}, J_G[1])$. Consequently (4.5) describes the derived zero stack of the cotangent section $(c_G, 0)$. In the monic resolution

$$[\, p^* E \otimes L^{-1} \xrightarrow{y - \Theta} p^* E \,]$$

multiplication by y is chain homotopic to multiplication by Θ. The torsor cocycle $\alpha_q y$ therefore represents the covector $\mathrm{Tr}_C(\alpha_q \,\mathrm{Tr}(\Theta\, d\Theta))$, including bundle directions by the trace identity preceding the proposition. It is dg_q. In a local splitting of the quotient, $X : E \to E \otimes A$ is the transverse field, and $\varepsilon : E \to K_C$ is the dual framing field. Write ψ for the cotangent field dual to Θ, and $\rho_{\mathrm{cot}}, \rho_{\mathrm{phys}}$ for the relation generators in the cotangent and physical models. The cone convention in (4.5) gives

$$\psi = -X, \qquad e = -\varepsilon, \qquad \rho_{\mathrm{cot}} = -\rho_{\mathrm{phys}}. \tag{4.6}$$

The last sign follows from the first two: they negate the moment expression, so they also negate its degree-minus-one relation variable. This proves the claimed equivalence. □

For the split bundle $L \oplus K_C L^{-1}$, the function is zero. For $N = K_C \oplus \mathcal{O}_C$ and the fixed positive quotient q_0, it is the function (4.2). These two uses have the same cone convention and the same dictionary (4.6).

4.2. **The full deformation complex and its trace.** We record a model that determines both the closed form and its orientation. It is useful to allow a framing bundle R in place of $\mathcal{O}_C$. On a curve and bundle chart, write V_0 for the curve pushdown and W for the two normal directions, with $\Delta = \det W$. The universal coefficient algebra has ordinary coordinates $\Theta, X, s, \varepsilon$, and all matrix entries of an odd coordinate ρ, with

$$|\rho| = -1, \qquad d\rho = [\Theta, X] + s\varepsilon.$$

For homogeneous coefficients and homogeneous matrix maps we use $(a\otimes f)(b\otimes h) = (-1)^{|f||b|}ab\otimes fh$. On the relative projective plane, let $[z_0 : z_1 : z_2]$ be the homogeneous coordinates, with $D = (z_0 = 0)$. The universal perfect complex has terms

$$M^{-1} = V_0\Delta^{-1}\mathcal{O}(-2), \quad M^0 = (V_0 \otimes W^\vee)\mathcal{O}(-1) \oplus R, \quad M^1 = V_0\mathcal{O}$$

and arrows

$$\begin{aligned} \alpha &= (z_0X - z_2,\ -(z_0\Theta - z_1),\ z_0^2\varepsilon), \\ \beta &= (z_0\Theta - z_1,\ z_0X - z_2,\ s), \qquad \gamma = -z_0^2\rho : M^{-1} \longrightarrow M^1. \end{aligned} \tag{4.7}$$

The matrix degree of γ is two and its coefficient degree is minus one. Hence it is part of the total differential, and $d\gamma + \beta\alpha = 0$. Projection to R gives a framing on $2D$: near $D = (z_0 = 0)$ the remaining two-variable Koszul complex is contractible, since one of z_1, z_2 is invertible.

For objects i, j, the strict morphism model $\widetilde{\mathcal{C}}_{ij}$ has generators

$$(g, r) \text{ in degree } 0, \qquad (a, b, u, v) \text{ in degree } 1, \qquad w \text{ in degree } 2.$$

The source and target types are $g, a, b, w : V_i \to V_j$, $r : R_i \to R_j$, $u : R_i \to V_j$, and $v : V_i \to R_j$, with the normal twists in (4.7). Its differential is

$$\begin{aligned} D(g, r) &= (g\Theta_i - \Theta_j g,\ gX_i - X_j g,\ gs_i - s_j r,\ r\varepsilon_i - \varepsilon_j g) - g\rho_i + \rho_j g, \\ D(a, b, u, v) &= aX_i - X_j a + \Theta_j b - b\Theta_i + u\varepsilon_i + s_j v, \qquad Dw = 0. \end{aligned} \tag{4.8}$$

The last two terms of $D(g, r)$ belong to the relation slot. Indeed, applying the second line to the four field entries gives $g([\Theta_i, X_i] + s_i\varepsilon_i) - ([\Theta_j, X_j] + s_j\varepsilon_j)g$, which is canceled by their coefficient differential.

The map to monad morphisms sends (g, r) to $\mathrm{diag}(g, g, g, r, g)$, a field to minus the derivative of (4.7), and w to the lower-to-top entry $-z_0^2 w$. Denote it by κ. Its field entries, in the order lower, middle, frame, top, are

$$(-z_0b, z_0a, -z_0^2v), \qquad (-z_0a, -z_0b, -u).$$

The only product of positive-degree generators that can be nonzero is

$$(a_2, b_2, u_2, v_2)(a_1, b_1, u_1, v_1) = w(-a_2b_1 + b_2a_1 - u_2v_1). \tag{4.9}$$

Degree-zero products and their actions are ordinary rectangular composition. Equations (4.8) and (4.9) follow by multiplication of the displayed matrices. Thus κ is a morphism of differential graded categories.

Lemma 4.2. *The morphism κ is a quasi-isomorphism on every universal chart and after arbitrary derived base change. The relative boundary-framing complex is its subcomplex $r = 0$.*

Proof. All entries in the matrix Hom complex have projective-plane twists between -2 and 2, whose higher cohomology vanishes. At zero fields the monad is $\mathcal{O} \otimes R \oplus i_*V_0[-1]$. The Koszul

resolution of the point i gives the four Hom blocks

source and target	nonzero cohomology
$i_*V_i[-1]$, $i_*V_j[-1]$	$\mathrm{Hom}(V_i, V_j) \otimes \bigwedge^k W \quad (k = 0, 1, 2)$
$\mathcal{O}R_i$, $i_*V_j[-1]$	$\mathrm{Hom}(R_i, V_j) \quad (k = 1)$
$i_*V_i[-1]$, $\mathcal{O}R_j$	$\mathrm{Hom}(V_i, R_j)\Delta \quad (k = 1)$
$\mathcal{O}R_i$, $\mathcal{O}R_j$	$\mathrm{Hom}(R_i, R_j) \quad (k = 0).$

For the first block, differentiation of the two Koszul equations gives the two exterior generators, and the lower-to-top map gives their exterior product. The two framing maps give respectively the restriction and the top Koszul class. Thus the displayed map κ identifies every generator with the corresponding class in this table.

Combine simultaneous dilation in the two normal directions with the central V_0-gauge action. Each of $\Theta, X, s, \varepsilon$ has weight one, and ρ has weight two. The cone of κ is a perfect equivariant complex on the whole affine derived chart. If its classical support were nonempty, the positive dilation would put the origin in its closure. Its derived fiber at the origin is zero by the calculation above. Derived Nakayama then makes that support empty. This argument is made before imposing a stability open, and its quasi-isomorphism can subsequently be pulled back to any derived base. Boundary restriction is exactly $(g, r) \mapsto r$, proving the last assertion. □

Write $\mathcal{C}_{ij}$ for $r = 0$. For arbitrary framing rank this means the fiber of the entire boundary restriction. For rank-one framing on the diagonal it is also the scalar trace fiber.

Lemma 4.3. *The proper relative Serre trace on $\mathcal{C}$ is the cochain pairing*

$$q_{\mathcal{C}}(g, w) = q_{\mathcal{C}}(w, g) = \mathrm{Tr}(gw),$$
$$q_{\mathcal{C}}((a, b, u, v), (a', b', u', v')) = -\mathrm{Tr}(ab' - ba') - \mathrm{Tr}(v'u) + \mathrm{Tr}(vu'). \tag{4.10}$$

The last two traces are taken on the respective framing spaces. After the shift $T = \mathcal{C}[1]$, it agrees under (4.6) *with the canonical critical pairing.*

Proof. Let π be the projection from the relative projective plane to its curve-and-bundle chart, and let π_D be its boundary restriction. Put $G_j = R\pi_* \,\mathrm{RHom}(M, M(-jD))$. The second-order framing gives $G_2 \simeq G_1$, because its boundary cone is $\mathrm{End}(R) \otimes R\pi_{D*}\mathcal{O}_D(-D) = 0$. The pairing in question is the specific product

$$G_1 \otimes G_1 \simeq G_2 \otimes G_1 \longrightarrow G_3 \xrightarrow{\text{Tr, Serre}} \Delta[-2].$$

We compute its sign in the ordered cover $(z_0 \neq 0, z_1 \neq 0, z_2 \neq 0)$. Put $x = z_1, y = z_2$, and order the five monad slots as lower, middle one, middle two, frame, top. At the origin the nonzero differential entries are

$$d_{21} = -y, \quad d_{31} = x, \quad d_{52} = -x, \quad d_{53} = -y.$$

The contractions away from the point are

$$h_x = \frac{E_{13} - E_{25}}{x}, \qquad h_y = -\frac{E_{12} + E_{35}}{y},$$

where E_{ij} is the elementary matrix. They satisfy $[d, h_x] = [d, h_y] = P_V$, the projector onto the unframed summand. Use the total differential $\delta_{\check{\mathrm{C}}\mathrm{ech}} + (-1)^p[d, -]$ in Čech degree p. The two normal cocycles in G_1 are $T = E_{31} - E_{52}$ and $X = -E_{21} - E_{53}$. A G_2 lift of T has components

$$T/z_0, \qquad b_{01} = Th_x/z_0, \quad b_{02} = Th_y/z_0, \quad b_{12} = 0, \qquad c_T = -h_x(b_{01} - b_{02}).$$

Explicitly,

$$b_{01} = \frac{E_{33} + E_{55}}{z_0 x}, \qquad b_{02} = -\frac{E_{32}}{z_0 y}, \qquad c_T = \frac{-E_{13} + E_{25}}{z_0 x^2} - \frac{E_{12}}{z_0 xy}.$$

Multiplication then gives $[d, b_{0i}] = -T/z_0$, $[d, c_T] + b_{01} - b_{02} = 0$, and

$$\operatorname{str}(c_T X) = -\frac{1}{z_0 xy}.$$

For the section cocycle the corresponding components are

$$-E_{54}/z_0^2, \quad -E_{24}/(z_0^2 x), \quad -E_{34}/(z_0^2 y), \quad c_u = -E_{14}/(z_0^2 xy).$$

The dual section in G_1 is $v_1 = -z_0 E_{41}$, and $\operatorname{str}(c_u v_1) = -1/(z_0 xy)$. For the gauge class the top component of its lift is $c_g = -h_x h_y / z_0^2 = E_{15}/(z_0^2 xy)$; multiplication by the relation $-z_0 E_{51}$ has supertrace $+1/(z_0 xy)$. Here the supertrace is $-m_{11} + m_{22} + m_{33} + m_{44} - m_{55}$. Reversing the two field factors reverses their sign, whereas reversing gauge and relation does not. The ordered projective Serre trace sends $1/(z_0 xy)$ to $+1$; this is the ordered residue convention of [55, Proposition 5.2.3]. Tensoring these calculations with arbitrary rectangular Hom spaces gives exactly (4.10).

There is also a direct differential check. If P is minus the field pairing in (4.10), cyclic trace gives $P(Dg, z) = \operatorname{Tr}(gDz)$. For opposite gauge inputs g, h, the additional ghost terms cancel as

$$\operatorname{Tr}((-g\rho_i + \rho_j g)h + g(-h\rho_j + \rho_i h)) = 0.$$

These identities verify the cochain relation with all ghosts.

To see that the origin calculation determines the entire universal pairing, assign generator degree i the section-module weight $-i$. The target $\Delta[-2]$ has degree two and weight -2. A pairing cochain of degree $-k$, with r ghosts and ordinary polynomial degree n, would satisfy

$$r = 2 + k - i - j, \qquad n + 2r = 2 - i - j, \quad \text{hence} \quad n + r + k = 0.$$

For $k > 0$ there is no such cochain; for $k = 0$ only constants on complementary generator degrees occur. Finite semifreeness and exact invariants for the reductive frame group justify this calculation in derived Hom. Thus the comparison and its descent are fixed by the origin computation. The resulting universal maps pull back to every test algebra.

Finally décalage, with s denoting suspension, gives $\omega(su, sv) = (-1)^{|u|} q_{\mathcal{C}}(u, v)$. The tangent field form is consequently $\operatorname{Tr}(d\Theta \wedge dX) + ds \wedge d\varepsilon$, the derivative of $-\operatorname{Tr}(X\, d\Theta) - \varepsilon\, ds$. For a coordinate of degree a, the canonical raw cross pairing of its cotangent generator with the base generator is $(-1)^{a+1}(-1)^a = -1$ times evaluation. The whole dual-half reversal in (4.6) therefore matches the gauge–relation block as well. □

4.3. **Closed forms and the determinant square.** We specify the closed lift of the trace just calculated. In the convention of [57, Theorem 2.12], the leading component of Ch_2 is $\frac{1}{2}\operatorname{Tr}(\mathrm{mult})$. We therefore use

$$\omega_{\mathrm{Perf}}^{\mathrm{Tr}} = 2\mathrm{Ch}_2. \tag{4.11}$$

On Y, the given volume has $K_Y = \mathcal{O}_Y(-3D)$. The framed mapping construction of [64, Theorem 3.3], with extra divisor D, applies: the lift of the specified framing from D to $2D$ is étale because $R\Gamma(D, \mathcal{O}_D(-D))[1] = 0$. Its framing isotropic path is zero, since the ideals of D and $2D$ multiply to zero on $3D$. The other path is the boundary triangle $K_Y \to \mathcal{O}_Y \to \mathcal{O}_{3D}$. This constructs a closed (-1)-shifted form whose leading pairing is precisely the proper product in Lemma 4.3. The factor in (4.11) does not change the coefficient $1/2$ of (4.2): the latter is its quadratic primitive coefficient.

Let $\mathrm{pr} : \mathcal{C}_{ij} \to \mathcal{B}_{ij}$ be the projection to the surface framing complex, forgetting b, v, w. It is a strict functor. Its kernel has b, v in degree one and w in degree two, and has square zero by (4.9). The pairing

$$(b, v), (a, u) \longmapsto \operatorname{Tr}(ba) + \operatorname{Tr}(vu), \qquad w, g \longmapsto \operatorname{Tr}(wg)$$

identifies this kernel with $\mathcal{B}_{ji}^{\vee} K_C[-2]$. Writing $\langle -, - \rangle$ for this pairing, the two bimodule identities are

$$\langle kh, b \rangle = \langle k, \mathrm{pr}(h) b \rangle, \qquad \langle hk, b \rangle = (-1)^{|h|(|k|+|b|)} \langle k, b\, \mathrm{pr}(h) \rangle.$$

Here k lies in the kernel and h, b are homogeneous composable entries. In generator degree two they are rectangular trace identities, using (4.9); all other cases vanish. Thus, after curve sections and Serre duality, there is a natural paired triangle

$$B_{ji}^{\vee}[-3] \longrightarrow C_{ij} \longrightarrow B_{ij}, \qquad B_{ij} = R\Gamma(C, \mathcal{B}_{ij}), \quad C_{ij} = R\Gamma(C, \mathcal{C}_{ij}). \tag{4.12}$$

It extends to finite twisted sums and cones by their block-matrix differentials. In particular it applies when a cofiber is a finite-length sheaf.

The leading pairing comparison of Lemma 4.3 extends to the full closed forms. Both forms have weight two for normal dilation. If δ is their difference and h_0 is its leading homotopy, then $\delta - Dh_0 \in F^3\mathrm{DR}$. For the Euler field $\mathcal{E}$, Cartan's formula gives

$$H_0 = h_0 + \tfrac{1}{2}\iota_{\mathcal{E}}(\delta - Dh_0) \in F^2\mathrm{DR}, \qquad DH_0 = \delta. \tag{4.13}$$

This uses the full de Rham mixed complex. In a formal presentation with invariant ideal J, the filtration $\prod_p J^{\max(m-p,0)}\Omega^p$ makes the operation continuous. Construct (4.13) on the equivariant formal diagram, with its transition maps, before restricting to a completion at a moving point.

For a diagonal object put $B = B_{ii}$. The half line is $\det B$; its square is given by (4.12), duality, and the canonical critical factor $(-1)^{\mathrm{vdim}\,\mathscr{P}_L}$ in [35, Example 3.20]. This defines the orientation used throughout the paper. The bimodule structure and the determinant double-sequence identity [35, Lemma 2.4] give its direct-sum and finite-flag compatibilities.

Proposition 4.4. *Let Ψ be the equivalence of Proposition 4.1, or one of the supported equivalences constructed below, and let M_A, M_B be the orientation lines of its source and target. Its specified half-line comparison $\gamma : M_A \to \Psi^* M_B$ satisfies*

$$o_A = \det(d\Psi^*)\, o_B\, \gamma^{\otimes 2}, \qquad d\Psi^* : \Psi^*\mathbb{L}_B \longrightarrow \mathbb{L}_A. \tag{4.14}$$

Here o_A, o_B are the orientation square maps and $\mathbb{L}_A, \mathbb{L}_B$ are cotangent complexes.

Proof. First consider (4.1). Write its paired triangle as $K \to C \to B$, with $h : K \xrightarrow{\sim} B^{\vee}[-3]$. The cotangent triangle is

$$B^{\vee}[-1] \longrightarrow C^{\vee}[-1] \longrightarrow K^{\vee}[-1].$$

Identify its last object with $B[2]$ by the inverse of $h^{\vee}[-1]$, using the same shifted double-dual convention as in [35, Appendix B.1]. The physical and critical triangles have identical first arrows. On the last term, the dual of the equality $h_{\mathrm{crit}}\Psi_K = h_{\mathrm{phys}}$, proved in Lemma 4.3, is

$$\Psi_K^{\vee}[-1]h_{\mathrm{crit}}^{\vee}[-1] = h_{\mathrm{phys}}^{\vee}[-1].$$

Consequently these are equivalent diagrams of fiber sequences. Taking their graded determinants, and inserting the same critical factor on both sides, proves (4.14). This keeps the direction of the cotangent derivative. For example, scaling h by c changes the raw square by $c^{\chi(B)}$, and negating the whole dual half changes its derivative determinant by $(-1)^{\chi(B)}$; these are the same factor. Dualizing and shifting inverts the full determinant map.

For a finite exact diagram, apply this proof to the paired bimodule triangles (4.12) on every entry of its double Hom grid. Naturality of duality and determinant double-sequence coherence give the asserted square for the resulting half-line map. The supported case and the identification of its full map with the derivative of the specified Ψ are spelled out in Lemma 4.7. □

4.4. **The fixed positive quotient.** We compare this orientation with the canonical one for $N = K_C \oplus \mathcal{O}_C$. For two fixed even-degree quotients $q_i : N \to L_i$, put $A_i = \ker q_i$, and assume their kernels are generically transverse. In the full deformation complex the dual-section and relation terms form a common isotropic subcomplex J. Its orthogonal quotient is

$$J^{\perp}/J = \mathrm{End}(E) \otimes N[-1],$$

and the two isotropic halves have quotients $\mathrm{End}(E) \otimes A_i[-1]$. These are filtrations of paired complexes.

Lemma 4.5. *The two orientation roots differ by $\mathcal{C}_{\mathrm{res}}^{r^2}$, where $\mathcal{C}_{\mathrm{res}}$ is one fixed quadratic line depending only on the two quotients. Choosing its quadratic trivialization once gives a square-preserving comparison, compatible with direct sums and flags. Its factors cancel canonically between a finite modification and its same-rank baseline.*

Proof. Refine the filtration by J, the two isotropic halves, and $J^{\perp}$. Determinant double-sequence coherence cancels the common hyperbolic part of their squares. The remaining ratio is the Serre linking form on the fixed finite module $N/(A_0 + A_1)$, of length $m = \deg L_0 + \deg L_1 - 2(g-1)$. Choose a local ordered frame with

$$A_1 = \mathcal{O}v_1, \quad A_0 = \mathcal{O}(uv_1 + t^a v_2), \quad v_1 \wedge v_2 = dt,$$

where u is a unit. Decomposing xv_2 along these two lines gives the linking form

$$b_{\mathrm{res}}(x, y) = \operatorname{Res} \frac{uxy\,dt}{t^a}.$$

Adding an integral lift does not change the residue. In the basis $1, t, \ldots, t^{a-1}$, its determinant is $(-1)^{a(a-1)/2}u(0)^a$. Tensoring with the trace form on $\operatorname{End}(E)$ gives

$$\det b_E = (-1)^{mr(r-1)/2}(\det b_{\mathrm{res}})^{r^2}.$$

The length m is even, so the additional sign is one. A change of frame acts by $\mathrm{Ad}(g)$, whose determinant over the finite support ring is one; its norm is also one. Thus this is an equality of quadratic line maps over nonreduced bases, independent of E and of all its fields. The difference of the critical factors has exponent $r^2(\deg L_1 - \deg L_0)$, which is even.

For a direct sum, opposite cross-Hom blocks give the factor $\mathcal{C}_{\mathrm{res}}^{2r_1r_2}$, contracted by its quadratic form. The interchange sign is $(-1)^{mr_1r_2} = 1$. The common determinant filtration gives coherence for further blocks and refinements. A modification and its baseline have the same rank, so their factors cancel by evaluation, including when the defect meets the residue divisor. □

4.5. **Support gluing and the derivative.** For the fixed-point calculation use the equivariant projection

$$q_K : K_C \oplus \mathcal{O}_C \longrightarrow K_C, \qquad S_K = \operatorname{Tot}_C K_C.$$

Transport the fixed positive-quotient orientation by Lemma 4.5. Both quotient degrees are even, and their kernels are generically transverse: otherwise the component $\mathcal{O}_C \to L$ of the positive quotient would vanish, forcing a surjection $K_C \to L$ contrary to $\deg L > \deg K_C$. The fixed quadratic factor has exponent r^2 and cancels between a modification and its same-rank baseline. All Hom, section, support, and duality maps below are equivariant for the torus H defined next. Their determinant grids are taken weight by weight: weight zero supplies the orientation of the fixed enhancement, and the moving weights supply the localization index.

Let $Z_\lambda \subset \operatorname{Tot}_C(K_C \oplus \mathcal{O}_C)$ be the monomial thickening specified by a partition λ, and let $I_0 = \operatorname{fib}(\mathcal{O} \to \mathcal{O}_{Z_\lambda})$. Use the volume-preserving torus

$$H : (\theta, z) \longmapsto (a\theta, a^{-1}z).$$

Consider a finite modification $0 \to \mathcal{O}_{Z_\lambda} \to F \to \widetilde{T} \to 0$. Ordinary H-fixed families of these modifications are identified by curve duality with quotients of the dualizing module ω_{Z_λ}, where $\omega_{Z_\lambda} = \operatorname{Hom}_C(\mathcal{O}_{Z_\lambda}, K_C)$.

Let two étale coordinate maps, over a smooth parameter scheme, identify selected neighborhoods of the finite defect support with neighborhoods on a reference curve. Their normal lifts have the form

$$(t, \theta, z) \longmapsto (f(t),\ \theta/f'(t),\ z),$$

so they preserve $dt \wedge d\theta \wedge dz$. The coordinate maps on ordinary quotient schemes will be denoted by A, B.

Lemma 4.6. *On the selected support branch there is a formal equivalence of the specified H-fixed physical enhancements whose truncation is BA^{-1} on every ordinary Artin base. It is constructed on all connective Artin bases by restriction and gluing with fixed determinant.*

Proof. For a rank-one perfect object define

$$\mathscr{A}_I = \mathrm{fib}(R\mathcal{H}om_Y(I,I) \xrightarrow{\mathrm{str}} \mathcal{O}_Y).$$

The identity has supertrace one, so this is the strict trace kernel in a locally free model. Near $\mathcal{O}_D$, the perfect-object stack equals the derived Picard stack formally: derived Nakayama and Tor-amplitude zero make the deformation an invertible module. Moreover $R\Gamma(Y,\mathcal{O}_Y) \simeq R\Gamma(D,\mathcal{O}_D)$. The determinant and boundary restriction square therefore identifies the boundary-framed formal problem with the fixed-determinant formal problem, including all higher automorphisms.

If Z is the underlying finite defect support and $U = Y - Z$, the modification problem is

$$\big(\mathrm{Perf}(Y) \times_{\mathrm{Pic}(Y)} \{\det I_0\}\big) \underset{\mathrm{Perf}(U)\times_{\mathrm{Pic}(U)}\{\det I_0|_U\}}{\times} \{I_0|_U\}. \tag{4.15}$$

It has deformation complex $R\Gamma_Z\mathscr{A}_I$. An étale map on the selected finite support is universally injective and proper there, hence a closed immersion onto that support. Delete the other sheets over it. This gives an elementary distinguished square with the open complement. Descent for Perf and Pic gives an equivalence of (4.15) on the two sides. Its inverse is the Mayer–Vietoris homotopy fiber of the difference of restriction maps; the same fiber on mapping complexes gives full faithfulness and the higher cells [65, Tag 08GS]. This construction uses the open complement in the derived ambient.

For the monomial baseline, its transverse ideal has scalar endomorphisms, local Ext^1 equal to the Hilbert tangent representation, and no higher local Ext, by the two-variable Hilbert–Burch resolution. The two tangent characters at a box are a^{h_b} and a^{-h_b}; none is zero-weight. After removing scalars, the H-invariant baseline complex on the complement and its overlaps is therefore acyclic. Curve sections and localization in curve points preserve these characters. For a square-zero extension of an Artin base, the obstructions, choices, and all automorphisms for forgetting the complement identification are shifts of this acyclic complex. Induction over small extensions and the Postnikov tower proves that forgetting it is a formal equivalence.

It remains to identify the ordinary map. In an ordinary Artin family, the curve pushdown E is locally free. The monomial annihilator acts by zero away from the defect. A local curve parameter, or its monic characteristic polynomial, is a nonzerodivisor on the free curve module over an Artin ring. The annihilator consequently acts by zero everywhere. The section gives $0 \to \mathcal{O}_{Z_\lambda} \to F \to \widetilde{T} \to 0$ with base-flat finite cokernel. Relative curve duality gives its quotient of ω_{Z_λ}. Selected-sheet pullback commutes with these operations and with the section. Hence the truncation is exactly BA^{-1} on ordinary Artin families. □

Lemma 4.7. *For the equivalence in Lemma 4.6, the algebraic half-line comparison defined by the four supported Hom and section maps has square* (4.14) *for the derivative of that same equivalence.*

Proof. Write the surface modification triangle as $G_0 \to G \to T$. On the equivariant surface S_K, put

$$B_G = \mathrm{fib}(\mathrm{RHom}_{S_K}(G,G) \to R\Gamma(S_K,G)), \quad M(G) = \det B_G, \quad M_H(G) = \det(B_G^H).$$

Superscripts $H, +, -$ denote the invariant, positive-weight, and negative-weight parts for H. The map is precomposition by the section. For the paired triangle (4.12), the positive-weight part of the physical cotangent complex has determinant

$$\det(\mathbb{L}^{>0}_{\mathrm{phys}}) = \det(B_G^-) \otimes \det(B_G^+).$$

The localization orientation $M(G) \otimes \det(\mathbb{L}_{\mathrm{phys}}^{>0})^{-1}$ therefore identifies with $M_H(G)$, using the graded regrouping maps of [35, Section 6.5]. The double Hom grid is H-equivariant before taking determinants. Its full determinant identity is

$$\frac{M(G)}{M(G_0)} = \det \mathrm{RHom}(T, G_0) \otimes \det \mathrm{RHom}(T, T) \otimes \det \mathrm{RHom}(G_0, T) \otimes (\det R\Gamma T)^{-1}. \tag{4.16}$$

Its comparison maps are supported adjunction on $\mathrm{RHom}(G_0, T)$, its Serre dual on $\mathrm{RHom}(T, G_0)$, the support-equivalence unit and counit on $\mathrm{RHom}(T, T)$, and composition of pushforwards on $R\Gamma T$. They are algebraic and compatible with base change. Taking weight zero in each grid entry and then the ordered determinant product defines γ on the lines M_H.

The physical triangle is $I_0 \to I \to J$ with $J = \widetilde{T}[-1]$, since the cofiber of the map of source sections is zero. Represent I as the twisted sum with differential

$$D_I = \begin{pmatrix} d_0 & \kappa_0 \\ 0 & d_J \end{pmatrix}.$$

For a degree-k endomorphism $M = \begin{pmatrix} a & b \\ c & e \end{pmatrix}$, its differential has blocks

$$\begin{aligned}
(\delta M)_{00} &= d_0 a - (-1)^k a d_0 + \kappa_0 c, \\
(\delta M)_{0J} &= d_0 b - (-1)^k b d_J + \kappa_0 e - (-1)^k a \kappa_0, \\
(\delta M)_{J0} &= d_J c - (-1)^k c d_0, \\
(\delta M)_{JJ} &= d_J e - (-1)^k e d_J - (-1)^k c \kappa_0.
\end{aligned}$$

Graded cyclicity cancels the extension terms in the supertrace. The strict trace kernel has a three-step filtration by the upper-right block, the trace-zero upper triangular blocks, and the whole kernel. Its graded pieces are

$$R\mathcal{H}om_Y(J, I_0), \qquad \mathscr{A}_{I_0} \oplus R\mathcal{H}om_Y(J, J), \qquad R\mathcal{H}om_Y(I_0, J). \tag{4.17}$$

The middle identification sends (a_0, e) to $(a_0 - \mathrm{str}_J(e)\,\mathrm{id}_{I_0}, e)$. The extension differential is retained in the filtered complex.

After supported sections and H-invariants, the baseline term in (4.17) is acyclic and the other terms are perfect with finite support, so their determinants are defined. Restriction and gluing preserve the displayed filtration. Consequently its determinant map is the determinant of the derivative of the functor in Lemma 4.6. To compare with the proper Hom grid, use the natural maps from supported to proper sections on its surviving finite-support terms. The possible trace-projector correction on $\mathrm{RHom}(J, J)$ is $1 - uv$, where $u(1) = \mathrm{id}_J$ and $v = \mathrm{str}_J$. After proper sections, these maps run between $R\Gamma(Y, \mathcal{O}_Y)$ and $R\Gamma(Y, R\mathcal{H}om_Y(J, J))$, both perfect over the parameter base. Their composite on the first complex is $vu = \mathrm{rk}(J)\,\mathrm{id} = 0$. The graded Sylvester identity $\det(1 - uv) = \det(1 - vu)$ gives determinant one, even on nonreduced bases.

The paired triangle (4.12), applied to the same exact arrow diagram, has surface halves

$$\mathrm{RHom}(T, G_0), \quad \mathrm{RHom}(T, T), \quad \mathrm{fib}(\mathrm{RHom}(G_0, T) \to R\Gamma T).$$

Its full half determinant is (4.16); its weight-zero half determinant is the same formula with invariants in every Hom and section term. Applying the proof of Proposition 4.4 to this grid identifies the square of γ with the derivative determinant just computed. The three full surface Hom Euler pairings are zero, and the section term has Euler characteristic $d = \mathrm{length}\, T$. If m_0, m_+, m_- are the Euler changes of the half complex in weight zero, positive weights, and negative weights, respectively, then

$$m_0 + m_+ + m_- = -d.$$

The fixed determinant square retains the dual-half reversal $(-1)^{m_0}$. The moving weight pieces retain their corresponding reversal factors; their product with the fixed factor is $(-1)^d$. These factors are independent of the even residue length in Lemma 4.5. □

4.6. **The comparison over the full parameter ring.** The preceding construction also compares the closed forms. On full endomorphism complexes, stable Koszul local cohomology gives the supported product. The boundary section is invertible on the support. Denote the maps from supported sections to $R\Gamma\,\mathrm{RHom}(I, I(-jD))$ by J_j: multiply by the inverse of the specified boundary section to the power j, then apply the support counit. Multiplicativity of this counit gives $J_i(a)J_j(b) = J_{i+j}(ab)$, with its composition homotopies in the enhanced category. The stable Koszul model represents these maps and products. Restricting the inputs to the trace kernel only after taking these products gives the proper pairing of Lemma 4.3. The selected étale maps preserve the volume. Their canonical bundle maps and duality counits therefore compare this pairing on the full moving family, including directions that move the support. The residue change-of-generators, localization, and base-change identities in [55, §§4,10] apply to these counits. Locally a monic characteristic polynomial in the curve coordinate, together with high powers of the two normal coordinates, gives a bounded flat Koszul model. Lifting its coefficients to a connective presentation gives the same construction after derived base change. Different lifts compute the same support functor and its already defined duality counit.

Normal dilation gives both closed forms weight two. Equation (4.13) consequently constructs their full closed homotopy on the common parameter diagram. The following argument upgrades this statement to the particular ordinary maps used in coefficient descent.

Proposition 4.8. *On a neighborhood of the punctual locus, the maps* A, B *identify the algebraic d-critical sections and the oriented vanishing-cycle coefficients. The comparison holds over the entire smooth moving-support and coordinate parameter space.*

Proof. The ordinary map A is étale. Its unique derived étale lift relative to the specified physical enhancement is identified with the lift in Lemma 4.6: their ordinary maps agree, and restriction–gluing identifies their lifting spaces and mapping complexes on all connective Artin test bases. On a selected completion this makes the A-side the fixed physical formal problem extended over $R = \mathbf{C}[[s_1, \dots, s_e]]$. Choose one minimal critical model at its reference point and extend scalars:

$$\mathcal{A}_R = R[[x_1, \dots, x_m]][e_1, \dots, e_m], \quad |e_j| = -1, \quad de_j = I_j(x) \in (x)^2. \tag{4.18}$$

Put $I = (I_1, \dots, I_m) \subset R[[x_1, \dots, x_m]]$. The existence of such a model is the formal critical-chart statement of [8]; it is chosen once over $\mathbf{C}$.

Represent the two specified maps on cofibrant complete models, calling the representatives α, β. Let f_A, f_B be their chart potentials and λ_A, λ_B the Liouville one-forms of their critical models. Set

$$F = \alpha^* f_A - \beta^* f_B, \qquad \ell_0 = \alpha^* \lambda_A - \beta^* \lambda_B.$$

The preceding closed homotopy makes $F - \ell_0 + H_0$ closed in the full relative de Rham complex. Its augmentation to R is zero: both potentials are normalized to vanish on their reduced critical loci, and the moving section is a critical family over the reduced ring R.

Give $x, e, d_{\mathrm{DR}}x, d_{\mathrm{DR}}e$ Euler degree one and R degree zero. On positive Euler degree v, $\iota_{\mathcal{E}}/v$ contracts the augmented de Rham differential. The internal differential raises that degree, by (4.18). The perturbation series

$$h_{\mathcal{E}}(1 + dh_{\mathcal{E}})^{-1}$$

therefore converges in the combined adic and Hodge filtration and contracts the full augmented complex. The zero- and one-form components of a primitive satisfy

$$F = d\Psi_0, \qquad -\ell_0 = d_{\mathrm{DR}}\Psi_0 + d\psi_1.$$

A degree-zero chart coordinate pulls back to a function of x over R; hence ℓ_0 has no $d_{\mathrm{DR}}e_j$ component. Degree constraints give

$$\Psi_0 = \sum_j L_j e_j, \qquad \psi_1 = \sum_{i,j} M_{ij} e_i \, d_{\mathrm{DR}} e_j + \sum_{i,j,k} N_{ijk} e_i e_j \, d_{\mathrm{DR}} x_k.$$

Comparing $d_{\mathrm{DR}}e_j$ coefficients yields $L_j + \sum_i M_{ij} I_i = 0$. Consequently

$$F = -\sum_{i,j} M_{ij} I_i I_j \in I^2 \quad \text{in } R[[x]]. \tag{4.19}$$

This is a single identity over the full parameter ring.

For original charts with acyclic Hessian variables y, put $J = (I, y)$. A class modulo J^2 has the form $f_0(x) + \sum y_a f_a(x)$. If its differential vanishes modulo J, its dy_a coefficients give $f_a \in I$. Thus (4.19) detects equality in the original charts, with their Hessian factors retained. The product parameter directions contribute the identical factor K_{par} to both half lines. The maps are over the parameter base, so their derivative matrices have an identity parameter block and introduce no additional determinant.

The two d-critical sections and the half-line map of Lemma 4.7 existed algebraically before completion. Faithful flatness of local completion detects equality of their sections of $\mathcal{O}/J^2$, and then equality of the two existing square maps on the reduced base. The identification of the critical cotangent determinant with the d-critical canonical line is [35, §§3.5–3.6], with the same Hessian and critical factor used above. Equation (4.14) is therefore the oriented d-critical comparison, proving the assertion about vanishing-cycle coefficients. □

5. Monodromic coefficients and localization

The intermediate comparisons take place in monodromic mixed Hodge modules. We give the descent details needed for the stack correspondences. For an oriented critical chart (U, f) the normalized coefficient is

$$\Phi_{(U,f)} = \phi_f^{\mathrm{mon}}(\mathbf{Q}_U)\Lambda^{-\dim U},$$

tensored with its orientation torsor. Vanishing cycles are normalized to be perverse exact and to satisfy $\phi_0 = \mathrm{id}$. All products in this formula are monodromic products. The ordinary Hodge category embeds exactly and symmetrically monoidally in this category [15, §2.1].

5.1. **An explicit coefficient projector.** Let $\pi : V \times \mathbf{A}^1 \to V$, and let B_V be the stable geometrically monodromic Hodge category on $V \times \mathbf{A}^1$. The constants form the image of π^*. Since $\mathrm{id} \to R\pi_*\pi^*$ is an equivalence, the functor

$$J_V F = \mathrm{cofib}(\pi^* R\pi_* F \to F) \tag{5.1}$$

is the reflector onto $\ker R\pi_*$. Indeed, $R\pi_* J_V F = 0$, and for $R\pi_* K = 0$ adjunction gives $\mathrm{Map}(J_V F, K) \simeq \mathrm{Map}(F, K)$. Thus the quotient by constants is this explicit kernel.

This projector also identifies the heart. If $j : \mathbf{G}_m \hookrightarrow \mathbf{A}^1$, $i : \{0\} \hookrightarrow \mathbf{A}^1$, the triangle

$$\mathbf{Q}_{\mathbf{A}^1} \longrightarrow i_* \mathbf{Q} \longrightarrow j_! \mathbf{Q}_{\mathbf{G}_m}[1]$$

and the change of coordinates $(u, v) \mapsto (u + v, u)$ give

$$J_V F \simeq R(\mathrm{id}_V \times +)_*(F \boxtimes j_! \mathbf{Q}_{\mathbf{G}_m}[1]). \tag{5.2}$$

The first map of the resulting triangle is exactly the counit in (5.1). Equation (5.2) is the perverse-exact projector of [15, §2.1, Lemmas 2.1–2.6]. Its image has the usual monodromic perverse heart. Both formulas commute with smooth pullback, and membership in the kernel is smooth local. They therefore descend through the ordinary smooth nerve of an Artin stack, using the Hodge six operations of [71]. The nilpotent and semisimple monodromy operators remain in this construction, as do the Hodge and weight filtrations.

5.2. **The integral localization map.** For a derived stack X, define the filtered and graded stacks by

$$\mathrm{Filt}(X) = \mathrm{Map}([\mathbf{A}^1/\mathbf{G}_m], X), \qquad \mathrm{Grad}(X) = \mathrm{Map}(B\mathbf{G}_m, X). \tag{5.3}$$

Restriction at zero and evaluation at one give $\mathrm{Grad}(X) \xleftarrow{\mathrm{gr}} \mathrm{Filt}(X) \xrightarrow{\mathrm{ev}} X$. We use the component specified by the chosen integral grading. For a one-parameter torus action μ, write X^μ for the fixed locus and $X^{\mu,+}, X^{\mu,-}$ for the attracting loci as the parameter tends to zero and infinity. Their limit and evaluation maps are

$$X^\mu \xleftarrow{p^\pm} X^{\mu,\pm} \xrightarrow{\eta^\pm} X.$$

Distinguish the functors

$$\mathrm{Loc}_\mu = (p^+)_*(\eta^+)^! \simeq (p^-)_!(\eta^-)^*, \qquad A_\mu = (p^+)_!(\eta^+)^* = \mathrm{Loc}_{\mu^{-1}}.$$

The first is used for filtered-stack attachment, and the second for compact attracting strata. At a fixed point, the index convention is

$$I_\mu = \big(\mathrm{rk}^+ H^0(T_X) - \mathrm{rk}^- H^0(T_X)\big) - \big(\mathrm{rk}^+ H^{-1}(T_X) - \mathrm{rk}^- H^{-1}(T_X)\big),$$

as in [35, Definition 7.21]. For the surface-half model (4.12), it is $\chi(B^-) - \chi(B^+)$.

Proposition 5.1. *For the oriented (-1)-shifted symplectic stacks used here, the integral map is an isomorphism in monodromic mixed Hodge modules*

$$\zeta_X : \Phi_{\mathrm{Grad}(X)}\Lambda^{-I_\mu} \xrightarrow{\sim} R\,\mathrm{gr}_*\,\mathrm{ev}^!\,\Phi_X. \tag{5.4}$$

It is compatible with products, smooth pullback, and matching refinements of an ordered finite flag. For a scheme with a torus action, its compact attracting version is

$$(p^+)_!(\eta^+)^*\Phi_X \simeq \Phi_{X^\mu}\Lambda^{I_\mu}. \tag{5.5}$$

Proof. The stacks in use are covered by finite-type quotient charts, are quasi-separated, and have affine stabilizers. Their orientations and their product and flag maps are supplied by (4.12) and Proposition 4.4. These are the stack hypotheses of [35, Corollary 7.19]. We describe the mixed-Hodge construction of its map; this also fixes the normalization.

On an equivariant critical chart, use the specialization exchange constructed by units and counits in [20, Theorem 3.7, Proposition 6.2]. Its staircase inverse follows from the adjunction triangle identities. Tensor it with the finite-torsor map of [35, (6.75)]. The normalization on a moving hyperbolic quadratic block is the compact Thom map of the same repelling axis as in [35, Theorem 7.9]. Switching that axis would introduce $(-1)^r$ in rank r. The scalar is checked on the normalized quadratic point, where the constructed maps lie in the perverse heart. All other maps are transformations in the Hodge six-functor system.

For a self-dual filtration, locally choose a splitting and write $q = q_{\mathrm{gr}}(1+N)$, with N filtration-raising and self-adjoint. The finite polynomial $(1+N)^{-1/2}$ orthogonalizes it. Orthogonal splittings form a torsor under the filtration-raising orthogonal unipotent group, an iterated affine-bundle torsor. Pullback to it is fully faithful, since its direct-image adjunction unit is an equivalence. The constructed map and all its comparison squares descend uniquely. The interpolation $\exp(t\log u)$ between splittings has determinant one, so it contributes no orientation character.

For smooth descent, let $q : U \to X$ be an equivariant smooth chart and let $q_\mu : U^\mu \to X^\mu$ be its fixed-locus map. Let the relative tangent of q have weight-zero, positive-weight, and negative-weight ranks d_0, d_+, d_-, and let $d = d_0 + d_+ + d_-$. For a smooth map h of relative dimension e, purity is $h^! = h^*[2e](e)$. Applied to the fixed and attracting maps, it gives

$$q_\mu^*\,\mathrm{Loc}_X \simeq \mathrm{Loc}_\mu\, q^*[2d_-](d_-) = \mathrm{Loc}_\mu\, q^*\Lambda^{-2d_-}.$$

With normalized smooth pullback $q^\dagger = q^*\Lambda^{-d}$, this becomes

$$q_\mu^\dagger\,\mathrm{Loc}_X \simeq \mathrm{Loc}_\mu\, q^\dagger\Lambda^{d_+-d_-}.$$

The indices obey $I_U = I_X + d_+ - d_-$, so both the shift and Tate factor of the local map descend. The common charts and their fixed overlaps are [35, Proposition 5.27, Corollary 5.28]. On every level p_m of their smooth nerve, the two ordinary pullbacks are the same invertible shift and twist of perverse heart objects. Hence their mapping spaces have $\pi_r = \mathrm{Hom}(A, B[-r]) = 0$ for $r > 0$. The smooth comparison squares, which are purity and Beck–Chevalley squares, give equal maps on double overlaps. Their compositions give the higher equalities. The discrete mapping spaces then descend the isomorphism and its inverse through the full nerve. Degeneracy compatibility is obtained by pullback of these identities, without applying smooth purity to a nonsmooth diagonal.

For two ordered cuts, write their block weights as a_i, b_i in the same order. Then $(a_i - a_j)(b_i - b_j) \geq 0$ on every cross-Hom block. This checks matching on the full cotangent complex, including relation and obstruction terms. Finitely many ordered cuts admit a common multi-torus chart by the construction of [35, Lemma 5.30]. On this locus the multi-filtered square is Cartesian and its excess map is the identity. Chart comparisons reduce to ordinary Beck–Chevalley by [35, Lemma 7.4]. Thus the product and matching-refinement constructions in [35, Propositions 7.10–7.12, Lemmas 7.14–7.15] are the same Hodge transformations, with determinant associators from (4.12). This proves (5.4). Using the inverse action changes the index sign and gives (5.5). □

Lemma 5.2. *Suppose a filtered correspondence $G \xleftarrow{q} F \xrightarrow{p} U$ occurring in* (5.4) *has classical evaluation a closed immersion $i : F_{\mathrm{cl}} \hookrightarrow U_{\mathrm{cl}}$. Let $h_U : U \to B$ and $h_G : G \to B$ satisfy $h_U p = h_G q$. Its induction map, formed from the adjoint of ζ, is, after that isomorphism,*

$$Rh_{U*}(i_* i^! \Phi_U \longrightarrow \Phi_U).$$

Its cofiber is $Rh_{U*} Rj_* j^* \Phi_U$, where j is the complementary open immersion.

Proof. Put $K_G = \Phi_G \Lambda^{-I}$, and let $c : q^* K_G \to p^! \Phi_U$ be the adjoint of $\zeta : K_G \to Rq_* p^! \Phi_U$. The triangle identity is $Rq_* c \circ \eta_q = \zeta$. Consequently induction is the composite

$$Rh_{G*} K_G \xrightarrow{Rh_{G*}\zeta} Rh_{U*} Rp_* p^! \Phi_U \xrightarrow{Rh_{U*}\epsilon_p} Rh_{U*} \Phi_U.$$

Here $p_* = p_!$ because its classical map is the stated closed immersion. The last arrow is therefore its counit. The closed/open localization triangle proves the cofiber assertion. No properness of h_U is needed. □

5.3. **One-sided constructibility on finite-type windows.** We will apply the preceding operations on bounded finite-type stack windows. If $V \to \mathcal{X}$ is a finite smooth atlas, $p_m : V_m \to \mathcal{X}$ its ordinary nerve, and $f_m : V_m \to B$ the maps to a finite-dimensional base, then a common lower standard bound b for $p_0^* K$ is a lower bound for every $p_m^* K$. Ordinary pullback is standard t-exact and Rf_{m*} preserves this lower bound. For $C^m = Rf_{m*} p_m^* K$, the normalized term $N^m C$ is a retract of C^m, and the totalization tower has fiber

$$\mathrm{fib}(\mathrm{Tot}_m C \to \mathrm{Tot}_{m-1} C) \simeq (N^m C)[-m] \in D^{\geq b+m}.$$

Thus each fixed degree depends on a finite partial totalization. Those finite terms have constructible cohomology. The finite amplitude difference between the standard and perverse structures on B gives finite constructible perverse cohomology and a one-sided bound. In a specified output range, finite convolutions can likewise be computed after a finite Postnikov truncation.

6. Harder–Narasimhan restrictions

We first construct the restriction maps that will be inverted in a perverse quotient. All constructions in this section are made separately on X_0 and X_1. The identification of either physical moduli problem with its critical presentation, including its orientation, is the one constructed in Section 4.

6.1. **The framed heart and its numerical bounds.** Let $X = \mathrm{Tot}_C(N)$ be either threefold in (3.8), and let $Y = \mathbf{P}_C(\mathcal{O}_C \oplus N)$ with infinity divisor $D_\infty = \mathbf{P}_C(N)$. In the usual tilted heart on Y consider

$$\mathscr{A}_X = \langle \mathcal{O}_Y, \mathrm{Coh}_{c,\le 1}(X)[-1] \rangle_{\mathrm{ex}}. \tag{6.1}$$

The framing is an isomorphism along D_∞, retained over the test base. In framing rank one, the weak slope of a framed object is a real parameter t. For every nonzero unframed coherent sheaf F, define $\mu(F[-1]) = \chi(F)/\operatorname{rk}_C(F)$ when $\operatorname{rk}_C(F) > 0$, and define $\mu(F[-1]) = +\infty$ when $\operatorname{rk}_C(F) = 0$. Weak semistability means $\mu(J) \le \mu(Q)$ for every exact sequence $0 \to J \to I \to Q \to 0$.

The heart criterion applies because $\mathrm{End}(\mathcal{O}_Y) = \mathbf{C}$ and

$$\mathrm{Hom}(\mathcal{O}_Y, F[-1]) = 0, \qquad \mathrm{Hom}(F[-1], \mathcal{O}_Y) = \mathrm{Ext}^1_Y(F, \mathcal{O}_Y) = 0.$$

The latter vanishing follows from codimension at least two and the local-to-global Ext spectral sequence. Thus the heart criterion of [70, Proposition 3.6] applies. To check noetherianity, consider a chain of epimorphisms. Its nonnegative framing and curve ranks eventually stabilize, so its remaining kernels are zero-dimensional objects shifted by $[-1]$. The ordinary H^1 sheaves form an eventually constant chain of quotients, after which the H^0 sheaves form an ascending chain inside the fixed reflexive hull of the first one. Coherence terminates this chain. The corresponding decreasing-chain condition for weak Harder–Narasimhan filtrations follows because a remaining nonzero zero-dimensional quotient has slope $+\infty$. These are the two chain conditions in [70, Proposition 2.12].

The boundary framing also removes infinitesimal Picard directions. For the projection $\pi : Y \to C$, relative projective-space cohomology gives

$$R\pi_* \mathcal{O}_Y(-D_\infty) = 0. \tag{6.2}$$

This equality is preserved by derived base change. Hence the trace deformation complex preserving the specified boundary trivialization vanishes.

Write $\mathscr{X}_{r,\chi}(c)$ for the physical window whose projected L-Higgs bundle is pure with minimum Euler slope greater than c. Its nonzero-section open is denoted by $\mathscr{X}^{\times}_{r,\chi}(c)$.

Lemma 6.1. *A compact cyclic structure sheaf $\mathcal{O}_Z$ of curve rank e satisfies*

$$\chi(\mathcal{O}_Z) \ge a_e.$$

Every final L-Higgs Harder–Narasimhan quotient of its pushdown is again the pushdown of a cyclic physical quotient. Consequently its minimum Euler slope is at least α_e. A positive-chamber semistable framed factor of curve rank e has pure projected bundle with minimum Euler slope at least α_e.

Proof. Write Hom_L for morphisms of L-Higgs bundles. The negative line $A = K_C L^{-1}$ preserves the projected Higgs filtration. At a strictly decreasing cut $E' \subset E$, the off-diagonal transverse field lies in $\mathrm{Hom}_L(E', (E/E') \otimes A) = 0$. In the nonsplit row, two local transverse lifts differ by a multiple of the projected field Θ; their off-diagonal blocks therefore agree. The vanishing is global. The same argument applies to an equal-slope subobject of a semistable factor.

Filter the cyclic module of $\mathcal{O}_Z$ by monomials of total degree at most j in the two normal variables and saturate this filtration on C. Let its positive successive ranks be $r_0, \dots, r_s$. They occur consecutively: a stage at which the cyclic span stops growing is already preserved by both variables. Thus $s \le e-1$, $\sum r_j = e$, and

$$\sum_{j=0}^{s} j r_j \le \frac{e(e-1)}{2}.$$

In either row, $\mathrm{Sym}^j(N^\vee)$ has a filtration with line quotients of degree at least $-j\ell$. Its minimum degree slope is therefore at least $-j\ell$. A rank-r_j quotient has degree at least $-j\ell r_j$, and saturation

increases degree. Therefore

$$\deg \pi_*\mathcal{O}_Z \geq -\frac{\ell e(e-1)}{2}, \qquad \chi(\mathcal{O}_Z) \geq -he - \frac{\ell e(e-1)}{2} = a_e.$$

By the negative-line argument, a final Higgs quotient is a physical quotient. A quotient of a cyclic structure sheaf is another cyclic structure sheaf. Applying the preceding bound to its rank $f \leq e$ gives slope at least $\alpha_f \geq \alpha_e$.

The coherent projection of an object with $H^0(I) = I_Z$ and $H^1(I) = Q$ is the exact sequence

$$0 \longrightarrow \pi_*\mathcal{O}_Z \longrightarrow E_I \longrightarrow \pi_*Q \longrightarrow 0. \tag{6.3}$$

It follows by applying $R\pi_*(-D_\infty)[1]$ to the truncation triangle and using (6.2). Finite-t semistability forces E_I to be pure. Indeed its maximal zero-dimensional submodule lifts uniquely through the coherent framing row because $\mathrm{Hom}_S(TK_S^{-1}, \mathcal{O}_S) = \mathrm{Ext}^1_S(TK_S^{-1}, \mathcal{O}_S) = 0$ on the smooth surface $S = \mathrm{Tot}_C(L)$. The lift is preserved by the field, and would give an unframed subobject of infinite slope. Every positive-rank quotient factor of Q has slope at least t. Equation (6.3) now gives the lower bound $\min(\alpha_e, t) = \alpha_e$ when $t > 0$. □

Purity and a lower minimum-slope bound are extension closed. Thus every extension of the good graded factors in the lemma lies in the prescribed window. This applies to the entire attracting fiber. The same assertion holds for the immediate negative chamber $t = -\epsilon$, $0 < \epsilon < 1/r$: a finite slope with denominator at most r that is at least $-\epsilon$ is nonnegative.

6.2. **Complete coherent flags.** The diagonal resolution of Section 4 realizes framing rank at most one by coherent diagrams on S with row

$$0 \longrightarrow \mathcal{O}_S \longrightarrow U \longrightarrow FK_S^{-1} \longrightarrow 0, \tag{6.4}$$

where F denotes the surface projection of the physical sheaf. The diagrams also include the compatible global field arrows. The nonsplit torsor changes the extra arrow but not this exact row. Exact sequences and flat families are termwise exact; in the split case this is also the construction of [69, Theorem 4.5.7 and Corollary 4.5.9].

Lemma 6.2. *For fixed graded Hilbert polynomials and a fixed position of the framing, compatible complete coherent flags in* (6.4) *have a representable projective evaluation map to the middle object. The receiver includes nonsaturated subobjects and torsion in the graded quotients.*

Proof. Use the fixed compactification $\overline{S} = \mathbf{P}_C(\mathcal{O}_C \oplus L)$. The compactly supported quotient F extends by zero, as do its field targets, and the rank-one part of U extends as $\mathcal{O}_{\overline{S}}$. Form the product of the relative projective flag-Quot schemes of these coherent terms. The flag of $\mathcal{O}$ is zero before the framing position and is all of $\mathcal{O}$ afterwards. Require preservation by every field arrow and containment of this framing by setting the composites to the universal quotient sheaves equal to zero. These are closed incidence conditions. For example, a common sufficiently positive twist on the bounded flag-Quot family makes them zero entries of matrices between base-change-compatible vector bundles.

The incidence conditions imply row exactness. Before the framed step, the flag term of U has support of dimension at most one, so it intersects the torsion-free $\mathcal{O}_{\overline{S}}$ trivially. Its map to the corresponding term of FK_S^{-1} is injective; the prescribed Hilbert polynomials make it an isomorphism on every fiber. After the framed step, divide by the contained $\mathcal{O}$ and apply the same argument. Flatness of the flag terms and quotients upgrades these fiberwise identifications to exact rows after arbitrary base change.

All changes from $\mathcal{O}$ are supported on $\operatorname{supp} F$. Thus the boundary flag is the fixed zero or $\mathcal{O}$ flag. The resulting closed incidence subscheme is projective over the middle object. Its field arrows are induced by those of the middle diagram, including on a nonsaturated inclusion. □

For the rank-zero framed semistable factor at finite t, set $a_0 = 0$. This factor is the vacuum $\mathcal{O}_Y$, with Euler characteristic $\chi_0 = 0$.

Lemma 6.3. *For fixed r, χ, c, a finite list of coherent receivers suffices for the Harder–Narasimhan decomposition in a sufficiently large positive chamber. A single chamber parameter can be chosen after the complete receivers, including their torsion quotients, have been formed.*

Proof. Choose a rational $\alpha \le \min(c, \alpha_r, 0)$. If E' is a pure rank-f subobject of the projected middle bundle, its torsion-free quotient has Euler characteristic at least $\alpha(r-f)$. Hence

$$\mu(E') \le \frac{\chi - \alpha(r-f)}{f}, \qquad b = \max_{1\le f\le r} \frac{\chi - \alpha(r-f)}{f}. \tag{6.5}$$

For $t > \max(b, 0)$ the first factor is framed; an earlier unframed factor would violate (6.5). Write its class as $(1, e, \chi_0)$ and the successive unframed classes as $(0, r_j, \eta_j)$. They satisfy

$$a_e \le \chi_0 \le \chi - \alpha(r-e), \qquad \eta_j > c r_j, \qquad \sum_j r_j = r - e, \qquad \chi_0 + \sum_j \eta_j = \chi. \tag{6.6}$$

This is a finite list. The strict inequalities are retained in the list; α is only a bounding device.

The projected window is bounded. Its Higgs Harder–Narasimhan numerical list is finite, and successive slopes in the ordinary bundle filtration of a semistable L-Higgs factor have gaps at most ℓ. In particular a rank-f semistable Higgs factor of Euler slope μ has underlying minimum degree slope at least $\mu + h - (f-1)\ell$. The coherent extensions and compatible arrows in (6.4) form finite-type families, being represented by bounded perfect relative Hom complexes. Form the projective receivers of Lemma 6.2 for the finite list (6.6).

Their graded coherent sheaves form bounded families. For a common large twist k, the length of a maximal zero-dimensional subsheaf $T_0(G)$ is at most $h^0(G(k))$, uniformly bounded. For each possible length, its possible positions are covered by the relative Quot scheme with the complementary Hilbert polynomial. This avoids requiring the maximal torsion subsheaf itself to be flat in the family. Absorb torsion backwards in each flag; each packet moves left through at most r positive-rank steps. Repeating the finite Quot construction still gives bounded families. If H is a pure Higgs subobject of $G/T_0(G)$, then

$$\chi(H(k)) \le h^0(H(k)) \le h^0((G/T_0(G))(k)) \le h^0(G(k)),$$

since $H^1(T_0(G)) = 0$. Thus all slopes occurring after torsion absorption and Higgs refinement have a common upper bound $b_{\mathfrak{fl}}$.

Choose a generic T above b, $b_{\mathfrak{fl}}$, and the final stable-pair thresholds for the finitely many framed classes. These thresholds can be obtained directly: if $t > \chi - a_r$, a positive-rank part of $H^1(I)$ would have Euler characteristic at least t, contradicting (6.3) and Lemma 6.1. The remaining $H^1(I)$ is zero-dimensional, and $\mathrm{Ext}^2(Q, \mathcal{O}_X) = 0$ on the smooth threefold lifts the truncation triangle to an ordinary pair. Absence of zero-dimensional subobjects makes its sheaf pure. Conversely the pure subobject slopes of a bounded stable-pair family are uniformly bounded. Thus T is a stable-pair chamber simultaneously for every receiver under consideration. □

6.3. **The closed-stratum counit.** Order the finite types first by increasing framed curve rank e, then by decreasing anchored tail polygon, choosing a linear extension of this order. Its vertices are

$$(e, \chi_0), \quad (e + r_1, \chi_0 + \eta_1), \quad \dots, \quad (r, \chi).$$

The initial height χ_0 is part of the polygon. Absorbing torsion of length d into the preceding flag member raises the preceding vertex by d and leaves later vertices fixed; this includes the initial vertex. Higgs refinement replaces the result by its concave majorant. If the framed prefix is unstable, its first framed factor has smaller curve rank and is earlier in the primary order. The uniform choice of T ensures that every unframed tail slope is less than T.

Consequently every bad displayed flag lies in an earlier true Harder–Narasimhan type. Delete the union of the earlier proper images and call the resulting open U_τ. Its receiver has exactly the good decreasing semistable flags. Conversely, a good flag is the unique Harder–Narasimhan flag. This identifies the open substacks, since an open substack is determined by its geometric points. Evaluation is now proper and universally injective. Its relative tangent is filtered by the forbidden Hom groups between decreasing semistable factors, and these vanish. It follows that evaluation is a representable closed immersion

$$i_\tau : S_\tau \hookrightarrow U_\tau. \tag{6.7}$$

This also gives the unique relative flag over nonreduced bases.

