THE HLRV CONJECTURE: A RECONSTRUCTED PROOF AND AUDIT EDITION

QIYUE TANG

Abstract. This edition reconstructs a proposed proof of the Hausel–Letellier–Rodriguez-Villegas formula for generic semisimple GL_n character varieties from the author's original materials. It replaces the moving-marking restriction step with a proof at fixed markings and makes the parameter-family comparison explicit. The accompanying audit records the exact scope of independent review, the validation of each repair, and the dependencies of every local implication.

hexagon:2611.00441v2

Contents

1. Introduction 4
2. The character-variety formula 4
3. The ordinary-curve construction 7
3.1. Notation and coefficients 7
3.2. Moduli and the two support categories 8
3.3. Established inputs 9
4. Critical models and their orientations 9
4.1. The framed critical model 9
4.2. The full deformation complex and its trace 14
4.3. Closed forms and the determinant square 21
4.4. The fixed positive quotient 24
4.5. Support gluing and the derivative 25
4.6. The comparison over the full parameter ring 28
5. Monodromic coefficients and localization 31
5.1. An explicit coefficient projector 31
5.2. The integral localization map 31
5.3. One-sided constructibility on finite-type windows 33
6. Harder–Narasimhan restrictions 34
6.1. The framed heart and its numerical bounds 34
6.2. Complete coherent flags 35
6.3. The closed-stratum counit 37
6.4. Vacuum removal and finite wall transport 38
6.5. The comparison arrow and its endpoints 39
7. The Hecke relation on the comparison arrow 42
7.1. Directed bounded presentations 42
7.2. Choices over the whole marked curve 43
7.3. The two derived words 43
7.4. A proper carrier and its supported class 46
7.5. Lifting the class to the entire arrow 47
7.6. Independence of the auxiliary presentation 48
7.7. Finite-band vanishing 50

Date: November 25, 2026; revised audit edition R8.
2020 *Mathematics Subject Classification.* Primary 14D20, 14C30; Secondary 14H60, 05E05.
Key words and phrases. Character variety, mixed Hodge polynomial, Higgs bundle, Macdonald polynomial.

8. The positive source identity 51
8.1. The unframed input and its hypotheses 52
8.2. Finite spectral addition and reduced spectral curves 53
8.3. The Hodge calculation at a nodal stratum 55
8.4. Purity and reconstruction of the source 57
9. The scalar critical primitive 58
9.1. The critical scheme on arbitrary ordinary bases 59
9.2. The normal Hessian and its determinant root 59
9.3. Pushforward and the scalar shift 60
10. Scalar packets and the global physical identity 61
10.1. Detection of the comparison error 61
10.2. The physical product over coprime packets 62
10.3. Stratified descent and induction 66
10.4. Restriction to the original nil coefficient 67
11. The fixed loci and their cohomology 67
11.1. The fixed functor 69
11.2. Closed attracting filtrations and punctual purity 70
11.3. Indices and the diagonal hook identity 71
11.4. Moving support and finite symmetric actions 73
11.5. Global purity and proof of the relative formula 75
11.6. Properness of the positive spectral map 76
11.7. The whole nil locus 76
12. Hodge recovery and the character-variety formula 77
12.1. The identity in full Hodge classes 78
12.2. Recovering perverse-graded Hodge structures 79
12.3. The hook expression 79
12.4. Proof of Theorem 3.1 80
13. Root curves and the positive comparison 80
13.1. Numerical classes and a positive length 80
13.2. A coarse positive quotient and bounded windows 82
13.3. Physical maps and their orientations on a root curve 85
13.4. The moving root Hecke correspondence 87
14. Positive root Hitchin support and canonical primitives 92
14.1. Full support at fixed markings in every genus 92
14.2. The independent unframed support input 97
14.3. Active profiles and the target Euler characteristic 98
14.4. The root and coarse surface halves 99
14.5. The critical scheme and the quadratic normal bundle 100
15. The global root Fock calculation 105
15.1. Fixed objects as finite quotients 105
15.2. Local indices, coordinate changes, and purity 106
15.3. Baseline and the global series 109
16. The football calculation and the Cauchy identity 110
16.1. Vanishing of higher-rank canonical primitives 110
16.2. The exact kernel and its signs 111
17. The punctual root vertex 112
17.1. A finite character and its root-invariant index 112
17.2. Stable alphabets and the normalized series 115
17.3. The LLT expansion and the one-letter value 115
18. Supersymmetrization at fixed paired heights 116

18.1. Signed letters and the finite polynomial 116
19. The signed lattice and its intertwiners 118
19.1. The five matrices 118
19.2. The exact finite path model 120
20. Fusion with its full null submodule 122
20.1. A rank-one identity for arbitrary incoming states 122
20.2. Moving an external row past a fused pair 123
21. The capacity of an arbitrary nested block 125
21.1. The raw embedding and its kernel 126
21.2. The zero-height boundary 127
22. The right ramp and all dominance inequalities 128
22.1. Successive disappearance on the right 129
22.2. A last auxiliary color when no zero plateau is present 129
22.3. A rectangle merged into the zero-height color 130
22.4. Every prefix and dominance support 131
23. Identification with the modified Macdonald polynomial 132
24. The Chern and perverse filtrations for arbitrary flags 133
24.1. Normalizations and the statement 134
24.2. Auxiliary markings and stability 135
24.3. A universal sheaf and generators on a residue fiber 135
24.4. Residue geometry in every genus 138
24.5. The Hecke action and its reduction 139
24.6. Relative Lefschetz theory and low genus 143
24.7. The zero-residue fiber and its centered shift 146
24.8. Blockwise Gysin transfer 146
25. The weight filtration at finite-order unitary parameters 147
25.1. The finite-order correspondence and universal classes 148
25.2. An algebraic finite-cover map on the entire Betti stack 151
25.3. The homogeneous weights of orbifold tautological classes 152
25.4. The coarse/root comparison is triangular 152
26. Variation of all generic eigenvalues 154
26.1. A relative full-flag family at an identity puncture 154
26.2. Uniform puncture and walk strata 155
26.3. Localization in mixed Hodge modules 157
26.4. Relative Springer sign extraction 158
27. Generic torsion parameters and the required degree 159
28. Recovery of the mixed Hodge polynomial 160
28.1. The full Hodge identity 161
28.2. Separation of the two degrees 161
28.3. The hook substitution 162
28.4. Proof of the main theorem 162
References 163

1. Introduction

The Hausel–Letellier–Rodriguez-Villegas conjecture predicts the compactly supported mixed Hodge polynomial of a generic semisimple GL_n character variety from a generating series of modified Macdonald polynomials [22, Conjecture 1.2.1]. The E-polynomial and Poincaré polynomial specializations are known [22, 41]; the full formula retains their joint weight and cohomological grading. Section 2 specifies the target, including arbitrary genus, rank, multiplicity profile and generic eigenvalues.

The proposed proof follows the refined stable-pair program of [15, 14]. It has three geometric inputs and one local combinatorial calculation. Their dependencies are separated below so that a local identity is not used as a substitute for a global comparison.

Ordinary-curve comparison (Sections 3–12). A sufficiently positive quotient of $K_C \oplus \mathcal{O}_C$ gives the critical comparison between two local Calabi–Yau threefolds. Sections 4–5 construct the physical model, its orientation and the coefficient formalism. Sections 6–7 construct stability-change maps and Hecke operators on the complete comparison arrow; a finite-band argument then removes its cofiber in the specified support quotient. Sections 8–12 supply the independent source support, primitive extraction, punctual calculation and global Fock assembly. The finite symmetric actions and both gradings are retained throughout.

Root geometry and the Cauchy identity (Sections 13–16). Rank, Euler characteristic and every isotropy multiplicity are retained. The positive support theorem is proved directly with the marked points fixed: the bad marked-fiber locus has one extra codimension in each integral equigeneric stratum. The resulting primitive calculation and two-pointed genus-zero pairing produce the geometric Cauchy identity.

Identification of the local vertex (Sections 17–23). The local symmetric function is defined from the deformation character. Signed lattice identities give a finite-capacity quotient for every nested block, including the zero-height boundary. This proves every dominance inequality with the other variables retained. The single triangular direction, the geometric Cauchy identity and the one-letter value identify the modified Macdonald polynomial.

Betti realization (Sections 24–28). The surface Hecke algebra of [23, Sections 4–7] supplies the operator construction; its geometric hypotheses in low genus, the zero-residue restriction and the partial-flag transfer are proved separately. Normalized universal curve and flag classes then compare centered perversity with Betti weight at finite-order unitary parameters, using [51]. Relative toric strata and Springer sign extraction [42] address variation of all generic eigenvalues. Finally, purity separates compactly supported degrees at a fixed perverse index, allowing the hook substitution and duality calculation of Section 28.

The support replacement and every other new proof expansion in this edition are recorded in the accompanying audit ledger. A conditional local check is kept distinct from verification of its upstream geometric hypotheses.

2. The character-variety formula

We prove the mixed Hodge formula of [22, 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 $\mu^p = (\mu^p_1, \dots, \mu^p_{\ell_p})$ 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

$$\prod_{p,a} \xi_{p,a}^{\mu^p_a} = 1, \qquad \prod_{p,a} \xi_{p,a}^{\nu^p_a} \neq 1 \tag{2.1}$$

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

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

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 M_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 M_p has right logarithmic image $(1 - \mathrm{Ad}_{M_p})\mathfrak{gl}_n$. The same induction shows commutation with every M_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 $X_1, \dots, X_m$ be independent alphabets. We use symmetric functions over $\mathbf{Q}(z, w)$, completed in total alphabet degree. For $X = (x_1, x_2, \dots)$, 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 $h_b = a_b + l_b + 1$. For partitions of the same size, write $\rho \le \lambda$ in dominance order when $\sum_{i=1}^{k} \rho_i \le \sum_{i=1}^{k} \lambda_i$ for every k.

For independent parameters q, t, the modified Macdonald function is

$$\widetilde{H}_\lambda[X;q,t] = t^{n(\lambda)} J_\lambda\Big[\frac{X}{1-t^{-1}};q,t^{-1}\Big], \qquad n(\lambda)=\sum_{i\ge 1}(i-1)\lambda_i. \tag{2.4}$$

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 basis with respect to dominance order and orthogonal for

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

The plethystic substitution in (2.4) is defined by $p_j[X/(1-t^{-1})] = p_j[X]/(1-t^{-j})$. In particular $\widetilde{H}_\lambda[1;q,t]=1$.

For a formal series F with constant term 1, its plethystic logarithm is

$$\operatorname{Log} F = \sum_{r\ge 1} \frac{\mu_{\mathrm{Mob}}(r)}{r} \log \psi_r(F), \qquad \psi_r(z,w,p_j[X_p]) = (z^r, w^r, p_{rj}[X_p]), \tag{2.5}$$

where μ_{Mob} is the Möbius function and log is the formal logarithm. Every coefficient of a prescribed positive alphabet degree is a finite expression. If an auxiliary rank variable R is introduced, then $\psi_r(R)=R^r$.

Define

$$\mathcal{H}^g_\lambda(z,w) = \prod_{b\in\lambda} \frac{(z^{2a_b+1}-w^{2l_b+1})^{2g}}{(z^{2a_b+2}-w^{2l_b})(z^{2a_b}-w^{2l_b+2})}, \tag{2.6}$$

$$\Omega_{g,m}(z,w;X_1,\dots,X_m) = \sum_\lambda \mathcal{H}^g_\lambda(z,w) \prod_{p=1}^m \widetilde{H}_\lambda[X_p;z^2,w^2], \tag{2.7}$$

$$\mathbb{H}_{\boldsymbol{\mu}}(z,w) = (z^2-1)(1-w^2)\left\langle \operatorname{Log}\Omega_{g,m}, \prod_{p=1}^m h_{\mu^p}[X_p]\right\rangle. \tag{2.8}$$

The empty partition contributes 1 to $\Omega_{g,m}$.

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 $H_c(Y;Q,u)$.

The primitive central-twist case is obtained by taking $m=1$, $\mu^1=(n)$, and $\xi_{1,1}$ a primitive nth root of unity.

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 every generic choice of eigenvalues.

3. The ordinary-curve construction

Let C be a smooth projective connected complex curve of genus $g \geq 0$, 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 $g \geq 1$, 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

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

We use the arm, leg, and hook lengths fixed in Section 2, and the factor $\mathcal{H}^g_\lambda(z, w)$ of (2.6). 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

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

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 $\alpha_q \in H^1(C, K_C L^{-2})$. The two threefolds

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

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 $\mathcal{A}^D_r = \bigoplus_{j=1}^{r} H^0(C, D^j)$. 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 [16]. 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 [39, Theorems 0.1 and 0.4]. The smooth-stack integrality theorem of [11, 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 [29, 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 [19]. The smooth descent and the comparison of orientations are included below. The conductor transversality criterion of [43, 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 [24, Corollary 1.3.2]. The $P = W$ theorem for coprime GL_n Higgs moduli is [40, Theorem 0.2]. The rational tautological classes have the Hodge types stated in [51]; 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 objects

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

The displayed equivalence is a full boundary framing. If a determinant trivialization is also specified, we retain its specified homotopy of agreement with the determinant of this framing. The rank-one formal identification with the fixed-determinant problem uses $R\Gamma(Y, \mathcal{O}_Y) \simeq R\Gamma(D, \mathcal{O}_D)$ and the derived Picard description near $\mathcal{O}_D$; it is not used to compare higher-rank framings. The stable-pair condition is imposed on the classical truncation. Its pure one-dimensional sheaves have locally free curve pushdown: curve torsion would generate a zero-dimensional submodule under the finite support algebra.

Here is the ambient category in which sums and exact diagrams are taken. For a connective Artin algebra A, put

$$\mathscr{F}_Y(\mathsf{A}) = \mathrm{Perf}(Y_{\mathsf{A}}) \overset{h}{\underset{\mathrm{Perf}(D_{\mathsf{A}})}{\times}} \mathrm{Perf}(C_{\mathsf{A}}). \tag{4.3}$$

An object is $(I, R, \eta : I|_D \simeq p_D^* R)$. A morphism includes the whole map $r : R_i \to R_j$ on C_{A} and a specified homotopy between its pullback and the boundary restriction. Thus its mapping complex is the homotopy pullback

$$\mathrm{RHom}_Y(I_i, I_j) \overset{h}{\underset{\mathrm{RHom}_D(I_i|_D, I_j|_D)}{\times}} \mathrm{RHom}_C(R_i, R_j).$$

This is a stable enhanced category. Its fixed-vector-bundle loci are used for moduli; they are not themselves asserted to be stable. In particular a cone has framing $\mathrm{cofib}(r)$, which can be a perfect complex. The stable-pair locus has $R = \mathcal{O}_C$.

Lemma 4.1 (Framed relative projective-line reconstruction)**.** *Let T be a derived base, $0 \to M \to V \to \mathcal{O}_T \to 0$ an extension by a line bundle, $\pi : P = \mathbf{P}_T(V) \to T$ in the lines convention, and $H = \mathcal{O}_P(\mathbf{P}_T(M))$. For a perfect object I framed by $P_0 \in \mathrm{Perf}(T)$ on $\mathbf{P}_T(M)$, set*

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

There is an equivalence, including mapping spaces, with maps of triangles

$$\begin{array}{ccccccc} P_0 & \longrightarrow & U & \longrightarrow & Q & \longrightarrow & P_0[1] \\ {\scriptstyle s}\downarrow & & \downarrow & & {\scriptstyle \mathrm{id}}\downarrow & & {\scriptstyle s[1]}\downarrow \\ G & \longrightarrow & GV^\vee & \longrightarrow & Q & \xrightarrow{c_G} & G[1]. \end{array}$$

Its inverse is

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

All these operations, their unit and counit, commute with derived base change and finite exact diagrams. On the open set where $G = 0$ they give a specified natural identification $I \simeq \pi^* P_0$, whose boundary restriction is the given framing.

Proof. The projective-bundle identities and diagonal resolution are

$$R\pi_* H = V^\vee, \qquad R\pi_* H^{-1} = 0, \qquad R\pi_* H^{-2} = M[-1], \tag{4.5}$$
$$0 \longrightarrow H_1^{-1} H_2^{-1} M^{-1} \longrightarrow \mathcal{O}_{P \times_T P} \longrightarrow \mathcal{O}_{\Delta_P} \longrightarrow 0.$$

The first arrow is the wedge of the two tautological lines. Applying this finite perfect kernel to IH^{-1} and tensoring by H gives (4.4). Adjunction gives the middle arrow $U \to GV^\vee$. Since $H|_{\mathbf{P}_T(M)} = M^{-1}$, restriction gives $U \to Q$, with fiber P_0. Conversely, push the displayed fiber down after tensoring by H^{-1} and by $H^{-2}M^{-1}$: the results are respectively $G[-1]$ and $U[-1]$.

These are the unit and counit of the diagonal kernel, not merely identifications of deformation groups. In cone coordinates, with the specified homotopy included, the upper extension and its map are

$$e_{\mathrm{map}} \in \mathrm{RHom}_T(Q, P_0)^1, \quad \varphi \in \mathrm{RHom}_T(Q, G)^0, \qquad d\varphi + s e_{\mathrm{map}} = c_G, \quad d e_{\mathrm{map}} = 0. \tag{4.6}$$

The connecting cochain uses the following cone convention. If a graded splitting of an extension has differential

$$d_B = \begin{pmatrix} d_A & b \\ 0 & d_C \end{pmatrix}, \qquad i(a) = (a, 0),$$

then $\mathrm{Cone}(i) = B \oplus A[1]$ has differential $\left(\begin{smallmatrix} d_B & i \\ 0 & -d_A \end{smallmatrix}\right)$. The chain equivalence

$$C \longrightarrow \mathrm{Cone}(i), \qquad c \longmapsto ((0, c), -bc)$$

identifies its projection to $A[1]$ with $-b$, while retaining the original inclusion $A \to B$ and quotient $B \to C$. For the signed cocycle in (4.1), dual local lifts differ by $-c_{ij}y$. The lower row therefore has off-diagonal Cech differential $-c_{ij}y$ and connecting cochain $c_G = +c_{ij}y$. This fixes the sign of the cotangent section without changing the framing, inclusion, or quotient. The formulas are constructions in the stable enhancement; in particular their maps and all higher cells are retained. The finite projective-space Cech calculation, or the projective-bundle theorem over derived bases, gives (4.5) after every derived base change. Projection formula and the same perfect diagonal kernel give the base-change identifications of its unit and counit. If $G = 0$, the upper triangle identifies U with P_0; this proves the last assertion with its stated boundary normalization. □

Compactify the torsor over S by $Z = \mathbf{P}_S(V)$, where $0 \to K_S \to V \to \mathcal{O}_S \to 0$ has class $\alpha_q y$, and put $H = \mathcal{O}_Z(D_\infty)$. Let $\mathscr{F}_Z^{\mathrm{fin}}(\mathsf{A})$ be the full stable subcategory of

$$\mathrm{Perf}(Z_{\mathsf{A}}) \underset{\mathrm{Perf}(S_{\mathsf{A}})}{\overset{h}{\times}} \mathrm{Perf}(C_{\mathsf{A}})$$

on which $G = R\pi_*(IH^{-1})[1]$ has support finite over C_{A}. Support here means the support on the classical truncation of a perfect complex. The framing maps r still come from C_{A}, not from arbitrary maps on S_{A}.

Lemma 4.2 (The two compactifications with full framing)**.** *There are natural inverse exact equivalences*

$$\Phi : \mathscr{F}_Y(\mathsf{A}) \simeq \mathscr{F}_Z^{\mathrm{fin}}(\mathsf{A}) : \Omega.$$

They preserve the entire boundary framing, arbitrary maps of framing complexes, and their higher homotopies. They commute with derived base change, finite sums and cones, and are natural for isomorphisms of the curve, the extension, and the framing data.

Proof. Write $A = K_C L^{-1}$, $B = \mathbf{P}_C(A) \subset D$, and $\bar{p} : \bar{S} = \mathbf{P}_C(\mathcal{O}_C \oplus L) \to C$. Let $T = \mathcal{O}_{\bar{S}}(-1)$ be the tautological line and $\Sigma = \mathbf{P}_C(L) = \bar{S} - S$. The inverse image of T under $\bar{p}^*(\mathcal{O}_C \oplus N) \to \bar{p}^*(\mathcal{O}_C \oplus L)$ is a bundle $\mathcal{E}$ with

$$0 \longrightarrow \bar{p}^* A \longrightarrow \mathcal{E} \longrightarrow T \longrightarrow 0.$$

The incidence morphism

$$b : W = \mathbf{P}_{\bar{S}}(\mathcal{E}) \longrightarrow Y, \qquad q : W \longrightarrow \bar{S}$$

is the blow-up of B. Indeed, in a splitting chart the ideal of B is (z_0, z_1) and the incidence equation is $z_0 v_1 = z_1 v_0$ in $Y \times \mathbf{P}^1$; these are the two usual blow-up charts. The description uses the quotient map itself and therefore glues without a choice of splitting. Its exceptional section $E_\infty = \mathbf{P}_{\bar{S}}(\bar{p}^* A)$ is $\bar{S}$, and its strict boundary transform $D' = q^{-1}\Sigma$ is D, with $b|_{D'} = \mathrm{id}_D$. Moreover $E_\infty \cap D' = B$ and $b^* D = E_\infty + D'$ as Cartier divisors. Put

$$\bar{V} = \mathcal{E} \otimes T^{-1}, \qquad M = \bar{p}^* A \otimes T^{-1}, \qquad \bar{H} = \mathcal{O}_W(E_\infty) = \mathcal{O}_{\mathbf{P}(\mathcal{E})}(1) \otimes q^* T.$$

Thus $0 \to M \to \bar{V} \to \mathcal{O}_{\bar{S}} \to 0$ and $\bar{H} = \mathcal{O}_{\mathbf{P}(\bar{V})}(1)$. On S the tautological line is trivialized by $(1, y)$, so $W|_S = Z$, $M|_S = K_S$, and $\bar{H}|_Z = H$. In particular no untwisted identification of $\mathcal{O}_{\mathbf{P}(\mathcal{E})}(1)$ with $\bar{H}$ is being made.

First compare Y with W. On W keep a framing by the same R on the whole divisor $E_\infty + D'$. Equivalently, for this fixed perfect object, keep its framings on E_∞, D' and their specified agreement on B. This equivalence of framing data follows by applying derived Hom from the perfect object to

$$0 \longrightarrow \mathcal{O}_{E_\infty + D'} \longrightarrow \mathcal{O}_{E_\infty} \oplus \mathcal{O}_{D'} \longrightarrow \mathcal{O}_B \longrightarrow 0;$$

it does not assert descent of arbitrary perfect objects along a closed cover. Pullback by b gives such a framed object $J = b^*I$.

Conversely let J have this full framing. The center $B_{\mathsf{A}} \to Y_{\mathsf{A}}$ is quasi-smooth of virtual codimension two. The perfect-complex blow-up theorem [28, Theorem C and §4.4] has precisely the two blocks

$$b^* \operatorname{Perf}(Y_{\mathsf{A}}), \qquad i_*(p_E^* \operatorname{Perf}(B_{\mathsf{A}}) \otimes \ell), \qquad \ell = \mathcal{O}_{E_\infty}(E_\infty) = \bar{p}^* A^{-1} \otimes T,$$

where $i : E_\infty \to W$ and $p_E : E_\infty \to B = C$. The order means that maps from the first block to the second vanish. Since $i^*J = p_E^*R$, for every $T_0 \in \operatorname{Perf}(B_{\mathsf{A}})$ we have

$$\operatorname{RHom}_W\big(J, i_*(p_E^*T_0 \otimes \ell)\big) = \operatorname{RHom}_{E_\infty}(p_E^*R, p_E^*T_0 \otimes \ell) = 0,$$

because $Rp_{E*}\ell = 0$. The two-block decomposition therefore places J in the first block. Its inverse there is the actual functor Rb_*J, and its counit $b^*Rb_*J \to J$ is an equivalence; the unit is induced by $\mathcal{O}_Y \simeq Rb_*\mathcal{O}_W$. These are statements about enhanced perfect categories over the derived base, not a classical essential-image assertion followed by an unproved extension to Artin algebras.

Restrict the counit to $D' \simeq D$ to frame Rb_*J on D. This recovers the specified exceptional framing too: both identifications of $J|_{E_\infty}$ with p_E^*R agree on B, and restriction from E_∞ to its section B induces an equivalence on the mapping spaces between pullbacks of perfect complexes on C, since $Rp_{E*}\mathcal{O}_{E_\infty} = \mathcal{O}_C$. Their agreement, including its higher homotopies, is therefore unique in the relevant homotopy fiber. The exact sequence above then supplies the full union framing.

Next compare this framed category on W with Z. For J on W put $\bar{G} = Rq_*(J\bar{H}^{-1})[1]$. It is perfect. Derived proper base change to Σ gives

$$\bar{G}|_\Sigma = Rp_{D*}(p_D^*R \otimes \mathcal{O}_D(-1) \otimes L^{-1})[1] = 0,$$

where $\bar{H}|_{D'} = \mathcal{O}_D(1) \otimes L$. The support of the perfect complex $\bar{G}$ is thus disjoint from Σ. It is proper over C_{A}, being closed in $\bar{S}$, and is contained in the affine C_{A}-scheme S; its classical support is consequently finite over C_{A}. Lemma 4.1 canonically identifies J with $q^*\bar{p}^*R$ away from this support.

For the inverse, start with (I, R, η) on Z and let K be the support of G. Its image in $\bar{S}$ is closed and misses Σ: a finite support over C is proper, and $\bar{S} \to C$ is separated. On $Z - \pi^{-1}K$ the same lemma supplies a natural isomorphism $I \simeq \pi^*p^*R$, inducing η on the boundary. Glue I on Z to $q^*\bar{p}^*R$ on $q^{-1}(\bar{S} - K)$ using this isomorphism. Denote the resulting perfect object by $\operatorname{Ext}(I)$. This is ordinary open descent in the enhanced category: explicitly it is the Mayer–Vietoris homotopy fiber of the two open direct images mapping to their intersection. It glues all maps and higher cells, not just isomorphism classes. The original framing on $E_\infty \cap Z$ glues to R near Σ, and D' has its fixed framing; their agreement on B gives the framing on $E_\infty + D'$.

Restriction and Ext are inverse by the same open descent unit and counit, using the canonical trivialization where $\bar{G} = 0$. This is an inverse also with the original D' framing: the framing on D' induced by that trivialization and the given one agree on B, because both agree there with the E_∞ framing. The restriction

$$\operatorname{RHom}_{D'}(p_D^*R_i, p_D^*R_j) \longrightarrow \operatorname{RHom}_B(R_i, R_j)$$

is an equivalence, since $Rp_{D*}\mathcal{O}_{D'} = \mathcal{O}_C$ and B is a section. Hence their compatible homotopy has contractible space of choices, and the union-framing comparison follows from the same exact sequence. Consequently

$$\Phi(I) = b^*I|_Z, \qquad \Omega(I) = Rb_* \operatorname{Ext}(I)$$

are the claimed inverse functors. Their two composites are the specified open-descent equivalences followed by the unit or counit of (b^*, Rb_*); the adjunction and descent triangle identities give their coherent inverse data.

The geometric construction is unchanged by derived extension of A. The maps b, q are proper, locally of finite presentation, and of finite Tor-amplitude; their direct images preserve perfect complexes and satisfy derived base change [28, §2.1.5 and §4.3.1]. The projective-line kernel has the same property. One may use a larger common finite support when making a base change or a finite diagram; refinement of the open cover gives the same extension with its canonical coherent identification. This proves naturality also when support shrinks. All functors are exact. Finite unions of finite supports remain finite, so finite sums, cones, and exact diagrams stay in the indicated categories. The construction is intrinsic in the extension $N \to L$, hence glues on curve charts and is equivariant for their automorphisms. □

Proposition 4.3. *For a fixed framing bundle R, the boundary-framed physical enhancement is naturally equivalent, on the indicated locus, to* $\operatorname{Crit}(\mathcal{P}_L^R, g_q)$*, where $\mathcal{P}_L^R$ is the derived stack of pairs $p^*R \to G$ and G has support finite over C. For $R = \mathcal{O}_C$ this is* $\operatorname{Crit}(\mathscr{P}_L, g_q)$*. The equivalence is defined over connective Artin algebras. Its ambient enhanced comparison respects maps, their higher homotopies, direct sums, and finite exact diagrams in* (4.3)*. For finite cones, the framing and surface data are allowed to be perfect complexes.*

Proof. Apply Lemma 4.2, followed by Lemma 4.1 with $M = K_S$ and $P_0 = p^*R$. Thus the actual comparison from the physical category is restriction of b^* followed by the displayed diagonal kernel; its inverse is that kernel's inverse followed by Rb_* Ext. Denote the physical-to-critical direction by Ψ. In particular both directions already have specified enhanced inverse data before considering their tangent complexes.

Put $J_G = \operatorname{fib}(p^*R \to G)$. With the framing fixed the tangent of $\mathcal{P}_L^R$ is $\operatorname{RHom}_S(J_G, G)$. Surface duality with proper support identifies its shifted cotangent with $\operatorname{RHom}_S(GK_S^{-1}, J_G[1])$. Hence (4.6), with its specified homotopy, is the derived zero stack of the cotangent section $(c_G, 0)$. For the monic resolution

$$P_G = [p^*EL^{-1} \xrightarrow{y-\Theta} p^*E]$$

the degree-minus-one identity supplies the homotopy between multiplication by y and by Θ. The torsor cocycle $\alpha_q y$ consequently represents the full covector $\operatorname{Tr}_C(\alpha_q \operatorname{Tr}(\Theta\, d\Theta))$. This includes bundle directions, since $\operatorname{Tr}(\Theta[b, \Theta]) = 0$, and is dg_q. The potential is unchanged when the source framing is R.

In the physical local coordinates and the cone convention above, the dictionary is

$$\psi = -X, \qquad e = -\varepsilon, \qquad \rho_{\mathrm{cot}} = -\rho_{\mathrm{phys}}. \tag{4.7}$$

Here e denotes the canonical cotangent coordinate, whereas e_{map} in (4.6) is the actual connecting cochain. In the unsigned monic entry it has coefficient $\mu = -e$; the cochain and this coefficient are different types. Lemma 4.5 identifies the actual section and universal projective-plane monad. Its explicit actual-tangent map (4.12) and the separate normal, section, and gauge evaluations (4.15)–(4.16) fix all three canonical coordinates before checking the moment relation. That calculation uses only the two preceding reconstruction lemmas, the displayed matrices, and Serre evaluation, independently of this proposition and of the later physical trace comparison. The indicated sheaf and stability conditions are imposed on the common classical truncation; the corresponding derived open loci are therefore the same. This proves the assertion with the exact and base-change compatibilities already constructed. □

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 (4.7).

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

$$\begin{aligned} \alpha &= (z_0 X - z_2,\ -(z_0 \Theta - z_1),\ z_0^2 \varepsilon), \\ \beta &= (z_0 \Theta - z_1,\ z_0 X - z_2,\ s), \qquad \gamma = -z_0^2 \rho : M^{-1} \longrightarrow M^1. \end{aligned} \tag{4.8}$$

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 (4.8). Its differential is

$$\begin{aligned} D(g, r) &= (g\Theta_i - \Theta_j g,\ gX_i - X_j g,\ gs_i - s_j r,\ r\varepsilon_i - \varepsilon_j g) - g\rho_i + \rho_j g, \\ D(a, b, u, v) &= aX_i - X_j a + \Theta_j b - b\Theta_i + u\varepsilon_i + s_j v, \qquad Dw = 0. \end{aligned} \tag{4.9}$$

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 (4.8), 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

$$(-z_0 b, z_0 a, -z_0^2 v), \qquad (-z_0 a, -z_0 b, -u).$$

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

$$(a_2, b_2, u_2, v_2)(a_1, b_1, u_1, v_1) = w(-a_2 b_1 + b_2 a_1 - u_2 v_1). \tag{4.10}$$

Degree-zero products and their actions are ordinary rectangular composition. Equations (4.9) and (4.10) follow by multiplication of the displayed matrices. Thus κ is a morphism of differential graded categories.

Lemma 4.4. *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
$i_*V_i[-1]$, $i_*V_j[-1]$	$\operatorname{Hom}(V_i,V_j)\otimes\bigwedge^k W \quad (k=0,1,2)$
$\mathcal{O}R_i$, $i_*V_j[-1]$	$\operatorname{Hom}(R_i,V_j) \quad (k=1)$
$i_*V_i[-1]$, $\mathcal{O}R_j$	$\operatorname{Hom}(V_i,R_j)\Delta \quad (k=1)$
$\mathcal{O}R_i$, $\mathcal{O}R_j$	$\operatorname{Hom}(R_i,R_j) \quad (k=0).$

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. □

Lemma 4.5 (Identification with the universal monad)**.** *The inverse of the geometric comparison just defined is represented on every curve-and-bundle chart by* (4.8)*. Its enhanced morphism map is the displayed κ, including rectangular maps and the framing component r. Thus the derivative of Ψ^{-1} is $\kappa[1]$ on the relative tangent models, and the derivative of Ψ is its inverse.*

Proof. It suffices to perform the following calculation over the universal coefficient algebra, before any stability restriction. Trivialize the two normal line bundles and write $t = y-\Theta$. On $Z = W|_S$ the blow-up map is

$$[u:v]\longmapsto[z_0:z_1:z_2]=[u:uy:v].$$

The pullback of the monad therefore has arrows

$$\alpha=(uX-v,\ ut,\ u^2\varepsilon),\quad \beta=(-ut,\ uX-v,\ s),\quad \gamma=-u^2\rho.$$

Use the ordered Cech representatives $1/(uv)$ in $H^1\mathcal{O}(-2)$, $(1/(u^2v),1/(uv^2))$ in $H^1\mathcal{O}(-3)$, and the analogous ordered three representatives in $H^1\mathcal{O}(-4)$. All relevant negative line bundles have only this cohomology, or are acyclic. Thus these give actual functorial pushforward complexes, with the total signs of the coefficient–matrix convention preceding (4.8).

The raw complex for $G = R\pi_*(MH^{-1})[1]$ is

$$K_1:\quad E^2\longrightarrow E^2,\qquad (a,b)\longmapsto(Xa-b,ta).$$

Its contraction to $P_G=[E\xrightarrow{t}E]$ is

$$p_1(a,b)=a,\quad p_1(P,Q)=Q,\qquad i_1(a)=(a,Xa),\quad i_1(Q)=(0,Q),\quad h_1(P,Q)=(0,-P).$$

For $U = R\pi_*(MH^{-2})[1]$ the raw complex is $K_2: E^3\to E^4\oplus R\to E$, with

$$d_1(a,b,c)=(Xa-b,Xb-c,ta,tb,\varepsilon a),$$
$$d_2(P,Q,T_0,T_1,r)=-tP+XT_0-T_1+sr,\qquad \gamma_2(a,b,c)=-\rho a.$$

It contracts to $U_0 = [E \xrightarrow{(t,\varepsilon)} E \oplus R]$. Here p_2 retains a, T_0, r; $i_2(a)$ has lower components (a, Xa, X^2a) and the additional degree-minus-one coefficient ρa in the T_1 slot;

$$\begin{aligned} i_2(T_0, r) &= (0, 0, T_0, XT_0 + sr, r), \\ h_2(P, Q, T_0, T_1, r) &= (0, -P, -XP - Q), \\ h_2(z) &= (0, 0, 0, -z, 0). \end{aligned}$$

These formulas satisfy $p_k i_k = 1$ and $dh_k + h_k d = 1 - i_k p_k$ in the total complex. For example the uncanceled T_1 entry in $di_2 - i_2 d$ is $[t, X] - s\varepsilon = -d\rho$; the coefficient differential of the displayed ρ entry cancels it. Its composite with $-T_1$ cancels γ_2. Thus the contraction includes the odd coordinate, and does not assume that X commutes with t.