Use U_τ also for the derived open of the physical window with this classical truncation. Choose ordered integral weights for this flag type, with weight zero on its unique framed graded piece. Let G_τ be the open product of its derived good graded-factor stacks. Using (5.3), define

$$\mathfrak{F}_\tau = \mathrm{Filt}(U_\tau)_\tau \underset{\mathrm{Grad}(U_\tau)}{\times} G_\tau, \qquad G_\tau \xleftarrow{q} \mathfrak{F}_\tau \xrightarrow{p} U_\tau.$$

The framing at infinity stays fixed in this mapping-stack problem. Proposition 4.1, applied to universal finite exact diagrams, identifies its truncation with the coherent flag rows and compatible fields constructed above. Every extension of the good factors lies in the cutoff window by Lemma 6.1. Its decreasing semistable flag is its Harder–Narasimhan flag, so it lies in the current open U_τ. Conversely a filtered object with these good graded terms has that coherent flag. These are open conditions on the common exact-diagram problem, determined by geometric fibers; hence the identification holds over derived test bases. The Rees construction realizes the chosen weights, and its classical evaluation is $t_0(p) = i_\tau$ in (6.7).

Let K_{G_τ} denote the product DT coefficient with the moving weight normalization in integral localization. The oriented localization map of Proposition 5.1, equivalently [35, Corollary 7.19] with its monodromic Hodge construction, is

$$\zeta : K_{G_\tau} \xrightarrow{\sim} Rq_* p^! \Phi_{U_\tau}. \tag{6.8}$$

The good product orientation is the one obtained from the same physical orientation; its moving determinant and all refinement maps are those of [35, Proposition 8.34].

Lemma 6.4. *Let h_{U_τ} and h_{G_τ} be the spectral maps to $\mathcal{A}_r^L$, using spectral addition on the graded factors of G_τ. After Hitchin direct image, the induction map of the good factor coefficient is the closed-stratum costalk counit*

$$Rh_{U_\tau,*} i_{\tau,*} i_\tau^! \Phi_{U_\tau} \longrightarrow Rh_{U_\tau,*} \Phi_{U_\tau}. \tag{6.9}$$

Its cofiber is the direct image from the next open.

Proof. Let $c : q^* K_{G_\tau} \to p^! \Phi_{U_\tau}$ be the adjoint of ζ. The triangular identity says exactly that $Rq_* c \circ \eta_q = \zeta$. Since $h_{U_\tau} p = h_{G_\tau} q$, the integrated induction map is

$$Rh_{G_\tau,*} K_{G_\tau} \xrightarrow{Rh_{G_\tau,*}\zeta} R(h_{U_\tau} p)_* p^! \Phi_{U_\tau} \xrightarrow{Rh_{U_\tau,*}\epsilon_p} Rh_{U_\tau,*} \Phi_{U_\tau}.$$

Mixed Hodge modules are insensitive to derived nilpotent thickenings of the underlying map. With $t_0(p) = i_\tau$, its second arrow is (6.9). The localization triangle for the complementary open j has cofiber $Rh_{U_\tau,*} Rj_* j^* \Phi_{U_\tau}$. Products, refinements, and restrictions use these same units and counits, so their comparison cells are retained. □

In the split row, apply $\phi_{g_\mathcal{A}}$ after these maps have been constructed. The trace potential is additive under finite spectral addition. Moreover the split coefficients are equivariant for a normal dilation under which $g_\mathcal{A}$ has weight two. After the degree-two Kummer cover they are locally acyclic in the function direction away from zero. Thus the relevant vanishing cycles of a sum are supported where every summand has value zero. Proper exchange for finite addition and monodromic Thom–Sebastiani give exactly the lower-rank source images used in $\mathscr{I}_r$.

6.4. **Vacuum removal and finite wall transport.** For $i = 0, 1$, let $\mathscr{X}^{\times}_{i,r,\chi}(c)$ be the physical nonzero-section window in the row X_i, with its map h_i to $\mathcal{A}^L_r$. Set

$$W^{\times}_{i,r,\chi}(c) = R(h_i)_*(\Phi_{\mathscr{X}^{\times}_{i,r,\chi}(c)} \otimes \Lambda^{\chi}).$$

We suppress i in assertions proved separately in both rows. Let $R_{r,\chi}$ denote the corresponding stable-pair restriction, with the same coefficient $\Phi\Lambda^{\chi}$.

Proposition 6.5. *For every $c \le c_r$, the restriction to the stable-pair chamber has cofiber in $D_{\mathscr{I}_r}$. The same is true for restriction to every generic $t > \max(\chi/r, 0)$ and for the nested-window restriction*

$$\rho_{\chi}(c_1, c_2) : W^{\times}_{r,\chi}(c_1) \longrightarrow W^{\times}_{r,\chi}(c_2), \qquad c_1 < c_2 \le c_r. \tag{6.10}$$

For the full ordinary source, the analogous statements hold in $D_{\mathscr{I}_r}$, using the independent rank-r positive unframed support theorem and induction on the framed rank.

Proof. First use the single sufficiently large T chosen in Lemma 6.3. The complete receiver of the vacuum subobject $\mathcal{O}$ is the whole locus $s = 0$. Its quotient is the full unframed object, and the dual framing field is its extension coordinate; it is retained. If the framing complex is $[E^{-1} \to E^0]$ of ranks a, b, then $b - a = \chi$. Giving the section weight -1 and its dual weight $+1$ yields index $a - b = -\chi$. The vacuum costalk of Φ is therefore the unframed coefficient tensored with $\Lambda^{\chi} = \mathbb{L}^{\chi/2}$. After the additional factor Λ^{χ} in the definition of $W^{\times}_{r,\chi}(c)$, the corresponding normalised vacuum term has factor $\Lambda^{2\chi} = \mathbb{L}^{\chi}$.

Immediately before the type $(\mathcal{O}, Q^{\mathrm{ss}})$, the complete vacuum receiver pulls back to its entire fiber over the semistable unframed quotient. Indeed every slope in every complete quotient is less than T; an unstable quotient has an earlier true refinement. Hence there is the restriction square from the complete vacuum to its semistable part and from the original window to this current open. Its vertical cofibers are built from types splitting r into at least two positive ranks. Taking cofibers of the square removes precisely the single rank-r unframed term. The three-by-three lemma, followed by the remaining costalk triangles, gives

$$\operatorname{cofib}(W^{\times}_{r,\chi}(c) \to R_{r,\chi}) \in D_{\mathscr{I}_r}. \tag{6.11}$$

It remains to transport from T. At a positive wall λ, the two-step types have classes

$$(1, e, \chi_0),\ (0, f, \eta), \qquad e + f = r, \quad \eta = \lambda f, \quad \chi_0 + \eta = \chi.$$

There are finitely many such types. The complete coherent receiver over a wall-semistable object already has a pure λ-semistable unframed factor. For a subobject, every HN slope is at most λ, and its average is λ. For a quotient, each finite HN slope is at least λ, and

$$\lambda f = \chi(Q) = \operatorname{length} T_0(Q) + \sum_j f_j \mu_j \ge \operatorname{length} T_0(Q) + \lambda f.$$

Thus torsion is absent and all slopes equal λ.

The framed factor is wall-semistable as well. In $0 \to J \to I \to Q[-1] \to 0$, a high-slope subobject of J is one of I, and a low-slope unframed quotient of J gives, after extension by Q, a low-slope quotient of I. For the opposite order, a high-slope subobject of J lifts to an extension by Q of high slope, and a low-slope quotient is already a quotient of I. Conversely, intersections and images of an arbitrary subobject in these diagrams show, by weighted averages, that every extension of good wall factors is wall-semistable. An adjacent unstable framed factor splits off a positive-rank factor of slope λ and has smaller framed rank. Ordering by that rank, the proper receivers and Lemma 6.4 give the two localization sequences from the wall stack to the adjacent chamber opens.

If $\lambda > \max(\chi/r, 0)$, the wall stack has no $s = 0$ object: its vacuum subobject would have slope λ and its unframed quotient slope χ/r. A rank-zero framed factor would be the vacuum and

would force $\lambda = \chi/r$. Therefore both ranks are strictly between 0 and r, so all closed terms lie in $D_{\mathscr{J}_r}$. The octahedral triangle transports (6.11) through the finitely many walls between t and T.

Two cutoffs $c_1 < c_2 \le c_r$ contain the same stable-pair open, by Lemma 6.1. The restrictions have their ordinary composite law. An octahedral triangle applied to their common stable-pair restriction proves (6.10). For the full source, the rank-r unframed term is allowed in $\mathscr{S}_r$; every other term is in that category by smaller-rank induction and finite addition. The same proof gives its stated restriction property. □

For later induction, note that

$$c_e - c_r = \frac{(r-e)(\ell re - 1)}{2re} > 0 \qquad (1 \le e < r).$$

A lower-rank window appearing at the rank-r cutoff is compared to its own core by its nested restriction and the octahedral triangle.

6.5. **The comparison arrow and its endpoints.** Let $\mathscr{P}_\chi(c) = \mathscr{P}_{r,\chi}(c)$ be the derived surface nonzero-section window and $h : \mathscr{P}_\chi(c) \to \mathcal{A}_r^L$ its Hitchin map. Put

$$\Omega_{\mathscr{P}} = \omega_{t_0 \mathscr{P}} \mathbb{L}^{(\chi + B_r)/2}, \qquad g = g_{\mathcal{A}} \circ h.$$

Let $\pi_g : \mathrm{Crit}(\mathscr{P}, g) \to \mathscr{P}$ be the projection from the derived critical stack to the surface-pair stack. The physical critical presentation supplies the local comparison $\delta_{\mathscr{P}}$, constructed in (7.15) and (7.16). The comparison arrow is

$$\Delta_\chi(c) : \phi_{g_{\mathcal{A}}} R h_* \Omega_{\mathscr{P}} \xrightarrow{e_h} R h_* \phi_g \Omega_{\mathscr{P}} \xrightarrow{R h_* \delta_{\mathscr{P}}} R(h\pi_g)_* \Phi_{\mathrm{Crit}(\mathscr{P},g)}. \tag{6.12}$$

Here e_h is the forward vanishing-cycle exchange. The map h need not be proper. Define

$$\Sigma_\chi(c) = \Delta_\chi(c) \mathbb{L}^{\chi/2}, \qquad K_\chi(c) = \mathrm{cofib}\, \Sigma_\chi(c).$$

All restriction maps in Proposition 6.5 act on this same arrow and its cofiber.

Lemma 6.6 (The split framed reflection)**.** *Let $\mathscr{X}_{r,\chi}^t$ be the full t-semistable framed stack on $X_0 = \mathrm{Tot}_S(K_S)$, with its specified boundary framing. The zero section is retained. For finite t, derived duality induces an isomorphism over $\mathcal{A}_r^L$*

$$\mathsf{D} : \mathscr{X}_{r,\chi}^t \xrightarrow{\sim} \mathscr{X}_{r,-\chi}^{-t}.$$

It identifies the oriented Donaldson–Thomas coefficients of Section 4. Write $h_{r,\chi}^t : \mathscr{X}_{r,\chi}^t \to \mathcal{A}_r^L$ for the Hitchin map. If

$$\mathcal{B}_{r,\chi}^t = R(h_{r,\chi}^t)_* (\Phi_{\mathscr{X}_{r,\chi}^t} \otimes \Lambda^\chi),$$

then the normalized identity is

$$\mathcal{B}_{r,\chi}^t \simeq \mathcal{B}_{r,-\chi}^{-t} \otimes \mathbb{L}^\chi. \tag{6.13}$$

For $0 < \epsilon < 1/r$, both intrinsic chamber stacks $\mathscr{X}_{r,\chi}^{\pm\epsilon}$ are contained in the core window $\mathscr{X}_{0,r,\chi}(c_r)$.

Proof. Let $\mathfrak{M}_{r,\chi}$ be the derived stack of properly supported pure one-dimensional sheaves F on $S = \mathrm{Tot}_C(L)$, finite over C, whose curve pushdown has rank r and Euler characteristic χ. Purity and finiteness refer to the geometric fibers. Its universal sheaf is denoted by $\mathcal{F}$. For the projection $p : S \times \mathfrak{M}_{r,\chi} \to \mathfrak{M}_{r,\chi}$, put

$$\mathcal{C} = R p_* \mathcal{F}, \qquad \mathcal{E} = \mathcal{C}^\vee.$$

Since the curve pushdown is a vector bundle, $\mathcal{C}$ is perfect of relative amplitude $[0,1]$, and $\mathcal{E}$ has amplitude $[-1,0]$. Define the two derived linear stacks

$$\mathfrak{P}_{r,\chi}^\dagger = \mathrm{Spec}_{\mathfrak{M}_{r,\chi}} \mathrm{Sym}(\mathcal{E}), \qquad \mathfrak{P}_{r,\chi}^\sharp = \mathrm{Spec}_{\mathfrak{M}_{r,\chi}} \mathrm{Sym}(\mathcal{E}^\vee[1]). \tag{6.14}$$

Their test-base descriptions are, respectively,

$$(F, s : \mathcal{O}_{S\times T} \to F), \qquad (F, \eta : F \to K_S \boxtimes \mathcal{O}_T[1]).$$

These descriptions include the mapping spaces. Relative duality for the proper support supplies the second description and commutes with derived base change. One may form these stacks on the projective surface $\overline{S} = \mathbf{P}_C(\mathcal{O}_C \oplus L)$ and then restrict to sheaves whose support misses its infinity divisor. Thus they are the corresponding opens of the pair and dual-pair stacks of [69, §4.4.1, Lemmas 4.4.1–4.4.2].

Write

$$D_S F = R\mathcal{H}om_S(F, K_S)[1].$$

On ordinary fibers it is $\mathcal{E}xt^1_S(F, K_S)$ and is pure. The monic spectral resolution identifies its curve data with $(E^\vee K_C, \Theta^{\mathsf{t}})$. In particular it changes χ to $-\chi$, preserves the characteristic polynomial, and reverses the Euler slopes of the Higgs filtration. Relative Serre duality gives the base-change-compatible identity

$$Rp_* D_S \mathcal{F} \simeq \mathcal{C}^\vee[-1]. \tag{6.15}$$

It follows that D_S identifies $\mathfrak{P}^\sharp_{r,\chi}$ with $\mathfrak{P}^\dagger_{r,-\chi}$: the dual-section map η becomes the section $D_S\eta : \mathcal{O} \to D_S F$.

The remaining step exchanges the two cotangent presentations over the same $\mathfrak{M}_{r,\chi}$. The physical pure stack is the classical truncation of $T^*[-1]\mathfrak{P}^\dagger_{r,\chi}$ by Proposition 4.1. The same coherent diagram has a second projection to $\mathfrak{P}^\sharp_{r,\chi}$: it retains the extension class of

$$0 \longrightarrow \mathcal{O}_S \longrightarrow U \longrightarrow F K_S^{-1} \longrightarrow 0.$$

Over this dual pair, its other datum is the map $U \to F$. The tangent triangle and surface duality identify the latter datum with the shifted cotangent fiber. Thus both projections describe the same classical physical moduli functor, including families over nonreduced bases. This is also the construction of [69, Proposition 4.4.3 and §4.5.6].

Here is the required enhancement from that classical statement to an oriented critical statement. On a derived smooth chart write $\mathfrak{M} = Z^{\mathrm{der}}(\sigma : Y \to V)$ and $\mathcal{E} = [\mathcal{E}^{-1} \xrightarrow{d} \mathcal{E}^0]$. The two shifted cotangent presentations have the common critical chart

$$Z = \mathrm{Tot}_Y\big((\mathcal{E}^0)^\vee \oplus \mathcal{E}^{-1} \oplus V^\vee\big), \qquad w(y,u,v,z) = \langle d(v), u\rangle + \langle \sigma(y), z\rangle.$$

The resulting identification of classical truncations is the natural one of [69, Lemma 4.6.13]; it is independent of the displayed presentation. It also identifies the d-critical sections, because both sections are represented by this same function w.

The orientation comparison is global. Put $\mathscr{K}_{\mathfrak{M}} = \det \mathbb{L}_{\mathfrak{M}}$. The relative cotangent triangles identify both orientation roots, after pullback to the common physical stack, with

$$\mathscr{K}_{\mathfrak{M}} \otimes \det \mathcal{E}. \tag{6.16}$$

For the sharp root use the graded determinant isomorphism $\det(\mathcal{E}^\vee[1]) \simeq \det \mathcal{E}$. In the displayed chart this is exactly the same line

$$K_Y \det V \det \mathcal{E}^0 (\det \mathcal{E}^{-1})^{-1} = K_Z$$

for the two presentations. If $a = \mathrm{rk}\, \mathcal{E}^{-1}$ and $b = \mathrm{rk}\, \mathcal{E}^0$, their virtual dimensions are $\mathrm{vdim}\, \mathfrak{M} + (b-a)$ and $\mathrm{vdim}\, \mathfrak{M} - (b-a)$. Hence their canonical critical square signs agree. The determinant comparison in (6.16) therefore preserves the orientation squares. An acyclic refinement adds the same equation-variable hyperbolic block on both sides; determinant contraction gives the same root map and square map. Smooth atlas changes contribute the same atlas determinant on both sides. These comparisons descend to the global determinant isomorphism. The cotangent lift of D_S carries the canonical cotangent orientation by its cotangent derivative. Together with Proposition 4.4, this identifies the specified physical orientations. Functoriality of the oriented vanishing-cycle construction now identifies the coefficients.

We check the chamber opens. On the split curve chart a pure physical object has data

$$(E, \Theta, X, s, \varepsilon), \quad \Theta : E \to EL, \quad X : E \to EA, \quad s : \mathcal{O}_C \to E, \quad \varepsilon : E \to K_C, \quad [\Theta, X] + s\varepsilon = 0.$$

With the usual shifted-dual signs the preceding reflection is

$$(E, \Theta, X, s, \varepsilon) \longmapsto (E^{\vee} K_C, \Theta^{\mathsf{t}}, X^{\mathsf{t}}, \varepsilon^{\mathsf{t}}, -s^{\mathsf{t}}). \tag{6.17}$$

The transformed moment expression is the negative transpose of the old one. Its square is identified with the identity by the gauge transformation $-\mathrm{id}_E$. Equivalently this is the duality $R\mathcal{H}om_Y(I, \mathcal{O}_Y)$ with its induced boundary framing; the coherent-diagram comparison is [69, Lemma 4.5.14].

An unframed subobject is an invariant curve submodule $E' \subset E$ annihilated by ε; an unframed quotient is an invariant quotient $E \twoheadrightarrow E''$ annihilating s. To test semistability it suffices to use pure quotients and saturated subobjects. Indeed, removing the zero-dimensional torsion from an unframed quotient lowers its Euler slope, while saturation of an unframed subobject raises its slope and preserves both field invariance and the condition $\varepsilon|_{E'} = 0$. The complementary framed factor then has pure curve pushdown. Nonzero zero-dimensional subobjects are excluded at every finite t.

Reflection exchanges these pure exact sequences, reversing arrows, and sends their unframed Euler slopes μ to $-\mu$. Thus the subobject inequality $\mu \le t$ becomes the quotient inequality $-t \le -\mu$, and conversely. This proves the stability equivalence for every finite t. Since stability defines open substacks, the all-base isomorphism above restricts to the asserted isomorphism of chamber moduli functors. The same classical chamber duality is stated in [69, Lemma 4.4.6].

Since Θ is transposed without a sign, reflection is over the identity of $\mathcal{A}_r^L$. Ordinary direct image carries the coefficient isomorphism to an isomorphism over this base. Finally $\Lambda^{\chi} = \Lambda^{-\chi} \otimes \mathbb{L}^{\chi}$, which proves (6.13).

It remains to justify the cutoff assertion. The coherent projection sequence (6.3) applies to every finite-t semistable object. Its cyclic term has minimum Euler slope at least α_r. Every positive-rank HN quotient of the other term has slope at least t. For $t = \pm\epsilon$ and $0 < \epsilon < 1/r$, a finite such slope with denominator at most r is nonnegative whenever it is at least $-\epsilon$. Purity, the cyclic bound, and extension closure therefore give $\mu_{\min} \ge \alpha_r > c_r$. This proves containment for both complete intrinsic chamber stacks. □

Lemma 6.7. *The core surface window is empty for $\chi < a_r$. There is a constant M_r, independent of perverse degree, such that $K_{\chi}(c_r) \in D_{\mathscr{I}_r}$ for $\chi > rM_r$. For the ordinary source modulo $\mathscr{S}_r$, the upper vanishing endpoint is instead $\chi > -a_r$.*

Proof. Since ℓ is even, α_r is an integer. A rational slope of denominator at most r below α_r is at most $\alpha_r - 1/r$. Thus $\mu_{\min} > c_r$ is equivalent to $\mu_{\min} \ge \alpha_r$. If $\chi < a_r$, the average slope is at most $\alpha_r - 1/r < c_r$, so the whole core is empty. The same lower bound applies to both vertices of (6.12).

A rank-r semistable L-Higgs bundle has

$$\mu_{\min,\deg}(E) \ge \chi/r + h - (r-1)\ell.$$

Consequently $H^1(C, E) = 0$ when $\chi/r > h + (r-1)\ell$. Increase this bound to the projective framing threshold. Consider a semistable physical object in the immediate chamber $t = \chi/r + \epsilon$, $0 < \epsilon < 1/r^2$. If its projected Higgs bundle E were unstable, its highest Higgs Harder–Narasimhan factor H would satisfy $\mu(H) > \chi/r$. The denominators are at most r, so $\mu(H) > t$. The preceding bound, applied to the semistable bundle H, gives $H^1(C, H) = 0$. Serre duality then gives $\varepsilon|_H = 0$ for the dual framing $\varepsilon : E \to K_C$. The moment relation makes the off-diagonal transverse field $H \to (E/H)A$ a Higgs morphism. In the nonsplit row the local transverse lifts differ by a multiple of Θ, so their off-diagonal blocks agree. Negativity of A and the decreasing Higgs slopes force this morphism to vanish. Thus H is an unframed physical subobject of slope greater than t, a contradiction. Consequently E is semistable, $H^1(C, E) = 0$, and $\varepsilon = 0$. If $s = 0$, the vacuum subobject and its full unframed quotient would give $t \le \chi/r$, a contradiction. Thus $s \ne 0$. The same negative-A Hom vanishing preserves every equal-slope subobject. The remaining framed

stability condition is the Joyce–Song condition: the nonzero section is contained in no proper Higgs subbundle of the same slope. Denote this pair space by $\mathscr{P}^{\mathrm{JS}}$. The ordinary pair stack is therefore the smooth stable framed open, projective over the positive semistable Hitchin moduli; the projective framing construction is described in [56, Definition 2.11 and (2.37)]. The relative framing map is smooth because $H^1(E) = 0$, and the positive Higgs stack is smooth. The framed stability kills stabilizers. Thus its Hitchin map is proper and its two critical presentations are $T^*[-1]\mathscr{P}^{\mathrm{JS}}$ and $\mathrm{Crit}(\mathscr{P}^{\mathrm{JS}}, g)$ with the orientations of Section 4. On this open δ and the proper exchange e_h are isomorphisms. Proposition 6.5 then gives the upper endpoint for K_χ.

The ordinary-source endpoint is needed only in the split row. Use the full source, with the zero section retained, and work modulo $\mathscr{S}_r$. Its restriction to a sufficiently large positive chamber is an isomorphism in this quotient by Proposition 6.5. The same proper-receiver argument transports it across every wall between that chamber and $-\epsilon$, where $0 < \epsilon < 1/r$. Here one permits the rank-zero framed factor $\mathcal{O}$ and its single rank-r unframed complement. The latter has allowed support by the independent Proposition 8.2; every other framed factor has smaller rank and is covered by the source induction. At the zero wall the purity argument is unchanged: if the unframed quotient has average zero, its nonnegative finite slopes and its torsion length force all slopes and that length to vanish. All walls in this bounded interval have the finite numerical lists and complete coherent receivers constructed above. The transport uses the independent rank-r unframed support theorem and the smaller-rank source induction.

Lemma 6.6 places both immediate chambers in the core and identifies their normalized direct images by (6.13). If $\chi > -a_r$, then $-\chi < a_r$, so the whole positive immediate chamber in class $-\chi$ is empty by the lower endpoint already proved. Consequently $\mathcal{B}^{-\epsilon}_{r,\chi} = 0$, and the full ordinary source belongs to $D_{\mathscr{S}_r}$. This proves the upper source endpoint. □

7. The Hecke relation on the comparison arrow

The next construction proves a degree-zero Weyl relation on the whole arrow (6.12), including its square homotopy. It will also supply the ordinary-source relation used to prove support before any Hilbert-scheme calculation.

7.1. **Directed bounded presentations.** Fix a finite interval of Euler characteristics, including its two neighboring values, and a core cutoff c. For $j = 0, 1, 2$, choose nested bounded smooth opens $\mathscr{B}_{\chi,j}$ of $\mathrm{Bun}_r(C)$. The level-zero open contains the bundles of the core. The level-one open contains the cutoff-$c - 1$ core, level zero, and every elementary neighbor of a level-zero target. Level two contains the cutoff-$c-2$ core, level one, and every neighbor of a level-one target. Each enlargement exists: elementary modifications over a bounded family are parametrized by a projective curve-Hecke scheme, whose image is bounded. The directed edges run only from level two to one or from level one to zero.

Choose a sufficiently large auxiliary divisor D of degree m after this finite system is fixed. On each $\mathscr{B}_{\chi,j}$, let E be the universal bundle and π the universal-curve projection. Relative Serre vanishing, including evaluation at the active marks, makes

$$V_\Theta = \pi_* \mathcal{E}nd(E)L(D), \qquad R_\Theta = \pi_*(\mathcal{E}nd(E)L(D)|_D),$$
$$V_s = \pi_* E(D), \qquad R_s = \pi_*(E(D)|_D)$$

vector bundles. The smooth ambient is

$$X_\chi = \mathrm{Tot}(V_\Theta \oplus V_s \oplus R_\Theta^\vee \oplus R_s^\vee), \qquad e = mr(r+1), \quad d_\chi = \chi + B_r + 2e. \tag{7.1}$$

Extend the functional α_q to a functional β_D on $H^0(L^2(2D))$. With residue-dual coordinates ξ, η, its function is

$$F_\chi = \langle \mathrm{res}\,\Theta, \xi \rangle + \langle \mathrm{res}\, s, \eta \rangle + \tfrac{1}{2}\beta_D(\mathrm{Tr}\,\Theta^2). \tag{7.2}$$

Let S_χ be the residue-zero locus, retaining the dual variables. Delete only the closed bad part of S_χ for the specified cutoff. Thus every nonzero-residue point remains in the ambient. For (6.12) also take the nonzero meromorphic-section open; on S_χ it is exactly the nonzero holomorphic-section condition.

For a Higgs elementary modification $0 \to E_- \to E_+ \to Q_x \to 0$,

$$\mu_{\min}(E_+) - 1 \le \mu_{\min}(E_-) \le \mu_{\min}(E_+). \tag{7.3}$$

Extend a pure quotient of E_- to E_+ and remove its torsion to prove the first inequality; intersect a minimum-slope quotient of E_+ with E_- to prove the second. The change of Euler characteristic is at most one. Near D the modification is an isomorphism. Hence residues agree at both ends: either both are zero or neither is. In the former case (7.3) supplies one unit of cutoff buffer; in the latter every ambient cutoff accepts the point. The section inclusion is an inclusion of vector bundles with line quotient, so its inverse image of the upper zero section is the lower zero section over every base.

It follows that each one-step space is the full inverse image of its target window along its proper arm. Properness is preserved by this base change. The directed bundle opens ensure that the source is in the required preceding level. In a two-step word ending at $(\chi, 0)$, both possible middle objects lie in level one and every initial object lies in level two. The operators are iterated in the quotient category using these three levels.

7.2. **Choices over the whole marked curve.** Choose three effective divisors D_1, D_2, D_3 of the same sufficiently large degree and with pairwise disjoint supports. Put

$$U_i = (C \setminus \operatorname{supp} D_i)^2 \subset C_x \times C_y.$$

These symmetric opens cover $C_x \times C_y$: two marked points can meet at most two of the three disjoint supports. Over U_i use D_i as the auxiliary divisor. Increase their common degree so that all the preceding vanishings hold also for the thick union $\Gamma_x + \Gamma_y$, including on the diagonal.

Let B_i be the affine space of extensions β_{D_i} of α_q, and set $T_i = U_i \times B_i$ and $T = \coprod_i T_i$. Thus $T \to C_x \times C_y$ is a smooth surjective presentation. Full frames along D_i are retained separately by their smooth action groupoids.

For a nonempty index set I, on $U_I = \bigcap_{i \in I} U_i$ use the common enlargement $D_I = \sum_{i \in I} D_i$. The even hyperbolic stabilization comparisons proved below identify the local arrows and their binary composition paths. For $I \subset J \subset K$, the comparison for $D_I \subset D_K$ agrees with the composite through D_J: after an affine torsor of splittings this is the direct-sum and determinant-contraction identity for the added acyclic blocks. Its higher compatibilities are the corresponding identities for successive direct sums. Consequently the local arrows and paths descend over this symmetric open cover before the proper curve trace is taken.

In a local lattice chart the modification is $D_x = \operatorname{diag}(t - x, 1, \dots, 1)$. Transport the full frame by $D_x|_D$, which is invertible. Then the two residues and their dual coordinates agree exactly. The intertwining equation $\Theta_+ D_x = D_x \Theta_-$ also gives

$$(t - x)^2 (\operatorname{Tr} \Theta_+^2 - \operatorname{Tr} \Theta_-^2) = 0.$$

The graph parameter is a relative non-zero-divisor, so it cancels even over a nonreduced base. Thus the trace potentials and characteristic polynomials agree on the entire curve. Both words have the same function and the same residue-zero ideal.

7.3. **The two derived words.** All stacks in this subsection are relative to the auxiliary parameter space T, while it is present. Put $S_D = \operatorname{Tot}_C L(D)$ and let $p : S_D \to C$ be its projection. The section source is the line bundle

$$R_D = p^* \mathcal{O}_C(-D), \qquad R\operatorname{Hom}_{S_D}(R_D, F) = R\operatorname{Hom}_C(\mathcal{O}_C(-D), p_* F). \tag{7.4}$$

Let $\mathfrak{P}_D$ be the derived stack of properly supported perfect spectral objects whose classical fibers are length-one sheaves. Its classical truncation is $P_D = S_D \times B\mathbf{G}_m$; a point records a support (x, λ) and a line ℓ. Denote its truncation map by $\iota_D : P_D \to \mathfrak{P}_D$.

In the stable category $\mathrm{Fun}(\Delta^1, \mathrm{Perf}(S_D))$, let $\mathfrak{H}^{\mathrm{der}}_{\chi,D}$ classify exact triangles

$$(R_D \to F_-) \longrightarrow (R_D \to F_+) \longrightarrow (0 \to Q),$$

whose source map is id_{R_D}, whose endpoint curve pushdowns are bundles of Euler characteristics $\chi, \chi+1$ in the prescribed bounded opens, and whose last term lies in $\mathfrak{P}_D$. The cofiber functor defines $c_Q : \mathfrak{H}^{\mathrm{der}}_{\chi,D} \to \mathfrak{P}_D$. Define the one-step stack by

$$\mathfrak{H}_{\chi,D} = \mathfrak{H}^{\mathrm{der}}_{\chi,D} \underset{\mathfrak{P}_D}{\overset{\mathbf{R}}{\times}} P_D. \tag{7.5}$$

Retain the common full frame along D and the common dual-residue coordinates. The endpoints identify these data along D, since the support of Q is disjoint from D. They supply smooth factors. Restrict the proper target window and take its full inverse image. The point map to P_D specifies the active mark; its restriction at x is denoted by $\mathfrak{H}_{\chi,x}$.

The cotangent triangle of (7.5) is

$$\mathbb{L}_{\mathfrak{P}_D}|_{\mathfrak{H}_{\chi,D}} \longrightarrow \mathbb{L}_{\mathfrak{H}^{\mathrm{der}}_{\chi,D}}|_{\mathfrak{H}_{\chi,D}} \oplus \mathbb{L}_{P_D}|_{\mathfrak{H}_{\chi,D}} \longrightarrow \mathbb{L}_{\mathfrak{H}_{\chi,D}}. \tag{7.6}$$

Its first component is the cotangent map of the specified cofiber functor. The following calculation identifies this map on the obstruction line.

The stack $\mathfrak{H}_{\chi,D}$ is smooth. To compute the obstruction map, use a local lattice chart. The arrow-field differential is

$$(B_-, B_+) \longmapsto D_x B_- - B_+ D_x.$$

Its (i,j) image is the ideal generated by the corresponding diagonal entries of D_x. Every entry is surjective except $(1,1)$, whose cokernel is the line $\mathcal{O}/(t-x)$ tensored by the Higgs line. For the spectral resolution $w - \Theta$, the cone of an arrow has differentials

$$E_- \xrightarrow{(w-\Theta_-, -D_x)} E_- \oplus E_+ \xrightarrow{(D_x, w-\Theta_+)} E_+.$$

Put $f = t - x$. The intertwining relation gives block matrices

$$\Theta_- = \begin{pmatrix} a & b \\ f\gamma & d \end{pmatrix}, \qquad \Theta_+ = \begin{pmatrix} a & fb \\ \gamma & d \end{pmatrix}.$$

Let $K(f, w-a)$ denote the three-term Koszul complex with maps $(w-a, -f)^{\mathsf{t}}$ and $(f, w-a)$ in degrees $-2, -1, 0$. An explicit map from the displayed cone to it is

$$\pi^0(z) = z_1, \qquad \pi^{-1}(u, v) = (u_1 - b v_{\mathrm{rest}}, v_1), \qquad \pi^{-2}(z) = z_1,$$

where v_{rest} comprises the last $r-1$ coordinates. Indeed,

$$\pi^0 d^{-1}(u, v) = f(u_1 - b v_{\mathrm{rest}}) + (w-a) v_1,$$
$$\pi^{-1} d^{-2}(z) = ((w-a) z_1, -f z_1).$$

The remaining $r-1$ unit blocks are contractible. Replacing $a(t)$ by $\lambda = a(x)$ is an elementary column operation, since $a(t) - a(x)$ is divisible by f. Thus the cone contracts to $K(t-x, w-\lambda)$. Its curvature is $\Theta_+ D_x - D_x \Theta_- = -(D_x \Theta_- - \Theta_+ D_x)$; the differential shift in the tangent complex $\mathrm{End}(\mathrm{Cone})[1]$ reverses this sign. Consequently the remaining arrow obstruction maps by positive evaluation modulo $t-x$ to the top obstruction class of the exterior point. Intrinsically this is the isomorphism

$$\mathrm{End}(\ell) \otimes T_x C \otimes L(D)_x \xrightarrow{\sim} \mathrm{Ext}^2_{S_D}(Q, Q) = \mathrm{End}(\ell) \otimes \bigwedge^2 T_{(x,\lambda)} S_D.$$

It is the derivative of c_Q, identified by the Koszul contraction. Since $P_D \to \mathfrak{P}_D$ induces isomorphisms in tangent degrees $-1, 0$, the tangent fiber of (7.6) removes this obstruction

without adding a tangent direction. The remaining field and section complexes have no H^1 by the choice of D, including evaluation at the mark, and are vector bundles. This proves smoothness. The map exists globally as the derivative of c_Q; the local calculation therefore commutes with coordinate changes and derived base change. The $r-1$ flag dimensions cancel the $r-1$ field-preservation conditions, leaving one-step dimension d_χ at a fixed mark.

Let $X_{\chi,j,T}$ denote the meromorphic ambient at bundle level j, including the parameter space T. Suppressing only D and the bundle levels in the next display, define the two words by

$$W_+ = \mathfrak{H}_{\chi,x} \overset{\mathbf{R}}{\underset{X_{\chi+1,T}}{\times}} \mathfrak{H}_{\chi,y}, \qquad W_- = \mathfrak{H}_{\chi-1,y} \overset{\mathbf{R}}{\underset{X_{\chi-1,T}}{\times}} \mathfrak{H}_{\chi-1,x}. \tag{7.7}$$

They are the upper and lower diagrams

$$E_1 \longrightarrow U \longleftarrow E_2, \qquad E_1 \longleftarrow J \longrightarrow E_2$$

in the stable arrow category, with identity maps on the source R_D. Let $\mathfrak{D}$ be the marked endpoint graph, on which the marks and complete endpoints agree.

For the upper word form $J = E_1 \times^{\mathbf{R}}_U E_2$. The square is bicartesian and its source is R_D. It has canonical cofiber identifications

$$\operatorname{cofib}(J \to E_1) \simeq \operatorname{cofib}(E_2 \to U) = Q_y, \qquad \operatorname{cofib}(J \to E_2) \simeq \operatorname{cofib}(E_1 \to U) = Q_x.$$

The stable pushout recovers U. Finite limits and colimits supply the inverse functors and their unit and counit on every connective base, including mapping spaces and higher homotopies. These identifications commute with the maps to $\mathfrak{P}_D \times \mathfrak{P}_D$ after transposing the labels. They consequently commute with the base change from $P_D \times P_D$ in (7.5).

Away from $\mathfrak{D}$, the map $U \to Q_x \oplus Q_y$ is surjective and its kernel is a curve bundle. Failure forces coincident quotient directions, hence coincident marks and complete endpoints, including fields, sections, and frames. Conversely a torsion submodule of the lower pushout injects into both exterior quotients and has the same consequence. The curve-bundle condition is open relative Tor-amplitude. Thus the two functors restrict to inverse equivalences of the required derived open stacks off $\mathfrak{D}$.

Their mutual collision obstruction is retained. Its local complex is

$$\operatorname{Hom}(\ell_x, \ell_y) \otimes \left[\mathcal{O} \xrightarrow{(u,v)^t} \mathcal{O}^2 \xrightarrow{(-v,u)} \mathcal{O}\right], \quad u = y - x, \quad v = \lambda_y - \lambda_x. \tag{7.8}$$

The standard Koszul contractions apply on $u \neq 0$ or $v \neq 0$; the complex is retained at $u = v = 0$.

7.4. **A proper carrier and its supported class.** Choose the degree of the full frame on D large enough that

$$H^0(C, \mathcal{H}om(E, F)(-D)) = 0 \tag{7.9}$$

for every ordered pair in the finite bounded families. This depends only on the degree, so holds uniformly for the moving, possibly nonreduced divisor. The full framed endpoints are separated algebraic spaces. To see the separatedness assertion, extend a generic framed isomorphism over a DVR after multiplying by a minimal power τ^m of the uniformizer. If $m > 0$, its special fiber vanishes on D and contradicts (7.9); hence $m = 0$. Its determinant is a section of a degree-zero line, nonzero on D, and therefore nowhere zero on the special fiber. It is an isomorphism, and uniqueness follows from torsion-freeness over the DVR. Fields and sections extend compatibly for the same reason. Universally, the diagonal can also be described as the zero locus of the boundary obstruction of the prescribed frame in $R^1\pi_*\mathcal{H}om(E, F)(-D)$; the vanishing of R^0 makes the extension unique. This proves the closed-diagonal statement over arbitrary bases and also removes stabilizers.

In the framed product of the initial level-two and final level-zero endpoints over T, take the scheme-theoretic closed union $\widetilde{Z}$ of the two word images and the marked graph $\widetilde{\mathfrak{D}}$. A word is proper over its final endpoint; its map to the product is proper because the initial endpoint is

separated. Each image, their union, and the graph are therefore proper over the final endpoint. Quotient by the common diagonal frame group $G = \mathrm{Res}_{D/T}\,\mathrm{GL}_r$:

$$Z = [\widetilde{Z}/G], \qquad \mathfrak{D} = [\widetilde{\mathfrak{D}}/G].$$

The right map $b : Z \to X_{\chi,0,T}$ is representable proper, as is checked after the final frame atlas. The parameter T is retained in both endpoints and in this proper pushforward. The intrinsic maps and homotopies will descend along $T \to C \times C$ before the proper curve trace is applied. The graph $\mathfrak{D}$ is smooth of dimension $q = d_\chi + \dim T - 1$. Smooth descent applies to G, including its unipotent part. The equalities of functions and residue coordinates descend to the scheme-theoretic images and their union. If a is the left map, then on the carrier itself

$$a^* F_{\mathrm{initial}} = b^* F_{\mathrm{final}}, \qquad a^{-1} S_{\mathrm{initial}} = b^{-1} S_{\mathrm{final}} \tag{7.10}$$

as functions and closed schemes.

For ordinary Hodge Borel–Moore theory use the mapping spectrum

$$\mathrm{BM}_q(Z) = \mathrm{Map}(\mathbf{Q}_Z, \omega_Z \mathbb{L}^q).$$

Each derived word is a derived base change of the regular diagonal of the smooth framed middle endpoint. Both are quasi-smooth of virtual dimension q. The fundamental class is the coordinate-class specialization, localization boundary, and inverse Thom map; its composition and base-change paths are those of [34, §2, Constructions 3.1, 3.4, 3.6 and Theorems 3.12–3.13]. In this regular-diagonal situation the binary path is also the double-deformation diagram of [19, Sections 3.2.19–3.2.21]. Swapping the two odd coordinate classes contributes -1, and swapping the two odd localization boundaries contributes -1. They cancel. This fixes the binary path, including its sign.

Transport the class by the nil-invariant equivalence induced by $t_0 W \to W$ and properly push the classical truncation to Z. The bicartesian equivalence above identifies the two classes off $\mathfrak{D}$, with these binary paths. Localization lifts their difference to $\mathrm{BM}_q(\mathfrak{D})$. Smooth purity identifies its underlying mapping space with

$$\mathrm{Map}(\mathbf{Q}, R\Gamma(\mathfrak{D}, \mathbf{Q})).$$

There is no negative cohomology, so the space is discrete with components $H^0(\mathfrak{D}, \mathbf{Q})$, one scalar per connected component. This reduction takes place in Hodge theory.

On the dense source open where $\Theta(y)$ is regular semisimple and every eigencomponent of $s(y)$ is nonzero, removal first is impossible. Creation followed by removal has one branch for each eigenline. On its jth chart the condition is exactly

$$(y - x) s_j(y) = 0. \tag{7.11}$$

The second factor is a unit, so the branch is smooth of dimension q. Its quasi-smooth virtual dimension is also q; hence the obstruction space vanishes there, and its fundamental class has complex multiplicity $+1$. Evaluation surjectivity for the meromorphic bundles makes this open meet every component of the smooth graph. The eigenline cover has degree r, and its trace is r. We conclude

$$[W_+] - [W_-] = r[\mathfrak{D}] \quad \text{in } \pi_0\,\mathrm{BM}_q(Z), \tag{7.12}$$

with the specified binary composition paths. The localization lift is locally constant, so the dense calculation determines it.

7.5. **Lifting the class to the entire arrow.** Let $X \xleftarrow{a} Z \xrightarrow{b} Y$ be such a carrier with b proper and X, Y smooth. If $d = \dim X$, dualizing purity and proper adjunction give

$$\begin{aligned} \mathrm{BM}_q(Z) &\simeq \mathrm{Map}(a^*\omega_X, b^!\omega_Y \mathbb{L}^{q-d}) \\ &\simeq \mathrm{Map}(b_* a^* \omega_X, \omega_Y \mathbb{L}^{q-d}). \end{aligned} \tag{7.13}$$

Pushing a class properly inserts the pullback unit followed by the proper counit. Thus (7.13) represents the original convolution maps and their binary paths.

Let $i_X : S_X \hookrightarrow X$ and $i_Y : S_Y \hookrightarrow Y$ be the residue-zero inclusions, and put $\Gamma_X = i_{X,*}i_X^!$ and $\Gamma_Y = i_{Y,*}i_Y^!$. Write K_X, K_Y for the ambient dualizing normalizations, absorbing the indicated Tate factors. Put $A_0 = b_*a^*\Gamma_X K_X$ and $A_1 = b_*a^*K_X$. By (7.10), A_0 is supported on S_Y. The mapping spectrum of squares to $\Gamma_Y K_Y \to K_Y$ is

$$\operatorname{Map}(A_0, \Gamma_Y K_Y) \underset{\operatorname{Map}(A_0,K_Y)}{\overset{h}{\times}} \operatorname{Map}(A_1, K_Y). \tag{7.14}$$

The first map to the base is an equivalence, since Γ_Y is the coreflector to supported objects. Consequently (7.14) projects equivalently to $\operatorname{Map}(A_1, K_Y)$. Its inverse produces the supported map and its square homotopy with contractible choice. It is additive and preserves the binary composition paths.

Write ϕ_X, ϕ_Y, ϕ_Z for vanishing cycles of the common function on the three spaces, and a_0, b_0 for the restrictions of a, b to its zero fibers. Apply vanishing cycles to the entire square. Its source is $\phi_X\Gamma_X K_X$. The critical correspondence uses

$$b_{0,*}a_0^*\phi_X K_X \longrightarrow b_{0,*}\phi_Z a^*K_X \xrightarrow{\sim} \phi_Y b_*a^*K_X \longrightarrow \phi_Y K_Y.$$

The first exchange is used only in its forward direction; only the proper exchange along b is inverted.

In the residue presentation, let V be the smooth ambient, $S = V|_{t_0\mathscr{P}}$, $i : S \hookrightarrow V$, and $p : S \to t_0\mathscr{P}$ the dual-residue vector bundle. Use $\mathscr{D}_\chi = \omega_V \mathbb{L}^\chi$. The supported arrow is

$$\phi_F(i_*i^!\mathscr{D}_\chi) \longrightarrow \phi_F \mathscr{D}_\chi. \tag{7.15}$$

Both vertices are supported on S, because derivatives in the dual residue variables are the residue equations. The source identification uses

$$i^!\omega_V = \omega_S = p^*\omega_{t_0\mathscr{P}}\mathbb{L}^{-e}, \qquad \phi_{p^*g}p^* \simeq p^*\phi_g, \qquad \mathrm{id} \simeq Rp_*p^*.$$

If $d_0 = B_r + 2e$, its pushdown is

$$Rp_*\phi_{p^*g}i^!\mathscr{D}_\chi = \phi_g\omega_{t_0\mathscr{P}}\mathbb{L}^{\chi-e} = \phi_g\Omega_{\mathscr{P}}\mathbb{L}^{\chi/2}\mathbb{L}^{-d_0/2}. \tag{7.16}$$

Its target is $R\pi_{g,*}\Phi_{\mathrm{Crit}(\mathscr{P},g)}\mathbb{L}^{\chi/2}\mathbb{L}^{-d_0/2}$. Thus (7.15) is the local δ in (6.12), with the same common factor at both vertices. Closed base change identifies its supported correspondence with the ordinary $\Omega_{\mathscr{P}}$ correspondence after this vector-bundle pushdown.

Finally the nonproper exchange e_h is the right mate of the forward pullback exchange. If $h : X \to B$ and $h_0 : X_0 \to B_0$ are the maps before and after the zero fiber, then on K the mate is

$$\phi_B h_*K \longrightarrow h_{0,*}h_0^*\phi_B h_*K \xrightarrow{\mathrm{Ex}_h} h_{0,*}\phi_X h^*h_*K \longrightarrow h_{0,*}\phi_X K.$$

The outer maps are the pullback unit and counit. Composition of forward exchanges and the adjunction triangle identities determine its composition homotopies. For the ordinary action of a correspondence u, use those same homotopies in

$$T(u) : h_{X,*}K_X \xrightarrow{\eta_a} (h_X a)_*a^*K_X = h_{Y,*}b_*a^*K_X \xrightarrow{u} h_{Y,*}K_Y,$$

its mate identities give

$$e_{h_X a} = (h_{X,*}e_a)(e_{h_X}a_*), \qquad e_a\phi_X(\eta_a) = a_*\,\mathrm{Ex}_a(\eta_a\phi_X).$$

They identify the square for e_h with the square already obtained for the local δ. The complete diagram has rows

$$\phi_{g_{\mathcal{A}}}Rh_*\Omega \longrightarrow Rh_*\phi_g\Omega \longrightarrow R(h\pi_g)_*\Phi,$$

and includes the two square homotopies. Composing the horizontal arrows is an exact functor on enhanced diagram categories. It takes (7.12) to a relation on precisely Δ, and then on its cofiber.

7.6. **Independence of the auxiliary presentation.** Enlarging the divisor adds an acyclic equation-variable block $[W \xrightarrow{d} E]$. On a splitting torsor, a triangular change of equations and its dual reduce the added potential to $\langle d(c), \zeta\rangle$ with support $c = 0$. Differences between extensions β_D belong to the residue ideal: the principal-part expansion of $\operatorname{Tr}\Theta^2$ expresses them as linear combinations of the residue coordinates. Translating the corresponding dual variables removes them.

For one hyperbolic pair write $c = z_2 + \mathrm{i}z_1$, $\zeta = z_2 - \mathrm{i}z_1$. Then $c\zeta = z_1^2 + z_2^2$. Projection of the positive real thimble to ζ preserves orientation, whereas projection to c reverses it. Thus the raw support-counit scalar for a rank-k block in the compact dimensional-reduction convention is $(-1)^k$. Equivalently this is $\binom{2k}{2} \equiv k \pmod 2$ in the duality normalization of [36, Lemma 3.4]. The grouped orientation form is

$$(-1)^{k(k-1)/2}(-\mathrm{i}/2)^k dc_1 \cdots dc_k\, d\zeta_1 \cdots d\zeta_k.$$

For invertible d, the contraction root contains $\det d$; its square is the Hessian determinant with this same parity. Changes of bundle frame act inversely on the dual half, so the total determinant is one. In our presentations $e = mr(r+1)$ and every enlargement rank are even. Hence the raw stabilization scalar is $+1$, and the existing local δ and its binary paths are preserved.

Splittings and potential extensions form affine torsors. On the fixed intrinsic endpoint objects their pullback is fully faithful, by $\mathrm{id} \simeq Rp_*p^*$ and the projection formula. The preceding transformations therefore descend the maps and their homotopies. To compare a previous one-point divisor to the two-point family, use their sum, which need only avoid the active mark of that one operator. Descend the one-step comparison first, then compose on the whole $C \times C$.

Changing bounded bundle opens is handled by supported open excision of (7.15). Both vertices are supported on the complete residue-zero window. Restriction to an open containing it is fully faithful on the objects and their mapping spaces. On a Hecke diagram first restrict the proper target, then take its full inverse image. If a further source excision is needed, delete the proper image of its bad closed inverse image; that image misses the target residue support. This gives a common refinement preserving the proper arm, fundamental class, coreflector square, and binary path. The three directed presentations consequently define the same intrinsic arrows and restriction squares.

Proposition 7.1. *There are degree-zero creation and removal morphisms on the arrows $\Sigma_\chi(c)$,*

$$E_\chi^c : \Sigma_\chi(c-1) \longrightarrow \Sigma_{\chi+1}(c), \qquad F_\chi^c : \Sigma_{\chi+1}(c-1) \longrightarrow \Sigma_\chi(c),$$

natural for the cutoff restrictions, with relation

$$F_\chi^c E_\chi^{c-1} - E_{\chi-1}^c F_{\chi-1}^{c-1} = r\rho_\chi(c-2, c). \tag{7.17}$$

The same holds on $K_\chi(c)$ and on the normalized full ordinary source

$$\mathcal{A}_\chi(c) = Rh_*\omega_{t_0\mathscr{P}_{r,\chi}^{\mathrm{all}}(c)}\mathbb{L}^{\chi+B_r/2}, \tag{7.18}$$

where the zero section is allowed.

Proof. At a fixed mark creation has target $\Sigma_{\chi+1}\mathbb{L}^{-1}$; removal already has degree zero. Indeed the ambient dimensions in (7.1) and (7.13), together with the common factor $\mathbb{L}^{-d_0/2}$ in (7.16), give these twists. The class relation on $C_x \times C_y$ is therefore

$$f_\chi^c(y)e_\chi^{c-1}(x) - e_{\chi-1}^c(x)f_{\chi-1}^{c-1}(y) = r\rho_\chi(c-2, c) \boxtimes \mathrm{cl}(\Delta_C).$$

Fix y and integrate the moving mark x by the proper trace of the whole curve $R\Gamma(C, \mathbf{Q}) \to \mathbb{L}$. It cancels $\mathbb{L}^{-1}$ and sends $\mathrm{cl}(y)$ to 1. This gives (7.17). The graph term is the inclusion of the level-zero endpoint into level two, so its action is the restriction ρ. Naturality follows by proper target base change and the preceding supported excision. The enhanced cofiber functor gives the same relation on K_χ.

For the ordinary source use the same smooth meromorphic ambient with $s = 0$ allowed. The regular-semisimple open used in (7.11) is still dense in every component. Apply the coreflector but not vanishing cycles. Its pushed coefficient is $\omega_{t_0\mathscr{P}}\mathbb{L}^{\chi-e}$, differing from (7.18) by the single common factor $\mathbb{L}^{-B_r/2-e}$. This cancels on every edge. The same trace gives the ordinary-source relation, without requiring the holomorphic pair stack itself to be smooth. □

7.7. **Finite-band vanishing.**

Lemma 7.2. *Let $V = \bigoplus_{a\le\chi\le b} V_\chi$ in a $\mathbf{Q}$-linear abelian category. If degree-zero morphisms E, F raise and lower χ and satisfy $[F, E] = r\,\mathrm{id}$ for an integer $r > 0$, then $V = 0$.*

Proof. E is nilpotent, and $[F, E^m] = \sum_{j=0}^{m-1} E^j[F, E]E^{m-1-j} = mrE^{m-1}$. Starting with $E^m = 0$ and decreasing m gives $\mathrm{id} = 0$. □

Theorem 7.3. *For every r, χ and $c \le c_r$, the cofiber of the original comparison arrow* (6.12) *belongs to $D_{\mathscr{I}_r}$.*

Proof. Fix a perverse degree j, and let $q_r : \mathrm{MMHM}(\mathcal{A}_r^L) \to \mathrm{MMHM}(\mathcal{A}_r^L)/\mathscr{I}_r$ denote the Serre-quotient functor on the perverse heart. Proposition 6.5 makes every cutoff restriction an isomorphism. Use its unique inverse to transport the operators in Proposition 7.1 to

$$V_\chi = q_r {}^p\mathcal{H}^j K_\chi(c_r).$$

Restriction naturality makes their relation $[F, E] = r\,\mathrm{id}$. Lemma 6.7 leaves a finite characteristic band, independent of j. Lemma 7.2 gives $V_\chi = 0$ for all χ. All cutoffs follow by their restrictions. Only the finite band and its adjacent values were needed for the three directed bundle levels; iteration of E now occurs inside the quotient category. □

Theorem 7.4. *Every perverse constituent of the normalized split stable-pair direct image belongs to $\mathscr{S}_r$.*

Proof. We prove by induction on r the support assertion for every full ordinary window $\mathcal{A}_\chi(c)$, $c \le c_r$, and for its complete chamber and wall coefficients with $t > -1/r$. The rank-zero framed factor is the vacuum. The positive unframed factors have allowed supports by Proposition 8.2. In the full-source Harder–Narasimhan restrictions, every smaller framed-rank factor has allowed supports by this induction. For $e < r$, the interval $t > -1/r$ is contained in $t > -1/e$; nested cutoff restriction compares a lower-rank factor at cutoff c with its own core. Finite spectral addition sends the concatenated partition supports to the corresponding Σ_ν. The only one-factor rank-r unframed term is exactly the independent positive input. Thus source restrictions are isomorphisms in the quotient by $\mathscr{S}_r$. The lower empty endpoint and the source reflection of Lemma 6.7 leave the band $[a_r, -a_r]$. Apply the ordinary-source relation of Proposition 7.1 and Lemma 7.2 in each perverse degree. This proves support for the full core source, and nested restriction gives it for every $c \le c_r$.

For $t > -1/r$, the slope-denominator argument of Lemma 6.6 and the cyclic bound place the complete chamber and wall stacks inside the core. Their restrictions from the full window are successive closed-stratum costalk triangles from the proper-receiver wall argument in the proof of Proposition 6.5. Each closed term has either a smaller positive framed rank, or the framed vacuum and a rank-r unframed factor. The former is covered by the induction hypothesis; the latter by Proposition 8.2. Finite addition and the Serre property place every closed term in $D_{\mathscr{I}_r}$. The triangles therefore give support for these chamber and wall coefficients, completing the induction. Dimensional reduction identifies the ordinary source (7.18) with the split physical coefficient. Taking a sufficiently large positive chamber gives the stable-pair assertion. □

Corollary 7.5. *Let $R_{0,r,\chi}$ and $R_{1,r,\chi}$ be the normalized stable-pair direct images from X_0 and X_1. Then*

$$[R_{1,r,\chi}] - [\phi_{g_{\mathcal{A}}} R_{0,r,\chi}] \in K_0(\mathscr{I}_r). \tag{7.19}$$

Proof. The available morphisms are the original window comparison $\phi W_0 \to W_1$ and the two restrictions $\phi W_0 \to \phi R_0$, $W_1 \to R_1$. Their cofibers lie in $D_{\mathscr{I}_r}$ by Theorem 7.3 and Proposition 6.5. They form a zigzag of isomorphisms in each perverse quotient. The stable-pair endpoints are bounded proper direct images by Proposition 11.5. Taking their Grothendieck classes and applying the Serre localization sequence gives (7.19). □

8. The positive source identity

We first compute the split stable-pair coefficient on X_0. Let $h_{0,r,\chi} : P_{r,\chi}(X_0) \to \mathcal{A}_r^L$ be the positive spectral map, and set

$$R^L_{r,\chi} = R(h_{0,r,\chi})_*(\Phi_{P_{r,\chi}(X_0)} \otimes \Lambda^\chi). \tag{8.1}$$

Dimensional reduction and the ordinary-source argument of Theorem 7.4 put every perverse cohomology object of (8.1) in $\mathscr{S}_r$.

Choose any Euler characteristic χ_* coprime to r. Write $M_L(r,\chi_*)$ for the smooth stable positive Higgs moduli space, and let $h_L : M_L(r,\chi_*) \to \mathcal{A}_r^L$ be its Hitchin map. Put

$$P_i^L = {}^p\mathcal{H}^i R(h_L)_* \mathbf{Q}_{M_L(r,\chi_*)}[B_r + 1]. \tag{8.2}$$

This Hodge module is independent of χ_* by [44, Theorem 0.4]. Define

$$b_r = \frac{\ell r(r+1)}{2} - rh, \qquad \gamma_r = 1 + rh + \frac{\ell r(r-1)}{2}. \tag{8.3}$$

Thus $\dim \mathcal{A}_r^L = b_r$, γ_r is the arithmetic genus of a rank-r spectral curve, and $b_r + \gamma_r = B_r + 1$. Let x be the formal variable recording Euler characteristic.

Theorem 8.1. *The complexes $R^L_{r,\chi}$ are ordinary pure Hodge complexes of weight zero. In the coefficientwise completed Grothendieck group of mixed Hodge modules on the positive Hitchin bases,*

$$1 + \sum_{r\ge1,\chi} R^r x^\chi [R^L_{r,\chi}] = \mathrm{Exp}_{\mathrm{add}}\left(\sum_{r\ge1} R^r \sum_i \sum_{u,v\ge0} x^{i+1+u+v} [P_i^L[-i-1-2v] \otimes \mathbf{T}^{v-B_r/2}] \right). \tag{8.4}$$

Only u, v are constrained to be nonnegative; i runs through the finite perverse range in (8.2)*. The exponential includes the derived symmetric-group actions.*

8.1. **The unframed input and its hypotheses.** We begin with the unframed support statement used in Section 6.

Proposition 8.2. *For any fixed Euler slope, every individual perverse direct-image coefficient of the split semistable unframed moduli stack has strict supports among the spectral-addition images Σ_ν. The rank-r primitive is the positive coarse intersection complex, whose proper Hitchin direct image is independent of the Euler characteristic as an object of $D^b\,\mathrm{MHM}(\mathcal{A}_r^L)$.*

Proof. Fix an Euler slope μ and let

$$\mathfrak{N}_\mu = \{0\} \sqcup \coprod_{r>0,\,\chi/r=\mu} \mathfrak{M}_L^{\mathrm{ss}}(r,\chi).$$

We apply the smooth linear-stack integrality theorem [9, Theorem 10.2.7] to this stack and verify its hypotheses.

The objects are compactly supported pure sheaves on $S = \mathrm{Tot}_C(L)$, or equivalently L-Higgs bundles on C. Semistable objects of slope μ form an abelian category. For example, if the image of a morphism between two such objects were not saturated, saturation would strictly raise its slope, while the quotient inequality for the source gives slope at least μ. This contradicts semistability of the target. Kernels, images, cokernels, and extensions therefore remain in the same category. Direct sum and the scalar action are its coherent operations.