Multiplication by u and v from K_2 to K_1 has lower components respectively $(a, b, c) \mapsto (a, b)$ and $(a, b, c) \mapsto (b, c)$, and middle components respectively $(P, Q, T_0, T_1, r) \mapsto (P, T_0)$ and (Q, T_1). Put $\bar{u} = p_1 u_* i_2$ and $\bar{v} = p_1 v_* i_2$. Then $\bar{u}$ is the projection $U_0 \to P_G$, whereas $\bar{v}$ has X on the lower term, (X, s) on the upper term, and the coefficient $+\rho$ from the lower to the upper term. Their comparison with the raw maps retains the homotopy

$$p_1 u_* = \bar{u} p_2, \qquad dh_v + h_v d = p_1 v_* - \bar{v} p_2, \qquad h_v = p_1 v_* h_2.$$

Explicitly $h_v(P, Q, T_0, T_1, r) = -P$ and $h_v(z) = -z$, in the lower and upper terms of P_G, respectively. These identities follow from $dh_2 + h_2 d = 1 - i_2 p_2$ and include the odd ρ term. The diagonal equation $u_1 v_2 - v_1 u_2$ gives the reconstruction arrow $v_2 u_* - u_2 v_*$, not the coordinatewise sum of these maps.

Normalize the identification of G with its monic resolution by $-p_1$, and the Q part of U by $-p_2$, leaving the framing R unchanged. The resulting triangle has

$$U = [E \xrightarrow{(t,-\varepsilon)} E \oplus R], \qquad f_v : U \longrightarrow Q \text{ the projection},$$

and $f_u : U \to G$ has lower component $-X$, upper component $(-X, s)$, and lower-to-upper coefficient $-\rho$. These formulas give the same map of triangles as Lemma 4.1, and its boundary fiber is the original R, with its original framing.

This normalization is fixed by the actual pair section. Indeed, when $\varepsilon = \rho = 0$ and the two fields commute, let K_F be the two-variable spectral Koszul resolution of F, with last augmentation $E \to F$ equal to the identity and differentials

$$(X - z, -(\Theta - y)), \qquad (\Theta - y, X - z).$$

The chain isomorphism from the affine monad to $\mathrm{fib}(R \xrightarrow{s} F)$ is $(-1, +1, +1, -1)$ on $(\text{lower}, \text{normal middle}, R, \text{top})$. It fixes the R projection. In the standard fiber differential, the R-to-top entry is $-s$, exactly the image of the monad entry $+s$ under this isomorphism. Its last component forces the minus sign in the monic identification of G above. In particular for $E = R = \mathcal{O}$, $s = 1$, and $\Theta = X = 0$, the result is the actual pair $\mathcal{O} \to \mathcal{O}_{(y,z)=(0,0)}$ with section $+1$ and boundary framing $+1$.

We now fix all three cotangent coordinates by the actual tangent map and Serre evaluation, independently of the physical trace comparison below. Keep the monic augmentation just fixed and write $A = K_C L^{-1}$, $P = P_G$, and $Q = PA^{-1}$. With pullbacks suppressed, the relevant terms and their cohomological degrees are

	degree -1	degree 0
P	EL^{-1}	E
Q	EK_C^{-1}	EA^{-1}

$d_P = d_Q = t.$

The specified section has the fiber model

$$\begin{aligned} J &= \mathrm{fib}(R \xrightarrow{s} P) = P[-1] \oplus R, \\ J^0 &= EL^{-1} \oplus R, \qquad J^1 = E, \qquad d_J = (-t, -s). \end{aligned} \tag{4.11}$$

Thus $J[1]$ has degrees $(-1,0)$ and differential (t,s). Projection $J \to R$ is the identity on the given framing. For homogeneous coefficient-free maps we use $d_{\mathrm{Hom}}h = d_P h - (-1)^{|h|}hd_J$ and $h[1] = (-1)^{|h|}h$. Coefficients are then inserted by the coefficient–matrix Koszul rule stated before (4.8). In particular a coefficient's degree is not substituted for the matrix degree in that rule.

Here is the map from actual infinitesimal parameters to $\mathrm{RHom}_S(J,P)$. For a variation $(\dot\Theta, \dot s)$ its degree-zero representative is

$$\begin{aligned} h_{\dot\Theta,\dot s}|_{J^0} &= (+\dot\Theta, -\dot s) : EL^{-1} \oplus R \longrightarrow E, \\ h_{\dot\Theta,\dot s}|_{J^1} &= 0. \end{aligned} \tag{4.12}$$

The first sign is the monic Kodaira–Spencer sign: over square-zero τ, the relation $(t - \tau\dot\Theta)\widetilde{v} = 0$ gives $t\widetilde{v} = \tau\dot\Theta v$ in the kernel of reduction, so the connecting class of this extension has entry $+\dot\Theta$ with the fixed augmentation. The second sign comes from differentiating the actual cross differential $-s$ in (4.11). The degree-minus-one gauge representative h_g has entries

$$h_g|_{J^0} = (g,0) : EL^{-1} \oplus R \to EL^{-1}, \qquad h_g|_{J^1} = g : E \to E.$$

Its differential, computed in these slots, is

$$d_{\mathrm{Hom}}h_g|_{J^0} = (tg - gt, -gs) = ([g,\Theta], -gs). \tag{4.13}$$

This is exactly the image under (4.12) of the actual action $(\dot\Theta, \dot s) = ([g,\Theta], gs)$ of $1 + \tau g$. It checks both signs with the original section and framing.

For comparison with the strict surface Hom model $\mathcal{B}$, its raw generators are g in degree zero and (a,u) in degree one, with $Dg = ([g,\Theta], gs)$. If σ denotes suspension, the map from $\mathcal{B}[1]$ to the monic Hom model is

$$\sigma g \longmapsto h_g, \qquad \sigma(a,u) \longmapsto h_{-a,-u}, \quad h_{-a,-u}|_{J^0} = (-a, +u).$$

Indeed $d_{\mathcal{B}[1]} = -D$, and its image is (4.13). Thus raw (a,u) and actual $(\dot\Theta, \dot s)$ are not being identified without signs. The monic elimination computes this surface model on the curve; these formulas give its representatives in $\mathrm{RHom}_S(J,P)$.

We next specify the unsigned entries used in the evaluation. For $\mu : E \to RK_C$, let $j_R(\mu) : Q^{-1} \to R$ be the coefficient μ in the degree-one Hom slot; it is not a cotangent coordinate. For $T : E \to EK_C$, let $j_G(T) : Q^{-1} \to P^0$ be the analogous degree-one slot. For $m : E \to EA$ let $\mathrm{diag}(m)$ act on both terms of $Q \to P$. A degree-zero cotangent cochain $\zeta = (\varphi, e_{\mathrm{map}}) : Q \to J[1]$ has representatives

$$\begin{aligned} \varphi &= \mathrm{diag}(m) + j_G(\lambda), \qquad e_{\mathrm{map}} = j_R(\mu), \\ |m| &= |\mu| = 0, \qquad |\lambda| = -1. \end{aligned} \tag{4.14}$$

Here $j_G(\lambda)$ has matrix degree one and coefficient degree minus one, hence total degree zero. Also $j_R(\mu)$ has total degree one as a cochain into R, but degree zero as the component into $R[1]$. Globally m takes values in $\mathrm{End}(E)A$, μ in $\mathrm{Hom}(E,R)K_C$, and λ in $\mathrm{End}(E)K_C$ with the indicated coefficient degree.

The unsigned monic Serre class $j_G(T)$ has trace $+\mathrm{Tr}_E(T)$. To fix the sign, use $[\mathcal{O}(-1) \xrightarrow{x_1} \mathcal{O}]$ in degrees $(-1,0)$ on the ordered cover $(x_0 \neq 0, x_1 \neq 0)$. The local degree-one entry $1/x_0$ on the first open and zero on the second has Cech difference $-1/x_0$. Its correction in Cech degree one and End degree zero is $\mathrm{diag}(-1/(x_0x_1), 0)$, whose commutator is $-1/x_0$. Thus it is closed for $\delta + (-1)^p[d,-]$. The supertrace of the correction is $+1/(x_0x_1)$, with ordered Serre trace $+1$ [45, Propositions 4.2.2 and 5.2.3]. For $t = y - \Theta$, the same calculation takes the y^{-1} coefficient of $\mathrm{Tr}_E(T(y-\Theta)^{-1})$, which is $\mathrm{Tr}_E(T)$. No commutation or diagonalization of the matrices is used. Composition with rectangular section maps is performed before this trace. Tensoring this finite perfect calculation with the coefficient algebra gives its derived-base version.

Write $\mathcal{E}_{\mathrm{Serre}}(h,\zeta)$ for the local K_C-valued Serre evaluation in tangent-first order: compose $h[1]\zeta : Q \to P[1]$ and take this normal trace. Curve Serre trace is applied afterwards. For the

actual degree-zero variation in (4.12), the lower-to-upper entry is $\dot\Theta m - \dot s\mu$, so

$$\begin{aligned}\mathcal{E}_{\mathrm{Serre}}(h_{\dot\Theta,\dot s},\zeta) &= \mathrm{Tr}_E(m\dot\Theta) - \mathrm{Tr}_R(\mu\dot s),\\ \mathcal{E}_{\mathrm{can}}(h_{\dot\Theta,\dot s},(\psi,e)) &= \mathrm{Tr}_E(\psi\dot\Theta) + \mathrm{Tr}_R(e\dot s).\end{aligned} \tag{4.15}$$

The second line is the canonical cotangent evaluation on the actual parameters, equivalently the even-coordinate part of the Liouville construction. The two independent blocks therefore give $m = \psi$ and $\mu = -e$. Notice that the trace in the first line contains the matrix coefficient μ, not the cochain e_{map}; the latter has the different type in (4.14).

The gauge block requires a separate suspension calculation. Take coefficient-free g, T first. The gauge map h_g has Hom degree -1 and $h_g[1] = -h_g$. The dual generator $j_G(T)$ has Hom degree $+1$ in $\mathrm{RHom}(Q, J[1])$: it goes from Q^{-1} to $(J[1])^0 = P^0$. Only its composite with the upper $-g$ component of $h_g[1]$ survives in the normal trace. Thus

$$\mathcal{E}_{\mathrm{Serre}}(h_g, j_G(T)) = -\mathrm{Tr}_E(Tg) = \mathcal{E}_{\mathrm{can}}(h_g, T^\vee). \tag{4.16}$$

Here $T^\vee$ denotes the gauge-dual generator of degree $+1$, normalized by $T^\vee(g) = \mathrm{Tr}_E(Tg)$. The last equality is canonical graded evaluation: in tangent-first order its symmetry factor is $(-1)^{(-1)(+1)} = -1$. Equivalently, in cotangent-first order both homogeneous evaluations are positive. Consequently the cross generator $j_G(T)$ is this canonical dual generator, not its negative. A point of the cotangent total space assigns to a gauge generator of degree -1 a coefficient of degree -1; hence the coefficient λ in (4.14) is ρ_{cot}.

This comparison also keeps track of odd coefficients. For example, if $|\xi| = -1$, then in coefficient-first matrix notation

$$\begin{aligned}h_g[1](\xi \otimes j_G(T)) &= (-1)^{(-1)(-1)}\xi \otimes (-h_g j_G(T)),\\ \mathcal{E}_{\mathrm{Serre}}(h_g, \xi \otimes j_G(T)) &= +\xi\,\mathrm{Tr}_E(Tg).\end{aligned}$$

The canonical evaluation has the identical sign, since moving ξ past the degree-minus-one first input contributes the same minus. Thus the homogeneous minus in (4.16) is not applied a second time to the odd coordinate. More generally graded tensor evaluation and this same coefficient–matrix rule extend the calculation to homogeneous entries of perfect framing complexes, their shifts, and rectangular blocks. This is also the graded evaluation in the local Liouville construction [47, Definition 1.20 and the proof of Proposition 1.21]; the Artin-stack cotangent construction is as in [29, Example 3.10].

For clarity, $\mathcal{E}$ is dual evaluation in the stated tangent-first order, whereas the retained raw-pairing convention uses the canonical two-form in cotangent-first order. Their signs are related as follows. If a base tangent generator has degree a, its unshifted generator has degree $a+1$ and the raw dual generator has degree $1-a$. On coefficient-free generators the cotangent-first value of the canonical form is $(-1)^a$ times dual evaluation; this is the tensor-shift sign in the Liouville construction. In particular it is positive on the normal and section tangent generators of degree zero, and negative on the gauge tangent generator of degree minus one. With raw generators u, v in base and dual order, the same suspension convention is

$$\begin{aligned}\omega(\sigma v, \sigma u) &= (-1)^{|v|}q(v,u),\\ q(v,u) &= (-1)^{1-a}(-1)^a\langle v,u\rangle = -\langle v,u\rangle.\end{aligned}$$

Here $|v| = 1-a$ is the raw generator degree; it is not the degree of a coefficient used to make a total cochain, and $\langle v,u\rangle$ is unsigned dual-on-base evaluation, not the tangent-first $\mathcal{E}$. This recovers the cotangent-first raw convention of the later trace comparison, after the actual generator map has been fixed. Arbitrary coefficients are then handled once by the Koszul rule, as in the ξ calculation above; no second reversal of the odd coefficient is made.

For the later cotangent determinant diagram we retain precisely [29, Appendix B.1, (B.2)]:

$$\beta_n : (F[n])^\vee \xrightarrow{\sim} F^\vee[-n], \qquad \beta_n = (-1)^n\alpha_n, \tag{4.17}$$

where α_n is the rotation-compatible comparison defined there. This dual/shift convention is distinct from the suspension of a homogeneous Hom map used above. In particular the inverse of $h^\vee[-1]$ in the later paired cotangent triangle is formed with this same β_n on both sides; no usual dual/shift convention is substituted on one side. The evaluations just computed precede that determinant operation.

We may now, after the three evaluations, write the dual-cone diagram in canonical coordinates as

$$U = [E \xrightarrow{(t,e)} E \oplus R], \qquad f_u = (\psi;\ (\psi, s);\ \rho_{\mathrm{cot}}), \qquad f_v = \mathrm{pr}_Q.$$

The twists are those of Q above. The standard cone connecting cochain is $-j_R(e)$, and the chain-map equation is $d\rho_{\mathrm{cot}} = [\Theta, \psi] + se$. Our actual monad image has U-entry $-\varepsilon$, diagonal $-X$, and cross coefficient $-\rho_{\mathrm{phys}}$. Hence, with all three types displayed,

$$e_{\mathrm{map}} = j_R(\varepsilon), \quad \mu = \varepsilon, \quad e = -\varepsilon, \qquad \psi = -X, \quad \rho_{\mathrm{cot}} = -\rho_{\mathrm{phys}}.$$

This proves (4.7) for the actual geometric functor. As a subsequent differential check, its cochain $\varphi = \mathrm{diag}(-X) + j_G(-\rho_{\mathrm{phys}})$ gives

$$d\varphi = j_G([\Theta, X] - d\rho_{\mathrm{phys}}) = -j_G(s\varepsilon),$$
$$se_{\mathrm{map}} = j_G(s\varepsilon).$$

On a pure odd algebra with $\Theta = X = s = \varepsilon = 0$ and closed ρ_{phys}, the only cross coefficient is $-\rho_{\mathrm{phys}}$. This checks the already normalized dictionary; it is not used to identify that coefficient as a canonical one.

Here is the morphism check, without replacing a rectangular map by a diagonal infinitesimal variation. General homogeneous cochains are checked by the retained dg differential and product calculation (4.9)–(4.10). A closed degree-k rectangular morphism, after shifting its source by $[1-k]$, is a degree-one twisting entry; higher blocks must satisfy the full Maurer–Cartan equation. These are the entries that form actual twisted diagrams. In the source cone diagram the off-diagonal data and their monad entries are as follows; all products use the stipulated coefficient–matrix Koszul rule:

Hom slot	off-diagonal source datum	off-diagonal monad entry
(g, r)	$d_E = g,\ d_R = r$	$\mathrm{diag}(g, g, g, r, g)$
a	$\Theta = -a$	$\alpha_2 = +ua,\ \beta_1 = -ua$
b	$X = -b$	$\alpha_1 = -ub,\ \beta_2 = -ub$
u_0	$s = -u_0$	$\beta_R = -u_0$
v_0	$\varepsilon = -v_0$	$\alpha_R = -u^2 v_0$
w	$\rho = w$	$\gamma = -u^2 w$

The source signs in this table come from the same cone differential: its off-diagonal t entry is $+a$, its U-to-R entry is $+v_0$, and the off-diagonal components of f_u are $+b, -u_0, -w$. Thus the table is a check on the source triangle diagram as well as on its target monad. The monad differential is linear in each listed field, so the table is exact for these arbitrary blocks, not just a first derivative at a point. Its entries are precisely (4.9)– (4.10) under κ.

Apply the displayed p_k, i_k, h_k to this whole graded block object. Their equations remain valid over the algebra of upper triangular blocks: $d\rho = -[t, X] + s\varepsilon$ is exactly its off-diagonal closed-morphism equation. For three or more blocks the same equation includes the composition and higher-homotopy terms. Thus the contractions compare the actual images of the whole finite twisted diagram, including each morphism and each higher cell. Equivalently, inserting these same h_k between composable maps is the explicit coherent comparison on the dg nerve. This proves an identification of enhanced functors, not only an agreement of their object values or tangent maps. Full faithfulness of b^* then identifies the descended monad and its maps with $Rb_*\,\mathrm{Ext}$ of these same triangle data. Lemma 4.4 independently verifies the full Hom quasi-isomorphism of this morphism map.

For a change of local splitting $X_j = X_i - c_{ij}\Theta$ and $z_{2,j} = z_{2,i} - c_{ij}z_1$, the two normal middle slots change by $\left(\begin{smallmatrix} 1 & c_{ij} \\ 0 & 1 \end{smallmatrix}\right)$. The columns α and rows β transform by this matrix and its inverse, respectively; γ and the framing are unchanged. These matrices satisfy the cocycle identity. Bundle frame changes obey the tensor and wedge rules already displayed. The intrinsic blow-up and diagonal functors supply all coherence on overlaps, including the units and counits. This proves global and equivariant descent of this identification. □

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

$$q_{\mathcal{C}}(g, w) = q_{\mathcal{C}}(w, g) = \operatorname{Tr}(gw),$$
$$q_{\mathcal{C}}((a, b, u, v), (a', b', u', v')) = -\operatorname{Tr}(ab' - ba') - \operatorname{Tr}(v'u) + \operatorname{Tr}(vu'). \tag{4.18}$$

The last two traces are taken on the respective framing spaces. After the shift $T = \mathcal{C}[1]$, it agrees under (4.7) *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 $G_j = R\pi_* \operatorname{RHom}(M, M(-jD))$. The second-order framing gives $G_2 \simeq G_1$, because its boundary cone is $\operatorname{End}(R) \otimes R\pi_{D*}\mathcal{O}_D(-D) = 0$. 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 $\delta_{\check{\text{C}}\text{ech}} + (-1)^p[d, -]$ 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 [45, Proposition 5.2.3]. Tensoring these calculations with arbitrary rectangular Hom spaces gives exactly (4.18).

There is also a direct differential check. If P is minus the field pairing in (4.18), cyclic trace gives $P(Dg, z) = \mathrm{Tr}(gDz)$. For opposite gauge inputs g, h, the additional ghost terms cancel as

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

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 $\mathrm{Tr}(d\Theta \wedge dX) + ds \wedge d\varepsilon$, the derivative of $-\mathrm{Tr}(X\, d\Theta) - \varepsilon\, ds$. 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 (4.7) 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 [47, Theorem 2.12], the leading component of Ch_2 is $\frac{1}{2}\mathrm{Tr}(\mathrm{mult})$. We therefore use

$$\omega_{\mathrm{Perf}}^{\mathrm{Tr}} = 2\mathrm{Ch}_2. \tag{4.19}$$

We give the fixed-source construction once, allowing the tame root curves used later. All parameter bases below are derived base changes of this fixed curve and fixed bundle.

Proposition 4.7 (The fixed-framing trace form)**.** *Let $\mathscr{C}$ be a fixed smooth proper tame Deligne–Mumford curve over $\mathbf{C}$, and let N be a fixed rank-two bundle with a specified isomorphism $\det N \simeq K_{\mathscr{C}}$. Put $Y = \mathbf{P}_{\mathscr{C}}(\mathcal{O} \oplus N)$ and $D = \mathbf{P}_{\mathscr{C}} N$. For $R \in \mathrm{Perf}(\mathscr{C})$, the fixed-framing fiber*

$$\mathrm{Map}(Y, \mathrm{Perf}) \underset{\mathrm{Map}(D,\mathrm{Perf})}{\overset{h}{\times}} \{p_D^* R\}$$

has a closed (-1)-shifted symplectic trace form normalized by $2\,\mathrm{Ch}_2$. Its specified paths are the boundary path for $K_Y \to \mathcal{O}_Y \to \mathcal{O}_{3D}$ and the zero framing path. Restriction from $2D$ to D is an equivalence over the open substack of pullback framings, on all derived test bases and including higher maps. The construction and its relative Hom and trace maps are compatible with base change and finite diagrams of framed perfect objects.

Proof. The determinant identification gives $K_Y \simeq \mathcal{O}_Y(-3D)$. For $Z = Y, D, 2D, 3D$, write $\mathfrak{M}_Z = \mathrm{Map}(Z, \mathrm{Perf})$. Each Z is a proper tame Deligne–Mumford stack with finite diagonal, including the nonreduced Cartier thickenings. By [67, Theorem A, Lemmas 4.4–4.5 and Remark 4.6], its quasi-coherent category is compactly generated and its compact objects are precisely its perfect objects. Exact tame coarse pushforward and proper coherent cohomology make $\mathrm{RHom}_Z(E, F)$ a perfect $\mathbf{C}$-complex for perfect E, F. Thus [66, Theorem 3.1] gives local geometricity of $\mathfrak{M}_Z$.

For every connective algebra B, [68, Propositions 3.24 and 4.6] identifies $\mathrm{Perf}(Z_B)$ with $\mathrm{Perf}(Z) \otimes \mathrm{Perf}(B)$ in small stable idempotent-complete categories. Every perfect object is therefore a retract of a finite iterated extension of external products $E \boxtimes P$. Projection formula and derived base change [68, Proposition 3.10] give

$$R\Gamma(Z_B, E \boxtimes P) = R\Gamma(Z, E) \otimes P,$$
$$R\,\mathrm{Hom}_{Z_B/B}(E \boxtimes P, F \boxtimes Q) = \mathrm{RHom}_Z(E, F) \otimes R\,\mathrm{Hom}_B(P, Q).$$

Finite cones and retracts show that these pushforwards and relative Hom complexes are perfect and commute with arbitrary derived base change. The unit on Z_B is compact. This also verifies O-compactness in [47, Definition 2.1] on every derived affine test base.

The same description verifies local finite presentation. For a filtered colimit $B = \operatorname{colim}_i B_i$, perfect modules and the finitely many attaching maps of a finite cell descend to some B_i. A retract descends after its idempotent and the relation $e^2 = e$ descend in the homotopy category; the perfect category is idempotent-complete, so the descended idempotent splits. Mapping spectra commute with the filtered colimit by the displayed formula, finite cones and retracts. Inverses and their identities descend as well. Taking cores therefore gives

$$\operatorname*{colim}_i \mathfrak{M}_Z(B_i) \simeq \mathfrak{M}_Z(B).$$

The relative Hom calculation and [66, Propositions 3.3–3.5] likewise give the locally geometric, locally finitely presented stacks of maps and finite flags needed below. Compatible Artin open neighborhoods are obtained by intersecting inverse images under the finitely many restriction maps; Tor-amplitude bounds can be enlarged when taking a cone. We use these opens, rather than asserting that an unbounded perfect-complex stack is a single Artin stack.

We next establish the full framing locus. The square-zero ideal of D in $2D$ is $J = \mathcal{O}_D(-D)$. The projective-line calculation gives

$$Rp_{D*}J = 0, \qquad Rp_{2D*}\mathcal{O}_{2D} = \mathcal{O}_{\mathscr{C}},$$

after every derived base change. For a perfect family F on D_T, the counit $p_{D_T}^* Rp_{D_T*}F \to F$ has a perfect cone commuting with base change. Its vanishing locus in T is open, since $D_T \to T$ is proper. Hence pullback framings form an open substack $\mathfrak{U}_D \subset \mathfrak{M}_D$. If $E \in \operatorname{Perf}(2D_T)$ has specified restriction $p_{D_T}^* R_T$, tensoring the structure triangle with E and pushing to $\mathscr{C}_T$ gives $Rp_{2D_T*}E \simeq R_T$. The counit $p_{2D_T}^* Rp_{2D_T*}E \to E$ is an equivalence on D_T; its perfect cone vanishes by derived Nakayama. The second displayed identity also gives full faithfulness. These actual adjunction maps identify

$$\mathfrak{M}_{2D} \underset{\mathfrak{M}_D}{\times} \mathfrak{U}_D \simeq \mathfrak{U}_D.$$

This describes the entire inverse image, including all derived test bases and higher maps, and gives the required actual étale locus.

Proper Deligne–Mumford duality and regular-embedding adjunction [69, Corollary 2.10, Theorem 2.22 and Theorem 3.8] give, by the Cartier Koszul resolution,

$$j_n^! K_Y[3] \simeq (K_Y \otimes \mathcal{O}_Y(nD))|_{nD}[2], \qquad \omega_{2D} = \mathcal{O}_{2D}(-D), \quad \omega_{3D} = \mathcal{O}_{3D}.$$

The trace and perfect-complex Serre duality identities extend to every derived affine base by the external-product, cone and retract argument above. The composite

$$[3D] : R\Gamma(3D, \mathcal{O}_{3D}) \xrightarrow{\partial} R\Gamma(Y, K_Y)[1] \xrightarrow{\mathrm{Tr}_Y} \mathbf{C}[-2]$$

is a degree-two O-orientation. The triangle $K_Y \to \mathcal{O}_Y \to \mathcal{O}_{3D}$ gives its boundary nullhomotopy. For every perfect E, duality identifies

$$R\Gamma(Y, E^\vee) \simeq \operatorname{fib}\big(R\Gamma(3D, E^\vee|_{3D}) \longrightarrow R\Gamma(Y, E)^\vee[-2]\big)\,,$$

also after derived base change. This is the nondegenerate boundary condition of [70, Definition 2.8]. Integration in [47, §2.1] and [70, Theorem 2.9], applied to the same universal class $2\,\mathrm{Ch}_2$, gives the closed form on $\mathfrak{M}_{3D}$ and the Lagrangian ambient restriction $\mathfrak{M}_Y \to \mathfrak{M}_{3D}$ with this specified path.

For the other Lagrangian, fix $f = p_D^* R$ and apply the symplectic argument of [56, Lemma 3.4] using the actual open equivalence just proved. On $3D$,

$$I_D = \mathcal{O}_{2D}(-D), \qquad I_{2D} = \mathcal{O}_D(-2D), \qquad I_D I_{2D} = 0.$$

In the fixed-$2D$-framing description of $\mathfrak{M}_{3D} \times_{\mathfrak{M}_D} \{f\}$, these ideal products give the zero isotropic path for the full negative-cyclic trace class. For $V = g^* T_{\mathrm{Perf}}$, the nondegeneracy map is

$$R\Gamma(2D, V \otimes \omega_{2D}) \longrightarrow R\Gamma(2D, V^\vee \otimes \omega_{2D})[2] \xrightarrow{\sim} R\Gamma(2D, V)^\vee.$$

The first arrow is induced by the 2-shifted symplectic trace form and the second by dimension-two duality, so this is an equivalence. The framing-fiber equivalence identifies the D- and $2D$-relative tangent fibers. This proves the Lagrangian assertion with its zero path.

On compatible Artin opens the two Lagrangians intersect in

$$U_Y \overset{h}{\underset{U_{3D}}{\times}} (U_{3D} \overset{h}{\underset{U_D}{\times}} \{f\}) \simeq U_Y \overset{h}{\underset{U_D}{\times}} \{f\}.$$

Their intersection has a closed (-1)-shifted form by [47, Theorem 2.9]. The universal class, proper trace, boundary path and zero framing path are the same on every open. Functoriality supplies coherent overlap identifications, and descent for closed forms glues the form on the whole fixed-framing fiber. Its leading pairing is the proper relative Serre product. The construction of relative Hom and trace uses the same restriction maps and finite diagrams. For varying framings or flags this last assertion concerns those maps; it does not assert a symplectic form on their total parameter stack. □

Apply the proposition to the ordinary curve $\mathscr{C} = C$. For the monad above the specified lift on $2D$ is its actual projection to R: the entry $\alpha_R = z_0^2 \varepsilon$ vanishes, and the kernel is the contractible Koszul complex near D. The entries of the morphism map κ into the framing contain z_0^2, except for the given framing map r. Thus the projection respects the full framing maps and their finite twisted diagrams. The leading pairing is exactly Lemma 4.6. The normalization (4.19) leaves the coefficient $1/2$ in (4.2) unchanged: that coefficient is the 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 (4.10). The pairing

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

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 (4.10); all other cases vanish. Thus, after curve sections and Serre duality, there is a natural paired triangle

$$B_{ji}^\vee[-3] \longrightarrow C_{ij} \longrightarrow B_{ij}, \qquad B_{ij} = R\Gamma(C, \mathcal{B}_{ij}), \quad C_{ij} = R\Gamma(C, \mathcal{C}_{ij}). \tag{4.20}$$

It extends to finite twisted sums and cones in (4.3) by their block-matrix differentials. The framing of a cone is the corresponding perfect cone on C; a fixed-vector-bundle slice is not claimed to be closed under this operation. In particular it applies when a cofiber is a finite-length sheaf.

The leading pairing comparison of Lemma 4.6 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 $\delta - D h_0 \in F^3 \mathrm{DR}$. For the Euler field $\mathcal{E}$, Cartan's formula gives

$$H_0 = h_0 + \tfrac{1}{2} \iota_{\mathcal{E}} (\delta - D h_0) \in F^2 \mathrm{DR}, \qquad D H_0 = \delta. \tag{4.21}$$

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 (4.21) 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 (4.20), duality, and the canonical critical factor $(-1)^{\mathrm{vdim}\, \mathscr{P}_L}$ in [29, Example 3.20]. This defines the orientation

used throughout the paper. The bimodule structure and the determinant double-sequence identity [29, Lemma 2.4] give its direct-sum and finite-flag compatibilities.

In the following square, Ψ is the physical-to-critical functor specified above. Lemma 4.5 identifies its inverse with the monad functor, so its tangent map is the inverse of $\kappa[1]$ and its cotangent map is the corresponding dual inverse. The canonical coordinates of this same functor are fixed by (4.15)–(4.16); the dual and shifted maps below use (4.17) on both sides. Thus the paired-triangle comparison uses the derivative of this same enhanced functor with those independently evaluated coordinates.

Proposition 4.8. *Let Ψ be the equivalence of Proposition 4.3, or one of the supported equivalences constructed below, and let M_A, M_B be the orientation lines of its source and target. In the supported case these are the baseline-relative orientations specified below; an absolute comparison also requires the identity of a common baseline or a separately specified square-preserving baseline isometry. Its specified half-line comparison $\gamma : M_A \to \Psi^* M_B$ satisfies*

$$o_A = \det(d\Psi^*)\, o_B\, \gamma^{\otimes 2}, \qquad d\Psi^* : \Psi^* \mathbb{L}_B \longrightarrow \mathbb{L}_A. \tag{4.22}$$

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

Proof. First consider (4.3). 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 [29, 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 4.6, 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 (4.22). 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 (4.20) 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 4.11. □

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 4.9. *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 $1, t, \dots, t^{a-1}$, 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 4.9. 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 $I_0 = \mathrm{fib}(\mathcal{O} \to \mathcal{O}_{Z_\lambda})$. 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 $\omega_{Z_\lambda} = \operatorname{Hom}_C(\mathcal{O}_{Z_\lambda}, K_C)$.

We make the baseline convention explicit. Write $M_H(G)$ for the surface-half line defined in Lemma 4.11, and put

$$M_{0,C} = M_H(G_0), \qquad \overline{M}_H(G) = M_H(G) \otimes M_{0,C}^{-1}. \tag{4.23}$$

The invariant physical baseline is acyclic, as proved below, so its full cotangent determinant has its canonical acyclic trivialization. Its square map makes $M_{0,C}$ a fixed quadratic line; it does not give a canonical trivialization of that half line. The relative orientation is the graded tensor quotient of the two square maps, using that acyclic full-line trivialization. At the coefficient level put $U_{0,C} = \Phi_{\{I_0\}^H}$ and $\overline{\Phi}_G = \Phi_G \otimes U_{0,C}^{-1}$. All supported comparisons between different baselines are first made on these relative roots and these baseline-normalized coefficients.

If both ends use the same fixed reference baseline, tensor the relative comparison with $\mathrm{id}_{M_{0,C}}$. For different baselines an absolute graded root comparison is asserted only after specifying square-preserving baseline isometries of the same parity. Such isometries, when wanted, are chosen through one fixed quadratic reference line, so identity and composition are automatic;

ordered product comparisons use their graded tensor products. Different-parity baseline lines are not identified. At the coefficient level one may instead fix a unit calibration $U_{0,C} \simeq \mathbf{Q}$ once per baseline, using separately the parity calibration of [29, §7.8]. Ordered products use the Thom–Sebastiani product of these fixed calibrations, not new independent choices.

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. Each selected branch is required to be universally injective on the entire finite defect support; the other sheets over its image are removed.

Lemma 4.10. *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

$$\big(\mathrm{Perf}(Y) \times_{\mathrm{Pic}(Y)} \{\det I_0\}\big) \underset{\mathrm{Perf}(U)\times_{\mathrm{Pic}(U)}\{\det I_0|_U\}}{\times} \{I_0|_U\}. \tag{4.24}$$

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 (4.24) 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. We use etale hyperdescent of perfect dg-modules, including mapping spaces [60, Lemma 3.1], and the natural derived determinant [50, Section 3.1, Definition 3.1]. For the underlying support square and its restriction maps, see [57, Tags 08GW, 08H9, 08GG, 08HB]. 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 tangent-character formula [44, Proposition 3.19] gives a^{h_b} and a^{-h_b} at a box; 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 4.11. *For the equivalence in Lemma 4.10, the four supported Hom and section maps define an algebraic comparison of the baseline-relative lines $\overline{M}_H$. Its square is* (4.22)*, with the relative square maps, for the derivative of that same equivalence. An absolute comparison is obtained by tensoring with the identity of a common fixed baseline, or with a separately specified square-preserving baseline isometry as in* (4.23)*.*

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 (4.20), 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 $M(G) \otimes \det(\mathbb{L}^{>0}_{\mathrm{phys}})^{-1}$ therefore identifies with $M_H(G)$, using the graded regrouping maps of [29, Section 6.5]. The double Hom grid is H-equivariant before taking determinants. Its full determinant identity is

$$\frac{M(G)}{M(G_0)} = \det \mathrm{RHom}(T,G_0) \otimes \det \mathrm{RHom}(T,T) \otimes \det \mathrm{RHom}(G_0,T) \otimes (\det R\Gamma T)^{-1}. \tag{4.25}$$

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 $\overline{\gamma}$ on the relative lines $\overline{M}_H$. This is the comparison furnished by the four maps alone. If the baseline is identical at both ends, define $\gamma = \overline{\gamma} \otimes \mathrm{id}_{M_{0,C}}$; for different baselines replace this identity only by a separately specified baseline isometry.

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

$$R\mathcal{H}om_Y(J,I_0), \qquad \mathscr{A}_{I_0} \oplus R\mathcal{H}om_Y(J,J), \qquad R\mathcal{H}om_Y(I_0,J). \tag{4.26}$$

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 (4.26) is acyclic and the other terms are perfect with finite support, so their determinants are defined. Restriction and gluing

preserve the universal triangle, composition, identity, and supertrace. They therefore preserve the displayed filtration, including its extension differential κ_0. Taking the tangent complex of the same restriction-and-gluing functor gives this filtered map. Consequently its determinant is the determinant of the derivative in Lemma 4.10. 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 (4.20), 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 (4.25); its weight-zero half determinant is the same formula with invariants in every Hom and section term. Applying the proof of Proposition 4.8 to this grid identifies the square of $\overline{\gamma}$, using the quotient of the two orientation square maps, with the derivative determinant just computed. The baseline full determinant is canceled by its acyclic trivialization; no half-line trivialization follows from that step. Tensoring with a specified square-preserving baseline comparison gives the same assertion for the corresponding absolute γ. 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 4.9. □

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. Locally, for generators f of the support ideal, its model is the tensor product of the complexes $[\mathcal{O}\to\mathcal{O}[f^{-1}]]$. These bounded flat complexes represent the support functor and commute with derived base change. 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 4.6. 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 [45, §§4,10] identify these counits on ordinary Noetherian models. Their derived compatibility uses perfect duality and the following flat models. 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 (4.21) 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.

Lemma 4.12 (The reference branch over the parameter space)**.** *Let S be a smooth coordinate-and-support parameter space and let $W\to\mathbb{A}^1\times S$ be the selected relative étale neighborhood*

used to define A*. The normal coordinates are transported by the cotangent lift, and the reference transverse monomial ideal is fixed. At a punctual reference quotient* q_0 *and a point* $s_0 \in S$*, the* A*-side supported, fixed-determinant,* H*-fixed formal problem is the scalar extension to* $R = \widehat{\mathcal{O}}_{S,s_0}$ *of the reference physical formal problem at* q_0*. This identifies the actual restriction–gluing functor on connective Artin test algebras, not only its ordinary truncation. Here the relative threefold is* $\mathrm{Tot}_W(K_{W/S} \oplus \mathcal{O}_W)$*, with its canonical relative volume. The completion is taken at* q_0 *in the full defect moduli, and then along the resulting constant section over* $\operatorname{Spf} R$*; it is not a completion inside a fiber of the support-cycle map. Nilpotent motions and splittings of the defect support are therefore retained.*

Proof. Translate the reference curve coordinate so that the chosen support section is $\tau = 0$. Relative formal étaleness gives

$$\widehat{W}_{\mathrm{section}} \simeq \operatorname{Spf} R[[\tau]].$$

After the cotangent lift, the formal transverse surface and volume are the reference ones over this ring: the normal transformation multiplies each transverse monomial by a unit and hence preserves its monomial ideal. The baseline and its section are therefore the pullbacks of the fixed reference baseline and section. This is a specified isomorphism of formal ringed spaces, with the specified volume and H-action. It induces the corresponding equivalence on perfect complexes, their maps, determinants and trace kernels after every connective base change. Thus it gives the constant reference description on the formal support neighborhood.

The pulled-back object at the section is the reference perfect object I_{q_0}, not only its baseline: this follows by transporting the actual quotient q_0, its relative curve dual, and its section cone. For a connective Artin R-algebra A, restriction under the displayed formal isomorphism sends this entire diagram to its reference diagram over A. The localization triangle and the acyclicity of the invariant trace-free baseline on the complement identify its supported lifting and mapping spaces with those of the physical problem. Induction over square-zero extensions and the Postnikov tower therefore gives an equivalence of these same formal functors on every such A. This proves the asserted scalar-extension description without deducing it from the ordinary truncation.

For the moving-coordinate family, write $\upsilon = \tau v(t,p)$ and $X = s\tau$ as in (11.24). Formal étaleness of x solves $x(t) = x(p) + s\tau$ uniquely for t over $R = \mathbb{C}[[p - p_0, s - s_0]]$, including when $s_0 = 0$. In relative differentials one has

$$d_{W/B}t = \frac{s}{x'(t)}d\tau, \qquad d_{W/B}\upsilon = \frac{y'(t)}{x'(t)}d\tau.$$

The second multiplier is a unit. Thus the relative cotangent change is

$$(\tau, \eta, z) \longmapsto \left(\upsilon, \eta\frac{x'(t)}{y'(t)}, z\right),$$

and it preserves $d\tau \wedge d\eta \wedge dz$ and the transverse monomial ideal over the full ring R. The first identity is not used to identify $K_{W/B}$ with the pullback of K_{U_t} at $s = 0$.

No fiberwise choice of a critical chart or square root is used. The ordinary quotient map produced by these operations is precisely the originally specified A, by the Artin-family identification in Lemma 4.10. □

Proposition 4.13. *On a neighborhood of the punctual locus, the maps* A, B *identify the algebraic d-critical sections and the baseline-normalized oriented vanishing-cycle coefficients* $\overline{\Phi}$*. The comparison holds over the entire smooth moving-support and coordinate parameter space. When the two reference branches have the same fixed baseline, its identity gives the comparison of the absolute coefficients as well. For different baselines absolute coefficient comparisons use the fixed unit calibrations specified after* (4.23)*.*

Proof. The ordinary map A is étale. Its relative derived étale lift is the restriction–gluing lift: the selected étale map induces the same perfect-complex functor on every connective Artin test

algebra, as in Lemma 4.12. That lemma, rather than ordinary étaleness by itself, identifies the A-side with the fixed reference physical formal problem extended over $R = \mathbb{C}[[s_1, \dots, s_e]]$. Choose its minimal critical model once at the reference point over $\mathbb{C}$, and then extend scalars:

$$\mathcal{A}_R = R[[x_1, \dots, x_m]][e_1, \dots, e_m], \quad |e_j| = -1, \quad de_j = I_j(x) \in (x)^2. \tag{4.27}$$

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 [9, Theorem 5.18]; it is chosen once over $\mathbf{C}$.

Represent the two specified maps on cofibrant complete models, calling the representatives α, β.

The de Rham complex in the following calculation is relative to R. The functions f_A, f_B are the actual functions representing the already defined smooth pullbacks of the absolute d-critical sections to the parameter products. The representative β is the composite of the specified restriction–gluing functor with the original second chart map; changing cofibrant presentations does not replace that functor or either ordinary map.

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, $h_{\mathcal{E}} = \iota_{\mathcal{E}}/v$ contracts the augmented de Rham differential. The internal differential raises that degree, by (4.27). 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

$$F = -\sum_{i,j} M_{ij} I_i I_j \in I^2 \quad \text{in } R[[x]]. \tag{4.28}$$

This is equality of the actual functions modulo I^2 in $R[[x]]$, including all parameter coefficients, and hence equality of their absolute d-critical classes. No conclusion is drawn by discarding absolute parameter differentials. In particular the relative shifted-symplectic model is not being asserted to be an absolute shifted-symplectic product with the smooth parameter base. The coefficient comparison is the one in [9, Example 5.22, equations (5.35)–(5.39)], applied to the constant reference model after extension to R.

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 (4.28) 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 relative half-line map of Lemma 4.11 existed algebraically before completion. With a common fixed reference baseline, its identity also defines the absolute

half-line map 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 [29, §§3.5–3.6], with the same Hessian and critical factor used above. Equation (4.22) is therefore the oriented d-critical comparison for the relative roots, proving the assertion for $\overline{\Phi}$. The coefficient map is constructed by algebraic critical-chart stabilization and its orientation torsors [10, Theorems 5.4(d), 6.9 and Section 6.4], with monodromic Thom–Sebastiani [16, Proposition 2.13(6)]. A common baseline is restored by its identity. Different baselines are restored only using the separately fixed coefficient units; this step does not assert an isomorphism of different-parity graded orientation lines. □

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 [16, §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 [16, §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 [63]. The monodromy operators and Hodge data are retained by this equivalence. On the kernel model the weight filtration is transported from the weight-stable Serre quotient [16, Section 2.1, Definition 2.7]; it is not the ambient weight filtration of its representative [16, Remark 2.8]. All weights below use this intrinsic quotient convention.

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

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

The functors below are defined on the coefficient category. Their displayed identifications are used on its subcategory of objects monodromic for the acting torus: the underlying constructible complexes are locally constant along its orbits, as in [29, Definition 7.1, Theorem 7.7]. This is a condition on the action on X, separate from the auxiliary monodromy in $\mathrm{MMHM}(X)$. On that subcategory, distinguish

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

The first is used for filtered-stack attachment, and the second for compact attracting strata.

Here the coefficients to which these identifications are applied are monodromic for the acting torus. For the equivariant charts of the filtered-stack construction [29, Proposition 3.29], the d-critical section is pulled back from the original stack and is invariant. In the compact applications, H from (11.1) and R from (11.11) preserve $dt\wedge d\theta\wedge dz$; their root-chart versions commute with the finite root groups and preserve the same volume form. Thus they preserve the physical d-critical section. On an invariant critical chart the coefficient before the orientation twist is equivariant, by smooth functoriality of vanishing cycles. The orientation transport tensors it with a finite sign local system. Its restriction to each orbit is locally constant; on a parametrized $\mathbf{G}_m$-orbit any sign monodromy is killed by the cover $t \mapsto t^2$. This proves the required torus monodromicity without choosing a global equivariant trivialization of the orientation. Equivariant restriction, the induced fixed-locus orientation, and finite root-group descent preserve this property, so it also holds for the successive compact localizations used below.

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 [29, Definition 7.21]. For the surface-half model (4.20), it is $\chi(B^-) - \chi(B^+)$.

Proposition 5.1. *On every finite-type window of the oriented (-1)-shifted symplectic stacks used here, whose classical truncation is a quasi-separated 1-Artin stack with affine stabilizers, 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 preserving its d-critical section, use the coefficient, monodromic for the acting torus as established above, and the induced fixed-locus orientation. 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 (4.20) and Proposition 4.8. These are the stack hypotheses of [29, 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 [19, Theorem 3.7, Proposition 6.2]. Its staircase inverse follows from the adjunction triangle identities. Tensor it with the finite-torsor map of [29, (6.75)]. The normalization on a moving hyperbolic quadratic block is the compact Thom map of the same repelling axis as in [29, 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

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

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

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

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 [29, 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 $\pi_r = \operatorname{Hom}(A, B[-r]) = 0$ 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 [29, 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 [29, Lemma 7.4]. Thus the product and matching-refinement constructions in [29, Propositions 7.10–7.12, Lemmas 7.14–7.15] are the same Hodge transformations, with determinant associators from (4.20). 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 apply the preceding operations on bounded finite-type stack windows with constructible coefficients K. If $V \to \mathcal{X}$ is a smooth atlas of finite type given by finitely many charts, $p_m : V_m \to \mathcal{X}$ 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

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

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 [62, 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 [62, 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

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

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 $f \le e$ gives slope at least $\alpha_f \ge \alpha_e$.

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. For the range $c \le c_r$ used in the restriction maps below, a positive-rank framed graded factor has minimum slope at least $\alpha_e \ge \alpha_r > c_r \ge c$, and the unframed good factors have minimum slope greater than c. Thus every extension of these good graded factors 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

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

where R is the fixed framing and F 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 [61, 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 L-Higgs subobject of the projected middle bundle, the torsion-free part of its Higgs 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

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

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 $\operatorname{Ext}^2(Q, \mathcal{O}_X) = 0$ 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.** For this attachment construction, assume $c \le c_r$. The numerical finiteness assertion of Lemma 6.3 itself holds for arbitrary fixed c.

For attachment retain the Harder–Narasimhan types, with strictly decreasing unframed tail slopes; equal adjacent slopes are grouped. Order these 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. The uniform choice of T puts every slope occurring in the unframed quotient, after torsion absorption and Higgs refinement, strictly below T. If the resulting framed prefix is unstable, its first Harder–Narasimhan factor is framed and has smaller curve rank, so the true type is earlier in the primary order. If that prefix is semistable, it is the first true Harder–Narasimhan factor. The true Harder–Narasimhan polygon of the remaining unframed quotient dominates the polygon of every lifted flag, and hence its concave majorant. Indeed, intersect a flag subobject with the Harder–Narasimhan filtration and apply the semistable slope inequality to each successive image; filling the available ranks in decreasing slope order gives the asserted polygon bound. If an unframed graded factor is unstable, take the preimage of its first Harder–Narasimhan subobject in the corresponding flag member. The resulting internal vertex lies strictly above that factor's displayed chord. Torsion absorption already raises an original vertex, including the initial vertex when torsion enters the framed prefix. Thus every bad flag with the same framed curve rank has a strictly higher true anchored tail polygon.

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

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

Thus a family factoring through the scheme-theoretic stratum S_τ has a unique relative flag, also over a nonreduced base.

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. Apply the exact equivalence of Lemma 4.2 and the diagonal reconstruction of Lemma 4.1 to the universal finite diagram, with its framing diagram on C. On the vector-bundle locus the monic reconstruction and Lemma 4.5 identify its truncation with the coherent flag rows and their compatible fields. 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 U_τ. Conversely a filtered object with these graded terms has that coherent flag. The sheaf, vector-bundle, and window conditions are imposed as opens on this common enhanced diagram problem; the identification therefore

persists 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 [29, 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 [29, 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. The split coefficients are equivariant for the cotangent action $(\Theta, X, s, \varepsilon, \rho) \mapsto (u\Theta, u^{-1}X, s, \varepsilon, \rho)$, 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 the source images in $\mathscr{I}_r$ for types with at least two positive-rank factors. The vacuum with one rank-r unframed factor is treated next.

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 $\mathcal{A}_r^L$. Set

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

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 $D_{\mathscr{S}_r}$, 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

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

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

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

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 $D_{\mathscr{I}_r}$. The octahedral triangle transports (6.11) through the finitely many walls between t and T.

Two cutoffs $c_1 < c_2 \leq c_r$ 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

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

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

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

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

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

The nested-window restrictions in Proposition 6.5 act on this same arrow and its cofiber. Chamber restrictions are used separately on the two physical coefficients. A restriction of the whole arrow to a chamber is used only when both physical chamber loci are inverse images of a common open in the surface-pair stack, as at the Joyce–Song endpoint below.

Lemma 6.6 (The split chamber reflection)**.** *Let $\mathscr{X}^t_{r,\chi}$ be the full finite-t semistable framed stack on $X_0 = \mathrm{Tot}_S(K_S)$, with its specified boundary framing and with the zero section retained. There is an isomorphism of classical moduli stacks over $\mathcal{A}^L_r$*

$$\mathsf{D} : t_0\mathscr{X}^t_{r,\chi} \xrightarrow{\sim} t_0\mathscr{X}^{-t}_{r,-\chi}.$$

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

Proof. Finite-t semistability gives pure curve pushdown by the purity argument in Lemma 6.1. The coherent-diagram description of Proposition 4.3 therefore uses a vector bundle E on C and fields

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

The maps below are maps of these coherent diagrams, including families and their specified boundary framing. They give

$$(E, \Theta, X, s, \varepsilon) \longmapsto (E^\vee K_C, \Theta^{\mathsf{t}}, X^{\mathsf{t}}, \varepsilon^{\mathsf{t}}, -s^{\mathsf{t}}), \tag{6.13}$$

with a gauge transformation g sent to $(g^{-1})^{\mathsf{t}}$. The transformed moment expression is the negative transpose of the old one, so it still vanishes. Bundle duality and transposition commute with arbitrary ordinary base change. Applying the map twice fixes the two fields and negates both framings; the gauge map $-\mathrm{id}_E$ identifies this square with the identity. Thus it defines an isomorphism of the classical moduli functors. Equivalently, this is the boundary-framed physical duality in the coherent-diagram description of [61, Lemma 4.5.14].

An unframed subobject is an invariant curve submodule annihilated by ε, and an unframed quotient is an invariant quotient annihilating s. To test semistability, it suffices to use saturated subobjects and pure quotients. Saturation raises the subobject's Euler slope, preserves both field invariance and the vanishing of ε, and gives a pure complementary bundle. Removing finite torsion from a quotient lowers its Euler slope and retains the condition on s. Zero-dimensional subobjects are excluded by finite-t semistability.

Duality exchanges these exact sequences and reverses their arrows. Serre duality on C gives $\chi(E^\vee K_C) = -\chi(E)$, and the same identity applies to every unframed bundle factor. Consequently the subobject inequality $\mu \le t$ becomes the quotient inequality $-t \le -\mu$, and conversely. This proves the stability equivalence for every finite t, including walls and $s = 0$, as in [61, Lemma 4.4.6]. Since Θ is transposed without a sign, the characteristic polynomial is unchanged. The isomorphism is therefore over the identity of the Hitchin base.

For the cutoff assertion, the coherent projection sequence (6.3) has cyclic term of minimum Euler slope at least α_r. Every positive-rank HN quotient of the other term has slope at least t. If $t = \pm\epsilon$ and $0 < \epsilon < 1/r$, every such finite slope is nonnegative, since its denominator is at most r. As $\alpha_r \le 0$, purity and extension closure give $\mu_{\min} \ge \alpha_r > c_r$. This proves containment of both full 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 $\mu_{\min} \geq \alpha_r$. 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

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

Consequently $H^1(C, E) = 0$ when $\chi/r > h + (r-1)\ell$. The projective Joyce–Song framing threshold can be chosen uniformly as well. Fix a degree-one line bundle M and write $\chi = \chi_a + rk$, with $0 \leq \chi_a < r$. Tensoring the underlying Higgs bundle gives $F = E \otimes M^{-k}$ with $\chi(F) = \chi_a$; a section $\mathcal{O}_C \to E$ becomes $M^{-k} \to F$. The finitely many fixed-characteristic semistable families are bounded, so their ordinary Joyce–Song constructions admit a common sufficiently large twist bound. This is used only for the projectivity of the ordinary pair space below, not as a tensor identification of the physical framed moduli. Increase the bound on χ/r accordingly. 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 $t \leq \chi/r$, a contradiction. Thus $s \neq 0$. 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 [46, Definition 2.11 and (2.33)]. 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. Both physical chamber opens are the inverse images of $\mathscr{P}^{\mathrm{JS}}$ under their critical projections: the preceding stability test is an equivalence in both directions, and equality on geometric fibers identifies these derived opens. The local counit defining $\delta_{\mathscr{P}}$ and the forward exchange e_h therefore restrict naturally to the same comparison arrow over $\mathscr{P}^{\mathrm{JS}}$. 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 classical moduli stacks after replacing χ by $-\chi$. If $\chi > -a_r$, then $-\chi < a_r$, so the positive immediate chamber in class $-\chi$ is empty by the lower endpoint. The negative immediate chamber in class χ is therefore empty as well, and its direct-image coefficient is zero. The preceding wall transport and split dimensional reduction place the full ordinary source in $D_{\mathscr{S}_r}$, proving the upper source endpoint. □

7. The Hecke relation on the comparison arrow

The next construction proves a degree-zero Weyl relation on the normalized whole arrow $\Sigma_\chi(c) = \Delta_\chi(c)\Lambda^\chi$ formed from (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 $\operatorname{Fun}(\Delta^1, \operatorname{Perf}(S_D))$, 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.

For clarity, put $M = L(D)$ and fix the source R_D and its identity by homotopy fibers. For a specified exterior object Q in P_D, a triangle in (7.5) is equivalently a map

$$(R_D \longrightarrow F_+) \longrightarrow (0 \longrightarrow Q)$$

whose underlying curve map is a quotient. Its fiber recovers the first arrow with source canonically R_D. The mapping complex of this map of arrows is

$$\mathrm{fib}\big(R\,\mathrm{Hom}_{S_D}(F_+, Q) \longrightarrow R\,\mathrm{Hom}_{S_D}(R_D, Q)\big).$$

The spectral resolution

$$[\,p^*(E_+ \otimes M^{-1}) \xrightarrow{w-\Theta_+} p^*E_+\,]$$

identifies this complex with

$$\Big[\mathrm{Hom}_C(E_+, Q) \xrightarrow{q\mapsto(q\Theta_+-\lambda q,\ qs_+)} \mathrm{Hom}_C(E_+, Q\otimes M) \oplus H^0(C, Q(D))\Big]$$

in degrees zero and one. Finite support and local freeness of E_+ make its terms vector bundles of ranks r and $r+1$. Thus the specified derived fiber product is the derived zero locus of these two equations on the generating-quotient open. This assertion includes the nullhomotopy of qs_+. The fiber construction, the resolution, and their maps are defined over the parameter ring, so the identification holds over every connective derived base.

The underlying curve quotient space is the smooth curve-Hecke space. Evaluation of the meromorphic field and section at the mark is surjective by the chosen vanishings. Composition with the quotient then gives surjections onto the two equation bundles, independently in the field and section directions. Consequently this derived zero locus is smooth. The common full frame and dual-residue coordinates add smooth factors.

There is also a coordinate-free description of the obstruction map in (7.6). Write the curve sequence as $0 \to E_- \xrightarrow{i} E_+ \xrightarrow{q} Q \to 0$. For an arrow-field error $\alpha : E_- \to E_+ \otimes M$, set

$$\alpha \longmapsto \partial(q\alpha) \ \in\ \mathcal{E}xt^1_C(Q, Q\otimes M),$$

where ∂ is the boundary obtained by applying $\mathcal{H}om_C(-, Q\otimes M)$ to that sequence. The map kills $iB_- - B_+i$ and induces an isomorphism

$$\mathrm{coker}(\mathcal{E}nd(E_-)\otimes M \oplus \mathcal{E}nd(E_+)\otimes M \longrightarrow \mathcal{H}om(E_-, E_+)\otimes M) \ \simeq\ \mathcal{E}xt^1_C(Q, Q\otimes M).$$

The spectral resolution and the scalar action on Q identify the last sheaf with $\mathcal{E}xt^2_{S_D}(Q, Q)$. This is the obstruction component of the derivative of the cofiber functor: on the cone resolution, an arrow-field error is the degree-two map from its first term to its last term. With the convention $iB_- - B_+i$, cone curvature has the opposite sign and the differential of End(Cone)[1] reverses it. Hence the obstruction map is positive evaluation in the ordered Koszul coordinates $(t-x, w-\lambda)$.

For $Q = x_*\ell$ this identifies the source and target globally with

$$\mathrm{End}(\ell)\otimes T_xC \otimes M_x.$$

The identification with $\bigwedge^2 T_{(x,\lambda)}S_D$ uses the base tangent followed by the vertical tangent. The boundary, cone, and spectral identifications are natural for isomorphisms and derived pullback. They therefore give the global derivative of the specified cofiber functor, independently of lattice coordinates.

A local lattice chart checks the sign explicitly. The arrow-field differential is

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

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

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

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

$$\pi^0 d^{-1}(u,v) = f(u_1-bv_{\mathrm{rest}}) + (w-a)v_1,$$
$$\pi^{-1}d^{-2}(z) = ((w-a)z_1, -fz_1).$$

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

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

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

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

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 $G = \operatorname{Res}_{D/T} \mathrm{GL}_r$:

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

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

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

as functions and closed schemes.

For ordinary Hodge Borel–Moore theory use the mapping spectrum

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

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 [27, §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 [18, 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 $t_0W \to W$ and properly push the classical truncation to Z. More explicitly, let $P_\pm$ be the binary normal-deformation paths from the two convolution morphisms to the properly pushed fundamental morphisms of $W_\pm$. Applying the functorial normal-deformation construction to the off-graph equivalence, with its unit and counit, gives a path Q between those fundamental morphisms. The off-graph comparison of convolution morphisms is, by definition, the concatenation $P_+ \cdot Q \cdot P_-^{-1}$. Thus localization is applied to the difference together with this specified nullhomotopy. It gives a point of

$$\mathrm{fib}\big(\mathrm{BM}_q(Z) \longrightarrow \mathrm{BM}_q(Z \setminus \mathfrak{D})\big) \simeq \mathrm{BM}_q(\mathfrak{D}).$$

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. On this source open, both words are empty after removing the graph. The specified composition and bicartesian homotopies consequently factor through the zero Borel–Moore spectrum of the empty correspondence. Their restriction is the canonical zero homotopy. The localization lift on this open is therefore the proper pushforward of the ordinary fundamental class of the eigenline cover, with scalar 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

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

The first map to the base is an equivalence, since Γ_Y is the coreflector to supported objects. Consequently (7.14) projects equivalently to $\mathrm{Map}(A_1, K_Y)$. 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

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

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

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

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

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

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 $\mathrm{Tr}\,\Theta^2$ 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 [30, 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 $+1$.

For clarity, these stabilizations act on whole diagrams. Put $A = \mathcal{E}nd(E)L \oplus E$, $\mathcal{V}_D = \pi_* A(D)$ and $\mathcal{R}_D = \pi_*(A(D)/A)$. For $D \subset D'$ the quotient of $[\mathcal{V}_D \to \mathcal{R}_D]$ in $[\mathcal{V}_{D'} \to \mathcal{R}_{D'}]$ is

$$[C_{DD'} \xrightarrow{\mathrm{id}} C_{DD'}], \qquad C_{DD'} = \pi_*(A(D')/A(D)).$$

Near D' all Hecke arrows are isomorphisms. Their quotient blocks are consequently this same identity complex. The prescribed evaluation vanishing at $\Gamma_x + \Gamma_y$ permits compatible polar lifts on the finite Hecke diagram. Indeed, for the upper word choose polar field and section lifts on its common upper bundle that vanish along $\Gamma_x + \Gamma_y$; they preserve, and lie in, both lower lattices. For the lower word choose such lifts on its common lower bundle; they extend to both upper lattices. The same vanishing proves surjectivity onto the new polar quotient, also when the thick divisor is $2\Gamma_x$. On a common affine torsor of such lifts and splittings, every endpoint and correspondence is stabilized by the same hyperbolic factor $H_C = C \oplus C^\vee$, and each correspondence has the identity correspondence on this factor.

Thom–Sebastiani, the support counit, and proper adjunction therefore identify the stabilized enhanced diagram with the original diagram tensored with the hyperbolic block counit. This identifies the binary normal-deformation paths as well as the vertices and squares. After the common normalization, the block counit is multiplication by $(-1)^{\mathrm{rk}\, C} = 1$ on the unit object. Its mapping space is discrete, so this is an identification of the block morphism used in that natural diagram transformation.

The orientation contractions are multiplicative. If ν_k is the displayed grouped rank-k form, then moving the first block's $d\zeta$'s past the second block's dc's gives

$$\nu_k \wedge \nu_l = \nu_{k+l}, \qquad \binom{k+l}{2} = \binom{k}{2} + \binom{l}{2} + kl.$$

Together with $\det(d_1 \oplus d_2) = \det(d_1)\det(d_2)$, this is the contraction identity for successive enlargements. The associativity and unit coherences of the same tensor construction provide its higher cells. Thus, for $I \subset J \subset K$, the comparison $u_{JK}u_{IJ} \simeq u_{IK}$ simultaneously carries the E- and F-squares and their two binary paths. Full faithfulness on mapping spaces of the fixed intrinsic endpoint diagrams descends these cells from the common splitting torsor. The frame action is handled by its smooth action groupoid; its inverse-dual action preserves the same orientation contraction.

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, which allows the zero section:*

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

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 normalized comparison cofiber*

$$K_\chi(c) = \mathrm{cofib}(\Delta_\chi(c)\Lambda^\chi)$$

formed from (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 $D_{\mathscr{S}_r}$. 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 [39, 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,*

$$1+\sum_{r\geq 1,\chi} R^r x^\chi [R^L_{r,\chi}] = \mathrm{Exp}_{\mathrm{add}}\left(\sum_{r\geq 1} R^r \sum_i \sum_{u,v\geq 0} x^{i+1+u+v}[P^L_i[-i-1-2v]\otimes \mathbf{T}^{v-B_r/2}]\right). \quad (8.4)$$

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 $D^b\,\mathrm{MHM}(\mathcal{A}^L_r)$.*

Proof. Fix an Euler slope μ and let

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

We apply the smooth linear-stack integrality theorem [11, 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

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

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

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

The stabilizer is $\prod_j \mathrm{GL}(V_j)$. The opposite summands $\mathrm{Hom}(V_i,V_j)\otimes \mathrm{Ext}^1(S_i,S_j)$ and $\mathrm{Hom}(V_j,V_i)\otimes \mathrm{Ext}^1(S_j,S_i)$ 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 $\mathrm{End}(V_i)$ 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 [11].

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 [11, §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 [39, 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 [11, 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 [11, 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 [11, §§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 [39, 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 $U_r^{\mathrm{red}} \subset \mathcal{A}_r^L$ be the reduced spectral open.

Lemma 8.4. *Restriction to U_r^{red} 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_r^{red}. 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*

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

For $m < 0$, the stable-pair locus and the relative Hilbert scheme are empty, and $R^L_{r,\chi}|_{U_r^{\mathrm{red}}} = 0$. For $m \geq 0$, write $H_{r,m} = \mathrm{Hilb}^m(\mathcal{D}_r/U_r^{\mathrm{red}})$ and $h_m : H_{r,m} \to U_r^{\mathrm{red}}$ for the relative Hilbert scheme of the universal spectral curve and its structure map. Then $H_{r,m}$ is smooth of dimension $b_r + m = B_r + \chi$, and the normalized coefficient on $H_{r,m}$ is $\mathbf{Q}_{H_{r,m}}[B_r] \otimes \mathbf{T}^{-B_r/2}$. Consequently

$$R^L_{r,\chi}|_{U_r^{\mathrm{red}}} \simeq R(h_m)_*\big(\mathbf{Q}_{H_{r,m}}[B_r] \otimes \mathbf{T}^{-B_r/2}\big).$$

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. Their local Kodaira–Spencer image is transverse to the conductor subspace, which is the linear support of the equigeneric tangent cone. The all-length criterion [43, 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. Every length-m subscheme of a reduced locally planar fiber deforms within that fiber to m distinct smooth points [7, Proposition 1.4]. The assertion is local at the finite support, so products of disjoint planar neighborhoods cover reducible or disconnected fibers as well. Near the distinct smooth points the relative Hilbert scheme is smooth of relative dimension m. Since its total space is smooth, its dimension is therefore $b_r + m = B_r + \chi$ at every point; for $m = 0$ this is immediate.

This dimension 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}. After the finite labeling cover, specialize to the normal bundle of the nodal stratum and pull back to the product of the frame torsors of the node-smoothing lines. Let $H_{\lim}$ denote the resulting limiting mixed variation on this torsor. The weight-graded objects below have their canonical descent to S_ν, with the labeling action retained. In particular,*

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

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 [48, §§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 [64, 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 $t'_e = u_e t_{\sigma(e)}$, the logarithmic-frame comparisons are exponentials of N_e, interpreted on local logarithmic covers; they transport the Hodge filtration. Each N_e lowers the untwisted weight filtration by two. On the weight graded, the comparison is therefore independent of the units and of the choices of logarithms, apart from the recorded graph permutation and the transition of V. The maps $N_I : \wedge^j H_{\lim} \to \wedge^j H_{\lim}(-|I|)$ are morphisms of mixed variations and are strict for F and W, so their graded images have the same descent. These images and the coordinate-order determinants give the stated equivariant descent data. Additivity for the weight filtrations and the normal-crossing hypercube computes the Grothendieck class of $\mathrm{IC}_0(\wedge^j H_{\mathrm{sm}})|_{S_\nu}$ from these descended graded terms. □

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 $\mathbf{C}[[x_0, y_0]]/(x_0 y_0)$. 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

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

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 [32, Theorem 5.6 and Proposition 5.7], restricted to the semistable locus as in [31, Proposition 3.6], uses the function

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

Its ordinary objects are commuting fields (E,θ,z) on $\operatorname{Tot}_C(K_C\oplus\mathcal{O}_C)$, 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

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

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 [57, 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 $u_0\chi_*+v_0r=1$. 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

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

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 [31, 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 [17, §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. In the convention of [10, Theorem 5.1(ii)], the normal quadratic form is $q_{\mathcal{N}} = \frac{1}{2}\mathrm{Hess}(g)$. Thus $2^{-c/2}\sqrt{u_Z}\,\delta$ trivializes its determinant square-root torsor, compatibly with base change. The relative Morse lemma in mixed Hodge modules [10, Definition 5.2, Theorem 5.4(d), Example 2.23] 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. For the symmetric-power substitution in Section 10, fix one such isomorphism in each rank. Their external tensor powers commute with permutations of equal factors. The physical packet-product comparison used there is constructed separately.

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

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

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). □

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

Lemma 10.2 (From categorical characters to geometric closed forms)**.** *For a derived stack V over $\mathbf{C}$, there is a natural morphism*

$$\eta_{V,2} : CC^-(\mathrm{Perf}(V)/\mathbf{C}) \longrightarrow NC^w(V/\mathbf{C})(2). \tag{10.5}$$

Here the source is categorical negative cyclic homology and the target is the geometric weight-two complex defined by affine-test descent. The composite with the categorical Chern character, on the space of perfect objects, is the geometric Ch_2*. This identification retains the paths from equivalences of objects and their higher homotopies. For an open $U \subset V$, the same construction gives a natural map on the restriction fibers. All complexes and maps are absolute over $\mathbf{C}$.*

Proof. For every derived affine test $u : \operatorname{Spec} B \to V$, compose exact pullback, affine Morita invariance, the natural affine HKR comparison, and the weight-two projection:

$$\begin{aligned} CC^-(\mathrm{Perf}(V)/\mathbf{C}) &\xrightarrow{u^*} CC^-(\mathrm{Perf}(B)/\mathbf{C}) \simeq CC^-(B/\mathbf{C}) \\ &\simeq NC(\mathrm{DR}(B/\mathbf{C})) \longrightarrow NC^w(B/\mathbf{C})(2). \end{aligned}$$

The affine comparison is natural as a mixed-complex comparison, as used in [47, Section 2.3]. The last arrow is the natural retraction to one weight, described before Definition 1.2 there; it does not identify the unweighted complex with the direct sum of all weights. These maps form a coherent cone over all derived affine tests. Their homotopy limit is

$$CC^-(\mathrm{Perf}(V)/\mathbf{C}) \longrightarrow \operatorname*{holim}_{(\operatorname{Spec} B \to V) \in (\mathrm{dAff}/V)^{\mathrm{op}}} NC^w(B/\mathbf{C})(2) \simeq NC^w(V/\mathbf{C})(2),$$

where the last equality is affine-test descent [47, Corollary 1.4 and Definitions 1.12–1.13]. No equivalence of the two cyclic theories on V is asserted.

Use the natural categorical character of [25, Theorem 6.9, Corollary 6.11 and Remark 6.12]. On each affine test its restriction to perfect-object spaces refines the same Toën–Vezzosi pre-character as [25, Theorem 6.5]. Under the natural comparison above, its weight-two part is precisely the geometric character in [47, Section 2.3]. Thus the identification is one of maps of spaces, not merely of their cohomology classes. It carries the path of any specified equivalence to the corresponding geometric path and retains all higher compatibilities. Restriction to an open and change of parameter base commute with the coherent cone. Taking the fiber of the restriction square proves the supported assertion. In particular a path trivializing a relative K-theory point is not forgotten. The construction keeps all parameter differentials because its coefficient base remains $\mathbf{C}$. □

Lemma 10.3 (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

$$\mathfrak{D}(Y, D_\infty) = \mathrm{fib}_{(\mathcal{O}_Y, \mathcal{O}_{D_\infty}, \mathrm{id})} \left(\mathrm{Perf}(Y) \longrightarrow \mathrm{Pic}(Y) \underset{\mathrm{Pic}(D_\infty)}{\overset{h}{\times}} \mathrm{Perf}(D_\infty) \right), \tag{10.6}$$

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 4.10. 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.