For two objects E, F in this category, the deformation complex is

$$R\Gamma\big(C, [\mathcal{H}om(E,F) \longrightarrow \mathcal{H}om(E,F) \otimes L]\big).$$

Riemann–Roch and surface duality give

$$\chi_S(E,F) = -\ell \operatorname{rk}(E) \operatorname{rk}(F), \qquad \operatorname{Ext}^2_S(E,F) = \operatorname{Hom}_S(F, E \otimes A)^\vee = 0. \tag{8.5}$$

The second equality uses $\deg A < 0$ and semistability at the same slope. Hence $\mathfrak{N}_\mu$ is smooth and its numerical Euler pairing is symmetric.

We also check the stronger pointwise orthogonality condition. At a closed polystable point write $E = \bigoplus_j S_j \otimes V_j$, where the S_j are pairwise nonisomorphic stable objects of ranks r_j. Then

$$\dim \operatorname{Ext}^1_S(S_i, S_j) = \delta_{ij} + \ell r_i r_j.$$

The stabilizer is $\prod_j \mathrm{GL}(V_j)$. The opposite summands $\operatorname{Hom}(V_i, V_j) \otimes \operatorname{Ext}^1(S_i, S_j)$ and $\operatorname{Hom}(V_j, V_i) \otimes \operatorname{Ext}^1(S_j, S_i)$ have equal multiplicity and admit an invariant symmetric cross-pairing, using the trace pairing on the two Hom spaces. For $i = j$, use trace on $\operatorname{End}(V_i)$ and any nondegenerate symmetric form on the multiplicity space. The gauge summands are orthogonal by the same trace pairing. Thus the stabilizer representation is pointwise almost orthogonal in the sense of [9].

The semistable spectral Quot construction supplies a good moduli space $p : \mathfrak{N}_\mu \to M_\mu$, with affine diagonal. At fixed (r, χ) the stack is of finite type: consecutive ordinary Harder–Narasimhan slopes of a semistable L-Higgs bundle differ by at most ℓ, which bounds the underlying bundles. Direct sum is quasi-compact. Indeed a decomposition over a fixed family is a tuple of complementary idempotents in its relative endomorphism space; the idempotent equations are closed, and at a fixed target rank there are only finitely many numerical decompositions. For each fixed finite character-support face, the graded-object stack is a finite product of these eigensheaf stacks. Its total-object morphism is therefore quasi-compact. This is the facewise finiteness condition of [9, §2.3.7]. The perfect cotangent complex has finitely many weights on each such face. The theorem is locally finite in the numerical components. We keep the union over all ranks when applying it, so that the stack remains closed under direct sum.

The torus Chern parameters can be chosen on the stack itself. For $p_0 \in C$, the line $\det(E|_{p_0})$ has scalar weight r; its rational Chern class divided by r restricts to the generator of $H^2(B\mathbf{G}_m, \mathbf{Q})$. On a graded stack the same construction for each eigensheaf gives a basis of torus parameters. Finite automorphisms of the face permute these lines.

By [44, Proposition 2.9], the rank-r stack is smooth irreducible of dimension B_r and its coarse space is irreducible of dimension $B_r + 1$, for every χ. On the dense smooth spectral locus, the stack is the relative Picard stack and has exactly the scalar stabilizer. Scalar rigidification is generically an isomorphism onto the coarse moduli space. The BPS identification in [9, Corollary 7.2.4 and Theorem 10.2.7] consequently identifies the primitive with the coarse intersection complex, also when (r, χ) is not coprime.

Let $\eta = (r, \chi)$ run through the nonzero numerical components with $\chi/r = \mu$, let M_η be their coarse moduli spaces, and let $\mathrm{Sym}_\oplus$ denote the symmetric-power operation formed using direct sum. With $\mathrm{IC}^{\mathrm{vir}}_Y = \mathrm{IC}_Y \otimes \vartheta^{-\dim Y}$, the theorem gives the cohomological induction isomorphism

$$\mathrm{Sym}_\oplus\left(\bigoplus_{\eta \neq 0} \mathrm{IC}^{\mathrm{vir}}_{M_\eta} \otimes H^*(B\mathbf{G}_m)_{\mathrm{vir}}\right) \xrightarrow{\sim} Rp_*\, \mathrm{IC}^{\mathrm{vir}}_{\mathfrak{N}_\mu}, \quad H^*(B\mathbf{G}_m)_{\mathrm{vir}} = \Lambda \otimes \bigoplus_{v \geq 0} \mathbf{T}^v[-2v]. \tag{8.6}$$

It is formed from the lowest-perverse primitive inclusion, the specified Chern parameters, and the smooth Hall induction maps. It supplies object-level support. The primitive normalized by scalar rigidification differs from the unnormalized first perverse coefficient by its half-Tate factor.

In the split physical threefold the unframed critical stack is $T^*[-1]\mathfrak{N}_\mu$. Negativity in (8.5) makes the ordinary transverse field zero; the derived statement follows from the physical model in

Section 4. Its half/dual-half triangle identifies the fixed orientation with the cotangent orientation. Numerical symmetry and [9, Lemma 10.2.9] make this orientation commutative. Dimensional reduction gives

$$\pi_*\Phi_{T^*[-1]\mathfrak{N}_\mu} = \omega_{\mathfrak{N}_\mu} \otimes \mathbb{L}^{B_r/2} = \mathrm{IC}^{\mathrm{vir}}_{\mathfrak{N}_\mu}.$$

Since $\mathfrak{N}_\mu$ is smooth, this is the zero-function critical chart. The critical and smooth induction maps agree by [9, §§6.2.3 and 8.1.7].

The coarse Hitchin maps are proper. Spectral addition is finite, as proved below, and coarse direct sum is finite as well: its composite with the Hitchin map is proper, while a polystable object has only finitely many allocations of its Jordan–Hölder multiplicities among a fixed list of summands. Properness and quasi-finiteness give finiteness. Finally [44, Theorem 0.4] identifies the positive coarse Hitchin direct image of IC with the full-support intermediate extensions of the exterior-power spectral variations, independently of χ, in $D^b\,\mathrm{MHM}$. Applying the proper Hitchin pushforward to (8.6) therefore gives precisely the supports Σ_ν. At each fixed perverse degree only finitely many degrees of the torus factors contribute. The single rank-r unframed coefficient belongs to $\mathscr{S}_r$. □

8.2. **Finite spectral addition and reduced spectral curves.**

Lemma 8.3. *For every rank partition $\nu = (r_1, \dots, r_s)$ of r, spectral addition $\prod_j \mathcal{A}^L_{r_j} \to \mathcal{A}^L_r$ is finite. Its image Σ_ν is irreducible and has a dense locus of reduced nodal curves with smooth irreducible components of the specified ranks. The number of cross-component nodes and the codimension of this image are both*

$$e_\nu = \ell \sum_{i<j} r_i r_j. \tag{8.7}$$

Proof. Give a coefficient of degree j weight j. Spectral addition is equivariant for these positive gradings. Its fiber over the origin is supported at the origin: over $\mathbf{C}(C)$ a product of monic polynomials equal to t^r has every factor a power of t. Let R be the positively graded source coordinate ring and $f_1, \dots, f_m$ the pullbacks of target coordinates. The quotient $R/(f_1, \dots, f_m)$ is finite-dimensional. Choose homogeneous lifts of a basis. Induction on degree, subtracting a linear combination of these lifts and then writing the remainder in $\sum_j f_j R$, proves that the lifts generate R as a module over $\mathbf{C}[f_1, \dots, f_m]$. This proves finiteness with the full scheme structure of the fiber retained.

For the chosen finite rank range and sufficiently positive L, the monic coefficient systems generate the first jets needed to choose smooth spectral curves, pairwise transverse intersections, and no triple intersections. These conditions are open and nonempty by the corresponding incidence conditions in the coefficient spaces. A smooth spectral curve is connected: its structure algebra has the filtration with graded terms $\mathcal{O}_C, L^{-1}, \dots, L^{-(r_j-1)}$, so its global functions are $\mathbf{C}$. It is therefore irreducible. Distinct generic factors produce a reduced nodal union. Two factors of ranks r_i, r_j intersect in $\ell r_i r_j$ points. The source of addition is irreducible; its generic fiber only orders equal-rank factors. Using (8.3) gives $b_r - \sum_j b_{r_j} = e_\nu$ and proves the assertions. □

Let $U^{\mathrm{red}}_r \subset \mathcal{A}^L_r$ be the reduced spectral open.

Lemma 8.4. *Restriction to U^{red}_r is injective on $K_0(\mathscr{S}_r)$. Finite spectral addition of external intersection complexes, followed by rational finite-group projectors, has no strict supports other than the corresponding image supports.*

Proof. A finite morphism is small onto the image of any irreducible support. The pushforward of an intersection complex is the intermediate extension of its generically finite pushforward local system; no constituent is supported only on a smaller boundary. External products and finite projectors preserve this assertion, by taking simple perverse constituents. By Lemma 8.3, every support allowed in $\mathscr{S}_r$ meets U^{red}_r. A simple object restricts to a nonzero simple object and

is recovered by Hodge intermediate extension. Distinct simple objects therefore have distinct restrictions. Finite-length Grothendieck groups are free on their simple classes, which proves injectivity. □

Lemma 8.5. *On U_r^{red} the split stable-pair problem is the relative Hilbert scheme of length*

$$m = \chi + \gamma_r - 1. \tag{8.8}$$

This Hilbert scheme has smooth total space of dimension $b_r + m = B_r + \chi$. Its normalized coefficient in (8.1) *is $\mathbf{Q}[B_r] \otimes \mathbf{T}^{-B_r/2}$.*

Proof. For a reduced spectral curve D, a physical stable-pair sheaf is a rank-one torsion-free D-module F with a transverse A-valued endomorphism. If $\nu : \widetilde{D} \to D$ is the normalization, then $\mathcal{E}nd_D(F) \hookrightarrow \nu_*\mathcal{O}_{\widetilde{D}}$. Every normalization component is finite of positive degree over C. Since $\deg A < 0$, one obtains $\mathrm{Hom}_D(F, F \otimes A) = 0$. This also holds over an arbitrary ordinary base T. Let G be the spectral sheaf on $S_T = \mathrm{Tot}_{C\times T}(L)$, with locally free curve pushdown E. Write $p_T : S_T \to C \times T$ and $\pi_T : C \times T \to T$. The monic resolution

$$[p_T^*(E \otimes L^{-1}) \xrightarrow{y-\Theta} p_T^*E] \longrightarrow G$$

computes derived Hom on the smooth relative surface. It identifies the proper relative complex with

$$R(\pi_T p_T)_* R\mathcal{H}om_{S_T}(G, G \otimes p_T^*A) \simeq R\pi_{T,*}[\mathcal{E}nd(E) \otimes A \xrightarrow{[\Theta,-]} \mathcal{E}nd(E) \otimes A \otimes L]. \tag{8.9}$$

The two terms on the right have degrees $0, 1$. Since they are vector bundles on the proper smooth relative curve, this is a perfect complex of Tor-amplitude in $[0, 2]$, compatible with derived base change. Its degree-zero cohomology on every geometric fiber is the Hom group just shown to vanish. The fiberwise Tor-amplitude criterion therefore improves the interval to $[1, 2]$. Pullback to any ordinary base preserves this interval, so the transverse field is uniquely zero in families. Surface duality for the Gorenstein spectral curve identifies the remaining stable pair with a length-m subscheme, giving (8.8).

The smoothness criterion follows from conductor transversality. Write $p : S \to C$ for the projection. Let $\mathfrak{c}_D$ be the conductor and $\mathcal{N}$ the normal line of D in S. Finite duality and adjunction give

$$\mathfrak{c}_D \otimes \mathcal{N} = \nu_*(\omega_{\widetilde{D}} \otimes \nu^* K_S^{-1}), \qquad H^1(D, \mathfrak{c}_D \otimes \mathcal{N})^\vee = H^0(\widetilde{D}, \nu^* K_S) = 0.$$

Here $K_S = p^*A$ is negative on each component. The monic equation gives $\mathcal{N} = p^*L^r|_D$ and $H^0(D, \mathcal{N}) = \bigoplus_{j=1}^r H^0(C, L^j)$. Thus the Hitchin tangent directions surject onto the conductor quotient. They are transverse to the equigeneric tangent space. The all-length criterion [51, Theorem 4.16(3)], valid also for reducible locally planar curves, proves smoothness of every relative Hilbert scheme on this open. At a nodal curve it gives independent node-smoothing parameters.

The dimension $B_r + \chi$ equals the virtual dimension of the quasismooth surface-pair model, so its restriction here is the smooth model. The split physical critical model is its oriented shifted cotangent. Dimensional reduction gives $\mathbf{Q}[B_r + \chi]\mathbf{T}^{-(B_r+\chi)/2}$; multiplication by Λ^χ leaves $\mathbf{Q}[B_r]\mathbf{T}^{-B_r/2}$, as asserted. □

8.3. **The Hodge calculation at a nodal stratum.** We give the local calculation retaining the Hodge variation and the finite permutation action. Let S_ν be a smooth dense nodal stratum in Σ_ν. Labeling its components, nodes, and branches gives a finite étale cover, and every construction retains the deck action. Let G be the dual graph, J its edge set, and V the weight-one variation of the smooth normalization family. For a graph H set

$$A_H = H^1(H, \mathbf{Q}) \oplus H_1(H, \mathbf{Q}) \otimes \mathbf{T}, \qquad B_H = H^0(H, \mathbf{Q}) \oplus H_0(H, \mathbf{Q}) \otimes \mathbf{T}.$$

The graph spaces include their permutation and branch-reversal actions. For a subset $I \subset J$, let E_I be its signed oriented-edge representation and P_I its unoriented permutation representation. Put

$$\epsilon_I = \det(E_I)^\vee \otimes \det(P_I). \tag{8.10}$$

When I is not deck-invariant, a sum over its orbit means the induced object from its stabilizer. Write $\wedge_{-y} W = \sum_j (-y)^j [\wedge^j W]$.

Lemma 8.6. *Let $H_{\mathrm{sm}} = R^1\pi_{\mathrm{sm},*}\mathbf{Q}$ be the weight-one variation of the smooth spectral family π_{sm}, and let $H_{\lim}$ be its limiting variation. Its associated weight graded is*

$$\mathrm{Gr}^W H_{\lim} = H^1(G,\mathbf{Q}) \oplus V \oplus H_1(G,\mathbf{Q}) \otimes \mathbf{T}. \tag{8.11}$$

For $I \subset J$ and $p = |I|$, view the product of monodromy logarithms N_I as a map into the $\mathbf{T}^p$-twisted target. Its image on $\wedge^j H_{\lim}$ is zero unless deleting I preserves the connected components of G. In the latter case its weight-graded image is

$$\det(E_I)^\vee \otimes \mathbf{T}^p \otimes \wedge^{j-p}(V \oplus A_{G\setminus I}). \tag{8.12}$$

Consequently, if IC_0 denotes intermediate extension normalized to agree with the unshifted variation on the smooth locus, then

$$\sum_j (-y)^j [\mathrm{IC}_0(\wedge^j H_{\mathrm{sm}})|_{S_\nu}] = \wedge_{-y} V\, K_G(y), \quad K_G(y) = \sum_{\substack{I\subset J \\ \pi_0(G\setminus I)=\pi_0(G)}} (y\mathbf{T})^{|I|}\epsilon_I \wedge_{-y} A_{G\setminus I}. \tag{8.13}$$

These are identities of Hodge classes with their deck actions.

Proof. Normalization and nearby cycles give (8.11) as a weight-graded statement. For an oriented edge e, write $a_e \in H^1(G)$ for its coedge. On the graded exterior algebra, N_e is wedge multiplication by a_e followed by contraction against the corresponding functional on $H_1(G)$, with its target Tate twist. The contractions vanish on $H^1(G)$ and on V. Hence

$$N_I = (-1)^{p(p-1)/2}(\bigwedge_{e\in I} a_e) \wedge \iota_{\wedge_{e\in I} a_e}.$$

Dependent coedges make this operator zero. The graph chain sequence identifies independence with $\pi_0(G \setminus I) = \pi_0(G)$. For independent coedges, extend them to a basis and use a dual basis of cycles. The image contains all selected coedges and no selected dual cycle; the remaining exterior factor is $V \oplus A_{G\setminus I}$. This description is canonical: changing a lift from the quotient by the selected coedges changes it by a term killed by wedging all those coedges. The selected wedge supplies $\det(E_I)^\vee$ and the target supplies $\mathbf{T}^p$. Strictness of monodromy for the weight and Hodge filtrations identifies the associated graded of its image with this image on the associated graded. This proves (8.12) without splitting the limiting mixed variation.

The normal-crossing Hodge hypercube has in cochain degree p the sum of these images tensored by $\det(P_I)$. This second determinant records the ordering of smoothing coordinates. It cancels the permutation sign in $\det(E_I)^\vee$, while retaining branch reversal. Setting $j = k + p$, the exterior and cochain signs give $(-y)^{k+p}(-1)^p = y^p(-y)^k$. This proves (8.13) on a product chart.

To justify the relative Hodge statement when smoothing lines are nontrivial, specialize first to the normal bundle of S_ν in an affine neighborhood in $\mathcal{A}_r^L$. The specialization functor and its zero-section comparison are in the Hodge category [58, §§2.30–2.31]. After labeling, pull back to the product of the frame torsors of the smoothing lines. The normal bundle is then a literal product with $\mathbf{A}^{|J|}$. The normal-crossing description [72, Theorems 4.3, 4.5 and Proposition 4.8] retains $F, W, N, \mathrm{can}, \mathrm{Var}$ and the coordinate-order compatibility. On an overlap $t'_e = u_e t_{\sigma(e)}$, logarithmic coordinate change acts by exponentials of N_e. Each logarithm lowers the untwisted weight by two, so its action on the weight graded is the identity apart from the recorded graph permutation and the transition of V. The same holds on every graded image by strictness. Thus

the displayed Hodge-graded terms descend and determine the Grothendieck class of the limiting mixed variation. Finite labeling descent is kept equivariant throughout. □

Lemma 8.7. *Let π_m be the relative length-m Hilbert scheme over S_ν and set $\mathcal{Z}_G(y) = \sum_{m\geq 0} y^m [R(\pi_m)_*\mathbf{Q}]$. Then*

$$\mathcal{Z}_G(y) = \wedge_{-y} V \sum_{I\subset J} (y\mathbf{T})^{|I|} \epsilon_I \frac{\wedge_{-y} A_{G\setminus I}}{\wedge_{-y} B_{G\setminus I}}. \tag{8.14}$$

Equivalently, if τ runs over partitions of the graph vertices obtained as connected components after deleting edges, and J_τ is the set of edges between its blocks, then

$$\mathcal{Z}_G(y) = \wedge_{-y} V \sum_{\tau} (y\mathbf{T})^{|J_\tau|} \epsilon_{J_\tau} \frac{K_{G\setminus J_\tau}(y)}{\wedge_{-y} B_{G\setminus J_\tau}}. \tag{8.15}$$

Proof. The punctual Hilbert scheme of length $m \geq 2$ at a node is a chain of $m-1$ projective lines. To see the components, write the completed node ring as $\mathbf{C}[[x_0, y_0]]/(x_0y_0)$. For $1 \leq i \leq m-1$ the ideals

$$(ux_0^{m-i} + vy_0^i,\ x_0^{m-i+1},\ y_0^{i+1}), \qquad [u:v] \in \mathbf{P}^1,$$

give one component. The endpoint monomial ideals identify successive components. Conversely the lowest pure powers of x_0, y_0 in a colength-m ideal leave at most one relation between the two branches, which gives exactly these ideals and their endpoints. Its H^0 is constant and its H^2 is $\mathbf{T}$ times the permutation representation on those lines. Branch reversal reverses the chain; its action on the top class of a fixed middle component is $+1$. For the branch character ϵ, with $\epsilon^2 = 1$, the resulting series is

$$P(y) = \frac{1}{1-y} + \frac{\mathbf{T}y^2}{1-y^2} + \frac{\mathbf{T}(1+\epsilon)y^3(1+y)}{(1-y^2)^2}.$$

Multiplication gives the identity

$$P(y)(1-y)(1-\epsilon y) = 1 - \epsilon y + \mathbf{T}y^2. \tag{8.16}$$

It remains valid in families: the components may be projective line bundles, but their relative H^0, H^2 and the intersection sections give the same variations. For an orbit of d nodes under a deck transformation, with accumulated branch sign ϵ, the corresponding elementary equality is

$$(1-\epsilon y^d)(1-\epsilon \mathbf{T}^d y^d) + \epsilon \mathbf{T}^d y^d = 1 - \epsilon y^d + \mathbf{T}^d y^{2d}. \tag{8.17}$$

This is a polynomial identity using only $\epsilon^2 = 1$.

Stratify the Hilbert scheme by the lengths at its nodes and the lengths on its smooth complement. Evaluate this stratification with compact-support pushforward; properness of the total Hilbert map then identifies it with the required direct image. Write $W = \mathbf{Q}^{V(G)}$ for the vertex permutation variation. The normalization Macdonald series is $\wedge_{-y} V / \wedge_{-y} (W \oplus W\mathbf{T})$. Removing the two preimages of a node multiplies this by $(1-y)(1-\epsilon y)$, and its punctual Hilbert factor is $P(y)$. Hence each node contributes the numerator in (8.16). Expand that numerator as $(1-\epsilon y)(1-\epsilon \mathbf{T} y) + \epsilon \mathbf{T} y$. For a graph H on the same vertex set, its chain sequence and the dual sequence give

$$0 \longrightarrow H_1(H) \longrightarrow E_H \longrightarrow W \longrightarrow H_0(H) \longrightarrow 0, \qquad [A_H] - [B_H] = [E_H \oplus E_H\mathbf{T}] - [W \oplus W\mathbf{T}].$$

The edge representation is orthogonally self-dual, so this identity includes the dual cochain sequence as written. Expanding over the selected node set I now yields (8.14). Formula (8.17) checks the expansion on every node orbit with every finite deck action. Characters detect only the finite graph representations here; the common factor $\wedge_{-y} V$ remains the Hodge variation. To group the terms, write uniquely $I = J_\tau \sqcup I'$, where $\tau = \pi_0(G \setminus I)$ and I' does not change the components of $G \setminus J_\tau$. Substitution of (8.13) yields (8.15). This argument includes loops, multiple edges, and disconnected normalizations. □

8.4. **Purity and reconstruction of the source.**

Proof of Theorem 8.1. On U_r^{red}, Lemma 8.5 gives a proper direct image of a smooth constant complex of weight zero. It is therefore pure. The weight pieces of the global perverse coefficients still belong to $\mathscr{S}_r$. Any piece with the wrong weight restricts to zero on U_r^{red}, hence is zero by Lemma 8.4. The same argument using simple intermediate extensions shows that these pure coefficients are ordinary Hodge modules. Thus $R^L_{r,\chi}$ is globally ordinary and pure of weight zero.

We now recover the perverse pieces on the reduced open. Before the normalization of Lemma 8.5, the length-m Hilbert direct image of $\mathbf{Q}[b_r+m]$ has perverse-degree-i piece pure of weight b_r+m+i. On a strict support of codimension e, its generic variation has weight and unshifted degree

$$\mathrm{wt} = m+i+e, \qquad \text{shift } [-m-e-i]. \tag{8.18}$$

Descend through the finite collection of supports by dimension. On a dense nodal stratum of a remaining support, the larger-support pieces have already been identified. Their restrictions are given by Lemma 8.6. Subtract them in (8.15). The remainder is the normalization Hilbert term of length $m-e$, with factor $\mathbf{T}^e\epsilon_J[-2e]$. Its cohomological degree $k=m-e+i$ has weight $k+2e=m+i+e$ and total shift $[-k-2e]=[-m-e-i]$, exactly as in (8.18). Distinct i have distinct weights. Equality of the full variation classes consequently gives equality in each perverse degree, including its sign. Polarizable pure variations are semisimple; therefore their classes determine their isomorphism types. Hodge intermediate extension recovers the corresponding perverse summands.

On a smooth connected spectral curve the Macdonald summand of exterior degree j is

$$\wedge^j H^1[-j-2v]\otimes\mathbf{T}^v, \qquad m=j+u+v, \quad u,v\geq 0.$$

The positive Hitchin support theorem identifies P_i^L with the intermediate extension of $\wedge^{\gamma_r+i}H^1[b_r]$. Setting $j=\gamma_r+i$ and $\chi=m-\gamma_r+1$ gives precisely the primitive

$$x^{i+1+u+v}P_i^L[-i-1-2v]\otimes\mathbf{T}^{v-B_r/2}.$$

For a spectral partition $\nu=(r_1,\dots,r_s)$, (8.7) gives

$$B_r-2e_\nu=\sum_j B_{r_j}, \qquad \gamma_r-1-e_\nu=\sum_j(\gamma_{r_j}-1). \tag{8.19}$$

Thus the cross-node normalization is exactly the product of the lower-rank normalizations. Exchanging equal-rank components of rank d reverses branches at ℓd^2 nodes. This sign is $+1$ since ℓ is even. The node-ordering determinants have already canceled in (8.10); the remaining action is the ordinary derived symmetric action, with its Koszul signs. The supportwise computation is therefore the right-hand side of (8.4) on U_r^{red}.

Both sides have perverse constituents in $\mathscr{S}_r$: on the left by Section 6, on the right by finite addition and Lemma 8.4. That lemma extends the equality globally. In fixed rank and Euler characteristic the perverse range and the solutions for u,v are finite, as are the rank partitions. All operations are consequently coefficientwise finite. The usual equivariant filtration of tensor powers of a short exact sequence, followed by the exact rational invariant projector, shows that these symmetric operations are well defined on the indicated Grothendieck groups. □

9. The scalar critical primitive

In this section $(r,\chi_*)=1$. Write $M_L=M_L(r,\chi_*)$ and $M_K=M_{K_C}(r,\chi_*)$. Let $h_K: M_K\to\mathcal{A}_r^{K_C}$ be the canonical Hitchin map. Its perverse coefficients are

$$P_i^K={}^p\mathcal{H}^i R(h_K)_*\mathbf{Q}_{M_K}[2N_r]. \tag{9.1}$$

They include every strict support of the canonical Hitchin map.

Proposition 9.1. *There is a closed immersion*

$$f_r : \mathcal{A}_r^{K_C} \times \mathbf{A}^1 \longrightarrow \mathcal{A}_r^L, \qquad (\theta, \lambda) \longmapsto b\theta + a\lambda \tag{9.2}$$

on characteristic coefficients. For the potential $g_{\mathcal{A}}$ in (3.13)*, one has*

$$\phi_{g_{\mathcal{A}}} P_i^L \cong (f_r)_*(P_i^K \boxtimes \mathbf{Q}_{\mathbf{A}^1}[1]) \otimes \mathbf{T}^{(B_r - 2N_r)/2}. \tag{9.3}$$

The right-hand side is ordinary; its auxiliary semisimple and nilpotent monodromies are respectively 1 *and* 0*.*

9.1. **The critical scheme on arbitrary ordinary bases.** The derived critical-stack identification [38, Theorem 5.6 and Proposition 5.7], restricted to the semistable locus as in [37, Proposition 3.6], uses the function

$$g(E, \Phi) = \frac{1}{2}\langle \alpha_q, \mathrm{Tr}(\Phi^2) \rangle = g_{\mathcal{A}} \circ h_L. \tag{9.4}$$

Its ordinary objects are commuting fields (E, θ, z) on $\mathrm{Tot}_C(K_C \oplus \mathcal{O}_C)$, with $\Phi = b\theta + az$. The sufficiently positive choice of L in the fixed rank range ensures the required stability restriction. We record why this gives the full critical scheme, including nonreduced test bases.

At a geometric point coprimality makes the local-threefold object stable. Its commuting endomorphism z is scalar by Schur's lemma. A subbundle is then invariant under both fields precisely when it is θ-invariant. Hence (E, θ) is a stable canonical Higgs bundle on each geometric fiber.

Let T be an arbitrary ordinary base for such a family and $\pi : C \times T \to T$. The relative perfect complex

$$\mathcal{C}_T = R\pi_*[\mathcal{E}nd(E) \xrightarrow{[\theta, -]} \mathcal{E}nd(E) \otimes K_C] \tag{9.5}$$

has the two sheaves in degrees $0, 1$. Identity and normalized trace split off $R\pi_*[\mathcal{O}_C \to K_C]$, with zero differential. On every geometric fiber, the trace-free summand $\mathcal{C}_T^0$ has $H^0 = 0$ by stability and $H^2 = 0$ by canonical Higgs Serre duality. Its only possible cohomology is therefore in degree one. Perfection and the fiberwise Tor-amplitude criterion give amplitude $[1, 1]$: locally it is one vector bundle, of rank $2h(r^2 - 1)$, in degree one. The formation of (9.5) commutes with derived ordinary base change, so the trace-free H^0 remains zero over every ordinary base. Consequently

$$z = \lambda \,\mathrm{id}_E, \qquad \lambda = \frac{1}{r}\,\mathrm{Tr}(z) \in \Gamma(T, \mathcal{O}_T). \tag{9.6}$$

Here $\pi_*\mathcal{O}_{C\times T} = \mathcal{O}_T$ is used to identify the scalar. The perfect-base-change and amplitude facts are the standard criteria recorded in [65, Tags 0A1G, 0G99 and 0656]. Conversely, a stable canonical Higgs family and a scalar λ give the required commuting family. These constructions are inverse on objects and morphisms and commute with base change.

The scalar gerbe is the same in both descriptions. Fix $p_0 \in C$ and choose integers u_0, v_0 with $u_0\chi_* + v_0 r = 1$. The line

$$\mathcal{D}_E = (\det R\pi_* E)^{u_0} \otimes (\det E|_{p_0})^{v_0}$$

has scalar weight one and neutralizes the stable $\mathbf{G}_m$-gerbe. Since the projection changes neither E nor this line, it neutralizes the same gerbe on both sides. The classical critical scheme on the smooth coarse space is therefore

$$Z := t_0\,\mathrm{Crit}(g : M_L \to \mathbf{A}^1) \cong M_K \times \mathbf{A}^1. \tag{9.7}$$

One can see the scheme structure directly on a local neutral chart $U \times B\mathbf{G}_m$. The critical cdga has degree-(-1) generators ξ_j with $d\xi_j = \partial_j g$, and a degree-(-2) scalar Lie generator with zero differential. Its degree-zero cohomology is $\mathcal{O}_U/(\partial_j g)$. The scalar Lie direction adds no degree-zero equation or coordinate. Thus (9.7) identifies the gradient ideal. The extra $\mathbf{A}^1$ in that equation parametrizes the commuting scalar λ.

9.2. **The normal Hessian and its determinant root.** The entire scheme Z lies in $g^{-1}(0)$. Indeed Higgs scaling has weight two on g, so $dg(\mathrm{Eu}) = 2g$ puts g in its critical ideal. The smooth critical scheme (9.7) has dimension $2N_r + 1$ in the smooth space M_L of dimension $B_r + 1$. Its normal bundle $\mathcal{N}$ has even rank

$$c = B_r - 2N_r = (\ell - 2h)r^2 - 2. \tag{9.8}$$

The kernel of the Hessian is the tangent space of the critical scheme. It follows that $\mathrm{Hess}(g) : \mathcal{N} \to \mathcal{N}^\vee$ is a nondegenerate quadratic form.

Choose a section of L with reduced divisor D and fix bases in the residual lines at its points. For a spectral sheaf $\widetilde{E}$, the monic resolution and $0 \to \mathcal{O}_C \to L \to L|_D \to 0$ give, functorially in families,

$$\begin{aligned} \det \mathrm{RHom}_S(\widetilde{E}, \widetilde{E}) &= \det R\Gamma \mathcal{E}nd(E) \otimes \left(\det R\Gamma(\mathcal{E}nd(E) \otimes L)\right)^{-1} \\ &= \bigotimes_{p \in D} \det(\mathcal{E}nd(E_p) \otimes L_p)^{-1}. \end{aligned} \tag{9.9}$$

The determinant of $V \otimes V^\vee$ is canonically trivial: the factors $(\det V)^r$ cancel with their duals. Equivalently, $\det \mathrm{Ad}(g) = 1$ is a polynomial identity on GL_r, valid over every base. With the fixed divisor section and residual bases, (9.9) gives a nowhere-vanishing section, functorial in E, of the smooth positive canonical bundle. The scalar Lie generator trivializes the relative determinant of its gerbe, so the section descends to $k_L \in \Gamma(M_L, K_{M_L})$. This is also the determinant construction of [37, Proposition 3.8].

Let ω_K be the canonical Higgs symplectic form. On Z use $k_Z = (\omega_K^{N_r}/N_r!) \wedge d\lambda$. The cotangent sequence for $Z \hookrightarrow M_L$ gives a trivialization

$$\delta = k_L|_Z / k_Z \quad \text{of } \det \mathcal{N}^\vee.$$

Consequently

$$\det \mathrm{Hess}(g) = u_Z \delta^{\otimes 2}, \qquad u_Z \in \Gamma(Z, \mathcal{O}_Z)^\times. \tag{9.10}$$

The units in (9.10) are constant. The coprime canonical Hitchin space is smooth of pure dimension $2N_r$; its Hitchin morphism is proper and flat of relative dimension N_r and has connected geometric generic Jacobian fiber [16, §2.2]. These facts also exclude vertical components: smooth irreducible components are open and closed, and their images under a proper flat morphism are closed and open in the connected base. Each would meet the generic fiber, whose connectedness permits only one. Hence M_K is integral. Its finite Stein algebra is generically the base algebra and lies in its function field; normality of the affine Hitchin base gives

$$(h_K)_* \mathcal{O}_{M_K} = \mathcal{O}_{\mathcal{A}_r^{K_C}}, \qquad \Gamma(M_K \times \mathbf{A}^1, \mathcal{O})^\times = \mathbf{C}^\times.$$

Choose a constant square root of u_Z. The section $\sqrt{u_Z}\,\delta$ trivializes the determinant square-root torsor of the normal quadratic bundle, compatibly with base change. The relative Morse lemma now gives the full coefficient isomorphism

$$\phi_g \mathbf{Q}_{M_L}[B_r + 1] \cong j_* \mathbf{Q}_Z[2N_r + 1] \otimes \mathbf{T}^{c/2}, \qquad j : Z \hookrightarrow M_L. \tag{9.11}$$

The even normal rank gives semisimple monodromy $+1$, and the quadratic model has zero nilpotent monodromy. Changing the constant square root changes this isomorphism by a global sign and does not change the Hodge object. The physical determinant maps supply compatibility with products.

9.3. **Pushforward and the scalar shift.** Choose a zero p of the nonzero section $b \in H^0(C, LK_C^{-1})$. Positivity guarantees its existence, and (3.7) gives $a(p) \neq 0$. Evaluation of the trace function $\mathrm{Tr}(b\theta + a\lambda)$ at p, divided by $ra(p)$, recovers λ. After subtracting the scalar $a\lambda$, the degree-j coefficient belongs to $b^j H^0(C, K_C^j) \subset H^0(C, L^j)$. These are closed linear conditions

after a polynomial coordinate change, and multiplication by b^j has a linear inverse on its image. This proves that (9.2) is a closed immersion over every ordinary base.

Proper vanishing-cycle exchange for h_L, followed by (9.11), gives (9.3) on taking perverse cohomology. External product with $\mathbf{Q}_{\mathbf{A}^1}[1]$ and the finite map f_r are perverse exact. This proves Proposition 9.1. In particular, the positive primitive in (8.4), after vanishing cycles and restriction to $\lambda = 0$, becomes

$$x^{i+1+u+v} P_i^K[-i-2v] \otimes \mathbf{T}^{v-N_r}. \tag{9.12}$$

The shift [1] of the scalar line is retained in this formula. The Tate exponent is $v - B_r/2 + (B_r - 2N_r)/2 = v - N_r$.

10. Scalar packets and the global physical identity

Write

$$R^K_{r,\chi} = R(h_{1,r,\chi})_*(\Phi_{P_{r,\chi}(X_1)} \otimes \Lambda^\chi), \tag{10.1}$$

where $h_{1,r,\chi}$ is the positive spectral map defined by $\Phi = b\theta + az$. The endpoint properness used here is established independently in Proposition 11.5. Corollary 7.5 gives

$$[R^K_{r,\chi}] - [\phi_{g_{\mathcal{A}}} R^L_{r,\chi}] \in K_0(\mathscr{I}_r). \tag{10.2}$$

Let $C_{r,\chi}$ denote the coefficients of

$$\begin{aligned} &1 + \sum_{r,\chi} R^r x^\chi [C_{r,\chi}] \\ &= \mathrm{Exp}_{\mathrm{add}} \Bigg(\sum_{r \geq 1} R^r \sum_i \sum_{u,v \geq 0} x^{i+1+u+v} \Big[(f_r)_* (P_i^K \boxtimes \mathbf{Q}_{\mathbf{A}^1}[1]) \\ &\qquad\qquad [-i-1-2v] \otimes \mathbf{T}^{v-N_r} \Big] \Bigg). \end{aligned} \tag{10.3}$$

This expression is defined in MMHM, with its symmetric projectors. We will prove $[R^K_{r,\chi}] = [C_{r,\chi}]$ by induction on rank.

10.1. **Detection of the comparison error.** Rescaled evaluation at p defines $\mathrm{ev}_{p,r} : \mathcal{A}^L_r \to \mathrm{Pol}_r$, the affine space of monic degree-r polynomials. Let T_r be the single-root locus for $\mathrm{ev}_{p,r}$ and let $U_r = \mathcal{A}^L_r \setminus T_r$. On a physical object, $\mathrm{ev}_{p,r}$ is the characteristic polynomial of z.

Lemma 10.1. *Every simple support in $\mathscr{I}_r$ meets U_r, and*

$$K_0(\mathscr{I}_r) \longrightarrow K_0(\mathrm{MMHM}(U_r)) \tag{10.4}$$

is injective.

Proof. The lift of a in the extension (3.7) gives $\alpha_q a = 0$. Consequently scalar translation preserves the function:

$$g(\Phi + at\,\mathrm{id}) - g(\Phi) = t\langle \alpha_q, a\,\mathrm{Tr}\,\Phi \rangle + \frac{rt^2}{2} \langle \alpha_q, a^2 \rangle = 0.$$

The corresponding physical action is $z \mapsto z + t$. It preserves invariant subbundles, slopes, sections, coherent flags, and the determinant orientations constructed from E. Thus it acts on the coefficient objects used to generate $\mathscr{I}_r$.

Each simple support of such an object is translation invariant. To justify this without choosing an equivariant structure on each simple, take the finite collection of supports of one dimension. The image of $\mathbf{G}_a \times Z$ for an irreducible member Z is irreducible and lies in their finite union. It contains Z at the identity and has no larger dimension, so it is contained in Z. Inverse translation gives equality.

Every generator of $\mathscr{I}_r$ is a finite-addition convolution of at least two positive-rank factors. For any points of their invariant simple supports, their evaluated spectra are finite nonempty sets. Translate the factors independently. Outside finitely many affine hyperplanes in the translation

parameters, those sets are disjoint. The image support therefore meets U_r. Finite pushforward of an external intersection complex has only its image support, by Lemma 8.4. Every simple in $\mathscr{I}_r$ thus restricts nontrivially to U_r. Intermediate extension recovers it, and distinct simples have distinct restrictions. The free bases of simple classes give (10.4). The exclusion of the single rank-r unframed factor from $\mathscr{I}_r$ is essential at precisely this step. □

10.2. **The physical product over coprime packets.**

Lemma 10.2 (Excision for separated scalar packets)**.** *Work on an ordinary coprime factorization chart for the characteristic polynomial of z, and fix an allocation of the Euler characteristic to its factors. The Bézout decomposition of a stable pair is induced by an equivalence of the formal physical enhancements on all connective Artin bases. Under this same equivalence the closed Calabi–Yau form is the sum of the forms of its packets, with the original trace normalization $2\,\mathrm{Ch}_2$. Consequently the ordinary decomposition map preserves the algebraic d-critical sections.*

Proof. Use the projective compactification Y and its boundary framing from Section 4. At a point of the factorization chart let $Z_1, \dots, Z_s \subset Y$ be the disjoint compact curve supports of its packets. The baseline on their complement is $\mathcal{O}_Y$. Put

$$U_0 = Y \setminus \bigcup_i Z_i, \qquad U_i = Y \setminus \bigcup_{j \neq i} Z_j.$$

The U_i cover Y, and $U_i \cap U_j = U_0$ for $i \neq j$. They contain the boundary divisor. These constructions also work relatively: on the coprime chart the resultants are units, so the packet supports remain disjoint over the entire ordinary base. In a connective Artin deformation the corresponding open derived subschemes are determined by these underlying open sets.

Keep the fixed determinant and the boundary framing with their determinant compatibility. More precisely, use the homotopy fiber

$$\mathfrak{D}(Y, D_\infty) = \mathrm{fib}_{(\mathcal{O}_Y, \mathcal{O}_{D_\infty}, \mathrm{id})} \left(\mathrm{Perf}(Y) \longrightarrow \mathrm{Pic}(Y) \underset{\mathrm{Pic}(D_\infty)}{\overset{h}{\times}} \mathrm{Perf}(D_\infty) \right), \tag{10.5}$$

where the map is determinant together with restriction. An object has a determinant trivialization τ and a boundary framing σ, with the specified compatibility $\det(\sigma) = \tau|_{D_\infty}$. Near $\mathcal{O}_{D_\infty}$, determinant identifies the formal perfect-object problem with the Picard problem. The projective bundle formulas give $R\Gamma(Y, \mathcal{O}_Y) \simeq R\Gamma(D_\infty, \mathcal{O}_{D_\infty})$. Thus the compatible fiber is the original boundary-framed formal enhancement. On U_0, near $\mathcal{O}_{U_0}$, derived Nakayama and Tor-amplitude zero likewise identify every perfect deformation with an invertible object, whose determinant is itself. Its compatible fixed-determinant fiber is contractible, including all higher automorphisms and the overlap trivialization.

Descent for Perf and Pic on the displayed open cover now splits the formal problem as a product. Explicitly, restrict I to U_i and glue it to $\mathcal{O}_Y$ on $Y \setminus Z_i$ using the canonical trivialization on U_0; call the result I_i. Conversely glue the restrictions of the I_i to the U_i. The two gluing constructions are inverse, because every overlap problem is the contractible one just identified. Their homotopy fibers on mapping complexes give the unit, counit, and higher compatibilities on every connective Artin base. The determinant trivializations, boundary framings, and their specified compatibility are retained in each homotopy fiber product. On ordinary families this construction is exactly the Bézout decomposition of the target and of its section; the framing source stays the specified $\mathcal{O}$. Thus it induces the particular algebraic decomposition map in the statement. The ordinary stable-pair condition is preserved on each packet, so the equivalence restricts to the physical formal enhancements.

For the closed-form comparison use the strict trace kernel $\mathscr{A}_I$ from Lemma 4.6. On U_0 the supertrace is the identity of $\mathcal{O}_{U_0}$, so $\mathscr{A}_I|_{U_0} = 0$. Excision therefore gives, compatibly with the

derivative of the preceding gluing functor,

$$R\Gamma(Y,\mathscr{A}_I)\simeq\bigoplus_i R\Gamma_{Z_i}(Y,\mathscr{A}_I)\simeq\bigoplus_i R\Gamma(Y,\mathscr{A}_{I_i}).$$

To compute traces, products are first taken in the full endomorphism complexes and only then restricted to these trace-kernel inputs. A cyclic word whose inputs belong to different supports has support in their empty intersection and is canonically zero. A word supported on Z_i is identified, by the restriction and localization units, with the same word for I_i. The duality counit used for the volume trace is natural for these open restrictions and supported excision. The identity respects the cyclic trace constructions, their mixed differentials, and their negative-cyclic homotopies. This finite direct sum over the support components commutes with Hodge completion.

Equivalently, the canonical identifications with $\mathcal{O}$ off the supports give relative perfect-object classes. The category of perfect objects supported on the disjoint union is the product of the categories supported on the Z_i. We use the natural spectrum-valued Chern character and its additivity homotopies [32, Theorem 6.5], together with localization of the mixed complex [33, Sections 1.5, 1.12 and 2.4]. Taking the relative fibers in the localization diagram, and then the finite direct sum over the disjoint supports, gives the homotopy

$$\mathrm{Ch}(I)-\mathrm{Ch}(\mathcal{O}_Y)=\sum_i\bigl(\mathrm{Ch}(I_i)-\mathrm{Ch}(\mathcal{O}_Y)\bigr)$$

in the relative negative-cyclic complex. The relative class is defined by its trivialization on the open complement, and the cyclic words above represent its supported trace. The constant object $\mathcal{O}_Y$ has no positive-degree form in the moduli directions. The fixed boundary and volume transgression of $2\,\mathrm{Ch}_2$ is linear. Hence the very gluing equivalence constructed above pulls back the closed form to the sum of the packet forms, with all its higher Hodge components.

The passage from a (-1)-shifted symplectic enhancement to its d-critical section is functorial for these equivalences; in critical charts it is the canonical section determined by the closed form, as in [35, §3.3, (3.10) and Theorem 3.23]. The description by differential forms and the critical-ideal quotient commutes with completion of the local chart rings. Here we complete the entire algebraic factorization chart at a closed point, allowing all packet and base parameters to vary in its Artin test objects. Differential forms on these completions mean the continuous completions of the algebraic de Rham and mixed complexes. The preceding formal equivalence therefore compares the completions of the two existing algebraic d-critical sections under the specified ordinary Bézout map. Their pullbacks agree on the completion of the whole parameter space. Faithful flatness of completion detects their equality in the coherent critical-section quotient $\mathcal{O}/I^2$, hence gives their algebraic equality on the chart. The algebraic determinant line maps used below can be checked by the same completion argument. The construction is natural in the full parameter base and under refinement of the coprime factors. □

Lemma 10.3. *Over a base carrying a coprime factorization of the characteristic polynomial of z, the physical stable-pair coefficient factors into its scalar-packet coefficients with one common framing object. With the normalization Λ^{χ}, this factorization is compatible with refinement and with the ordinary derived permutation action. It holds over all ordinary bases of that factorization chart.*

Proof. If $P(t)=\prod_j P_j(t)$ with pairwise invertible resultants, Bézout identities produce polynomial idempotents in z. They split $E=\bigoplus_j E_j$, both fields, and the section. The construction is functorial over the base, including nilpotent bases. It decomposes the target of the map $\mathcal{O}\to\bigoplus_j F_j$, preserving its single framed source.

Put $S_K = \mathrm{Tot}_C(K_C)$ and let G_j be the surface projection of F_j. Let T be the factorization base and let $\pi_S : S_K \times T \to T$, $\pi_X : X_1 \times T \to T$ be the projections. For two distinct packets set

$$C_{ij} = R\pi_{S,*}R\mathcal{H}om_{S_K\times T}(G_i, G_j), \qquad A_{ij} = z_j(-) - (-)z_i : C_{ij} \to C_{ij}.$$

The Bézout identity contracts the sheaf-level monic resolution, so A_{ij} is an equivalence of the entire perfect complex. Thus

$$D_{ij} := \mathrm{fib}(A_{ij}) = R\pi_{X,*}R\mathcal{H}om_{X_1\times T}(F_i, F_j) \tag{10.6}$$

is canonically contractible in every degree.

We describe the orientation map explicitly. For the surface framing half use the determinant of

$$\mathrm{fib}\big(R\pi_{S,*}R\mathcal{H}om_{S_K\times T}(G, G) \longrightarrow R\pi_{S,*}G\big).$$

Its filtration under $G = G_1 \oplus G_2$ has the two framed diagonal terms and the off-diagonal terms C_{12}, C_{21}. The filtration retains the section differential coupling these terms. Surface duality identifies $C_{21} = C_{12}^\vee[-2]$, and the physical paired triangle of Section 4 identifies $D_{21} = D_{12}^\vee[-3]$ with the same trace form. Over the coprime factorization chart, write $\mathfrak{P}$ and $\mathfrak{P}_j$ for the formal physical enhancements of the total object and its jth packet. Let $\Psi : \mathfrak{P}_1 \times \mathfrak{P}_2 \to \mathfrak{P}$ be the gluing equivalence inverse to Bézout decomposition. The product denotes the coprime open in the absolute packet-moduli product, completed in all deformation and parameter directions. The relative Hom complexes above are evaluated on the test base T. Let M, M_j be their surface-half determinant lines and o, o_j their orientation square maps. The graded determinant double-triangle identity gives

$$\Psi^* M \simeq (M_1 \boxtimes M_2) \otimes \det D_{12}.$$

The canonical acyclic trivialization $\tau_{12} : \det D_{12} \to \mathcal{O}$ is represented by the graded determinant $\det A_{12}$ under $\det D_{12} = \det C_{12} \otimes (\det C_{12})^{-1}$. The root comparison in the direction of Proposition 4.4 is therefore

$$\gamma_{12} = (\mathrm{id} \otimes \tau_{12})^{-1} : M_1 \boxtimes M_2 \longrightarrow \Psi^* M.$$

The cross summand of the cotangent complex is

$$D_{12}^\vee[-1] \oplus D_{21}^\vee[-1] \simeq D_{12}^\vee[-1] \oplus D_{12}[2].$$

Each summand has determinant $\det D_{12}$, and the derivative contracts the cross determinant by $\tau_{12}^{\otimes 2}$. In a local free presentation this is the graded comparison for the cross Hessian

$$\begin{pmatrix} 0 & A_{12}^{\mathrm{t}} \\ A_{12} & 0 \end{pmatrix}.$$

The cohomological shifts are those inherited from C_{12}. Consequently

$$o_1 \boxtimes o_2 = \det(d\Psi^*) \circ \Psi^* o \circ \gamma_{12}^{\otimes 2}, \qquad d\Psi^* : \Psi^* \mathbb{L}_{\mathfrak{P}} \longrightarrow \mathbb{L}_{\mathfrak{P}_1 \times \mathfrak{P}_2}. \tag{10.7}$$

Lemma 10.2 identifies the original algebraic d-critical sections for this same gluing map. The paired determinant comparison above is its derivative comparison, with the algebraic root γ_{12} and square (10.7). Consequently it identifies the original oriented vanishing-cycle coefficients.

Refining into more packets uses direct sums of the same cross complexes. Exact-filtration coherence of graded determinants and their canonical acyclic trivializations give all regrouping maps. The triangular changes which separate the common framing have determinant one on the chosen half and preserve its trace pairing. This proves product coherence over the whole factorization chart.

Finally, Riemann–Roch gives

$$\chi(C_{ij}) = -2hr_i r_j, \qquad \kappa_j := \chi(\mathrm{fib}(\mathrm{RHom}_{S_K}(G_j, G_j) \to R\Gamma G_j)) = -2hr_j^2 - \chi_j \equiv \chi_j \pmod 2.$$

The opposite cross contractions have even parity. The raw oriented DT interchange has the extra character $(-1)^{\kappa_i\kappa_j}$. For raw degrees p, q, after the normalization its exponent is $pq + \kappa_i\kappa_j + \chi_i q + \chi_j p$.

The ordinary derived exchange of the shifted degrees has exponent $(p+\chi_i)(q+\chi_j)$. These agree modulo two. The fixed positive-quotient orientation comparison is the same paired determinant comparison in Section 4; its opposite residue blocks have even length and contribute no additional permutation character. Hence the normalized packet factorization has exactly the derived symmetric action. □

10.3. **Stratified descent and induction.**

Theorem 10.4. *For every rank and Euler characteristic,*

$$[R^K_{r,\chi}] = [C_{r,\chi}]. \tag{10.8}$$

Thus (10.3) *is the full monodromic Hodge class identity for the original physical stable-pair coefficients on* X_1.

Proof. First apply vanishing cycles to (8.4). Spectral addition is finite and the trace potential is additive, so proper exchange and monodromic Thom–Sebastiani apply. The coefficients are equivariant for Higgs scaling, under which $g_{\mathcal{A}}$ has weight two. After the Kummer cover they are locally acyclic in the function direction away from zero. Hence the sum-function vanishing cycles on the external product are supported where every factor value is zero. Proposition 9.1 now identifies the result with (10.3). In particular, (10.2) becomes

$$[R^K_{r,\chi}] - [C_{r,\chi}] \in K_0(\mathscr{I}_r). \tag{10.9}$$

Suppose (10.8) is known in all smaller positive ranks. Stratify U_r by the multiplicities of its distinct evaluated roots. For a multiplicity partition $\nu = (r_1, \dots, r_s)$, $s \geq 2$, denote the stratum by $\mathscr{R}_\nu$ and its ordered-root cover by $p_\nu : \widetilde{\mathscr{R}}_\nu \to \mathscr{R}_\nu$. This is a finite étale torsor under the group G_ν permuting roots of equal multiplicity. These covers are used separately on the root-multiplicity strata.

On T_d the unique root is the regular trace coordinate λ. Let

$$W_\nu = \{(x_1, \dots, x_s) \in \prod_j T_{r_j} : \lambda_i \neq \lambda_j \ (i \neq j)\}.$$

Spectral addition with the ordered roots defines a finite map $v_\nu : W_\nu \to \widetilde{\mathscr{R}}_\nu$. It is the closed coefficient-matching locus in a base change of the finite addition map; its image can be proper. Proper base change and Lemma 10.3 give the G_ν-equivariant expression

$$p_\nu^*(R^K_{r,\chi}|_{\mathscr{R}_\nu}) \cong (v_\nu)_* \left(\bigoplus_{\sum_j \chi_j = \chi} \boxtimes_j (R^K_{r_j,\chi_j}|_{T_{r_j}})\Big|_{W_\nu} \right). \tag{10.10}$$

Every Euler-characteristic allocation is retained, including summands permuted by G_ν.

The coefficients $C_{r,\chi}$ have the same expression. Each of their primitive factors lies on one scalar by (9.3); grouping its symmetric powers according to distinct scalar roots gives (10.10) with C in place of R^K. This uses the induction identity for finite symmetric projectors and keeps the within-root stabilizers. Equalities within a repeated-root packet are used on reductions, which determine the Hodge modules.

The induction hypothesis is an equality of full Grothendieck classes, and this is sufficient here. The tensor power of a short exact sequence has a finite equivariant filtration whose graded terms are induced tensor products of the subobject and quotient. Degree shifts carry their sign representations. Thus tensor and symmetric powers preserve the lower-rank equalities with the full G_ν action. Rational finite invariants are exact, so these equalities descend through p_ν. Finite closed/open localization combines the root strata and gives equality on all of U_r. By (10.9) and Lemma 10.1, the equality holds globally. The initial rank has $\mathscr{I}_1 = 0$, which starts the induction.

All expressions are finite in each (r,χ). Indeed stable pairs and the coefficients $C_{r,\chi}$ vanish for $\chi < a_r$; for a rank partition,

$$\sum_j a_{r_j} - a_r = \ell \sum_{i<j} r_i r_j \geq 0.$$

Hence every χ_j in (10.10) is bounded below and, with fixed sum, bounded above. There are finitely many partitions and perverse indices, completing the coefficientwise argument. □

10.4. **Restriction to the original nil coefficient.** Let $J_r = \mathrm{ev}_{p,r}^{-1}(t^r) \subset \mathcal{A}_r^L$ and let $N_{r,\chi} \subset P_{r,\chi}(X_1)$ be its inverse image. It is the characteristic-nil locus of z. Its coefficient is the restriction of the original ambient coefficient,

$$K_{r,\chi} = \Phi_{P_{r,\chi}(X_1)}|_{N_{r,\chi}} \otimes \Lambda^{\chi}. \tag{10.11}$$

The map $h_{K,\mathrm{nil}}$ records the canonical field θ. On reductions, $h_1|_{N_{r,\chi}} = f_r|_{\lambda=0} \circ h_{K,\mathrm{nil}}$; simultaneous triangularization of commuting θ, z proves this equality on geometric points, hence on the reduced schemes. Hodge modules and their six operations are unchanged by nilpotent thickenings. Properness in Proposition 11.5 gives the required proper base change.

Corollary 10.5. *Put $R^{\mathrm{nil}}_{r,\chi} = R(h_{K,\mathrm{nil}})_* K_{r,\chi}$. In the coefficientwise completed Grothendieck group of monodromic Hodge modules on the canonical Hitchin bases,*

$$1 + \sum_{r,\chi} R^r x^{\chi} [R^{\mathrm{nil}}_{r,\chi}] = \mathrm{Exp}_{\mathrm{add}} \left(\sum_{r\geq 1} R^r \sum_i \sum_{u,v\geq 0} x^{i+1+u+v} [P_i^K[-i-2v] \otimes \mathbf{T}^{v-N_r}] \right). \tag{10.12}$$

After compact integration, if

$$Z_{\mathrm{nil}}(R,x) = 1 + \sum_{r,\chi} R^r x^{\chi} [R\Gamma_c(N_{r,\chi}, K_{r,\chi})],$$

then

$$[R^r] \operatorname{Log} Z_{\mathrm{nil}}(R,x) = \frac{x\mathbf{T}^{-N_r}}{(1-x)(1-x\mathbf{T})} \sum_i (-x)^i [R\Gamma_c(\mathcal{A}_r^{K_C}, P_i^K)]. \tag{10.13}$$

Once the ordinary nil cohomology theorem 11.2 is proved, this last equality is also an equality in the completed $K_0(\mathrm{MHS})$.

Proof. Under finite addition, the reduced inverse image of J_r is the product of the lower J_{r_j}. Over a field, each monic factor of t^r is a power of t; this suffices for the assertion about reductions. Restrict (10.3) by proper base change, use (9.12), and identify its closed image with the canonical Hitchin base. This gives (10.12) with the ambient restriction (10.11).

Compact integration commutes with these external products, finite proper addition, and finite projectors. The plethystic logarithm extracts the primitive. Summing $u, v \geq 0$ gives $x/((1-x)(1-x\mathbf{T}))$ and the shift $[-i-2v]$ gives $(-1)^i$, proving (10.13) in $K_0(\mathrm{MMHS})$. The scalar primitives are ordinary by Proposition 9.1; the nil cohomology is ordinary by Theorem 11.2. Let e be the inclusion of 1 in the auxiliary affine line, and let $\overline{J}$ be the representative functor from the monodromic quotient to $\ker R\pi_*$ constructed by the projector in Section 5. If $i : \{0\} \hookrightarrow \mathbf{A}^1$ and $j : \mathbf{G}_m \hookrightarrow \mathbf{A}^1$, the ordinary object i_*F has representative $F \boxtimes j_! \mathbf{Q}_{\mathbf{G}_m}[1]$ in this model. Hence $e^* \overline{J}[-1]$ sends it back to F. The monodromic condition gives local constancy on $\mathbf{G}_m$, so the shifted fiber functor is exact, as in [15, §2.1]. For these two ordinary endpoints it is therefore a left inverse to the ordinary embedding. It proves injectivity of the induced map on their Grothendieck groups and returns the equality to $K_0(\mathrm{MHS})$. The multiplicative operations take place in MMHM, and the exact fiber functor recovers the two ordinary endpoints. □

11. The fixed loci and their cohomology

We now compute the second endpoint of the comparison. Throughout this section $X = X_1 = \mathrm{Tot}_C(K_C \oplus \mathcal{O}_C)$, with fiber coordinates (θ, z) and the volume-preserving action

$$H : \ (\theta, z) \longmapsto (u\theta, u^{-1}z). \tag{11.1}$$

Retain the characteristic-nil locus $N_{r,\chi}$ and its ambient coefficient $K_{r,\chi}$ from (10.11).

For a partition λ, denote by Z_λ the monomial thickening of the zero section with basis $z^i\theta^j$ indexed by the boxes (i, j) of λ, using zero-based row and column indices. Put

$$\omega_\lambda = \mathrm{Hom}_C(\mathcal{O}_{Z_\lambda}, K_C), \qquad D_{\lambda,d} = \mathrm{Quot}^H_{Z_\lambda}(\omega_\lambda, d). \tag{11.2}$$

Here ω_λ is the dualizing module of Z_λ. The integer a is the H-localization index relative to the baseline pair $\mathcal{O}_X \to \mathcal{O}_{Z_\lambda}$. It is locally constant on $D_{\lambda,d}$. Write $D_{\lambda,d,a}$ for its pieces, $f_{d,a} : D_{\lambda,d,a} \to \mathrm{Sym}^d C$ for the support map, and

$$L_{\lambda,d} = \bigoplus_a Rf_{d,a,*}(\Phi_{d,a} \otimes \Lambda^a[d]). \tag{11.3}$$

Here the enhancement and $\Phi_{d,a}$ are inherited from the fixed problem for physical stable pairs. The identification with a Quot scheme in (11.2) concerns its truncation.

For integers $r_b \geq 0$, indexed by $b \in \lambda$, set

$$d(\mathbf{r}) = \sum_b h_b r_b, \quad R(\mathbf{r}) = \sum_b r_b, \quad A(\mathbf{r}) = \sum_b (a_b + 1) r_b, \quad L(\mathbf{r}) = \sum_b l_b r_b, \tag{11.4}$$
$$e_b = l_b - a_b - 1, \qquad E(\mathbf{r}) = \sum_b e_b r_b, \qquad C_{\mathbf{r}} = \prod_{b \in \lambda} \mathrm{Sym}^{r_b} C.$$

The map

$$\mu_{\mathbf{r}} : C_{\mathbf{r}} \longrightarrow \mathrm{Sym}^{d(\mathbf{r})} C, \qquad (D_b)_b \longmapsto \sum_b h_b D_b \tag{11.5}$$

is finite: it is proper and has only finitely many allocations of each support multiplicity in a fiber.

Theorem 11.1. *The complex $L_{\lambda,d}$ is pure of intrinsic weight d. For each k there is a noncanonical isomorphism of pure perverse monodromic Hodge modules*

$${}^p\mathcal{H}^k L_{\lambda,d} \simeq \bigoplus_{\substack{d(\mathbf{r})=d \\ R(\mathbf{r})-2A(\mathbf{r})=k}} (\mu_{\mathbf{r}})_* \mathbf{Q}_{C_{\mathbf{r}}}[R(\mathbf{r})] \otimes \vartheta^{E(\mathbf{r})}. \tag{11.6}$$

Equivalently, the following identity holds in the full Grothendieck group with its symmetric-power operations:

$$\sum_{d \geq 0} v^d [L_{\lambda,d}] = \mathrm{Sym}_{\mathrm{add}} \left(\sum_{b \in \lambda} v^{h_b} (\delta_{h_b})_* \mathbf{Q}_C \otimes \vartheta^{e_b}[2(a_b + 1)] \right), \tag{11.7}$$

where $\delta_j : C \to \mathrm{Sym}^j C$ sends p to jp.

Theorem 11.2. *Write $\|\lambda\|^2 = \sum_i \lambda_i^2$ and $d_\lambda = \chi + h\|\lambda\|^2$. For every j there is a noncanonical isomorphism of ordinary pure Hodge structures*

$$H^j_c(N_{r,\chi}, K_{r,\chi}) \simeq \bigoplus_{\substack{\lambda \vdash r \\ d_\lambda \geq 0}} \bigoplus_{d(\mathbf{r}) = d_\lambda} H^{j-2L(\mathbf{r})}(C_{\mathbf{r}}, \mathbf{Q})(-L(\mathbf{r})). \tag{11.8}$$

The weight of each summand is j. The auxiliary monodromy is the identity and its nilpotent part is zero.

11.1. **The fixed functor.** The support of a compact pure one-dimensional sheaf on X is finite over C: it is proper and affine over C, and a fiber cannot contain a positive-dimensional proper subscheme of an affine plane. Its pushdown is a vector bundle, since curve torsion would give a zero-dimensional subsheaf. The fiberwise vector-bundle criterion gives the same assertion for flat Artin families.

The section of a stable pair fixes its equivariant structure uniquely. At the generic point of C it is cyclic, also over an Artin base by Nakayama's lemma. Its annihilator is a point of $\mathrm{Hilb}^r(\mathbf{A}^2)^H$. Each difference-weight summand of $\mathbf{C}[\theta, z]$ is a rank-one module over $\mathbf{C}[\theta z]$. An H-homogeneous ideal of finite colength is supported at the origin, so each such summand is a power submodule in θz. This gives the monomial ideals indexed by partitions. Let t_1, t_2 be the geometric tangent characters of (z, θ). At a partition ideal the two tangent characters attached to a box are

$$t_1^{-l_b} t_2^{a_b+1}, \qquad t_1^{l_b+1} t_2^{-a_b}. \tag{11.9}$$

For $(t_1, t_2) = (u^{-1}, u)$ their weights are $h_b, -h_b$. Thus the fixed tangent is zero. A local Artinian ring with $\mathfrak{m}/\mathfrak{m}^2 = 0$ is the ground field, so this fixed Hilbert scheme is reduced. The generic annihilator of any Artin family is consequently one constant partition ideal.

Every element of this ideal annihilates the pushdown after passing to the generic point of C. Multiplication by a nonzero curve function is injective on a vector bundle over an Artin base. Hence it annihilates the whole sheaf. We obtain

$$0 \longrightarrow \mathcal{O}_{Z_\lambda} \longrightarrow F \longrightarrow Q \longrightarrow 0. \tag{11.10}$$

The quotient is finite and flat over the base, by the fiberwise injection criterion applied to the two curve bundles. Curve duality sends (11.10) to a length-d quotient of ω_λ. Conversely, the kernel of a finite flat quotient of ω_λ is a curve bundle and duality recovers (11.10). These operations commute with base change. They identify the fixed moduli functor with (11.2) on all ordinary Artin bases, with the physical enhancement retained as specified above.

Use $\mathbf{P}^1$ as a reference curve and restrict the support of the defect to its affine coordinate t. The second volume-preserving torus is

$$R: \ (t, \theta, z) \longmapsto (vt, v^{-1}\theta, z). \tag{11.11}$$

The dualizing module has a basis

$$M = \bigoplus_{(i,j)\in\lambda} \mathbf{C}[t] e_{ij}, \qquad \theta e_{ij} = e_{i,j-1}, \qquad z e_{ij} = e_{i-1,j},$$
$$\mathrm{wt}_{H,R}(t) = (0, -1), \qquad \mathrm{wt}_{H,R}(e_{ij}) = (j - i, -j - 1).$$

The last -1 is the dualizing K_C factor. On this affine reference problem let $T_3 = (\mathbf{G}_m)^3$ scale (z, θ, t) independently. The subgroup $H \times R$ is its volume-preserving subtorus, and $T_3/(H \times R)$ is the residual one-dimensional torus. A T_3-fixed kernel has the form $\bigoplus t^{m_{ij}} \mathbf{C}[t] e_{ij}$, where the nonnegative heights m_{ij} weakly increase to the right and down.

Lemma 11.3. *The common H, R-fixed points of the length-d Quot scheme are the finitely many such height arrays of total size d. Their physical fixed tangent complexes are acyclic, and their normalized coefficients are $\mathbf{Q}$.*

Proof. An invariant Quot tangent map $K \to M/K$ must have the form

$$t^{m_{ij}} e_{ij} \longmapsto \sum_k c_{ij,k} t^{m_{ij}-k} e_{i+k,j+k}, \qquad 0 \le m_{ij} - k < m_{i+k,j+k}. \tag{11.12}$$

For $k \le 0$ the displayed inequality contradicts monotonicity. For $k > 0$ apply θ. Every target term remains nonzero after this operation: its new box lies in λ and its height is at least m_{ij}. Thus the coefficient propagates to column $j - 1$. At column zero the source has zero θ-image whereas every target still survives; the coefficient is zero. Induction on j proves $\mathrm{Hom}(K, M/K)^{H,R} = 0$.

A common fixed quotient is supported at $d[0]$ and is killed by t^d, by Cayley–Hamilton. It lies in a projective finite-module Quot fiber. The residual torus has a limit there, which is monomial. A nonconstant residual orbit would have positive-dimensional closure through this limit, contradicting the vanishing of its fixed tangent. This proves the classification and reducedness.

Differentiating the Artin-family identification above identifies the invariant tangent with $\mathrm{Ext}^1(I,I)_0^{H,R}$. Stability kills the trace-free Ext^0, and proper-support Calabi–Yau duality kills the invariant Ext^2 and Ext^3. The derived fixed tangent is therefore acyclic. For $d=0$ the same assertion follows before taking curve cohomology: the trace-free local self-Ext complex of the partition ideal has only the nonzero hook characters in (11.9). Its normalized oriented coefficient is the unit. The orientation comparisons of Section 4 give the unit for every common fixed point. □

11.2. **Closed attracting filtrations and punctual purity.** We use the compact form of oriented hyperbolic localization. For a torus action μ, write $i^+ : Y^+ \to Y$ and $p^+ : Y^+ \to Y^0$. With the orientation from Section 4, it is

$$p^+_!(i^+)^*\Phi_Y \simeq \Phi_{Y^0} \otimes \Lambda^{I_\mu}. \tag{11.13}$$

Here I_μ is the index of Proposition 5.1, namely

$$(\mathrm{rk}^+ H^0T_Y - \mathrm{rk}^- H^0T_Y) - (\mathrm{rk}^+ H^{-1}T_Y - \mathrm{rk}^- H^{-1}T_Y).$$

This is the localization theorem for the inverse action in the convention $p_*\eta^!$ of [35]. On a smooth zero-potential chart the index is $d_+ - d_-$. Compact contraction contributes Λ^{2d_+}, the ambient normalization contributes $\Lambda^{-d_0-d_+-d_-}$, and the fixed normalization is Λ^{-d_0}. Their quotient is exactly the exponent in (11.13). The chart maps, orientation torsors, and their descent in MMHM were fixed in Section 4.

Lemma 11.4. *Let Y be a finite-type separated $\mathbf{G}_m$-scheme, let V be an affine representation of strictly positive weights, and suppose that $Y \to V$ is proper and equivariant. If Y has a linearized ample line bundle, it has a finite closed filtration whose differences are entire attracting schemes over open-and-closed unions of fixed components. This description holds on arbitrary ordinary bases.*

Proof. A power of the ample line gives a relative projective embedding. Choose homogeneous generators of the finite equivariant module defining it. Their free cover yields a closed equivariant embedding $Y \hookrightarrow \mathbf{P}(W) \times V$ for a finite-dimensional representation W. Order its distinct weights $w_1 < \cdots < w_s$ and intersect Y with $\mathbf{P}(\bigoplus_{w \geq w_i} W_w) \times V$. These intersections form a descending closed filtration. In the ith difference, the weight-w_i projective coordinates generate the unit ideal and the lower coordinates vanish. The extension of the orbit is explicitly

$$([v_w], v) \longmapsto ([t^{w-w_i}v_w], t \cdot v). \tag{11.14}$$

It lies in Y over every ordinary ring: each defining equation vanishes after inverting t, and $R[t] \to R[t,t^{-1}]$ is injective.

Conversely, take an equivariant extension with limit in the weight-w_i fixed block. Locally on its base a weight-w_i coordinate is nonzero at zero. The inverse image of the complement of its invariant projective chart is closed, invariant, and disjoint from zero. Any nonzero geometric orbit in it would specialize to zero, so it is empty. In this chart a coordinate of relative weight k is a polynomial $a(t)$ satisfying $a(st) = s^k a(t)$. Comparing coefficients over $R[t,s,s^{-1}]$ gives $a(t) = t^k a(1)$ for $k \geq 0$ and $a(t) = 0$ for $k < 0$. Thus the extension is exactly (11.14). Separatedness gives uniqueness. The intersections of the fixed locus with the distinct projective weight blocks are open and closed. This proves the assertion. □

For the Quot problem on $\mathbf{A}^1$, let $p(t)$ be the characteristic polynomial of the support. A quotient is killed by $p(t)$ and is therefore a quotient of the vector bundle $M/p(t)M$. Stability under t, θ, z and the characteristic-polynomial condition are closed incidence equations in a quotient Grassmannian. Consequently its support map to $\mathrm{Sym}^d \mathbf{A}^1$ is projective. The monomial basis supplies a linearization; the base weights are $1, \dots, d$. Lemma 11.4 and Lemma 11.3 apply. Each compact attracting layer has coefficient Λ^{I_R}, whose only cohomology has degree and intrinsic weight I_R. In the localization long exact sequence the boundary maps go between pure weights differing by one, and hence vanish. Strictness of weights gives

$$H_c^j(D_{\lambda,d,a}(\mathbf{A}^1), \Phi_{d,a}) \text{ is a sum of copies of } \vartheta^j. \tag{11.15}$$

The untwisted coefficient $\Phi_{d,a}$ is monodromically Verdier self-dual. Properness of the support map and the positive-base contraction adjunction [7, Lemma 6] identify ordinary cohomology with the cohomology of its punctual fiber at $d[0]$, with the restricted ambient coefficient. If D_{punct} denotes that fiber, these identifications are

$$R\Gamma(D_{\mathrm{punct}}, \Phi_{d,a}|_{D_{\mathrm{punct}}}) \simeq R\Gamma(D_{\lambda,d,a}(\mathbf{A}^1), \Phi_{d,a}) \simeq R\Gamma_c(D_{\lambda,d,a}(\mathbf{A}^1), \Phi_{d,a})^\vee.$$

Thus degree j on the punctual fiber is dual to compact degree $-j$. Formula (11.15) makes it a sum of ϑ^j.

11.3. **Indices and the diagonal hook identity.** We compute the index in arbitrary shape. Let L_c be the number of boxes of content $c = j - i$. With rows indexed from zero, put

$$\mathcal{B}_\lambda = \{\lambda_i - i - 1 : i \geq 0\}, \quad \sigma(c) = \begin{cases} 1 & c \in \mathcal{B}_\lambda, \\ -1 & c \notin \mathcal{B}_\lambda, \end{cases} \quad d_c = \sum_{j-i=c} m_{ij}. \tag{11.16}$$

Use geometric tangent characters (t_1, t_2, t_3) for (z, θ, t). Their restrictions to H, R are (u^{-1}, uv^{-1}, v), so a monomial $t_1^{d_1} t_2^{d_2} t_3^{d_3}$ has lexicographic weight $(d_2 - d_1, d_3 - d_2)$. Set

$$\kappa = t_1 t_2 t_3, \quad \mathsf{Z}_\lambda = \sum_{(i,j)\in\lambda} t_1^{-i} t_2^{-j},$$

$$\mathsf{Z}_m = \sum_{(i,j)\in\lambda} \sum_{\alpha=1}^{m_{ij}} t_1^{-i} t_2^{-j} t_3^{\alpha}.$$

The bar inverts characters. Put $P = (1-t_1)(1-t_2)(1-t_3)$. The three-variable Koszul resolution gives $\chi(A, B) = \overline{A}BP$ on finite-support differences. For the stable-pair class $I = \mathcal{O} - F$, expansion of $-\chi(I, I)_0$ gives

$$V(F) = F - \kappa\overline{F} - PF\overline{F}.$$

Indeed, the trace-free part removes $\chi(\mathcal{O}, \mathcal{O})$, while $\chi(\mathcal{O}, F) = F$ and Serre duality gives $\chi(F, \mathcal{O}) = -\kappa\overline{F}$. The baseline and modified characters are

$$F_0 = \frac{\mathsf{Z}_\lambda}{1 - t_3^{-1}}, \qquad F = F_0 + \mathsf{Z}_m.$$

Subtracting before evaluating characters gives

$$V(F) - V(F_0) = \mathsf{Z}_m - \kappa\overline{\mathsf{Z}_m} - P\big(F_0\overline{\mathsf{Z}_m} + \overline{F_0}\mathsf{Z}_m + \mathsf{Z}_m\overline{\mathsf{Z}_m}\big).$$

Use $(1 - t_3)F_0 = -t_3\mathsf{Z}_\lambda$ and $(1 - t_3)\overline{F_0} = \overline{\mathsf{Z}_\lambda}$ to obtain the finite supported character

$$\begin{aligned} v_m = \mathsf{Z}_m - \kappa\overline{\mathsf{Z}_m} \\ - (1-t_1)(1-t_2)\big(-t_3\mathsf{Z}_\lambda\overline{\mathsf{Z}_m} + \overline{\mathsf{Z}_\lambda}\mathsf{Z}_m + (1-t_3)\mathsf{Z}_m\overline{\mathsf{Z}_m}\big). \end{aligned} \tag{11.17}$$

Every term is a finite-support difference of Hom complexes. On $H \times R$, Calabi–Yau duality writes the virtual tangent as $\mathrm{Ext}^1 - (\mathrm{Ext}^1)^\vee$ after removing automorphisms. Let $\mathrm{sgn}_{\mathrm{lex}}(\omega)$ be the

sign of the first nonzero coordinate of ω, and zero at $(0,0)$. If c_ω is the coefficient of the weight $\omega=(\omega_1,\omega_2)$ in v_m, the changes from the baseline indices are therefore

$$I_{H,R}(m)-I_{H,R}(0)=\tfrac{1}{2}\sum_\omega c_\omega \operatorname{sgn}_{\mathrm{lex}}(\omega),$$
$$I_H(m)-I_H(0)=\tfrac{1}{2}\sum_\omega c_\omega \operatorname{sgn}(\omega_1).$$

The zero-H part supplies I_R, so $I_{H,R}=I_H+I_R$. In the calculation below all indices are measured relative to the baseline. Abbreviate $\operatorname{sgn}(D,e)=\operatorname{sgn}_{\mathrm{lex}}(D,e)$ and define

$$\begin{aligned}A(D,e)=\ &\operatorname{sgn}(D,e)-\operatorname{sgn}(D-1,e)\\ &-\operatorname{sgn}(D+1,e-1)+\operatorname{sgn}(D,e-1).\end{aligned}$$

It vanishes outside $D=-1,0,1$; at these three values it equals, respectively, $-1-\operatorname{sgn}(e-1)$, $\operatorname{sgn}(e)+\operatorname{sgn}(e-1)$, and $1-\operatorname{sgn}(e)$. For a legal height increase at (i,j), put $\alpha=m_{ij}+1$, $c=j-i$, and $q=t_1^{-i}t_2^{-j}t_3^{\alpha}$. Replacing Z_m by Z_m+q in (11.17) gives

$$\begin{aligned}\Delta v=\ &q-\kappa\overline{q}\\ &-(1-t_1)(1-t_2)\big[-t_3\mathsf{Z}_\lambda\overline{q}+\overline{\mathsf{Z}_\lambda}q+(1-t_3)(q\overline{\mathsf{Z}_m}+\mathsf{Z}_m\overline{q}+1)\big].\end{aligned}$$

Pair opposite weights in the half-sign evaluation. For the box (l,k) put $D=k-l-c$ and $e=\alpha+j-k$. Its baseline contributes $-A(D,e)$, and its occupied levels contribute

$$\sum_{\beta=1}^{m_{lk}}\big(-A(D,e-\beta)+A(D,e-\beta+1)\big)=A(D,e)-A(D,e-m_{lk}).$$

The term $q-\kappa\overline{q}$ contributes $\operatorname{sgn}(-c,\alpha+j)$, and the constant term $-P$ contributes 1. Adding these contributions gives

$$\Delta I=\operatorname{sgn}(-c,\alpha+j)+1-\sum_{(l,k)\in\lambda}A(k-l-c,\alpha+j-k-m_{lk}). \tag{11.18}$$

On diagonal c, columns below, equal to, and above j contribute $2,1,-2$. On diagonal $c+1$, only columns above j contribute, each 2; on diagonal $c-1$, only columns below j contribute, each -2. These statements include the endpoint signs $0,1$ and follow from the legal-increment inequalities for southeast and northwest heights. If $s=\min(i,j)$, the four counts are $s, L_c-1-s, L_{c+1}-s-\mathbf{1}_{c<0}, s+\mathbf{1}_{c>0}$. Thus the sum in (11.18) is

$$1+2(L_{c+1}-L_c)+2\mathbf{1}_{c=0}. \tag{11.19}$$

Counting row endpoints gives $L_c-L_{c+1}=\mathbf{1}_{c\in\mathcal{B}_\lambda}-\mathbf{1}_{c<0}$. Since $\alpha+j>0$, (11.18) reduces to $\Delta I=\sigma(c)$.

Restricting (11.17) to H kills its quadratic term in Z_m. A legal increment has difference $F_c(u)-F_c(u^{-1})$, with

$$F_c(u)=u^{-c}\big(1-(1-u^{-1})(1-u)\overline{\mathsf{Z}_\lambda}\big).$$

The kernel $2\operatorname{sgn}(D)-\operatorname{sgn}(D-1)-\operatorname{sgn}(D+1)$ is supported at $D=-1,1$ with values $-1,1$. Therefore

$$\begin{aligned}\Delta I_H&=\operatorname{sgn}(-c)+L_{c-1}-L_{c+1}=\tfrac{1}{2}\big(\sigma(c-1)+\sigma(c)\big),\\ I_H&=\tfrac{1}{2}\sum_c d_c\big(\sigma(c-1)+\sigma(c)\big),\quad I_R=\tfrac{1}{2}\sum_c d_c\big(\sigma(c)-\sigma(c-1)\big).\end{aligned} \tag{11.20}$$

The first index is constant on each graded-dimension component. Passing from compact affine cohomology to punctual cohomology changes I_R to $-I_R$. After the factor $\Lambda^a[d]$ in (11.3), the punctual exponent and degree are

$$J=I_H-I_R=\sum_c d_c\sigma(c-1),\qquad \text{degree}=J-d. \tag{11.21}$$

We use the diagonal-trace rim-hook identity [68, Theorem 0.2 and Section 6]. In the present notation it asserts that reverse plane partitions are enumerated, with their individual diagonal sums, by independently repeated rim hooks. The hook of a box of content c, arm a, and leg l occupies the diagonal interval $[c-l, c+a]$. That interval contains $l+1$ Maya members and a nonmembers; its right endpoint is a member and the preceding left endpoint is a nonmember. Hence

$$\sum_{v=c-l}^{c+a} \sigma(v-1) = l - a - 1 = e_b, \qquad e_b - h_b = -2(a_b+1). \tag{11.22}$$

Equations (11.21)–(11.22) and the already proved punctual purity give the punctual graded cohomology object as the symmetric object on the generators

$$v^{h_b}\vartheta^{e_b}[2(a_b+1)], \qquad b \in \lambda. \tag{11.23}$$

The intrinsic weight $k+d$ separates degree k before any global symmetric-power operation is taken.

11.4. **Moving support and finite symmetric actions.** The punctual calculation must descend with its Hodge structure and its finite permutation actions. Let x, y be two étale coordinates on an open set $U \subset C$. On a neighborhood V of the diagonal in $U_t \times U_p$, discard the other sheets of the two étale equivalence relations. Then

$$X = x(t) - x(p), \qquad Y = y(t) - y(p) = Xv(t,p),$$

where v is a unit. Over $B = U_p \times \mathbf{A}^1_s$ form

$$W = \operatorname{Spec}\big(\mathcal{O}(V)[s,\tau]/(X - s\tau)\big), \qquad \upsilon = \tau v(t,p). \tag{11.24}$$

It has two étale maps to $\mathbf{A}^1_\tau \times B$ and $\mathbf{A}^1_\upsilon \times B$. At $s=1$ they are the original coordinate maps; at $s=0$ their ratio is the linear change $\upsilon = (dy/dx)(p)\tau$. Cotangent lifting preserves the relative threefold volume and the transverse monomial module. For a fixed d, the support-cycle fiber d times the distinguished section is the constant punctual Quot scheme over B. This holds on nonreduced bases: formal étaleness and Cayley–Hamilton identify the quotients killed by the dth power of the section coordinate.

The supported fixed-determinant gluing construction in Proposition 4.8 applies to these two maps. On the completed parameter ring $R = \mathbf{C}[[p-p_0, s-s_0]]$ its reference branch is constant. The full-parameter calculation there gives equality of the two potentials modulo the square of the common critical ideal, for these same maps. Its algebraic four-Hom line map has square equal to the determinant of the same cotangent derivative. Faithful flatness of completion therefore identifies the existing algebraic oriented critical structures over B. Tensor the line map with id_{K_B} to retain the smooth-base factor.

Proper pushforward of the punctual comparison gives a map between pullbacks of one punctual cohomology complex over $U \times \mathbf{A}^1$. Pullback from U is fully faithful, by the affine-line cohomology adjunction. The two endpoint maps are equal. At $s=0$ this is the linear coordinate action. The physical orientation map gives an action over $\mathbf{G}_m$, and its induced maps on cohomology objects are morphisms of constant variations. They equal the identity at 1, and hence everywhere on the connected group. Thus all moving coordinate transitions are identity on the punctual Hodge variations.

On an ordered configuration space the finite defect decomposes canonically by support idempotents, also over nilpotent bases: $T = \bigoplus_i T_i$. Use a single baseline G_0 and the triangles $G_0 \to G_i \to T_i$. Put

$$M(G) = \det \operatorname{fib}\big(\mathrm{RHom}_{S_K}(G,G) \longrightarrow R\Gamma(S_K, G)\big).$$

Use the equivariant surface and transported orientation of Lemma 4.7. The following grid is H-equivariant before taking determinants. Its weight-zero part compares the fixed orientations;

its moving parts supply the index and permutation factors below. The supported Hom and section triangles give the line map

$$M(G)M(G_0)^{-1} \simeq \bigotimes_i (M(G_i)M(G_0)^{-1}). \tag{11.25}$$

The cross complexes $\mathrm{RHom}(T_i, T_j)$ vanish canonically for $i \neq j$; the volume trace splits by the same idempotents. Taking weight zero in this grid before determinants gives the oriented fixed-coefficient comparison, with the determinant functor's regrouping maps. The merge morphism is étale on the labeled branch, since the idempotents lift uniquely over infinitesimal extensions. When two reference curves have baselines of different parity, divide first by their baseline point coefficients. The quadratic-point calibration in [36, Lemma 3.4] and [35, Section 7.8] is a rational unit map of the associated orientation torsors. The compact Thom and product-thimble maps have scalar $+1$. This fixes the cancellation in (11.25). The ordered-configuration variations are consequently constant; their cohomology still retains the topology of the configuration space.

We check the symmetry used in their descent. In the J-model of monodromic Hodge modules, let $r(x) = x^2$, let $\iota(x) = -x$ be its deck involution, and put

$$\vartheta = e_- r_* \mathbf{Q}_{\mathbf{A}^1}[1], \qquad e_- = (1 - \iota)/2.$$

Its global cohomology vanishes. Its kth convolution power is the full-sign summand of $(\sum x_i^2)_* \mathbf{Q}_{\mathbf{A}^k}[k]$. A smooth quadratic fiber retracts equivariantly onto S^{k-1}; its vanishing cohomology has determinant character. A coordinate permutation has its sign on this cohomology, and the shifted [1] factors have the same Koszul sign. The product is $+1$. These objects lie in the perverse heart, where rational realization is faithful. This proves that ϑ is even and Λ is odd, as stipulated in (3.11).

For a relative defect packet, let m_0, m_+, m_- be the Euler ranks of the surface half in weights zero, positive, and negative. The four-Hom comparison gives

$$m_0 + m_+ + m_- = -d, \qquad a = m_- - m_+, \qquad \kappa = m_0 + 2m_+, \qquad \kappa + a = -d. \tag{11.26}$$

The graded orientation exchange contributes parity κ. After the normalization $\Lambda^a[d] = \vartheta^a[d-a]$, the extra exchange exponent for two packet colors is

$$\kappa_i\kappa_j + (d_i - a_i)(d_j - a_j) = 2d_id_j + 2a_ia_j.$$

For repeated packets the remaining exponent is

$$\kappa - a^2 + d^2 = m_0(m_0 + 1) + 2m_+ + 2m_0(m_+ + m_-) + 4m_+m_-.$$

Both are even for all integer Euler ranks. The permutation group relations themselves come from determinant and Thom–Sebastiani coherence. The odd finite-defect correction $(-1)^d$ is multiplicative and has ratio one under exchange and refinement. Thus the finite symmetric actions are the ordinary derived actions on (11.23). In particular both charge twists $F_d \mapsto F_d \otimes \vartheta^{-d}$ and $F_d \mapsto F_d \otimes \vartheta^d[-2d]$ are symmetric monoidal.

Thom–Sebastiani products are formed in MMHM. Choose local critical functions vanishing on their reduced critical loci and shrink to the zero-critical-value situation of [59, Theorem 5.3].

11.5. **Global purity and proof of the relative formula.** Write $A_a = Rf_{d,a,*}\Phi_{d,a}$ before twisting by its index. The global support map is proper, since the dualizing Quot scheme on the finite thickening of the projective curve is projective. At a support cycle $\sum d_ip_i$, proper base change and (11.25) identify its coefficient with the external product of the punctual coefficients, summed over the compatible local indices. The punctual result shows that every stalk $H^j(i_x^*A_a)$ is pure of intrinsic weight j. It follows that $\mathrm{wt}(A_a) \leq 0$. More explicitly, on a smooth stratum of dimension s, a lisse perverse coefficient is a variation shifted by $[s]$ and its perverse weight is its fiber weight plus s. Applying this to a stratification adapted to the weight filtration proves

the bound. The untwisted $\Phi_{d,a}$ is self-dual, and proper pushforward preserves duality. Thus the lower weight bound also holds:

$$\mathrm{wt}(A_a) = 0, \qquad \mathrm{wt}(L_{\lambda,d}) = d. \tag{11.27}$$

We have used self-duality before the twist by Λ^a.

Pure polarizable monodromic Hodge modules are semisimple. One can see this directly in the weight-stable Serre quotient defining MMHM: a pure quotient object of weight w has the pure ambient lift $\mathrm{Gr}_w^W M$. The latter is semisimple, and its image under an exact Serre quotient remains semisimple. This also splits the same-weight extensions used in the punctual purity argument.

Stratify $\mathrm{Sym}^d C$ by support multiplicities. On the ordered finite cover of a stratum, the preceding coordinate and product comparisons give the constant punctual objects with their computed stabilizer actions. The right-hand side of (11.7) has the same description: it assigns box occupancies at each support point, with total weighted multiplicity $\sum h_b r_b$. Taking finite rational invariants is exact. Equivariant tensor-power filtrations of short exact sequences show that passage to full K_0 classes preserves these projectors. For a finite group quotient of C^s, the invariant direct image of its constant object is the constant object of the symmetric product, including at diagonals: the unit followed by the projector is an isomorphism on the finite orbit fibers. Therefore finite closed/open localization on the base gives

$$[L_{\lambda,d}] = \left[\bigoplus_{d(\mathbf{r})=d} (\mu_{\mathbf{r}})_* \mathbf{Q}_{C_{\mathbf{r}}} \otimes \vartheta^{E(\mathbf{r})}[2A(\mathbf{r})] \right]. \tag{11.28}$$

Each box on the right has weight $e_b + 2(a_b+1) = h_b$. Hence both sides of (11.28) are pure of weight d. The component of weight $d+k$ in their classes is exactly $(-1)^k[{}^p\mathcal{H}^k]$. Semisimplicity identifies the corresponding perverse objects. A term on the right has perverse degree $R(\mathbf{r}) - 2A(\mathbf{r})$ because $\mu_{\mathbf{r}}$ is finite and $C_{\mathbf{r}}$ is smooth of dimension $R(\mathbf{r})$. This proves Theorem 11.1.

Twisting by the symmetric charge ϑ^{-d} changes the box generator to

$$v^{h_b}(\delta_{h_b})_* \mathbf{Q}_C(a_b+1)[2(a_b+1)]. \tag{11.29}$$

Thus the constituents of $\vartheta^{-d} L_{\lambda,d}$ lie in the ordinary Hodge-module image, with trivial auxiliary monodromy. This conclusion retains the $H^1(C)$ and all its exterior powers under proper integration.

11.6. **Properness of the positive spectral map.** The positive spectral map $P_{r,\chi}(X) \to \mathcal{A}_r^L$ records the characteristic polynomial of $az+b\theta$. We will use its properness both for the comparison endpoint and for a closed filtration of the whole nil locus.

Proposition 11.5. *The positive spectral maps on $P_{r,\chi}(X_0)$ and $P_{r,\chi}(X_1)$ are proper. The restriction of the latter to $N_{r,\chi,\mathrm{red}}$ is equivariant for* (11.1) *and the positive grading on $\mathcal{A}_r^L$, and admits an H-linearized ample line bundle.*

Proof. The closed immersions f_s of (9.2) and the finite spectral addition of Lemma 8.3 show that, for each rank partition, the map

$$\prod_i (\mathcal{A}_{r_i}^{K_C} \times \mathbf{A}^1) \longrightarrow \mathcal{A}_r^L \tag{11.30}$$

is finite.

On every compact irreducible curve in X, the function z is constant, and its projection to $\mathrm{Tot}(K_C)$ is a spectral divisor finite over C. Thus the finitely many maps (11.30) parameterize all compact effective cycles of rank r, with their multiplicities. Given a DVR family in $\mathcal{A}_r^L$ and a generic stable pair, after a finite extension its generic cycle lifts to one of these finite maps. Its lift extends over the integral-closure DVR and has a limiting cycle entirely in X. The projective stable-pair moduli on $Y = \mathbf{P}_C(\mathcal{O}_C \oplus K_C \oplus \mathcal{O}_C)$ extend the generic pair. Separatedness of its

projective cycle identifies the limit with the cycle just constructed. Purity of the limiting sheaf prevents an additional isolated support component at infinity. The extended pair is therefore supported in X. Boundary avoidance descends through the finite extension, proving properness.

For $X_0 = \mathrm{Tot}_C(L \oplus A)$, the pullback of A to the normalization of any compact irreducible curve has negative degree. Its transverse coordinate is consequently zero. Every compact cycle lies in $\mathrm{Tot}_C(L)$ and is a finite spectral divisor. The same DVR argument, now with the ordinary positive spectral coefficients as the cycle parameters, proves properness for this row as well.

At a geometric nil point, the commuting θ, z are simultaneously triangularizable over the curve's generic field after extension, and the diagonal entries of z vanish. Hence

$$\mathrm{char}(az + b\theta) = \mathrm{char}(b\theta), \quad \mathrm{char}(au^{-1}z + bu\theta) = \mathrm{char}(bu\theta). \tag{11.31}$$

These are identities of regular functions on the reduced nil locus, and give the asserted positive-base equivariance. The H action extends to Y. An H-linearized polarization induces linearizations on the projective section factor and the Plücker line in the stable-pair GIT construction [66]. It commutes with the framing linear group. Passing to the invariant graded ring and a divisible Veronese gives a linearized ample line on the projective stable-pair compactification. Its restriction gives the required line on the nil locus. □

11.7. **The whole nil locus.** Proposition 11.5 and Lemma 11.4 give a finite closed filtration of $N_{r,\chi,\mathrm{red}}$ by whole attracting schemes over unions of complete fixed components. These are the reductions of the ambient attractors. Indeed, on an arbitrary ordinary ring the jth characteristic coefficient of z in an attracting family has Laurent form $u^{-j}c_j$. Extension through zero forces $c_j = 0$, by injectivity of $R[u] \to R[u, u^{-1}]$. On a reduced test base the orbit extension is reduced and therefore factors through $N_{r,\chi,\mathrm{red}}$; the converse is immediate. The universal reduced attractors and their limit maps consequently agree. Hodge modules and their six operations are insensitive to nilpotent thickenings, so (11.13) computes the layers with exactly the coefficient (10.11).

The baseline Euler characteristic and H-index are

$$\chi_0(\lambda) = \chi\left(\bigoplus_{(i,j)\in\lambda} K_C^{-j}\right) = -h \sum_{(i,j)\in\lambda} (2j+1) = -h\|\lambda\|^2, \tag{11.32}$$

$$I_0(\lambda) = \sum_{b\in\lambda} \chi(K_C^{a_b+1}) = h\sum_{b\in\lambda}(2a_b+1) = h\|\lambda\|^2 = -\chi_0(\lambda). \tag{11.33}$$

The positive transverse hook line is $K_C^{a_b+1}$, explaining the second formula. For fixed λ, the defect is $d = \chi - \chi_0(\lambda)$. The corresponding complete attracting layer has compact cohomology

$$\begin{aligned} R\Gamma(D_{\lambda,d,a}, \Phi_{d,a} \otimes \Lambda^{\chi+I_0+a}) &= R\Gamma(D_{\lambda,d,a}, \Phi_{d,a} \otimes \Lambda^{d+a}), \\ \bigoplus_a R\Gamma(D_{\lambda,d,a}, \Phi_{d,a} \otimes \Lambda^{d+a}) &= R\Gamma(\mathrm{Sym}^d C, \vartheta^d[-2d]L_{\lambda,d}). \end{aligned} \tag{11.34}$$

The latter is pure of complex weight zero by (11.27). The symmetric charge changes each generator of (11.7) to

$$\vartheta^{h_b+e_b}[2(a_b+1) - 2h_b] = \mathbf{T}^{l_b}[-2l_b]. \tag{11.35}$$

Hence the perverse cohomology and the proper cohomology of the normalized layer are ordinary pure Hodge objects, with summands $(\mu_{\mathbf{r}})_*\mathbf{Q}_{C_{\mathbf{r}}}(-L(\mathbf{r}))[-2L(\mathbf{r})]$.

For fixed r, χ there are finitely many partitions, and each partition fixes d. Each remaining projective Quot scheme has finitely many connected components; its index and the weight of the linearized ample line are locally constant. Every filtration layer therefore selects an open-and-closed union of complete components. Its proper cohomology is a direct summand of (11.34), so has the same purity and ordinary coefficient properties. In the compact-cohomology sequences of successive closed filtration steps, the boundary maps join weights differing by one and vanish.

The remaining extensions of a fixed weight split in the pure polarizable heart. This gives (11.8), including its degree and auxiliary monodromy assertions, and proves Theorem 11.2.

Remark 11.6. The perverse degree in (11.6) belongs to the defect support map to $\mathrm{Sym}^d C$. The Hitchin perverse degree is extracted from the source primitive using x in Section 12. Theorem 11.2 supplies the ordinary Hodge cohomology and its purity in each degree.

12. Hodge recovery and the character-variety formula

We combine the two calculations before taking any numerical realization. All formal variables in this section are even. For a bounded Hodge complex K, its class is $[K] = \sum_j (-1)^j [H^j(K)]$. Let

$$\zeta_C(y) = \sum_{m\geq 0} [H^*(\mathrm{Sym}^m C)]\, y^m = \frac{\bigwedge_{-y} H^1(C)}{(1-y)(1-\mathbf{T}y)}. \tag{12.1}$$

The last identity is the symmetric-product formula for a curve [41]; it is an identity of Hodge classes, with the derived symmetric action on $H^1(C)$ giving the exterior numerator.

12.1. **The identity in full Hodge classes.** Define the series of the original restricted physical coefficients by

$$Z_{\mathrm{nil}}(R,x) = 1 + \sum_{r\geq 1} \sum_{\chi} R^r x^\chi [R\Gamma_c(N_{r,\chi}, K_{r,\chi})]. \tag{12.2}$$

Theorem 11.2 gives

$$Z_{\mathrm{nil}}(R,x) = \sum_{\lambda} R^{|\lambda|} x^{-h\|\lambda\|^2} \prod_{b\in\lambda} \zeta_C(x^{h_b}\mathbf{T}^{l_b}). \tag{12.3}$$

The degreewise ordinariness and purity proved above place these classes in $K_0(\mathrm{MHS})$.

Choose an integer d coprime to r with $\zeta = \exp(2\pi i d/r)$, and set $\chi_* = d - r(g-1)$. Let M_K be the smooth moduli space of stable K_C-Higgs bundles of rank r and degree d, equivalently of Euler characteristic χ_*. This is the degree matching the fixed twist in (3.1). The source formula permits this choice of its coprime parameter. Let $h_K : M_K \to \mathcal{A}_r^{K_C}$ be its Hitchin map, and put

$$P_i^K = {}^p\mathcal{H}^i R h_{K,*} \mathbf{Q}_{M_K}[2N_r], \qquad A_i = [R\Gamma_c(\mathcal{A}_r^{K_C}, P_i^K)]. \tag{12.4}$$

The scalar primitive of Section 9, restricted to the characteristic-nil locus, is

$$x^{i+1+u+v} P_i^K[-i-2v] \otimes \mathbf{T}^{v-N_r}, \qquad u,v \geq 0. \tag{12.5}$$

Proper base change and compact integration of the finite spectral addition maps preserve its symmetric-power operations. Taking the plethystic logarithm of the source identity, summing u, v, and using the shift in (12.5) give

$$[R^r]\,\mathrm{Log}\, Z_{\mathrm{nil}}(R,x) = \frac{x\mathbf{T}^{-N_r}}{(1-x)(1-x\mathbf{T})} \sum_i (-x)^i A_i. \tag{12.6}$$

Both sides are now ordinary. The exact left inverse to the ordinary embedding into the monodromic category gives an injection on their Grothendieck groups, so the source identity gives (12.6) in $K_0(\mathrm{MHS})$.

Equivalently, define

$$F_r(x) = x^{-1}\mathbf{T}^{N_r}(1-x)(1-x\mathbf{T})[R^r]\,\mathrm{Log}\, Z_{\mathrm{nil}}(R,x). \tag{12.7}$$

Then

$$F_r(x) = \sum_i (-x)^i [R\Gamma_c(\mathcal{A}_r^{K_C}, P_i^K)]. \tag{12.8}$$

The variable x records Euler characteristic; the source primitive identifies its exponent i with the Hitchin perverse coefficient in (12.8).

All these operations are coefficientwise finite. For the finite rank range chosen at the start, the positive-source coefficients and the stable-pair endpoints vanish below a_r. For $r = \sum r_i$,

$$\sum_i a_{r_i} - a_r = \ell \sum_{i<j} r_i r_j \geq 0. \tag{12.9}$$

At a fixed (r, χ), lower bounds on every χ_i consequently give upper bounds as well. Each primitive has a finite perverse range, and its fixed sum $u+v$ has finitely many choices. The geometric series in (12.1) have positive x-degrees after their baseline shift, and there are finitely many partitions in fixed rank. This defines the positive-rank and lower-bounded-characteristic completion in which (12.3)–(12.8) hold.

12.2. **Recovering perverse-graded Hodge structures.** The moduli space M_K is smooth semiprojective: properness of h_K and positive scaling of its base give a limit of every scaling orbit, and the fixed locus lies in the projective nil fiber. Thus $H^D(M_K)$ is pure of weight D by [28, Corollary 1.3.2]. Smooth duality gives the same purity for $H^D_c(M_K)$. This is independent of the stable-pair purity proved in Section 11. The proper decomposition theorem in Hodge modules makes $H^j_c(\mathcal{A}_r^{K_C}, P_i^K)$ a direct summand of $H^{j+2N_r+i}_c(M_K)$. It is therefore pure of weight $2N_r + i + j$. Distinct j have distinct weights, so (12.8) implies

$$[H^j_c(\mathcal{A}_r^{K_C}, P_i^K)] = (-1)^{i+j}\,\mathrm{Gr}^W_{2N_r+i+j}([x^i]F_r(x)). \tag{12.10}$$

The right side belongs to the Grothendieck group of pure polarizable Hodge structures of that weight. This category is semisimple, so its class determines its isomorphism class. Formula (12.10) recovers the pure Hodge structure $H^j_c(\mathcal{A}_r^{K_C}, P_i^K)$ for every i, j.

Normalize the perverse Leray filtration by $Rh_{K,*}\mathbf{Q}[N_r]$, so its index k corresponds to the perverse object $P^K_{k-N_r}$. The ordinary piece in degree D is

$$V_{D,k} = H^{D-N_r-k}(\mathcal{A}_r^{K_C}, P^K_{k-N_r}). \tag{12.11}$$

Duality gives $\mathbb{D}P_i^K = P^K_{-i}(2N_r)$, whence

$$V_{D,k} \simeq H_c^{N_r+k-D}(\mathcal{A}_r^{K_C}, P^K_{N_r-k})^\vee \otimes \mathbf{T}^{2N_r}. \tag{12.12}$$

In particular the compact indices and weight corresponding to (D, k) are

$$i_c = N_r - k, \qquad j_c = N_r + k - D, \qquad 2N_r + i_c + j_c = 4N_r - D. \tag{12.13}$$

Write $\beta_{D,k} = \dim V_{D,k}$. Equations (12.3), (12.10), and (12.12) determine these dimensions as well as the underlying pure Hodge structures.

12.3. **The hook expression.** Apply the weight character only at this point. Write s for the character of ϑ, so $\mathbf{T}$ has character s^2. Since $H^1(C)$ has weight one and dimension $2g$,

$$\zeta_C(y; s) = \frac{(1-sy)^{2g}}{(1-y)(1-s^2y)}. \tag{12.14}$$

Write $Z_{\mathrm{nil}}(R, x; s)$ for the weight character of $Z_{\mathrm{nil}}(R, x)$. Formula (12.3) gives

$$Z_{\mathrm{nil}}(R, x; s) = \sum_\lambda R^{|\lambda|} x^{-h\|\lambda\|^2} \prod_{b\in\lambda} \zeta_C(x^{h_b} s^{2l_b}; s). \tag{12.15}$$

Let $E_c(P_i^K; s)$ denote the alternating weight character of $R\Gamma_c(\mathcal{A}_r^{K_C}, P_i^K)$. Then (12.6) says

$$[R^r]\,\mathrm{Log}\, Z_{\mathrm{nil}}(R, x; s) = \frac{xs^{-2N_r}}{(1-x)(1-xs^2)} \sum_i (-x)^i E_c(P_i^K; s). \tag{12.16}$$

At each fixed rank all expressions are rational functions in these variables. An identity in the indicated Laurent completion is thus an identity of rational functions. We may now substitute

$$x = z^{-2}, \qquad s = zw. \tag{12.17}$$

The substitution is taken in the field of rational functions.

For a box with arm a and leg l, its zeta argument is $x^{a+l+1}s^{2l} = w^{2l}/z^{2a+2}$. Its contribution is

$$z^{-2h(2a+1)}\frac{(z^{2a+1}-w^{2l+1})^{2g}}{(z^{2a+2}-w^{2l})(z^{2a}-w^{2l+2})}. \tag{12.18}$$

The baseline contributes $z^{2h\|\lambda\|^2}$. The identity $\sum_{b\in\lambda}(2a_b+1) = \|\lambda\|^2$ cancels these powers for every partition. Hence $Z_{\mathrm{nil}}(R, z^{-2}; zw) = \Omega(R; z, w)$ from (3.3). Multiplying (12.16) by $(z^2-1)(1-w^2)$ gives

$$\overline{H}_r(z,w) = (zw)^{-2N_r}\sum_i(-z^{-2})^i E_c(P_i^K; zw). \tag{12.19}$$

The signs are additive inverses in the lambda ring. The Adams operations fix -1, so $\psi_d(-s) = -s^d$.

The compact indices in (12.13) have total sign $(-1)^{i_c+j_c} = (-1)^D$ and weight $4N_r - D$. Substitution in (12.19) gives the finite formula

$$\overline{H}_r(z,w) = \sum_{D,k}\beta_{D,k}(-1)^D z^{2k-D}w^{2N_r-D}. \tag{12.20}$$

12.4. **Proof of Theorem 3.1.** Under nonabelian Hodge theory, the $P = W$ theorem [45, Theorem 0.2] identifies the normalized filtration with

$$P_kH^D(M_K, \mathbf{Q}) = W_{2k}H^D(\mathcal{M}_{g,r}, \mathbf{Q}) = W_{2k+1}H^D(\mathcal{M}_{g,r}, \mathbf{Q}). \tag{12.21}$$

This equality of filtrations under the cohomological identification implies $\dim \mathrm{Gr}^W_{2k} H^D(\mathcal{M}_{g,r}) = \beta_{D,k}$. The normalized tautological-generator argument in the following paragraph gives the Hodge–Tate assertion, and hence $H(\mathcal{M}_{g,r}; q, t) = \sum_{D,k}\beta_{D,k}q^kt^D$. For each summand of (12.20),

$$(t\sqrt{q})^{2N_r}(-1)^D(\sqrt{q})^{2k-D}\left(-\frac{1}{t\sqrt{q}}\right)^{2N_r-D} = q^kt^D.$$

This proves Theorem 3.1.

The argument also determines the multiplicities in the rational Tate decomposition of [61, Corollary 1]. We recall its Hodge-theoretic explanation. Shende's theorem places each rational tautological generator of weight $2k$ in $F^k \cap \overline{F}^k \cap W_{2k}$. Tautological generation for coprime Higgs moduli [42, Theorem 7], including the determinant factor for GL_r, shows that products of these classes span cohomology. They give a surjection of Hodge structures from a finite direct sum of $\mathbf{Q}(-k)$ to $H^D(\mathcal{M}_{g,r})$. A quotient of this semisimple object is semisimple. Consequently

$$H^D(\mathcal{M}_{g,r}, \mathbf{Q}) \simeq \bigoplus_k \mathbf{Q}(-k)^{\oplus\beta_{D,k}}. \tag{12.22}$$

Formula (12.10) supplies the multiplicities in this splitting.

Corollary 12.1. *The rational function $\overline{H}_r(z,w)$ is a polynomial, $\overline{H}_r(z,-w)$ has nonnegative integer coefficients, and*

$$\overline{H}_r(z,w) = \overline{H}_r(w,z).$$

Proof. The smooth affine variety $\mathcal{M}_{g,r}$ has dimension $2N_r$. Artin vanishing applied to $\mathbf{Q}[2N_r]$ gives $H^D = 0$ for $D > 2N_r$ [3, Section 4.1]. The ordinary weight bound for smooth varieties gives $2k \geq D$ whenever $\beta_{D,k} \neq 0$. Both exponents in (12.20) are nonnegative; replacing w by $-w$ cancels the displayed sign. Relative hard Lefschetz for the proper Hitchin map gives

$$\beta_{D,k} = \beta_{D+2(N_r-k),\,2N_r-k}.$$

It interchanges the two exponents in (12.20) and preserves the parity of D. □

13. Root curves and the positive comparison

13.1. **Numerical classes and a positive length.** Let C be a smooth projective connected complex curve of genus g, let $p_1, \dots, p_m$ be distinct points, and let $s_1, \dots, s_m$ be positive integers. Throughout this section

$$\nu : \mathcal{X} = C[\sqrt[s_1]{p_1}, \dots, \sqrt[s_m]{p_m}] \longrightarrow C$$

is the tame root stack, E_p is its reduced root divisor, and $K = K_{\mathcal{X}}$. Thus $\nu^* p = s_p E_p$. A local chart at p is $[\operatorname{Spec} \mathbf{C}[t]/\mu_{s_p}]$, with coarse parameter $u = t^{s_p}$. All equivariant complexes below are formed on this chart and then passed to invariants. In particular, invariants are taken before determinants, traces, or cohomology of a complex.

Choose the character convention for which a bundle E has profile $\boldsymbol{\mu} = (\mu_{p,c})_{p, 0 \le c < s_p}$ and age

$$a(\boldsymbol{\mu}) = \sum_p \sum_{c=0}^{s_p - 1} \frac{c}{s_p} \mu_{p,c}.$$

Every row of a rank-r profile has sum r. With this convention the age of the canonical character at p is $(s_p - 1)/s_p$. Put

$$\kappa = \sum_p \left(1 - \frac{1}{s_p}\right), \qquad k = \deg_{\mathrm{orb}} K = 2g - 2 + \kappa.$$

The numerical class retained in all Hall constructions is $(r, \chi, \boldsymbol{\mu})$. The root-curve Riemann–Roch formula reads

$$\deg_{\mathrm{orb}} E = \chi(E) - r(1 - g) + a(\boldsymbol{\mu}). \tag{13.1}$$

It follows either by writing a root line locally as an ordinary line times a fractional divisor, or by filtering a parabolic bundle by its character eigenspaces and applying ordinary Riemann–Roch to the successive elementary modifications.

The additive charge used to choose bounded windows is

$$q(E) = \deg_{\mathrm{orb}} E - \frac{k}{2} \operatorname{rk} E = \chi(E) + a(\boldsymbol{\mu}) - \frac{r\kappa}{2}. \tag{13.2}$$

For a finite sheaf T, this means $q(T) = \deg_{\mathrm{orb}} T$. A nonzero finite sheaf has strictly positive charge: a simple at a root of order s has charge $1/s$, whatever its character. Its ordinary Euler characteristic can be zero. In particular, ordinary Euler characteristic alone cannot be used as a length function for a root Quot scheme.

Lemma 13.1. *Fix a bundle E_0, an endpoint profile $\boldsymbol{\mu}$, and an integer χ. The finite quotients, or finite extensions,*

$$0 \longrightarrow E_0 \longrightarrow E \longrightarrow T \longrightarrow 0, \qquad \chi(E) = \chi, \quad \text{profile}(E) = \boldsymbol{\mu},$$

have a fixed nonnegative modified length, positive when $T \neq 0$. They form a bounded projective Quot problem. The same statement holds in families and includes nonreduced finite modules.

Proof. At a root point use the character convention in which an endpoint line has character $z^{-\gamma}$, $z^s = 1$. Derived restriction of the exact sequence gives

$$[E|_p] - [E_0|_p] = (1 - z)[T_p] \quad \text{in } \mathbf{Z}[z]/(z^s - 1). \tag{13.3}$$

The kernel of multiplication by $1 - z$ is $\mathbf{Z}(1 + z + \cdots + z^{s-1})$. The invariant multiplicity of the displayed regular representation is one. Hence the endpoint characters and the invariant Euler increment determine the full representation class of a punctual defect, once its support is fixed. For an arbitrary distribution of support, additivity gives

$$\deg_{\mathrm{orb}} T = \chi(T) + \sum_p \big(a(E|_p) - a(E_0|_p)\big). \tag{13.4}$$

This formula can also be checked on a local interval of length m: if $\gamma = (m+j) \bmod s$ and $r_0 = \lfloor (m+j)/s \rfloor$, then

$$\frac{m}{s} = r_0 - \left\lfloor \frac{j}{s} \right\rfloor + \frac{\gamma - (j \bmod s)}{s}.$$

Every finite module has a filtration by such simple intervals, so this verification proves the formula without a monomial assumption.

For a global integral measure of orbifold length, set $S = \prod_p s_p$ and use the generating bundle

$$\mathcal{G} = \bigoplus_{0 \le a_p < s_p} \bigotimes_p \mathcal{O}_{\mathcal{X}}(a_p E_p). \tag{13.5}$$

It has rank S and contains every character at p exactly S/s_p times. Consequently

$$N(T) := \chi(T \otimes \mathcal{G}^\vee) = S \deg_{\mathrm{orb}} T = S \left(\chi(T) + \sum_p (a(E|_p) - a(E_0|_p)) \right). \tag{13.6}$$

It is a nonnegative integer and is zero only for $T = 0$. The right side is fixed in the statement. The root Quot construction, using $\mathcal{G}$ and a coarse ample line, therefore gives a projective parameter space with fixed modified Hilbert polynomial. Equivariance, a fixed character class, and the action of any additional finite algebra are closed conditions, with the character class selecting open and closed components. This can be seen directly on an affine chart: a finite cover module is killed by the characteristic polynomial of multiplication by t on the whole module, and its quotients are closed subschemes of a finite-module Grassmannian. Taking invariants is exact, so all the assertions commute with base change. To pass from extensions to quotients, apply $\mathcal{H}om(-, K)$ to the displayed exact sequence. With $T^\dagger = \mathcal{E}xt^1(T, K)$ this gives

$$0 \longrightarrow E^\vee K \longrightarrow E_0^\vee K \longrightarrow T^\dagger \longrightarrow 0.$$

Local duality preserves the cover length at each root point and hence the orbifold length. Serre duality gives $\chi(E^\vee K) = -\chi(E)$, so $\chi(T^\dagger) = \chi(T)$. The endpoint character c becomes $s_p - 1 - c$. For base change, work on an equivariant curve chart over a base ring R_0. Let M be the finite locally free R_0-module representing T, with coordinate action t_M. It has the equivariant resolution

$$0 \longrightarrow R_0[t] \otimes_{R_0} M\, dt \xrightarrow{d} R_0[t] \otimes_{R_0} M \longrightarrow M \longrightarrow 0, \qquad d(m\, dt) = tm - t_M m.$$

Injectivity and exactness follow by division by the monic matrix $t - t_M$. Dualizing into the relative canonical line identifies the sole cohomology, in degree one, with $M^\vee = \mathrm{Hom}_{R_0}(M, R_0)$, with coordinate action $t_M^\vee$. It is finite locally free and commutes with arbitrary base change. Exact invariants preserve these identifications on the root chart. Thus Serre duality identifies the extension problem with the fixed-length quotient problem for $E_0^\vee K$ in families. □

Serre duality has to retain the profile. Write $\mu_{p,c}^\vee = \mu_{p,s_p-1-c}$. Then

$$\chi(E^\vee K) = -\chi(E), \qquad a(\boldsymbol{\mu}) + a(\boldsymbol{\mu}^\vee) = r\kappa, \qquad q(E^\vee K) = -q(E). \tag{13.7}$$

The last equality is why the centered charge is useful. The formal variable x will nevertheless always record χ.

13.2. **A coarse positive quotient and bounded windows.** Fix a rank bound R. Choose a line bundle L_0 on C of sufficiently large even degree ℓ, put $L = \nu^* L_0$, and choose a surjection

$$0 \longrightarrow A = KL^{-1} \longrightarrow K \oplus \mathcal{O}_{\mathcal{X}} \xrightarrow{(b,a)} L \longrightarrow 0. \tag{13.8}$$

Here a and b have no common zero. At a root point the equivariant section b has the required nontrivial character, so a is chosen nonzero there. Away from the roots a general choice has disjoint zero divisors. Enlarge ℓ so that $\deg_{\mathrm{orb}} A < 0$ and every vanishing assertion below holds in ranks at most R. The fact that L is a coarse pullback is essential: its action on the regular-character elementary quotient in the Hecke correspondence is scalar.

Let S_0 be a common multiple of the root orders and define

$$H = \max(k/2, 0), \quad \alpha_r = -H - \frac{\ell(r-1)}{2}, \quad A_r = r\alpha_r, \quad \epsilon_r = \frac{1}{4S_0 r}, \quad c_r = \alpha_r - \epsilon_r.$$

All centered charges lie in $(2S_0)^{-1}\mathbf{Z}$. Thus replacing a weak lower slope bound α_r by the strict cutoff c_r does not alter the allowed numerical classes.

Lemma 13.2. *The bounded comparisons of Lemmas 6.1 and 6.3 hold on $\mathcal{X}$ with the full numerical class and the charge* (13.2)*. In particular:*

(1) *a rank-r cyclic object has $q \geq A_r$;*
(2) *the cutoffs are nested under positive-rank factors;*
(3) *the unframed source and the comparison error are supported in finite intervals of q, independent of perverse degree.*

Proof. Filter a cyclic quotient by normal monomials, as in Lemma 6.1. If r_j are its consecutive positive multiplicities, then $\sum_j r_j = r$ and $\sum_j j r_j \leq r(r-1)/2$. A monomial of index j contributes degree at least $-j\ell$. Hence $\deg_{\mathrm{orb}} E \geq -\ell r(r-1)/2$, and subtracting $rk/2$ gives $q(E) \geq A_r$. The same argument applies to a final positive Higgs quotient: the negative A-field preserves its projected Harder–Narasimhan filtration. In a fixed profile this implies the concrete lower bound

$$\chi(E) \geq A_r - a(\boldsymbol{\mu}) + r\kappa/2.$$

For $e < r$,

$$c_e - c_r = (r-e)\left(\frac{\ell}{2} - \frac{1}{4S_0 er}\right) > 0,$$

which proves the required nesting.

There is one additional point for coherent wall factors. If T_0 is their torsion and the positive-rank factors have slopes $\mu_j \geq \lambda$, equality at the wall gives

$$\lambda f = q(T_0) + \sum_j f_j \mu_j \geq q(T_0) + \lambda f.$$

Lemma 13.1 forces $T_0 = 0$. Thus the argument eliminates also the nontrivial-character simples with zero ordinary Euler characteristic. A finite-rank-zero framed factor at finite framing parameter is the vacuum: on a smooth surface chart a finite sheaf T has $\mathrm{Hom}(T, K_S^{-1}) = \mathrm{Ext}^1(T, K_S^{-1}) = 0$; exact finite-group invariants preserve these vanishings.

The physical reflection constructed below is

$$X^t_{r,\chi,\boldsymbol{\mu}} \simeq X^{-t}_{r,-\chi,\boldsymbol{\mu}^\vee}, \qquad B^t_{r,\chi,\boldsymbol{\mu}} \simeq B^{-t}_{r,-\chi,\boldsymbol{\mu}^\vee} \otimes \mathbb{L}^\chi.$$

Together with the lower bound it gives the finite source interval $[A_r, -A_r]$. It does not identify two equal profiles unless the profile is self-dual.