For the corresponding relative classes, let A be a connective Artin $\mathbf{C}$-algebra, put $Y_A = Y \times \operatorname{Spec} A$, and write $Z = \coprod_i Z_i$. Let K denote nonconnective algebraic K-theory and put $CC^-(V/\mathbf{C}) = CC^-(\mathrm{Perf}(V)/\mathbf{C})$ for the categorical negative-cyclic complex. Define the supported fibers

$$\begin{aligned} \mathrm{K}_Z(Y_A) &= \mathrm{fib}\big(\mathrm{K}(Y_A) \to \mathrm{K}(U_{0,A})\big), \\ CC^-_Z(Y_A) &= \mathrm{fib}\big(CC^-(Y_A/\mathbf{C}) \to CC^-(U_{0,A}/\mathbf{C})\big). \end{aligned} \tag{10.7}$$

For a perfect object J, write $k(J)$ for its point in the infinite loop space of K-theory. The compatible determinant fiber gives the trivialization $\alpha : I|_{U_{0,A}} \simeq \mathcal{O}_{U_{0,A}}$ and hence the point

$$(k(I) - k(\mathcal{O}), k(\alpha)) \in \Omega^\infty \mathrm{K}_Z(Y_A).$$

The second entry is the specified path trivializing the restriction of the first entry. For each packet let $\alpha_i : I_i|_{U_{0,A}} \simeq \mathcal{O}_{U_{0,A}}$ be the trivialization induced by the same gluing construction. The

natural spectrum-valued Chern character [25, Theorems 6.5, 6.9, Corollary 6.11 and Remark 6.12] and localization of the mixed complex [26, Sections 1.5, 1.12 and 2.4] give the commutative diagram of fiber sequences

$$\begin{array}{ccccc} \mathrm{K}(\mathrm{Perf}_Z(Y_A)) & \longrightarrow & \mathrm{K}(Y_A) & \longrightarrow & \mathrm{K}(U_{0,A}) \\ \downarrow{\scriptstyle \mathrm{Ch}_Z} & & \downarrow{\scriptstyle \mathrm{Ch}} & & \downarrow{\scriptstyle \mathrm{Ch}} \\ CC_Z^-(Y_A) & \longrightarrow & CC^-(Y_A/\mathbf{C}) & \longrightarrow & CC^-(U_{0,A}/\mathbf{C}). \end{array} \tag{10.8}$$

Here Perf_Z is the category of perfect complexes supported on Z. Supported open descent identifies it with $\prod_i \mathrm{Perf}_{Z_i}(Y_A)$, and restriction to U_i is its ith projection. This is the same restriction used by the gluing functor, so it sends the displayed point to the point defined by (I_i, α_i). The finite product equivalence and the Chern character therefore give the homotopy

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

in $CC_Z^-(Y_A)$, naturally in A, packet refinements, and permutations. The cyclic complexes are taken over $\mathbf{C}$, so all parameter differentials are retained. Apply the map on supported fibers from Lemma 10.2. It sends this specified categorical character homotopy to the geometric relative Ch_2 homotopy, including the paths from α and α_i. The cyclic words above represent its supported trace. The constant object $\mathcal{O}_Y$ has no positive-degree form in the moduli directions. Apply the fixed geometric boundary and volume transgression to twice this geometric class. Linearity then shows that 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 [29, §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 10.4. *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

$$D_{ij} := \mathrm{fib}(A_{ij}) = R\pi_{X,*}R\mathcal{H}om_{X_1\times T}(F_i, F_j) \tag{10.9}$$

is canonically contractible in every degree.

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

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

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 4.8 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

$$o_1 \boxtimes o_2 = \det(d\Psi^*) \circ \Psi^* o \circ \gamma_{12}^{\otimes 2}, \qquad d\Psi^* : \Psi^* \mathbb{L}_{\mathfrak{P}} \longrightarrow \mathbb{L}_{\mathfrak{P}_1 \times \mathfrak{P}_2}. \tag{10.10}$$

Lemma 10.3 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 (10.10). 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 10.5. *For every rank and Euler characteristic,*

$$[R^K_{r,\chi}] = [C_{r,\chi}]. \tag{10.11}$$

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

$$[R^K_{r,\chi}] - [C_{r,\chi}] \in K_0(\mathscr{I}_r). \tag{10.12}$$

Suppose (10.11) 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 10.4 give the G_ν-equivariant expression

$$p_\nu^*(R^K_{r,\chi}|_{\mathscr{R}_\nu}) \cong (v_\nu)_* \left(\bigoplus_{\sum_j \chi_j = \chi} \boxtimes_j (R^K_{r_j,\chi_j}|_{T_{r_j}}) \Big|_{W_\nu} \right). \tag{10.13}$$

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 (10.13) 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 (10.12) 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 (10.13) 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,

$$K_{r,\chi} = \Phi_{P_{r,\chi}(X_1)}|_{N_{r,\chi}} \otimes \Lambda^{\chi}. \tag{10.14}$$

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

$$1 + \sum_{r,\chi} R^r x^{\chi} [R_{r,\chi}^{\mathrm{nil}}] = \mathrm{Exp}_{\mathrm{add}} \left(\sum_{r \geq 1} R^r \sum_i \sum_{u,v \geq 0} x^{i+1+u+v} [P_i^K[-i-2v] \otimes \mathbf{T}^{v-N_r}] \right). \tag{10.15}$$

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

$$[R^r] \operatorname{Log} Z_{\mathrm{nil}}(R, x) = \frac{x \mathbf{T}^{-N_r}}{(1-x)(1-x\mathbf{T})} \sum_i (-x)^i [R\Gamma_c(\mathcal{A}_r^{K_C}, P_i^K)]. \tag{10.16}$$

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 (10.15) with the ambient restriction (10.14).

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 (10.16) 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 [16, §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 (10.14).

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 each fixed pair (C,λ), retain the actual baseline point coefficient $U_{0,C}$, and fix its unit calibration once as in (4.23). Local support comparisons and factorization are first formed after division by this point object. Restoring the single fixed calibration gives comparisons for the inherited coefficients in (11.3); no new baseline choice is made on a coordinate overlap or refinement.

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 [44, Proposition 3.19]

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

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 and identify the fixed functor with (11.2) on all ordinary Artin bases. For fixed (r, χ) they define a morphism from the finite disjoint union of the projective schemes D_{λ, d_λ} to the separated ordinary fixed locus. It is proper and is an isomorphism on geometric points and their Artin neighborhoods. Hence it is étale and radicial, and therefore an isomorphism. Only the ordinary truncation is identified here; the physical enhancement remains the inherited one 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

$$R\colon\ (t, \theta, z) \longmapsto (vt, v^{-1}\theta, z). \tag{11.11}$$

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 $k \leq 0$ 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

$$(\mathrm{rk}^+ H^0 T_Y - \mathrm{rk}^- H^0 T_Y) - (\mathrm{rk}^+ H^{-1} T_Y - \mathrm{rk}^- H^{-1} T_Y).$$

This is the localization theorem for the inverse action in the convention $p_*\eta^!$ of [29]. 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 $\operatorname{Sym}^d \mathbf{A}^1$ 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 [29, Proposition 7.2][8, 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

$$I_{H,R}(m) - I_{H,R}(0) = \tfrac{1}{2} \sum_\omega c_\omega \,\mathrm{sgn}_{\mathrm{lex}}(\omega),$$
$$I_H(m) - I_H(0) = \tfrac{1}{2} \sum_\omega c_\omega \,\mathrm{sgn}(\omega_1).$$

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 $\mathrm{sgn}(D,e) = \mathrm{sgn}_{\mathrm{lex}}(D,e)$ and define

$$\begin{aligned} A(D,e) = {} & \mathrm{sgn}(D,e) - \mathrm{sgn}(D-1,e) \\ & - \mathrm{sgn}(D+1,e-1) + \mathrm{sgn}(D,e-1). \end{aligned}$$

It vanishes outside $D = -1, 0, 1$; at these three values it equals, respectively, $-1 - \mathrm{sgn}(e-1)$, $\mathrm{sgn}(e) + \mathrm{sgn}(e-1)$, and $1 - \mathrm{sgn}(e)$. 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 $\mathsf{Z}_m + q$ 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 $\mathrm{sgn}(-c, \alpha + j)$, and the constant term $-P$ contributes 1. Adding these contributions gives

$$\Delta I = \mathrm{sgn}(-c, \alpha + j) + 1 - \sum_{(l,k) \in \lambda} A(k - l - c, \alpha + j - k - m_{lk}). \tag{11.18}$$

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 $2\,\mathrm{sgn}(D) - \mathrm{sgn}(D-1) - \mathrm{sgn}(D+1)$ is supported at $D = -1, 1$ with values $-1, 1$. Therefore

$$\begin{aligned} \Delta I_H &= \mathrm{sgn}(-c) + L_{c-1} - L_{c+1} = \tfrac{1}{2}\big(\sigma(c-1) + \sigma(c)\big), \\ I_H &= \tfrac{1}{2} \sum_c d_c \big(\sigma(c-1) + \sigma(c)\big), \quad I_R = \tfrac{1}{2} \sum_c d_c \big(\sigma(c) - \sigma(c-1)\big). \end{aligned} \tag{11.20}$$

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 [59, 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

$$W = \operatorname{Spec}(\mathcal{O}(V)[s,\tau]/(X - s\tau)), \qquad \upsilon = \tau v(t,p). \tag{11.24}$$

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 as well. The characteristic-polynomial equation for the support cycle implies, by Cayley–Hamilton, that the quotient is killed by the dth power of the section coordinate. Formal étaleness identifies these finite neighborhoods and the quotient diagrams, and the characteristic-polynomial equations are retained under this identification. The annihilation condition alone is not substituted for the support-cycle condition. Indeed, the characteristic-polynomial equations give zero positive power traces. Since $\upsilon = h(\tau)$ has zero constant term and invertible linear coefficient, Newton identities give the same equations for υ; the inverse coordinate change proves the converse.

The supported fixed-determinant gluing construction in Proposition 4.13 applies to these two maps. Both targets are the same affine reference curve in the same $\mathbf{P}^1$, with the same transverse monomial baseline. Thus the relative four-Hom comparison is tensored with the identity of that fixed reference baseline line and point coefficient. 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

$$M(G) = \det \mathrm{fib}\big(\mathrm{RHom}_{S_K}(G, G) \longrightarrow R\Gamma(S_K, G)\big).$$

Use the equivariant surface and transported orientation of Lemma 4.11. 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. For different reference curves, divide first by their actual baseline DT point objects, as in (4.23); do not identify baseline graded lines of different parity. Fix each baseline unit calibration once. The parity conversion is a separate fixed choice: [29, §7.8, (7.30)–(7.32)] chooses a map $a : \Phi_{\mathrm{pt},o_{\mathrm{odd}}} \simeq \mathbf{Q}$ satisfying $\mathrm{TS} \circ (a^{-1} \otimes a^{-1}) = \mathrm{id}_{\mathbf{Q}}$. The underlying complex orientation-line map for two odd factors induces a map of μ_2-torsors and hence a rational map of their sign local systems; it is not multiplication of rational coefficients by i. For the specified quadratic block, [30, Lemma 3.4] identifies this map with the compact Thom and product-thimble calibration of scalar $+1$. Use the monodromic Hodge realization of these same maps from Section 5. Ordered products use the Thom–Sebastiani products of the already fixed units, so associativity gives all regrouping and refinement maps. Identity comparisons are identities, and comparisons through a common reference compose. This supplies the baseline cancellation in (11.25); it does not produce a canonical isometry between arbitrary baseline quadratic lines. 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 [49, 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 $\mathrm{wt}(A_a) \le 0$. 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 $\sum h_br_b$. 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. Choose on $Y = \mathbf{P}_C(\mathcal{O}_C \oplus K_C \oplus \mathcal{O}_C)$ a polarization $\mathcal{O}_Y(k) \otimes \pi^* M$, with $k > 0$ and M sufficiently positive. Its tautological factor is trivial on X, so a compact pair of invariants (r, χ) has Hilbert polynomial $r \deg(M) n + \chi$. Given a DVR R, a generic stable pair, and an extension of its positive spectral coefficients over R, the projective stable-pair moduli on Y extend the pair over R. After a finite extension of the fraction field, its generic cycle lifts to one of the finite maps (11.30); choose a DVR R' extending the valuation. Properness of that finite map extends the lift over R'. Its monic spectral factors, together with their scalar z coordinates, define a closed family $Z \subset X_{R'}$, proper over R', containing the generic support. A power of the ideal of Z annihilates the generic sheaf. Its image in the extended sheaf vanishes as well: it is generically zero and the sheaf is R'-flat. Thus the whole extended support lies in Z, including every zero-dimensional defect. Boundary avoidance descends along the faithfully flat map $\operatorname{Spec} R' \to \operatorname{Spec} R$. The positive coefficients of the extension are the prescribed ones by separatedness, and uniqueness follows from the projective moduli. This proves the valuative criterion.

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 [58]. 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 (10.14).

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 $q_m : C^m \to \mathrm{Sym}^m C$ be the finite quotient by S_m. For the unshifted constant Hodge complexes, the adjunction map gives

$$\mathbf{Q}^H_{\mathrm{Sym}^m C} \simeq \left(Rq_{m,*}\mathbf{Q}^H_{C^m}\right)^{S_m}.$$

Indeed, the finite direct image $Rq_{m,*}\mathbf{Q}^H_{C^m}[m]$ is a Hodge module. Its averaging idempotent $e_m = (m!)^{-1}\sum_{\sigma\in S_m}\sigma$ splits in this abelian category. The rational realization of $\operatorname{im} e_m$ is $\mathbf{Q}_{\mathrm{Sym}^m C}[m]$: each fiber is one S_m-orbit and has no higher rational cohomology. The shifted adjunction map induces this isomorphism on rational realizations. Conservativity of rational realization, followed by shifting back, proves the displayed isomorphism. Consequently, Hodge

Künneth and the geometric permutation action give

$$H^j(\mathrm{Sym}^m C, \mathbf{Q}) \simeq \bigoplus_{\substack{a,b,c\geq 0\\ a+b+c=m\\ b+2c=j}} \bigwedge^{b} H^1(C,\mathbf{Q}) \otimes \mathbf{T}^c$$

as rational Hodge structures [36, Theorem 1 and §1.12]. Here the action on the Künneth factors uses the Koszul sign, so the degree-one factor contributes an exterior power. With the alternating-class convention above, and with $\bigwedge_{-y} H^1(C) := \sum_{b=0}^{2g}(-y)^b[\bigwedge^b H^1(C)]$, this yields

$$\zeta_C(y) = \sum_{m\geq 0} [H^*(\mathrm{Sym}^m C)]\, y^m = \frac{\bigwedge_{-y} H^1(C)}{(1-y)(1-\mathbf{T}y)}. \tag{12.1}$$

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

$$P_i^K = {}^p\mathcal{H}^i Rh_{K,*}\mathbf{Q}_{M_K}[2N_r], \qquad A_i = [R\Gamma_c(\mathcal{A}_r^{K_C}, P_i^K)]. \tag{12.4}$$

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.** Let $h_M : M \to \mathcal{A}$ be proper, with M smooth semiprojective of dimension $2N$, and put $P_i = {}^{p}\mathcal{H}^i Rh_{M,*}\mathbf{Q}_M[2N]$. If $F(x) = \sum_i (-x)^i [R\Gamma_c(\mathcal{A}, P_i)]$, the extraction and duality below apply to these data. For the present application we retain the notation $(M_K, \mathcal{A}_r^{K_C}, N_r, P_i^K, F_r)$ from (12.8).

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 [24, Corollary 1.3.2]. Smooth duality gives the same purity for $H_c^D(M_K)$. This is independent of the stable-pair purity proved in Section 11. The proper decomposition theorem in Hodge modules makes $H_c^j(\mathcal{A}_r^{K_C}, P_i^K)$ a direct summand of $H_c^{j+2N_r+i}(M_K)$. It is therefore pure of weight $2N_r + i + j$. Distinct j have distinct weights, so (12.8) implies

$$[H_c^j(\mathcal{A}_r^{K_C}, P_i^K)] = (-1)^{i+j} \operatorname{Gr}^W_{2N_r+i+j}\big([x^i]F_r(x)\big). \tag{12.10}$$

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 $H_c^j(\mathcal{A}_r^{K_C}, P_i^K)$ 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 $2g$,

$$\zeta_C(y; s) = \frac{(1 - sy)^{2g}}{(1 - y)(1 - s^2 y)}. \tag{12.14}$$

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

$$[R^r] \operatorname{Log} Z_{\mathrm{nil}}(R, x; s) = \frac{x s^{-2N_r}}{(1 - x)(1 - xs^2)} \sum_i (-x)^i E_c(P_i^K; s). \tag{12.16}$$

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

$$\overline{H}_r(z,w)=(zw)^{-2N_r}\sum_i(-z^{-2})^iE_c(P_i^K;zw). \tag{12.19}$$

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

$$\overline{H}_r(z,w)=\sum_{D,k}\beta_{D,k}(-1)^Dz^{2k-D}w^{2N_r-D}. \tag{12.20}$$

12.4. **Proof of Theorem 3.1.** Under nonabelian Hodge theory, the $P=W$ theorem [40, 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 [51, 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 [38, 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 $D>2N_r$ [4, Section 4.1]. The ordinary weight bound for smooth varieties gives $2k\geq D$ whenever $\beta_{D,k}\neq0$. 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

$$\nu:\mathcal{X}=C[\sqrt[s_1]{p_1},\dots,\sqrt[s_n]{p_m}]\longrightarrow C$$

is the tame root stack, E_p is its reduced root divisor, and $K=K_{\mathcal{X}}$. Thus $\nu^*p=s_pE_p$. A local chart at p is $[\mathrm{Spec}\,\mathbf{C}[t]/\mu_{s_p}]$, 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. With E required to be a bundle, 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. In families, T is finitely presented and flat over the parameter scheme, with proper zero-dimensional fibers, and projectivity is relative to that scheme. Nonreduced finite modules are included.

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 additivity gives the formula for every finite module.

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 [34, Proposition 4.20]. Equivariance, a fixed character class, and invariance of the kernel under a fixed algebra action on the source 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, and a is chosen nonzero at every root point. If $s_p > 1$, the map $b : K \to L$ vanishes on the residual gerbe because its source and target inertia characters differ. 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 regular-character elementary correspondence below imposes that the L-field act by one common scalar on all character lines of the quotient. Since L is a coarse pullback, its inertia is trivial and this condition is equivariant; it is an additional condition on the correspondence.

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.

In the following window and nilpotence arguments the full ordinary source means the normalized object $\mathcal{A}_\chi(c)$ of (7.18), with the full root class retained. Likewise the comparison cofiber means

$$K_\chi(c) = \operatorname{cofib} \Sigma_\chi(c), \qquad \Sigma_\chi(c) = \Delta_\chi(c) \otimes \Lambda^\chi,$$

with the profile suppressed from the notation. All quotient-category statements below concern these normalized objects.

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) *for each fixed profile, the normalized full ordinary framed source $\mathcal{A}_\chi(c_r)$ modulo $\mathscr{S}_r$ and the normalized comparison cofiber $K_\chi(c_r)$ modulo $\mathscr{I}_r$ vanish outside finite intervals of q in every perverse degree, with endpoints independent of that degree. Here $\mathscr{S}_r$ and $\mathscr{I}_r$ are the allowed-support and lower-rank Serre categories with the full root numerical class retained.*

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 $\operatorname{Hom}(T, K_S^{-1}) = \operatorname{Ext}^1(T, K_S^{-1}) = 0$; exact finite-group invariants preserve these vanishings.

The source statement is established simultaneously with the source-support argument of Corollary 13.6, by induction on the framed rank. The induction includes every Euler characteristic and profile, every cutoff $c \leq c_r$, and the complete chamber and wall coefficients with $t > -1/(2S_0 r)$. For $0 < e < r$, this parameter interval is contained in $t > -1/(2S_0 e)$, and the nesting already proved compares a rank-e factor at the present cutoff with its own core.

The same-rank unframed input is independent of this induction: Lemma 14.4 applies the smooth linear-stack integrality theorem to all ranks at fixed centered slope before extracting the desired coefficient. Together with finite spectral addition, it places every positive unframed coefficient in the allowed-support category, without using Corollary 13.6.

The proper coherent receivers and their costalk triangles now transport the full ordinary source, after split dimensional reduction, to any of the stated chamber coefficients modulo $\mathscr{S}_r$. A closed term either contains a framed factor of rank $0 < e < r$, covered by the induction hypothesis, or is the framed vacuum with its rank-r semistable unframed complement, covered by the preceding

independent input. Finite spectral addition preserves the allowed partition supports. Torsion at the zero wall is absent by the positive-mass calculation above. The bounded lists and projective receivers are supplied by the root versions of Lemma 6.3 and Lemma 13.1. Thus all of these restriction maps become isomorphisms after perverse cohomology and passage to the quotient by $\mathscr{S}_r$.

For the source endpoint only classical chamber duality is needed. The field transformation (6.13) is equivariant on a root chart, preserves the boundary framing, and reverses centered charge by (13.7). The saturated-subobject and pure-quotient argument therefore gives

$$t_0 X^{t}_{r,\chi,\boldsymbol{\mu}} \simeq t_0 X^{-t}_{r,-\chi,\boldsymbol{\mu}^\vee}.$$

Choose $0 < \varepsilon < 1/(2S_0 r)$. In the coherent projection sequence, the cyclic part satisfies the lower bound already proved. Every positive-rank HN quotient of the other part has centered slope at least t. A finite slope of rank at most r which is at least $-\varepsilon$ is nonnegative, by the charge lattice. Hence both complete immediate chambers lie in their cores. If $q > -A_r$, the reflected class has charge $-q < A_r$, and its positive immediate chamber is empty. Classical duality makes the negative immediate chamber of the original class empty. The preceding Hall transport modulo $\mathscr{S}_r$ then puts the full ordinary source in $D_{\mathscr{S}_r}$. Together with the empty lower core this gives the band $[A_r, -A_r]$ for its image in the quotient, uniformly in the perverse degree. No comparison of coefficients on two nonempty chambers is used. Duality changes the profile unless that profile is self-dual.

For the upper endpoint of the comparison cofiber, a semistable positive L-Higgs bundle E of rank f satisfies

$$\mu^{\mathrm{orb}}_{\min}(E) \geq \frac{q(E)}{f} + \frac{k}{2} - (f-1)\ell.$$

Indeed a gap greater than ℓ in its underlying bundle HN filtration would make the upper part Higgs invariant. Serre duality therefore gives $H^1(E) = 0$ when $q(E)/f > k/2 + (f-1)\ell$. Choose M_r larger than $\max\{0, k/2 + (r-1)\ell\}$ and the projective-framing threshold, and take $t = q/r + \delta$ with $0 < \delta < 1/(4S_0 r^2)$. If the projected Higgs bundle of a t-semistable physical object with $q > rM_r$ were unstable, its highest Higgs HN factor H would have centered slope greater than q/r by at least $1/(2S_0 rf)$, where $f = \operatorname{rk} H \leq r$. Thus its slope is greater than t. The displayed estimate, applied to the semistable factor H, gives $H^1(H) = 0$ and hence zero dual framing on H. The moment relation makes its off-diagonal A-field a Higgs morphism, which vanishes by negativity of A and decreasing HN slopes. It would therefore be an unframed physical subobject of slope greater than t, a contradiction.

The underlying positive Higgs bundle is consequently semistable, its dual framing vanishes, and its section is nonzero. Its remaining framed stability is the Joyce–Song condition. The resulting ordinary pair scheme is smooth, its forgetful morphism to the positive semistable moduli is projective, and the comparison on this chamber is an isomorphism.

Here is the root justification of the Joyce–Song endpoint. Put $S = \prod_p s_p$, and use the generating bundle $\mathcal{G}$ of (13.5). Averaging the ages over its character summands gives, for every root bundle E,

$$\chi(E \otimes \mathcal{G}^\vee) = S\big(\deg_{\mathrm{orb}} E + r(1-g) - r\kappa\big) = S\big(q(E) - r\kappa/2\big).$$

Compactify S_L to the fixed smooth projective root surface $\overline{S} = \mathbf{P}_{\mathcal{X}}(\mathcal{O} \oplus L)$, with projection π and infinity divisor D_∞. Use $\pi^*\mathcal{G}$ as generating sheaf and a coarse ample polarization $\mathcal{O}(D_\infty) \otimes \pi^*\nu^* H$, where $\deg H = d$ is sufficiently large. For a spectral sheaf F whose support avoids infinity, with $\pi_* F = E$, its modified Hilbert polynomial is

$$P_{\pi^*\mathcal{G}}(F)(m) = Srd\, m + S\big(q(E) - r\kappa/2\big).$$

The reduced modified slope is therefore $(q(E)/r - \kappa/2)/d$; this holds also for every spectral subsheaf. For a sufficiently small generic positive constant η, the η-pair inequality is consequently

equivalent to centered-slope semistability of E, together with the condition that no proper equal-slope Higgs subbundle contain the nonzero section. Indeed the discrete slope gaps preserve all strict inequalities, while an equal-slope subobject containing the section has the larger term $\eta/(Sfd)$ when its rank is $f < r$.

Apply [35, Theorem 4.24], in its nondegenerate-pair scope, on this fixed $\overline{S}$, with fixed source $\mathcal{O}_{\overline{S}}$. Its family convention is [35, Definition 2.40]. Choosing η in the first positive chamber makes semistability equal stability, as in [35, Theorems 4.4 and 4.26]. For sufficiently small η, all underlying sheaves are semistable, so forgetting the section maps this projective pair scheme to the separated coarse sheaf moduli and is projective. The locus of sheaves whose support avoids D_∞ is saturated for S-equivalence: support is the union of the supports of the Jordan–Hölder factors. Restrict to this open and to the specified root class. Its inverse image is precisely the Joyce–Song pair space above. Thus its morphism to $M^L_{r,\chi,\boldsymbol{\mu}}$ is projective by base change. This uses no relative version of Lin's theorem.

Finally, on the positive semistable stack,

$$\operatorname{Ext}^2_{S_L}(F,F) = \operatorname{Hom}_L(E, E\otimes A)^\vee = 0,$$

so that stack is smooth. In the charge range already established, $H^1(\mathcal{X}, E) = 0$ on every fiber. Proper perfect base change on the root curve makes $R\pi_{\mathrm{univ},*}E$ a vector bundle in degree zero, compatible with arbitrary ordinary base change. The stack of sections is its total space, and the Joyce–Song locus is open. An endomorphism killing the nonzero section cannot be invertible. If it were nonzero, semistability would make its nonzero proper kernel an equal-slope Higgs subbundle containing the section, contradicting the Joyce–Song condition. Thus the fixed-section stabilizers are trivial. The pair scheme is therefore smooth. Its Hitchin map is proper, since the positive coarse Hitchin map is proper. With the physical presentations and orientations of Proposition 13.3, the local comparison and proper vanishing-cycle exchange on this chamber are isomorphisms.

The root proper-receiver argument transports the comparison cofiber from the core to this chamber modulo $\mathscr{I}_r$: all intervening walls have $\lambda > \max\{q/r, 0\}$, so a rank-zero framed factor is excluded and every positive-rank factor has rank less than r. The resulting cofiber belongs to $D_{\mathscr{I}_r}$ for $q > rM_r$. Thus its image in the quotient has the band $[A_r, rM_r]$. Both endpoint constructions concern objects and restriction maps, so their bounds are independent of the 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 $S_L = \operatorname{Tot}_{\mathcal{X}} L$ and a compactly supported spectral sheaf G, let

$$B_{\mathcal{X}}(G) = \operatorname{fib}\big(\operatorname{RHom}_{S_L}(G,G) \longrightarrow R\Gamma(S_L, G)\big), \qquad M_{\mathcal{X}}(G) = \det B_{\mathcal{X}}(G). \tag{13.9}$$

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 4.3 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 physical orientation is $M_{\mathcal{X}}$ in* (13.9)*. The supported determinant and parameter comparisons inherit the baseline-relative convention of* (4.23)*: take the relevant torus invariants, divide the half-line by its actual baseline half-line, and divide the coefficient by its actual baseline DT point object. A common fixed reference baseline is restored by its identity. Comparisons between different root or ordinary curves use the separately fixed coefficient units, or specified same-parity baseline isometries when an absolute graded root comparison is wanted. No different-parity graded baseline lines are identified.*

Proof. The construction first uses the intrinsic incidence blow-up of $\mathbf{P}(\ker q)$, the full-framing descent and open extension of Lemma 4.2, and their actual adjunction and descent units and counits. It then uses the relative diagonal of $\mathbf{P}^1$, the monic spectral resolution, and the monad identification of Lemma 4.5. These operations are functorial for equivariant sheaves on every root-curve chart; the blow-up theorem used above is itself a theorem for derived Artin stacks and quasi-smooth centers. Thus the full enhanced functors and their inverse data descend before taking any invariant summands. In characteristic zero, finite-group invariants are exact and commute with extension of a parameter algebra on which the root group acts trivially. For perfect complexes over that algebra, averaging identifies the dual of the invariant summand with the invariant summand of the dual. Evaluation against the framing, composition, and relative Serre duality with its canonical twist and proper trace supply the paired diagram; the invariant-summand identity applies to its parameter-perfect terms. 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 global closed-form construction is Proposition 4.7, applied intrinsically to the fixed tame curve $\mathcal{X}$ and fixed physical bundle N, with $\det N = K_{\mathcal{X}}$. It verifies the mapping stack, finiteness, all-base framing and duality hypotheses also for the nonreduced boundary thickenings. Its two specified paths and class $2\,\mathrm{Ch}_2$ are the ones just used in the proper trace calculation. The monad projection and its morphism map are equivariant, so they realize the same full framing lift. With (13.10), its leading pairing is the root Serre pairing above. Thus an intrinsic global closed form has been constructed before the following Cartan comparison; it is not inferred from invariance of a leading pairing.

The closed-form comparison is equally equivariant. The normal Euler vector field commutes with the root group. In the notation of the proof of 4.13, 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. Take proper sections on $\overline{Y} = \mathbf{P}_{\mathcal{X}}(\mathcal{O}_{\mathcal{X}} \oplus N)$, where $N = L \oplus A$ or $N = K \oplus \mathcal{O}_{\mathcal{X}}$ is the physical rank-two bundle. Over a parameter base B, write $p_B : \overline{Y} \times B \to B$ and let $J = \widetilde{T}[-1]$ be the finite-support physical cofiber. Both

$$Rp_{B*}\mathcal{O}_{\overline{Y}\times B}, \qquad Rp_{B*}R\mathcal{H}om_{\overline{Y}\times B}(J, J)$$

are perfect. The first is $R\Gamma(\mathcal{X}, \mathcal{O}_{\mathcal{X}}) \otimes \mathcal{O}_B$; the second is a proper pushforward of a relative perfect complex with finite support. The unit u from the first complex to the second and the graded supertrace v in the reverse direction satisfy $vu = \mathrm{rk}(J)\,\mathrm{id} = 0$ already in the sheaf-level trace diagram. Finite free presentations therefore give

$$\det(1 - uv) = \det(1 - vu) = 1.$$

The unit, trace, and supported-to-proper counit commute with base change and exact root-group invariants. This is the determinant correction for the derivative of the same supported gluing functor, proving the baseline-relative supported comparison. The four-Hom identity compares $M_{\mathcal{X}}(G)/M_{\mathcal{X}}(G_0)$, and, after taking torus invariants in each entry, the corresponding relative fixed half-lines. Acyclicity of the fixed physical baseline trivializes its full determinant, not its half root. Use only the common-baseline identity or the separately specified baseline data from (4.23) to restore absolute objects.

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.

Lemma 13.4 (The enhanced moving regular quotient at a root)**.** *Fix an equivariant curve chart at a root of order s and a bundle of character multiplicities μ_c, with $\sum_c \mu_c = r > 0$. The moving regular-character quotient problem used in* (13.12) *has a smooth total space. On the meromorphic presentations with the jet vanishings used below, its one-step field-and-section correspondence has a canonical smooth derived presentation. It restricts to the ordinary correspondence away from the root. Its description and its identification with the exact-triangle correspondence commute with arbitrary derived base change; subsequent fiber products are taken as derived fiber products.*

Proof. Relabel the character indices cyclically so that multiplication by the chart coordinate t sends the cth character to the $(c+1)$st. Write the equivariantly trivialized bundle as $E = \bigoplus_c \mathcal{O}[t] \otimes V_c$, where $\dim V_c = \mu_c$. A regular quotient consists of character lines T_c, maps

$$z_c : T_c \longrightarrow T_{c+1}, \qquad q_c : V_c \longrightarrow T_c, \qquad c \in \mathbf{Z}/s\mathbf{Z}.$$

The maps q_c must generate $T = \bigoplus_c T_c$ as an $\mathcal{O}[t]$-module. This is an open condition: the vectors $t^j q_c(v)$, $0 \le j < s$, must span its finite locally free underlying base module. In frames for the T_c, the equation $t^s = a$ is the single equation

$$a = \prod_{c \in \mathbf{Z}/s\mathbf{Z}} z_c.$$

Thus, after eliminating a, the framed quotient space is the generating open U in

$$\mathbf{A}^s_{(z_c)} \times \bigoplus_c V_c^*.$$

It is smooth and irreducible. The group $G = (\mathbf{G}_m)^s$ changes the frames of the character lines. Its action on U is free: an automorphism which fixes all z_c and all q_c fixes every generator of T, and is consequently the identity over the base ring. The quotient U/G, identified with the fine quotient functor, is therefore smooth. The locus where every z_c is invertible is nonempty and dense in U, and hence in U/G. Consequently the schematic closure of the ordinary-moving-point incidence is this whole regular-quotient space, with multiplicity one. This also applies when some V_c are zero.

For $s = 1$ this presentation is $\mathbf{A}^1_a \times \mathbf{P}(V_0^*)$. At $a = 0$, several z_c may vanish, and T may split into several character intervals; all these quotients are included by the generating-open condition. The fiber $\prod z_c = 0$ can be singular. The construction above uses its smooth total space. The kernel $E_- \to E$ is locally free on every geometric curve fiber, since a submodule of a free module over the regular curve local ring is free. The quotient is flat over the parameter base. Equivalently, the monic resolution

$$[\, \mathcal{O}[t] \otimes T\, dt \xrightarrow{t - Z} \mathcal{O}[t] \otimes T\,]$$

is a relative perfect resolution of T as a curve module; the symbol dt retains the character of t in the first term. The resulting kernel is perfect and has vector-bundle fibers; the fiberwise Tor-amplitude criterion makes it a relative vector bundle. This argument retains kernels whose finite quotient has torsion or nonzero nilpotent t-action.

We next specify the enhancement, rather than infer it from classical smoothness. Put $S_D = \mathrm{Tot}_{\mathcal{X}} L(D)$ and $R_D = p^*\mathcal{O}_{\mathcal{X}}(-D)$. The smooth exterior-object stack on this chart is

$$P_D^{\mathrm{reg}} = [\, \mathbf{A}^s_{(z_c)} \times \mathbf{A}^1_\lambda \,/\, G\,], \qquad a = \prod_c z_c,$$

with its universal curve module T and scalar spectral action $w = \lambda$. It has a canonical morphism to the derived stack of perfect spectral sheaves. Define the one-step enhancement by the exact triangles in the stable arrow category with cofiber $(0 \to T)$, base changed along this specified morphism, and with the open condition that the curve pushdowns at the two ends are bundles. The source R_D and its identity map are fixed by homotopy fibers over the specified source object

and map; automorphisms of the source are not included. In the following calculation, the monic spectral resolution is

$$[\,p^*(E \otimes L(D)^{-1}) \xrightarrow{w-\Phi} p^*E\,].$$

This is a different resolution from the monic curve resolution above. Here $L(D)$ has trivial root inertia. The scalar condition on the quotient is imposed as part of the exterior object; trivial inertia alone does not force a common scalar on distinct character intervals.