For the other endpoint use the elementary Higgs slope estimate

$$\mu^{\mathrm{orb}}_{\min}(E) \geq \frac{q(E)}{r} + \frac{k}{2} - (r-1)\ell.$$

It follows by inserting the successive Harder–Narasimhan quotients: a gap larger than ℓ makes the upper part L-Higgs invariant. Thus $H^1(E) = 0$ as soon as $q(E)/r > k/2 + (r-1)\ell$. Choose adjacent framing chambers with distance less than $1/(4S_0 r^2)$. If an underlying positive Higgs object at this endpoint were unstable, its highest Harder–Narasimhan part would have vanishing H^1, zero dual framing, and be preserved by the negative A-field. It would destabilize the framed object. The remaining framing map is therefore a smooth projective Joyce–Song framing over the semistable positive moduli, and its comparison is an isomorphism. This gives an upper endpoint for the error. The proof is about objects and correspondences, so the same endpoint works in every perverse degree. □

13.3. **Physical maps and their orientations on a root curve.** We spell out the root descent of the physical construction, because descent of determinant lines alone would not identify the maps used in the comparison. For $S_L = \mathrm{Tot}_{\mathcal{X}} L$ and a compactly supported spectral sheaf G, let

$$B_{\mathcal{X}}(G) = \mathrm{fib}\big(\mathrm{RHom}_{S_L}(G,G) \longrightarrow R\Gamma(S_L,G)\big), \qquad M_{\mathcal{X}}(G) = \det B_{\mathcal{X}}(G). \tag{13.9}$$

The arrow is composition with the section. The same definition with two entries gives the off-diagonal halves.

Proposition 13.3. *The critical equivalence, paired deformation triangle, orientation square, supported determinant comparison, and parameter square of Proposition 4.1 are valid over $\mathcal{X}$. They are equivalences and commutative diagrams of the enhanced complexes, compatible with arbitrary Artin base change, direct sums, and finite-support modifications. Their orientation is $M_{\mathcal{X}}$ in* (13.9)*.*

Proof. The construction of 4.1 uses the relative diagonal resolution of $\mathbf{P}^1$, the monic two-term spectral resolution, evaluation against the framing, composition of Hom complexes, and the relative Serre trace. Each is functorial for equivariant sheaves on a curve chart. In characteristic zero the invariant functor is exact, commutes with tensoring by a trivial base algebra, and commutes with the dual of a perfect complex. Apply it to this complete diagram. In particular the triangle is

$$B_{\mathcal{X}}(G_j,G_i)^{\vee}[-3] \longrightarrow C_{\mathcal{X}}(G_i,G_j) \longrightarrow B_{\mathcal{X}}(G_i,G_j).$$

This establishes the pairing and the equivalence of derived mapping spaces, including maps between objects over nonreduced Artin bases. The local Calabi–Yau coordinates (t,y,z) have root weights $(1,-1,0)$, so $dt\wedge dy\wedge dz$ is invariant.

The trace has exactly the ordinary normalization. With $u=t^s$, the pullback of $F(u)du$ is $sF(t^s)t^{s-1}dt$. Stack integration contributes $1/s$ times cover integration. These two factors cancel, giving

$$\nu_* K_{\mathcal{X}} = K_C \tag{13.10}$$

with its Serre trace, and residue one for the pullback $s\,dt/t$ of du/u. No scalar square root of a root order enters this construction.

The closed-form comparison is equally equivariant. The normal Euler vector field commutes with the root group. In the notation of the proof of 4.8, the Cartan homotopy is

$$H_0 = h_0 + \tfrac{1}{2}\iota_{\mathrm{Eu}}(\delta - Dh_0).$$

It may be taken before invariants. The derivative of this same cotangent equivalence gives the determinant-square map after invariants. This proves compatibility with the specified maps in the square. The virtual dimension is recomputed on the invariant complex.

For $0\to G_0\to G\to T\to 0$, use the four-Hom grid with entries $\mathrm{RHom}(T,G_0)$, $\mathrm{RHom}(T,T)$, $\mathrm{RHom}(G_0,T)$, and $R\Gamma T$. Its cones, traces, and derivative determinant maps are again equivariant. The traceless correction has determinant one: in a finite free presentation it is $\det(1-uv) = \det(1-vu)$, and the latter is one on the physical rank-zero cofiber. This proves the supported comparison.

Two sign points require retaining more information than in the ordinary curve. First, on the canonical surface the Euler form is symmetric and its self-pairings are even, since the Ext^1 pairing is symplectic. Hence

$$\chi(B(G)) - \chi(B(G_0)) = 2\chi_{S_K}(G_0,T) + \chi_{S_K}(T,T) - \chi(T) \equiv -\chi(T) \pmod 2. \tag{13.11}$$

Individual surface pairings need not vanish at a root. The Euler characteristic on the right is computed after invariants.

Second, for two distinct scalar packets write $C_{21} = C_{12}^{\vee}[-2]$ and $A_{21} = -A_{12}^{\vee}$ for their cross differential. Put $L_{12} = \det C_{12}$, of parity $c = \chi(C_{12}) \bmod 2$. The two cross cones have

determinant lines $L_{12} \otimes L_{12}^{\vee}$ and $L_{12}^{\vee} \otimes L_{12}$. Graded interchange contributes $(-1)^c$, while $\det(-A_{12}^{\vee}) = (-1)^c \det A_{12}$. Consequently their acyclic trivializations satisfy

$$\tau_{21} \circ \beta = \tau_{12}.$$

This identity is functorial over the parameter ring; it follows degree by degree for free complexes and then from determinant coherence for exact filtrations. Taking inverses gives precisely the comparison of the two specified orientation roots. Thus there is no extra character in packet permutations even when a cross-profile Euler pairing is odd.

Finally, a change of positive quotient has residual quadratic space $(\mathrm{End}(V) \otimes R_{\mathrm{res}})^{\mu_s}$. Its equivariant change-of-frame group has reductive part $\prod_c \mathrm{GL}(V_c)$ and connected unipotent radical. Its determinant character into $\{1, -1\}$ is trivial. The residual orientation line is therefore constant for a fixed profile; ordered color blocks and the graded determinant functor give its direct-sum compatibility. Ordinary elementary modifications preserve this profile. Root modifications use the ratio of the two constant lines together with the same determinant maps. For the comparison with the canonical projection, the residual module is $K/(aA) = K|_{D_a}$, where $D_a = (a = 0)$ avoids the root points and has even length ℓ. The ordinary even-residue interchange calculation therefore applies directly. These constructions glue by equivariant descent, completing the proof. □

13.4. **The moving root Hecke correspondence.** Fix an ordinary point $y \in C$. The creation point x moves on the whole coarse curve C. An elementary quotient is annihilated by the pullback of the ideal of its coarse graph, has invariant Euler characteristic one, and has regular character at a root. More explicitly, over a base algebra R near a root of order s it is an R-flat module over

$$R[t]/(t^s - a) \tag{13.12}$$

whose character summands all have rank one; a is the moving coarse coordinate. Multiplication by t need not be zero when $a = 0$. Require that the L-field act as a scalar λ on this quotient and that the section vanish in it. Its kernel is locally free by the fiberwise torsionfree criterion on the regular root chart. Moreover

$$[E_+|_p] - [E_-|_p] = (1 - z)[\mathrm{Reg}] = 0,$$

so the profile is unchanged and both q and χ increase by one. The correspondence is the schematic closure of this incidence over ordinary x, with the fundamental multiplicities of that dense incidence.

Proposition 13.4. *With these elementary quotients the whole-arrow Weyl relation of Proposition 7.1 holds on the root curve:*

$$F_{\chi}^{c} E_{\chi}^{c-1} - E_{\chi-1}^{c} F_{\chi-1}^{c-1} = r\, \rho_{\chi}(c-2, c). \tag{13.13}$$

It holds before taking a cone, perverse cohomology, or a Grothendieck class, and is compatible with the physical orientation of Proposition 13.3.

Proof. Use the three bounded bundle levels of the ordinary construction, with a one-unit buffer in q. For their finite collection of Harder–Narasimhan types choose a coarse divisor D disjoint from the active graphs and large enough that

$$H^1(\mathrm{End}(E)L(D)(-J)) = H^1(E(D)(-J)) = 0$$

for the thick union J of the graphs. Exact ν_* and ordinary Serre vanishing give this uniformly. The framed bundle stack maps smoothly to its J-jets: the obstruction lies in H^2 on a curve. A moving root jet has frame group the units in an invariant finite locally free endomorphism algebra, an open smooth group scheme. The two displayed vanishings make the field and section restrictions jointly surjective vector bundle maps. The residue variables at D are independent of the active jets. Hence every correspondence chart is a smooth pullback from the local jet incidence.

For $x \neq y$, the two ordered Hecke words have a common four-bundle receiver. If $U \to T_x, T_y$ are their quotients, put

$$E_1 = \ker(U \to T_x), \quad E_2 = \ker(U \to T_y), \quad J_0 = \ker(U \to T_x \oplus T_y).$$

Then $U = E_1 \amalg_{J_0} E_2$. The derived cross-Homs of the two quotients vanish, so this bicartesian diagram and its fields and sections are unique up to a contractible space of choices. The trace functions agree as functions on the whole scheme. The residue support equations agree as well, since they are restrictions along relative Cartier divisors.

On the smooth jet presentation the variables for x and y are independent. Flat base change commutes with the schematic closure of the ordinary-x incidence. In particular, if A_x, A_y are the jet algebras and R_x, R_y their incidence algebras, then the two orders have the same derived tensor product, namely the smooth pullback of $R_x \otimes_{\mathbf{C}} R_y$ over $A_x \otimes_{\mathbf{C}} A_y$. There is no cross Tor term. Their refined Borel–Moore cycles therefore agree off $x = y$, including when x is a root point.

For clarity, the proper carrier on which this comparison is made can be constructed without using a special-fiber cycle. For a fixed final pair (E, Φ, s) let $H = \nu^*(x + y)$ and let σ_H be its canonical section. Consider the intermediate kernels

$$E(-H) \subset G \subset E(H), \quad \chi(G) = \chi(E), \quad \text{profile}(G) = \boldsymbol{\mu}, \quad \Phi(G) \subset G \otimes L, \tag{13.14}$$

with $\sigma_H s$ factoring through $G(D)$. Relative to $\mathcal{G}$ the finite module $E(H)/E(-H)$ has modified length $4rS$ and $E(H)/G$ has length $2rS$. These lengths are constant also at a root and on the collision diagonal. Lemma 13.1 makes the carrier projective; the last conditions are closed. Each Hecke word maps properly to it, by multiplying its initial bundle by σ_H.

The difference of their cycles is supported on the endpoint graph $x = y$, $G = \sigma_H E$. The prescribed degree makes its coefficient a degree-zero class on each graph component. At a dense point with regular semisimple Φ_y and nonzero section in each eigenline, removal first is absent and creation then removal has r transverse branches, each of multiplicity $+1$. Thus the coefficient is r. Jet surjectivity proves that this computation meets every graph component. The collision takes place at the ordinary point y, so its orientation is exactly the local orientation computed in Proposition 7.1.

This cycle equality gives the full arrow equality. In the notation of the supported coreflector (7.14), for an object A_0 supported on the residue support S_Y, the natural map

$$\mathrm{Map}(A_0, \Gamma_{S_Y} K_Y) \longrightarrow \mathrm{Map}(A_0, K_Y)$$

is an equivalence of mapping spaces. It preserves the homotopy just constructed. Forward vanishing-cycle exchange, its proper inverse, and the right mate therefore lift the entire difference arrow and its cofiber.

To descend from a chosen residue divisor, choose two sufficiently large ordinary divisors D_1, D_2 with disjoint supports, both avoiding the roots and y. Their admissible chart opens cover the moving curve. Compare both with $D_1 + D_2$. The same H^1 vanishings choose splittings zero on the active jets, with an affine space of choices. The added hyperbolic block has rank $\deg(D' - D)r(r + 1)$, which is even, and its counit is $+1$. Full frame changes are handled by smooth descent. The enhanced coherences glue before taking proper trace over C. That trace cancels the L^{-1} factor and takes the class of y to one. This proves (13.13). □

Corollary 13.5. *The conclusions of Theorems 7.3 and 7.4 hold for the full root numerical class. All their arrows retain the specified physical orientations and finite permutation actions.*

Proof. Use the finite intervals of Lemma 13.2 and the whole-arrow relation (13.13). The nilpotence argument of Lemma 7.2 is an argument in a bounded enhanced diagram; it depends only on these two facts and on the nonzero scalar r. It applies separately to every profile and to its dual

profile. Lemma 13.1 gives the finite graded-face data required for the bounded Hall diagrams. Proposition 13.3 supplies all the root orientation squares used in their maps. □

14. Positive root Hitchin support and canonical primitives

14.1. **Full support in every genus.** For a fixed profile write $M^L_{r,\chi,\boldsymbol{\mu}}$ for the coarse moduli of semistable L-Higgs bundles and

$$h_L : M^L_{r,\chi,\boldsymbol{\mu}} \longrightarrow \mathcal{A}^L_r = \bigoplus_{j=1}^{r} H^0(C, L_0^j)$$

for the positive Hitchin map. The weights are rational; replacing the root orders by common multiples realizes any rational weights. The profile keeps its ordered character labels under this replacement. We use IC with its perverse normalization. Set

$$\begin{aligned} b_r &= \dim \mathcal{A}^L_r = \frac{\ell r(r+1)}{2} + r(1-g), \\ \gamma_r &= 1 + r(g-1) + \frac{\ell r(r-1)}{2}, \qquad B_r = \ell r^2. \end{aligned} \tag{14.1}$$

Thus $b_r + \gamma_r = B_r + 1$. The character terms in $\chi(\operatorname{End} E)$ and $\chi(\operatorname{End} E \otimes L)$ cancel, because L has trivial inertia.

Theorem 14.1. *For sufficiently large ℓ the following statements hold for every $g \geq 0$, every profile, and every Euler characteristic. The moduli $M^L_{r,\chi,\boldsymbol{\mu}}$ is irreducible and has a dense stable locus. Every simple constituent of*

$$^{p}\mathcal{H}^i Rh_{L*}\,\mathrm{IC}_{M^L_{r,\chi,\boldsymbol{\mu}}}$$

has full support $\mathcal{A}^L_r$. These perverse Hodge modules are independent of χ, with the identifications obtained from the degree-independent Jacobian variation on the smooth spectral locus and intermediate extension.

Proof. First suppose that the parabolic numerical type admits weights for which semistability is stability. This always holds after the auxiliary full-flag refinement used at the end of the proof; a divisible partial-flag type need not admit such weights itself. Choose weights in such a chamber. The obstruction space is dual to $\operatorname{Hom}_{\mathrm{Higgs}}(E, E \otimes A)$, which vanishes because $A = KL^{-1}$ has negative degree. The stable coarse space is therefore smooth of dimension $B_r + 1$. The same argument works when the distinct marked points vary over $U = \operatorname{Conf}_m(C)$ and proves smoothness relative to U. The Hitchin map is proper: after the spectral construction the objects are semistable sheaves with fixed modified Hilbert polynomial on the projective finite spectral curve, and the valuative extension followed by semistable reduction is unique on the coarse moduli. Tameness and the generating bundle (13.5) give the same construction for root sheaves.

We require a uniform fiber-dimension bound. Choose a reduced divisor $D' \in |L_0 K_C^{-1}|$ containing all marked points. On the nilpotent fiber the weakly parabolic residues preserve their partial flags and are nilpotent. Refine each flag to a full flag strictly lowered by its residue, and add full flags at the other points of D'. Perturb the weights by sufficiently small generic numbers. This preserves stability of the old objects: the possible subprofiles are finite and their integral subdegrees give a positive gap from zero in each strict inequality. The resulting full strongly parabolic nilpotent moduli maps onto the original nilpotent fiber.

Its dimension is at most

$$1 + r^2(g-1) + \deg(D')\frac{r(r-1)}{2} = \gamma_r. \tag{14.2}$$

Here is the isotropic argument, which also covers $g = 0, 1$. Resolve a nilpotent stratum by a filtration on which the field strictly lowers degree. A tangent deformation of the bundle preserves this filtration, whereas the field deformation lowers it. Their composite is strictly filtration-lowering and has trace zero. Thus the Liouville one-form on the strongly parabolic cotangent

moduli restricts to zero on this resolution. Its differential, the trace symplectic form, also restricts to zero. On a dense smooth locus of each image component this makes the component isotropic, so its dimension is at most half the dimension of the stable strongly parabolic moduli. Riemann–Roch with a full flag at each point of D' gives that half-dimension as the left side of (14.2). Filtration ranks and degrees are bounded on the semistable locus, so finitely many such strata suffice. This is also the filtration-and-trace proof of the parabolic nilpotent isotropy statement in [5, Section 3.2]; it uses no lower bound on the genus. Properness and positive-weight scaling of the Hitchin base imply that every fiber has dimension at most γ_r.

Let P be the identity component of the ordinary spectral Picard acting by tensor product. Its multidegree on every reduced spectral component is zero. It therefore preserves parabolic degree, every subobject inequality, and the root-character profile. Stabilizers are affine. Indeed, on the normalization of every reduced component the sheaf has positive rank. A line bundle stabilizing it has a torsion power by taking its determinant, so the stabilizer image in the product of the normalization Jacobians is finite. The kernel of the map from the spectral Picard to that product is affine. Its Tate module has the polarization induced from the normalization Jacobians.

We use the weak abelian support inequality in the following precise form. For a proper map from a smooth space with a polarizable relative Picard action and affine stabilizers, and with fiber dimension bounded by the arithmetic genus γ_r, a strict support Z satisfies

$$\operatorname{codim} Z \leq \delta_Z,$$

where δ_Z is the dimension of the affine part of the generic Picard. Equivalently its abelian dimension β_Z satisfies $\beta_Z \leq \gamma_r - \operatorname{codim} Z$. We also use the integral spectral-curve inequality $\beta_W \geq \gamma_r - b_r + \dim W$ on an integral locus. Both follow from the exterior-algebra action of the Tate module and the integral δ-regularity estimate; these are the support inputs in [11, 12] and [44]. Their numerical comparison is given here to specify their application to moving partial flags.

Let W be the projection of $Z \subset \mathcal{A}_r^L \times U$ to $\mathcal{A}_r^L$. At its generic point write the spectral polynomial as a product of distinct integral factors of degrees $n_1, \ldots, n_s$, with multiplicities $m_1, \ldots, m_s$, so that $r = \sum_i m_i n_i$. The finite multiplication map gives a finite cover of W with projections W_i to the integral degree-n_i Hitchin bases. The abelian part is the product of their normalization Jacobians. The two preceding inequalities give

$$\sum_i (\gamma_{n_i} - b_{n_i} + \dim W_i) \leq \beta_W \leq \gamma_r - b_r + \dim W \leq \gamma_r - b_r + \sum_i \dim W_i. \tag{14.3}$$

The last step uses only $\dim W \leq \sum_i \dim W_i$; the projections need not be independent. Since $\gamma_n - b_n = 1 + n(2g - 2 - \ell)$, cancellation gives

$$1 - s \geq \left(r - \sum_i n_i \right) (\ell - 2g + 2). \tag{14.4}$$

As $\ell > 2g - 2$, this forces $s = 1$ and $m_1 = 1$. The integral inequality then forces equality throughout the dimension estimate for Z, so $Z = W \times U$.

Now move the marked points away from the branch and singular loci of the generic curve in W. Over this open set the choice of profile is the finite etale allocation cover with fiber

$$\prod_p \mathfrak{S}_r / \prod_c \mathfrak{S}_{\mu_{p,c}}. \tag{14.5}$$

The spectral sheaf is locally a line on each of these branches; prescribing its characters changes neither its coarse Jacobian nor its ordinary Euler characteristic. The direct image is therefore the ordinary positive direct image tensored with this finite allocation variation. Here the ordinary integral-locus support statement is needed in every genus. In its precise form, for a smooth projective curve of any genus, an effective twist of degree $> 2g-2$, and any integer spectral degree, the proper compactified Jacobian over the integral spectral locus has smooth total space and

every perverse direct-image constituent is the intermediate extension from the smooth spectral locus. This is the elliptic support theorem in [11, Sections 2 and 4, Theorem 7.1]; the all-base coprime form is [11, Theorem 9.1] with the correction in [12]. Its genus hypotheses therefore apply also for $g = 0, 1$. The support assertion for rational Hodge modules follows from this underlying rational perverse support assertion and the pure Hodge decomposition theorem.

One can also reduce the arbitrary-degree integral statement to its coprime form locally, which makes its use here transparent. The smooth locus of the universal integral spectral curve is a smooth surjective morphism. Etale locally around any point of its base, choose a section through a smooth point avoiding all marked branches. Its divisor defines a relative degree-one line bundle. Tensoring by an integral power identifies the compactified Jacobians of any two Euler characteristics. Integral spectral support makes all their rank-one torsionfree sheaves stable, so this is an isomorphism of the whole moduli problem in this neighborhood, with no stability wall. Choose the second Euler characteristic coprime to r and apply the all-genus coprime support theorem.

The allocation cover (14.5) is finite etale in the neighborhood under consideration; after a further etale pullback it is a disjoint union of copies of that neighborhood. On each copy root descent is precisely the ordinary compactified-Jacobian problem just described, and tensoring at the chosen ordinary section leaves all root characters unchanged. Etale pullback and finite-group projectors commute with intermediate extension. Thus this allocation cannot produce a proper strict support. It follows that $W = \mathcal{A}_r^L$, proving full support over $\mathcal{A}_r^L \times U$.

Restriction to fixed marked points preserves that conclusion. The source is smooth over U. A nonzero covector in a U-direction has nonzero pullback, so the proper direct-image estimate for singular support shows that the slice at a fixed point of U is noncharacteristic. Its shifted pullback is perverse exact and takes the intermediate extension just found to the intermediate extension of the restricted variation. Thus it creates no new strict support.

We next prove the irreducibility needed at nongeneric weights, without assuming the existence of stable underlying bundles. The stack of root bundles of a fixed rank, degree, and profile is a partial-flag fibration over the ordinary bundle stack, hence is irreducible. On a Harder–Narasimhan stratum let $E_1, \dots, E_v$ be the semistable pieces in strictly descending order of slope. Its codimension is

$$c = \sum_{i<j} -\chi_{\mathcal{X}}(E_i, E_j) = \sum_{i<j} h^0(\mathrm{Hom}(E_j, E_i)K) \geq 0. \tag{14.6}$$

The excess dimension of the L-Higgs fiber is

$$h^1(\mathrm{End}(E)L) = \hom(E, EA) \leq \sum_{i<j} h^0(\mathrm{Hom}(E_j, E_i)A).$$

For every nonzero summand on the right of (14.6), choose an ordinary point where one of its K-valued sections is nonzero, and a section of L_0 vanishing there. Multiplication embeds the corresponding A-valued space into a proper subspace of the K-valued space. Therefore the excess is strictly less than c if $c > 0$; if $c = 0$, the excess is zero. Thus every positive codimension stratum contributes dimension strictly less than B_r to the Higgs stack.

Only finitely many such strata meet Higgs semistability. Successive slope gaps are at most ℓ, since a larger gap makes the upper part Higgs invariant; the total spread is at most $(r-1)\ell$. Fixed total degree and the finite set of profiles then bound all the degrees. The Higgs stack is smooth of pure dimension B_r. There are stable points for every χ and every profile: take a smooth connected spectral curve unramified over the marks, allocate the characters, and take a line of the desired Euler characteristic. Its integral spectral support permits no proper invariant generic subspace. Every component therefore meets the unique codimension-zero Harder–Narasimhan stratum of the irreducible bundle stack. On that stratum the Higgs stack is a vector bundle and

its semistable open is irreducible. The whole semistable stack, and hence its coarse space, is irreducible with dense stable locus. The codimension-zero stratum can be unstable on $\mathbf{P}^1$.

Finally, at a weight wall introduce an auxiliary ordinary full invariant flag and perturb its weights by a sufficiently small generic amount. The new semistable objects are old semistable objects, and their forgetful map $\pi : M_{\mathrm{new}} \to M_{\mathrm{old}}$ is projective over the Hitchin base. Over the smooth unramified spectral locus it is finite etale of degree $r!$. Its image is dense and closed, hence all of the irreducible old space. The trivial summand of the generic permutation local system extends, by the decomposition theorem, to a direct summand $\mathrm{IC}_{M_{\mathrm{old}}}$ of $R\pi_*\mathbf{Q}_{M_{\mathrm{new}}}[B_r+1]$. Full support of the new direct image therefore implies full support of the old IC direct image.

On the smooth spectral allocation cover the fibers are Jacobian torsors of varying degree. Their cohomology is the exterior algebra of the same H^1 variation and is independent of the degree. The allocation system also depends only on the profile. Full support and uniqueness of pure intermediate extension now give the asserted isomorphisms of perverse Hodge modules for all χ, including wall classes. □

14.2. **Active profiles and the target Euler characteristic.** Fix a target profile $\boldsymbol{\mu}$ of rank n. A profile $\boldsymbol{\beta}$ is called active if $\beta_{p,c} \le \mu_{p,c}$ for all p, c, and the sums $\sum_c \beta_{p,c}$ have the same value r at every marked point. The active set is finite and is closed under taking summands of an active total.

Lemma 14.2. *Within a prescribed target determinant hyperplane, one can choose rational weights for which every active profile is stable whenever it is semistable at Euler characteristic* $\chi_* = 1$.

Proof. Let the weights be $\alpha_{p,c}$ and write the target hyperplane as $\sum_{p,c} \mu_{p,c}\alpha_{p,c} = -d_0$. For an active rank-$r$ profile $\boldsymbol{\beta}$ and a proper rank-r' subprofile $\boldsymbol{\beta}'$, the slope wall at $\chi_* = 1$ is

$$(r\boldsymbol{\beta}' - r'\boldsymbol{\beta}) \cdot \boldsymbol{\alpha} = r' - r\chi'. \tag{14.7}$$

In a bounded weight chamber only finitely many integers χ' occur. The sum of the entries of its normal at each marked point is zero, whereas that of the target normal is n. Consequently a nonzero wall normal is not proportional to the target normal. If the wall normal is zero, its right side is nonzero because $0 < r' < r$. No wall contains the target hyperplane. A finite union of its proper rational hyperplanes cannot cover an open weight chamber. A rational point in the complement has the claimed property. □

The Euler characteristic used in this auxiliary choice is not in general the degree-zero parabolic Euler characteristic. For finite-order unitary monodromies $\exp(-2\pi i\alpha_{p,c})$ the latter is

$$\chi_{\mathrm{target}} = n(1-g) - \sum_{p,c} \mu_{p,c}\alpha_{p,c}. \tag{14.8}$$

It is an integer by the determinant condition. The use of $\chi_* = 1$ below is followed by the perverse-Hodge-module transport of Theorem 14.1. Applying the same vanishing-cycle functor to that isomorphism transports the canonical primitive to (14.8). Receiver profiles and dual profiles occurring in a comparison need not be active. They use the all-profile positive support theorem; no simultaneous genericity assertion for this infinite set is required.

14.3. **The root and coarse surface halves.**

Lemma 14.3. *On the reduced nodal spectral locus whose nodes avoid the marks and whose marked fibers are unramified with distinct eigenvalues, passage to the allocation cover identifies the root stable-pair deformation half with its ordinary coarse counterpart. This is an identification of the section arrow, the Serre pairing, and the physical determinant-square maps.*

Proof. Fix the character on each marked spectral branch. A rank-one torsionfree spectral sheaf is a line near these branches. Removing its fixed fractional divisor makes its inertia trivial, hence gives a unique coarse torsionfree sheaf G. Conversely pullback and tensor with this fractional divisor recover the root sheaf F. With residues in $0, \dots, s_p - 1$, its divisor line pushes forward to $\mathcal{O}$. Sections $\mathcal{O} \to F$ therefore correspond exactly to sections $\mathcal{O} \to G$. Purity and finite cokernel are preserved, as is ordinary Euler characteristic. Exact root descent makes these statements valid over every Artin base.

This equivalence also identifies the deformation half, a stronger assertion than equivalence of the stable-pair schemes. Near a branch $y = f(u)$ with character c the monic resolution is

$$[\mathcal{O}\chi^c L^{-1} \xrightarrow{\cdot, y - f(u),} \mathcal{O}\chi^c].$$

The characters cancel in self-Hom, giving the ordinary monic Hom complex after invariants. The Hom complex between two different branches is contractible, since their eigenvalue difference is a unit. Thus

$$\mathrm{RHom}_{\mathrm{Tot}_{\mathcal{X}}\, L}(F, F) \simeq \mathrm{RHom}_{\mathrm{Tot}_C\, L_0}(G, G).$$

Moreover $R\Gamma(\mathcal{X}, F) = R\Gamma(C, G)$, and the map induced by the section agrees. These are derivatives of the descent functor in families, so they identify the two fibers in (13.9). Equation (13.10) identifies their dualities and traces. Their determinants therefore carry the same derivative square map. The split potential is zero on both sides, and the nonsplit potential is the same trace quadratic form. There is no additional character-dependent half-Tate factor.

At a node the curve is ordinary. The node deformation, edge orientation, graph contraction, and punctual Hilbert factor are consequently the ones proved in 8.6 and 8.7. Allocations restrict to spectral components and their profiles add. Taking the finite allocation-group projectors preserves the identified maps. This proves the claim with all finite actions retained. □

Proposition 14.4. *The positive source formula of Theorem 8.1 holds profile by profile on $\mathcal{X}$. Its primitive is the positive perverse IC object of Theorem 14.1, and its scalar denominator is $x/((1-x)(1-x\mathbf{T}))$.*

Proof. The smooth-stack integrality input used in Theorem 8.1 has the same hypotheses here. The Euler difference is $-\ell r s$, since the character corrections cancel for a coarse twist; the semistable stack is smooth; good moduli and bounded graded-face data come from the generating-bundle construction and Lemma 13.1. Stabilizers are the same products of general linear groups on stable factors, with the same almost-orthogonality. Irreducibility and the dense stable locus in Theorem 14.1 identify the primitive with the coarse IC object.

On the good nodal open, Lemma 14.3 identifies the relative Hilbert source including its Hodge variation and determinant maps. The fractional root zeros have ordinary Euler increment zero. The Hilbert length is therefore still $\chi + \gamma_r - 1$, and the two ordinary free directions give the same denominator $x/((1-x)(1-x\mathbf{T}))$. The finite-addition restriction injection of 8.4 and 8.3 applies to the allocation covers. Together with full support it extends the identity from this open to the whole base. Only the finite set of active profiles enters a fixed target coefficient.

Scalar error detection in 10.1 extends to the root curve as follows. Coprime polynomial idempotents split the spectral bundles over every base, so their profiles add componentwise. Independent translations of at least two positive-rank packets give the same detecting family. The unique one-factor rank-r unframed vacuum is removed by the complete receiver before forming the error. Let $\alpha \in H^1(\mathcal{X}, KL^{-2})$ be the extension class of (13.8). The translation $\Phi \mapsto \Phi + ac\,\mathrm{id}$ preserves its trace potential: the lift of a to the $\mathcal{O}$ summand gives $\alpha a = 0$. The cross-profile determinant comparison in Proposition 13.3 identifies its permutation maps. Corollary 13.5 therefore supplies the source comparison needed for this extension. □

14.4. **The critical scheme and the quadratic normal bundle.** For the extension class α of (13.8), define on the smooth stable positive space

$$g(E,\Phi) = \tfrac{1}{2}\langle \alpha, \operatorname{Tr}(\Phi^2)\rangle = g_{\mathcal{A}} \circ h_L. \tag{14.9}$$

For a stable numerical class put

$$d_{\boldsymbol{\mu}} = 2N_{\boldsymbol{\mu}} = 2(g-1)r^2 + 2 + \sum_p \left(r^2 - \sum_c \mu_{p,c}^2 \right). \tag{14.10}$$

The following assertion includes the case of an empty canonical moduli space.

Proposition 14.5. *If semistability is stability in the given class, the full classical critical scheme of* (14.9) *is*

$$\operatorname{Crit}(g) = M^K_{r,\chi,\boldsymbol{\mu}} \times \mathbf{A}^1. \tag{14.11}$$

The isomorphism sends (E,θ,λ) *to* $(E, b\theta + a\lambda \operatorname{id})$. *It is an isomorphism on arbitrary ordinary test schemes, including nonreduced schemes. If the right side is nonempty, it is smooth of dimension* $d_{\boldsymbol{\mu}} + 1$ *and its normal Hessian has even rank* $B_r - d_{\boldsymbol{\mu}}$.

Proof. We verify the gradient ideal directly. On a smooth local universal chart of the stable positive moduli use the following two-term complexes, with displayed terms in degrees zero and one:

$$C_\Phi = [\operatorname{End} E \xrightarrow{[\Phi,-]} \operatorname{End} E \otimes L],$$

$$D_\Phi = [\operatorname{End} E \otimes A \xrightarrow{[\Phi,-]} \operatorname{End} E \otimes K],$$

$$B_\Phi = [\operatorname{End} E \otimes (K \oplus \mathcal{O}) \xrightarrow{[\Phi,-]} \operatorname{End} E \otimes (K \oplus \mathcal{O})L],$$

$$C^L_\Phi = [\operatorname{End} E \otimes L \xrightarrow{[\Phi,-]} \operatorname{End} E \otimes L^2].$$

There is an exact sequence of these complexes

$$0 \longrightarrow D_\Phi \longrightarrow B_\Phi \longrightarrow C^L_\Phi \longrightarrow 0. \tag{14.12}$$

The tautological Φ is a closed degree-zero section of C^L_Φ. Choose local splittings of (13.8). Their differences represent α, and the connecting class of Φ is

$$\delta\Phi = [\alpha\Phi, 0] \in \mathbb{H}^1(D_\Phi).$$

The Serre trace pairing identifies this class with dg. Indeed its pairing with a field variation is $\langle \alpha, \operatorname{Tr}(\Phi\,\delta\Phi)\rangle$; with a transition-function variation it is the same Cech cocycle paired against the degree-zero part of C_Φ. These are precisely the two terms in the variation of (14.9). Thus bundle variations, as well as field variations, occur in this identification.

On every geometric fiber stability and negativity give

$$H^0(C_\Phi) = \mathbf{C}\operatorname{id}, \quad H^2(C_\Phi) = 0, \qquad H^0(D_\Phi) = 0, \quad H^2(D_\Phi) = \mathbf{C}.$$

The middle dimensions are constant by Riemann–Roch. The relative complexes are perfect. On this smooth universal chart, represent them locally by finite free complexes. Constant fiber dimensions give constant-rank differentials on this reduced parameter space; elementary row and column reduction splits off their acyclic summands. The remaining complexes have zero differentials and locally free terms. Pulling these split complexes back to any ordinary test scheme shows that their cohomology is locally free and commutes with arbitrary ordinary base change. In particular

$$H^0(C_\Phi) = \mathcal{O}\operatorname{id}, \quad H^1(C_\Phi) = T_{M^L}, \qquad H^1(D_\Phi) = \Omega_{M^L}. \tag{14.13}$$

The assertion $H^0(D_\Phi) = 0$ remains true over nonreduced bases.

Let $T \to \operatorname{Crit}(g)$ be affine. Its gradient is zero, so $\delta\Phi = 0$ in the relative hypercohomology. The exact sequence (14.12) gives a unique closed lift $(\theta, z) \in H^0(B_\Phi)$; uniqueness uses $H^0(D_\Phi) = 0$.

The same construction on affine opens glues over an arbitrary T. Closedness says $[\Phi,\theta]=[\Phi,z]=0$, and $\Phi=b\theta+az$. Since a,b generate the unit ideal, these equations imply $[\theta,z]=0$. By (14.13), $z=\lambda\,\mathrm{id}$ over T, not only at its geometric points.

A saturated subbundle is Φ-invariant if and only if it is θ-invariant. The nontrivial implication follows from $\Phi=b\theta+a\lambda\,\mathrm{id}$: multiplication by the nonzero section b kills the image of θ in the torsionfree quotient, hence that image is zero. Thus stability is preserved. Conversely a stable canonical Higgs family and a scalar give a closed lift and hence a family in the gradient zero scheme. These constructions are inverse on objects and morphisms and commute with base change.

The scalar gerbe does not alter the degree-zero critical equations. On a neutral chart its scalar Lie generator has degree -2 and zero differential; the degree -1 generators have differentials $\partial_j g$. Hence degree-zero cohomology is the ordinary quotient by the gradient ideal. At $\chi_*=1$ the line $\det R\pi_*E$ has scalar weight one and neutralizes the gerbe globally. More generally, if $\gcd(r,\chi,(\mu_{p,c})_{p,c})=1$, an integral tensor combination of $\det E|_{p_0}$, $\det R\pi_*E$, and the determinants of the character eigenspaces does the same. Their respective scalar weights are $r,\chi,\mu_{p,c}$. Even without this global choice, the End complexes and all the preceding isomorphisms descend from the local charts. Thus (14.11) holds on the coarse scheme.

The canonical Higgs deformation complex is Serre self-dual. Stability makes its tracefree H^0 and H^2 zero, so the coarse space is smooth symplectic. Riemann–Roch gives its dimension (14.10). Since (14.11) identifies the entire critical scheme, the kernel of the Hessian is its tangent bundle. The induced normal quadratic form is nondegenerate. Its rank is $B_r-d_{\boldsymbol{\mu}}$, which is even: ℓ is even and $r^2-\sum_c\mu_{p,c}^2\equiv r-\sum_c\mu_{p,c}=0$ modulo two. If the canonical moduli is empty, the same functorial construction proves that the critical scheme is empty, so its vanishing-cycle object is zero. □

Lemma 14.6. *In the situation of Proposition 14.5, with j denoting the closed critical inclusion,*

$$\phi_g\mathbf{Q}_{M^L}[B_r+1]\simeq j_*\mathbf{Q}_{M^K\times\mathbf{A}^1}[d_{\boldsymbol{\mu}}+1]\otimes\mathbf{T}^{(B_r-d_{\boldsymbol{\mu}})/2}. \tag{14.14}$$

The auxiliary semisimple monodromy is one and its nilpotent logarithm is zero. The object on the right is understood to be zero when M^K is empty.

Proof. Choose a section of L_0 with reduced divisor D disjoint from the marks and fix bases in the residual lines. The monic resolution gives

$$\det\mathrm{RHom}_{S_L}(\widetilde{E},\widetilde{E})=\det R\Gamma\,\mathrm{End}\,E\otimes(\det R\Gamma(\mathrm{End}\,E\otimes L))^{-1}=\bigotimes_{p\in D}\det(\mathrm{End}(E_p)\otimes L_p)^{-1}.$$

Each factor is canonically trivial in its E_p dependence: $\det(V\otimes V^\vee)=1$, equivalently $\det\mathrm{Ad}(g)=1$ as a polynomial identity over every base. With the chosen residual bases this trivializes K_{M^L}, compatibly with the scalar gerbe. The canonical Higgs symplectic form and $d\lambda$ trivialize $K_{M^K\times\mathbf{A}^1}$. Their ratio is a section δ of the determinant of the conormal bundle. Thus

$$\det\mathrm{Hess}(g)=u\,\delta^{\otimes2}$$

for a regular unit u on the critical scheme.

This unit is constant on every connected component. The canonical Hitchin map is proper over a positive-weight affine base. Scaling (θ,λ) therefore extends the orbit of every point across zero. A regular unit restricted to this copy of $\mathbf{A}^1$ is constant, and its value lies among the finitely many constant values on the connected components of the proper zero fiber. The image of a connected component under u lies in this finite set and hence consists of one value. Choose a constant square root on each component. This trivializes the determinant square-root torsor of the normal quadratic bundle. The relative Morse lemma proves (14.14). Even normal rank gives monodromy $+1$, and a quadratic Morse model has no nilpotent monodromy. Product

compatibility is the derivative determinant compatibility of Proposition 13.3; it does not depend on independent choices of sign for individual packets. □

Let $\mathcal{A}^K_{\boldsymbol{\mu}}$ denote the canonical Hitchin base with the pole and vanishing conditions determined by the profile. On characteristic coefficients the map

$$f_{r,\boldsymbol{\mu}} : \mathcal{A}^K_{\boldsymbol{\mu}} \times \mathbf{A}^1 \longrightarrow \mathcal{A}^L_r, \qquad (\theta, \lambda) \longmapsto b\theta + a\lambda \tag{14.15}$$

is a closed immersion on the canonical spectral image. To see this, evaluate the trace at an ordinary zero of b at which a is nonzero; this recovers λ. Subtracting $a\lambda$ makes the degree-j coefficient belong to the closed linear subspace $b^j H^0(\mathcal{X}, K^j)$, with its profile conditions. Multiplication by b^j has a linear inverse on this image. This argument works over arbitrary base algebras.

Normalize

$$P^L_i = {}^p\mathcal{H}^i Rh_{L*}\mathbf{Q}_{M^L}[B_r + 1], \qquad P^K_i = {}^p\mathcal{H}^i Rh_{K*}\mathbf{Q}_{M^K}[d_{\boldsymbol{\mu}}].$$

Proper vanishing-cycle exchange and (14.14) give

$$\phi_{g_{\mathcal{A}}} P^L_i \simeq (f_{r,\boldsymbol{\mu}})_*(P^K_i \boxtimes \mathbf{Q}_{\mathbf{A}^1}[1]) \otimes \mathbf{T}^{(B_r - d_{\boldsymbol{\mu}})/2}. \tag{14.16}$$

External product with $\mathbf{Q}_{\mathbf{A}^1}[1]$ and the finite map $f_{r,\boldsymbol{\mu}}$ are perverse exact. If M^K is empty this equation says the whole left side is zero.

Theorem 14.7. *Let $Z^{\mathcal{X}}_{\mathrm{nil}}$ be the normalized nil series of the framed physical coefficient, with R recording rank, x recording ordinary Euler characteristic, and X_p recording the ordered root characters. In a target active class one has*

$$[R^r X^{\boldsymbol{\mu}}] \operatorname{Log} Z^{\mathcal{X}}_{\mathrm{nil}} = \frac{x\mathbf{T}^{-N_{\boldsymbol{\mu}}}}{(1-x)(1-x\mathbf{T})} \sum_i (-x)^i [R\Gamma_c(\mathcal{A}^K_{\boldsymbol{\mu}}, P^K_i)]. \tag{14.17}$$

The P^K_i can be taken at the target Euler characteristic (14.8)*. They are zero if that canonical moduli is empty. The identity is an identity of classes of ordinary graded Hodge objects with their finite allocation actions, obtained from perverse-Hodge-module identities.*

Proof. Apply Proposition 14.4 at the simultaneous generic auxiliary value $\chi_* = 1$ of Lemma 14.2. The split/nonsplit comparison and the scalar critical calculation identify its primitive. On restricting the scalar line to zero, the terms are exactly

$$x^{i+1+u+v} P^K_i[-i-2v] \otimes \mathbf{T}^{v-N_{\boldsymbol{\mu}}}, \qquad u, v \geq 0.$$

The scalar shift [1] in (14.16) is retained. Taking compact cohomology and summing u, v gives (14.17); the only cohomological sign is $(-1)^i$. Root-character variables are even. Symmetric powers and Adams operations are the ones on the Hodge category, as in (12.6).

To obtain the asserted target object, use Theorem 14.1 to identify the positive perverse IC objects at $\chi_* = 1$ and at (14.8). Apply the same base vanishing-cycle functor and use (14.16) on both sides. Generic determinant data make semistability stability at the target degree. This also follows numerically: a proper subprofile of parabolic degree zero would give the forbidden proper product of monodromy eigenvalues equal to one. Thus the two canonical primitive objects are identified. If the target is empty, (14.16) makes both primitives zero. This proves the formula without a coprimality assumption on (r, χ). □

15. The global root Fock calculation

15.1. **Fixed objects as finite quotients.** On the physical threefold $Y = \mathrm{Tot}_{\mathcal{X}}(K \oplus \mathcal{O})$ let $H = \mathbf{G}_m$ act by $(\theta, z) \mapsto (u\theta, u^{-1}z)$. For a partition λ, write its boxes as $b = (i_b, j_b)$ with indices beginning at zero, with j_b the column index. Let $a_b, l_b, h_b = a_b + l_b + 1$ be its arm, leg, and hook. The fixed monomial thickening has baseline pushdown

$$E_{0,\lambda} = \bigoplus_{b \in \lambda} K^{-j_b}. \tag{15.1}$$

We use the same framed complex and physical normalization as in (11.3).

Lemma 15.1. *The H-fixed stable-pair functor of shape λ is the H-equivariant finite-quotient functor of $\omega_{Z_\lambda} = \mathrm{Hom}_{\mathcal{X}}(\mathcal{O}_{Z_\lambda}, K)$. The equivalence holds in Artin families with the physical enhancement and its orientation. Its components of fixed Euler characteristic and endpoint profile are projective over their coarse support cycles.*

Proof. A compact one-dimensional support in Y is finite over $\mathcal{X}$. Otherwise a positive-dimensional proper fiber would give, after a finite root-chart pullback, a complete curve in an affine plane, which is impossible. Purity implies that the pushdown is torsionfree, hence locally free on the regular curve chart. Flat families satisfy the same conclusion by the fiberwise locally free criterion.

The section fixes the H-linearization. At the generic point its cyclic annihilator is a monomial partition ideal. This remains true over an Artin base: the anti-diagonal Hilbert tangent weights are the nonzero weights $\pm h_b$, so there is no invariant infinitesimal deformation of that ideal. A nonzero coarse curve function is a nonzerodivisor on the pulled-back free bundle over an Artin algebra. Generic annihilation consequently implies global annihilation. Thus the pair is the extension

$$0 \longrightarrow \mathcal{O}_{Z_\lambda} \longrightarrow F \longrightarrow T \longrightarrow 0 \tag{15.2}$$

with finite H-equivariant T. Curve duality converts this to an H-equivariant finite quotient of ω_{Z_λ}; dualizing back is inverse and commutes with base change. The equivariant monic resolutions and section maps identify the physical complexes as in Proposition 13.3.

The baseline and endpoint characters with $\chi(T)$ determine the nonnegative modified length (13.6). Therefore Lemma 13.1 supplies a projective finite Quot problem. The support map records the cycle of $\nu_*(T \otimes \mathcal{G}^\vee)$ in $\mathrm{Sym}^{N(T)} C$. At an ordinary point its multiplicity is S times the ordinary length; at a root point it is (S/s_p) times the cover length. Fixing this cycle and the endpoint profile determines the possible local representation classes by (13.3). Its fiber is a finite union of products of the precise ordinary and root punctual Quot problems. □

There is a useful local refinement of the last proof. In a root chart the full cover module is killed by the characteristic polynomial of multiplication by t. Equivariance can force coefficients of that polynomial to vanish but cannot remove its monic leading term. Quotients of the finite module obtained by imposing this polynomial form a Grassmannian with closed conditions for multiplication, the partition relations, and the root action. Hence the affine support map is projective. Its support coordinates have strictly positive weights for the second torus R of (11.11), and its Plucker line is linearized. It is the full cover characteristic polynomial which is used here, not the characteristic polynomial on the invariant summand.

15.2. **Local indices, coordinate changes, and purity.** At a root of order s a monomial quotient is specified by an allowed height array $\mathbf{m} = (m_b)_{b\in\lambda}$, with the monotonicity conditions imposed by the two partition-module multiplications. Define

$$\begin{aligned} \gamma_b &= (m_b + j_b) \bmod s, \qquad & r_b &= \left\lfloor \frac{m_b + j_b}{s} \right\rfloor, \\ d_s(\mathbf{m}) &= \sum_b \left(r_b - \left\lfloor \frac{j_b}{s} \right\rfloor \right), & d_0(\lambda, s) &= \sum_b \left(j_b - \left\lfloor \frac{j_b}{s} \right\rfloor \right). \end{aligned} \tag{15.3}$$

Here d_s is the ordinary Euler increment. The full orbifold length, which bounds this array, is

$$\sum_b \frac{m_b}{s} = d_s(\mathbf{m}) + \sum_b \frac{\gamma_b}{s} - \sum_b \frac{j_b \bmod s}{s}.$$

Let $J(\mathbf{m}) = I_H - I_R$ be the finite index difference of the invariant physical tangent complex. The finite character calculation in Proposition 17.1, carried out on this Quot problem, gives

$$d_s(\mathbf{m}) + J(\mathbf{m}) = 2\mathcal{L}(\mathbf{m}) \in 2\mathbf{Z}. \tag{15.4}$$

For reference, in the column formulation of that calculation,

$$J = -d_0 + \sum_b \sigma_b(r_b - j_b) + 2\,\mathrm{dinv}_\lambda(\boldsymbol{\gamma}), \qquad \mathcal{L} = \mathrm{dinv}_\lambda(\boldsymbol{\gamma}) + \sum_{\sigma_b=1} (r_b - j_b).$$

Proposition 15.2. *The normalized punctual root coefficient is ordinary and degreewise pure. Its cohomology class is*

$$V_{\lambda,s}(x, \mathbf{T}; X) = \sum_{\mathbf{m}} x^{d_s(\mathbf{m})} \mathbf{T}^{\mathcal{L}(\mathbf{m})} \prod_{b \in \lambda} X_{\gamma_b}. \tag{15.5}$$

The corresponding derived summand is $\mathbf{T}^{\mathcal{L}(\mathbf{m})}[-2\mathcal{L}(\mathbf{m})]$. *The local variations and their orientation maps are independent of root coordinates.*

Proof. The common H, R tangent space is the root-invariant part of the ordinary tangent $\mathrm{Hom}(K_0, T)$, where K_0 and T are the kernel and quotient. The column propagation in Lemma 11.3 proves that this latter tangent has no common weight-zero vector, for every allowed height array. The residual torus has a limit in the projective finite-module fiber, so these exhaust the fixed locus. Tracefree Ext^0 vanishes by the framing, and threefold duality then annihilates the fixed Ext^2 and Ext^3 as well. The fixed derived tangent is acyclic and its normalized oriented coefficient is the unit.

The positive R-weights on the affine support base, its projectivity, and its linearized Plucker line give the closed attracting filtration of Lemma 11.4. Compact oriented localization assigns one half-Tate object to each layer; its cohomological degree equals its intrinsic weight. Connecting maps between adjacent layers join different pure weights and vanish. The full compact cohomology is therefore degreewise pure. Verdier self-duality of the untwisted coefficient and contraction of the positive support base identify ordinary cohomology of the punctual fiber with the dual compact cohomology. This applies to the whole punctual fiber and its ambient coefficient.

Its normalized summand is $\vartheta^{d_s+J}[-(d_s+J)]$. Equation (15.4) turns it into the ordinary object stated in the proposition, with trivial auxiliary monodromy. Its integer Tate exponent can be negative; neither purity nor the bound (13.6) requires it to be nonnegative. Taking classes gives (15.5).

A root-coordinate change has the form $f(t) = at + \sum_{j\geq 1} c_j t^{1+js}$. On every required finite jet it has an algebraic deformation

$$f_v(t) = v^{-1} f(vt) = at + \sum_{j\geq 1} c_j v^{js} t^{1+js}$$

to its linear part. It lifts to the volume-preserving transformation $(t, y, z) \mapsto (f(t), y/f'(t), z)$. By Proposition 13.3 this is a family of oriented complex maps. The induced maps of proper punctual cohomology are maps between constant variations over the parameter line, so the endpoint maps agree. The linear part belongs to a connected volume-preserving torus and acts trivially on cohomology, with the specified orientation action. The root deck rotation belongs to the same torus. The local-system monodromy around a root is therefore trivial. Only ordinary movable packets are permuted; root packets sit at fixed marked points. Their possible cross-profile signs are already identified in Proposition 13.3. □

Let G_λ denote the normalized punctual series of Definition 17.2. The variables are related by

$$q = x^{-1}, \qquad t = x\mathbf{T}, \qquad V_{\lambda,s}(x, \mathbf{T}; X) = x^{d_0(\lambda,s)} G_\lambda[X; q, t]. \tag{15.6}$$

One first retains the finite root alphabet and its root order in this notation. The explicit local expression proves independence of the choice of sufficiently separated rational character labels and of an enlarged denominator. We use this stability below when passing through finite degree bounds.

Proposition 15.3. *For each shape, endpoint profile, and Euler characteristic, the complete normalized fixed component has ordinary degreewise pure compact cohomology. The same is true for the complete nil coefficient.*

Proof. First fix a locally constant H-index. Push the untwisted self-dual physical coefficient along the proper support map of Lemma 15.1. Its fibers are finite unions of products of the exact punctual problems. Proper base change, Proposition 15.2, and the disjoint-support factorization show that every stalk H^j is pure of intrinsic weight j. The stalk criterion gives complex weights at most zero. Proper self-duality gives weights at least zero. Hence this direct image is pure of complex weight zero on the proper symmetric power of C. This reasoning is made before stratifying the support into the punctured curve and the root points. The coordinate argument in Proposition 15.2 identifies the movable punctual variations, so it is valid over all support strata.

We record also why the entire physical spectral map is proper, which is needed to include complete attracting components. A compact irreducible curve in $\mathrm{Tot}_{\mathcal{X}}(K \oplus \mathcal{O})$ has constant z and projects to a finite root spectral divisor in $\mathrm{Tot}_{\mathcal{X}} K$. The closed immersion (14.15) and finite spectral addition give finite parameter covers of all its compact cycle types. Compactify in $\mathbf{P}_{\mathcal{X}}(\mathcal{O} \oplus K \oplus \mathcal{O})$. A coarse polarization and the pullback of (13.5) give a fixed modified Hilbert polynomial for fixed rank, ordinary Euler characteristic, and root profiles. The projective stable-pair moduli with this polynomial exists: the projective Deligne–Mumford stable-pair theorem asserts that, for a generating bundle, coarse ample polarization, fixed source $\mathcal{O}$, and stability parameter of degree at least one, the pure one-dimensional pairs with zero-dimensional cokernel form a projective fine moduli space [40, Lemmas 2.24 and 2.26, Theorems 4.24 and 5.20]. These are precisely the conditions here.

For a DVR family of positive spectral coefficients, lift its compact cycle after a finite base extension to one of the finite cycle covers. The lift extends and gives a limiting cycle entirely in the affine chart. A projective stable-pair limit exists. Apply separatedness of coarse projective Chow to $\nu_*(F \otimes \mathcal{G}^\vee)$. This sheaf is flat by tameness, and its support is exactly the coarse support of F because $\mathcal{G}$ is generating. Every generic horizontal curve has positive multiplicity, namely S times the original multiplicity. Its limiting cycle must be the explicit affine cycle. A curve at infinity would contribute positively and is therefore impossible; purity excludes an isolated zero-dimensional support there. The limit stays in the affine chart. This proves the required properness. Using merely $\nu_* F$ would not suffice, because it can annihilate a nontrivial root character. For the split row negativity of A gives the same argument.

On the characteristic-nil locus, simultaneous triangularization on reductions gives $\mathrm{char}(az + b\theta) = \mathrm{char}(b\theta)$. The H-weights on this base are positive. The projective linearized polarization gives a finite closed filtration by entire attractors over complete projective fixed components, exactly as in Lemma 11.4. A term $u^{-j} c_j$ extends at $u = 0$ only when $c_j = 0$, so these are precisely the reduced ambient attractors. Applying the first part of the proof to each complete component and the parity (15.4) makes all normalized layers ordinary and pure of weight zero. Their boundary maps vanish by weights, and same-weight pure Hodge objects are semisimple. This proves the assertion for the whole nil coefficient. □

15.3. **Baseline and the global series.** The pushdown formula is

$$\nu_* K^{-j} = K_C^{-j} \otimes \mathcal{O}_C \left(- \sum_p \left(j - \left\lfloor \frac{j}{s_p} \right\rfloor \right) p \right). \tag{15.7}$$

Writing $h = g - 1$, ordinary Riemann–Roch gives

$$\chi_0^{\mathcal{X}}(\lambda) = -h\|\lambda\|^2 - \sum_p d_0(\lambda, s_p), \qquad I_0(\lambda) = -\chi_0^{\mathcal{X}}(\lambda). \tag{15.8}$$

For the second equality compute $I_0 = \sum_b \chi(K^{a_b+1})$. Locally its correction is

$$\left\lfloor \frac{(a+1)(s-1)}{s} \right\rfloor = a - \left\lfloor \frac{a}{s} \right\rfloor .$$

In each row the multiset of arm lengths equals that of column indices. Together with $\sum_b(2a_b+1) = \|\lambda\|^2$ this proves both identities in (15.8).

For a curve or an open curve V let

$$\zeta_V(v) = \sum_{d\geq 0} [R\Gamma_c(\mathrm{Sym}^d V, \mathbf{Q})] v^d$$

be its derived compact-cohomology zeta function in the Hodge lambda ring. If $U = C \setminus \{p_1, \dots, p_m\}$, then

$$\zeta_U(v) = \zeta_C(v)(1-v)^m. \tag{15.9}$$

This is the symmetric-power identity for $[R\Gamma_c(U)] = [R\Gamma(C)] - m$; it retains cohomological signs and all Adams operations.

Theorem 15.4. *The normalized nil series is*

$$Z^{\mathcal{X}}_{\mathrm{nil}} = \sum_{\lambda} R^{|\lambda|} x^{\chi_0^{\mathcal{X}}(\lambda)} \prod_{b\in\lambda} \zeta_U(x^{h_b}\mathbf{T}^{l_b}) \prod_{p=1}^{m} V_{\lambda,s_p}(x, \mathbf{T}; X_p). \tag{15.10}$$

Equivalently, with the punctual series G_λ,

$$Z^{\mathcal{X}}_{\mathrm{nil}} = \sum_{\lambda} R^{|\lambda|} x^{-h\|\lambda\|^2} \prod_{b\in\lambda} \zeta_U(x^{h_b}\mathbf{T}^{l_b}) \prod_{p=1}^{m} G_\lambda[X_p; x^{-1}, x\mathbf{T}]. \tag{15.11}$$

Both identities are coefficientwise identities of ordinary graded Hodge classes, with their specified finite actions. Every coefficient has a finite-type moduli interpretation.

Proof. The fixed functor is the Quot functor of Lemma 15.1. Split its support into ordinary movable points in U and the fixed root points. Disjoint support idempotents lift uniquely over nilpotent bases, all cross-Hom complexes vanish, and the trace splits with the same idempotents. Thus its physical coefficient factors by the supported maps of Proposition 13.3. On U the calculation is Theorem 11.1, including its ordinary symmetric-group projectors, and gives $\prod_b \zeta_U(x^{h_b}\mathbf{T}^{l_b})$. At p it gives (15.5). Equation (15.8) supplies the baseline exponent. This proves (15.10). The complete fixed components are pure by Proposition 15.3, proved before this support stratification. Substituting (15.6) cancels exactly the root contribution to the baseline and gives (15.11).

All operations are performed coefficientwise in rank, endpoint profiles, and ordinary Euler characteristic. The rank bounds the number of partitions, and (13.6) fixes a nonnegative modified length for each coefficient. Consequently the sums and symmetric operations involve finite-type projective support problems. This also specifies the completion in which the plethystic logarithm in Theorem 14.7 is taken. □

Corollary 15.5. *After the local identification of Theorem 23.1, the series is*

$$Z^{\mathcal{X}}_{\mathrm{nil}} = \sum_{\lambda} R^{|\lambda|} x^{-(g-1)\|\lambda\|^2} \prod_{b\in\lambda} \zeta_C(x^{h_b}\mathbf{T}^{l_b}) \prod_{p=1}^{m} \widetilde{H}_\lambda[X_p; x^{-1}, x\mathbf{T}]. \tag{15.12}$$

Proof. That theorem gives, with $D_\lambda = \prod_b(1 - x^{h_b}\mathbf{T}^{l_b})$, $G_\lambda = D_\lambda^{-1}\widetilde{H}_\lambda$. Each puncture contributes one inverse D_λ, which cancels its factor in (15.9). Equation (15.11) now gives the claim. □

16. The football calculation and the Cauchy identity

The local recognition will use a pairing identity. We obtain it from a genus-zero geometry in which the canonical primitive can be calculated directly. This section uses the series G_λ of Definition 17.2, before its identification with a Macdonald function.

16.1. **Vanishing of higher-rank canonical primitives.** Take the football

$$\mathcal{F}_s = \mathbf{P}^1[\sqrt[s]{0}, \sqrt[s]{\infty}] \simeq [\mathbf{P}^1/\mu_s].$$

Its canonical degree is $-2/s$. All positive comparison and moving-Hecke constructions use $L = \nu^*\mathcal{O}_{\mathbf{P}^1}(\ell)$ with ℓ sufficiently large and even. The noncoarse canonical line is introduced only through (13.8) and the trace potential.

Lemma 16.1. *A stable canonical Higgs bundle on $\mathcal{F}_s$ has rank one and zero Higgs field. Consequently, for generic active weights at $\chi_* = 1$, every positive primitive of rank greater than one has zero vanishing-cycle image.*

Proof. If the underlying bundle of a semistable canonical Higgs bundle is unstable, its highest Harder–Narasimhan subbundle is Higgs invariant: its slope is strictly larger than every slope of the quotient tensored with the negative line K. It would destabilize the Higgs bundle. Thus the underlying bundle is semistable, and negativity gives $\mathrm{Hom}(E, E \otimes K) = 0$. Its field is zero.

Pull the bundle back to $\mathbf{P}^1$. The Birkhoff–Grothendieck splitting and its Harder–Narasimhan filtration are preserved by the finite group. Split successive extension retractions and average them over μ_s; hence the decomposition by line degree is equivariant. An equal-degree piece is $\mathcal{O}(d) \otimes V$, where the finite-dimensional μ_s representation V splits into characters. It is therefore a direct sum of equivariant lines of the same root slope. A semistable bundle of rank greater than one has a proper summand of equal slope and cannot be stable. This proves the first assertion.

In any fixed total rank bound choose rational weights generic for the finite active set at $\chi_* = 1$, using the wall avoidance of Lemma 14.2, without a prescribed determinant hyperplane. Clear their denominators to obtain s. Semistability is then stability in the required classes. The canonical moduli in rank greater than one is empty. Proposition 14.5 identifies the entire positive critical scheme with its product with $\mathbf{A}^1$, so the critical scheme is empty as well. Its vanishing-cycle object is therefore zero. □

Lemma 16.2. *For each rank-one endpoint profile and fixed Euler characteristic the canonical moduli on $\mathcal{F}_s$ is a point. Its normalized primitive is*

$$Rp_1[X]p_1[Y]\frac{x}{(1-x)(1-x\mathbf{T})}. \tag{16.1}$$

Proof. The endpoint characters and ordinary Euler characteristic determine one root line bundle on $\mathbf{P}^1$. Its canonical field is zero because $H^0(K) = 0$. The positive moduli is the affine space $H^0(\mathcal{F}_s, L)$ of dimension $\ell + 1$. The function (14.9) on it is a symmetric quadratic form. The connecting map of (13.8) is

$$H^0(L) \longrightarrow H^1(A) = H^0(L)^\vee.$$

Since $H^0(K) = 0$ and $H^0(\mathcal{O}) = \mathbf{C}$, its kernel is exactly $\mathbf{C}a$. The critical line is $\mathbf{A}^1$ and the nondegenerate normal rank is ℓ. It is even, so its vanishing cycles are the ordinary $\mathbf{T}^{\ell/2}$-twist, with semisimple monodromy one and nilpotent monodromy zero. The determinant root is a fixed constant root of this quadratic space and is independent of both endpoint characters.

Here $d_{\boldsymbol{\mu}} = 0$. Theorem 14.7 and the two free Euler directions therefore give $x/((1-x)(1-x\mathbf{T}))$ for each pair of characters. Summing these pairs gives $p_1[X]p_1[Y]$, proving (16.1). □

It follows that the full football nil series is

$$Z^{\mathcal{F}_s}_{\mathrm{nil}} = \mathrm{Exp}\left[Rp_1[X]p_1[Y]\frac{x}{(1-x)(1-x\mathbf{T})}\right]. \tag{16.2}$$

This calculation may be made in any finite rank and character degree. Rational avoidance uses only finitely many walls, and the denominator s is chosen after that avoidance. Independence of the normalized local expression from an enlarged denominator and from the separation of the character labels, as proved in the local calculation, then identifies these finite calculations.

Taking all degree bounds proves the symmetric-function identity for arbitrary finite alphabets and then in the completed symmetric-function ring.

16.2. **The exact kernel and its signs.** On the football, $U = \mathbf{G}_m$ and

$$\zeta_{\mathbf{G}_m}(v) = \frac{1-v}{1-\mathbf{T}v}.$$

Since $g - 1 = -1$, the geometric Fock identity (15.11) gives

$$Z_{\mathrm{nil}}^{\mathcal{F}_s} = \sum_{\lambda} R^{|\lambda|} x^{\|\lambda\|^2} \prod_{b\in\lambda} \frac{1-x^{h_b}\mathbf{T}^{l_b}}{1-x^{h_b}\mathbf{T}^{l_b+1}} G_\lambda[X; x^{-1}, x\mathbf{T}] G_\lambda[Y; x^{-1}, x\mathbf{T}]. \tag{16.3}$$

Thus both sides of the required pairing identity have been calculated from the same physical coefficient.

Theorem 16.3. *Put $q = x^{-1}$, $t = x\mathbf{T}$, and*

$$D_\lambda(q,t) = \prod_{b\in\lambda}(1-q^{-a_b-1}t^{l_b}), \qquad F_\lambda[X;q,t] = D_\lambda(q,t)G_\lambda[X;q,t].$$

For the normalized punctual series,

$$\sum_{\lambda} R^{|\lambda|} \frac{F_\lambda[X;q,t]F_\lambda[Y;q,t]}{\prod_{b\in\lambda}(q^{a_b+1}-t^{l_b})(q^{a_b}-t^{l_b+1})} = \mathrm{Exp}\left[\frac{Rp_1[X]p_1[Y]}{(q-1)(1-t)}\right]. \tag{16.4}$$

Proof. Equate (16.2) and (16.3), and substitute $F_\lambda = D_\lambda G_\lambda$. The coefficient of $F_\lambda[X]F_\lambda[Y]$ on the resulting left side is

$$c_\lambda = \frac{x^{\|\lambda\|^2}}{\prod_{b\in\lambda}(1-x^{h_b}\mathbf{T}^{l_b})(1-x^{h_b}\mathbf{T}^{l_b+1})}. \tag{16.5}$$

For each box,

$$q^{a+1} - t^l = x^{-a-1}(1-x^{a+l+1}\mathbf{T}^l),$$
$$q^a - t^{l+1} = x^{-a}(1-x^{a+l+1}\mathbf{T}^{l+1}).$$

The identity $\sum_b(2a_b+1) = \|\lambda\|^2$ therefore gives

$$c_\lambda^{-1} = \prod_b (q^{a_b+1}-t^{l_b})(q^{a_b}-t^{l_b+1}).$$

Also $x/((1-x)(1-x\mathbf{T})) = 1/((q-1)(1-t))$. This proves (16.4) with its powers and signs.

For clarity, the power-sum form of its pairing is

$$\langle p_\rho, p_\tau\rangle_{\mathcal{F}} = \delta_{\rho\tau} z_\rho \prod_i (q^{\rho_i}-1)(1-t^{\rho_i}). \tag{16.6}$$

Expanding the plethystic exponential indeed gives the reciprocal of these diagonal entries. If

$$\widetilde{w}_\lambda = \prod_b (q^{a_b}-t^{l_b+1})(t^{l_b}-q^{a_b+1}),$$

then $c_\lambda^{-1} = (-1)^{|\lambda|}\widetilde{w}_\lambda$. On homogeneous degree n, (16.6) is $(-1)^n$ times the usual modified star pairing. Accordingly the standard star Cauchy identity is converted to this one by the ordinary replacement $R \mapsto -R$ in homogeneous degree. This is ordinary scalar negation of the rank variable, not additive plethystic negation of an alphabet. These observations fix the pairing convention used in the local triangular recognition. □

17. The punctual root vertex

We compute the punctual contribution before identifying it with a symmetric function. All Young diagrams in this section use zero-based coordinates: a box is $b = (i, j)$, its content is $c_b = j - i$, and λ_i is the length of row i. Its arm and leg are denoted by a_b and l_b. A reverse plane partition is a nonnegative integer array weakly increasing to the right and down. Write $\mathrm{RPP}(\lambda)$ for the set of such arrays.

17.1. **A finite character and its root-invariant index.** Let $s \geq 1$ be the order of the root stabilizer. The three geometric tangent characters and their two-torus restrictions are

$$(t_1, t_2, t_3) = (u^{-1}, uv^{-1}, v).$$

Thus $t_1^{d_1} t_2^{d_2} t_3^{d_3}$ has weight $(d_2 - d_1, d_3 - d_2)$, and root invariants are precisely the terms whose second coordinate is divisible by s. For $m \in \mathrm{RPP}(\lambda)$, put

$$\mathsf{Z}_\lambda = \sum_{(i,j)\in\lambda} t_1^{-i} t_2^{-j}, \qquad \mathsf{Z}_m = \sum_{(i,j)\in\lambda} \sum_{\alpha=1}^{m_{ij}} t_1^{-i} t_2^{-j} t_3^{\alpha}, \qquad \kappa = t_1 t_2 t_3.$$

The bar inverts characters. The Koszul resolution gives $\chi(A, B) = \overline{A}B(1 - t_1)(1 - t_2)(1 - t_3)$ on finite-support differences. For the stable-pair complex $I = \mathcal{O} - F$, removal of the trace and Serre duality give

$$V(F) = F - \kappa\overline{F} - (1 - t_1)(1 - t_2)(1 - t_3)F\overline{F}.$$

Subtract the character of $F_0 = \mathsf{Z}_\lambda/(1 - t_3^{-1})$ from that of $F = F_0 + \mathsf{Z}_m$. Since $(1 - t_3)F_0 = -t_3\mathsf{Z}_\lambda$, this subtraction is finite and equals

$$\begin{aligned} v_m = {} & \mathsf{Z}_m - \kappa\overline{\mathsf{Z}_m} \\ & - (1 - t_1)(1 - t_2)\big(-t_3\mathsf{Z}_\lambda\overline{\mathsf{Z}_m} + \overline{\mathsf{Z}_\lambda}\mathsf{Z}_m + (1 - t_3)\mathsf{Z}_m\overline{\mathsf{Z}_m}\big). \end{aligned} \tag{17.1}$$

In particular, the invariant operation below is applied to a finite Laurent character; it does not require a choice of expansion of F_0.

If $n_{D,e}$ is the coefficient of $u^D v^e$ in (17.1), define the relative indices

$$\begin{gathered} I_{\mathrm{lex}}(m) = \frac{1}{2} \sum_{D,e\,:\,,s|e} n_{D,e} \operatorname{sgn}_{\mathrm{lex}}(D, e), \qquad I_H(m) = \frac{1}{2} \sum_{D,e\,:\,,s|e} n_{D,e} \operatorname{sgn}(D), \\ J(m) = 2I_H(m) - I_{\mathrm{lex}}(m). \end{gathered} \tag{17.2}$$

The lexicographic sign is the sign of the first nonzero coordinate, and is zero at the origin. These indices are measured from $m = 0$. The root character of the generator at $b = (i, j)$ is the residue of $m_b + j$. Accordingly set

$$\begin{gathered} \ell_b = m_b + j_b = sr_b + \gamma_b, \quad 0 \leq \gamma_b < s, \\ d_s(m) = \sum_b \left(r_b - \left\lfloor \frac{j_b}{s} \right\rfloor \right), \quad d_0(\lambda, s) = \sum_b \left(j_b - \left\lfloor \frac{j_b}{s} \right\rfloor \right). \end{gathered} \tag{17.3}$$

The Maya set and its signs are

$$\mathcal{B}_\lambda = \{\lambda_i - i - 1 : i \geq 0\}, \qquad \sigma(c) = 2\mathbf{1}_{c\in\mathcal{B}_\lambda} - 1, \qquad \sigma_b = \sigma(c_b - 1). \tag{17.4}$$

Zero rows are included in this definition. Let S_λ consist of ordered pairs (a, b) in which b is strictly southeast of a and $c_b = c_a$. Let A_λ consist of pairs in which b is strictly below and weakly to the right of a, and $c_b = c_a - 1$. Define

$$\operatorname{dinv}_\lambda(\gamma) = \sum_{(a,b)\in S_\lambda} \mathbf{1}_{\gamma_b > \gamma_a} + \sum_{(a,b)\in A_\lambda} \mathbf{1}_{\gamma_b < \gamma_a}. \tag{17.5}$$

Proposition 17.1 (Closed index)**.** *For every partition, every root order, and every legal height array,*

$$J(m) = -d_0(\lambda, s) + \sum_{b\in\lambda} \sigma_b(r_b - j_b) + 2\,\mathrm{dinv}_\lambda(\gamma). \tag{17.6}$$

Proof. Write $\epsilon(e) = \mathbf{1}_{s|e}$. The finite-difference kernel obtained by applying the first expression in (17.2) to (17.1) is

$$\begin{aligned} A_s(D,e) = \epsilon(e)\big(&\mathrm{sgn}_{\mathrm{lex}}(D,e) - \mathrm{sgn}_{\mathrm{lex}}(D-1,e)\big) \\ &+ \epsilon(e-1)\big(-\mathrm{sgn}_{\mathrm{lex}}(D+1,e-1) + \mathrm{sgn}_{\mathrm{lex}}(D,e-1)\big). \end{aligned}$$

It is zero outside $D = -1, 0, 1$. At these three values it is, respectively,

$$\begin{aligned} A_s(-1,e) &= -\epsilon(e-1)(1+\mathrm{sgn}(e-1)), \\ A_s(0,e) &= \epsilon(e)(1+\mathrm{sgn}(e)) + \epsilon(e-1)(\mathrm{sgn}(e-1)-1), \\ A_s(1,e) &= \epsilon(e)(1-\mathrm{sgn}(e)). \end{aligned} \tag{17.7}$$

The H-kernel A_s^H is obtained by replacing every $\mathrm{sgn}_{\mathrm{lex}}(D,e)$ by $\mathrm{sgn}(D)$; its three values are $-\epsilon(e-1)$, $\epsilon(e)-\epsilon(e-1)$, and $\epsilon(e)$.

Here is the finite-character calculation leading to this kernel. Raise the height at $a = (i,j)$ by one, and let $\alpha = m_a + 1$. Set $z = t_1^{-i}t_2^{-j}t_3^{\alpha}$. Replacing Z_m by $\mathsf{Z}_m + z$ in (17.1) gives

$$\begin{aligned} \Delta v = z - \kappa\overline{z} \\ - (1-t_1)(1-t_2)\big[-t_3\mathsf{Z}_\lambda\overline{z} + \overline{\mathsf{Z}_\lambda}z + (1-t_3)(z\overline{\mathsf{Z}_m} + \mathsf{Z}_m\overline{z} + 1)\big]. \end{aligned}$$

For another box (l,k) put $D = k - l - (j-i)$ and $e = \alpha + j - k$. Its baseline contributes $-A_s(D,e)$; its occupied levels contribute the telescoping sum

$$\sum_{\beta=1}^{m_{lk}} \big(-A_s(D, e-\beta) + A_s(D, e-\beta+1)\big) = A_s(D,e) - A_s(D, e-m_{lk}).$$

The source contributes $\epsilon(\alpha+j)\,\mathrm{sgn}_{\mathrm{lex}}(-c_a, \alpha+j)$. The constant $-P = -(1-t_1)(1-t_2)(1-t_3)$ contributes 1 if $s = 1$, and -1 otherwise. The corresponding H-constants are 0 and -1. The moving self term always has $(D,e) = (0,1)$; its kernel is exactly the relevant constant, so the two cancel before summing the other boxes. This also treats $s = 1$.

Construct m from zero in reverse row-major order. During the construction of the height at (i,j), every southeast box has its final height and every northwest box still has height zero. Every box on one of the three contributing diagonals is comparable to (i,j). Substitution in $2A_s^H - A_s$ gives the following contributions to ΔJ, in addition to the source $f(c_a)\epsilon(\alpha+j)$, where $f(c) = 1$ for $c < 0$ and $f(c) = -1$ for $c \geq 0$:

D	northwest box (l,k)	southeast box (l,k)
0	$2\epsilon(\alpha+j-k-1)$	$-2\epsilon(\alpha+j-k-m_{lk})$
1	$-2\epsilon(\alpha+j-k)$	0
-1	0	$2\epsilon(\alpha+j-k-m_{lk}-1)$.

The endpoint signs in this table follow directly from legality. For $D = 0$, the northwest and southeast kernel arguments are at least 2 and at most -1, respectively. For $D = 1$ they are at least 1 and at most -1, and for $D = -1$ at least 2 and at most 0. Thus the possible zero argument occurs precisely where the accompanying sign factor vanishes.

Using $\sum_{\alpha=1}^{M}\epsilon(\alpha+a)=\lfloor (M+a)/s\rfloor-\lfloor a/s\rfloor$ and combining the two appearances of each pair yields

$$\begin{aligned} J(m)=&\sum_b f(c_b)\left(\left\lfloor\frac{\ell_b}{s}\right\rfloor-\left\lfloor\frac{j_b}{s}\right\rfloor\right)\\ &+2\sum_{(a,b)\in S_\lambda}\left(\left\lfloor\frac{\ell_b-\ell_a-1}{s}\right\rfloor-\left\lfloor\frac{j_b-j_a-1}{s}\right\rfloor\right)\\ &-2\sum_{(a,b)\in A_\lambda}\left(\left\lfloor\frac{\ell_b-\ell_a}{s}\right\rfloor-\left\lfloor\frac{j_b-j_a}{s}\right\rfloor\right). \end{aligned} \tag{17.8}$$

We identify both its linear part and its constant. Give the attacking graph the directed edges $b\to a$ for $(a,b)\in S_\lambda$ and $a\to b$ for $(a,b)\in A_\lambda$. For a box $b=(i,j)$,

$$\sigma_b=f(c_b)+2\big(\mathrm{out}(b)-\mathrm{in}(b)\big). \tag{17.9}$$

Indeed, northwest same-diagonal boxes number $\min(i,j)$, and northwest boxes of content c_b+1 number $\min(i,j+1)$. Their difference is $-\mathbf{1}_{i>j}$. The difference between the numbers of southeast boxes of content c_b-1 and c_b is

$$\#\{h\ge 1:\lambda_{i+h}=j+h\}=\mathbf{1}_{c_b-1\in\mathcal{B}_\lambda}.$$

There is at most one such h, since row lengths decrease while $j+h$ increases; a Maya equality with row index at most i is impossible because (i,j) lies in the diagram. These two differences prove (17.9). Hence the coefficient of r_b in (17.8) is exactly σ_b.

For the constant use the legal array $m_b^0=(s-1)j_b$, for which $r_b=j_b$ and all γ_b vanish. Put $A_s(k)=k-\lfloor k/s\rfloor$. Substitution in (17.8) gives

$$\frac{J(m^0)+d_0}{2}=\sum_{c_b<0}A_s(j_b)+\sum_{(a,b)\in S_\lambda}A_s(j_b-j_a-1)-\sum_{(a,b)\in A_\lambda}A_s(j_b-j_a). \tag{17.10}$$

In the last sum, move the northwest box one step left whenever its column is positive. This is a bijection onto S_λ, and preserves the displayed argument of A_s. Its inverse is legal because the southeast box forces the requisite box in the longer northwest row. The remaining pairs have northwest column zero and correspond bijectively to boxes of negative content, with argument j_b. Thus (17.10) vanishes. Finally,

$$\left\lfloor\frac{\ell_b-\ell_a-1}{s}\right\rfloor=r_b-r_a-\mathbf{1}_{\gamma_b\le\gamma_a},\qquad \left\lfloor\frac{\ell_b-\ell_a}{s}\right\rfloor=r_b-r_a-\mathbf{1}_{\gamma_b<\gamma_a}.$$

Subtracting the baseline in (17.8) proves (17.6). □

17.2. **Stable alphabets and the normalized series.** The inequalities on m have a useful equivalent form. The integers $\ell_b=m_b+j_b$ increase strictly in rows and weakly in columns. Consequently r is a reverse plane partition; within an equal-r level, γ increases strictly in rows and weakly in columns. Conversely these conditions imply the inequalities on ℓ and hence on $m=\ell-j$. Nonnegativity follows from the first column and the successive strict row increments. Thus this description depends only on the order of the alphabet, not on its cardinality.

Let q,t be algebraically independent. For an ordered alphabet $X_s=x_0+\cdots+x_{s-1}$ set $X_s^\gamma=\prod_b x_{\gamma_b}$.

Definition 17.2 (Punctual vertex)**.** The normalized root vertex is the stable homogeneous series whose value on X_s is

$$G_\lambda[X_s;q,t]=\sum_{\substack{r\in\mathrm{RPP}(\lambda),\ 0\le\gamma_b<s\\ \gamma\text{ row-strict and column-weak on equal-}r\text{ levels}}} X_s^\gamma(qt)^{\mathrm{dinv}_\lambda(\gamma)}\prod_{\sigma_b=-1}q^{j_b-r_b}\prod_{\sigma_b=1}t^{r_b-j_b}. \tag{17.11}$$

We use the scalars and normalized functions

$$D_\lambda(q,t)=\prod_{b\in\lambda}(1-q^{-a_b-1}t^{l_b}),\qquad F_\lambda[X;q,t]=D_\lambda(q,t)G_\lambda[X;q,t]. \tag{17.12}$$

For the empty diagram every expression is 1.

Here v,Θ denote independent bookkeeping variables for the Euler increment and the index, respectively. Set $q=\Theta/v$ and $t=v\Theta$. By Proposition 17.1, a summand satisfies

$$v^{d_s-d_0}\Theta^{J+d_0}=(qt)^{\mathrm{dinv}_\lambda(\gamma)}\prod_{\sigma_b=-1}q^{j_b-r_b}\prod_{\sigma_b=1}t^{r_b-j_b}. \tag{17.13}$$

In the geometric variables $v=x\Theta$, $\mathbf{T}=\Theta^2$, the punctual exponent is $\mathcal{L}=(d_s+J)/2$. Therefore

$$V_{\lambda,s}(x,\mathbf{T};X_s)=x^{d_0(\lambda,s)}G_\lambda[X_s;x^{-1},x\mathbf{T}],\qquad D_\lambda(x^{-1},x\mathbf{T})=\prod_b(1-x^{a_b+l_b+1}\mathbf{T}^{l_b}). \tag{17.14}$$

Lemma 17.3 (Rationality and stability)**.** *For each finite alphabet, the coefficients of G_λ belong to $\mathbf{Q}(q,t)$. Specialization of an alphabet variable to zero deletes that letter. The rational expressions are compatible with the height completion in* (17.13)*.*

Proof. Fix a color assignment and divide its legal height arrays into their finitely many total weak-order types. On one type the distinct levels are $h_1<\cdots<h_m$, where $h_1=u_1\ge 0$ and $h_i-h_{i-1}=1+u_i$ with $u_i\ge 0$. Equalities or strictness forced by a neighboring color either exclude the type or are automatic on it. Increasing u_i translates a nonempty union U_i of its last height blocks. The weight ratio is

$$q^{-a_i}t^{b_i},\qquad a_i=\#\{b\in U_i:\sigma_b=-1\},\quad b_i=\#\{b\in U_i:\sigma_b=1\}.$$

Since $a_i+b_i>0$, this is not 1. The sum on the type is a monomial divided by $\prod_i(1-q^{-a_i}t^{b_i})$. There are finitely many types and color assignments. Every increase of u_i raises the original height degree by $|U_i|$, so the same geometric sums converge in that completion. Deletion of a letter retains the order of the others and all defining comparisons, proving stability. □

17.3. **The LLT expansion and the one-letter value.** For $r\in\mathrm{RPP}(\lambda)$ put $\lambda_{\le h}=\{b:r_b\le h\}$ and $\rho_h=\lambda_{\le h}/\lambda_{<h}$. These are skew diagrams. Transpose their nonempty components, retain their coordinates, and list them in decreasing order of h; call the resulting tuple $\nu(r)$. Let $\mathcal{G}_\nu[X;T]$ denote the inversion LLT polynomial of this tuple. Specifically, for components $i<j$ an equal-content pair is inverted if the entry in component i is larger; for $i>j$ a pair with the first content one larger is inverted under the same comparison. The content of a transposed box is measured so that it equals the original $j-i$. This specifies all attacks and agrees with the usual inversion convention.

Proposition 17.4. *With*

$$W_\lambda(r;q,t)=\prod_{\sigma_b=-1}q^{j_b-r_b}\prod_{\sigma_b=1}t^{r_b-j_b},\qquad \kappa(r)=\#\{(a,b)\in S_\lambda:r_a=r_b\},$$

one has

$$G_\lambda[X;q,t]=\sum_{r\in\mathrm{RPP}(\lambda)}W_\lambda(r;q,t)(qt)^{\kappa(r)}\mathcal{G}_{\nu(r)}[X;qt]. \tag{17.15}$$

In particular $G_\lambda\in\mathbf{Q}(q,t)\otimes\Lambda_{|\lambda|}$. Furthermore,

$$G_\lambda[1;q,t]=D_\lambda(q,t)^{-1},\qquad F_\lambda[1;q,t]=1. \tag{17.16}$$

Proof. Transposition changes the levelwise row-strict, column-weak condition into the ordinary semistandard condition. On different levels the southeast member of a same-content pair lies in the earlier component, and gives exactly the first term in (17.5). For an adjacent-content pair the northwest member lies in the later component and has content one larger; its inversion is the second term. All possible attacks on these diagonals are comparable, so this list is exhaustive.

On the same level a same-content southeast box has a strictly larger entry: an order-convex path inside that level contains a horizontal step. Every such pair contributes one. A same-level pair in A_λ contributes zero. Hence $\mathrm{dinv} = \mathrm{inv}_{\mathrm{LLT}} + \kappa(r)$ and (17.15) follows.

The symmetry theorem for inversion LLT polynomials [22, Theorem 3.3] gives symmetry term by term in the height completion. Lemma 17.3 then places the stable symmetric function in the stated field.

For one letter, the equal-level strict row condition says $r_b = j_b + u_b$ for a unique $u \in \mathrm{RPP}(\lambda)$, and $\mathrm{dinv} = 0$. Thus the one-letter sum is the ordinary diagonal-weighted reverse-plane-partition series. Under the rim-hook insertion bijection, each box supplies an independent nonnegative multiplicity; its hook traverses a_b+1 negative Maya signs and l_b positive signs. The diagonal-hook calculation (11.22), applied to the weights q^{-1} and t, gives the factor $(1-q^{-a_b-1}t^{l_b})^{-1}$ for this multiplicity. Multiplying the independent geometric series proves (17.16). □

18. Supersymmetrization at fixed paired heights

We next prove the complete dominance condition for $G_\lambda[(t-1)X; q, t]$. It is useful to make the finite polynomial at each paired height explicit. Put $T = qt$ and write E_λ for the directed graph used in (17.9).

18.1. **Signed letters and the finite polynomial.** For $X = x_1 + \cdots + x_s$, use the signed residue order

$$\bar{1} < \bar{2} < \cdots < \bar{s} < s < \cdots < 2 < 1.$$

A positive letter η contributes tx_η and a negative letter $\bar{\eta}$ contributes $-x_\eta$. At a fixed height, positive equalities are permitted in columns and negative equalities in rows. On an attacking edge, a larger letter is an inversion, and an equal negative pair is also an inversion. These are the signed standardization rules for the plethystic difference $tX - X$.

Supersymmetrization follows from the following tableau rule. Standardize each group of equal positive letters in the order that permits column equality, and each group of equal negative letters in the opposite order, which permits row equality. Ordinary strict comparison after this standardization counts an equal pair precisely when the letter is negative. The descent enumerator of a standardized tableau is thereby transformed by $p_m[X] \mapsto (t^m - 1)p_m[X]$. Summing over the standardizations gives the signed rule for every LLT summand in (17.15); this is the superization identity of [22, Section 4]. Since standardization fixes the positions of all boxes, the factors $W_\lambda(r; q, t)$ are retained individually.

If ε_b is 1 for a positive letter and 0 for a negative letter, introduce the paired height $p_b = r_b + \varepsilon_b$. The full ordered integer position of the letter is

$$v_b = \begin{cases} 2sp_b - \eta_b, & \varepsilon_b = 1, \\ 2sp_b + \eta_b - 1, & \varepsilon_b = 0, \end{cases} \tag{18.1}$$

where $\eta_b \in \{1, \ldots, s\}$ is its absolute color. Different p-blocks are strictly ordered. Thus $p \in \mathrm{RPP}(\lambda)$, a box with $p = 0$ is negative, and all remaining restrictions concern adjacent boxes of the same paired height.

For fixed p and absolute colors, form the forced-positive set P, forced-negative set N, and the set U of unforced boxes by the following rules:

adjacent boxes of equal p	color comparison	forced sign
L, R in one row	$\eta_L \le \eta_R$	R negative
L, R in one row	$\eta_L > \eta_R$	L positive
U, D in one column	$\eta_U < \eta_D$	D negative
U, D in one column	$\eta_U \ge \eta_D$	U positive.

(18.2)

All zero-height boxes are added to N. A conflicting pair of forced signs means that the contribution is zero.

Write $\mathrm{out}_{>\eta_b}(b)$ for the number of outgoing neighbors whose absolute color exceeds η_b, and similarly for other inequalities and incoming neighbors. Define

$$\begin{aligned} D_0(\eta) &= \#\{u \to v \text{ in } E_\lambda : \eta_u \geq \eta_v\}, \\ B_b(\eta) &= \mathbf{1}_{\sigma_b=-1} + \mathrm{out}_{>\eta_b}(b) - \mathrm{in}_{\geq\eta_b}(b) \\ &= \mathbf{1}_{c_b\geq 0} - \mathrm{out}_{\leq\eta_b}(b) + \mathrm{in}_{<\eta_b}(b), \\ M_\lambda(p;q,t) &= \prod_{\sigma_b=-1} q^{j_b-p_b} \prod_{\sigma_b=1} t^{p_b-j_b}. \end{aligned} \tag{18.3}$$

The equality of the two formulas for B_b is (17.9).

Definition 18.1. For a nonnegative reverse plane partition p, put

$$K_{\lambda,p}[X;T] = \sum_{\eta:\lambda\to\{1,\dots,s\}} X^\eta T^{D_0(\eta)} \prod_{b\in P} T^{B_b(\eta)} \prod_{b\in N}(-1) \prod_{b\in U}(T^{B_b(\eta)}-1), \tag{18.4}$$

where $X^\eta = \prod_b x_{\eta_b}$, with the convention just stated for conflicting forced signs. This is a finite homogeneous polynomial in the x_i with coefficients in $\mathbf{Z}[T,T^{-1}]$.

Lemma 18.2. *The exact plethystic identity is*

$$G_\lambda[(t-1)X;q,t] = \sum_{p\in\mathrm{RPP}(\lambda)} M_\lambda(p;q,t) K_{\lambda,p}[X;qt]. \tag{18.5}$$

Equivalently, put all positive letters before all negative letters, in the order

$$s < \cdots < 1 < \bar{1} < \cdots < \bar{s}. \tag{18.6}$$

Then $K_{\lambda,p}$ is the signed-tableau sum on the equal-p skew components with weight

$$X^\eta (-1)^{\#\{b:\varepsilon_b=0\}} T^{D'(\eta,\varepsilon)+\#\{b:c_b\geq 0,\ \varepsilon_b=1\}}, \tag{18.7}$$

where D' is the inversion count on E_λ in (18.6)*, including equal negative pairs. The zero-height component is entirely negative.*

Proof. Inspection of the intervals with endpoints $2sp-\eta$ and $2sp+\eta-1$ proves (18.2). For example, if $\eta_L \leq \eta_R$, a positive letter at R lies at or below either allowed letter at L, with the positive equality forbidden in a row; hence R is negative. If $\eta_L > \eta_R$, the analogous obstruction forces L positive. The column argument reverses the equality rule. These are all restrictions, since the tableau conditions are adjacent conditions. Consequently every unforced sign may be chosen independently.

For an edge $u \to v$, its old inversion indicator is ε_u if $\eta_u < \eta_v$, and $1-\varepsilon_v$ if $\eta_u \geq \eta_v$. Thus

$$D = D_0 + \sum_b \big(\mathrm{out}_{>\eta_b}(b) - \mathrm{in}_{\geq\eta_b}(b)\big)\varepsilon_b.$$

Replacing r_b by $p_b - \varepsilon_b$ in the height weight and including the letter factor t^{ε_b} gives

$$M_\lambda(p;q,t)(-1)^{\#\text{negative}} T^{D+\sum_{\sigma_b=-1}\varepsilon_b}.$$

The exponent of each positive sign is exactly $B_b(\eta)$. Summing a forced positive, forced negative, or free sign gives, respectively, T^{B_b}, -1, or $T^{B_b}-1$. This proves (18.5).

Changing the residue order to (18.6) changes the inversion on an edge $u \to v$ by $\varepsilon_u - \varepsilon_v$. Therefore

$$D + \sum_{\sigma_b=-1}\varepsilon_b = D' + \sum_b (\mathbf{1}_{\sigma_b=-1} + \mathrm{out}(b) - \mathrm{in}(b))\varepsilon_b = D' + \sum_{c_b\geq 0}\varepsilon_b.$$

The equal-p order in (18.1) is already (18.6), proving the last assertion. Every rearrangement is finite in each original height degree by Lemma 17.3. □

19. The signed lattice and its intertwiners

In this section $\tau = T^{-1}$ is the vertex parameter; it is distinct from the Macdonald parameter t. There are finitely many ordered path colors. A state is a subset of these colors, identified with its zero-one vector. Sums and differences of states in an exponent are vector sums and differences. Define

$$\phi(U,V) = \sum_{i<j} U_i V_j, \qquad (z;\tau)_d = \prod_{r=0}^{d-1} (1 - z\tau^r).$$

All vector spaces below have the subset states as a basis and are defined over a rational function field in τ and the row parameters. We use P_a for a positive row of parameter a and N_b for a negative row of parameter b.

19.1. **The five matrices.** At an ordinary vertex let A, C be the vertical input and output, and I, K the horizontal input and output. The two local weights are

$$\begin{aligned} L_a(A,I;C,K) &= \mathbf{1}_{A+I=C+K}\mathbf{1}_{A\cap I=\varnothing} a^{|K|}\tau^{\phi(K,A+I)}, \\ L'_b(A,I;C,K) &= \mathbf{1}_{A+I=C+K}\mathbf{1}_{I\subset C} b^{|K|}\tau^{\phi(K,C-I)}. \end{aligned} \tag{19.1}$$

At a crossing write the incoming states from bottom to top as (I,J), and the outgoing states from bottom to top as (L,K). The following notation includes the parameters of the incoming lower and upper rows:

$$\begin{aligned} R_{x,y}(I,J;L,K) &= \mathbf{1}_{I+J=K+L}\mathbf{1}_{K\subset J}(-1)^{|J|-|K|}(x/y)^{|J|}(y/x;\tau)_{|J|-|K|}\tau^{\phi(J,K-J)}, \\ R'_{b,a}(I,J;L,K) &= \mathbf{1}_{I+J=K+L}\mathbf{1}_{I\cap J=\varnothing}\frac{(a/b)^{|L|}\tau^{\phi(L,I+J)}}{(-a/b;\tau)_{|I|+|J|}}, \\ R''_{b,c}(I,J;L,K) &= \mathbf{1}_{I+J=K+L}\mathbf{1}_{J\subset K}(c/b)^{|L|}(c/b;\tau)_{|K|-|J|}\tau^{\phi(L,K-J)}. \end{aligned} \tag{19.2}$$

Their types are, respectively,

$$P_xP_y \longrightarrow P_yP_x, \qquad N_bP_a \longrightarrow P_aN_b, \qquad N_bN_c \longrightarrow N_cN_b.$$

An inadmissible state has weight zero. These formulas are the five signed LLT matrices of [21, Section 3], with the lower-to-upper convention displayed explicitly.

For a word $w = \epsilon_1(z_1)\cdots\epsilon_m(z_m)$ in row types, write $M_w(A,C)$ for its one-column operator: sum over all intermediate vertical subsets and multiply the L or L' weights from bottom to top. This definition makes every statement about moving a crossing past a column an equality of finite matrices.

Lemma 19.1 (Local intertwining and the auxiliary braid)**.** *Let $\mathscr{R}_{\epsilon\eta}$ be the appropriate matrix in* (19.2)*. For the types $++, -+, --$,*

$$\mathscr{R}_{\epsilon\eta}M_{\epsilon\eta}(A,C) = M_{\eta\epsilon}(A,C)\mathscr{R}_{\epsilon\eta} \tag{19.3}$$

for every A, C. Crossings on disjoint factors commute. The two three-crossing paths in each of the following displays are equal, with all row parameters retained:

$$\begin{aligned} N_bP_aP_c &\longrightarrow P_aN_bP_c \longrightarrow P_aP_cN_b \longrightarrow P_cP_aN_b, \\ N_bP_aP_c &\longrightarrow N_bP_cP_a \longrightarrow P_cN_bP_a \longrightarrow P_cP_aN_b; \\ N_cN_bP_a &\longrightarrow N_bN_cP_a \longrightarrow N_bP_aN_c \longrightarrow P_aN_bN_c, \\ N_cN_bP_a &\longrightarrow N_cP_aN_b \longrightarrow P_aN_cN_b \longrightarrow P_aN_bN_c. \end{aligned} \tag{19.4}$$

The same-sign matrices are generically invertible and take the two-row vacuum to itself. The mixed matrix has, in each binary total-color sector B, one-dimensional image.

Proof. The finite identities (19.3) are the three Yang–Baxter identities for the five displayed matrices [21, Propositions 3.2–3.4]. To specify also the auxiliary braid required here, we record its common degeneration. The colored fermionic weights $W_z(A, B; C, D \mid r, s)$ satisfy the full finite identity [1, Proposition 5.1.4]

$$\sum_{I_2,J_2,K_2} W_{\chi/\gamma}(I_1, J_1; I_2, J_2 \mid r, s) W_{\chi/z}(K_1, J_2; K_2, J_3 \mid r, h) W_{\gamma/z}(K_2, I_2; K_3, I_3 \mid s, h)$$
$$= \sum_{I_2,J_2,K_2} W_{\gamma/z}(K_1, I_1; K_2, I_2 \mid s, h) W_{\chi/z}(K_2, J_1; K_3, J_2 \mid r, h) W_{\chi/\gamma}(I_2, J_2; I_3, J_3 \mid r, s). \quad (19.5)$$

All six exterior states are arbitrary binary vectors and every sum is finite. The precise degeneration to (19.2) is as follows. Put $\alpha = -S^2$. To P_x assign evaluation parameter x and spin $(x/\alpha)^{1/2}$; to N_y assign evaluation parameter α and spin $(-y/\alpha)^{-1/2}$. As $S \to 0$, the positive-positive, negative-positive, and negative-negative crossings tend to R, R', R'', respectively [21, Lemma 3.1]. These three limits have no state-dependent gauge. The gauges used for a vertical L-line are unnecessary in this auxiliary identity. Substitute these assignments simultaneously in (19.5). Taking its finite sums term by term gives the first and second pair of paths in (19.4). In particular, each mixed crossing has the displayed type $NP \to PN$.

For the remaining assertions one can read the matrices directly. On a conserved block, R has $K \subset J$ and R'' has $J \subset K$. Ordering the upper states by inclusion makes these matrices triangular; their diagonal entries are nonzero powers of a spectral ratio. Their vacuum entries are 1. In the mixed case $B = I \sqcup J$ is binary and the coefficient depends on the output split but not on the input split. Its image is therefore one dimensional in each such sector, as asserted. These observations also fix the normalizations in (19.3). □

19.2. **The exact finite path model.** For fixed p, let $\alpha_h = \{b : p_b < h\}$ and $\beta_h = \{b : p_b \le h\}$. Use one path color for each nonempty plateau, in decreasing order of h. The lowest plateau, when it has height zero, is consequently the last path color. In original column j, the path for plateau h starts at the Maya site $\alpha'_{h,j} - j$ and ends at $\beta'_{h,j} - j$. Its step for box (i, j) crosses site

$$u = i - j = -c_b. \quad (19.6)$$

For consecutive heights the endpoints sew because $\beta_h = \alpha_{h+1}$; empty levels may be inserted. Thus in original column j the colored segments concatenate from $-j$ to $\lambda'_j - j$.

Read the alphabet (18.6) from bottom to top. A positive letter a uses row P_a and a negative letter $\bar{a}$ uses row N_a, so the row order is

$$P_s \cdots P_1 N_1 \cdots N_s. \quad (19.7)$$

In a positive row use L, and in a negative row use L'. These rules follow directly from the ordinary path encoding of a transposed tableau: the positive part is row-weak and column-strict after transposition, whereas the negative part is row-strict and column-weak. The binary constraints in (19.1) express exactly these two strip rules. The vertical boundary consists of the just-specified Maya states; both horizontal boundaries are empty. Only finitely many sites differ from the constant tails, whose weights are 1.

We recall explicitly the normalization between inversion and these local weights. If $C(f)$ denotes the sum of the ϕ-exponents in (19.1) for a signed filling f, then

$$D'(f) + C(f) = m_{\lambda,p}, \quad (19.8)$$

where the right side depends only on the boundary diagrams. Here the within-component same-content contribution is included in D'. One way to compute $m_{\lambda,p}$ without any normalization choice is to use $|\lambda|$ distinct negative labels in row-major order; then $D' = |S_\lambda|$ and $m_{\lambda,p} = |S_\lambda| + C(f)$.

The pointwise identity (19.8) can be checked directly, including signed ties. Write $\rho_h = \beta_h / \alpha_h$ for a plateau. Fix two heights $h > h'$; their path colors occur in this order in ϕ. Fix a box $b \in \rho_h$

at site $u = i_b - j_b$, and let r be the row of its signed letter in (18.6). For the lower plateau put

$$a_0(u) = \mathbf{1}_{u\in\{\alpha'_{h',j}-j:j\geq 0\}}, \qquad n_v^{<r} = \#\{a \in \rho_{h'} : i_a - j_a = v,\ f(a) < r\},$$

where the comparison is in the row order of the signed alphabet; define $n_v^{\leq r}$ similarly and let n_v count all boxes of $\rho_{h'}$ at site v. Conservation $A + I = C + K$ through the preceding rows shows that the incoming vertical bit of this lower color at site u is $a_0(u) + n_{u-1}^{<r} - n_u^{<r}$, its outgoing vertical bit is $a_0(u) + n_{u-1}^{\leq r} - n_u^{\leq r}$, and its incoming horizontal bit is $n_{u-1}^{\leq r} - n_{u-1}^{<r}$. The contribution of b and this lower color to the local exponent is therefore

$$C_{b,h'} = \begin{cases} a_0(u) + n_{u-1}^{\leq r} - n_u^{<r}, & r \text{ positive}, \\ a_0(u) + n_{u-1}^{<r} - n_u^{\leq r}, & r \text{ negative}. \end{cases}$$

Every relevant cross-plateau pair is comparable in the original diagram, so the lower-height box is northwest. The same-content attacks have the higher box as source; the adjacent-content attacks have the lower box, at site $u - 1$, as source. Consequently their signed inversion contribution is

$$D'_{b,h'} = \begin{cases} n_u^{<r} + n_{u-1} - n_{u-1}^{\leq r}, & r \text{ positive}, \\ n_u^{\leq r} + n_{u-1} - n_{u-1}^{<r}, & r \text{ negative}. \end{cases}$$

Thus $D'_{b,h'} + C_{b,h'} = a_0(u) + n_{u-1}$ in both cases, a quantity determined entirely by the plateau boundaries. Summing over all b and all pairs $h > h'$ accounts for every summand of C and every cross-plateau attack. Within one plateau, signed strip monotonicity makes each same-content pair contribute one to D' and each adjacent-content pair contribute zero; there is no same-color summand in ϕ. Hence $D'(f) + C(f)$ is independent of the signed filling. Evaluation on the all-negative row-major filling identifies this constant with the specified $m_{\lambda,p}$ and proves (19.8).

Proposition 19.2 (Lattice realization)**.** *Let $Z_{\lambda,p}(X;\tau)$ be the above finite path partition function with row parameters*

$$a_i(u) = \begin{cases} x_i/\tau, & u \leq 0, \\ x_i, & u > 0, \end{cases} \qquad b_i(u) = -x_i. \tag{19.9}$$

When $p = 0$ occurs, impose the vacuum Maya state of that last color between the positive and negative row blocks. Then

$$K_{\lambda,p}[X;T] = T^{m_{\lambda,p}} Z_{\lambda,p}(X;T^{-1}). \tag{19.10}$$

Proof. The path encoding is a bijection with the signed fillings in Lemma 18.2. The product of the row monomials supplies $X^\eta(-1)^{\#\text{negative}}$. The local exponent is $\tau^{C(f)} = T^{D'(f)-m_{\lambda,p}}$ by (19.8). A positive box has $c_b \geq 0$ precisely when its site is $u \leq 0$; the extra factor τ^{-1} in (19.9) supplies the remaining T for that box. This proves (19.10). The zero-height restriction says that its positive skew shape is empty, which is exactly the imposed midpoint vacuum. □

Proposition 19.3 (Fixed-height symmetry)**.** *For every p, the polynomial $K_{\lambda,p}(X;T)$ is symmetric in X. A chosen set of k variables can be moved to the innermost k matched pairs in* (19.7)*, while retaining all other rows.*

Proof. Exchange two neighboring variable pairs by exchanging their positive rows with R and their negative rows with R''. Equation (19.3) moves these crossings through every column. On the left all positive parameters are multiplied by the same τ^{-1}; hence their ratios are unchanged. The positive crossing is therefore identical on the two sides of the parameter cut. The negative crossing is identical there as well. The empty horizontal boundaries are fixed, with coefficient 1, so the complete partition function is unchanged.

If a zero-height color is present, the positive rows must preserve its vacuum state. Equivalently, their horizontal occupations of that color all vanish. A same-sign positive crossing preserves the sum of these occupations in the two rows, and a negative crossing does not act on them. Thus

the paired permutation commutes with this restriction and the same argument applies. The boundary scalar in (19.10) is independent of the variables. Adjacent pair exchanges generate all variable permutations, including one that places any chosen set innermost. □

20. Fusion with its full null submodule

20.1. **A rank-one identity for arbitrary incoming states.** For a binary set B define

$$v_B(a,b)=\sum_{I\subset B} a^{|I|}b^{|B-I|}\tau^{\phi(I,B)}\,|I,B-I\rangle,\qquad Z_d(a,b)=\prod_{r=0}^{d-1}(b+a\tau^r). \tag{20.1}$$

The coefficient sum of v_B is $Z_{|B|}$: each color in B contributes $b+a\tau^r$, where r is the number of larger colors of B. Let $\mathcal{K}(a,b)$ be the vector space with basis e_B, embedded in P_aN_b by $e_B\mapsto v_B(a,b)$. This embedding is defined whenever $b\neq 0$, including all specializations below, because its all-negative coordinate is $b^{|B|}$.

Lemma 20.1 (Raw two-row identity)**.** *The incoming horizontal states I,J need not be disjoint. Every nonzero product $L_a(A,I;M,K)L'_b(M,J;C,L)$ has $K\cap L=\varnothing$. Writing $D=K\sqcup L$, for every $K\subset D$ one has*

$$\begin{aligned}&\sum_M L_a(A,I;M,K)L'_b(M,J;C,D-K)\\ &\qquad=\mathbf{1}_{A\cap I=\varnothing}\mathbf{1}_{J\subset C}\mathbf{1}_{A+I+J=C+D}\,\tau^{\phi(D,C-J)}a^{|K|}b^{|D-K|}\tau^{\phi(K,D)}.\end{aligned} \tag{20.2}$$

Thus every fixed incoming split has its outgoing vector on the line spanned by $v_D(a,b)$.

Proof. The negative rule gives $L\subset M$, while the positive rule gives $M\cap K=\varnothing$, proving disjointness. The intermediate state is forced to be $M=A+I-K$. Under the three displayed conditions it is binary for every $K\subset D$. Indeed if $D_i=1$, conservation and $J_i\le C_i$ force $A_i+I_i=1$ and $C_i=J_i$; if $D_i=0$, there is no split choice at that color. The two exponents satisfy

$$\phi(K,A+I)+\phi(D-K,C-J)=\phi(D,C-J)+\phi(K,A+I-C+J)=\phi(D,C-J)+\phi(K,D).$$

The monomial factors are $a^{|K|}b^{|D-K|}$, proving (20.2) in every number of colors. □

Lemma 20.2 (Parameter jump)**.** *Suppose $A+B=C+D$ with B,D binary. A two-row column of parameters (a,b), applied to the incoming vector $v_B(a_0,b_0)$, has outgoing D-component $\mathcal{J}(A,B;C,D)v_D(a,b)$, where*

$$\mathcal{J}(A,B;C,D)=b_0^{|B|}\tau^{\phi(D,C-B)}\prod_{i\in B-C}\left(\frac{a_0}{b_0}\tau^{h_i}\right)\prod_{i\in B-D}\left(1+\frac{a_0}{b_0}\tau^{h_i}\right), \tag{20.3}$$

$$h_i=\#\{j\in B: j>i\}+\#\{j\in D: j<i\}.$$

If conservation fails, the coefficient is zero. Here $B-C$ and $B-D$ denote set differences, whereas $C-B$ in ϕ denotes vector difference.

Proof. Apply (20.2) with $J=B-I$ and sum the coefficient of $|I,B-I\rangle$ in $v_B(a_0,b_0)$. The allowed input splits are exactly

$$B-C\subset I\subset B-A.$$

The free colors are $B-D$, as is checked by the conserved binary possibilities for (A_i,B_i,C_i,D_i). The input exponent, after removing the common output factor, is

$$\phi(D,C-B)+\phi(D,I)+\phi(I,B)=\phi(D,C-B)+\sum_{i\in I}h_i.$$

Each forced color contributes $(a_0/b_0)\tau^{h_i}$, and each free color contributes $1+(a_0/b_0)\tau^{h_i}$. This is (20.3). □

Proposition 20.3 (Finite capacity as a quotient)**.** *At $b = -a\tau^{\ell}$ with $a \neq 0$, where ℓ is a nonnegative integer, the span*

$$\mathcal{H}^{>\ell}(a,b) = \operatorname{span}\{e_B : |B| > \ell\} \subset \mathcal{K}(a,b)$$

is preserved by every vertical matrix element of a paired column. Its quotient is spanned by $|B| \leq \ell$.

More precisely, take nonnegative integers $\ell_1 \leq \ell_0$ and parameters $b_i = -a_i\tau^{\ell_i}$ with $a_i \neq 0$. For fixed vertical input A and output C, the identity gluing followed by the first paired column with parameters (a_1, b_1) defines a map

$$T^{A,C}_{1\leftarrow 0} : \mathcal{K}(a_0,b_0) \longrightarrow \mathcal{K}(a_1,b_1), \qquad e_B \longmapsto \sum_{\substack{D \text{ binary} \\ A+B=C+D}} \mathcal{J}_{a_0,b_0}(A,B;C,D)e_D,$$

where $\mathcal{J}_{a_0,b_0}$ denotes the coefficient $\mathcal{J}$ in (20.3)*. This map descends to*

$$\mathcal{K}(a_0,b_0)/\mathcal{H}^{>\ell_0}(a_0,b_0) \longrightarrow \mathcal{K}(a_1,b_1)/\mathcal{H}^{>\ell_1}(a_1,b_1).$$

Proof. Write $\iota_i(e_B) = v_B(a_i,b_i)$. The parameter-jump identity gives the equality of raw two-row maps

$$\iota_1 T^{A,C}_{1\leftarrow 0} = M_{P_{a_1}N_{b_1}}(A,C)\iota_0.$$

Here the matrix entry of the one-column operator on the right, from $|I,J\rangle$ to $|K,L\rangle$, is

$$\sum_M L_{a_1}(A,I;M,K)L'_{b_1}(M,J;C,L).$$

For a transition with $|B| > \ell_0$ and $|D| \leq \ell_1$, scan the ordered colors from smallest to largest. The integer

$$\#\{\text{unscanned colors in } B\} + \#\{\text{scanned colors in } D\}$$

starts above ℓ_0 and ends at or below ℓ_0. It therefore has a downward crossing to height ℓ_0. Such a crossing occurs at a color $i \in B - D$, and its new height is $h_i = \ell_0$. At $b_0 = -a_0\tau^{\ell_0}$ the corresponding factor in (20.3) is zero. Hence $T^{A,C}_{1\leftarrow 0}(\mathcal{H}^{>\ell_0}) \subset \mathcal{H}^{>\ell_1}$, proving descent. Taking identical input and output parameters proves invariance at a fixed capacity. The argument uses the unnormalized vectors v_B, including when $Z_{|B|} = 0$. □

Corollary 20.4 (An innermost pair)**.** *For a matched pair with left parameters $(x/\tau, -x)$ and right parameters $(x, -x)$, every left fused state of degree at least two is annihilated by the cut-and-right-tail observer with empty horizontal boundary. This is an identity for each choice of vertical boundary states, including when further rows surround the pair.*

Proof. The left capacity is one, and its high submodule is invariant up to the cut. The first right column realizes the map of Proposition 20.3 with $(a_0,b_0) = (x/\tau,-x)$, $\ell_0 = 1$, and $(a_1,b_1) = (x,-x)$, $\ell_1 = 0$. Thus (20.3) sends a high state into the new capacity-zero high submodule, which remains invariant through the right tail. The empty horizontal boundary kills it. The argument is valid separately for every intermediate vertical state, so additional rows may be summed afterward. □

20.2. **Moving an external row past a fused pair.** The positive train and negative train are the two-crossing maps

$$\begin{aligned} \mathcal{B}^{+}_{a,b;c}: &\quad P_aN_bP_c \xrightarrow{R'_{b,c}} P_aP_cN_b \xrightarrow{R_{a,c}} P_cP_aN_b, \\ \mathcal{B}^{-}_{c;a,b}: &\quad N_cP_aN_b \xrightarrow{R'_{c,a}} P_aN_cN_b \xrightarrow{R''_{c,b}} P_aN_bN_c. \end{aligned} \tag{20.4}$$

Only the outside spectral parameter occurs in their possible mixed-crossing denominators.

Lemma 20.5 (Image preservation at resonance)**.** *The trains restrict to maps*

$$\mathcal{K}(a,b)P_c \longrightarrow P_c\mathcal{K}(a,b), \qquad N_c\mathcal{K}(a,b) \longrightarrow \mathcal{K}(a,b)N_c.$$

They intertwine columns, also at $b=-a\tau$, for a generic independent outside parameter c.

Proof. For generic a,b, the image of $R'_{b,a}$ is exactly $\mathcal{K}(a,b)$: on total set B its output is $v_B(a,b)/Z_{|B|}(a,b)$. The two auxiliary braid identities (19.4) show that both trains preserve this image. Equation (19.3) gives column intertwining.

To specialize, use the unnormalized generators $v_B(a,b)$. Different B occupy different total-color sectors, and the all-negative coefficient $b^{|B|}$ is nonzero at $b=-a\tau$. They therefore give a regular constant-rank subbundle through this specialization. The external trains are regular there over the rational function field in c. In each output sector extract the putative scalar by dividing its all-negative coordinate by $b^{|D|}$. Proportionality of every other coordinate to this scalar is a rational identity regular at the specialization, so it continues to hold. The same reasoning applies to intertwining. The specialization uses only unnormalized v_B and regular external trains; the normalized own-pair crossing may have a pole at this parameter. □

Lemma 20.6 (High-submodule descent for both trains)**.** *At $b=-a\tau$, both maps in Lemma 20.5 preserve the high submodule $\mathcal{H}=\operatorname{span}\{e_B : |B|\geq 2\}$ of the inner pair. At $b=-a$, they preserve the nonempty-state submodule and induce the identity on the inner vacuum quotient.*

Proof. By Lemma 20.5, it suffices to prove that every high input has zero coefficient in the vacuum and singleton output sectors. Rescaling the inner basis permits $a=1$. First take $b=-\tau$ and positive outside parameter $c=z$. Write the input split as $I\subset B$, $J=B-I$, and the outside state as H. To read the output inner set D, of degree $d\leq 1$, choose its positive component empty. Its negative component is then D, whose coefficient in v_D is nonzero. The intermediate outside positive state is $L=H+J-D$. The two crossing formulas give, before dividing by that nonzero output coefficient,

$$(-1)^{|J|}\tau^{d-|H|}\tau^{\phi(I,B)+\phi(L,D)}\frac{(z;\tau)_{|H|+|J|-d}}{(z/\tau;\tau)_{|H|+|J|}}. \tag{20.5}$$

The nonzero-state restrictions are $J\cap H=\varnothing$, $I\cap H\subset D$, and binary L. They will also be used in the cases below.

For $d=0$, conservation forces $B\cap H=\varnothing$, and all $I\subset B$ are allowed. The Pochhammer ratio is

$$\frac{1-z\tau^{|H|+|J|-1}}{1-z/\tau}.$$

Consequently the split sum is a linear combination of

$$\prod_{j\in B}(\tau^{r_j}-1), \qquad \prod_{j\in B}(\tau^{r_j}-\tau), \qquad r_j=\#\{b\in B : b>j\}.$$

The largest color kills the first product; since $|B|\geq 2$, the second-largest kills the second.

For $D=\{i\}$ the Pochhammer ratio in (20.5) is $1/(1-z/\tau)$, independent of the split. Up to a constant, the exponent assigned to a free color $j\in I$ is $r_j-\mathbf{1}_{j<i}$. If $i\in H-B$ and $B\cap H=\varnothing$, all colors of B are free. If the largest lies above i, its exponent is zero; otherwise the second-largest has exponent zero. If $i\in B-H$, it is forced into J; if $B\cap H=\{i\}$, it is forced into I. In either case the exponents on the other colors are exactly the numbers of larger colors in $B-\{i\}$, so the largest remaining color has exponent zero. The free split product therefore vanishes. Two repeated colors are incompatible with a singleton output. This proves the positive assertion.

For the negative train put $c=-u\tau$. Use the same input split and again read output positive component empty and negative component D. The intermediate state is $H+I$; the restrictions

are $H \cap I = \varnothing$, $D \subset H + I$, and binary outside output $H + B - D$. The summand obtained from (19.2) is

$$(-\tau)^{|J|}\tau^{\phi(I,B)+\phi(D,H+I)}u^{-d}\frac{(u^{-1};\tau)_{|H|+|I|-d}}{(u^{-1}/\tau;\tau)_{|H|+|I|}}. \tag{20.6}$$

For $d = 0$ the two resulting products have factors $\tau^{r_j} - \tau$ and $\tau^{r_j+1} - \tau$. The second-largest and largest colors, respectively, make them zero. For $D = \{i\}$ the Pochhammer ratio is constant, and each free color has exponent $r_j + \mathbf{1}_{j>i}$, with factor $\tau^{r_j+\mathbf{1}_{j>i}} - \tau$. If i is outside B, the largest color above i, or the second-largest if all are below i, has exponent one. If $i \in B$, its split is forced and the remaining exponents are the numbers of larger colors in $B - \{i\}$ plus one. Their largest color again has exponent one. The free set is nonempty because $|B| \geq 2$. This proves negative descent.

For capacity zero put $b = -a$ instead. Only a vacuum output is visible. In either train the two Pochhammer factors cancel, leaving, up to a nonzero common factor,

$$\sum_{I \subset B}(-1)^{|B-I|}\tau^{\phi(I,B)} = \prod_{j \in B}(\tau^{r_j} - 1).$$

This vanishes for $B \neq \varnothing$. For $B = \varnothing$ the surviving coefficient is 1 and the outside state is unchanged. This proves the final assertion. □

We record an explicit quotient matrix, which will also make the genericity of the positive train transparent. Let the inner pair have $a = 1, b = -\tau$ and let the outside positive parameter be z. On a binary total set $S = \{i_1 < \cdots < i_d\}$ use basis 0 for inner vacuum and outside S, and basis j for inner $\{i_j\}$ and outside $S - \{i_j\}$. If $B_S(z)$ is the induced train and $C = (1 - z/\tau)B_S(z)$, its entries are

$$\begin{aligned} C_{00} &= \tau^{-d} - z/\tau, & C_{0j} &= z(\tau - 1)/\tau, \\ C_{i0} &= -\tau^{i-1-d}, & C_{ii} &= \tau^{i-d}, \\ C_{ij} &= (\tau - 1)\tau^{i-1-d} \quad (j < i), & C_{ij} &= 0 \quad (j > i). \end{aligned} \tag{20.7}$$

For $d = 0$ the matrix is 1. To verify the formula, choose output inner positive state empty in the two terms of the input singleton in (20.5), and divide by its negative coefficient $-\tau$. The other output split has the same scalar by Lemma 20.5. Factoring τ^{i-d} from row $i > 0$ leaves lower-right block $Q = I + (1 - \tau^{-1})N$, where N has entries 1 strictly below the diagonal. Since $Q^{-1}(1, \ldots, 1)^{\mathsf{t}} = (1, \tau^{-1}, \ldots, \tau^{-(d-1)})^{\mathsf{t}}$, its Schur complement is $\tau^{-d}(1 - z/\tau)$. Hence

$$\det B_S(z) = \tau^{-d(d+1)/2}(1 - z/\tau)^{-d}. \tag{20.8}$$

If one color i is repeated in the total charge, the block is one dimensional and has coefficient $\tau^{\#\{h \in H : h < i\} - |H| + 1}/(z - \tau)$, where H is the outside set. These formulas are consequences of the trains and do not require an inverse mixed matrix.