On the chart above, this is the derived zero locus of

$$q\Phi - \lambda q = 0, \qquad q\sigma = 0. \tag{13.13}$$

Here σ denotes the meromorphic section, to distinguish it from the root order s.

To check this assertion as an assertion about derived functors, use the monic spectral resolution of the upper spectral sheaf F_+ and apply $R\mathcal{H}om_{S_D}(-,T)$. Proper sections and exact invariants identify its mapping complex with

$$[\,\mathrm{Hom}_{\mathcal{X}}(E,T) \xrightarrow{q\mapsto q\Phi-\lambda q} \mathrm{Hom}_{\mathcal{X}}(E,T\otimes L(D))\,]$$

in degrees zero and one. The finite support makes these Hom spaces vector bundles on the parameter chart, and there are no higher curve-cohomology terms. Moreover

$$R\,\mathrm{Hom}_{S_D}(R_D,T) = \Gamma(\mathcal{X},T(D))$$

is concentrated in degree zero and has rank one. Taking the fiber of composition with σ adds the second equation in (13.13) in degree one. A morphism of stable arrows with this specified cofiber determines its exact triangle by the functorial fiber construction. This proves the displayed derived zero-locus presentation, including mapping spaces and the section homotopy. Both monic resolutions, these fibers, and the invariant projector are defined before base change, so the identification holds after arbitrary connective derived base change.

The jet restriction map for Φ surjects onto $\mathrm{Hom}_{\mathcal{X}}(E,T\otimes L(D))$. Indeed, locally E is equivariantly free, q is surjective, and the chosen jet vanishing makes the global restriction onto the relevant finite jet surjective. The independent section restriction surjects onto $\Gamma(T(D))$. The derivative of the two equations in these field and section directions is therefore surjective. Their derived zero locus is smooth and has the expected dimension. Adding the dual-residue variables adds smooth factors, because D avoids the active graph. In fact these solutions form affine and vector bundles over the quotient chart with its scalar λ. Hence the ordinary-moving-point locus is still dense after imposing the field and section equations and after restricting a nonempty window open. The prescribed schematic closure with its fundamental multiplicities is the classical truncation of this specified smooth enhancement.

The asserted smoothness is that of the total moving correspondence over $\mathbf{C}$. The morphism to the moving coarse coordinate $a = \prod_c z_c$ need not be smooth. All presentations and exact-triangle identifications commute with derived base change, but no smoothness of the fixed-root fiber is asserted. For the two words considered here, the other mark y is fixed and ordinary. Thus the product over the moving-curve base of the moving correspondence and the constant y-correspondence is smooth, and its derived pullback along the regular diagonal of the smooth framed endpoint is quasi-smooth.

Changes of equivariant frames and curve coordinates act on the universal exact triangle and its monic resolutions. The identifications just made consequently descend from these charts. Under the ordinary-point identification the exterior stack is $S_D \times B\mathbf{G}_m$ and the construction is precisely (7.5). □

Proposition 13.5. *With these elementary quotients the whole-arrow Weyl relation of Proposition 7.1 holds on the root curve:*

$$F_\chi^c E_\chi^{c-1} - E_{\chi-1}^c F_{\chi-1}^{c-1} = r\,\rho_\chi(c-2,c). \tag{13.14}$$

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.

Lemma 13.4 supplies the smooth enhanced one-step correspondences at a root. Form each word as the derived fiber product over its framed smooth middle endpoint, retaining the moving coarse coordinate. Each word is therefore a derived pullback of a regular diagonal and carries the quasi-smooth fundamental morphism and binary composition path used in Section 7. This construction is performed on the total moving family, including its root fibers.

On $x \neq y$, the coarse supports of T_x and T_y are disjoint. All their derived cross-Homs vanish. The fiber and pushout constructions in the stable arrow category therefore give inverse equivalences between the two words, over every derived base, and preserve their maps to the exterior-object stacks. In the jet presentation, the two ordered zero loci are the same derived zero locus of the direct sum of the two independent equation bundles. The permutation of these two bundles identifies the regular-diagonal fundamental morphisms with the binary path fixed in Section 7: the coordinate-class sign and localization-boundary sign cancel. The units and counits of the functorial fiber–pushout equivalence specify the homotopy of these morphisms. Thus the properly pushed classes agree off $x = y$ with a specified homotopy, including over the root fibers.

Here the fixed removal point y is ordinary. Choose an ordinary open neighborhood V of y. Over $x \in V$, both active quotients are length-one quotients at ordinary points. The argument of Section 7 therefore gives the same derived fiber–pushout equivalence on the whole complement of the endpoint graph, including the locus $x = y$ with unequal complete endpoints. The opens $C \setminus \{y\}$ and V cover the moving curve. On their overlap both constructions are the same functorial stable fiber and pushout construction; their universal-property coherence identifies the fundamental morphisms and the specified binary paths. They consequently define the off-graph homotopy on the entire carrier. Localization can now be taken with support on the endpoint graph itself.

For clarity, the proper carrier on which this comparison is made can be constructed without using a special-fiber cycle. For a fixed final meromorphic endpoint, write $\Phi : E \to E \otimes L(D)$ and $s \in H^0(E(D))$. Let $H = \nu^*(x+y)$ and let σ_H be its canonical section. Consider the intermediate kernels

$$E(-H) \subset G \subset E(H), \quad \chi(G) = \chi(E), \quad \mathrm{profile}(G) = \boldsymbol{\mu}, \quad \Phi(G) \subset G \otimes L(D), \tag{13.15}$$

with $\sigma_H s$ factoring through $G(D)$. Near D, multiplication by σ_H identifies E with G; use this identification for the full frame and the dual-residue variables. The field and section residues then agree. 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 this Quot receiver projective; the last conditions are closed. Each Hecke word maps properly to it by multiplying its initial bundle by σ_H: the intervening bundles form flags in the same finite module, with closed field and section conditions.

Use the actual endpoint carrier of Section 7: in the framed product of the level-two initial and level-zero final ambients, take the scheme-theoretic closed union of the two word images and the complete-endpoint graph, retaining their common parameters. The directed levels contain all word endpoints. Full framing makes the initial ambient separated, so the word maps to this product and their images are proper over the final ambient. Descend the union by the common diagonal frame group. This supplies both endpoint maps on the carrier. Off H the fields coincide; cancellation of its relative Cartier equation on the locally free curve modules proves equality of the trace potentials over every base. Together with the transported full frame and dual-residue variables, this gives equality of the functions and of the residue-zero closed schemes. These identities descend to the scheme-theoretic images and their union, giving (7.10) for this root carrier.

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.

The preceding equality is obtained in the enhanced Hodge Borel–Moore mapping spectrum, with its localization lift and specified binary paths. Smooth purity on the endpoint graph reduces that lift to a degree-zero locally constant class, as in Section 7; the ordinary collision calculation determines it to be r. Apply the identification (7.13), the supported coreflector (7.14), and then the forward vanishing-cycle exchange and the proper exchange. These are the same exact functors on the enhanced diagrams as in Section 7. They carry the fundamental morphisms and their specified paths to the physical comparison arrow and its cofiber. Proposition 13.3 identifies their orientation maps. This proves the full-arrow relation with the indicated supported square homotopy.

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 $\mathbb{L}^{-1}$ factor and takes the class of y to one. This proves (13.14). □

Corollary 13.6. *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. For the normalized ordinary source, proceed by the framed-rank induction specified in Lemma 13.2, including all profiles, Euler characteristics, cutoffs, and complete chamber and wall coefficients with $t > -1/(2S_0 r)$. The rank-zero framed factor is the vacuum. All same-rank unframed coefficients have the independent support input identified in that lemma.

Fix a profile and a perverse degree. Pass to the quotient by $\mathscr{S}_r$ for $\mathcal{A}_\chi(c)$, or by $\mathscr{I}_r$ for $K_\chi(c) = \mathrm{cofib}(\Delta_\chi(c) \otimes \Lambda^\chi)$. The proper-receiver restrictions become isomorphisms in the relevant

quotient; their unique inverses transport the operators of (13.14) to the core cutoff. Restriction naturality gives $[F, E] = r\,\mathrm{id}$. The regular elementary quotient preserves the profile and increases both q and χ by one, while Lemma 13.2 leaves a finite band. Lemma 7.2 therefore makes every such perverse quotient object zero. This proves the comparison assertion and source support for the core windows, and nested restriction gives all $c \le c_r$.

The complete chamber and wall coefficients in the induction range lie in the core by the same charge-lattice and cyclic-bound argument. Their proper-receiver triangles have only smaller positive framed ranks or the framed vacuum and a rank-r unframed factor. The induction hypothesis and the independent unframed input place all closed terms in $D_{\mathscr{I}_r}$, so these triangles complete the source-support induction. A sufficiently large positive chamber gives the stable-pair conclusion. Lemma 13.1 supplies the finite graded-face and projective-receiver data, and Proposition 13.3 retains the specified orientation and finite permutation maps throughout. The dual profile is needed only to obtain an empty endpoint; the operators and the nilpotence argument each preserve the original profile. □

14. Positive root Hitchin support and canonical primitives

14.1. **Full support at fixed markings in every genus.** Keep the marked points fixed. For a profile $\boldsymbol{\mu}$ 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; common multiples of the root orders realize these weights while retaining the ordered character labels. We use perverse normalization for IC. Set

$$\begin{aligned} b_r &= \dim \mathcal{A}^L_r = \frac{\ell r(r+1)}{2} + r(1-g), \\ \gamma_r &= 1 + r(g-1) + \frac{\ell r(r-1)}{2}, \qquad B_r = \ell r^2. \end{aligned} \tag{14.1}$$

Use the same formulas with n in place of r for a lower rank. Thus $b_r + \gamma_r = B_r + 1$. The character contributions to $\chi(\mathrm{End}\, E)$ and $\chi(\mathrm{End}\, E \otimes L)$ cancel because L has trivial inertia.

Theorem 14.1. *For sufficiently large ℓ, in every genus and for every profile and Euler characteristic, $M^L_{r,\chi,\boldsymbol{\mu}}$ is irreducible with a dense stable locus. Every simple constituent of*

$${}^{p}\mathcal{H}^i Rh_{L*}\,\mathrm{IC}_{M^L_{r,\chi,\boldsymbol{\mu}}}$$

has support $\mathcal{A}^L_r$. These perverse Hodge modules are independent of χ: the identifications of the Jacobian cohomology variations over smooth spectral curves etale at all the fixed marks extend by intermediate extension. The statement includes rational weight walls, with IC *on the possibly singular coarse space.*

Here sufficiently large includes the positivity conditions of Section 13, negativity of KL^{-1}, and existence of the reduced auxiliary divisors below, also after adding one ordinary marked point. We further impose $\ell > 0$ and $\ell \ge 2g$. The new spectral argument itself only needs

$$e_n := n(\ell - 2g + 2) \ge 4 \qquad (n \ge 2). \tag{14.2}$$

These inequalities do not replace the other positivity requirements.

Lemma 14.2 (The full spectral deformation family)**.** *Let $S = \mathrm{Tot}_C(L_0)$ and let $\overline{S}$ be its smooth projective ruled compactification, with infinity section S_∞. The space $\mathcal{A}^L_n$ is the open locus, in the scheme of effective divisors of the spectral Neron–Severi class on $\overline{S}$, consisting of divisors disjoint from S_∞.*

Proof. Disjointness is open in the proper universal divisor family. A divisor in this open locus contains no ruling fiber, since every such fiber meets S_∞. Its projection to C is therefore finite of degree n, and is flat: a Cartier curve on the smooth surface is Cohen–Macaulay and has no component in a ruling fiber. For a flat family of such divisors over T, the local criterion for flatness gives finite flatness of rank n over $C \times T$. Locally on C its algebra is a quotient of $\mathcal{O}_{C\times T}[y]$. The characteristic polynomial of multiplication by y is monic of degree n; Cayley–Hamilton gives a surjection from its monic polynomial algebra onto the given algebra. Both are locally free of rank n, so this is an isomorphism. Changing the trivialization of L_0 transforms its jth coefficient as a section of L_0^j. The construction is inverse, in families, to the universal equation $y^n + a_1 y^{n-1} + \cdots + a_n = 0$. In particular, a spectral equigeneric component is open in a component of the full equigeneric family in this Neron–Severi class, even when $g > 0$. □

Lemma 14.3 (The extra codimension at a fixed mark)**.** *Let $\mathcal{A}_n^{\mathrm{int}}$ be the integral spectral locus and put*

$$\delta(a) = \gamma_n - g(\widetilde{\Sigma}_a), \qquad V_{n,d} = \{a \in \mathcal{A}_n^{\mathrm{int}} : \delta(a) = d\}.$$

Under (14.2)*, every nonempty irreducible component of $V_{n,d}$ has codimension d in $\mathcal{A}_n^L$ and contains a dense nodal open subset. If B_p is the closed locus where $\Sigma_a \to C$ is not etale above a fixed point $p \in C$, then*

$$\operatorname{codim}_{\mathcal{A}_n^L}(V_{n,d} \cap B_p) \geq d + 1. \tag{14.3}$$

The same bound holds for the union of the B_p over finitely many fixed points. For $n = 1$, every spectral curve is a section, so $\delta = 0$ and all the B_p are empty. Codimension of a constructible subset means codimension of its closure.

Proof. Let V be a component and $\varphi : T \to \overline{S}$ the normalization map of its general member, with $h = g(T) = \gamma_n - d$. Since this curve avoids infinity,

$$K_{\overline{S}}|_S = \pi^*(K_C \otimes L_0^{-1}), \qquad -K_{\overline{S}} \cdot \Sigma_a = e_n.$$

The line bundle $N = \omega_T \otimes \varphi^* \omega_{\overline{S}}^{-1}$ has degree $2h - 2 + e_n \geq 2h + 2$, hence is nonspecial and globally generated. Dedieu–Sernesi's published Corollary 2.7 applies: for a component of the full equigeneric family on a smooth projective complex surface, these conditions and this degree bound make its general member nodal.[1] Lemma 14.2 verifies the full-family hypothesis. No normal-bundle interpretation of N is assumed before this application.

On the nodal open subset, φ is an immersion and its normal bundle is $N_\varphi = N$. Deformations of this map, with its smooth source allowed to vary, have tangent and obstruction spaces $H^0(T, N_\varphi)$ and $H^1(T, N_\varphi)$. The second is zero. The map has no automorphisms, since it is birational onto its image. Near a nodal curve its image identifies this deformation space with the nodal equigeneric family: transverse double intersections persist, and normalization gives the inverse. Consequently

$$\dim V = h^0(T, N_\varphi) = h - 1 + e_n = b_n - d.$$

For every effective divisor D_T of length at most three,

$$H^1(T, N_\varphi(-D_T))^\vee = H^0(T, \varphi^* \omega_{\overline{S}}(D_T)) = 0, \tag{14.4}$$

because the last bundle has negative degree $-e_n + \deg D_T$. Thus two normal values, or a first normal jet at one point, can be prescribed independently and extend to actual deformations. These are deformations in the spectral base by Lemma 14.2.

Label a node and its two normalization preimages q_1, q_2, a finite etale choice on the nodal family. Its distinct tangent lines make the map from the ambient tangent plane to $(N_\varphi)_{q_1} \oplus (N_\varphi)_{q_2}$ an isomorphism. Surjectivity of two-point evaluation in (14.4) therefore moves the node in any

[1] T. Dedieu and E. Sernesi, *Equigeneric and equisingular families of curves on surfaces*, Publ. Mat. **61** (2017), 175–212, Corollary 2.7, p. 190; doi:10.5565/PUBLMAT_61117_07. The author preprint numbers this result Corollary 2.11.

ambient direction, allowing its preimages to move along their branches. Its projection to C is a submersion. The condition that this labeled node lie above p has codimension one.

For a smooth ramification point, take surface coordinates (x, y) with x a base coordinate. Its tangent is vertical; using $z = y$ on the branch writes the map as $z \mapsto (x(z), z)$. On the universal normalized curve over the nodal family, ramification above p is given by

$$x(z) = x(p), \qquad x'(z) = 0.$$

First-jet evaluation in (14.4) varies the constant and linear terms of x independently. These equations therefore have independent differentials, including at higher ramification. Their incidence has codimension two in a space of dimension $\dim V + 1$, and its image has dimension at most $\dim V - 1$. Nodes and these incidences exhaust failure of etaleness on the nodal locus. The nonnodal complement has smaller dimension as well. Since B_p is closed (the evaluated monic discriminant vanishes), it cannot contain V. This proves (14.3). There are finitely many equigeneric components and finitely many prescribed marks; taking their union retains the bound. No additivity for several marked conditions is asserted. □

Proof of Theorem 14.1. First assume the given weights make semistability equal stability. A divisible partial-flag type may admit no such weights; the auxiliary flag below treats that case without this assumption. Put $A = KL^{-1}$. The obstruction space is dual to $\mathrm{Hom}_{\mathrm{Higgs}}(E, E \otimes A)$, which vanishes by semistability and $\deg_{\mathrm{orb}} A < 0$. Rigidifying the scalar stabilizer gives a smooth coarse space of pure dimension $B_r + 1 = b_r + \gamma_r$. The spectral construction and semistable reduction, with the tame root generating bundle (13.5) and the projective-stack moduli construction of [34, Section 6], give a quasi-projective coarse moduli space and a proper Hitchin map. It is projective: its graph is closed in the product of the quasi-projective source with the affine base, and a proper quasi-projective morphism is projective.

Fiber dimensions. Choose a reduced divisor $D' \in |L_0 K_C^{-1}|$ containing all the marks. In the nilpotent fiber the weakly parabolic residues are nilpotent and preserve the partial flags. Refine each to a full flag strictly lowered by its residue, and add such flags at the remaining points of D'. Small generic perturbations preserve old stable objects: the finitely many subprofiles and rational weights put nonzero slope differences in a discrete set bounded away from zero. The locus of refined stable strongly parabolic nilpotent objects whose old object is stable therefore surjects onto the original nilpotent fiber.

The refined nilpotent cone has dimension at most

$$1 + r^2(g-1) + \deg(D')\frac{r(r-1)}{2} = \gamma_r. \tag{14.5}$$

Indeed, stratify by nilpotent filtration types and lift to the corresponding filtered-bundle spaces. Bundle deformations preserve the filtration and the field strictly lowers it, so their trace pairing is zero. The pullback of the Liouville one-form and hence of its symplectic differential vanishes. Each image is isotropic at a general smooth point, giving half the stable strongly parabolic dimension, namely (14.5) by Riemann–Roch. Boundedness gives finitely many filtration types. This is the filtration proof in [6, Section 3.2]; its parabolic statement explicitly permits $g = 0, 1$ whenever stable objects exist. Properness, upper semicontinuity of fiber dimension, and the contracting scaling action on the Hitchin base now give $\dim h_L^{-1}(a) \le \gamma_r$ for every a.

The weak abelian support inequality. Let $P = \mathrm{Pic}^0(\Sigma/\mathcal{A}_r^L)$ be the ordinary connected spectral Picard. The universal finite flat spectral curve has direct image $\bigoplus_{j=0}^{r-1} L_0^{-j}$ on $C \times \mathcal{A}_r^L$. Thus $h^0(\mathcal{O}_{\Sigma_a}) = 1$, $h^2(\mathcal{O}_{\Sigma_a}) = 0$, and $h^1(\mathcal{O}_{\Sigma_a}) = \gamma_r$. The relative Picard is a smooth commutative group scheme with connected fibers of dimension γ_r, including at nonreduced curves; see [12, Section 5]. Tensor product preserves every subobject slope and every root profile because a member of P has degree zero on every reduced component and has trivial inertia. Its stabilizers are affine: after restriction to each normalized reduced component, the spectral sheaf modulo

torsion has positive rank. Taking determinants makes a stabilizing line bundle torsion on that component. Its image in the product of normalization Jacobians is finite, and the kernel of the Picard map is affine. The relative Tate module of this same ordinary spectral Picard is polarizable on the whole base, not just fiberwise.[2]

We have verified the hypotheses of [39, Section 1.1 and Theorem 1.8]: a proper morphism with quasi-projective source, a smooth connected relative group of dimension γ_r, $\dim M = b_r + \gamma_r$, affine stabilizers, and a polarizable Tate module. Its additional truncation hypothesis follows here from smoothness and the fiber bound:

$$\tau_{>2\gamma_r} Rh_{L*}(\mathrm{IC}_M[-\dim M]) = \tau_{>2\gamma_r} Rh_{L*}\mathbf{Q}_M = 0.$$

Consequently every strict support Z satisfies

$$\operatorname{codim} Z \le \delta_Z^{\mathrm{aff}}, \qquad \beta_Z \le \gamma_r - b_r + \dim Z, \tag{14.6}$$

where β_Z and $\delta_Z^{\mathrm{aff}} = \gamma_r - \beta_Z$ are the generic abelian and affine dimensions of P.

Integral support and the fixed marks. At the generic point of Z factor the spectral polynomial into distinct integral factors of degrees n_i, with multiplicities m_i, so that $r = \sum_{i=1}^{s} m_i n_i$. The finite spectral multiplication map gives a finite cover of Z, whose projections have closures Z_i in the integral rank-n_i bases. The abelian part is the product of their normalization Jacobians. The unmarked codimension assertion in Lemma 14.3, including its rank-one case, gives the usual integral δ-regularity inequality on each Z_i. Together with (14.6), this gives

$$\sum_i (\gamma_{n_i} - b_{n_i} + \dim Z_i) \le \beta_Z \le \gamma_r - b_r + \dim Z \le \gamma_r - b_r + \sum_i \dim Z_i. \tag{14.7}$$

Only $\dim Z \le \sum_i \dim Z_i$ is used here, precisely the finite-factor dimension inequality of [13, Section 1]. Since $\gamma_n - b_n = 1 + n(2g - 2 - \ell)$, cancellation yields

$$1 - s \ge \left(r - \sum_i n_i \right) (\ell - 2g + 2). \tag{14.8}$$

Its left side is nonpositive and its right side nonnegative; $\ell > 2g - 2$ forces $s = 1$ and $m_1 = 1$. Thus Z is generically integral. If its generic δ is d, then (14.6) gives $\operatorname{codim} Z \le d$. If Z were contained in the union of the fixed B_p, Lemma 14.3 would instead give $\operatorname{codim} Z \ge d + 1$. Hence every support meets

$$\mathcal{A}_r^{\mathrm{good}} = \mathcal{A}_r^{\mathrm{int}} \setminus \bigcup_p B_p.$$

Compactified Jacobians on an actual open set. Over $\mathcal{A}_r^{\mathrm{good}}$, allocation of the prescribed characters to the distinct marked branches is a finite etale cover with fiber

$$\prod_p \mathfrak{S}_r / \prod_c \mathfrak{S}_{\mu_{p,c}}. \tag{14.9}$$

After this cover, tensoring by the prescribed fractional divisors removes inertia and descends the spectral sheaf to an ordinary rank-one torsion-free sheaf. This identifies each sheet, locally etale, with an ordinary compactified Jacobian family. Its Euler characteristic differs by a constant determined by the profile; integral spectral support makes every such sheaf stable.

We use [12, Sections 2 and 4, Theorem 7.1] in its precise scope: a smooth connected projective base curve of any genus, an effective twist of degree $> 2g - 2$, characteristic zero, the integral spectral locus, and arbitrary integer degree. Its compactified Jacobian has smooth total space and its perverse direct image is the intermediate extension from smooth spectral curves. The coprimality condition is imposed only later for the all-base Theorem 9.1, corrected in [13].

[2] M. A. A. de Cataldo, *A support theorem for the Hitchin fibration: the case of* SL_n, Compos. Math. **153** (2017), 1316–1347, Theorem 3.3.1; arXiv:1601.02589. The theorem concerns the ordinary GL_n spectral Picard over its entire base; the induced form on the product of normalization Jacobians is weighted by the positive component multiplicities. It therefore applies unchanged to the present Picard group scheme.

Alternatively, its arbitrary-degree use follows etale locally by choosing a section of the smooth locus of the universal integral curve, away from the marked branches. Its degree-one divisor allows tensoring to an Euler characteristic coprime to r, equivalently an underlying ordinary Higgs-bundle degree coprime to r, where that all-base theorem applies. There is no stability wall on the integral locus.

Trivialize the allocation cover etale locally. On each sheet the preceding ordinary theorem applies; etale pullback and finite etale pushforward preserve intermediate extension. Therefore no proper support can meet $\mathcal{A}_r^{\mathrm{good}}$. Every support does meet it, so all supports equal $\mathcal{A}_r^L$. The same assertion holds for pure Hodge modules by their decomposition theorem and the faithful underlying rational-perverse-sheaf functor.

Irreducibility for arbitrary weights. The stack of root bundles of fixed rank, Euler characteristic and profile is an iterated partial-flag fibration over the ordinary bundle stack, hence is irreducible. For an underlying bundle Harder–Narasimhan type with semistable pieces $E_1, \dots, E_v$ in strictly descending slopes, the stratum has codimension

$$c = \sum_{i<j} -\chi_{\mathcal{X}}(E_i, E_j) = \sum_{i<j} h^0(\mathrm{Hom}(E_j, E_i)K) \geq 0. \tag{14.10}$$

Here $\mathrm{Hom}(E_i, E_j) = 0$ for $i < j$ and Serre duality gives the equality. The excess dimension of the Higgs fiber is bounded by

$$h^1(\mathrm{End}(E)L) = \dim \mathrm{Hom}(E, E \otimes A) \leq \sum_{i<j} h^0(\mathrm{Hom}(E_j, E_i)A).$$

For each nonzero summand on the right of (14.10), choose an ordinary point where a K-valued section is nonzero and a nonzero section of L_0 vanishing there. Positivity guarantees the latter. Multiplication embeds the corresponding A-valued space into a proper subspace of the K-valued space. The excess is therefore $< c$ if $c > 0$, and zero if $c = 0$.

Only finitely many such types meet Higgs semistability: a successive slope gap $> \ell$ makes the upper piece Higgs invariant, so all gaps are at most ℓ; fixed total degree and finitely many subprofiles then bound the factor degrees. The semistable Higgs stack is smooth of pure dimension B_r, by the same obstruction vanishing, also at strictly semistable objects. Every positive codimension bundle stratum contributes dimension $< B_r$. Every component therefore meets the unique open Harder–Narasimhan stratum. Over that stratum $h^1(\mathrm{End}(E)L) = 0$, so the Higgs stack is a vector bundle and its nonempty semistable open is irreducible. The entire semistable stack and its coarse space are consequently irreducible. Stable objects exist in every class: take a smooth connected spectral curve etale at the marks, allocate the characters, and choose a spectral line of the required Euler characteristic. Integral support excludes a proper invariant generic subspace. The stable locus is thus dense. This argument allows the open bundle stratum to be unstable in genus zero.

Weight walls. Add one ordinary marked point q with a full invariant flag and rational weights $0 < \epsilon_1 < \cdots < \epsilon_r$ sufficiently small. For the old fixed class, nonzero slope differences are uniformly separated from zero. Choose the new weights so that their possible contribution is smaller than this gap. New semistability then implies old semistability, and every old stable object stays stable under every invariant full-flag refinement. For a rank-r' subobject the auxiliary wall has normal vector $r\mathbf{1}_I - r'\mathbf{1}$, where $|I| = r'$; this is nonzero for $0 < r' < r$. Only finitely many walls meet the small weight chamber, so rational weights can be chosen off all of them. The new moduli is smooth and all its semistable objects are stable.

Forgetting the flag defines $\pi : M_{\mathrm{new}} \to M_{\mathrm{old}}$ over the same Hitchin base. Both Hitchin maps are projective and the old space is separated; the closed graph of π therefore proves it projective. Over smooth spectral curves etale also at q, the map is finite etale of degree $r!$, the orderings of distinct eigenlines. Its closed image contains a dense open subset of the irreducible old space, hence is all of it. Both spaces have dimension $B_r + 1$. The trivial summand of the generic

permutation local system, split by trace divided by $r!$, and the decomposition theorem give

$$\mathrm{IC}_{M_{\mathrm{old}}} \quad \text{as a direct summand of} \quad R\pi_*\mathbf{Q}_{M_{\mathrm{new}}}[B_r+1].$$

The generic-weight proof applies to this enlarged fixed set of marks. Its full-support direct image contains the old IC direct image as a direct summand; full support follows at the wall.

Finally, on the smooth spectral allocation cover the fibers are torsors for the same Jacobian family. Translation acts trivially on cohomology, so their variations of Hodge structure are the exterior powers of the same H^1 variation, independently of degree. The allocation local system depends only on the profile. Full support and uniqueness of intermediate extension extend these identifications in each perverse degree over $\mathcal{A}_r^L$. No splitting of the entire derived direct image is required. □

14.2. **The independent unframed support input.**

Lemma 14.4. *At any fixed centered slope, every perverse direct-image coefficient of the split semistable unframed root stack belongs to the allowed-support category $\mathscr{S}_r$, with its full numerical class retained. Its rank-r primitive is the positive coarse intersection complex of Theorem 14.1. This assertion does not use a framed-source comparison.*

Proof. Fix a centered slope λ and keep the entire linear stack

$$\mathfrak{N}_\lambda = \{0\} \sqcup \coprod_{q(E)/\operatorname{rk}(E)=\lambda} \mathfrak{M}_L^{\mathrm{ss}}(r,\chi,\boldsymbol{\beta})$$

over all positive ranks and full root numerical classes. It is closed under direct sum. The same-slope category is abelian: the image of a morphism has slope at least λ by the quotient inequality, whereas its saturation has slope at most λ by semistability of the target. Nonzero torsion in the saturation quotient would strictly increase centered degree by Lemma 13.1. Thus images are saturated; kernels, cokernels and extensions remain in this category. Direct sum and scalar multiplication are coherent operations; objects graded by a finite-rank torus split into their finitely supported eigensheaves. Thus this is a linear moduli stack in the sense of [11, §2.4.2].

For two objects E, F, the Higgs Hom complex and root Riemann–Roch give

$$\chi_{\mathrm{Tot}_{\mathcal{X}}\,L}(E,F) = -\ell \operatorname{rk}(E)\operatorname{rk}(F), \qquad \operatorname{Ext}^2(E,F) = \operatorname{Hom}(F, E\otimes KL^{-1})^\vee = 0.$$

The first identity holds because the character corrections cancel under the coarse twist $L = \nu^* L_0$; the second uses the negative orbifold degree of KL^{-1}. Hence the stack is smooth. The extension fiber is $\operatorname{Ext}^1 / \operatorname{Ext}^0$, so the same bilinear formula is its Euler pairing in [11, §2.4.7]. At a polystable point $\bigoplus_i S_i \otimes V_i$, with pairwise distinct stable factors of ranks r_i,

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

Pair the opposite $\operatorname{Hom}(V_i, V_j)$ blocks by trace and their equal-dimensional multiplicity spaces by a perfect pairing; use a symmetric form on each diagonal multiplicity space. The gauge blocks have the same invariant trace pairing. This verifies pointwise almost-orthogonality exactly as in Proposition 8.2.

The generating-bundle spectral Quot construction supplies the good moduli space with affine diagonal. For each numerical class, successive underlying Harder–Narasimhan gaps are at most ℓ; otherwise the upper subbundle is Higgs invariant. Fixed total centered degree and finitely many subprofiles then bound the underlying bundles. Each component is of finite type. At fixed rank and slope only finitely many full numerical classes occur, since the profile determines χ from (13.2). A decomposition over a bounded family is a tuple of complementary idempotents in its relative endomorphism space. Their equations give a quasi-compact direct-sum map. The same argument on each fixed finite character-support face gives quasi-compact graded points, with finitely many cotangent weights, as in [11, §2.3.7]. At an ordinary point p_0, the class

$c_1(\det E|_{p_0})/r$ is the scalar torus parameter; the analogous classes for eigensheaves give the face parameters, permuted by the finite face automorphisms.

Theorem 14.1 proves irreducibility and a dense stable locus in every component. Over smooth spectral curves etale at the marks, the objects are spectral lines with their allocated characters, and their only automorphisms are scalars. Scalar rigidification is therefore generically an isomorphism onto the coarse space. The hypotheses of [11, Theorem 10.2.7 and Corollary 7.2.4] are now verified. Apply the theorem to the whole $\mathfrak{N}_\lambda$ before extracting a rank coefficient. It gives the analogue of (8.6) with the positive coarse IC primitives, the same $H^*(B\mathbf{G}_m)_{\mathrm{vir}}$, and the actual smooth Hall induction maps with their derived finite permutation actions.

In the split physical unframed stack, the negative field preserves the positive-Higgs Harder–Narasimhan filtration. Total semistability therefore forces positive-Higgs semistability, and the negative field vanishes by the same Hom inequality. The surface Hom calculation and the equivariant construction of Proposition 13.3, now without the section arrow, give the derived identification $T^*[-1]\mathfrak{N}_\lambda$. Its half is the surface Hom complex and its paired triangle gives the standard cotangent orientation. This physical comparison precedes and does not use root Hall comparison. The symmetric Euler pairing and [11, Lemma 10.2.9] make the orientation commutative. Dimensional reduction, with the normalization of Proposition 8.2, identifies its coefficient and induction maps with those of the smooth stack; the map compatibility is [11, §§6.2.3 and 8.1.7].

The coarse Hitchin maps are proper by Theorem 14.1. Spectral addition is finite by Lemma 8.3; retaining the finite root-profile allocations preserves finiteness. Coarse direct sum is also finite: it is proper over the Hitchin base and a polystable object has only finitely many allocations of its stable factors to fixed numerical classes. Proper Hitchin pushforward of the cohomological induction isomorphism therefore has only allowed partition supports. At any fixed perverse degree, only finitely many degrees of the torus factors contribute. The single rank-r unframed coefficient is included. The degree-independent perverse IC identification is the one in Theorem 14.1. No framed chamber, window or source-comparison conclusion enters this proof. □

14.3. **Active profiles and the target Euler characteristic.** Fix a target profile $\boldsymbol{\mu}$ of rank n. A profile $\boldsymbol{\beta}$ of specified rank $0 \le r \le n$ is called active if $\beta_{p,c} \le \mu_{p,c}$ for all p, c and $\sum_c \beta_{p,c} = r$ at every marked point. The active set is finite and is closed under taking summands of an active total.

Lemma 14.5. *Let $H = \{\boldsymbol{\mu} \cdot \boldsymbol{\alpha} = -d_0\}$, where $d_0 \in \mathbf{Z}$, and let $U \subset H$ be a nonempty relatively open part of a bounded admissible ordering chamber, with fixed auxiliary identity coordinates omitted. One can choose rational weights in U for which every positive-rank active profile is stable whenever it is semistable at ordinary 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

$$(r\boldsymbol{\beta}' - r'\boldsymbol{\beta}) \cdot \boldsymbol{\alpha} = r' - r\chi'. \tag{14.11}$$

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. Their finitely many proper rational intersections cannot cover the nonempty relatively open set U; rational points are dense in the complement. If H is zero-dimensional, each wall misses it. The chosen point has the claimed property. □

The auxiliary ordinary Euler characteristic need not equal the ordinary Euler characteristic of the parabolic-degree-zero target. For the generic finite-order unitary target monodromies

$\exp(-2\pi i\alpha_{p,c})$, the latter is

$$\chi_{\text{target}} = n(1-g) - \sum_{p,c} \mu_{p,c}\alpha_{p,c}. \tag{14.12}$$

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 (14.12). 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.

14.4. **The root and coarse surface halves.**

Lemma 14.6. *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

$$[\mathcal{O}\chi^c L^{-1} \xrightarrow{y-f(u)} \mathcal{O}\chi^c].$$

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. Use the actual surface-pair functor $\mathcal{D} = R\nu_{S,*}$. The monic comparison and composition with the section identify its derivative on the fibers in (13.9); the exact invariant projector commutes with derived Artin base change. The dual map is its Serre-dual cotangent map. Locally, for $e_c = t^{-c}$, the invariant section generator $t^c e_c$ pairs with $t^{s-1-c}e_c^\vee dt$ to give $t^{s-1}dt = du/s$ under (13.10). This canonical-twisted comparison is independent of c; it is not raw pushforward of an untwisted dual Hom block.