21. The capacity of an arbitrary nested block

Fix k matched variables and use the nested order $P_k \cdots P_1 N_1 \cdots N_k$. The left parameters are $a_i = x_i/\tau$, $b_i = -x_i = -\tau a_i$; the right parameters are $a_i = x_i$, with b_i unchanged. The horizontal degree of a state is the sum of its occupations over these $2k$ rows. A right observer is a linear functional obtained by applying a finite word of right-column matrices and then the empty horizontal covector. We may fix all vertical states throughout such a word. Additional left columns before the cut are allowed as part of the observer. All statements below hold before summing any fixed outer row data.

21.1. **The raw embedding and its kernel.** Let $\mathcal{K}_i = \mathcal{K}(a_i, b_i)$, $\mathcal{H}_i = \operatorname{span}\{e_B : |B| \geq 2\}$, and $\mathcal{F}_i = \mathcal{K}_i/\mathcal{H}_i$. Define a map

$$E_k : \mathcal{K}_1 \otimes \cdots \otimes \mathcal{K}_k \longrightarrow P_k \cdots P_1 N_1 \cdots N_k \tag{21.1}$$

recursively. The map E_1 is the embedding (20.1). For $k > 1$, expand the last factor as $P_k N_k$, move P_k successively to the left of $\mathcal{K}_{k-1}, \ldots, \mathcal{K}_1$ using $\mathcal{B}^+$, and apply E_{k-1} to the remaining $k-1$ fused factors. In formulas, with neighboring positions followed through the successive swaps,

$$E_k = (\mathrm{id}_{P_k} \otimes E_{k-1} \otimes \mathrm{id}_{N_k})(\mathcal{B}^+_{1,k} \cdots \mathcal{B}^+_{k-1,k} \otimes \mathrm{id}_{N_k}). \tag{21.2}$$

The rightmost indicated train acts first. The final expansion of each remaining $\mathcal{K}_i$ is part of E_{k-1}; thus (21.2) is an equality of typed linear maps, with the tensor positions changing at each successive train.

Lemma 21.1. *The map E_k preserves total color charge, takes tensor vacuum to raw vacuum, and intertwines every left-column matrix of the adjacent-pair and nested models. Every left-vacuum reachable vector of the nested model lies in* $\operatorname{im} E_k$.

Proof. Conservation holds at every local weight and crossing. Their vacuum entries are 1. Lemma 20.5 and (19.3) allow each crossing in (21.2) to pass through a left column, giving intertwining by induction. Apply it successively to a word of left columns and use the vacuum assertion. This expresses the raw vector as E_k of the corresponding adjacent-pair vector; no injectivity is required. □

Proposition 21.2 (Kernel of the right pairing)**.** *Every cut-and-right-tail observer annihilates*

$$E_k \left(\sum_{i=1}^{k} \mathcal{K}_1 \otimes \cdots \otimes \mathcal{H}_i \otimes \cdots \otimes \mathcal{K}_k \right). \tag{21.3}$$

The assertion holds for each vertical matrix element, with arbitrary surrounding positive and negative rows.

Proof. Induct on k. For $k = 1$ use Corollary 20.4. For an old factor $i < k$, the positive trains in (21.2) preserve $\mathcal{H}_i$ by Lemma 20.6. The inductive E_{k-1} therefore has invisible image in the inner nested block. This vanishing is an identity for all its vertical boundary states, so the outer P_k, N_k rows preserve it after those states are summed.

For the new factor $\mathcal{H}_k$, simultaneously permute the matched variable k to the innermost position. Denote the same-sign intertwiner by Ξ. Its positive and negative components are generically invertible by Lemma 19.1. Their spectral ratios are unchanged at the cut, so Ξ intertwines the entire left-and-right network and fixes the empty horizontal boundary.

An equivalent expression for E_k first applies E_{k-1} and expands $\mathcal{K}_k$, then moves P_k left through the old negative block and subsequently through the old positive block. Its equivalence with (21.2) follows by the first identity in (19.4), one old pair at a time. The crossings of Ξ that move P_k back through the positive block cancel the last positive-positive crossings of this expression. The remaining crossings can be ordered, by the second identity in (19.4) and commutation on disjoint strands, into

$$N_j P_k N_k \longrightarrow P_k N_j N_k \longrightarrow P_k N_k N_j \qquad (j = k-1, \ldots, 1).$$

These are precisely the negative trains $\mathcal{B}^-_{b_j;a_k,b_k}$. Every one preserves $\mathcal{H}_k$. The resulting high pair is adjacent and innermost, hence invisible by Corollary 20.4. Since Ξ preserves the pairing, the original image is invisible as well. This proves the induction. □

Theorem 21.3 (Left capacity)**.** *At any section in the left region, the pairing of a left-vacuum reachable state with its cut-and-right observer is zero in horizontal degree greater than k. More generally the same vanishing holds on* $\operatorname{im} E_k$.

Proof. By Proposition 21.2, the pairing on $\operatorname{im} E_k$ factors through

$$\bigotimes_{i=1}^{k} \mathcal{K}_i \Big/ \sum_i \mathcal{K}_1 \otimes \cdots \otimes \mathcal{H}_i \otimes \cdots \otimes \mathcal{K}_k = \bigotimes_{i=1}^{k} \mathcal{F}_i.$$

Each factor has degree zero or one, and every map used preserves total degree. This quotient has no degree greater than k. Lemma 21.1 places every reachable state in its image. The argument is over a generic rational function field. Finite-network coefficients are polynomials in the row parameters, so the vanishing extends to all specializations by clearing denominators. □

21.2. **The zero-height boundary.** Suppose the last path color g has height zero. Its positive part is the vacuum Maya diagram. Its vertical bit below the positive block is one for $u \le 0$ and zero for $u > 0$. Let Π project onto the subspace in which every positive horizontal g-bit vanishes. Since these bits are nonnegative, this is also the projection onto total positive g-occupation zero.

Lemma 21.4 (One projected pair)**.** *For the adjacent pair the outgoing fused lines are $\widehat{v}_D = \Pi v_D$. At left capacity one their degree-at-least-two part is invariant and invisible to the right observer. Their coefficient sums are*

$$\epsilon(\widehat{v}_D) = \begin{cases} Z_{|D|}(a,b), & g \notin D, \\ b\, Z_{|D|-1}(a\tau, b), & g \in D. \end{cases} \tag{21.4}$$

Proof. Apply (20.2) and then Π. If $g \in D$, it is forced into the negative component. As the largest color, it gives

$$\widehat{v}_D(a,b) = b\, v_{D-\{g\}}(a\tau, b) \quad \text{with a negative } g\text{-bit appended.}$$

This proves (21.4). At $b = -a\tau$ the remaining colors see capacity zero.

On a left column the vertical input g-bit is one. For a projected incoming vector containing g, the positive rule forces its intermediate bit to one, and negative conservation forces the outgoing horizontal g-bit to remain one. The other colors have effective parameters $(a\tau, b)$; Proposition 20.3 keeps every nonempty set of them in its capacity-zero high submodule. If the input omits g, ordinary high-to-low vanishing applies and survives output projection.

At the first right column the positive vertical g-bit is zero. An incoming negative g-arrow can only be absorbed vertically, and $C_g - J_g = 0$. When the old high vector contains g, its remaining colors form a nonempty capacity-zero vector $v_{D-\{g\}}(a\tau, b)$. Its effective old parameters $(a\tau, b) = (x, -x)$ equal the new right parameters, so its vacuum coefficient vanishes by (20.3). Positive-degree right output is an ordinary capacity-zero high vector and is killed by the empty boundary. When the old vector omits g, the ordinary jump proof applies. This proves the assertion pointwise in the other vertical colors. □

Proposition 21.5 (Postponing the single projection)**.** *The left-capacity theorem holds with the zero-height restriction. Its raw comparison map is*

$$\widehat{E}_k = \Pi E_k \tag{21.5}$$

on the full tensor product of unprojected $\mathcal{K}_i$. For the right pairing it factors through $\bigotimes_i \mathcal{F}_i$.

Proof. Let $M_L(A, C)$ be the unprojected nested left-column operator and m the vertical g-bit at the midpoint. Conservation through the positive block gives

$$\text{positive output count} = \text{positive input count} + 1 - m \;\ge\; \text{positive input count}.$$

Thus

$$\Pi M_L = \Pi M_L \Pi. \tag{21.6}$$

On the Π-space, a surviving output forces $m = 1$, which is exactly the midpoint vacuum constraint. The projected column is $\widehat{M}_L = \Pi M_L \Pi$. Intertwining gives

$$\widehat{M}_L \widehat{E}_k = \Pi M_L \Pi E_k = \Pi M_L E_k = \Pi E_k M_{\text{adjacent}}.$$

The projected map takes vacuum to vacuum, hence contains every projected reachable state in its image.

The kernel induction of Proposition 21.2 survives the final projection. For an old high factor apply the full positive trains first; they preserve that factor. The final projector factors into its parts on the new outer positive row and on the old positive block. The inductive projected inner vanishing applies for every vertical boundary.

For the new high factor use the simultaneous same-sign permutation Ξ before Π. It preserves total positive g-occupation, so

$$\Xi\Pi = \Pi\Xi. \tag{21.7}$$

Cancel its positive-positive crossings as in the unprojected proof. The remaining full negative trains preserve the ordinary high submodule. Only at the end apply Π; the innermost vectors become $\widehat{v}_B$ with $|B| \geq 2$, invisible by Lemma 21.4. Equation (21.7) transfers this vanishing to the original pairing. Thus the same tensor quotient proves capacity k.

On the right the bottom g-bit is zero. An incoming Π-state has positive output count $-m$, forcing $m = 0$ and that count zero. The ordinary right column on this subspace is the projected right column, as required for all observers in the proof. □

The proof applies with a translated vacuum: if its occupied sites are $u \leq s$ and the parameter cut is translated to s, the same statements hold on that left region. It also applies to any initial state in $\mathrm{im}(\Pi E_k)$; reachability from an earlier vertical boundary was used only to establish membership in this image.

22. The right ramp and all dominance inequalities

Let $h(u)$ denote the total horizontal occupation of the chosen k pairs immediately after site u. We prove the observable bound

$$h(u) \leq r_k(u) := \max\{k - \max(u, 0), 0\}. \tag{22.1}$$

Here a bound means that the corresponding complete coefficient, with fixed outside data, is zero when the inequality fails. For $u \leq 0$ this is Theorem 21.3 and Proposition 21.5.

22.1. **Successive disappearance on the right.**

Lemma 22.1. *On the right tail, the innermost pair is observably vacuum after the first column and is the vertical identity thereafter. The next pair has this property after the second column, and pair i after column i. In particular, all nonempty horizontal states beyond column k are annihilated by the right boundary. The assertion includes the zero-height restriction.*

Proof. At the first right site, the adjacent inner pair has parameters $(x_1, -x_1)$. Equation (20.2) puts every incoming raw state on its fused output lines. Their positive-degree span is the invariant capacity-zero submodule and is killed by the empty right boundary. Thus its observable output is vacuum. With both input and output vacuum the pair has coefficient $\mathbf{1}_{A=C}$, so it is the vertical identity at each subsequent site. At site two the next pair is consequently adjacent. Repeat the same argument.

For the zero-height color, a right positive row has input vertical bit zero and horizontal bit zero; both its outputs are therefore zero. In the negative row an incoming horizontal bit can only be deposited vertically, giving $C_g = J_g$ and horizontal output zero. The last-color contribution to $\phi(D, C - J)$ vanishes. The same argument applies on all other colors and at each successive pair. Since every step is an identity for each vertical matrix element, surrounding rows are retained throughout. □

22.2. **A last auxiliary color when no zero plateau is present.**

Proposition 22.2. *If there is no zero-height plateau, the bound* (22.1) *holds at every positive site.*

Proof. Fix $1 \le s \le k$. Adjoin one last path color g. Before the first right site its horizontal state and all its vertical bits are zero. Thus extending the earlier left network has no effect, and its state lies in $\operatorname{im} E_k$ with no g-occupation. Encode the first s original right columns as extended left columns. At each encoded column give g bottom vertical bit one and top bit zero, and require its positive horizontal occupations to vanish. The projection can be postponed until the end by (21.6). The projected left-capacity argument applies to this initial state in $\operatorname{im}(\Pi E_k)$.

We give a right observer that selects the required auxiliary path history. Number the negative rows $1, \dots, k$ from bottom to top. At a negative vertex its single-color equations are

$$A_g + J_g = C_g + L_g, \qquad J_g \le C_g.$$

For a column with bottom bit one and top bit zero, let $J = \{j_1 < \cdots < j_m\}$ and $L = \{l_1 < \cdots < l_{m+1}\}$ be the incoming and outgoing horizontal row positions. Propagating the vertical bit through the rows gives exactly

$$l_1 \le j_1 < l_2 \le j_2 < \cdots < l_m \le j_m < l_{m+1}. \tag{22.2}$$

To see both necessity and sufficiency, the incoming vertical one can first exit horizontally at l_1; the next vertical one must enter from j_1, so $l_1 \le j_1 < l_2$, and the argument repeats. When an exit and entry share a row the full binary vertex is allowed. This also proves uniqueness of the intervening vertical bits. If the output is $\{1, \dots, m+1\}$, the input is forced to be $\{1, \dots, m\}$. Hence a final prefix $\{1, \dots, s\}$ forces the successive prefixes after all encoded columns.

Append k true extended right columns with bottom g-bit zero, top word $0^{k-s}1^s$, and final empty horizontal boundary. Positive g-bits remain zero automatically. For a bottom-zero, top-zero column with s particles, the same binary equations give

$$j_1 < l_1 \le j_2 < l_2 \le \cdots \le j_s < l_s. \tag{22.3}$$

Every particle advances at least one negative row. After $k - s$ such columns, the last particle fits in row k only if its initial row is at most s. Since there are s distinct particles, the initial set must be $\{1, \dots, s\}$. Applying the same bound at intermediate times forces the successive blocks $\{r+1, \dots, r+s\}$ for $0 \le r \le k - s$. The last s top-one columns then force their successive exits. The resulting right-column word, followed by the empty boundary, selects the prefix.

Along the unique selected history, each positive encoded vertex has $A_g + I_g = 1$, and every negative vertex, in the encoded region and observer, has $C_g = J_g$. As g is last, its positive contribution to the exponent is exactly $|K_{\mathrm{old}}|$, changing the parameter x_i/τ into x_i. The negative exponent on old colors is unchanged. The new color contributes a nonzero monomial; in the complete encoded-and-observer rectangle each negative row has s new horizontal steps, so that monomial is $\prod_{i=1}^{k} b_i^s$. Thus the extended coefficient is this nonzero multiple of the original coefficient, for every choice of the old vertical and horizontal states.

An arbitrary original right observer may be used in these k columns. If it has fewer columns, append vacuum identity columns. If it has more, Lemma 22.1 makes the portion after the first k columns the vertical identity on the observable quotient; it merely contributes its prescribed Kronecker factors. Consequently the construction proves the identity for every original observer.

At the section after the encoded columns, the extended degree is $h(s) + s$. Left capacity gives $h(s) + s \le k$ for a nonzero coefficient. This proves (22.1) for $1 \le s \le k$; Lemma 22.1 proves it for $s \ge k$. □

22.3. **A rectangle merged into the zero-height color.**

Proposition 22.3. *The same right ramp holds when the zero-height plateau is present, with its positive-vacuum restriction.*

Proof. Fix the intermediate vertical boundary just above the chosen k negative rows, and let ν be the zero-height partition at that boundary in original, untransposed coordinates. Its negative-only filling is an ordinary semistandard tableau with entries $1, \dots, k$. Hence $\ell(\nu) \le k$ whenever its coefficient is nonzero. After row i its partial shape $\nu(i)$ has at most i rows. Its transposed Maya state therefore has no occupied site greater than i; its outgoing horizontal sites in row i are at most $i-1$.

For an integer $s > 0$, translate the positive vacuum of this one last color to charge s, with occupied sites $u \le s$. At the top of the k negative rows add the occupied sites $k+1, \dots, k+s$. They are disjoint from the old state. After subtracting the charge s, the transposed target partition is $(k^s, \nu'_0, \nu'_1, \dots)$; its untransposed partition is

$$\widetilde{\nu} = \nu + (s^k). \tag{22.4}$$

The first s full-height columns of a semistandard tableau of this shape in k letters are forced to contain $1, 2, \dots, k$ from top to bottom. Deleting them gives a tableau of ν. Conversely every entry in row i of a tableau of ν is at least i, so prefixing these forced columns respects the row inequalities. This is a bijection, including every partial shape and every original label.

We now track its physical edges. Transposition gives partial shapes $(i^s, \nu(i)')$. The charge translation cancels the row-index shift of every old transposed box. Thus every old box remains at the same physical Maya site. After negative row i, the extra vertical sites are

$$\{i+1, \dots, i+s\}; \tag{22.5}$$

the extra outgoing horizontal sites in that row are $\{i, \dots, i+s-1\}$. These sets are disjoint from the old vertical and horizontal sites, respectively. Their incoming horizontal sites are shifted one step to the right, so at every negative vertex the added occupations satisfy $C_g = J_g$. Merging them with the old zero-height paths therefore leaves every edge binary. A meeting of the two paths can give the allowed full state $A_g = J_g = C_g = L_g = 1$.

Use auxiliary positive parameters x_i/τ through site s, and x_i afterward; keep negative parameters $b_i = -x_i$. The last color is projected to its charge-s positive vacuum. This is the one-projected-color left-capacity model with its cut translated to s. At sites $1, \dots, s$ the extra positive vertical bit is one and its horizontal bits vanish. Its contribution to $\phi(K, A+I)$ on old colors is $|K_{\mathrm{old}}|$, converting x_i/τ into the original right parameter x_i. At every negative vertex, the equality $C_g = J_g$ leaves the second argument of $\phi(L, C-J)$ unchanged. The new horizontal bit is last in color order, so it contributes no further term to ϕ. Each negative row has exactly s added boxes. The auxiliary coefficient is therefore exactly

$$\left(\prod_{i=1}^{k} b_i^s\right) \times \text{the original coefficient.} \tag{22.6}$$

This is a bijection of path configurations and a pointwise equality of their weights.

At the section after site s, the added horizontal arrows occupy the rows i for which $i \le s \le i+s-1$, namely $1, \dots, \min(k, s)$. The translated left capacity gives

$$h(s) + \min(k, s) \le k.$$

This is (22.1) for every $s > 0$. The multiplier in (22.6) is nonzero over the generic field, so division is legitimate. If further variables surround the selected block, hold its boundary ν fixed, apply this inner coefficient identity, and then multiply and sum the original outside coefficients. The changed auxiliary boundary is used only inside that identity. □

22.4. **Every prefix and dominance support.**

Theorem 22.4 (All-prefix capacity)**.** *For every partition λ, every nonnegative $p \in \mathrm{RPP}(\lambda)$, and every choice of k alphabet variables, with all other variables retained,*

$$\deg_{\{x_{i_1},\dots,x_{i_k}\}} K_{\lambda,p}[X;T] \le \lambda_0 + \cdots + \lambda_{k-1}. \tag{22.7}$$

Equivalently, $K_{\lambda,p}$ is supported on m_μ with $\mu \trianglelefteq \lambda$.

Proof. Proposition 19.3 places the selected variables in the innermost k pairs. Fix the intermediate shapes on the boundary of this block for each plateau. They determine which skew subsets of the original diagram carry selected letters. Conservation fixes their total horizontal flux $h(u)$ on each Maya diagonal before any internal signs are summed. Thus the preceding coefficientwise vanishings apply to each such contribution.

Let $d_\lambda(u)$ be the number of boxes of λ with $i - j = u$. Necessarily $h(u) \le d_\lambda(u)$. The boxes on this diagonal start in row $\max(u, 0)$; hence its intersection with the first k rows has size

$$\min\{d_\lambda(u), \max(k - \max(u, 0), 0)\}.$$

Every box carrying a selected letter contributes one to precisely one such horizontal flux. The left capacity and Propositions 22.2–22.3 therefore give

$$\begin{aligned}\deg_{\text{selected}} &= \sum_u h(u) \\ &\le \sum_u \min\{d_\lambda(u), r_k(u)\} = \lambda_0 + \cdots + \lambda_{k-1}.\end{aligned}$$

This proves (22.7) with every outside variable retained. For a monomial sort its exponents in decreasing order and apply these inequalities to its k largest exponents. Fixed-p symmetry then gives all dominance inequalities $\sum_{i=1}^{k} \mu_i \le \sum_{i=1}^{k} \lambda_i$, where the latter display uses the conventional one-based notation for parts. This is exactly $\mu \trianglelefteq \lambda$. □

Corollary 22.5 (Complete one-sided triangularity)**.** *Over $\mathbf{Q}(q,t)$,*

$$F_\lambda[(t-1)X;q,t] \in \mathrm{span}\{m_\mu[X] : \mu \trianglelefteq \lambda\}. \tag{22.8}$$

The same statement holds with Schur functions in place of monomial symmetric functions.

Proof. Apply Theorem 22.4 to each finite polynomial in (18.5). Its scalar $M_\lambda(p;q,t)$ and the factor D_λ change no alphabet exponents. The height sum is valid coefficientwise in the completion and is rational by Lemma 17.3; hence the vanishing holds over $\mathbf{Q}(q,t)$. The Kostka matrix is unitriangular for dominance, so a downward dominance ideal has the same span in the monomial and Schur bases. □

23. Identification with the modified Macdonald polynomial

We combine the full two-alphabet Cauchy identity of Theorem 16.3 with the complete triangle (22.8).

Theorem 23.1. *For every partition λ,*

$$D_\lambda(q,t) G_\lambda[X;q,t] = \widetilde{H}_\lambda[X;q,t]. \tag{23.1}$$

Consequently the root vertex is

$$V_{\lambda,s}(x, \mathbf{T}; X_s) = x^{d_0(\lambda,s)} \frac{\widetilde{H}_\lambda[X_s; x^{-1}, x\mathbf{T}]}{\prod_{b \in \lambda} (1 - x^{a_b + l_b + 1} \mathbf{T}^{l_b})}. \tag{23.2}$$

Proof. All functions are over $K = \mathbf{Q}(q,t)$. The geometric two-point calculation in Theorem 16.3, with the functions of Definition 17.2, gives

$$\sum_{\lambda} R^{|\lambda|} c_\lambda F_\lambda[X]F_\lambda[Y] = \operatorname{Exp}\left[\frac{Qp_1[X]p_1[Y]}{(q-1)(1-t)}\right], \tag{23.3}$$

where

$$c_\lambda^{-1} = \prod_{b\in\lambda}(q^{a_b+1}-t^{l_b})(q^{a_b}-t^{l_b+1}). \tag{23.4}$$

The plethystic exponential on the right is the reproducing kernel for the nondegenerate pairing

$$\langle p_\rho, p_\sigma\rangle_{\mathcal{K}} = \delta_{\rho\sigma} z_\rho \prod_{r\in\rho}(q^r-1)(1-t^r). \tag{23.5}$$

Indeed its logarithm is $\sum_{r\geq 1} R^r p_r[X]p_r[Y]/(r(q^r-1)(1-t^r))$.

Fix degree n. Let A be the square matrix whose columns are the power-sum coordinates of the F_λ, and let B be the diagonal matrix of (23.5). If $C = \operatorname{diag}(c_\lambda)$, comparison of coefficients in (23.3) gives

$$ACA^{\mathsf{t}} = B^{-1}. \tag{23.6}$$

Both B and C are invertible. Taking determinants proves A invertible, so the F_λ form a basis. Inversion of (23.6) then gives $A^{\mathsf{t}}BA = C^{-1}$. Thus they are orthogonal, with nonzero norms c_λ^{-1}.

Let U be the invertible plethystic map $Uf[X] = f[(t-1)X]$. Transport the pairing by U^{-1}. Since $Up_r = (t^r-1)p_r$, the new pairing in degree n is

$$\begin{aligned}\langle p_\rho, p_\sigma\rangle_U &= \delta_{\rho\sigma} z_\rho \prod_{r\in\rho}\frac{-(1-q^r)}{1-t^r}\\ &= t^{-n}\delta_{\rho\sigma} z_\rho \prod_{r\in\rho}\frac{1-q^r}{1-t^{-r}}.\end{aligned} \tag{23.7}$$

It is a nonzero degreewise scalar multiple of the ordinary Macdonald pairing with parameters (q,t^{-1}). The vectors UF_λ are orthogonal and, by (22.8), triangular in dominance.

To verify their leading terms, order the partitions of n by any total order extending dominance, from smaller to larger. Their monomial coefficient matrix is triangular and invertible. Every diagonal coefficient $[m_\lambda]UF_\lambda$ is consequently nonzero. Divide each vector by that coefficient to obtain a monic family f_λ. Compare it with the monic Macdonald family $P_\lambda[X;q,t^{-1}]$ in (2.8). Induct in the chosen total order. The difference $f_\lambda - P_\lambda$ lies in the span of earlier monomials and therefore of earlier f_μ, whose triangular block is invertible. By induction those earlier vectors are the corresponding P_μ. Both terms in the difference are orthogonal to them. Their span has nonzero orthogonal norms and is nondegenerate, so the difference is zero. This argument includes incomparable partitions, since the total order only enlarges the collection of earlier terms.

Under the modified convention of the statement, homogeneity gives

$$U\widetilde{H}_\lambda[X;q,t] = t^{n+n(\lambda)}J_\lambda[X;q,t^{-1}]. \tag{23.8}$$

The integral form J_λ is a nonzero scalar multiple of P_λ. Therefore UF_λ and $U\widetilde{H}_\lambda$ span the same line. Inverting U gives $F_\lambda = a_\lambda \widetilde{H}_\lambda$. The one-letter calculation (17.16) gives $F_\lambda[1] = 1$, and the modified normalization gives $\widetilde{H}_\lambda[1] = 1$. Hence $a_\lambda = 1$. This proves (23.1); substituting (17.14) proves (23.2). □

24. The Chern and perverse filtrations for arbitrary flags

We establish the comparison needed for the Betti realization. Throughout this section, C is a smooth projective complex curve of genus $g \geq 0$, D is a nonempty reduced divisor, and $\mu^p = (\mu_1^p, \dots, \mu_{\ell_p}^p)$ is a composition of n at each $p \in D$. Ordering the parts amounts to ordering the parabolic weights; the final statements depend only on the associated partitions. Let $M_{\mathrm{Dol},\mu}$ be the moduli space of stable strongly parabolic Higgs bundles of rank n, fixed degree d_0, these

flag types, and numerically generic rational weights. Thus semistability coincides with stability. We work with rational cohomology.

Put

$$N_\mu = (g-1)n^2 + 1 + \frac{1}{2}\sum_{p\in D}\left(n^2 - \sum_a (\mu_a^p)^2\right). \tag{24.1}$$

When the moduli space is nonempty, it is smooth of dimension $2N_\mu$. Its Hitchin morphism is projective. All Hitchin bases below can be embedded in the same space of characteristic coefficients. A closed embedding of bases does not change the centered perverse filtration, because its direct image is perverse exact.

24.1. **Normalizations and the statement.** If $h : Y \to A$ is projective and Y is smooth of dimension d_Y, our centered convention is

$$P_j H^i(Y) = \operatorname{im}\left[H^{i-d_Y}\left(A, {}^p\tau_{\le j} Rh_* \mathbf{Q}_Y[d_Y]\right) \longrightarrow H^i(Y)\right]. \tag{24.2}$$

For a Hitchin system of relative dimension N, the unit has centered perverse degree $-N$. This convention is useful because pullback and Gysin maps over a common base have dimension shifts that cancel in the flag correspondences below.

There is a universal bundle $\mathcal{E}$ on the moduli stack times C, and there are universal graded block bundles $\mathcal{V}_{p,a}$. Choose $x \in C \setminus D$ and set

$$u = \frac{1}{n}c_1(\mathcal{E}_x), \qquad \operatorname{ch}^{\mathrm{nor}}(\mathcal{E}) = e^{-u}\operatorname{ch}(\mathcal{E}), \qquad \operatorname{ch}^{\mathrm{nor}}(\mathcal{V}_{p,a}) = e^{-u}\operatorname{ch}(\mathcal{V}_{p,a}). \tag{24.3}$$

These expressions are unchanged when the universal bundle is tensored by a line bundle from the moduli stack. They therefore define rational classes on the coarse moduli space. Equivalently, use the scalar gerbe and take the coefficient independent of its degree-two generator. The normalized Chern classes of a block are obtained from its normalized Chern character by the Newton identities.

Give

$$\int_C \operatorname{ch}_k^{\mathrm{nor}}(\mathcal{E})\,\gamma, \quad \gamma \in H^*(C), \qquad c_k^{\mathrm{nor}}(\mathcal{V}_{p,a})$$

Chern index k. Define $C_b H^*(M_{\mathrm{Dol},\mu})$ to be the span of products of these classes whose indices sum to at most b, including the empty product in index zero. Negative steps are zero. Numerical constants, even if they occur as degree-zero components of a Chern character, are assigned index zero as scalar multiples of the unit.

Theorem 24.1. *For every $g \ge 0$ and every collection of partial flags, the above normalized classes generate the cohomology, and*

$$P_j H^*(M_{\mathrm{Dol},\mu}) = C_{j+N_\mu} H^*(M_{\mathrm{Dol},\mu}). \tag{24.4}$$

The assertion holds in every degree d_0 for which the stability parameter is numerically generic, including noncoprime degrees.

We use the surface Hecke algebra and its polynomial reduction in the precise form of [26, §§4–7]. Those constructions allow every genus. Their geometric hypotheses, tautological generation, passage to zero residues, and passage to partial flags are proved below; these are needed separately from the operator construction.

24.2. **Auxiliary markings and stability.** When $M_{\mathrm{Dol},\mu}$ is empty, both sides of (24.4) vanish. We henceforth take it to be nonempty. A marking with the single block (n) and weight zero is harmless: strong parabolicity there says that the residue is zero, which says exactly that the Higgs field is regular at that point. Add such markings until $|D| \ge 3$ if $g = 0$, and $|D| \ge 1$ otherwise. Then

$$L = K_C(D), \qquad \deg L > 0. \tag{24.5}$$

The moduli problem and its Hitchin map have not changed.

Refine each block to a complete flag. Replace its repeated weight by nearby strictly ordered rational weights, chosen sufficiently close to the original ones. This can be done uniformly for all Higgs bundles of the prescribed type. Indeed, for a subrank $0 < r < n$ and subprofile ν, the difference of parabolic slopes has the form

$$\frac{e}{r} - \frac{d_0}{n} + \sum_{p,a} \alpha_{p,a} \left(\frac{\nu_{p,a}}{r} - \frac{\mu_a^p}{n} \right), \qquad e \in \mathbf{Z}.$$

There are finitely many pairs (r, ν). For a fixed pair these values are a translate of a discrete subgroup of $\mathbf{Q}$. Those sufficiently close to zero thus form a finite set; numerical genericity excludes zero. A sufficiently small weight perturbation changes none of their signs. Values farther from zero also retain their signs, since the weight contribution to the perturbation is uniformly bounded. Consequently refinement and forgetting preserve stability on the locus where each original graded block has zero induced residue.

Write X for the stable full-flag *weakly* parabolic moduli space: the Higgs residue preserves each flag. Write X_0 for the locus on which the ordered residue eigenvalues vanish. Its Higgs residues strictly lower the complete flags. If

$$N_F = (g-1)n^2 + 1 + |D| \binom{n}{2}, \tag{24.6}$$

then $\dim X_0 = 2N_F$.

24.3. **A universal sheaf and generators on a residue fiber.** We first prove the generation statement that the Hecke argument will require. Choose distinct ordered residues $\xi = (\xi_{p,i})$ with $\sum_{p,i} \xi_{p,i} = 0$. The regular spectral correspondence is realized on the surface obtained by blowing up $\mathbb{P}(L \oplus \mathcal{O}_C)$ at the n prescribed residue points above each $p \in D$. Denote this smooth projective surface by Z_ξ and its exceptional curves by $E_{p,i}$. Removing the infinity section and the strict transforms of the marked fibers gives a symplectic surface S_ξ. Its compact spectral curves are finite over C. The spectral class is

$$\Sigma = n\pi^* C_0 - \sum_{p,i} E_{p,i}, \qquad \Sigma \cdot E_{p,i} = 1. \tag{24.7}$$

Here the zero-section class is pulled back as a divisor class; the formula therefore also covers a zero residue. The blowup spectral correspondence [39, §2.2] identifies the fixed-residue Higgs moduli X_ξ with an open space of pure one-dimensional sheaves of class Σ and Euler characteristic $d_0 + n(1-g)$ on Z_ξ.

Lemma 24.2. *The space X_ξ has a projective compactification carrying a universal spectral sheaf. Its pure cohomology is generated by the curve and full-flag Chern classes. These assertions hold for every d_0.*

Proof. Let C_∞ be the infinity section and let F'_p be the strict transform of the fiber over p. If F is the pullback of the class of a point of C, choose rational $\kappa \gg 0$ and sufficiently small rational $\epsilon > 0$ and set

$$A = C_\infty + \kappa F + \epsilon \sum_{p \in D} F'_p = C_\infty + (\kappa + |D|\epsilon) F - \epsilon \sum_{p,i} E_{p,i}.$$

The ruled-surface divisor $C_\infty + \kappa F$ is ample for $\kappa \gg 0$. After pullback to its blowup, subtracting a sufficiently small positive multiple of the exceptional divisors is ample; this proves the assertion for A. Both C_∞ and the F'_p are removed from S_ξ, so $A|_{S_\xi} = \kappa F|_{S_\xi}$. Keep A fixed. If an interior spectral subsheaf has rank r over C and profile ν, then

$$\Sigma' = r\pi^* C_0 - \sum_{p,i} \nu_{p,i} E_{p,i}, \qquad A \cdot \Sigma' = \kappa r.$$

For a rational twisting class β, twisted Gieseker stability for one-dimensional sheaves compares the slopes

$$\frac{\chi(F')+\beta\cdot\Sigma'}{A\cdot\Sigma'}=\frac{1}{\kappa}\left(\frac{\deg E'}{r}+1-g+\beta\cdot\pi^*C_0+\frac{1}{r}\sum_{p,i}\alpha'_{p,i}\nu_{p,i}\right),\qquad \alpha'_{p,i}=-\beta\cdot E_{p,i}.$$

The common additive term and the positive factor cancel in every comparison. Initially take $\beta_0=\sum\alpha_{p,i}E_{p,i}$. Small changes of β preserve the entire interior stability chamber by the preceding discrete-slope argument.

There are only finitely many stability walls in a bounded neighborhood of β_0. To see this, a destabilizing effective subcycle satisfies $0<A\cdot\Sigma'<A\cdot\Sigma$. After clearing denominators in A, the possible cycles occur in finitely many projective Chow varieties of bounded degree, with finitely many numerical classes. The wall equation

$$(A\cdot\Sigma)(\chi'+\beta\cdot\Sigma')=(A\cdot\Sigma')(\chi+\beta\cdot\Sigma) \tag{24.8}$$

bounds the integer χ' for each of those classes. A same-cycle proper subsheaf has a positive-length quotient and a smaller Euler characteristic, so cannot give equality. Finally, a wall could be identically satisfied as β varies only if $\Sigma'=a\Sigma$ numerically for $0<a<1$. Intersection with any $E_{p,i}$ would give the impossible equality $\Sigma'\cdot E_{p,i}=a\in(0,1)$ between an integer and such a rational number. We may therefore choose rational β in the original chamber for which every semistable sheaf, including a boundary sheaf, is stable. The projective moduli construction for twisted Gieseker stability [39, §§2.2 and 3.2] gives a projective stable-sheaf space $\mathcal{P}_\xi$.

On its scalar gerbe, let $\mathcal{F}$ be the universal sheaf. For one exceptional divisor E, the determinant of cohomology

$$\mathcal{L}_E=\det Rp_*(\mathcal{F}\otimes^L q^*\mathcal{O}_E)$$

has scalar weight

$$\chi(F\otimes^L\mathcal{O}_E)=\Sigma\cdot E=1.$$

This follows directly from surface Riemann–Roch, since $\mathrm{ch}_0(F)=0$ and $\mathrm{ch}(\mathcal{O}_E)=1-e^{-[E]}$. It remains valid at every boundary point. Thus $\mathcal{F}\otimes p^*\mathcal{L}_E^{-1}$ descends to a universal sheaf on $\mathcal{P}_\xi\times Z_\xi$. In particular no divisibility condition on d_0 was used.

Take the closure of X_ξ in $\mathcal{P}_\xi$ and resolve that closure while leaving its smooth open part X_ξ unchanged. This gives a smooth projective compactification $\overline{X}_\xi$ of the required dimension.

We recall precisely the graph formula for sheaves that is used here [42, §4]. For stable sheaves of a fixed numerical class, with one factor supported on a symplectic open surface, the relative Hom groups off the graph vanish, as do their Serre-dual Ext^2 groups. The top Chern class of the negative relative RHom complex is the class of the graph. A proper first factor may be singular, provided the graph lies in its smooth locus; the formula then gives its Borel–Moore class and pulls back to a resolution. Here the hypothesis follows because tensoring an interior sheaf by K_{Z_ξ} does not change it, and stability kills Hom between distinct sheaves of the same reduced Hilbert polynomial. Consequently, if p_{13} integrates over Z_ξ,

$$[\Gamma]=c_{2N_F}(-Rp_{13*}\,\mathrm{RHom}(\mathcal{F}_{12},\mathcal{F}_{32}))\quad\text{ on }\overline{X}_\xi\times X_\xi. \tag{24.9}$$

Here Γ is the graph of the open inclusion $j:X_\xi\hookrightarrow\overline{X}_\xi$.

Write its Künneth expansion as $[\Gamma]=\sum_s a_s\otimes b_s$. The graph acts by restriction:

$$j^*\eta=\sum_s\left(\int_{\overline{X}_\xi}\eta a_s\right)b_s.$$

Its right-hand components therefore span the image of $H^*(\overline{X}_\xi)\to H^*(X_\xi)$, which is the lowest-weight, or pure, part of the cohomology of this smooth open variety. Grothendieck–Riemann–Roch applied to (24.9) expresses every b_s in the universal Chern components. The blowup decomposition of $H^*(Z_\xi)$ reduces these components to those from C and the exceptional divisors.

Restriction to $E_{p,i}$ is the corresponding flag quotient line. In making this identification one retains $\mathrm{ch}(\mathcal{O}_E) = 1 - e^{-[E]}$; its surface point-class terms vanish on the right factor X_ξ after multiplication by its universal spectral Chern character: this spectral sheaf vanishes at a fixed point of the removed infinity section. This proves the unnormalized generation assertion.

The universal sheaf was normalized by one flag line, so that line is trivial on the coarse space. After (24.3) its first Chern class is $-u$. Thus u is recovered from a normalized flag generator. Multiplication by e^u is a finite polynomial in each cohomological degree, proving that the normalized generators also generate the pure part. □

24.4. **Residue geometry in every genus.** Let

$$A_D = \{(\xi_{p,i}) \in \mathbf{A}^{n|D|} : \sum_{p,i} \xi_{p,i} = 0\}.$$

The ordered-residue map $r : X \to A_D$ has fiber X_ξ. The ordered-residue Hitchin base is an equivariant affine product $A_D \times A_0$, with residue coordinates of weight one. Here is the cohomological verification, including low genus. For $i \geq 2$,

$$H^1(C, L^i(-D)) = 0. \tag{24.10}$$

For $g = 0$, its degree is $i(|D| - 2) - |D| \geq -1$; for $g = 1$ it is $(i-1)|D| > 0$; for $g \geq 2$ it is greater than $2g - 2$. For $i = 1$, the residue theorem gives precisely the displayed total-trace relation. Choose linear splittings of the evaluation maps for the characteristic coefficients. Subtract from the ith coefficient the chosen lift of the degree-i elementary symmetric polynomial in the residues. These homogeneous polynomial changes of coordinates give the stated equivariant product.

The weak and fixed-residue deformation complexes, in degrees zero and one, are

$$K_{\mathrm{weak}} = [\mathrm{ParEnd}(E) \xrightarrow{[\theta,\,\cdot\,]} \mathrm{ParEnd}(E) \otimes K_C(D)],$$

$$K_{\mathrm{fix}} = [\mathrm{ParEnd}(E) \xrightarrow{[\theta,\,\cdot\,]} \mathrm{SParEnd}(E) \otimes K_C(D)].$$

Parabolic Serre duality identifies $\mathbb{H}^2(K_{\mathrm{weak}})^\vee$ with the commuting strongly parabolic endomorphisms. Stability makes a commuting parabolic endomorphism scalar; a strictly lowering scalar at a nonempty marked set is zero. Thus $\mathbb{H}^2(K_{\mathrm{weak}}) = 0$ and X is smooth. The same duality gives $\mathbb{H}^2(K_{\mathrm{fix}}) = \mathbf{C}$, dual to the scalar identity. The residue map gives an exact sequence of complexes

$$0 \longrightarrow K_{\mathrm{fix}} \longrightarrow K_{\mathrm{weak}} \longrightarrow \bigoplus_{p,i} \mathbf{C}[-1] \longrightarrow 0.$$

Its connecting map to $\mathbb{H}^2(K_{\mathrm{fix}})$ is the sum of the residue traces, since its dual includes the scalar identity into all diagonal endomorphisms. Hence the end of the long exact sequence is

$$\mathbb{H}^1(K_{\mathrm{weak}}) \longrightarrow \mathbf{C}^{n|D|} \xrightarrow{\sum} \mathbf{C} \longrightarrow 0.$$

The image of dr is therefore exactly the tangent space of A_D, so r is smooth.

If X_0 is nonempty, the image of this smooth map is an open subset of A_D containing zero and stable under positive scaling. Its complement is closed and scaling invariant. Every nonempty such complement would contain zero, by taking the limit under scaling. Thus r is surjective. The projective Hitchin morphism and its positive-weight base make X semiprojective.

We use smooth-fiber semiprojective contraction [28, Corollary 1.3.3], whose target is an affine line. Here all ordered residues have weight one. For a nonzero ξ, take the line ℓ_ξ through 0 and ξ in A_D. Its inverse image X_{ℓ_ξ} is smooth, since r is a submersion, and is a closed invariant semiprojective subvariety of X. It has the same core as X and X_0: the core is the zero Hitchin fiber, by properness and positive weights on the Hitchin base. The inclusion of the core in a smooth semiprojective variety is a homotopy equivalence. Therefore $H^*(X) \to H^*(X_{\ell_\xi})$ is an isomorphism. The affine-line fiber theorem now gives isomorphisms from this latter cohomology

to both $H^*(X_\xi)$ and $H^*(X_0)$. If A_D is a point, the assertion is immediate. All restriction maps are algebraic, so the resulting isomorphisms respect mixed Hodge structures. Thus

$$H^*(X) \xrightarrow{\sim} H^*(X_\xi), \qquad H^*(X) \xrightarrow{\sim} H^*(X_0). \tag{24.11}$$

These groups are pure, and the restrictions carry the universal curve and flag classes to the corresponding classes. Lemma 24.2 proves that all three are generated by the normalized tautological classes.

Let $U \subset X$ be the open locus of integral spectral curves with distinct marked residues. Choose ξ outside all hyperplanes

$$\sum_{p,i} \nu_{p,i}\xi_{p,i} = 0, \qquad \sum_i \nu_{p,i} = r, \quad 0 < r < n, \quad \nu_{p,i} \in \{0,1\}. \tag{24.12}$$

A proper spectral factor would give a saturated invariant subbundle of rank r and the forbidden relation by the residue theorem for its trace. A repeated factor is excluded by the distinct residues at any marking. Thus $X_\xi \subset U$. The composite

$$H^*(X) \longrightarrow H^*(U) \longrightarrow H^*(X_\xi)$$

is an isomorphism, so the first map is injective. For a smooth open inclusion, the lowest-weight part of the open cohomology is the restriction of the lowest-weight part of the original space; this also follows by choosing a common smooth projective compactification. Hence

$$H^*(X) = H^*_{\mathrm{pure}}(U) = H^*_{\mathrm{taut}}(U). \tag{24.13}$$

24.5. **The Hecke action and its reduction.** We describe the geometric inputs to the operator argument in enough detail to fix the lines, gerbe, and normalization involved. On the elliptic stack let $\mathcal{F}$ be the universal spectral sheaf. The ordered branch above (p,i) has graph $\Gamma_{p,i}$ and is smooth and unramified over p. Define

$$\mathcal{D}_{p,i} = (\Gamma_{p,i})_*\mathcal{O}, \qquad \mathcal{L}_{p,i} = \Gamma_{p,i}^*\mathcal{F}. \tag{24.14}$$

The degree-one modification is the exact sequence

$$0 \longrightarrow \mathrm{Mod}^*_{p,i}\,\mathcal{F}_{d-1} \longrightarrow \mathcal{F}_d \longrightarrow \mathcal{L}_{p,i} \otimes \mathcal{D}_{p,i} \longrightarrow 0. \tag{24.15}$$

Equivalently $\mathcal{L}_{p,i} = \mathrm{Hom}(\mathcal{F}, \mathcal{D}_{p,i})^\vee$. The skyscraper in this Hom is unweighted; this distinction is needed to obtain the eigenline.

The scalar automorphism acts on $\mathcal{L}_{p,i}$ with weight one. Choosing one branch and tensoring by its inverse splits the scalar gerbe on U and supplies a universal spectral sheaf. Its first Chern class $v = c_1(\mathcal{L}_{p,i})$ is the tautological gerbe coordinate used for polynomial reduction.

The surface for the Hecke construction is $\mathrm{Tot}\,L$. The surface Hecke operator is defined from the correspondence of colength-one inclusions $F' \subset F$. Its middle space is $\mathbb{P}(\mathcal{F})$, interpreted as the zero locus of $\mathcal{E}_1^*(1)$ in $\mathbb{P}(\mathcal{E}_0)$ for a locally free resolution $\mathcal{E}_1 \to \mathcal{E}_0 \to \mathcal{F}$. If z is the first Chern class of its quotient line, then

$$T_k(\gamma)(a) = p_{+*}\bigl(z^k p_S^*\gamma\, p_-^* a \cap [\mathbb{P}(\mathcal{F})]^{\mathrm{vir}}\bigr).$$

The tautological multiplication operators have the curve expression

$$\psi_k(\gamma) = k! \int_C (\mathrm{ch}_k(\mathcal{E}) + (1-g)\omega\,\mathrm{ch}_{k-1}(\mathcal{E}))\gamma, \qquad \gamma \in H^*(C), \tag{24.16}$$

where ω is the point class and negative Chern indices give zero. Indeed, their surface definition uses $\mathrm{ch}(\mathcal{F})\,\mathrm{Td}(S)$; Grothendieck–Riemann–Roch for the spectral projection sends this expression to $\mathrm{ch}(\mathcal{E})\,\mathrm{Td}(C)$, and $\mathrm{Td}(C) = 1+(1-g)\omega$. In particular, $\psi_1(\omega) = c_1(\mathcal{E}_x)$ and $\psi_1(1) = d_0+n(1-g)$. Multiplication by e^{-u} gives the normalized series. The nonzero leading coefficient $k!$ and the lower-index terms in (24.16) give the Chern filtration of (24.3). The degree-one exact sequence on the correspondence gives the basic relation

$$[\psi_k(\gamma), T_l(\eta)] = kT_{k+l-1}(\gamma\eta). \tag{24.17}$$

Brackets here and below are graded brackets when both operators are odd. The virtual projective-bundle formula evaluates $T_k(\gamma)1$; together with (24.17) this determines the action on the tautological ring. The surface Hecke relations of [26, §§4–5] are relations between these correspondences and their virtual classes, so pull back to the parabolic elliptic stack.

Put $y_{p,i} = c_1(\mathcal{L}_{p,i})$ and let $X_{p,i}$ denote branch modification. On its own branch the line changes by the cotangent line of the spectral curve at that branch. The unramified projection identifies it with the fixed line T_p^*C, whose first Chern class on the parameter space is zero. Other branch lines are unchanged. Thus the extended relations are

$$\begin{aligned} &[X_{p,i}, y_{p',i'}] = 0, && [\psi_k(1), X_{p,i}] = k y_{p,i}^{k-1} X_{p,i}, \\ &[\psi_k(\gamma), X_{p,i}] = 0 \quad (\deg\gamma > 0), && [y_{p,i}, T_0(1)] = X_{p,i}. \end{aligned} \tag{24.18}$$

For the last equality, equality of the moving Hecke point and the marked branch is one transverse Cartier equation on the smooth universal spectral curve near that branch. Comparing the two eigenlines gives this divisor with multiplicity one; the virtual pullback is its ordinary divisor class. This establishes the coefficient and excludes an extra multiplicity.

Let $H_{n,D}$ be the extension of $H^*(C, \mathbf{Q})$ by degree-two symbols $\tau_{p,i}$, with relations

$$\sum_{i=1}^{n} \tau_{p,i} = \omega, \qquad \tau_{p,i} H^{>0}(C) = 0, \qquad \tau_{p,i}\tau_{q,j} = 0.$$

Before reduction define

$$\psi_k(\tau_{p,i}) = y_{p,i}^k, \qquad T_k(\tau_{p,i}) = y_{p,i}^k X_{p,i}, \qquad q_1(\gamma) = T_0(\gamma).$$

These assignments respect the relations: the sum of the eigenline powers is the point component of the universal Chern character, and the analogous identity for T_k follows by decomposing the Hecke correspondence over p into its ordered branches. This is the parabolic algebra of [26, Definition 7.8 and Theorem 7.9].

For the degree comparison, let U_d denote the ordered elliptic space in degree d, over the same ordered Hitchin base, and put $R_d = H^*_{\mathrm{taut}}(U_d)$ with the normalized Chern filtration. Thus the space previously denoted by U is U_{d_0}. On the moduli stack over U_d, a scalar gerbe, write $u_d = c_1(\mathcal{E}_{d,x})/n$ and $\ell_{p,i}^{(d)} = c_1(\mathcal{L}_{p,i}^{(d)}) - u_d$.

Lemma 24.3. *The degree operator $Q_d : R_{d-1} \to R_d$ induced by $T_0(\omega)$ is an isomorphism, and*

$$Q_d C_b R_{d-1} = C_b R_d, \qquad Q_d P_j R_{d-1} = P_j R_d.$$

Moreover, $R_d = H^*_{\mathrm{pure}}(U_d)$ for every d. The normalized maps Q_d/n and their inverses preserve the unit.

Proof. Fix $p \in D$. The ordered branch $\Gamma_{p,i}$ is a relative Cartier divisor in the smooth locus of the universal spectral curve. Tensoring by $\mathcal{O}(-\Gamma_{p,i})$ gives an isomorphism $\mathrm{Mod}_{p,i}^{(d)} : U_d \to U_{d-1}$ over the Hitchin base, with inverse tensoring by $\mathcal{O}(\Gamma_{p,i})$. Integral spectral support makes all these sheaves stable. Consequently $X_{p,i}^{(d)} = (\mathrm{Mod}_{p,i}^{(d)})^*$ is a unital cohomology-ring isomorphism in the opposite direction.

The modification is supported at p, so $X_{p,i}^{(d)} u_{d-1} = u_d$. On its own branch the eigenline changes by the fixed line T_p^*C, and the other eigenlines are unchanged. Hence every normalized flag root is preserved. Pushing forward (24.15) to C gives

$$X_{p,i}^{(d)} \operatorname{ch}^{\mathrm{nor}}(\mathcal{E}_{d-1}) = \operatorname{ch}^{\mathrm{nor}}(\mathcal{E}_d) - e^{\ell_{p,i}^{(d)}\omega}.$$

Thus, for $k \geq 1$ and $a_{k,\gamma}^{(d)} = \int_C \operatorname{ch}_k^{\mathrm{nor}}(\mathcal{E}_d)\gamma$,

$$X_{p,i}^{(d)} a_{k,\gamma}^{(d-1)} = a_{k,\gamma}^{(d)} - \frac{(\ell_{p,i}^{(d)})^{k-1}}{(k-1)!} \int_C \omega\gamma. \tag{24.19}$$

The correction has Chern index at most $k-1$, including index zero when $k=1$. Since $X_{p,i}^{(d)}$ is a ring map, this proves the filtration bound for every product of generators. Its inverse has the same formula with a plus sign, and therefore $X_{p,i}^{(d)} C_b R_{d-1} = C_b R_d$. These algebraic isomorphisms also identify the pure cohomologies. The equality $R_{d_0} = H^*_{\mathrm{pure}}(U_{d_0})$ already proved in (24.13) consequently holds in every degree.

The Hecke correspondence over p is the disjoint union of its n ordered branches, each with multiplicity one. Hence

$$Q_d = \sum_{i=1}^{n} X_{p,i}^{(d)}, \qquad Q_d(1) = n\,1.$$

Fix a branch i_0. By (24.19), each ring automorphism $(X_{p,i_0}^{(d)})^{-1} X_{p,i}^{(d)}$ is the identity modulo one lower Chern step. Thus

$$(X_{p,i_0}^{(d)})^{-1} Q_d = n\,\mathrm{id} + N_d, \qquad N_d C_b R_{d-1} \subset C_{b-1} R_{d-1}.$$

The finite-dimensionality and exhaustiveness of the Chern filtration make N_d nilpotent. The finite geometric series for $(n\,\mathrm{id}+N_d)^{-1}$ shows that Q_d and Q_d^{-1} both preserve C.

Each branch modification is an isomorphism over the same base, so it strictly preserves the centered perverse filtration, with no dimension shift. Therefore $Q_d P_j \subset P_j$, while any single branch isomorphism gives $\dim P_j R_{d-1} = \dim P_j R_d$. The injectivity of Q_d makes this inclusion an equality. The assertion about the unit follows from $Q_d(1) = n\,1$. □

Here is the precise algebraic reduction input we shall use from [26, §§6–7]. For the above pure one-dimensional surface-sheaf correspondence, assume that colength-one subsheaves remain in the elliptic moduli, that a weight-one tautological line splits the scalar gerbe, and that cohomological positive-degree classes on the coarse space are nilpotent. The degree identifications used in the polynomiality construction are given on the stack by $T_0(\omega)/n$. On normalized tautological cohomology they induce the maps Q_d/n of Lemma 24.3. Weyl reduction with

$$y = \psi_1(\omega)/n = c_1(\mathcal{E}_x)/n \tag{24.20}$$

produces operators $D_{a,b}(\gamma)$ on the reduced tautological cohomology. We use $D_{a,b}(\gamma)$ for the operators denoted $\widetilde{D}_{a,b}(\gamma)_{\mathrm{red}}$ in [26, Section 6.6 and Proposition 7.5]; $q_1(\omega) = T_0(\omega)$ retains its unreduced meaning as an operator between successive degrees. Reduction replaces every curve and flag Chern series by e^{-y} times that series. The difference $y - v$ descends to the coarse space and is nilpotent, so the splitting by (24.14) gives the same reduced operators. The stack operator satisfies $T_0(\omega)(u_{d-1}a) = u_d T_0(\omega)(a)$ for every tautological class a on the source stack. Hence its quotient by $y = u$ induces the operator Q_d of Lemma 24.3. The normalized maps Q_d/n fix the unit. This is the quotient identification of reduced cohomology in [26, Section 6.6]. For the operator argument now fix

$$d_* = n(g-1), \qquad U_* = U_{d_*}, \qquad R_* = H^*_{\mathrm{pure}}(U_*) = R_{d_*}.$$

The branch isomorphisms identify U_* with U_{d_0} over the ordered Hitchin base. In this degree $\psi_1(1) = d_* + n(1-g) = 0$, and the three relations needed below are

$$\begin{aligned} [D_{2,0}(1), D_{a,b}(\gamma)] &= -2b D_{a+1,b-1}(\gamma), \\ [D_{1,1}(1), D_{a,b}(\gamma)] &= (a-b) D_{a,b}(\gamma), \\ [D_{0,2}(1), D_{a,b}(\gamma)] &= 2a D_{a-1,b+1}(\gamma). \end{aligned} \tag{24.21}$$

The extended relations (24.18) add the flag generators to this reduction. These are statements about the raw operator construction and its algebraic reduction.

All the stated hypotheses now hold. An integral spectral curve has no proper positive-rank spectral subcurve, so its rank-one torsion-free sheaves and their colength-one modifications are

stable. The weight-one line was constructed above. The tautological cohomology of the coarse space is finite dimensional and its positive-degree ideal is nilpotent. Hence

$$e = \tfrac{1}{2}D_{0,2}(1), \qquad h = -D_{1,1}(1), \qquad f = -\tfrac{1}{2}D_{2,0}(1) \tag{24.22}$$

is an $\mathfrak{sl}_2$-triple, and

$$[h, \psi_k(\gamma)] = k\psi_k(\gamma), \qquad [h, y_{p,i}] = y_{p,i} \tag{24.23}$$

after the indicated normalization. All these operators and exact commutators are now taken on R_*. We will establish the filtration comparison there, and then use Lemma 24.3 to return to degree d_0.

24.6. **Relative Lefschetz theory and low genus.** We give the filtration comparison explicitly. We use two elementary consequences of relative hard Lefschetz [26, Theorem 3.6 and Propositions 3.7–3.9]. First, a relatively ample class and the centered filtration form a Lefschetz structure, and restriction to an open base is a morphism of such structures. Second, if $a : Y \to Z$ is over a common base, both structure maps are projective and both spaces are smooth, then on pure cohomology

$$a^*P_j \subset P_{j+\dim Z-\dim Y}, \qquad a_*P_j \subset P_{j+\dim Z-\dim Y}. \tag{24.24}$$

The Gysin map retains its usual cohomological degree $2(\dim Z - \dim Y)$. These inequalities follow either from the perverse truncations of the corresponding cohomological correspondences or from their projective compactifications and the nilpotent filtration for a pulled-back ample base class.

Every Hecke correspondence above preserves the spectral curve. Applying (24.24), with the virtual-class version of [26, Proposition 3.9] for the Hecke projections, and using the degree of the inserted class gives, for $\gamma \in H^i(C)$,

$$D_{a,b}(\gamma)P_j \subset P_{j+2b-2+i}.$$

The degree maps Q_d and their inverses strictly preserve the centered perverse filtration by Lemma 24.3. Polynomial extraction and Weyl reduction therefore preserve the displayed bounds. It follows that the operators (e, h, f) on R_* in (24.22) have perverse degrees $(2, 0, -2)$.

A relatively ample class α on U_* is pure and, by Lemma 24.3, tautological. The following linear-algebra argument takes place on R_* and permits the same comparison in genus zero and one. Write $y(\gamma) = \psi_1(\gamma)$ and $x(\gamma) = D_{1,0}(\gamma)$ for positive-degree curve or flag classes. Put $\langle \gamma, \eta \rangle = \int_C \gamma\eta$ for curve classes. The reduced relations give

$$[y(\gamma), x(\eta)] = n\langle \gamma, \eta \rangle \quad (\gamma, \eta \in H^1(C)), \qquad [e, x(\gamma)] = y(\gamma). \tag{24.25}$$

The first bracket vanishes if one argument is a flag class. Operators x have h-degree -1, and operators y have degree 1. Exponential conjugation by the nilpotent flag operators adds arbitrary flag terms to e while preserving the h-filtration. For the odd curve generators, the operators $x(\gamma)y(\eta)$ generate the ordinary $\mathfrak{gl}(H^1(C))$ action. They preserve the h-filtration. The flag operators $x(\tau_{p,i})$ have perverse degree zero. For odd curve classes, $x(\gamma)y(\eta)$ has perverse degree $-1 + 1 = 0$. Their conjugations and inverse conjugations therefore also preserve P. In our normalization $\psi_2(1) = 2e$. Choose a symplectic basis $\gamma_1, \dots, \gamma_{2g}$, and put

$$y_i = y(\gamma_i), \qquad \Omega_y = \sum_{i=1}^{g} y_i y_{i+g}, \qquad J = \begin{pmatrix} 0 & I_g \\ -I_g & 0 \end{pmatrix}.$$

The odd generators satisfy

$$[\Omega_y, x(\gamma)] = -ny(\gamma), \qquad [\Omega_y, x(\gamma)y(\eta)] = -ny(\gamma)y(\eta), \qquad [e, x(\gamma)y(\eta)] = y(\gamma)y(\eta).$$

For example, the first identity follows by applying the graded product rule to each $y_i y_{i+g}$ and using (24.25); the two nonzero pairings in a symplectic pair give $-n$ times the corresponding

component of $y(\gamma)$. Thus $K = e + \Omega_y/n$ is fixed by this linear-group action, as are all flag classes. Write the pure degree-two class as

$$\alpha = aK + \sum_{i<j} B_{ij} y_i y_j + \sum_{p,i} C_{p,i} y_{p,i},$$

where B is the antisymmetric matrix with upper entries B_{ij}. Replacing α by $\alpha + \lambda\psi_2(1)$ changes its coefficients to

$$a_\lambda = a + 2\lambda, \qquad B_\lambda = B - (2\lambda/n)J.$$

For $g > 0$,

$$\det B_\lambda = (2/n)^{2g}\lambda^{2g} + \text{a polynomial of degree at most } 2g-1, \tag{24.26}$$

since $\det J = 1$. For $g = 0$, the determinant is the empty determinant 1. Thus $a_\lambda \neq 0$ and B_λ is nondegenerate outside a finite set of parameters in every genus.