For the nonsplit row use the ordinary quotient obtained by exact pushforward of (13.8). Its extension class $\alpha_0 \in H^1(C, K_C L_0^{-2})$ is the image of α under the same trace-compatible identification. The two functions then satisfy $g_{\mathcal{X}} = \mathcal{D}^* g_C$; both split functions are zero. Through the specified physical critical equivalences, use the critical lift $\mathrm{Crit}(\mathcal{D})$ as the physical comparison. Its derivative identifies the paired triangles, so Proposition 4.8 gives the same derivative square map. The equivalent half complexes have equal virtual dimension and parity. 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 14.7. *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. Lemma 14.4 supplies the smooth-stack integrality input, its actual induction maps and finite permutation actions. Its primitive is the coarse IC object of Theorem 14.1.

On the good nodal open, Lemma 14.6 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. The degreewise reconstruction uses the same purity and weight separation as (8.18), on each allocation sheet and then under its finite-group projectors.

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 13.6 therefore supplies the source comparison needed for this extension. □

14.5. **The critical scheme and the quadratic normal bundle.** For the extension class α of (13.8), define on the smooth stable positive space

$$g(E,\Phi)=\tfrac{1}{2}\langle\alpha,\mathrm{Tr}(\Phi^2)\rangle=g_{\mathcal{A}}\circ h_L. \tag{14.13}$$

For a stable numerical class put

$$d_{\boldsymbol{\mu}}=2N_{\boldsymbol{\mu}}=2(g-1)r^2+2+\sum_p\left(r^2-\sum_c\mu_{p,c}^2\right). \tag{14.14}$$

The following assertion includes the case of an empty canonical moduli space.

Proposition 14.8. *If semistability is stability in the given class, the full classical critical scheme of* (14.13) *is*

$$\mathrm{Crit}(g)=M^K_{r,\chi,\boldsymbol{\mu}}\times\mathbf{A}^1. \tag{14.15}$$

The isomorphism sends (E,θ,λ) to $(E,b\theta+a\lambda\,\mathrm{id})$. 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:

$$C_\Phi=[\mathrm{End}\,E\xrightarrow{[\Phi,-]}\mathrm{End}\,E\otimes L],$$
$$D_\Phi=[\mathrm{End}\,E\otimes A\xrightarrow{[\Phi,-]}\mathrm{End}\,E\otimes K],$$
$$B_\Phi=[\mathrm{End}\,E\otimes(K\oplus\mathcal{O})\xrightarrow{[\Phi,-]}\mathrm{End}\,E\otimes(K\oplus\mathcal{O})L],$$
$$C_\Phi^L=[\mathrm{End}\,E\otimes L\xrightarrow{[\Phi,-]}\mathrm{End}\,E\otimes L^2].$$

There is an exact sequence of these complexes

$$0\longrightarrow D_\Phi\longrightarrow B_\Phi\longrightarrow C_\Phi^L\longrightarrow 0. \tag{14.16}$$

The tautological Φ is a closed degree-zero section of C_Φ^L. 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. Represent a tangent class by a Cech hypercocycle (u_{ij}, v_i), where u_{ij} varies the bundle transitions and v_i the field. In compatible frames $v_j - v_i = [\Phi, u_{ij}]$, so $\operatorname{Tr}(\Phi(v_j - v_i)) = \operatorname{Tr}(\Phi[\Phi, u_{ij}]) = 0$. Thus $\operatorname{Tr}(\Phi_i v_i)$ glues. Pairing $[\alpha\Phi, 0]$ with this full cocycle gives

$$\langle \alpha, \operatorname{Tr}(\Phi v)\rangle = d\Big(\tfrac{1}{2}\langle \alpha, \operatorname{Tr}(\Phi^2)\rangle\Big)\,(u, v).$$

This computes the differential on all bundle and field deformations; a pure transition deformation with $v = 0$ contributes zero on both sides.

On every geometric fiber stability and negativity give

$$H^0(C_\Phi) = \mathbf{C}\,\mathrm{id}, \quad H^2(C_\Phi) = 0, \qquad H^0(D_\Phi) = 0, \quad H^2(D_\Phi) = \mathbf{C}.$$

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

$$H^0(C_\Phi) = \mathcal{O}\,\mathrm{id}, \quad H^1(C_\Phi) = T_{M^L}, \qquad H^1(D_\Phi) = \Omega_{M^L}. \tag{14.17}$$

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 (14.16) 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 (14.17), $z = \lambda\,\mathrm{id}$ over T.

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 (14.15) 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 (14.14). Since (14.15) 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 14.9. *In the situation of Proposition 14.8, with j denoting the closed critical inclusion,*

$$\phi_g \mathbf{Q}_{M^L}[B_r+1] \simeq j_* \mathbf{Q}_{M^K\times \mathbf{A}^1}[d_{\boldsymbol{\mu}}+1] \otimes \mathbf{T}^{(B_r-d_{\boldsymbol{\mu}})/2}. \tag{14.18}$$

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

$$\det \mathrm{RHom}_{S_L}(\widetilde{E},\widetilde{E}) = \det R\Gamma \operatorname{End} E \otimes (\det R\Gamma(\operatorname{End} E \otimes L))^{-1} = \bigotimes_{p\in D} \det(\operatorname{End}(E_p)\otimes L_p)^{-1}.$$

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

$$\det \mathrm{Hess}(g) = u\,\delta^{\otimes 2}$$

for a regular unit u on the critical scheme.

This unit is constant on every connected component. The algebraic closed-immersion calculation in (14.19) is independent of this Morse argument. Since $h_L \circ j = f_{r,\boldsymbol{\mu}} \circ (h_K \times \mathrm{id})$ is proper, base change along this closed immersion makes $h_K \times \mathrm{id}$ proper. Its base has positive scaling weights. Thus scaling (θ,λ) extends every orbit morphism across zero. A unit pulled back along this morphism from $\mathbf{A}^1$ 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 of u on each connected component. Put $c = B_r - d_{\boldsymbol{\mu}}$. With the usual second-derivative convention, the normal quadratic form in [10, Theorem 5.1(ii), Definition 5.2] is $q = \frac{1}{2}\mathrm{Hess}(g)$, so $2^{-c/2}\sqrt{u}\,\delta$ is a square root of $\det q$. This trivializes the determinant square-root torsor of the normal quadratic bundle. The relative Morse lemma [10, Theorem 5.4(d), Example 2.23] proves (14.18). Changing a chosen root changes this isomorphism by a locally constant sign. Even normal rank gives semisimple monodromy $+1$, and the quadratic model has zero nilpotent logarithm.

Fix these isomorphisms for the finitely many profiles used in a given target coefficient and use their external tensor products. Naturality of the derived symmetry makes these identifications compatible with packet permutations; interchanging the even-rank normal quadratic blocks contributes no further orientation sign. This assertion concerns external products of the scalar primitive objects. The actual physical product on a coprime scalar-factorization chart is the separate enhanced gluing map of Lemmas 10.3 and 10.4, with its root descent and cross-profile derivative determinants supplied by Proposition 13.3. □

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

$$f_{r,\boldsymbol{\mu}} : \mathcal{A}^K_{\boldsymbol{\mu}} \times \mathbf{A}^1 \longrightarrow \mathcal{A}^L_r, \qquad (\theta,\lambda) \longmapsto b\theta + a\lambda \tag{14.19}$$

is a closed immersion. To see this on the coefficient space, 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 complex 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 (14.18) give

$$\phi_{g_{\mathcal{A}}} P^L_i \simeq (f_{r,\boldsymbol{\mu}})_*(P^K_i \boxtimes \mathbf{Q}_{\mathbf{A}^1}[1]) \otimes \mathbf{T}^{(B_r-d_{\boldsymbol{\mu}})/2}. \tag{14.20}$$

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 14.10. *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. For each positive-rank active profile one has*

$$[R^r X^{\boldsymbol{\mu}}] \operatorname{Log} Z^{\mathcal{X}}_{\mathrm{nil}} = \frac{x\mathbf{T}^{-N_{\boldsymbol{\mu}}}}{(1-x)(1-x\mathbf{T})} \sum_i (-x)^i [R\Gamma_c(\mathcal{A}^K_{\boldsymbol{\mu}}, P_i^K)]. \tag{14.21}$$

Initially take the P_i^K at the auxiliary Euler characteristic $\chi_* = 1$. For the prescribed target profile with generic determinant data, they may instead be taken at its integral Euler characteristic (14.12)*. They are zero if the corresponding 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 14.7 at the simultaneous generic auxiliary value $\chi_* = 1$ of Lemma 14.5. We first recover the full nonsplit physical coefficient, by the rank induction of Theorem 10.5, now retaining the root profile. Fix the target and the finite active set above; choose its simultaneous auxiliary weights at $\chi_* = 1$. The induction is over all active profiles and all integer χ. Every packet profile remains active. Receiver and dual profiles need not be active: only their all-profile support and comparison inputs are used, without requiring simultaneous genericity for them. Write $R^K_{r,\chi,\boldsymbol{\mu}}$ for that coefficient and $C_{r,\chi,\boldsymbol{\mu}}$ for the candidate obtained from the positive source by vanishing cycles and (14.20). Finite spectral addition and the weight-two scaling argument used in that theorem justify Thom–Sebastiani at zero. Corollary 13.6 and the bounded-endpoint argument of Corollary 7.5 give only

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

On the locus of at least two distinct evaluated scalar roots, polynomial idempotents commute with inertia and split the full numerical class. We need the actual packet gluing, not only its map on numerical classes. Here is the root-stack version of Lemmas 10.3 and 10.4.

Use $\overline{Y} = \mathbf{P}_{\mathcal{X}}(\mathcal{O}_{\mathcal{X}} \oplus K_{\mathcal{X}} \oplus \mathcal{O}_{\mathcal{X}})$ and its boundary divisor D_∞. The packet supports Z_i are disjoint over the whole coprime factorization chart, since all pairwise resultants are units. Put $U_0 = \overline{Y} \setminus \bigcup_i Z_i$ and $U_i = \overline{Y} \setminus \bigcup_{j\neq i} Z_j$. These are open substacks, with $U_i \cap U_j = U_0$, and all contain D_∞. Use the compatible determinant-and-boundary homotopy fiber (10.6) on this stack. The projective bundle formulas give

$$R\Gamma(\overline{Y}, \mathcal{O}) \simeq R\Gamma(D_\infty, \mathcal{O}) \simeq R\Gamma(\mathcal{X}, \mathcal{O}).$$

Thus, as in that fiber's proof, fixing the compatible determinant does not change the boundary-framed formal deformation problem. On U_0 a perfect Artin deformation of the baseline $\mathcal{O}$ is invertible, and determinant is the identity on this formal Picard problem. Its compatible fixed-determinant fiber is contractible, including the overlap trivialization and its higher automorphisms. Descent for perfect complexes and line bundles on this open cover therefore gives mutually inverse formal gluing functors on every connective Artin base. Their ordinary truncations are the polynomial-idempotent decomposition and its inverse. Each packet is glued to the same baseline $\mathcal{O}$ on its complement: the original section has one common framed source, never a direct sum of independent framing sources. These restrictions and gluing units preserve the stable-pair open condition, the complete ordered inertia profile, and all parameter and packet-permutation maps.

The closed-form comparison follows under this same gluing map. In (10.8) use categorical $CC^-(\mathrm{Perf}(V)/\mathbf{C})$ and its restriction fiber. Perfect complexes on these quasi-compact tame root stacks satisfy open localization, up to idempotent completion. Compact generation with supports and Thomason localization apply to their finite linearly reductive quotient charts and glue; the same argument works after the connective Artin base changes here, whose small etale

topoi are unchanged [67, Theorem A, Theorem 3.12, Lemma 4.10 and Example 9.4]. Keller's mixed-complex localization and then homotopy S^1-fixed points give the categorical negative-cyclic fiber sequence.

Write $Z = \coprod_i Z_i$. Supported open descent identifies $\mathrm{Perf}_Z(\overline{Y}_A)$ with $\prod_i \mathrm{Perf}_{Z_i}(\overline{Y}_A)$, with projections given by the very restrictions in the gluing functor. They send the relative point $(k(I)-k(\mathcal{O}), k(\alpha))$ to the points $(k(I_i)-k(\mathcal{O}), k(\alpha_i))$. The natural categorical character therefore gives their additive homotopy in the supported fiber, including its specified paths and all refinement and permutation compatibilities. Apply Lemma 10.2 to $\overline{Y}_A$ and $U_{0,A}$. Its map on restriction fibers carries that homotopy to the geometric relative Ch_2 homotopy with the same paths. This uses perfect complexes on the root stack itself and a natural map to geometric cyclic theory, without a coarse-invariant identification. All complexes remain over $\mathbf{C}$, so parameter differentials are retained; the finite support sum commutes with Hodge completion.

Now apply the geometric boundary transgression and proper Serre trace of Proposition 13.3 to twice this geometric class. The overlap trivializations restrict to the fixed boundary framing, and (13.10) retains $2\,\mathrm{Ch}_2$. Linearity gives the sum of the actual closed packet forms under the specified gluing map. The strict trace kernel is zero on U_0; its supported summands also identify the leading trace pairing and make every mixed-support cyclic word zero. Completion of the whole algebraic factorization chart, followed by faithful flatness in its coherent critical-section quotient, compares the algebraic d-critical sections as in Lemma 10.3.

For the orientation, form the cross complexes C_{ij} and $A_{ij} = z_j(-) - (-)z_i$ on the canonical root surface. The sheaf-level Bézout contraction makes $D_{ij} = \mathrm{fib}(A_{ij})$ canonically acyclic on all parameter bases. Root Serre duality and the paired deformation triangle identify the opposite cross complexes. The determinant double-triangle calculation in (10.10) is therefore the derivative comparison for this gluing functor, with the same canonical acyclic trivialization. Its opposite cross Euler characteristic can be odd; the graded interchange sign and the sign of $-A_{ij}^{\vee}$ cancel exactly as proved in Proposition 13.3. The self-Euler form of a compactly supported object on the canonical root surface is even, by its symplectic Ext^1 pairing. Hence the framed half has parity $\chi_i \bmod 2$, and normalization by Λ^{χ_i} gives precisely the ordinary derived packet permutation action. Exact-filtration coherence gives all refinements and stabilizers, with the common-baseline identity used throughout. Thus the original oriented coefficient has the root-profile version of (10.13). The positive-source candidate has the same decomposition by finite spectral addition, vanishing cycles and the scalar calculation, with both Euler and character-profile allocations retained on both sides. Each packet has smaller positive rank. The equivariant tensor filtration in the proof of Theorem 10.5 therefore carries the lower-rank class identities to the same finite allocation-group class, including repeated-packet stabilizers. Exact finite invariants and finite closed/open localization give equality on this whole many-root locus. The scalar error-detection argument in the last paragraph of Proposition 14.7 is injective on $K_0(\mathscr{I}_r)$; it uses $\alpha a = 0$ and the removal of the one-factor rank-r unframed vacuum. Hence the error is zero globally. The induction starts with $\mathscr{I}_1 = 0$ and is finite coefficientwise by Lemmas 13.1 and 13.2.

Now use the ambient nil restriction, finite addition and proper base change as in Corollary 10.6. On geometric points, a product of monic characteristic polynomials equals t^r only if every factor is the corresponding power of t. Hence the reduced inverse image is the product of the reduced nil loci. This identifies the actual nil series with the symmetric exponential of the scalar candidate. On restricting the scalar line to zero, the terms are exactly

$$x^{i+1+u+v} P_i^K[-i-2v] \otimes \mathbf{T}^{v-N_{\boldsymbol{\mu}}}, \qquad u, v \geq 0.$$

The scalar shift [1] in (14.20) is retained. Taking compact cohomology and summing u, v gives (14.21); 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). This calculation first lies in $K_0(\mathrm{MMHS})$. Proposition 15.3, proved from punctual geometry and the closed index

independently of this primitive formula, makes the actual nil cohomology ordinary. The exact ordinary retraction used in Corollary 10.6 then returns the equality of these ordinary endpoints to $K_0(\mathrm{MHS})$.

For the prescribed target class, use Theorem 14.1 to identify the positive perverse IC objects at $\chi_* = 1$ and at (14.12). Apply the same base vanishing-cycle functor and use (14.20) 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. Apply the left inverse $i_0^* f_{r,\boldsymbol{\mu}}^*[-1]$ and undo the common Tate twist, where i_0 is the scalar zero section. This identifies the two canonical perverse objects with their allocation actions. If the target is empty, (14.20) 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 quotient functor carries the enhancement and orientation inherited from the physical fixed-pair problem. The equivalence holds in Artin families with these structures. 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, the root action, and equality with the prescribed full-cover characteristic polynomial. The last equations retain the exact support fiber over nonreduced bases; annihilation alone does not replace them. 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

$$\begin{aligned} \gamma_b &= (m_b + j_b) \bmod s, \qquad & r_b &= \left\lfloor \frac{m_b + j_b}{s} \right\rfloor, \\ d_s(\mathbf{m}) &= \sum_b \left(r_b - \left\lfloor \frac{j_b}{s} \right\rfloor \right), \quad & d_0(\lambda, s) &= \sum_b \left(j_b - \left\lfloor \frac{j_b}{s} \right\rfloor \right). \end{aligned} \tag{15.3}$$

Here d_s is the ordinary Euler increment. The full orbifold length, which bounds this array, is

$$\sum_b \frac{m_b}{s} = d_s(\mathbf{m}) + \sum_b \frac{\gamma_b}{s} - \sum_b \frac{j_b \bmod s}{s}.$$

Both indices below are measured relative to the zero-height baseline; I_R is the second-torus index on the H-fixed physical problem. Thus $I_{\mathrm{lex}} = I_H + I_R$. 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,

$$J = -d_0 + \sum_b \sigma_b(r_b - j_b) + 2\,\mathrm{dinv}_\lambda(\boldsymbol{\gamma}), \qquad \mathcal{L} = \mathrm{dinv}_\lambda(\boldsymbol{\gamma}) + \sum_{\sigma_b=1} (r_b - j_b).$$

For comparison with a local root reference or an ordinary reference curve, first divide by the actual baseline DT point object, with its fixed root profile. Retain the once-for-all unit calibration of that baseline and use its Thom–Sebastiani products for ordered factorizations, as in (4.23). This convention changes no Euler-characteristic or half-Tate normalization and makes no identification of different-parity graded baseline roots.

Proposition 15.2. *The normalized punctual root coefficient is ordinary and degreewise pure. Its cohomology class is*

$$V_{\lambda,s}(x, \mathbf{T}; X) = \sum_{\mathbf{m}} x^{d_s(\mathbf{m})} \mathbf{T}^{\mathcal{L}(\mathbf{m})} \prod_{b\in\lambda} X_{\gamma_b}. \tag{15.5}$$

After a noncanonical splitting of the graded cohomology, each allowed array contributes $\mathbf{T}^{\mathcal{L}(\mathbf{m})}[-2\mathcal{L}(\mathbf{m})]$. *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.

For clarity, let $d = \chi - \chi_0^{\mathcal{X}}(\lambda)$ and let a be the relative H-index on a fixed component D, so its absolute H-index is $I_0 + a$. On a punctual root packet, $a = I_H$ with the relative convention of (17.2). The ambient nil coefficient is the restricted physical coefficient tensored with Λ^{χ}, as in (10.14). Compact H-localization gives

$$R\Gamma(D, \Phi_D \otimes \Lambda^{\chi+I_0+a}) = R\Gamma(D, \Phi_D \otimes \Lambda^{d+a}), \qquad I_0 = -\chi_0^{\mathcal{X}}(\lambda).$$

The baseline equality is proved in (15.8). On a punctual root packet the preceding compact-to-ordinary duality changes its relative R-index from I_R to $-I_R$. Its exponent is consequently

$$d_s + a - I_R = d_s + J,$$

and its normalized derived 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 maps of the baseline-normalized complexes. Both ends use the same fixed local root reference baseline, so its identity restores the absolute coefficient map. 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 normalized series $G_\lambda = x^{-d_0(\lambda,s)} V_{\lambda,s}$ depends on the used character labels only through their order. Increasing embeddings into a larger root alphabet, with unused letters set to zero, therefore preserve G_λ. The raw series $V_{\lambda,s}$ retains its displayed factor $x^{d_0(\lambda,s)}$. We use the stability of G_λ 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. The index and Euler twist $\Lambda^k = \vartheta^k[-k]$ preserves complex weight zero. For this normalized direct image, (15.4) and the movable packet calculation make the auxiliary semisimple monodromy one on every cohomology stalk. Its non-unit semisimple-monodromy summand is therefore the zero object. Proper integration gives weight j in degree j; the nilpotent logarithm $N : H^j \to H^j(-1)$ has target weight $j+2$ and hence vanishes. Thus the normalized complete-component cohomology is ordinary [10, Section 2.9].

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 (14.19) 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 [35, Lemmas 2.24 and 2.26, Theorems 4.24 and 5.20]. These are precisely the conditions here.

Given a generic pair over a DVR R with prescribed positive spectral coefficients, first extend it over R in the projective fine stable-pair moduli just constructed. Write $\overline{F}$ for its R-flat sheaf. Independently, after a finite extension of the fraction field and passage to a dominating DVR R', lift the generic compact cycle to one of the finite spectral parameter covers. Finiteness of that cover extends the lift over R'. No finiteness assertion about a normalization is needed here.

Over R', the lifted monic canonical spectral polynomials, placed at their lifted scalar z-values, define a finite union Z of root spectral curves. This union is finite over $\mathcal{X}_{R'}$, hence proper over R'; its immersion in the projective compactification is therefore closed, and it lies entirely in the affine chart. Let $\mathcal{J}_Z$ be its coherent ideal in that compactification. The generic support of $\overline{F}_{R'}$ is contained in Z, so coherence gives an integer $M > 0$ with $\mathcal{J}_Z^M \overline{F}_{\mathrm{Frac}(R')} = 0$. The image of the multiplication map $\mathcal{J}_Z^M \otimes \overline{F}_{R'} \to \overline{F}_{R'}$ is therefore R'-torsion. Since $\overline{F}_{R'}$ is R'-flat, this image is zero. Thus $\mathcal{J}_Z^M \overline{F}_{R'} = 0$ on the whole family, including its nonreduced transverse structure. In particular no support, of any dimension, can appear at infinity.

Boundary avoidance descends along the faithfully flat map $R \to R'$, so the original extended pair is supported in the affine chart. Its support is finite over $\mathcal{X}_R$; purity on each fiber and the fiberwise locally free criterion make its curve pushdown locally free. Its rank and root profile are unchanged. Flatness preserves ordinary Euler characteristic, and the extended spectral coefficients agree with the prescribed ones: both are morphisms from the original DVR to the separated affine coefficient space and agree on its generic point. This proves properness over the original DVR. 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. The resulting degree-j groups remain pure, with semisimple auxiliary monodromy one; the same weight argument kills N. This proves ordinary degreewise purity 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 baseline-normalized physical coefficient factors by the supported maps of Proposition 13.3. Retain the global baseline point object once, and restore it with its fixed unit calibration after forming the relative products. The local root and ordinary reference comparisons use the same fixed calibrations as in (4.23); they do not identify the graded baseline roots of different curves. 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 14.10 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 splitting type and the 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 14.5, 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 14.8 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. Summing over the endpoint profiles, the normalized rank-one 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 (14.13) 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 14.10 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). □

Fix finite ordered alphabets X, Y and a rank bound N. Choose rational weights avoiding the slope walls at $\chi_* = 1$ for every profile on these alphabets of rank at most N, and then clear denominators to obtain a common root order S. Only these selected characters are active; set all unused root-character variables to zero. Here Theorem 14.10 is used at the auxiliary $\chi_* = 1$, without prescribed-target determinant transport. The preceding lemmas give the exponential identity below modulo R^{N+1} on these active alphabets.

The ordered formula (17.11) is unchanged by inserting unused letters or increasing the gaps between selected characters. In passing to (15.11), the two local factors $x^{d_0(\lambda,S)}$ cancel the root correction in $\chi_0^{\mathcal{F}_S}(\lambda) = \|\lambda\|^2 - 2d_0(\lambda, S)$. Thus the normalized Fock expression is independent of S after the stated specialization. It follows that the same identity holds for arbitrary finite alphabets, including the full character alphabets of any fixed football $\mathcal{F}_s$:

$$Z_{\mathrm{nil}}^{\mathcal{F}_s} = \operatorname{Exp}\left[Rp_1[X]p_1[Y]\frac{x}{(1-x)(1-x\mathbf{T})}\right]. \tag{16.2}$$

Taking all rank bounds, and using compatibility with deletion of letters and the independent symmetry in Proposition 17.4, gives the identity in the completed symmetric-function ring. No genericity assertion is made for all characters of one fixed root order.

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. The baseline-relative comparison and once-for-all unit restoration of Proposition 13.3 are used on both sides: the actual football baseline point object is retained once per global object, with ordered product calibrations inherited by Thom–Sebastiani. Passing from an ordinary reference curve to the football does not identify their graded baseline orientation lines.

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,

$$\sum_{\lambda} R^{|\lambda|}\frac{F_\lambda[X;q,t]F_\lambda[Y;q,t]}{\prod_{b\in\lambda}(q^{a_b+1}-t^{l_b})(q^{a_b}-t^{l_b+1})} = \operatorname{Exp}\left[\frac{Rp_1[X]p_1[Y]}{(q-1)(1-t)}\right]. \tag{16.4}$$

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

$$\mathsf{Z}_\lambda=\sum_{(i,j)\in\lambda}t_1^{-i}t_2^{-j},\qquad \mathsf{Z}_m=\sum_{(i,j)\in\lambda}\sum_{\alpha=1}^{m_{ij}}t_1^{-i}t_2^{-j}t_3^{\alpha},\qquad \kappa=t_1t_2t_3.$$

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^Dv^e in (17.1), define the relative indices

$$\begin{gathered}I_{\mathrm{lex}}(m)=\frac{1}{2}\sum_{D,e\,:,s|e}n_{D,e}\,\mathrm{sgn}_{\mathrm{lex}}(D,e),\qquad I_H(m)=\frac{1}{2}\sum_{D,e\,:,s|e}n_{D,e}\,\mathrm{sgn}(D),\\ J(m)=2I_H(m)-I_{\mathrm{lex}}(m).\end{gathered} \tag{17.2}$$

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

$$\ell_b = m_b + j_b = sr_b + \gamma_b, \quad 0 \le \gamma_b < s,$$
$$d_s(m) = \sum_b \left(r_b - \left\lfloor \frac{j_b}{s} \right\rfloor \right), \quad d_0(\lambda, s) = \sum_b \left(j_b - \left\lfloor \frac{j_b}{s} \right\rfloor \right). \tag{17.3}$$

The Maya set and its signs are

$$\mathcal{B}_\lambda = \{\lambda_i - i - 1 : i \ge 0\}, \qquad \sigma(c) = 2\mathbf{1}_{c \in \mathcal{B}_\lambda} - 1, \qquad \sigma_b = \sigma(c_b - 1). \tag{17.4}$$

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,*

$$J(m) = -d_0(\lambda, s) + \sum_{b \in \lambda} \sigma_b (r_b - j_b) + 2 \operatorname{dinv}_\lambda(\gamma). \tag{17.6}$$

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

$$\begin{aligned} A_s(D, e) = {} & \epsilon(e)\big(\operatorname{sgn}_{\mathrm{lex}}(D, e) - \operatorname{sgn}_{\mathrm{lex}}(D - 1, e)\big) \\ & + \epsilon(e - 1)\big(-\operatorname{sgn}_{\mathrm{lex}}(D + 1, e - 1) + \operatorname{sgn}_{\mathrm{lex}}(D, e - 1)\big). \end{aligned}$$

It is zero outside $D = -1, 0, 1$. At these three values it is, respectively,

$$\begin{aligned} A_s(-1, e) &= -\epsilon(e - 1)(1 + \operatorname{sgn}(e - 1)), \\ A_s(0, e) &= \epsilon(e)(1 + \operatorname{sgn}(e)) + \epsilon(e - 1)(\operatorname{sgn}(e - 1) - 1), \\ A_s(1, e) &= \epsilon(e)(1 - \operatorname{sgn}(e)). \end{aligned} \tag{17.7}$$

The H-kernel A_s^H is obtained by replacing every $\operatorname{sgn}_{\mathrm{lex}}(D, e)$ by $\operatorname{sgn}(D)$; 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 $\epsilon(\alpha + j) \operatorname{sgn}_{\mathrm{lex}}(-c_a, \alpha + j)$. 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

$$\#\{h \geq 1 : \lambda_{i+h} = j + h\} = \mathbf{1}_{c_b - 1 \in \mathcal{B}_\lambda}.$$

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,

$$\left\lfloor \frac{\ell_b - \ell_a - 1}{s} \right\rfloor = r_b - r_a - \mathbf{1}_{\gamma_b \leq \gamma_a}, \qquad \left\lfloor \frac{\ell_b - \ell_a}{s} \right\rfloor = r_b - r_a - \mathbf{1}_{\gamma_b < \gamma_a}.$$

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 r-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 Euler degree d_s by $|U_i|$. Under $q = \Theta/v$, $t = v\Theta$, each denominator has the form $1 - v^{a_i + b_i} \Theta^{b_i - a_i}$, with $a_i + b_i > 0$. Thus the same sums converge in $\mathbf{Q}(\Theta)((v))$, and the injective rational-field substitution identifies their rational expressions with these expansions. 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 each nonempty level diagram, retain its coordinates, and list these diagrams 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 is the row-minus-column content convention of [21, Section 3] on the transposed diagrams.

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 [21, 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 [59, Theorem 6.6 and proof of Theorem 0.2], 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 $tX-X$.

For each fixed r, the color sum in (17.15) is symmetric. Apply [21, Proposition 4.2 and Section 10, (81)–(82)] to its fundamental quasisymmetric expansion. Standardizing equal positive letters in content reading order and equal negative letters in the opposite order gives, after transposition, the column and row equality rules above. Superization sends $f[X]$ to $\omega_B f[A+B]$; the ordinary variable assignments $a_i=tx_i$, $b_i=-x_i$ then send

$$p_m\longmapsto t^m p_m[X]+(-1)^{m-1}(-1)^m p_m[X]=(t^m-1)p_m[X].$$

The cited rule permits the chosen total order of signed letters. Standardization fixes every box position, so each factor $W_\lambda(r;q,t)$ is retained.

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 $\eta_b \in \{1, \dots, s\}$ 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

$$\begin{aligned} D_0(\eta) &= \#\{u \to v \text{ in } E_\lambda : \eta_u \ge \eta_v\}, \\ B_b(\eta) &= \mathbf{1}_{\sigma_b = -1} + \mathrm{out}_{>\eta_b}(b) - \mathrm{in}_{\ge \eta_b}(b) \\ &= \mathbf{1}_{c_b \ge 0} - \mathrm{out}_{\le \eta_b}(b) + \mathrm{in}_{<\eta_b}(b), \\ M_\lambda(p;q,t) &= \prod_{\sigma_b = -1} q^{j_b - p_b} \prod_{\sigma_b = 1} t^{p_b - j_b}. \end{aligned} \tag{18.3}$$

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

$$X^\eta (-1)^{\#\{b : \varepsilon_b = 0\}} T^{D'(\eta,\varepsilon) + \#\{b : c_b \ge 0,\ \varepsilon_b = 1\}}, \tag{18.7}$$

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 $\eta_L \le \eta_R$, 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 $\eta_u \ge \eta_v$. Thus

$$D = D_0 + \sum_b \big(\mathrm{out}_{>\eta_b}(b) - \mathrm{in}_{\ge\eta_b}(b)\big)\varepsilon_b.$$

Replacing r_b by $p_b - \varepsilon_b$ in the height weight and including the letter factor t^{ε_b} gives

$$M_\lambda(p;q,t)(-1)^{\#\mathrm{negative}} T^{D+\sum_{\sigma_b=-1}\varepsilon_b}.$$

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).

For an edge $u \to v$, its inversion indicator in the old residue order minus its indicator in (18.6) is $\varepsilon_u - \varepsilon_v$. Therefore

$$D + \sum_{\sigma_b=-1}\varepsilon_b = D' + \sum_b (\mathbf{1}_{\sigma_b=-1} + \mathrm{out}(b) - \mathrm{in}(b))\varepsilon_b = D' + \sum_{c_b\ge 0}\varepsilon_b.$$

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 [20, 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 assertion holds for the three-positive paths

$$\begin{aligned} P_aP_bP_c &\longrightarrow P_bP_aP_c \longrightarrow P_bP_cP_a \longrightarrow P_cP_bP_a,\\ P_aP_bP_c &\longrightarrow P_aP_cP_b \longrightarrow P_cP_aP_b \longrightarrow P_cP_bP_a. \end{aligned} \tag{19.5}$$

Furthermore, as typed maps on P_xP_y,

$$R_{y,x}\, R_{x,y} = \mathrm{id}_{P_xP_y}\,. \tag{19.6}$$

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 [20, 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]

$$\begin{aligned} &\sum_{I_2,J_2,K_2} W_{\chi/\gamma}(I_1, J_1; I_2, J_2 \mid r, s) W_{\chi/z}(K_1, J_2; K_2, J_3 \mid r, h) W_{\gamma/z}(K_2, I_2; K_3, I_3 \mid s, h)\\ &\quad= \sum_{I_2,J_2,K_2} W_{\gamma/z}(K_1, I_1; K_2, I_2 \mid s, h) W_{\chi/z}(K_2, J_1; K_3, J_2 \mid r, h) W_{\chi/\gamma}(I_2, J_2; I_3, J_3 \mid r, s). \end{aligned} \tag{19.7}$$

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 [20, 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 (19.7). 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 (19.5), assign to every positive strand P_x the evaluation parameter x and spin $(x/\alpha)^{1/2}$, with the same α on all three strands. Apply [1, Proposition 5.1.4, equation (5.1.7)] and take $\alpha \to 0$. The third limit in [20, Lemma 3.1] is precisely the positive-positive matrix in (19.2), in the displayed lower-to-upper convention. There is no state-dependent scalar in this limit. All internal states are binary and the sums are finite, so the limit gives (19.5) term by term, initially over the rational function field in the independent row parameters.

For completeness, (19.6) also follows directly from (19.2). Put $r = x/y$. If J and N are the upper input and final upper output, a nonzero composite requires $N \subset J$. Write $W = J - N$, $d = |W|$, and write its intermediate upper state as $K = N \sqcup U$, $U \subset W$. Summing over U of each cardinality e gives the composite coefficient

$$(-1)^d \tau^{-\phi(N,W)-\binom{d}{2}} \sum_{e=0}^{d} \begin{bmatrix} d \\ e \end{bmatrix}_\tau r^{d-e} (r^{-1};\tau)_{d-e} (r;\tau)_e.$$

Here $\begin{bmatrix} d \\ e \end{bmatrix}_\tau = \sum_{|U|=e} \tau^{\phi(W-U,U)}$. The elementary finite τ-binomial convolution identifies the sum with $(1;\tau)_d$. It vanishes for $d > 0$ and equals 1 for $d = 0$. Conservation then fixes the lower output, proving the identity. The calculation allows a color to occur in both incoming rows.

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

$$u = i - j = -c_b. \tag{19.8}$$

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

$$P_s \cdots P_1 N_1 \cdots N_s. \tag{19.9}$$

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

$$D'(f) + C(f) = m_{\lambda,p}, \tag{19.10}$$

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 (19.10) 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

$$a_0(u) = \mathbf{1}_{u \in \{\alpha'_{h',j} - j : j \ge 0\}}, \qquad n_v^{<r} = \#\{a \in \rho_{h'} : i_a - j_a = v,\ f(a) < r\},$$

where the comparison is in the row order of the signed alphabet; define $n_v^{\le r}$ 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 $a_0(u) + n_{u-1}^{\le r} - n_u^{\le r}$, and its incoming horizontal bit is $n_{u-1}^{\le r} - n_{u-1}^{<r}$. The contribution of b and this lower color to the local exponent is therefore

$$C_{b,h'} = \begin{cases} a_0(u) + n_{u-1}^{\le r} - n_u^{<r}, & r \text{ positive}, \\ a_0(u) + n_{u-1}^{<r} - n_u^{\le r}, & r \text{ negative}. \end{cases}$$

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

$$D'_{b,h'} = \begin{cases} n_u^{<r} + n_{u-1} - n_{u-1}^{\le r}, & r \text{ positive}, \\ n_u^{\le r} + n_{u-1} - n_{u-1}^{<r}, & r \text{ negative}. \end{cases}$$

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 (19.10).

Proposition 19.2 (Lattice realization)**.** *Let $Z_{\lambda,p}(X;\tau)$ be the above finite path partition function with row parameters*

$$a_i(u) = \begin{cases} x_i/\tau, & u \le 0, \\ x_i, & u > 0, \end{cases} \qquad b_i(u) = -x_i. \tag{19.11}$$

When $p = 0$ occurs, impose the vacuum Maya state of that last color between the positive and negative row blocks. Then

$$K_{\lambda,p}[X;T] = T^{m_{\lambda,p}} Z_{\lambda,p}(X;T^{-1}). \tag{19.12}$$

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 (19.10). A positive box has $c_b \ge 0$ precisely when its site is $u \le 0$; the extra factor τ^{-1} in (19.11) supplies the remaining T for that box. This proves (19.12). 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* (19.9)*, 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 (19.12) 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*

$$\begin{aligned}\sum_M L_a(A,I;M,K)L'_b(M,J;C,D-K)\\ =\mathbf{1}_{A\cap I=\varnothing}\mathbf{1}_{J\subset C}\mathbf{1}_{A+I+J=C+D}\,\tau^{\phi(D,C-J)}a^{|K|}b^{|D-K|}\tau^{\phi(K,D)}.\end{aligned} \tag{20.2}$$

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 A,B,C,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 $\mathcal{J}(A,B;C,D)v_D(a,b)$. Put $F=B\setminus C$, $E=B\setminus D$, and $N=B\cap A$. Then*

$$\mathcal{J}(A,B;C,D)=a_0^{|F|}b_0^{|N|}\tau^{\phi(D,C\setminus B)+\phi(F,B)}\prod_{i\in E}(b_0+a_0\tau^{h_i}), \tag{20.3}$$

$$h_i=\#\{j\in B: j>i\}+\#\{j\in D: j<i\}.$$

If conservation fails, the coefficient is zero. Set differences in this formula are regarded as their binary vectors in ϕ.

Proof. Apply (20.2) with $J=B-I$ and sum the coefficient of $|I,B-I\rangle$ in $v_B(a_0,b_0)$. Binary conservation partitions B into the disjoint sets F, E, and N: their colors are respectively forced into I, free, and excluded from I. Thus the allowed splits are $I=F\sqcup S$ for $S\subset E$. After removing the common output factor, the exponent is

$$\phi(D,C-B)+\phi(D,I)+\phi(I,B)=\phi(D,C-B)+\sum_{i\in I}h_i.$$

Here $C-B$ is a vector difference. Since $C-B=(C\setminus B)-F$ and $\sum_{i\in F}h_i=\phi(F,B)+\phi(D,F)$, the forced colors give the exponent displayed in (20.3). Their scalar factor is $a_0^{|F|}$, the excluded colors give $b_0^{|N|}$, and summing each free choice gives $b_0+a_0\tau^{h_i}$. This proves the polynomial identity without division by b_0. □

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 $|B| \le \ell$.

More precisely, take nonnegative integers $\ell_1 \le \ell_0$ *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 $|D| \le \ell_1$, 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 $i \in B \setminus D$, and its new height is $h_i = \ell_0$. At $b_0 = -a_0\tau^{\ell_0}$ the corresponding factor $b_0 + a_0\tau^{h_i}$ 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 an adjacent unprojected 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. If the right word is empty, the empty horizontal covector kills it directly. Otherwise its first column realizes 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 and is killed by the empty horizontal boundary. This argument fixes each pair column's exterior vertical states and sums its internal state between the two rows. Fixing any surrounding-row history therefore gives the same zero; summing those histories preserves it. □

20.2. **Moving an external row past a fused pair.** The positive train and negative train are the two-crossing maps

$$\begin{aligned} \mathcal{B}^{+}_{a,b;c} &: \quad P_aN_bP_c \xrightarrow{R'_{b,c}} P_aP_cN_b \xrightarrow{R_{a,c}} P_cP_aN_b, \\ \mathcal{B}^{-}_{c;a,b} &: \quad N_cP_aN_b \xrightarrow{R'_{c,a}} P_aN_cN_b \xrightarrow{R''_{c,b}} P_aN_bN_c. \end{aligned} \tag{20.4}$$

Their possible mixed-crossing denominators are $(-c/b; \tau)_d$ and $(-a/c; \tau)_d$, respectively. For an independent outside parameter c, neither is forced to vanish at either inner-pair specialization below.

Lemma 20.5 (Image preservation at resonance)**.** *The trains restrict to maps*

$$\mathcal{K}(a,b)P_c \longrightarrow P_c\mathcal{K}(a,b), \qquad N_c\mathcal{K}(a,b) \longrightarrow \mathcal{K}(a,b)N_c.$$

At fixed row parameters they intertwine columns, also at $b=-a\tau^{\ell}$ for $\ell\in\{0,1\}$ and $a\neq 0$, with a generic independent outside parameter c. Transport across the parameter cut below uses the same-sign permutation Ξ, whose ratios are unchanged.

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 either specialization. More precisely, the map

$$p|I,J\rangle = \mathbf{1}_{I=\varnothing}b^{-|J|}e_J$$

is a regular left inverse to $e_B \mapsto v_B(a,b)$. The fused module is therefore a split subbundle of constant rank. The external trains are regular there over the rational function field in c. For each fixed outside output, extract the scalar of a binary inner sector D from its all-negative coordinate by dividing by $b^{|D|}$. Proportionality of all other coordinates is a regular rational identity and thus specializes; nonbinary inner sectors vanish by the same identity principle. Applying the left inverse to the raw column-intertwining identity gives the restricted identity at both specializations. Only unnormalized v_B and regular external trains are used; the normalized own-pair crossing may have a pole there. □

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 $d = 0$, $B_\varnothing(z) = (1)$ and $C = (1 - z/\tau)$. To verify the formula, choose output inner positive state empty in the two terms of the input singleton in (20.5), and divide by $b^{|D|}$: this is 1 for vacuum output and $-\tau$ for singleton output. 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. Each column's exterior vertical input and output may be fixed; its vertical states between rows are summed. 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. Write $\iota_i(e_B) = v_B(a_i, b_i)$ for the embedding (20.1); thus $E_1 = \iota_1$. For $k > 1$, expand the last factor by ι_k 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,

$$\begin{aligned} E_k = (\mathrm{id}_{P_k} \otimes E_{k-1} \otimes \mathrm{id}_{N_k}) \\ \circ (\mathcal{B}^+_{1,k} \cdots \mathcal{B}^+_{k-1,k} \otimes \mathrm{id}_{N_k}) \circ (\mathrm{id} \otimes \iota_k). \end{aligned} \tag{21.2}$$

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 the typed identity $M_{\mathrm{nested}} E_k = E_k M_{\mathrm{adjacent}}$ by induction, where M_{adjacent} acts on the tensor product of fused pairs. 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 moving the new positive strand through the crossing word defining E_{k-1}. Moving it through an old mixed crossing uses the first identity in (19.4); moving it through an old positive-positive crossing uses (19.5). Crossings on disjoint strands commute. This proves equality of the two typed maps one crossing of E_{k-1} at a time. Choose the positive component of Ξ to use the inverse crossings in (19.6), so that the stated positive-positive cancellation is exact, with scalar 1.

The crossings of Ξ that move P_k back through the positive block cancel the last positive-positive crossings of this expression. Write M_j for the crossing of N_j, P_k and D_j for that of N_j, N_k. The remaining chronological word is

$$M_{k-1}, \ldots, M_1, D_{k-1}, \ldots, D_1.$$

Move D_j left past $M_{j-1}, \dots, M_1$; these pairs of crossings have disjoint strands. The resulting word consists of the adjacent pairs $M_{k-1}, D_{k-1}, \dots, M_1, D_1$, each acting as

$$N_j P_k N_k \longrightarrow P_k N_j N_k \longrightarrow P_k N_k N_j.$$

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. For the correspondingly permuted observer, $O = O^{\Xi}\Xi$, so 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 pullback of each right observer along 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, so this quotient has no degree greater than k. Since E_k preserves degree, every homogeneous vector in its image has a homogeneous preimage of the same degree. Its pairing therefore vanishes above degree k. Lemma 21.1 places every reachable state in that image. The image assertion uses the generic rational function field. Coefficients of a fixed finite raw network are polynomial in the row parameters over $\mathbf{Q}[\tau^{\pm 1}]$. Their vanishing therefore extends to every matched specialization with τ invertible: $b_i = -x_i$, with positive parameters x_i/τ on the left and x_i on the right. □

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)**.** *With the physical last-color boundary just specified, the adjacent pair has outgoing fused lines $\widehat{v}_D = \Pi v_D$. At left capacity one their degree-at-least-two part is invariant and invisible to the right observer. Write ϵ for the sum of basis coefficients; the empty horizontal covector selects only vacuum. The 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

$$\text{positive output count} = \text{positive input count} + 1 - m \;\geq\; \text{positive input count}.$$

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. On a surviving left history each outer positive row has zero incoming and outgoing horizontal g-bit, so it passes the physical vertical bit 1 to the inner positive block. On the right it similarly passes the bit 0. The inductive projected inner vanishing therefore applies with its required g-boundary, for every value of the other vertical colors.

For the new high factor use the simultaneous same-sign permutation Ξ before Π. It preserves total positive g-occupation, so

$$\Xi\Pi = \Pi\Xi, \tag{21.7}$$

where the projector on each side acts on the corresponding permuted row word. 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 surviving outer positive rows pass the same physical vertical bit to the innermost pair. Its 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. Here columns are numbered from the first right site, and the states are outputs of the preceding right columns. 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 is

$\{1,\dots,s\}$. The same bound at intermediate times, followed by the s top-one exits, forces the complete history

$$S_t = \{i : 1 \le i \le k,\ i \le t \le i+s-1\} \qquad (0 \le t \le s+k),$$

where t counts all encoded and observer columns. This also satisfies the two interlacing rules, proving existence and uniqueness.

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 $|K_{\text{old}}|$, 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 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 one-based 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 $\delta C_g = \delta J_g$. 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, $\delta C_g = \delta J_g$ 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

$$\begin{aligned} \deg_{\mathrm{selected}} &= \sum_u h(u) \\ &\le \sum_u \min\{d_\lambda(u), r_k(u)\} = \lambda_0 + \cdots + \lambda_{k-1}. \end{aligned}$$

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. For completeness, put $z=q^{-1}$, $T=qt$ and $J_\lambda=\sum_{(i,j)\in\lambda} j$. The paired summand has z-degree $|p|-J_\lambda$, and its remaining coefficient is in $\mathbf{Q}(T)$; hence only finitely many heights occur in each degree. The embedding $\mathbf{Q}(q,t)\simeq\mathbf{Q}(z,T)\hookrightarrow\mathbf{Q}(T)((z))$ is injective. Rationality from Lemma 17.3 therefore carries the coefficientwise vanishing back to $\mathbf{Q}(q,t)$. 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 $K=\mathbf{Q}(q,t)$. Use the full two-alphabet identity (16.4) for the functions of Definition 17.2, and write

$$c_\lambda^{-1}=\prod_{b\in\lambda}(q^{a_b+1}-t^{l_b})(q^{a_b}-t^{l_b+1}). \tag{23.3}$$

The right side of (16.4) is the reproducing kernel for the nondegenerate pairing (16.6).

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 (16.6). If $C=\operatorname{diag}(c_\lambda)$, comparison of coefficients in (16.4) gives

$$ACA^{\mathsf{t}}=B^{-1}. \tag{23.4}$$

Both B and C are invertible. Taking determinants proves A invertible, so the F_λ form a basis. Inversion of (23.4) 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

$$\begin{aligned}\langle p_\rho,p_\sigma\rangle_U&=\delta_{\rho\sigma}z_\rho\prod_{r\in\rho}\frac{-(1-q^r)}{1-t^r}\\&=t^{-n}\delta_{\rho\sigma}z_\rho\prod_{r\in\rho}\frac{1-q^r}{1-t^{-r}}.\end{aligned} \tag{23.5}$$

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 (2.4). Its existence, dominance triangularity and full orthogonality, including incomparable partitions, are recalled and proved in [71, §2]. 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

$$U\widetilde{H}_\lambda[X;q,t] = t^{n+n(\lambda)}J_\lambda[X;q,t^{-1}]. \tag{23.6}$$

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 $\mu^p = (\mu^p_1, \dots, \mu^p_{\ell_p})$ 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

$$N_\mu = (g-1)n^2 + 1 + \frac{1}{2}\sum_{p\in D}\left(n^2 - \sum_a (\mu^p_a)^2\right). \tag{24.1}$$

When the moduli space is nonempty, it is smooth of dimension $2N_\mu$. Its Hitchin morphism is projective by Lemma 24.1. 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.

Lemma 24.1 (Projectivity for rational parabolic weights)**.** *Fix a smooth projective curve, a reduced marking divisor, a line bundle L, a rank, a degree, and rational parabolic weights and multiplicities. The characteristic-coefficient map on the GIT moduli space of S-equivalence classes of semistable weakly parabolic L-Higgs bundles is projective. For numerically generic weights this is the stable coarse moduli space. The strongly parabolic locus has a projective Hitchin map as well. The same holds when the induced map on each graded block is prescribed to equal $\lambda_{p,a}\,\mathrm{id}$, with fixed $\lambda_{p,a} \in L|_p$. For $L = K_C(D)$ these are scalar graded residues via the residue identification $L|_p \simeq \mathbb{C}$. No genus or coprimality hypothesis is required.*

Proof. Choose denominators s_p for the weights, and write $\nu : \mathfrak{C} = C[\sqrt[s_p]{p} : p \in D] \to C$, with root divisors q_p. A parabolic bundle is equivalently a bundle E on $\mathfrak{C}$: its filtration is $\nu_*(E(-lq_p))$, and a character-c summand has weight c/s_p. Weakly parabolic fields correspond to maps $E \to E \otimes \nu^* L$. These assertions follow by writing the invariant modules in each root chart and their filtration-preserving transition matrices.

Put

$$a = \prod_p s_p, \qquad \mathcal{G} = \bigoplus_{0\leq b_p < s_p} \mathcal{O}_{\mathfrak{C}}\left(\sum_p b_p q_p\right), \qquad \kappa = 1 - g - \sum_p \frac{s_p - 1}{s_p}.$$

Every inertia character occurs in $\mathcal{G}$, so it is a generating bundle. For a rank-r bundle E the local formula $\nu_*\mathcal{O}((c-b)q_p) = \mathcal{O}_C(\lfloor (c-b)/s_p \rfloor p)$ gives

$$\chi(E \otimes \mathcal{G}^\vee) = a\big(\mathrm{pardeg}\, E + r\kappa\big). \tag{24.2}$$

Indeed, among $b = 0, \dots, s_p - 1$ exactly $s_p - 1 - c$ floors are negative. This also covers zero weights and unused characters.

Use the fixed smooth projective root ruled surface $\overline{\mathfrak{S}} = \mathbf{P}_{\mathfrak{C}}(\mathcal{O} \oplus \nu^* L)$, and pull $\mathcal{G}$ back to it. Choose a very ample coarse polarization $H = C_\infty + \pi_C^* B$, with $b = \deg B > 0$ sufficiently large.

On $\mathrm{Tot}(\nu^*L)$ its restriction is π_C^*B. A pure one-dimensional spectral sheaf F finite over $\mathfrak{C}$ has locally free pushforward E; its modified Hilbert polynomial is

$$P_{\mathcal{G}}(F,m) = a\big(rbm + \mathrm{pardeg}\, E + r\kappa\big).$$

Spectral subsheaves correspond to invariant subsheaves of E. After division by the leading coefficient their slopes differ from parabolic slopes by the same constant and positive factor. Saturation preserves invariance and increases slope; a proper full-rank subsheaf has a positive-length quotient and strictly smaller polynomial. Thus the two notions of (semi)stability agree.

Apply the projective pure-sheaf GIT construction of [34, Proposition 6.17 and Theorems 6.20–6.22] on this fixed surface. Over the affine coefficient space $A = \bigoplus_{i=1}^{n} H^0(C, L^i)$, require the universal monic spectral equation to annihilate the universal quotient. This is a closed invariant condition on the projective Quot construction. Its support misses infinity and is finite over $\mathfrak{C}$. The resulting relative quotient is therefore projective over A, and its semistable objects are precisely pure spectral sheaves. Their pushforwards in a flat family are vector bundles, so their inertia multiplicities are locally constant. The prescribed multiplicities select an open-and-closed part. This part is GIT-saturated, since multiplicities are additive in Jordan–Hölder filtrations; the polynomial then fixes the prescribed degree.

Annihilation alone does not prescribe the characteristic polynomial. Its actual coefficients define an invariant algebraic map to A, which descends to the quotient. Impose equality with the parameter in A; the diagonal of A is closed. The resulting projective family is exactly the weakly parabolic Higgs moduli and its characteristic-coefficient map.

Numerical genericity makes every semistable object stable, for the chosen degree. Vanishing, or equality to the prescribed scalar map, on each graded block is an invariant closed condition on the Quot atlas. Exactness of invariants for the reductive GIT group identifies the restricted quotient with a closed subscheme of the original quotient. This gives the strong and prescribed-scalar assertions. Full-flag ordered residues give a map to a separated ordered coefficient base over A; its graph is closed in the base change of the projective family. Thus the ordered map, and its residue and line slices, are projective as well. □

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

$$P_j H^i(Y) = \mathrm{im}\left[H^{i-d_Y}\left(A, {}^{p}\tau_{\le j} Rh_* \mathbf{Q}_Y[d_Y]\right) \longrightarrow H^i(Y)\right]. \tag{24.3}$$

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

$$u = \frac{1}{n} c_1(\mathcal{E}_x), \qquad \mathrm{ch}^{\mathrm{nor}}(\mathcal{E}) = e^{-u}\,\mathrm{ch}(\mathcal{E}), \qquad \mathrm{ch}^{\mathrm{nor}}(\mathcal{V}_{p,a}) = e^{-u}\,\mathrm{ch}(\mathcal{V}_{p,a}). \tag{24.4}$$

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

$$\int_C \mathrm{ch}_k^{\mathrm{nor}}(\mathcal{E})\,\gamma, \quad \gamma \in H^*(C), \qquad c_k^{\mathrm{nor}}(\mathcal{V}_{p,a})$$

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 24.2. *For every $g \geq 0$ and every collection of partial flags, the above normalized classes generate the cohomology, and*

$$P_j H^*(M_{\mathrm{Dol},\mu}) = C_{j+N_\mu} H^*(M_{\mathrm{Dol},\mu}). \tag{24.5}$$

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 [23, §§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 (24.5) 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 $|D| \geq 3$ if $g = 0$, and $|D| \geq 1$ otherwise. Then

$$L = K_C(D), \qquad \deg L > 0. \tag{24.6}$$

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

$$\frac{e}{r} - \frac{d_0}{n} + \sum_{p,a} \alpha_{p,a} \left(\frac{\nu_{p,a}}{r} - \frac{\mu^p_a}{n} \right), \qquad e \in \mathbf{Z}.$$

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

$$N_F = (g-1)n^2 + 1 + |D| \binom{n}{2}, \tag{24.7}$$

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

$$\Sigma = n\pi^* C_0 - \sum_{p,i} E_{p,i}, \qquad \Sigma \cdot E_{p,i} = 1. \tag{24.8}$$

Here the zero-section class is pulled back as a divisor class; the formula therefore also covers a zero residue. The blowup spectral correspondence [33, §2.2], applied independently in the disjoint charts at the marked points, 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. In a chart at p, the blowup equation is $y - \xi_{p,i} = xt$; the distinct residue differences are units. The local sheaf constructions and their inverses therefore commute with parameter base change and agree on the unchanged complements of these charts.

Lemma 24.3. *The space X_ξ has a projective compactification carrying a universal spectral sheaf. The pure cohomology of X_ξ 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

$$(A \cdot \Sigma)(\chi' + \beta \cdot \Sigma') = (A \cdot \Sigma')(\chi + \beta \cdot \Sigma) \tag{24.9}$$

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 pure one-dimensional twisted-semistable sheaves [65, Theorem 7.1], with twisting class $-\beta$ in the convention of [65, Definition 5.1], therefore 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

$$\mathcal{L}_E = \det Rp_*(\mathcal{F} \otimes^L q^* \mathcal{O}_E)$$

is defined by a perfect complex: use the two-term resolution of $\mathcal{O}_E$, the base-flat universal family, and proper pushforward from the smooth surface. It 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.

If $2N_F = 0$, the graph and generation assertions can be checked without the positive-dimensional diagonal formula. Suppose first that X_ξ is nonempty and choose $F \in X_\xi$. Smoothness of this zero-dimensional symplectic moduli space gives $\mathrm{Ext}^1(F,F) = 0$, whereas stability and Serre duality give $\mathrm{Hom}(F,F) = \mathrm{Ext}^2(F,F) = \mathbb{C}$. Thus $\chi(F,F) = 2$. For any other stable sheaf G in the compactification, stability and the fact that $F \otimes K_{Z_\xi} \simeq F$ would give $\mathrm{Hom}(F,G) = \mathrm{Ext}^2(F,G) = 0$, hence $\chi(F,G) = -\dim \mathrm{Ext}^1(F,G) \le 0$. This contradicts Riemann–Roch, since F and G have the same numerical class. The compactification therefore consists of the single point F; it is reduced because F lies in its smooth interior. Its graph class is $c_0 = 1$, and the generation assertion is immediate. The empty case is vacuous. We may henceforth assume $2N_F \ge 2$.

We apply [38, Lemma 4 and §4] directly on the smooth product $\overline{X}_\xi \times X_\xi$, using the pulled-back universal families. Put $m = 2N_F \ge 2$. The relative RHom complex in (24.10) is perfect of amplitude $[0,2]$; represent its shift by one as $V_{-1} \xrightarrow{g} V_0 \xrightarrow{f} V_1$. All sheaves have the same twisted reduced Hilbert polynomial, so Hom in both directions vanishes off the graph. Serre duality gives the same vanishing for Ext^2, using the canonical-bundle trivialization of the interior family.

The two cokernel hypotheses in Markman's lemma hold scheme-theoretically. Near the graph, nonzero Hom between the universal stable families is the Isom functor, a $\mathbf{G}_m$-torsor over their equality locus. Because the resolution is an isomorphism over X_ξ, this locus is the actual graph Γ. After eliminating an invertible block of the first differential, its remaining column has entries $s_1, \dots, s_a$. The nonzero Hom functor has equations $s_i z = 0$ with z invertible; hence the s_i generate the ideal of Γ. Thus $\mathrm{coker}(g^\vee)$ is the pushforward of a line bundle on Γ. Swapping the two families and using relative Serre duality gives the same assertion for $\mathrm{coker}(f)$. Moreover $\ker(g) = 0$, since it is a subsheaf of a vector bundle supported in positive codimension. The virtual rank is $m - 2$ by the Riemann–Roch calculation below. Markman's lemma, with m even, gives the graph class with coefficient one. Consequently, if p_{13} integrates over Z_ξ,

$$[\Gamma] = c_{2N_F}(-Rp_{13*}\,\mathrm{RHom}(\mathcal{F}_{12}, \mathcal{F}_{32})) \quad \text{on } \overline{X}_\xi \times X_\xi. \tag{24.10}$$

Here Γ is the graph of the open inclusion $j : X_\xi \hookrightarrow \overline{X}_\xi$. Surface Riemann–Roch gives $\Sigma^2 = 2N_F - 2$ and virtual rank $2N_F - 2$ for the negative relative RHom complex; the class in (24.10) is its Chern class of index $2N_F$.

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 (24.10) expresses every b_s in the universal Chern components. For the ruled surface, the summand complementary to $\pi^*H^*(C)$ carries the infinity-section class. The universal spectral sheaf is disjoint from that section, so multiplication of its Chern character by this class is zero. The blowup decomposition of $H^*(Z_\xi)$ therefore leaves only the curve and exceptional-divisor components. 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 (24.4) 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. For $i \geq 2$,

$$H^1(C, L^i(-D)) = 0. \tag{24.11}$$

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

$$K_{\text{weak}} = [\text{ParEnd}(E) \xrightarrow{[\theta,\cdot]} \text{ParEnd}(E) \otimes K_C(D)],$$

$$K_{\text{fix}} = [\text{ParEnd}(E) \xrightarrow{[\theta,\cdot]} \text{SParEnd}(E) \otimes K_C(D)].$$

Parabolic Serre duality identifies $\mathbb{H}^2(K_{\text{weak}})^\vee$ 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 $\mathbb{H}^2(K_{\text{weak}}) = 0$ and X is smooth. The same duality gives $\mathbb{H}^2(K_{\text{fix}}) = \mathbf{C}$, dual to the scalar identity. The residue map gives an exact sequence of complexes

$$0 \longrightarrow K_{\text{fix}} \longrightarrow K_{\text{weak}} \longrightarrow \bigoplus_{p,i} \mathbf{C}[-1] \longrightarrow 0.$$

Its connecting map to $\mathbb{H}^2(K_{\text{fix}})$ 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

$$\mathbb{H}^1(K_{\text{weak}}) \longrightarrow \mathbf{C}^{n|D|} \xrightarrow{\sum} \mathbf{C} \longrightarrow 0.$$

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 [24, 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

$$H^*(X) \xrightarrow{\sim} H^*(X_\xi), \qquad H^*(X) \xrightarrow{\sim} H^*(X_0). \tag{24.12}$$

These groups are pure, and the restrictions carry the universal curve and flag classes to the corresponding classes. Lemma 24.3 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

$$\sum_{p,i} \nu_{p,i}\xi_{p,i} = 0, \qquad \sum_{i} \nu_{p,i} = r, \quad 0 < r < n, \quad \nu_{p,i} \in \{0,1\}. \tag{24.13}$$

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

$$H^*(X) = H^*_{\mathrm{pure}}(U) = H^*_{\mathrm{taut}}(U). \tag{24.14}$$

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

$$\mathcal{D}_{p,i} = (\Gamma_{p,i})_*\mathcal{O}, \qquad \mathcal{L}_{p,i} = \Gamma_{p,i}^*\mathcal{F}. \tag{24.15}$$

The degree-one modification is the exact sequence

$$0 \longrightarrow \mathrm{Mod}^*_{p,i}\,\mathcal{F}_{d-1} \longrightarrow \mathcal{F}_d \longrightarrow \mathcal{L}_{p,i} \otimes \mathcal{D}_{p,i} \longrightarrow 0. \tag{24.16}$$

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

$$\psi_k(\gamma) = k! \int_C (\mathrm{ch}_k(\mathcal{E}) + (1-g)\omega\,\mathrm{ch}_{k-1}(\mathcal{E}))\gamma, \qquad \gamma \in H^*(C), \tag{24.17}$$

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 (24.17) give the Chern filtration of (24.4). The degree-one exact sequence on the correspondence gives the basic relation

$$[\psi_k(\gamma), T_l(\eta)] = kT_{k+l-1}(\gamma\eta). \tag{24.18}$$

Brackets here and below are graded brackets when both operators are odd. The virtual projective-bundle formula evaluates $T_k(\gamma)1$; together with (24.18) this determines the action on the tautological ring. The surface Hecke relations of [23, §§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

$$\begin{aligned} &[X_{p,i}, y_{p',i'}] = 0, && [\psi_k(1), X_{p,i}] = k y_{p,i}^{k-1} X_{p,i}, \\ &[\psi_k(\gamma), X_{p,i}] = 0 \quad (\deg\gamma > 0), && [y_{p,i}, T_0(1)] = X_{p,i}. \end{aligned} \tag{24.19}$$

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.

The coefficient ring is the one in [23, Definition 7.8]. We verify the algebra input explicitly.

Lemma 24.4. *The displayed tautological and branch operators satisfy the open-surface Hecke algebra relations with coefficient ring $H_{n,D}$. Its surface parameters satisfy $c_1^2 - c_2 = 0$ and $c_1\Delta = 0$, so the exact Lie relations of* [23, Theorem 5.6] *apply to this enlarged action.*

Proof. Let $\overline{S}$ be a smooth ruled compactification of $S = \mathrm{Tot}_C L$, put $\overline{H} = H^*(\overline{S}, \mathbf{Q})$, and let $J = \ker(\overline{H} \to H^*(C, \mathbf{Q}))$. Write δ for the point class of $\overline{S}$, with integral one. Let V be the degree-two vector space generated by $v_{p,i}$ subject to $\sum_i v_{p,i} = 0$ for every p. Give $V^\vee$ degree two as well. On the graded vector space

$$\widehat{H} = \overline{H} \oplus V \oplus V^\vee$$

retain the product on $\overline{H}$, let its unit act as the identity, and put

$$\overline{H}^{>0}(V \oplus V^\vee) = 0, \qquad VV = V^\vee V^\vee = 0, \qquad v\alpha = \alpha(v)\delta.$$

These rules are associative. Integration on $\overline{H}$, extended by zero, gives a nondegenerate graded Frobenius pairing on $\widehat{H}$. For dual bases $v_a, v_a^\vee$, its diagonal and formal surface parameters are

$$\widehat{\Delta} = \Delta_{\overline{S}} + \sum_a (v_a \otimes v_a^\vee + v_a^\vee \otimes v_a), \qquad \widehat{c}_1 = c_1(\overline{S}), \qquad \widehat{c}_2 = c_2(\overline{S}) + 2\dim(V)\delta.$$

The diagonal is balanced for every element of $\widehat{H}$, its contraction with integration is the unit, and $\widehat{\Delta}^2 = \widehat{c}_2\widehat{\Delta}$. Indeed, mixed old/new terms in the square vanish, and the new terms contribute $2\dim(V)\delta \otimes \delta$. Thus these data satisfy the hypotheses of [23, §5.3] before any quotient is taken.

Since J has positive degree and contains δ, the subspace $\widehat{I} = J \oplus V^\vee$ is a graded ideal. The quotient is exactly $H_{n,D}$ under

$$\tau_{p,i} = \omega/n + v_{p,i}.$$

The new summands of $\widehat{\Delta}$ vanish in the quotient; the images of $\widehat{\Delta}, \widehat{c}_1, \widehat{c}_2$ are therefore the usual open-surface data $\Delta_S, c_1(S), 0$. The class Δ_S is a scalar multiple of $\omega \otimes \omega$. In particular every $\tau_{p,i}$ annihilates it, $c_1(S)\Delta_S = 0$, and $c_1(S)^2 = 0$. The graded-ideal quotient construction of [23, §5.3] supplies the corresponding formal open-surface algebra and Fock action. For the rank-zero surface sheaves used here, the formal operator in [23, §4.6 and Corollary 5.5] is

$$T_l^F(\beta)f = \mathcal{Q}(\beta u^{l+1} R(f)), \qquad l \geq 0.$$

The tautological ideal generated by $p_k(\widehat{I})$, $k \geq 1$, is stable: the R-formula preserves it up to coefficients in $\widehat{I}$, and $h_j(\widehat{I})$ lies in this ideal for $j \geq 1$. Only positive powers of u occur in the displayed formula. Thus h_0 is never used, and no augmentation is required on the quotient.

It remains to identify this action with the actual operators. For the original curve classes this is the surface projective-bundle calculation above. For a flag symbol, (24.19) and the identity

$$T_l(\gamma) = \frac{1}{l+1}[\psi_{l+1}(\gamma), T_0(1)] \qquad (\gamma \in H^*(C))$$

give, for $k, l \geq 0$, $[\psi_k(\alpha), T_l(\beta)] = kT_{k+l-1}(\alpha\beta)$ for all $\alpha, \beta \in H_{n,D}$, including pairs of flag symbols. In the latter case both sides vanish because every $X_{p,i}$ commutes with every branch Chern root. When $k = 0$ the commutator vanishes; a negative-index term with zero coefficient is omitted. Also $T_l(\tau_{p,i})1 = y_{p,i}^l$. In the formal symmetric-function calculation, the vanishing $\Delta_S\tau_{p,i} = 0$ makes $(fg)(\tau_{p,i}) = 0$ whenever f, g have positive degree. Since $c_1(S)\tau_{p,i} = 0$, this gives

$$h_{l+1}(\tau_{p,i}) = \frac{p_{l+1}(\tau_{p,i})}{l+1} = \psi_l(\tau_{p,i}) = y_{p,i}^l.$$

These are exactly the commutator and vacuum formulas [23, (4.9)–(4.10), §4.6 and Corollary 5.5], with the rank-zero surface-sheaf index shift. Commuting an operator past the tautological generators reduces its action to its value on the unit. Hence evaluation of the formal Fock space intertwines the actual operators on the entire enlarged tautological ring. Its Hecke relations, and then the exact Lie relations under the displayed vanishing conditions, follow. The degree and gerbe hypotheses for the subsequent reduction are checked below. □

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}. For every d, this space is the restriction of the semistable projective Hitchin construction of Lemma 24.1 to the ordered integral spectral locus. Every such spectral sheaf is stable, so U_d is smooth and projective over that base even when the chosen weights have walls in degree d on the full Hitchin base. 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 24.5. *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 (24.16) 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 $a_{k,\gamma}^{(d)} = \int_C \mathrm{ch}_k^{\mathrm{nor}}(\mathcal{E}_d)\gamma$,

$$X_{p,i}^{(d)} a_{k,\gamma}^{(d-1)} = a_{k,\gamma}^{(d)} - \frac{(\ell_{p,i}^{(d)})^{k-1}}{(k-1)!} \int_C \omega\gamma. \tag{24.20}$$

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 (24.14) 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 (24.20), 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 [23, §§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 24.5. Weyl reduction with

$$y = \psi_1(\omega)/n = c_1(\mathcal{E}_x)/n \tag{24.21}$$

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 [23, 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 (24.15) 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 24.5. The normalized maps Q_d/n fix the unit. This is the quotient identification of reduced cohomology in [23, 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

$$\begin{aligned} [D_{2,0}(1), D_{a,b}(\gamma)] &= -2bD_{a+1,b-1}(\gamma), \\ [D_{1,1}(1), D_{a,b}(\gamma)] &= (a-b)D_{a,b}(\gamma), \\ [D_{0,2}(1), D_{a,b}(\gamma)] &= 2aD_{a-1,b+1}(\gamma). \end{aligned} \tag{24.22}$$

The extended relations (24.19) add the flag generators to this reduction. These are statements about the raw operator construction and its algebraic reduction.