Scale the original triple and conjugate by the flag operators to make its raising operator $a_\lambda e + \sum C_{p,i} y_{p,i}$. In the K decomposition its alternating part is $-(a_\lambda/n)J$. The group $\mathrm{GL}(H^1(C))$ carries this nondegenerate alternating form to B_λ, while fixing K and the flag terms. This gives a triple with raising operator $\alpha + \lambda\psi_2(1)$ and the same h-filtration. A sufficiently small permissible rational λ preserves relative ampleness. We obtain a triple (e', h', f') with $e' = \alpha + \lambda\psi_2(1)$, the same increasing h-filtration, and perverse degrees $(2, 0, -2)$.

On $\mathrm{Gr}^P R_*$, relative hard Lefschetz provides the centered grading h_P and a triple with the same e'. The operator h' has degree zero, so it commutes with h_P. We recall the needed finite-dimensional uniqueness fact: two $\mathfrak{sl}_2$-triples with the same raising operator and commuting semisimple grading operators have the same grading operator. One proof decomposes a representation into its raising-operator Jordan strings. A centralizer of the raising operator is the product of the general linear groups on equal-length primitive strings, extended by maps raising the string filtration. Conjugate grading operators differ by the latter maps; commutation with the original grading forces every positive-filtration term to vanish. The remaining product centralizes the grading. Applying this fact gives $h' = h_P$.

The unit of R_* has h-degree $-N_F$, the negative arithmetic genus of the integral spectral curve. Equation (24.23) therefore identifies its increasing eigenspace filtration with $C_{j+N_F} R_*$. We have proved

$$P_j R_* = C_{j+N_F} R_*.$$

Let $\mathcal{T}_{d_0,d_*} : R_* \to R_{d_0}$ be the composition of the normalized successive-degree maps Q_d/n, using their inverses when necessary. Lemma 24.3 gives

$$P_j R_{d_0} = \mathcal{T}_{d_0,d_*}(P_j R_*) = \mathcal{T}_{d_0,d_*}(C_{j+N_F} R_*) = C_{j+N_F} R_{d_0}.$$

Thus, for the original space $U = U_{d_0}$,

$$P_j H^*_{\mathrm{pure}}(U) = C_{j+N_F} H^*_{\mathrm{pure}}(U). \tag{24.27}$$

Choose a relatively ample class on the original space X and restrict it to U. Restriction from X is an isomorphism onto the pure part by (24.13), and is a morphism of the corresponding Lefschetz structures. Such morphisms are strict: kernels and cokernels inherit relative Lefschetz isomorphisms on their associated gradeds, so the two filtrations on the image agree. The same restriction identifies the normalized Chern generators. Hence (24.27) proves $P = C$ on the original space X in degree d_0.

24.7. **The zero-residue fiber and its centered shift.** We now check that the second restriction in (24.11) and its inverse preserve centered perversity.

Let $q = \dim A_D$, $d_X = \dim X$ and $d_0' = \dim X_0 = d_X - q$. Use the scaling compactifications

$$\overline{X} = ((X \times \mathbf{A}^1) \setminus (h^{-1}(0) \times \{0\}))/\mathbf{G}_m, \qquad \overline{X}_0 = ((X_0 \times \mathbf{A}^1) \setminus (h_0^{-1}(0) \times \{0\}))/\mathbf{G}_m.$$

Away from the excluded zero fiber the scaling stabilizers are finite. These are smooth proper orbifolds; rationally their coarse spaces are smooth for the duality and Lefschetz statements being used. The Hitchin maps extend to weighted projective bases. Write L_B and $L_{B,0}$ for the pulled-back weighted hyperplane classes. The complement of the zero Hitchin fiber in $\overline{X}$ is the orbifold line bundle over its boundary. Purity and the localization sequence imply

$$H^*(X) = H^*(\overline{X})/L_B H^*(\overline{X}), \qquad H^*(X_0) = H^*(\overline{X}_0)/L_{B,0} H^*(\overline{X}_0). \tag{24.28}$$

Indeed, restriction to the boundary is surjective by purity, and its Gysin map composed with restriction is multiplication by the boundary class L_B. The resulting short exact localization sequence gives the first quotient, and gives the second with the identical geometric argument for X_0.

The inclusion $\iota : \overline{X}_0 \hookrightarrow \overline{X}$ is the regular zero locus of the q weight-one residue sections. Consequently

$$\iota_* \iota^* = a L_B^q \quad (a \in \mathbf{Q}^\times). \tag{24.29}$$

The isomorphism modulo L_B in (24.11) shows that $\operatorname{im} \iota_* \subset \operatorname{im} L_B^q$. Write $\beta = \iota^*\gamma + L_{B,0}\beta_1$ and push forward:

$$\iota_*\beta = a L_B^q \gamma + L_B \iota_* \beta_1.$$

Repeat this lifting q times. Every term is divisible by L_B^q. It follows for every $a' \geq 0$ that

$$(\iota^*)^{-1}(\ker L_{B,0}^{a'} + \operatorname{im} L_{B,0}) \subset \ker L_B^{a'+q} + \operatorname{im} L_B. \tag{24.30}$$

Indeed, apply (24.29) to a lift of an element of the left side. The error in $\operatorname{im} L_{B,0}$ pushes forward to $\operatorname{im} L_B^{q+1}$ by the preceding inclusion, and subtracting its preimage leaves an element killed by $L_B^{a'+q}$.

On the quotient by the image of a nilpotent operator L_B, its canonical filtration at index $b \geq 0$ is represented by $\ker L_B^{b+1}$; at negative indices it is zero. The projective perverse/nilpotent-filtration formula represents $P_j H^i(X_0)$ by this filtration at index $j + d_0' - i$. Equation (24.30) replaces its kernel exponent by

$$(j + d_0' - i + 1) + q = j + d_X - i + 1,$$

which is precisely the exponent for $P_j H^i(X)$. Thus the inverse of restriction preserves centered P. It commutes with a chosen relative ample class, so it is a morphism of Lefschetz structures; strictness gives a filtered isomorphism. We obtain

$$P_j H^*(X_0) = C_{j+N_F} H^*(X_0) \tag{24.31}$$

in every genus and every numerically generic degree.

24.8. **Blockwise Gysin transfer.** Return to the original partial flags and put

$$c = \sum_{p,a} \binom{\mu_a^p}{2}, \qquad N_F = N_\mu + c. \tag{24.32}$$

Inside X_0, impose zero induced residue on each original graded block and call the resulting locus Z. It has two maps

$$X_0 \xleftarrow{\ i\ } Z \xrightarrow{\ \pi\ } M_{\mathrm{Dol},\mu}.$$

The stability perturbation already proved makes π the product of complete flag bundles in the graded block bundles; its relative dimension is c. The residue condition is the zero of the rank-c bundle of strictly lowering endomorphisms inside those blocks. Both Z and X_0 are smooth and their dimensions differ by c, so this section is regular. Its normal bundle is T_π^*.

The partial-flag moduli space $M_{\mathrm{Dol},\mu}$ is smooth semiprojective: scaling the Higgs field has a limit by properness of the Hitchin map, and its fixed locus is contained in the proper zero fiber. Its cohomology is pure by [28, Corollary 1.3.2]. The flag-bundle formula gives purity of $H^*(Z)$ as well, so the pure-cohomology bounds (24.24) apply to both maps.

Let $y_{p,a,b}$ be the normalized Chern roots of the refined block and set

$$\Delta_\mu = \prod_{p,a} \prod_{b<b'} (y_{p,a,b} - y_{p,a,b'}).$$

The Euler class of T^*_π is $\epsilon\Delta_\mu$, where $\epsilon \in \{1, -1\}$ depends only on the flag convention. Define

$$A = \pi_* i^*, \qquad B = i_* \pi^*. \tag{24.33}$$

Self-intersection and integration on the flag fibers give

$$AB = \delta_\mu \,\mathrm{id}, \qquad \delta_\mu = \pm \prod_{p,a} \mu^p_a! \neq 0. \tag{24.34}$$

In ordinary cohomological degree A lowers by $2c$ and B raises by $2c$; their algebraic Gysin maps carry the corresponding Tate twists. In centered perversity the shifts cancel. In fact

$$\dim M_{\mathrm{Dol},\mu} = 2N_\mu, \quad \dim Z = 2N_\mu + c, \quad \dim X_0 = 2N_\mu + 2c,$$

so (24.24) gives $+c$ for i^*, i_* and $-c$ for π^*, π_*. Both composites in (24.33) preserve centered P.

For tautological polynomials, B is multiplication by $\epsilon\Delta_\mu$ after lifting the symmetric block polynomial. It therefore raises Chern index by c. Conversely, flag integration is the ordinary divided difference

$$\pi_* F = \pm \frac{\mathrm{Alt}_{\prod_{p,a} S_{\mu^p_a}} F}{\Delta_\mu}, \tag{24.35}$$

expressed as a polynomial symmetric in each block. It lowers Chern index by c. Formula (24.34) also proves partial tautological generation: express Bv in full-flag generators and apply $\delta_\mu^{-1} A$.

If v has partial Chern index at most b, then Bv has full index at most $b + c$ and, by (24.31), perversity at most $b + c - N_F = b - N_\mu$. Apply A to obtain one inclusion in (24.4). Conversely, if $v \in P_j$, then $Bv \in P_j(X_0)$ has a tautological expression of index at most $j + N_F$. Applying A lowers this to $j + N_\mu$, proving the reverse inclusion. This proves Theorem 24.1 using pullback, regular Gysin, and flag integration.

25. The weight filtration at finite-order unitary parameters

We now compare the preceding normalization with the Betti mixed Hodge structure. Suppose that the distinct eigenvalues at each puncture are of finite order and have modulus one. Write them as

$$\xi_{p,a} = \exp(-2\pi i \alpha_{p,a}), \qquad \alpha_{p,a} = c_{p,a}/s_p \in \mathbf{Q} \cap [0, 1). \tag{25.1}$$

Let

$$\mathcal{C} = C[\sqrt[s_p]{p} : p \in D], \qquad \nu : \mathcal{C} \longrightarrow C$$

be the proper root curve. We use the convention that its root line $\mathcal{O}_{\mathcal{C}}(q_p)$ has character one and degree $1/s_p$. The associated orbifold bundle has parabolic degree zero, so its coarse degree is

$$d_0 = -\sum_{p,a} \mu^p_a \alpha_{p,a}. \tag{25.2}$$

The determinant condition makes this an integer.

The tame nonabelian Hodge correspondence in this situation identifies the Betti character variety $M_{B,\mu}(\xi)$ with the moduli of degree-zero stable orbifold Higgs bundles on $\mathcal{C}$, or, equivalently, the strongly parabolic Higgs moduli with the indicated weights and coarse degree (25.2). We use its topological identification of the universal bundles and of their isotropy summands. This is the bundle-compatible form of the correspondence: a harmonic bundle gives the same underlying smooth orbifold bundle on its flat and Higgs realizations, and restriction to an isotropy representation gives the stated graded block bundles. It applies to stable bundles in every genus; see [63, Section 6, Corollary 6.4] for the tame correspondence and [4] for its extension to prescribed polar data.

25.1. **An algebraic finite-cover map on the entire Betti stack.** The Betti moduli stack is a scalar $\mathbf{G}_m$-gerbe over its smooth irreducible-character space. On the stack, evaluation at an ordinary point gives an algebraic rank-n bundle. Its class $u = c_1(\mathcal{E}_x)/n$ restricts to the standard generator of $H^2(B\mathbf{G}_m)$, and is a Tate class of weight two. Rational Leray–Hirsch for this gerbe gives

$$H^*(\mathcal{M}_B, \mathbf{Q}) = H^*(M_B, \mathbf{Q})[u] \tag{25.3}$$

compatibly with mixed Hodge structures. This follows by filtering the Leray complex by powers of the global class u, whose restrictions give a basis of the cohomology of every fiber. The normalized expressions (24.3) are the elements invariant under the scalar-gerbe coaction and hence lie in the coarse factor.

If necessary, add three auxiliary root markings of nontrivial orders with trivial local monodromy and profile $(n, 0, \dots, 0)$. They do not change the moduli space. Trivial inertia makes the bundle descend there; locally, if $z = t^s$, invariant regular differentials are multiples of $t^{s-1}dt = dz/s$, so invariant Higgs fields descend to ordinary regular Higgs fields. These additions make the root curve a good orbifold. The finite-cover theorem for proper orbifold curves [2, Proposition 5.7] then supplies a representable finite *étale* map

$$f : C' \longrightarrow \mathcal{C} \tag{25.4}$$

of degree $a > 0$, with C' a smooth projective curve. In particular the restriction of a local system defines

$$R_f : \mathrm{Loc}_{\mathrm{GL}_n}(\mathcal{C}) \longrightarrow \mathrm{Loc}_{\mathrm{GL}_n}(C'). \tag{25.5}$$

This is an algebraic morphism of the entire representation stacks. Choose finite presentations of the two fundamental groups. Every generator of the covering group is a word in the base generators, so restriction is given by regular matrix-word functions on representation schemes, equivariant for simultaneous conjugation. An irreducible representation may become reducible upon restriction; the target of (25.5) includes those representations.

25.2. **The homogeneous weights of orbifold tautological classes.** We recall and explain the precise universal-class weight statement [61]. For a topological space Y with a finite model, the Betti local-system stack has universal classes

$$\int_\gamma \eta \in H^{2k-j}(\mathrm{Loc}_G(Y)), \qquad \gamma \in H_j(Y, \mathbf{Q}), \quad \eta \in H^{2k}(BG, \mathbf{Q})$$

when η is a characteristic class. They lie in

$$F^k \cap \overline{F^k} \cap W_{2k}. \tag{25.6}$$

Thus each nonzero rational such class defines a copy of $\mathbf{Q}(-k)$ in the indicated cohomology group. The statement concerns the entire local-system stack, without stability restrictions.

Here is the reason for the degree-independent weight in (25.6). Replace Y by a finite simplicial model and use the constant simplicial algebraic scheme whose components in each simplicial degree are points indexed by that model. Its cochain complex has weight zero. Evaluation of a local system on the vertices and transport on the edges gives an algebraic simplicial evaluation to BG. The Chern class on BG has type (k, k) and weight $2k$; its pullback has these same bounds in the simplicial mixed Hodge complex. Pairing with a degree-j topological cycle changes cohomological degree by j, while the cochain factor being paired has weight zero. It therefore does not change the weight or the two Hodge bounds. Simplicial descent gives (25.6). Finally, on W_{2k-1} the two bounds $p, q \geq k$ are incompatible with $p + q \leq 2k - 1$, so a nonzero rational class satisfying them is a pure Tate class of weight $2k$.

Apply this result on C' and pull back algebraically by (25.5). Rational homological transfer for (25.4) satisfies

$$f_* f^! = a \,\mathrm{id} : H_*(\mathcal{C}, \mathbf{Q}) \longrightarrow H_*(\mathcal{C}, \mathbf{Q}).$$

Naturality of the universal bundle consequently gives

$$\int_\gamma \mathrm{ch}_k(\mathcal{E}_{\mathcal{C}}) = \frac{1}{a} R_f^* \left(\int_{f^!\gamma} \mathrm{ch}_k(\mathcal{E}_{C'}) \right). \tag{25.7}$$

This proves the homogeneous Tate weight $2k$ for every orbifold universal-curve component. Only the pullback on representation stacks carries mixed Hodge structures in this argument; no transfer of a mixed Hodge structure through an analytic Riemann–Hilbert map is being used. The topological transfer supplies the equality of universal cohomology classes in (25.7).

The eigenbundle at a marking is algebraic on the fixed-monodromy stack: it is the image of the polynomial spectral projector

$$\prod_{b \neq a} \frac{X_p - \xi_{p,b}\,\mathrm{id}}{\xi_{p,a} - \xi_{p,b}}.$$

Its kth Chern class is therefore a Tate class of weight $2k$. Normalization by e^{-u} preserves this statement at total Chern index k, since $u^r \,\mathrm{ch}_{k-r}$ has weight $2r + 2(k-r) = 2k$. The normalized invariant classes descend via (25.3), which is an identity in mixed Hodge structures. All normalized orbifold and block generators on the coarse Betti space are thus homogeneous of the required weight.

25.3. **The coarse/root comparison is triangular.** The Dolbeault generators in Theorem 24.1 use the coarse curve bundle, while (25.7) uses the orbifold bundle. We compute their difference. Let $\mathcal{E}_0 = \nu_*\mathcal{E}$ and let $\mathcal{V}_{p,c}$ denote the multiplicity bundle of character c, with $0 \le c < s_p$. On a local root chart the map $\nu^*\mathcal{E}_0 \to \mathcal{E}$ is injective. For a character-c summand, its quotient has a filtration with c factors supported on the root gerbe. Rationally, the characters of a finite group have Chern character equal to their ranks and the normal character has first Chern class zero. The root gerbe has divisor class $[q_p] = \nu^*[p]/s_p$. The local filtration and the pushforward formula therefore give

$$\mathrm{ch}(\mathcal{E}) = \nu^* \,\mathrm{ch}(\mathcal{E}_0) + \sum_{p,c} \frac{c}{s_p} \nu^*[p] \,\mathrm{ch}(\mathcal{V}_{p,c}). \tag{25.8}$$

For example, the formula for the root line itself is $\mathrm{ch}(\mathcal{O}(cq_p)) = 1 + (c/s_p)\nu^*[p]$; higher curve powers vanish. The same equality holds after multiplication by e^{-u}.

Point components away from D and components paired with $H^1(C)$ are unchanged. The integrated curve component satisfies

$$\int_{\mathcal{C}} \mathrm{ch}_k^{\mathrm{nor}}(\mathcal{E}) = \int_C \mathrm{ch}_k^{\mathrm{nor}}(\mathcal{E}_0) + \sum_{p,c} \frac{c}{s_p} \,\mathrm{ch}_{k-1}^{\mathrm{nor}}(\mathcal{V}_{p,c}). \tag{25.9}$$

The right-hand correction has smaller Chern index. The two sets of generators are related by inverse triangular substitutions, and so define exactly the same increasing multiplicative filtration. This is also why individual coarse generators need not be assigned homogeneous Betti weight: for instance $\int_C c_1(\mathcal{E}_0) = d_0$ is a scalar of weight zero. In index one, (25.9) cancels this scalar against the weighted isotropy ranks, consistently with orbifold degree zero.

Proposition 25.1. *At finite-order unitary generic parameters, the Betti character variety is Hodge–Tate and*

$$C_b H^*(M_{B,\mu}) = W_{2b} H^*(M_{B,\mu}) = W_{2b+1} H^*(M_{B,\mu}). \tag{25.10}$$

Here C is transported from the normalized coarse Dolbeault generators by the bundle-compatible nonabelian Hodge identification.

Proof. Theorem 24.1 supplies generation. The inverse triangular change (25.9) replaces the generators by homogeneous Tate orbifold and block generators of weight twice their Chern indices, without changing C.

In a fixed cohomological degree only finitely many positive-degree monomials are needed. Give their free supercommutative algebra the direct-sum mixed Hodge structure in which an index-k

generator is $\mathbf{Q}(-k)$. Evaluation into cohomology is a surjective morphism of mixed Hodge structures. Strictness for W identifies its image in weights at most $2b$ with the span of monomials of index at most b. All source weights are even and all source Hodge types are diagonal. The same holds for the quotient, proving (25.10) and the Hodge–Tate assertion. This argument allows odd cohomological degrees: a class paired with $H_1(C)$ can have odd cohomological degree and still have even Tate weight. □

Theorem 25.2. *Let $g \geq 0$, let the semisimple conjugacy classes have arbitrary multiplicity partitions μ, and assume their eigenvalues are generic and of finite order on the unit circle. Under tame nonabelian Hodge theory, with degree* (25.2)*, one has*

$$P_j H^*(M_{\mathrm{Dol},\mu}) = W_{2(j+N_\mu)} H^*(M_{B,\mu}) = W_{2(j+N_\mu)+1} H^*(M_{B,\mu}). \tag{25.11}$$

The Betti cohomology is Hodge–Tate. The corresponding compactly supported assertion is obtained by Poincaré duality, including its degree shift and Tate twist by $\dim M_{B,\mu} = 2N_\mu$.

Proof. Genericity excludes a proper subobject of parabolic degree zero: the exponential of its parabolic degree would give the forbidden proper-subprofile monodromy product. Thus the weights are numerically generic for this degree. Theorem 24.1 and Proposition 25.1 apply and give (25.11). For a smooth d-dimensional space, the duality being used is

$$H^i_c(M) \simeq H^{2d-i}(M)^\vee(-d).$$

In particular a weight-w summand on the right before duality becomes weight $2d - w$, which fixes the compact-support normalization. □

26. Variation of all generic eigenvalues

For the fixed profile μ, let B_μ be the open subset of

$$T_\mu = \left\{ (\xi_{p,a}) \in \prod_{p,a} \mathbf{G}_m : \prod_{p,a} \xi_{p,a}^{\mu_a^p} = 1 \right\} \tag{26.1}$$

on which the eigenvalues at every puncture are distinct and every proper same-rank subprofile has nonunit total product. Let $f : \mathcal{X}_\mu \to B_\mu$ be the family of geometric character varieties. Genericity makes the conjugation stabilizers precisely the scalars; thus the usual representation quotient gives this algebraic smooth family. We establish constancy of the weights, separately in each cohomological degree.

Theorem 26.1. *For every i, w, the sheaf*

$$\mathrm{Gr}^W_w \, R^i f_! \mathbf{Q}_{\mathcal{X}_\mu} \tag{26.2}$$

is a local system on B_μ. Its fibers are Hodge–Tate, with locally constant Hodge numbers. Consequently the compactly supported mixed Hodge polynomial is constant on each nonempty generic part of a connected component of T_μ.

We use the fixed braid and walk construction of [50, §§5–8]. Its published local-constancy statement varies the last full-flag eigenvalues with the earlier conjugacy classes fixed. We verify here that its strata and maps vary algebraically when *all* the earlier eigenvalues vary with fixed multiplicities.

26.1. **A relative full-flag family at an identity puncture.** Add a final identity puncture. Put $T_1 = \{t \in (\mathbf{C}^*)^n : \prod_i t_i = 1\}$, and let $V \subset B_\mu \times T_1$ be the open locus where the enlarged tuple, including the ordered eigenvalues t, is generic. It contains $B_\mu \times \{(1, \dots, 1)\}$. The Grothendieck–Springer family $\mathcal{X}_{\mathrm{flag}} \to V$ is the geometric quotient of tuples

$$(b, t, A_1, B_1, \dots, A_g, B_g, X_1, \dots, X_m, Y, F),$$

where $X_p \in \mathcal{C}_p(b)$, F is a full flag, $YF_i \subset F_i$, the action of Y on F_i/F_{i-1} is t_i, and

$$\prod_{i=1}^{g}[A_i, B_i] \prod_{p=1}^{m} X_p\, Y = I_n.$$

The quotient is by simultaneous PGL_n conjugation. Write $X_{\mathrm{flag},(b,t)}$ for the fiber. At $t = (1, \dots, 1)$ the matrix Y is unipotent; its closed sublocus $Y = I_n$ is the full flag bundle over $\mathcal{X}_{\mu,b}$. The preceding semisimple matrices have fixed block centralizers

$$L_p = \prod_a \mathrm{GL}_{\mu_a^p}$$

and fixed associated parabolics. All their Bruhat index sets and unipotent dimensions are consequently independent of the parameters.

The upper triangular matrix at the last puncture remains a variable. These defining equations are regular through $t = (1, \dots, 1)$; inversion of the commutator change $u \mapsto utu^{-1}t^{-1}$ is used only on the distinct-diagonal comparison locus.

26.2. **Uniform puncture and walk strata.**

Proposition 26.2. *Let $f_{\mathrm{flag}} : \mathcal{X}_{\mathrm{flag}} \to V$ be the family just defined. It has a finite filtration $\varnothing = Z_{-1} \subset Z_0 \subset \cdots \subset Z_s = \mathcal{X}_{\mathrm{flag}}$ by relative closed subvarieties such that, for $f_a : Z_a \setminus Z_{a-1} \to V$, every $\mathcal{H}^i(Rf_{a!}\mathbf{Q})$ is a local system with constant Tate weight gradeds.*

Proof. We describe the relative coordinate check. Write $C_p(b)$ for the diagonal matrix with eigenvalue $\xi_{p,a}(b)$ in its prescribed block. On a puncture Bruhat cell, the unipotent coordinate change is

$$N_p \longrightarrow N_p, \qquad z \longmapsto C_p(b)^{-1} z C_p(b) z^{-1}, \tag{26.3}$$

where N_p is the unipotent radical between distinct blocks. Order its root coordinates by height. The output coordinate at each root equals the input coordinate times the relevant eigenvalue ratio minus one, plus a polynomial in earlier coordinates. The leading coefficient is invertible on B_μ. Successive substitution therefore gives an algebraic inverse over the entire fixed-type open. No difference within a repeated eigenvalue block is inverted. All other puncture coordinate changes use fixed permutations and unipotent factorizations, so are also algebraic over this base.

After these changes, the parameter dependence in each puncture Bruhat choice $\pi = (\pi_p)$ occurs in the target

$$C_\pi(b, t) = t^{-1} \prod_p \pi_p(C_p(b)^{-1}) \tag{26.4}$$

of a fixed braid-variety torus map. The handle charts in the braid construction involve the same fixed groups and braid words and no earlier eigenvalue parameter; they can be used unchanged as algebraic charts over V.

There is a single finite filtration by these strata over V. The first Bruhat maps take values in fixed finite Bruhat posets. For a braid word β, the walk map records the Bruhat cells of *all partial products*:

$$\rho(\beta)^0 \longrightarrow W^{\ell(\beta)+1}.$$

It is continuous for the product Bruhat order; inverse images of downward-closed subsets are closed. Its nonempty fibers are the walk strata. These maps are defined before fixing the diagonal target. Pullback along (26.4) preserves all the closed unions. Composing the two finite-poset stratifications and taking a linear extension of their order gives a finite closed filtration over V.

Integral torus splitting and quotient stabilizers.

For a fixed walk, its stratum is equivariantly an affine space times a torus. The torus map is encoded by the oriented surface Σ with marked set A associated to the walk. Its integral cocharacter map is the boundary homomorphism

$$H_1(\Sigma, A; \mathbf{Z}) \longrightarrow \ker\big(H_0(A; \mathbf{Z}) \to H_0(\Sigma; \mathbf{Z})\big). \tag{26.5}$$

The construction identifies the target with the determinant-one cocharacter lattice; fixed signs give only a fixed translate of its torus. Genericity removes a disconnected-surface stratum: each proper union of its components would give the product-one relation for a proper same-rank subprofile.

When Σ is connected, the exact homology sequence of (Σ, A) proves that (26.5) is surjective over $\mathbf{Z}$. The target is free and saturated. Choose an integral splitting for each of the finitely many walks. The corresponding torus map is therefore a product projection with a connected torus as kernel. Its pullback along any $C_\pi(b,t)$ is a product of the parameter space with this fixed kernel torus. A rational surjectivity statement would not establish this conclusion; the integral splitting excludes a finite-isogeny ambiguity.

The torus quotients in the walk construction introduce no varying finite stabilizer. Their weights are coordinate differences. On a connected walk, requiring these differences to be trivial forces equality of the involved torus coordinates, so the stabilizer group consists of the scalars. The nonzero-coordinate slices are indexed by finitely many choices independent of (b,t) and remove that scalar torus. The remaining affine factors are affine or vector bundles of fixed ranks over the product torus strata. The full-flag cover is a Zariski locally trivial Borel-modulo-scalars torsor. Quotienting its torus produces an affine bundle of rank $\binom{n}{2}$. Compactly supported direct image for such a bundle is given by its canonical trace isomorphism with the degree shift and Tate twist of its rank; global triviality of the bundle is unnecessary.

Thus every stratum of the full-flag family has a compactly supported direct-image complex whose cohomology sheaves are local systems with constant Tate weight gradeds. More explicitly, an affine bundle of rank a contributes $\mathbf{Q}(-a)[-2a]$, while a split torus of rank b contributes the constant complex $H_c^*((\mathbf{C}^*)^b, \mathbf{Q})$. These descriptions are relative algebraic descriptions, including their mixed Hodge structures. This proves the proposition. □

26.3. **Localization in mixed Hodge modules.** Apply compactly supported direct image to the successive localization triangles of Proposition 26.2. Complexes with locally constant cohomology are closed under cones: maps of local systems have locally constant kernels, images, and cokernels, and extensions of local systems are local systems. Induction therefore proves local constancy for the full-flag direct image.

For the weight assertion perform the same induction in the category of mixed Hodge modules on the smooth parameter base. A smooth mixed Hodge module is an admissible variation, shifted by the dimension of the base. Smooth mixed Hodge modules are closed under subquotients and extensions; this follows on underlying perverse sheaves from the preceding local-system observation and on Hodge filtrations from strictness. The stratum complexes have only smooth Tate graded pieces. All cohomology objects of the successive cones remain smooth, and strictness of the weight filtration gives local systems for every Gr^W. Their fibers remain Hodge–Tate. The localization morphisms, including all boundary maps, are retained throughout this argument.

26.4. **Relative Springer sign extraction.** Let $c_n = \binom{n}{2}$. The symmetric-group action on the Grothendieck–Springer direct image comes from the universal Grothendieck–Springer correspondence. Every map in that correspondence is algebraic over the preceding eigenvalue parameters. Restrict to $t = (1, \dots, 1)$. The identity-monodromy locus is the full flag bundle over $\mathcal{X}_\mu$; its inclusion and flag integration are morphisms over B_μ. The Springer sign identification [50, Proposition 8.5 and Corollary 8.11] is the resulting morphism and gives, fiberwise,

$$H_c^i(X_{\mathrm{flag},(b,1)})^{\mathrm{sgn}} \simeq H_c^{i-2c_n}(\mathcal{X}_{\mu,b})(-c_n). \tag{26.6}$$

The shift can also be read directly on the identity flag bundle: its relative top cohomology is $\mathbf{Q}(-c_n)$ in degree $2c_n$, and this is the sign representation in the Springer normalization.

Because the correspondence is obtained by base change from the universal one, (26.6) is an identification of relative direct images over B_μ, compatible with mixed Hodge modules. The

projector

$$e_{\mathrm{sgn}} = \frac{1}{n!} \sum_{\sigma \in S_n} \mathrm{sgn}(\sigma)\, \sigma$$

is a horizontal idempotent preserving W. Its image in each weight-graded local system is again a local system. After undoing the explicit shift and twist in (26.6), we obtain (26.2). Its Tate assertion follows from that for the full-flag family.

Each nonempty generic open in a component of T_μ is the complement of finitely many proper character hypersurfaces in a connected complex torus translate. It is irreducible and connected. The ranks of the local systems, and their Tate Hodge numbers, are therefore constant on it. This completes the proof of Theorem 26.1.

27. Generic torsion parameters and the required degree

The comparison at finite-order unitary parameters reaches every generic determinant component. We record the exact statement, including the numerical degree required in the Higgs realization.

Theorem 27.1. *Every connected component of T_μ meeting B_μ contains a generic finite-order unitary point. Such a point can be chosen with rational weights $\alpha_{p,a} \in [0,1)$ for which*

$$S = \sum_{p,a} \mu_a^p \alpha_{p,a} \in \mathbf{Z}, \qquad d_0 = -S, \qquad \chi_0 = n(1-g) - S, \tag{27.1}$$

and

$$\gcd(n, \chi_0, (\mu_a^p)_{p,a}) = 1. \tag{27.2}$$

More generally, in a chosen nonempty generic real weight chamber on the age hyperplane $\mu \cdot \alpha = S$, they may avoid any prescribed finite family of rational affine hyperplanes none of which contains that age hyperplane. In particular, for any fixed finite collection of active profiles, these weights may simultaneously avoid every proper parabolic-slope wall at auxiliary Euler characteristic 1.

Proof. Let d be the gcd of all positive entries of μ. The connected components of (26.1) are

$$T_{\mu,\zeta} = \left\{ \prod_{p,a} \xi_{p,a}^{\mu_a^p/d} = \zeta \right\}, \qquad \zeta \in \mu_d. \tag{27.3}$$

Indeed, the exponent vector divided by d is primitive, so its kernel is a connected torus. Each component is a torsion translate of that kernel.

If ζ has order $e < d$, the integral subprofile $\nu = (e/d)\mu$ has common rank en/d strictly between zero and n, and its eigenvalue product is $\zeta^e = 1$. Hence this component has no generic points. Thus permissible components have ζ of order d.

On a component that meets the generic open, the compact unitary torus translate is Zariski dense. Its torsion points are dense in its ordinary topology. Every excluded equal-eigenvalue or proper-subprofile condition is a proper character equation on this component. Its intersection with the compact torus has empty interior. A finite union of these proper closed subsets cannot fill that compact torus, so their complement contains a relatively open set and hence a torsion point. This gives a generic finite-order unitary point in the required component.

Choose a local logarithm branch near a generic unitary point and write $\xi_{p,a} = e^{-2\pi i \alpha_{p,a}}$. Shrink the neighborhood to preserve the ordering of the distinct weights. The determinant relation makes $S = \sum \mu_a^p \alpha_{p,a}$ a fixed integer on this neighborhood. Thus the allowed weights form a nonempty relative open in the rational affine hyperplane

$$H_S = \{\alpha : \mu \cdot \alpha = S\}. \tag{27.4}$$

Its component is fixed by $\zeta = e^{-2\pi i S/d}$, so the preceding argument implies $\gcd(S,d) = 1$. A coordinate corresponding to an auxiliary identity marking is fixed at zero and may be omitted throughout.

Consider an active profile β of common rank r and a proper subprofile β' of common rank $0 < r' < r$. At auxiliary ordinary Euler characteristic one, an equal-slope wall is

$$v \cdot \alpha = r' - r\chi', \qquad v = r\beta' - r'\beta, \quad \chi' \in \mathbf{Z}. \tag{27.5}$$

At every puncture, $\sum_a v_{p,a} = 0$. If a nonzero wall contained H_S, its normal would be a scalar multiple of μ. Taking the coordinate sum at any varying puncture would give $0 = an$, hence $a = 0$, a contradiction. If $v = 0$, the wall is empty: the equality $r\chi' = r'$ has no integral solution for $0 < r' < r$. Auxiliary identity coordinates already have v-coordinate zero, so omitting them does not affect this argument. If every original puncture has one block, all these normals vanish and all the walls are empty.

In the bounded weight box there are only finitely many integers χ' relevant to a fixed profile pair. The given finite set of active profiles therefore contributes finitely many proper affine hyperplanes in H_S. The original genericity exclusions are also finitely many affine integer hyperplanes in this box. Removing them leaves a nonempty relative open in H_S. Adding any prescribed finite family of rational affine hyperplanes that does not contain H_S preserves this conclusion. Rational points are dense in this rational affine hyperplane; choose one inside the chosen ordering chamber. It is a finite-order unitary point and has all the required simultaneous avoidance properties.

Parabolic degree zero now forces exactly (27.1). A proper equal-slope subbundle would have an integral weighted subprofile sum and hence a product-one local monodromy, contradicting genericity. If an integer $h > 1$ divided n, χ_0, and all the μ_a^p, then it would divide $S = n(1-g)-\chi_0$. The proper subprofile μ/h would have eigenvalue product $e^{-2\pi iS/h} = 1$, another contradiction. This proves (27.2). □

The auxiliary Euler characteristic one in (27.5) chooses chambers for the finite active coefficient construction. For the character variety in Theorem 2.2, positive-source degree transport is applied to χ_0 from (27.1), followed by its canonical vanishing-cycle realization. Theorem 26.1 carries the resulting weight-graded numbers to every generic parameter in that determinant component. Theorem 27.1 supplies this construction in each component.

28. Recovery of the mixed Hodge polynomial

We now combine the root-curve primitive with the punctual calculation. We retain Hodge classes until the two independent cohomological indices have been recovered. The ordinary-curve argument has already established this procedure in Section 12; we give the formulas with all puncture data present.

28.1. **The full Hodge identity.** In each generic determinant component, choose finite-order unitary eigenvalues as in Theorem 27.1, avoiding also the finite active-profile walls of Lemma 14.2. Write $\xi_{p,a} = \exp(-2\pi\mathrm{i}\alpha_{p,a})$ with $\alpha_{p,a} \in \mathbf{Q} \cap [0,1)$. Put

$$S = \sum_{p,a} \mu_a^p \alpha_{p,a} \in \mathbf{Z}, \qquad d_0 = -S, \qquad \chi_0 = n(1-g) - S, \qquad N = \frac{1}{2} d_{\boldsymbol{\mu}}.$$

Let M be the corresponding stable strongly parabolic Higgs moduli space of coarse degree d_0, and let $h : M \to \mathcal{A}$ be its proper Hitchin morphism. By Theorem 25.2, its cohomology is identified with that of the Betti variety, with the specified perverse and weight filtrations. Define

$$P_i = {}^p\mathcal{H}^i(Rh_*\mathbf{Q}_M[2N]), \qquad A_i = [R\Gamma_c(\mathcal{A}, P_i)]. \tag{28.1}$$

All sums over i are finite.

For a symmetric series in the m independent alphabets, let $\mathrm{Coeff}_{\boldsymbol{\mu}}$ mean Hall pairing with $\prod_p h_{\mu^p}[X_p]$. Equivalently, it is the coefficient of the prescribed monomial of multiplicities μ^p in each alphabet. These two descriptions agree because h_ρ and m_ρ are dual for the Hall pairing. There is no factorial in this coefficient operation.

Use R for the formal rank variable in the root stable-pair series Z_{nil}. The Fock formula of Theorem 15.4, followed by Theorem 23.1, is

$$Z_{\mathrm{nil}}(R,x;X_1,\dots,X_m)=\sum_{\lambda}R^{|\lambda|}x^{-(g-1)\|\lambda\|^2}\prod_{b\in\lambda}\zeta_C(x^{h_b}\mathbf{T}^{l_b})\prod_{p=1}^{m}\widetilde{H}_\lambda[X_p;x^{-1},x\mathbf{T}]. \tag{28.2}$$

Here ζ_C is the full Hodge-class series (12.1). The root baseline and removed-point factors have already canceled before this expression is formed. The primitive theorem, with χ_0 as above, gives

$$[R^n]\operatorname{Coeff}_{\boldsymbol{\mu}}\operatorname{Log}Z_{\mathrm{nil}}=\frac{x\mathbf{T}^{-N}}{(1-x)(1-x\mathbf{T})}\sum_i(-x)^iA_i. \tag{28.3}$$

This is Theorem 14.7, after positive-source transport to χ_0 and application of the canonical critical functor.

The coefficientwise finiteness proved for the positive source and the full root-character bounds make both sides meaningful in the same completed Grothendieck group. In any fixed rank and profile, there are finitely many possible lower-rank decompositions and perverse indices. The two nonnegative free indices in the scalar primitive have a prescribed sum in a fixed power of x.

If M is empty, the right side of (28.3) is zero. The hook substitution below then gives $\mathbb{H}_{\boldsymbol{\mu}}=0$, so the formulas hold with zero cohomology. In the separation of degrees that follows, assume M is nonempty.

28.2. **Separation of the two degrees.** The space M is smooth semiprojective. Its cohomology in degree D is pure of weight D, and smooth duality gives the same assertion for its compactly supported cohomology. The proper decomposition theorem identifies

$$H_c^j(\mathcal{A},P_i)\quad\text{as a direct summand of}\quad H_c^{2N+i+j}(M).$$

It is therefore pure of weight $2N+i+j$. If

$$F_{\boldsymbol{\mu}}(x)=x^{-1}\mathbf{T}^N(1-x)(1-x\mathbf{T})[R^n]\operatorname{Coeff}_{\boldsymbol{\mu}}\operatorname{Log}Z_{\mathrm{nil}},$$

then (28.3) yields

$$[H_c^j(\mathcal{A},P_i)]=(-1)^{i+j}\operatorname{Gr}^W_{2N+i+j}([x^i]F_{\boldsymbol{\mu}}(x)). \tag{28.4}$$

The right side determines the pure Hodge structure, since the category of pure polarizable Hodge structures of a fixed weight is semisimple. Thus this extraction retains both i and j.

The centered perverse objects in (28.1) correspond to the nonnegative perverse index k by $i=k-N$. Put

$$V_{D,k}=H^{D-N-k}(\mathcal{A},P_{k-N}),\qquad \beta_{D,k}=\dim V_{D,k}. \tag{28.5}$$

Since $\mathbb{D}P_i=P_{-i}(2N)$, Verdier duality gives

$$V_{D,k}\simeq H_c^{N+k-D}(\mathcal{A},P_{N-k})^\vee\otimes\mathbf{T}^{2N}. \tag{28.6}$$

The compact indices corresponding to (D,k) are consequently

$$i_c=N-k,\qquad j_c=N+k-D,\qquad 2N+i_c+j_c=4N-D.$$

By Theorem 25.2,

$$\beta_{D,k}=\dim\operatorname{Gr}^W_{2k}H^D(\mathcal{M}_{\boldsymbol{\mu}},\mathbf{Q}),\qquad H(\mathcal{M}_{\boldsymbol{\mu}};Q,u)=\sum_{D,k}\beta_{D,k}Q^ku^D. \tag{28.7}$$

The Hodge–Tate property follows in the same theorem from the normalized orbifold and flag Chern generators.

28.3. **The hook substitution.** Let s denote the weight character of the half-Tate object, so $\mathbf{T}$ has character s^2. Then

$$\zeta_C(y;s)=\frac{(1-sy)^{2g}}{(1-y)(1-s^2y)}.$$

Take this character only after (28.4). The substitution

$$x=z^{-2},\qquad s=zw \tag{28.8}$$

also sends the two local Macdonald parameters (x^{-1},xs^2) to (z^2,w^2). A box of arm a and leg l contributes

$$z^{-2(g-1)(2a+1)}\frac{(z^{2a+1}-w^{2l+1})^{2g}}{(z^{2a+2}-w^{2l})(z^{2a}-w^{2l+2})}.$$

The power outside the fraction cancels the baseline in (28.2), because $\sum_{b\in\lambda}(2a_b+1)=\|\lambda\|^2$. Hence the resulting series is the rank-marked version of $\Omega_{g,m}$ in (2.5).

If $E_c(P_i;s)$ is the alternating weight character of A_i, multiplying the transformed primitive identity by $(z^2-1)(1-w^2)$ gives

$$\mathbb{H}_{\boldsymbol{\mu}}(z,w)=(zw)^{-2N}\sum_i(-z^{-2})^iE_c(P_i;zw). \tag{28.9}$$

For the compact indices in (28.6), the total sign is $(-1)^{i_c+j_c}=(-1)^D$ and the weight is $4N-D$. Consequently

$$\mathbb{H}_{\boldsymbol{\mu}}(z,w)=\sum_{D,k}\beta_{D,k}(-1)^Dz^{2k-D}w^{d_{\boldsymbol{\mu}}-D}. \tag{28.10}$$

All signs in the plethystic calculations are additive signs in the lambda ring. An ordinary evaluation of the final rational function at negative z or w is made after these calculations.

28.4. **Proof of the main theorem.**

Proof of Theorem 2.2. First consider the finite-order unitary parameters fixed above. For each term of (28.10), since $d=2N$ is even,

$$(u\sqrt{Q})^d(-1)^D(\sqrt{Q})^{2k-D}\left(-\frac{1}{u\sqrt{Q}}\right)^{d-D}=Q^ku^D.$$

Thus

$$H(\mathcal{M}_{\boldsymbol{\mu}};Q,u)=(u\sqrt{Q})^d\mathbb{H}_{\boldsymbol{\mu}}\left(\sqrt{Q},-\frac{1}{u\sqrt{Q}}\right). \tag{28.11}$$

Smooth Poincare duality then gives

$$\begin{aligned}H_c(\mathcal{M}_{\boldsymbol{\mu}};Q,u)&=Q^du^{2d}H(\mathcal{M}_{\boldsymbol{\mu}};Q^{-1},u^{-1})\\&=(u\sqrt{Q})^d\mathbb{H}_{\boldsymbol{\mu}}(Q^{-1/2},-u\sqrt{Q}).\end{aligned}$$

Transposing every summation partition in (2.5) exchanges z and w, since the modified Macdonald transpose symmetry and the hook product exchange simultaneously. It follows that $\mathbb{H}_{\boldsymbol{\mu}}(z,w)=\mathbb{H}_{\boldsymbol{\mu}}(w,z)$. The hook factors and the squared Macdonald parameters also give $\mathbb{H}_{\boldsymbol{\mu}}(-z,-w)=\mathbb{H}_{\boldsymbol{\mu}}(z,w)$. This proves (2.9) at the chosen parameters.

Theorem 27.1 supplies such a generic finite-order unitary point in every nonempty generic determinant component. Theorem 26.1 makes every compactly supported weight-graded rank locally constant on that component. The same argument, or the relative Hodge–Tate stratification used there, keeps the off-diagonal Hodge numbers zero. Applying the formula separately in each component proves it for every generic tuple. The expression obtained is the same function of $g,n,\boldsymbol{\mu}$ in each case, which proves eigenvalue independence across the components as well.

Equation (28.10) also proves the polynomiality assertion. The ordinary cohomology of a smooth affine variety of complex dimension d vanishes above degree d, and its weights in degree D are at least D. Thus $D\le d$ and $2k\ge D$ whenever $\beta_{D,k}\neq0$. Both exponents in (28.10) are nonnegative. Replacing z by $-z$ removes its sign, since $(-1)^{D+2k-D}=1$. The coefficients are

the nonnegative integers $\beta_{D,k}$. This proves $\mathbb{H}_{\boldsymbol{\mu}}(-z,w) \in \mathbf{Z}_{\geq 0}[z,w]$. If the variety is nonempty, its nonzero H^0 contributes a positive multiple of w^d. No term has larger w-degree. The symmetry in z, w proves the asserted degree in both variables. For a monomial $z^a w^b$ in $\mathbb{H}_{\boldsymbol{\mu}}$, formula (2.9) has cohomological degree $d+b$ and weight $d-a+b$. Equality of these two degrees holds exactly when $a = 0$. Keeping precisely these terms gives (2.10). The empty case gives the zero polynomial by the same primitive identity. The displayed ordinary and compact formulas show directly that no choice of square root remains. □

Corollary 28.1. *The central-twist formula* (3.6) *holds for every* $g \geq 0$ *and* $n \geq 1$.

Proof. Take one puncture, $\mu^1 = (n)$, and monodromy $\zeta^{-1} I_n$, where ζ is a primitive nth root of unity. This is generic. Pairing a homogeneous symmetric function of degree n with h_n is its one-letter evaluation. Since $\widetilde{H}_\lambda[1; z^2, w^2] = 1$, the specialization of (2.6) is precisely $\overline{H}_n$. Formula (28.11) is therefore (3.6). For $g = 0, n > 1$, the defining variety is empty; the formula includes this case with value zero. □

References

[1] A. Aggarwal, A. Borodin, and M. Wheeler, *Colored fermionic vertex models and symmetric functions*, arXiv:2101.01605v2 [math.CO], 2021.

[2] K. Behrend and B. Noohi, *Uniformization of Deligne–Mumford curves*, arXiv:math/0504309 (2005), Proposition 5.7.

[3] A. A. Beilinson, J. Bernstein and P. Deligne, *Faisceaux pervers*, Astérisque **100** (1982), 5–171; Numdam digital edition.

[4] O. Biquard and P. Boalch, *Wild non-abelian Hodge theory on curves*, Compos. Math. **140** (2004), no. 1, 179–204.

[5] I. Biswas, N. L. Gammelgaard, and M. Logares, *Bohr–Sommerfeld Lagrangians of moduli spaces of Higgs bundles*, J. Geom. Phys. **94** (2015), 179–184. doi:10.1016/j.geomphys.2015.04.005.

[6] A. Borri, *Untwisting the twisted character variety via a stacky curve*, arXiv:2609.24584v1 (2026).

[7] T. Braden, *Hyperbolic localization of intersection cohomology*, Transform. Groups **8** (2003), no. 3, 209–216; arXiv:math/0202251v5.

[8] C. Brav, V. Bussi and D. Joyce, *A 'Darboux theorem' for derived schemes with shifted symplectic structure*, arXiv:1305.6302v4 (2018).

[9] C. Bu, B. Davison, A. Ibáñez Núñez, T. Kinjo and T. Pădurariu, *Cohomology of symmetric stacks*, arXiv:2502.04253v2 (2025).

[10] E. Carlsson and F. Rodriguez-Villegas, *Vertex operators and character varieties*, Adv. Math. **330** (2018), 38–60; arXiv:1603.09267v1.

[11] P.-H. Chaudouard and G. Laumon, *Un théorème du support pour la fibration de Hitchin*, Ann. Inst. Fourier **66** (2016), no. 2, 711–727. doi:10.5802/aif.3023.

[12] P.-H. Chaudouard and G. Laumon, *Addendum à "Un théorème du support pour la fibration de Hitchin"*, Ann. Inst. Fourier **67** (2017), no. 3, 1005–1008. doi:10.5802/aif.3103.

[13] W.-Y. Chuang, D.-E. Diaconescu, R. Donagi and T. Pantev, *Parabolic refined invariants and Macdonald polynomials*, Comm. Math. Phys. **335** (2015), no. 3, 1323–1379; arXiv:1311.3624.

[14] W.-Y. Chuang, D.-E. Diaconescu and G. Pan, *BPS states and the $P = W$ conjecture*, in *Moduli Spaces*, Cambridge University Press, 2014, 132–150; arXiv:1202.2039v2.

[15] B. Davison and S. Meinhardt, *Cohomological Donaldson–Thomas theory of a quiver with potential and quantum enveloping algebras*, arXiv:1601.02479v5 (2020).

[16] M. A. A. de Cataldo, J. Heinloth and L. Migliorini, *A support theorem for the Hitchin fibration: the case of* GL_n *and* K_C, arXiv:1906.09582v2 (2021).

[17] M. A. A. de Cataldo, T. Hausel and L. Migliorini, *Topology of Hitchin systems and Hodge theory of character varieties: the case* A_1, Ann. of Math. **175** (2012), no. 3, 1329–1407, doi:10.4007/annals.2012.175.3.7.

[18] M. A. A. de Cataldo, D. Maulik and J. Shen, *Hitchin fibrations, abelian surfaces, and the $P = W$ conjecture*, J. Amer. Math. Soc. **35** (2022), 911–953; arXiv:1909.11885v2.

[19] F. Déglise, F. Jin and A. A. Khan, *Fundamental classes in motivic homotopy theory*, arXiv:1805.05920v3 (2021).

[20] P. Descombes, *Hyperbolic localization in Donaldson–Thomas theory*, arXiv:2506.22400v1 (2025).

[21] A. Gitlin and D. Keating, *A vertex model for supersymmetric LLT polynomials*, arXiv:2110.10273v1 [math.CO], 2021.

[22] J. Haglund, M. Haiman, and N. Loehr, *A combinatorial formula for Macdonald polynomials*, arXiv:math/0409538v1 [math.CO], 2004.
[23] T. Hausel, E. Letellier and F. Rodriguez-Villegas, *Arithmetic harmonic analysis on character and quiver varieties*, Duke Math. J. **160** (2011), no. 2, 323–400; arXiv:0810.2076v3.
[24] T. Hausel, E. Letellier and F. Rodriguez-Villegas, *Arithmetic harmonic analysis on character and quiver varieties II*, Adv. Math. **234** (2013), 85–128; arXiv:1109.5202v1.
[25] T. Hausel, E. Letellier and F. Rodriguez-Villegas, *Positivity for Kac polynomials and DT-invariants of quivers*, Ann. of Math. **177** (2013), 1147–1168; arXiv:1204.2375v1.
[26] T. Hausel, A. Mellit, A. Minets and O. Schiffmann, $P = W$ *via* $\mathcal{H}_2$, arXiv:2209.05429v3 (2026).
[27] T. Hausel and F. Rodriguez-Villegas, *Mixed Hodge polynomials of character varieties*, with an appendix by N. M. Katz, Invent. Math. **174** (2008), 555–624; arXiv:math/0612668v3.
[28] T. Hausel and F. Rodriguez-Villegas, *Cohomology of large semiprojective hyperkähler varieties*, Astérisque **370** (2015), 113–156; arXiv:1309.4914v1.
[29] T. Hausel and M. Thaddeus, *Mirror symmetry, Langlands duality, and the Hitchin system*, Invent. Math. **153** (2003), no. 1, 197–229, doi:10.1007/s00222-003-0286-7.
[30] N. J. Hitchin, *Stable bundles and integrable systems*, Duke Math. J. **54** (1987), no. 1, 91–114.
[31] N. J. Hitchin, *The self-duality equations on a Riemann surface*, Proc. London Math. Soc. (3) **55** (1987), no. 1, 59–126, doi:10.1112/plms/s3-55.1.59.
[32] M. Hoyois, S. Scherotzke and N. Sibilla, *Higher traces, noncommutative motives, and the categorified Chern character*, Adv. Math. **309** (2017), 97–154; arXiv:1511.03589v3.
[33] B. Keller, *On the cyclic homology of exact categories*, J. Pure Appl. Algebra **136** (1999), no. 1, 1–56, doi:10.1016/S0022-4049(97)00152-7.
[34] A. A. Khan, *Intersection theory on derived stacks*, arXiv:1909.01332v2 (2026).
[35] T. Kinjo, H. Park and P. Safronov, *Cohomological Hall algebras for* 3*-Calabi–Yau categories*, arXiv:2406.12838v2 (2026).
[36] T. Kinjo, *Dimensional reduction in cohomological Donaldson–Thomas theory*, Compos. Math. **158** (2022), 123–167; arXiv:2102.01568v2.
[37] T. Kinjo and N. Koseki, *Cohomological χ-independence for Higgs bundles and Gopakumar–Vafa invariants*, J. Eur. Math. Soc. **28** (2026), 619–671; arXiv:2112.10053v2.
[38] T. Kinjo and N. Masuda, *Global critical chart for local Calabi–Yau threefolds*, Int. Math. Res. Not. **2024** (2024), no. 5, 4062–4093; arXiv:2112.10052v2.
[39] J. C. Lee and S. Lee, *Generators for the cohomology of the moduli space of irregular parabolic Higgs bundles*, arXiv:2402.05380v2 (2024).
[40] Y. Lin, *Moduli spaces of semistable pairs on projective Deligne–Mumford stacks*, Math. Res. Lett. **30** (2023), no. 4. doi:10.4310/MRL.2023.v30.n4.a7; arXiv:2006.16530v5.
[41] I. G. Macdonald, *Symmetric products of an algebraic curve*, Topology **1** (1962), no. 4, 319–343, doi:10.1016/0040-9383(62)90019-8.
[42] E. Markman, *Generators of the cohomology ring of moduli spaces of sheaves on symplectic surfaces*, J. Reine Angew. Math. **544** (2002), 61–82; arXiv:math/0009109v4.
[43] D. Maulik, J. Shen and Q. Yin, *Perverse filtrations and Fourier transforms*, Acta Math. **234** (2025), 1–69; arXiv:2308.13160v3.
[44] D. Maulik and J. Shen, *Cohomological χ-independence for moduli of one-dimensional sheaves and moduli of Higgs bundles*, Geom. Topol. **27** (2023), 1539–1586; arXiv:2012.06627v3.
[45] D. Maulik and J. Shen, *The $P = W$ conjecture for* GL_n, Ann. of Math. **200** (2024), no. 2, 529–556; arXiv:2209.02568v2.
[46] A. Mellit, *Poincaré polynomials of moduli spaces of Higgs bundles and character varieties (no punctures)*, Invent. Math. **221** (2020), 301–327; arXiv:1707.04214.
[47] A. Mellit, *Integrality of Hausel–Letellier–Villegas kernels*, Duke Math. J. **167** (2018), no. 17, 3171–3205; arXiv:1605.01299v2.
[48] A. Mellit, *Poincaré polynomials of character varieties, Macdonald polynomials and affine Springer fibers*, Ann. of Math. **192** (2020), no. 1, 165–228; arXiv:1710.04513.
[49] A. Mellit, *Cohomology rings of character varieties*, arXiv:2507.12454v1 (2025).
[50] A. Mellit, *Toric stratifications of character varieties*, Publ. Math. Inst. Hautes Études Sci. **142** (2025), 153–240, doi:10.1007/s10240-025-00158-0.
[51] L. Migliorini, V. Shende and F. Viviani, *A support theorem for Hilbert schemes of planar curves, II*, Compos. Math. **157** (2021), 835–882; arXiv:1508.07602v2.
[52] S. Monavari, *Double nested Hilbert schemes and the local stable pairs theory of curves*, Compos. Math. **158** (2022), no. 9, 1799–1849, doi:10.1112/S0010437X22007606.

[53] S. Monavari, *The refined local Donaldson–Thomas theory of curves*, Geom. Topol. **30** (2026), no. 9, 3423–3484; arXiv:2506.14359v3.
[54] S. Mozgovoy and O. Schiffmann, *Counting Higgs bundles and type A quiver bundles*, Compos. Math. **156** (2020), 744–769; arXiv:1705.04849.
[55] S. Nayak and P. Sastry, *Grothendieck duality and transitivity II: traces and residues via Verdier's isomorphism*, arXiv:1903.01783v2 (2019).
[56] T. Pădurariu and Y. Toda, *Topological K-theory of quasi-BPS categories for Higgs bundles*, arXiv:2409.10800v1 (2024).
[57] T. Pantev, B. Toën, M. Vaquié and G. Vezzosi, *Shifted symplectic structures*, arXiv:1111.3209v4 (2013).
[58] M. Saito, *Mixed Hodge modules*, Publ. Res. Inst. Math. Sci. **26** (1990), no. 2, 221–333, doi:10.2977/prims/1195171082.
[59] M. Saito, *Thom–Sebastiani theorem for Hodge modules*, arXiv:2605.05782v1 (2026).
[60] O. Schiffmann, *Indecomposable vector bundles and stable Higgs bundles over smooth projective curves*, Ann. of Math. **183** (2016), no. 1, 297–362, doi:10.4007/annals.2016.183.1.6.
[61] V. Shende, *The weights of the tautological classes of character varieties*, Int. Math. Res. Not. IMRN (2017), no. 22, 6832–6840; arXiv:1411.4975v2.
[62] C. T. Simpson, *Higgs bundles and local systems*, Publ. Math. Inst. Hautes Études Sci. **75** (1992), 5–95, doi:10.1007/BF02699491.
[63] C. T. Simpson, *Harmonic bundles on noncompact curves*, J. Amer. Math. Soc. **3** (1990), no. 3, 713–770.
[64] T. Spaide, *Shifted symplectic and Poisson structures on spaces of framed maps*, arXiv:1607.03807v1 (2016).
[65] The Stacks Project Authors, *The Stacks Project*, stacks.math.columbia.edu; Tags 08GS, 0A1G, 0G99 and 0656.
[66] J. Stoppa and R. P. Thomas, *Hilbert schemes and stable pairs: GIT and derived category wall crossings*, arXiv:0903.1444v3 (2009).
[67] T. Su, *Cell decomposition and dual boundary complexes of character varieties*, Adv. Math. **498** (2026), 111043; arXiv:2307.16657v5.
[68] R. Sulzgruber, *Inserting rim-hooks into reverse plane partitions*, arXiv:1710.09695v2 (2018).
[69] Y. Toda, *Categorical Donaldson–Thomas theory for local surfaces*, arXiv:1907.09076v4 (2021).
[70] Y. Toda, *Curve counting theories via stable objects I. DT/PT correspondence*, arXiv:0902.4371v3 (2009).
[71] S. Tubach, *Mixed Hodge modules on stacks*, Forum Math. Sigma **13** (2025), e175; arXiv:2407.02256v3.
[72] J. Wildeshaus, *Modules de Hodge mixtes à croisements normaux*, C. R. Math. Acad. Sci. Paris **337** (2003), no. 7, 467–472, doi:10.1016/j.crma.2003.08.001.

Zhili College, Tsinghua University, Beijing 100084, China
Email address: tangqy24@mails.tsinghua.edu.cn