Here the algebraic scope can be checked without using any assertion from [23, Section 8]: Section 7 of that paper explicitly starts with $g \geq 0$. For $S = \operatorname{Tot} L$, the identities $H^*(S) = H^*(C)$ give $c_1(S)^2 - c_2(S) = 0$ and $c_1(S)\Delta_S = 0$, as required by its Theorem 5.6. The five geometric assumptions in Sections 4 and 6 are also satisfied. Purity of the spectral sheaves gives (A1), and integral spectral support gives (A2), since colength-one subsheaves have the same integral support and remain stable. The weight-one branch line splits the gerbe and is tautological, giving (A3) and (A5). For (A4), tensoring by $\pi_C^*\mathcal{O}_C(x)$ preserves the spectral support and the ordered branches and identifies degree d with degree $d+n$; here $\eta = \pi_C^*[x]$ and $\psi_0(\eta) = n$. Alternatively, the successive-degree invertibility required for polynomiality is already supplied by Lemma 24.5. The tautological cohomology of the coarse space is finite dimensional and its positive-degree ideal is nilpotent. Hence

$$e = \tfrac{1}{2}D_{0,2}(1), \qquad h = -D_{1,1}(1), \qquad f = -\tfrac{1}{2}D_{2,0}(1) \tag{24.23}$$

is an $\mathfrak{sl}_2$-triple, and

$$[h, \psi_k(\gamma)] = k\psi_k(\gamma), \qquad [h, y_{p,i}] = y_{p,i} \tag{24.24}$$

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 24.5 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 [23, 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

$$a^*P_j \subset P_{j+\dim Z-\dim Y}, \qquad a_*P_j \subset P_{j+\dim Z-\dim Y}. \tag{24.25}$$

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. We spell out why neither a singular Hecke space nor the two algebraic reductions change the filtration bound. The virtual class of the length-one correspondence has dimension one more than its source and target. As an algebraic cycle it is a rational sum of cycles of that dimension. Resolve each cycle and apply (24.25) to its two projections, which are projective over the common Hitchin base. Each projection contributes -1 to the centered bound. The insertion $z^b\gamma$ is pure of cohomological degree $2b+i$, so its cup product contributes $2b+i$. This is the virtual-cycle interpretation of [23, Proposition 3.9], and gives

$$T_b(\gamma)P_j \subset P_{j+2b-2+i}.$$

On the split scalar gerbes, extend the coarse filtration to $R_d[v_d]$ by giving the weight-one coordinate v_d perverse degree two. On a length-one correspondence write $p_-^*v_- = v_+ + t$ and $z = v_+ + z_0$, where t, z_0 are pure degree-two coarse classes. An input av_-^l of perverse degree at most j has $a \in P_{j-2l}$. In the expansion of $(v_+ + t)^l(v_+ + z_0)^b$, a term with total coarse exponent k has gerbe exponent $b + l - k$. The preceding cycle calculation bounds its output by

$$j - 2l + (2k + i - 2) + 2(b + l - k) = j + 2b - 2 + i.$$

This proves the same bound on the gerbe. Changing from v_d to y preserves the filtration, since $y - v_d$ is a pure degree-two coarse class. Multiplication by $\psi_b(\gamma)$ has the same bound by [23,

Proposition 3.7]. The unreduced formula

$$D^{\mathrm{raw}}_{m,b}(\gamma) = \frac{b!}{(m+b)!}(-\operatorname{ad}_{T_0(1)})^m \psi_{m+b}(\gamma)$$

therefore has bound $2b-2+i$, independently of m.

Put $\mathsf{Q} = T_0(\omega)/n$. Lemma 24.5 shows that Q and its inverse have perverse degree zero. All coefficients of the polynomial $\mathsf{Q}^{-m} D^{\mathrm{raw}}_{m,b}(\gamma)$ consequently have the same bound, by polynomial interpolation. In the notation of [23, Section 6.5], the extra normalization is $\mathsf{U} = \mathsf{Q}\exp(\theta/n)$, where θ is the nilpotent linear coefficient extracted from $\mathsf{Q}^{-m} D^{\mathrm{raw}}_{m,0}(\omega)$. Its perverse degree is zero, so both exponentials preserve P; extracting the coefficients using U gives the same bounds. Finally $\partial_y = -\widetilde{D}_{1,0}(1)$ has perverse degree -2, whereas y has degree two. In the locally finite reduction formula of [23, Lemma 6.8],

$$F_{\mathrm{red}} = \sum_{a,b\geq 0} \frac{y^a}{a!b!}(\operatorname{ad}_y^b(-\operatorname{ad}_{\partial_y})^a F)\partial_y^b,$$

the added filtration shifts sum to zero in every term. The quotient by y has the original coarse filtration: its constant inclusion and evaluation at $y=0$ are filtered inverses on the quotient. Thus the reduced operators used here satisfy

$$D_{a,b}(\gamma)P_j \subset P_{j+2b-2+i}.$$

For a flag symbol $\tau_{p,i}$ the same argument uses $i=2$. It follows that (e,h,f) on R_* has perverse degrees $(2,0,-2)$.

A relatively ample class α on U_* is pure and, by Lemma 24.5, 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

$$[y(\gamma), x(\eta)] = n\langle\gamma,\eta\rangle \quad (\gamma,\eta \in H^1(C)), \qquad [e, x(\gamma)] = y(\gamma). \tag{24.26}$$

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 (24.26); 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 $g > 0$,

$$\det B_\lambda = (2/n)^{2g}\lambda^{2g} + \text{a polynomial of degree at most } 2g-1, \tag{24.27}$$

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 $(2, 0, -2)$. The conjugation may be performed over $\mathbb{C}$; equality of the resulting rational filtrations descends from that scalar extension.

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 use the following finite-dimensional uniqueness fact. Suppose (E, H_1, F_1) and (E, H_2, F_2) are two $\mathfrak{sl}_2$-triples and $[H_1, H_2] = 0$. Then $A = H_2 - H_1$ commutes with E and has weight zero for ad_{H_1} on $\mathrm{End}(V)$. In a finite-dimensional $\mathfrak{sl}_2$-module a weight-zero vector killed by the raising operator lies in a trivial summand. Hence $[F_1, A] = 0$. Applying the same argument to the second triple gives $[F_2, A] = 0$. The commuting semisimple operators H_1, H_2 make A semisimple, and each A-eigenspace is preserved by both triples. On such an eigenspace both H_i have trace zero, since $H_i = [E, F_i]$. Thus every eigenvalue of A is zero, proving $H_1 = H_2$. Applied on $\mathrm{Gr}^P R_*$ this gives $h' = h_P$.

The integral spectral curve has arithmetic genus

$$p_a(\Sigma) = 1 + (g-1)n^2 + |D|\, n(n-1)/2 = N_F.$$

The integral Hitchin base is smooth, and its compactified Jacobian has relative dimension N_F. The unit map from the constant complex of the base consequently places the unit in centered perverse step $-N_F$: after shifting by the dimension of the total space, its source is the perverse constant complex of the base shifted by $[N_F]$. On the smooth spectral locus it is the nonzero degree-zero fiber class, so it belongs to no smaller step. The comparison $h' = h_P$ already proved, and preservation of the increasing h-filtration under the conjugations above, give h-degree $-N_F$ for the unit. Tautological generation gives $R_*^0 = \mathbf{Q}$; all degree-zero generators are numerical constants. Equation (24.24) and its flag-class counterpart now identify the 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 24.5 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 $U = U_{d_0}$,

$$P_j H^*_{\mathrm{pure}}(U) = C_{j+N_F} H^*_{\mathrm{pure}}(U). \tag{24.28}$$

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 (24.14), 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 (24.28) 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 (24.12) 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

$$H^*(X) = H^*(\overline{X})/L_B H^*(\overline{X}), \qquad H^*(X_0) = H^*(\overline{X}_0)/L_{B,0} H^*(\overline{X}_0). \tag{24.29}$$

Indeed, the complement of the core retracts to the smooth proper orbifold boundary by its line-bundle projection, so its cohomology is pure. Lowest-weight surjectivity from the compactification gives surjectivity onto the boundary. 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

$$\iota_* \iota^* = a L_B^q \quad (a \in \mathbf{Q}^\times). \tag{24.30}$$

The isomorphism modulo L_B in (24.12) 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

$$(\iota^*)^{-1}(\ker L_{B,0}^{a'} + \operatorname{im} L_{B,0}) \subset \ker L_B^{a'+q} + \operatorname{im} L_B. \tag{24.31}$$

Indeed, write $\iota^* \alpha = \beta + L_{B,0} \beta_1$ with $L_{B,0}^{a'} \beta = 0$, and use the preceding inclusion to write $\iota_* \beta_1 = L_B^q \gamma$. Projection formula and (24.30) give

$$L_B^{a'+q}(\alpha - a^{-1} L_B \gamma) = a^{-1} \iota_* L_{B,0}^{a'} \beta = 0.$$

This changes α only by an element of $\operatorname{im} L_B$.

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 (24.31) 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

$$P_j H^*(X_0) = C_{j+N_F} H^*(X_0) \tag{24.32}$$

in every genus and every numerically generic degree.

24.8. **Blockwise Gysin transfer.** Return to the original partial flags and put

$$c = \sum_{p,a} \binom{\mu_a^p}{2}, \qquad N_F = N_\mu + c. \tag{24.33}$$

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 fibrations in the graded blocks; its relative dimension is c. If the partial-flag scalar gerbe is nontrivial, the block bundles are taken on that gerbe. Their flag varieties and their endomorphism bundles have trivial scalar action and descend to the coarse spaces. The normalized full-flag roots also descend in rational cohomology and restrict to the usual Chern roots on each flag fiber. Their standard monomials give a Leray–Hirsch basis, so the rational flag-integration formula below requires no untwisted universal bundle on the partial-flag coarse space. 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 [24, Corollary 1.3.2]. The flag-bundle formula gives purity of $H^*(Z)$ as well, so the pure-cohomology bounds (24.25) 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

$$A = \pi_* i^*, \qquad B = i_* \pi^*. \tag{24.34}$$

Self-intersection and integration on the flag fibers give

$$AB = \delta_\mu \,\mathrm{id}, \qquad \delta_\mu = \pm \prod_{p,a} \mu_a^p! \neq 0. \tag{24.35}$$

Writing $M = M_{\mathrm{Dol},\mu}$, their cohomological degrees and Tate twists are

$$A : H^k(X_0) \longrightarrow H^{k-2c}(M)(-c), \qquad B : H^k(M) \longrightarrow H^{k+2c}(X_0)(c).$$

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 (24.25) gives $+c$ for i^*, i_* and $-c$ for π^*, π_*. Both composites in (24.34) 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

$$\pi_* F = \pm \frac{\mathrm{Alt}_{\prod_{p,a} S_{\mu_a^p}} F}{\Delta_\mu}, \tag{24.36}$$

expressed as a polynomial symmetric in each block. It lowers Chern index by c. Formula (24.35) 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 (24.32), perversity at most $b + c - N_F = b - N_\mu$. Apply A to obtain one inclusion in (24.5). 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 24.2 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

$$\mathcal{C} = C[\, \sqrt[s]{p} : p \in D], \qquad \nu : \mathcal{C} \longrightarrow C$$

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

$$d_0 = -\sum_{p,a} \mu_a^p \alpha_{p,a}. \tag{25.2}$$

The determinant condition makes this an integer.

25.1. **The finite-order correspondence and universal classes.** We use the nonabelian Hodge correspondence only on the stable locus. The homeomorphism of coarse moduli spaces will be denoted by

$$\Psi : M_{\mathrm{Dol},\mu}^{\mathrm{top}} \longrightarrow M_{B,\mu}(\xi)^{\mathrm{top}}.$$

The comparison of universal classes requires a continuous family construction, which we now describe.

First add three auxiliary root markings of orders at least two, with trivial inertia and parabolic weight zero, and write $\mathcal{C}^+$ for the resulting root curve. This leaves the moduli problems unchanged. Indeed, a bundle with trivial inertia at a new marking descends across that marking. In an invariant local frame, with $z = t^s$, an invariant regular differential is $t^{s-1}h(t^s)\,dt = h(z)\,dz/s$, so regular Higgs fields descend as well. Finite-group averaging makes this descent compatible with families. Local systems descend by making the new inertia generators act trivially. The ordinary-point fiber, the original inertia summands, and scalar automorphisms are preserved.

By [3, Theorem 4.2, Remark 4.3, and Proposition 5.7], we may write

$$\mathcal{C}^+ = [Y/\Gamma]$$

with Y a connected smooth projective curve and Γ a finite group. Here $Y \to \mathcal{C}^+$ is a representable finite étale Γ-torsor; the map to the coarse curve may be ramified. For genus zero the three new nontrivial root markings put us in Case C of the proof of that proposition. No global finite-quotient presentation of the original root curve is being assumed.

We check the harmonic-type hypothesis without using the correspondence. If (E, θ) is a stable degree-zero Higgs bundle on $\mathcal{C}^+$, the first Harder–Narasimhan term of its pullback to Y, if destabilizing, is Γ-invariant by uniqueness and descends to a destabilizing Higgs subbundle of E. Thus the pullback is semistable. Its maximal polystable slope-zero Higgs subbundle, or socle, is also Γ-invariant. If it were proper, it would descend to a proper slope-zero subbundle of E, contrary to stability. The pullback is therefore polystable. Degree zero gives the vanishing of the rational first Chern class, and higher Chern classes vanish on a curve. Thus the objects have harmonic type in the sense of [55, Definitions 9.1 and 9.6].

The correspondence and the coarse homeomorphism Ψ now follow from [55, Theorem 9.7]. Proposition 9.9 of that reference, with $\lambda = 0$, identifies regular Higgs fields on the root curve with parabolic Higgs fields whose residue is zero on every parabolic graded piece. The total logarithmic residue is allowed to be strictly lowering.

Lemma 25.1. *Under Ψ, the normalized topological characteristic classes of the universal orbifold bundle and of its inertia multiplicity bundles agree. The scalar-normalizing class $u = c_1(\mathcal{E}_x)/n$ also agrees on the common scalar-gerbe model used to define these normalized classes.*

Proof. Choose an ordinary point x away from all markings, a lift $y_0 \in Y$, and its free Γ-orbit O. Use frames at every point of O. Let $R_{\eta,O}$ denote the usual fine parameter scheme of framed semiharmonic Higgs bundles, or of framed flat bundles, on Y, for $\eta = \mathrm{Dol}, \mathrm{DR}$. In this proof we work in its harmonic subspace, hence with polystable degree-zero Higgs bundles or semisimple flat bundles. The full parameter schemes and their added-frame bundles have the usual analytic topology; the framed representability and added-frame description are also recorded in [55, Theorem 6.5].

Write $R^h_{\eta,O}$ for the subspace on which there exists *one* harmonic metric making every frame at O unitary. Its value in the frame at y_0 fixes the harmonic metric uniquely. The requirement is not separate existence of a possibly different metric for each frame.

The harmonic correspondence gives a homeomorphism

$$R^h_{\mathrm{Dol},O} \simeq R^h_{\mathrm{DR},O}.$$

For completeness, the passage from one frame to the finite set O uses exactly the convergence argument of [54, Corollary 7.10, Lemma 7.13, and proof of Lemma 7.16]. After the metric-preserving bundle identifications supplied there, the operators converge on a fixed smooth bundle. All the additional frames lie in finite products of $U(n)$, so a common subsequence of them converges. The operator-and-frame convergence gives convergence in the usual fine parameter and frame-bundle topologies by [53, Lemma 5.10, Corollary 5.11, and Theorem 5.12 and its proof]. In particular, the compactness identifications and the standard local trivializations of the universal family differ by maps converging to the identity in C^0. This controls evaluation at all the finitely many points of O simultaneously. Separatedness identifies the limit with the prescribed framed object. Applying this to every subsequence proves continuity, and the same argument in the opposite direction proves the displayed homeomorphism.

We also need the universal bundles, rather than only the parameter spaces. Let $\mathcal{V}_\eta$ be the universal bundle on $R^h_{\eta,O} \times Y$, with its usual topological bundle structure. Define F^h_η to be the subset of its full frame bundle consisting of one additional frame at $y \in Y$ which is unitary for the uniquely normalized harmonic metric. Initially this is merely a subspace of the ordinary frame bundle; we do not assume that the harmonic metrics form a continuous family.

The same metric-preserving convergence proves that $F^h_\eta \to R^h_{\eta,O} \times Y$ is proper. Namely, if (p_i, y_i) converges, pass to the fixed limiting smooth bundle as above and choose a smooth unitary reference frame near the limit of y_i. An arbitrary unitary frame at y_i is then represented by a matrix in the compact group $U(n)$ and has a convergent subsequence. The operator-and-frame convergence used above applies also to this frame: allowing y_i to move only introduces evaluation at convergent points in that local trivialization. The C^0 comparison with the standard universal family trivializations in [53, Lemma 5.10, Corollary 5.11, and proof of Theorem 5.12] shows that this convergence is in the ordinary full frame bundle. Equivalently, one can allow the point of evaluation to vary by base change to the product with Y and use its tautological section. The limit is a frame for the prescribed limiting object and is unitary for its normalized harmonic metric. This proves sequential properness, hence properness here, since the parameter and frame spaces are metrizable subspaces of finite-dimensional analytic spaces. It also proves closedness of F^h_η in the full frame bundle.

This properness supplies the omitted continuity check for the metric. In a local continuous complex frame, let H_i be the positive Hermitian matrix of the normalized harmonic metric at (p_i, y_i). The columns of $A_i = H_i^{-1/2}$ form a unitary frame. Properness gives a convergent subsequence in the full frame bundle. Its limit A is invertible and unitary for the limiting metric, so

$$H_i = (A_i A_i^*)^{-1} \longrightarrow (AA^*)^{-1} = H.$$

The same reasoning applies to every subsequence. The metrics are therefore continuous, and F^h_η is the usual continuous $U(n)$-reduction of the universal frame bundle.

Repeating the two-direction convergence argument with this additional, variable frame gives a homeomorphism

$$F^h_{\mathrm{Dol}} \simeq F^h_{\mathrm{DR}}$$

over the preceding base homeomorphism times Y. It commutes with the right $U(n)$ action. Associated bundles for the standard representation consequently identify the entire universal topological bundles on Y. This step is what turns pointwise harmonic correspondence into

a comparison of families on the whole curve. Framed Riemann–Hilbert then gives the same assertion on the Betti side; see [54, Theorem 7.1].

The finite group Γ acts by pullback on Y and by permutation of O. All the preceding constructions commute with this action and with changes of unitary frames. A fixed point of the Γ action on the fine multi-framed space gives descent isomorphisms preserving the frames. They are unique, since an automorphism preserving one frame is the identity. Their cocycle identities therefore follow from uniqueness. Conversely, a Γ-linearization and a frame at y_0 give the compatible frames at all points of O. The same normalized harmonic metric is Γ-invariant: its pullback by a descent isomorphism is a harmonic metric with the same value in a frame at y_0, hence is equal to it. Thus restriction to the fixed subspaces is continuous descent, not an inference from an equivalence of untopologized categories.

Restrict now to the fixed inertia type and to objects stable on $[Y/\Gamma]$. The resulting two harmonic framed spaces are identified by a $U(n)$-equivariant homeomorphism; denote their common space by P. Here $U(n)$ changes the frame at x, and hence all the compatible frames at O, diagonally. The universal bundle identification just constructed is Γ-equivariant, so it descends to an identification of $U(n)$-equivariant topological orbifold bundles on $P \times \mathcal{C}^+$. Every inertia summand is the image of the same finite-group isotypical projector in this bundle. The identified bundles therefore include all inertia multiplicity bundles.

We spell out the passage to the coarse spaces. On the stable locus the automorphisms are the scalars, and harmonic metrics are unique up to a positive scalar. Consequently each fiber of $P \to M_\eta$ is $PU(n)$. The compactness argument above, applied after restriction to Y, shows that this map is proper over compact subsets of M_η. Thus $P/PU(n) \simeq M_\eta$ with its usual topology, as in [54, Lemma 7.17] and [55, proof of Theorem 8.8]. Let $R^{\mathrm{st}}_{\eta,\Gamma}$ be the ordinary complex-framed space for the stable descended objects. It is a $PGL_n(\mathbf{C})$-torsor over the stable coarse space. The continuous inclusion $P \to R^{\mathrm{st}}_{\eta,\Gamma}$ induces, by the quotient topology just proved, a continuous section of $R^{\mathrm{st}}_{\eta,\Gamma}/PU(n) \to M_\eta$. The space P is the pullback of the principal $PU(n)$-bundle $R^{\mathrm{st}}_{\eta,\Gamma} \to R^{\mathrm{st}}_{\eta,\Gamma}/PU(n)$ along this section. Thus $P \to M_\eta$ is a continuous $PU(n)$-reduction.

Use the common Borel construction

$$\mathscr{H} = EU(n) \times_{U(n)} P.$$

The equivariant inclusions and the group inclusion $U(n) \hookrightarrow \mathrm{GL}_n(\mathbf{C})$ give the actual pullback maps

$$\jmath_\eta : \mathscr{H} \longrightarrow E\,\mathrm{GL}_n(\mathbf{C}) \times_{\mathrm{GL}_n(\mathbf{C})} R^{\mathrm{st}}_{\eta,\Gamma}.$$

These are the classifying-space maps of the continuous quotient stack maps $[P/U(n)] \to [R^{\mathrm{st}}_{\eta,\Gamma}/\,\mathrm{GL}_n(\mathbf{C})]$. The latter quotient presents the original stable complex moduli stack. Its universal bundle pulls back to the associated bundle constructed above from the $U(n)$-equivariant bundle on P; the same holds for every inertia multiplicity bundle. The maps $p_\eta : \mathscr{H} \to M_\eta$ have fiber $BU(1)$. With the standard universal-bundle convention, the central $U(1)$ acts with weight one on the ordinary-point fiber and on each inertia multiplicity bundle, on both sides of the comparison. The identified universal bundles on $\mathscr{H}$ therefore give the same class $u = c_1(\mathcal{E}_x)/n$, whose restriction to $BU(1)$ is its standard degree-two generator. Rational Leray–Hirsch gives

$$H^*(\mathscr{H}, \mathbf{Q}) \simeq H^*(M_\eta, \mathbf{Q})[u],$$

so p^*_η is injective. This use of a common Borel model does not assert an equivalence of all semistable topological moduli stacks. The replacement of scalar $\mathbf{C}^*$ by $U(1)$ preserves the class u, since its positive-real factor is contractible.

The two pulled-back Chern characters and the class u are equal, so the pulled-back expressions $e^{-u}\,\mathrm{ch}(\mathcal{E})$ and $e^{-u}\,\mathrm{ch}(\mathcal{V}_{p,a})$ are equal as well. The normalized coarse classes are detected by the injective maps p^*_η, and $p_B = \Psi \circ p_{\mathrm{Dol}}$ under the common construction. This proves their equality

under Ψ. Descending away the auxiliary trivial-inertia roots preserves these identities, and proves the lemma for the original root curve. □

This argument compares topological bundles and their rational characteristic classes. The mixed Hodge weights on the Betti side are supplied by the algebraic restriction map and the tautological weight theorem below.

25.2. **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 (24.4) are the elements invariant under the scalar-gerbe coaction and hence lie in the coarse factor.

For the finite-cover, transfer, and coarse/root calculations in the remainder of this section, we use the augmented root curve $\mathcal{C}^+$ constructed above and abbreviate it again by $\mathcal{C}$. The coarse curve C is unchanged. In these calculations D includes the auxiliary root markings; their orders remain fixed, their inertia is required to be trivial on the moduli space under consideration, and their sole nonzero isotropy block is the character-zero block of rank n. Thus their parabolic weight is zero, and adding them changes neither d_0 nor N_μ. At the start of Section 26, the marking and profile notation again refers to the original punctures and profiles; the auxiliary parameters are never varied.

To justify this temporary notation, let $\rho : \mathcal{C}^+ \to \mathcal{C}_{\mathrm{orig}}$ forget the auxiliary roots. Rationally, $R\rho_*\mathbf{Q} \simeq \mathbf{Q}$: at a new root marking the additional topological fiber is the classifying stack of a finite group, whose positive degree rational cohomology vanishes. Hence ρ induces the canonical rational cohomology and homology identifications, compatible with the maps to C. On the locus of trivial new inertia, the universal bundle is the pullback of the original one; the original inertia summands and the ordinary-point normalizing class are also preserved. The same identities hold for families by the descent proved above. In the subsequent coarse/root formula, every auxiliary term has character $c = 0$ and therefore coefficient $c/s_p = 0$.

We now take $C' = Y$ and use the actual finite étale Γ-torsor of degree $a = |\Gamma|$,

$$f : C' = Y \longrightarrow \mathcal{C} = \mathcal{C}^+. \tag{25.4}$$

Restriction is defined on the entire local-system stack of this augmented curve:

$$R_f : \mathrm{Loc}_{\mathrm{GL}_n}(\mathcal{C}) \longrightarrow \mathrm{Loc}_{\mathrm{GL}_n}(C'). \tag{25.5}$$

After applying this map and the transfer identity, we restrict to the prescribed original local monodromies and to trivial inertia at the auxiliary markings. This locus is identified with the original Betti moduli problem by ρ. The map obtained by composing f with ρ is not used as an étale cover. The restriction morphism (25.5) is algebraic on 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.

25.3. **The homogeneous weights of orbifold tautological classes.** We recall and explain the precise universal-class weight statement [51]. 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. The algebraic pullback on representation stacks preserves mixed Hodge structures. 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

$$\prod_{b \neq a} \frac{M_p - \xi_{p,b}\,\mathrm{id}}{\xi_{p,a} - \xi_{p,b}}.$$

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.

25.4. **The coarse/root comparison is triangular.** The Dolbeault generators in Theorem 24.2 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 $0 \leq c < s_p$. 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 25.2. *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 identification of Lemma 25.1.

Proof. Theorem 24.2 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 25.3. *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. If the variety is nonempty, its complex dimension is $d = 2N_\mu$, and the corresponding compactly supported assertion is obtained by Poincaré duality, including the cohomological degree shift and Tate twist by d. For an empty variety both cohomological assertions are zero identities.

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 24.2 and Proposition 25.2 apply and give (25.11). For a smooth d-dimensional space, the duality being used is

$$H_c^i(M) \simeq H^{2d-i}(M)^\vee(-d).$$

In particular a weight-w summand before duality becomes weight $2d - w$. Under this pairing $P_j H_c^i$ is the annihilator of $P_{-j-1} H^{2d-i}$. Since $d = 2N_\mu$ and all Betti weights are even,

$$P_j H_c^i(M_{\mathrm{Dol},\mu}) = W_{2(j+N_\mu)} H_c^i(M_{B,\mu}) = W_{2(j+N_\mu)+1} H_c^i(M_{B,\mu}).$$

This transports the perverse filtration across the topological identification; it does not identify the two mixed Hodge structures. □

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 [42, §§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 $B_\mu \times \{(1, \dots, 1)\}$. Here genericity imposes the proper-subprofile product exclusions and permits repeated entries of t. The Grothendieck–Springer family $\mathcal{X}_{\mathrm{flag}} \to V$ is the geometric quotient of tuples

$$(b, t, A_1, B_1, \dots, A_g, B_g, M_1, \dots, M_m, Y, F),$$

where $M_p \in \mathcal{C}_p(b)$, F is a full flag, $YF_i \subset F_i$, the action of Y on F_i/F_{i-1} is t_i, and

$$\prod_{i=1}^{g} [A_i, B_i] \prod_{p=1}^{m} M_p \, Y = I_n.$$

The quotient is by simultaneous PGL_n conjugation. Write $f_{\mathrm{flag}} : \mathcal{X}_{\mathrm{flag}} \to V$ for the projection and $X_{\mathrm{flag},(b,t)}$ for its 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.

Let $B_n \subset \mathrm{GL}_n$ be the upper triangular Borel, T_n its diagonal torus, U_n its unipotent radical, and Z_n the scalar center. Put $c_n = \dim U_n = \binom{n}{2}$. Let $\mathcal{R}_{B_n}$ be the relative representation space with the final flag fixed to the standard flag. Its last matrix has the unique form $Y = ut$, with $u \in U_n$, and its relation is

$$\prod_{i=1}^{g} [A_i, B_i] \prod_{p=1}^{m} M_p \, u = t^{-1}. \tag{26.3}$$

The defining equations hold over all of V, including $t = (1, \dots, 1)$.

Genericity makes every tuple irreducible, so $\mathcal{R}_{B_n} \to \mathcal{X}_{\mathrm{flag}}$ is a principal B_n/Z_n bundle. This bundle is Zariski locally trivial. Indeed, B_n/Z_n is an extension of a split torus by a unipotent group filtered by additive groups. Torus torsors are locally products of line-bundle frame torsors;

after trivializing them on affine opens, the successive additive torsors split by vanishing of quasi-coherent first cohomology. Define

$$\widetilde{\mathcal{X}}_{\mathrm{flag}} = \mathcal{R}_{B_n}/(T_n/Z_n), \qquad \epsilon : \widetilde{\mathcal{X}}_{\mathrm{flag}} \longrightarrow \mathcal{X}_{\mathrm{flag}}, \qquad \widetilde{f} = f_{\mathrm{flag}}\epsilon. \tag{26.4}$$

The local trivializations identify the fiber of ϵ with $B_n/T_n \simeq U_n \simeq \mathbf{A}^{c_n}$. Thus ϵ is a Zariski locally trivial affine-space bundle of rank c_n on the entire parameter space V.

26.2. **Uniform puncture and walk strata.**

Proposition 26.2. *The family $\widetilde{f} : \widetilde{\mathcal{X}}_{\mathrm{flag}} \to V$ has a finite filtration $\varnothing = Z_{-1} \subset Z_0 \subset \cdots \subset Z_s = \widetilde{\mathcal{X}}_{\mathrm{flag}}$ 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. Here and below a model for a handle includes its torus surface factor. When $g > 0$ the fixed model is the convolution of braid factors and the marked surface factors of [42, §7.4, especially (7.9)–(7.10) and Remark 7.6]. It is not necessary to identify it with an ordinary positive braid variety alone. For example, at rank one each handle still contributes two multiplicative coordinates.

Put $R = \mathcal{O}(V)$. All of the following changes of coordinates are made over this ring. For a puncture, the conjugacy family is $B_\mu \times G/L_p$: its inverse identification with the matrix family is algebraic by the spectral projectors, because all differences of distinct block eigenvalues are units on B_μ. Choose the fixed parabolic containing L_p and its unipotent radical N_p. In a root coordinate x_ρ, ordered by root height, the map

$$z \longmapsto C_p^{-1} z C_p z^{-1} \tag{26.5}$$

has the form

$$x_\rho \longmapsto (\rho(C_p)^{-1} - 1)x_\rho + Q_\rho((x_\eta)_{\mathrm{ht}(\eta)<\mathrm{ht}(\rho)}).$$

The choice of ρ or its inverse only changes the displayed ratio. Every coefficient of x_ρ is a unit: a root of N_p joins different eigenvalue blocks. Recursive substitution gives the inverse over R. The other puncture changes use multiplication in fixed unipotent groups, fixed permutations, and diagonal conjugation. They are relative algebraic isomorphisms as well.

For a handle, use its two fixed Bruhat indices (σ, τ), and write its two torus coordinates as $a, b \in T_n$. After the unipotent factors have been moved as in [42, §7.4], the intervening torus factor is

$$\tau^{-1}(a)\, b\, a^{-1}\, \sigma^{-1}(b^{-1}).$$

This expression and all its changes of coordinates have no puncture eigenvalue parameter. It is the moment map of the fixed marked surface associated to (σ, τ), inserted between the four braid factors. The complete handle contribution, including these four factors, is equivariant for the identity permutation: its permutation is

$$(\sigma\tau)(\tau^{-1}\sigma^{-1}\tau\sigma)(\sigma^{-1}\tau^{-1}) = 1.$$

Gluing all handles and puncture factors therefore gives a fixed convolution model of total permutation 1. All varying puncture eigenvalues occur only in its target

$$C_\pi(b, t) = t^{-1} \prod_p \pi_p(C_p(b)^{-1}). \tag{26.6}$$

This argument uses the upper-triangular variable $Y = ut$ already present in the full-flag family. No inverse of $u \mapsto utu^{-1}t^{-1}$ is used near $t = 1$.

For each fixed first-stage Bruhat choice, record the Bruhat cells of the successive partial products in the convolution model. Torus factors preserve these cells. The resulting finite walk map is algebraic as a map to the finite Bruhat stratification, with closed inverse images of lower sets. The first-stage maps have the same property. Refining these two finite stratifications lexicographically and choosing a linear extension gives a finite closed filtration. Make the target

pullback over V before descending by the residual torus. On this generic pullback the quotient is a torsor for T_n/Z_n, so invariant closed subsets descend to closed subsets. This proves the claimed closed filtration on $\widetilde{\mathcal{X}}_{\mathrm{flag}}$, uniformly over V.

The permutation of the individual toric walk surface. The total equivariance permutation of the complete convolution must also be compared with the boundary permutation after the walk reduction. Read the factors in the order of multiplication used in the walk, adjoining each new factor on the left. Write p for the current Bruhat index and Π for the equivariance permutation of the current partial convolution. The toric surface accumulated by the walk has permutation

$$\theta = p^{-1}\Pi.$$

For braid factors this is the induction in [42, Corollary 5.25]: an up or down step changes (p, Π) to $(\tau_i p, \tau_i \Pi)$ and introduces no toric piece, whereas a stay step leaves p unchanged and introduces the elementary toric piece with permutation $p^{-1}\tau_i p$.

The same invariant is preserved by a diagonal surface factor. Let $h : S \to T_n$ be its moment map and π_S its equivariance permutation. For a partial Bruhat factorization $F = updv$, with $d \in T_n$, one has

$$h(x)F = \big(h(x)uh(x)^{-1}\big)p\,\big(p^{-1}(h(x))d\big)v.$$

Thus its insertion leaves the Bruhat index p unchanged and relabels the surface by p^{-1}. Its toric permutation is consequently $p^{-1}\pi_S p$. The convolution rule for equivariance therefore gives

$$\theta_{\mathrm{new}} = (p^{-1}\pi_S p)\theta = p^{-1}\pi_S\Pi = p^{-1}\Pi_{\mathrm{new}},$$

as required. These relabelings are fixed on each walk and are independent of all eigenvalue parameters.

At the endpoint, the target equation makes the complete matrix product upper triangular: it is an upper-unipotent factor times the diagonal target $C_\pi(b, t)$. The identity is the Bruhat index of the Borel, so $p_{\mathrm{final}} = 1$. The complete convolution has $\Pi_{\mathrm{total}} = 1$ by the puncture-pair cancellation and the complete handle calculation above. Hence the individual toric walk surface has

$$\theta_{\mathrm{final}} = p_{\mathrm{final}}^{-1}\Pi_{\mathrm{total}} = 1.$$

Its boundary permutation is therefore the identity. Combined with the fact that every boundary circle carries marked points, this is exactly the assertion needed to conclude that each boundary circle has one black point and that $\psi(e_i)$ traverses that entire circle once.

For a walk in the complete convolution model let (Σ, A) be its glued marked surface and put $\Lambda = H_1(\Sigma, A; \mathbb{Z})$. The handle pieces are included in this definition of Σ. These are the marked surfaces of [42, Definition 6.1]: every component has boundary and every boundary component has alternating black and white marked points. The elementary braid pieces satisfy this condition. A handle piece is the finite cover of a once-punctured torus in [42, Proposition 7.4]; each component covers the base and each boundary circle covers its boundary circle, hence contains preimages of both marked points.

Only the prescribed convolution gluing of [42, §6.3, Proposition 6.5] is used. It identifies the outgoing boundary arc at black point i on the left with the incoming arc at black point i' on the right, retaining the identified point as a black boundary point. The resulting marked set is precisely the right-hand set A'; all of its points remain on the boundary. Each left component has such a gluing point and hence joins a component carrying a retained point of A'; each right component already carries one. Thus every resulting component meets A'. The remaining boundary arcs still alternate black and white vertices. Induction over the convolution excludes unmarked closed components; this is not a claim about arbitrary gluing of surfaces.

The torus moment map has, up to a fixed sign translate, cocharacter homomorphism

$$\partial : \Lambda \longrightarrow \ker\big(\mathbb{Z}^A \longrightarrow H_0(\Sigma; \mathbb{Z})\big). \tag{26.7}$$

This follows by gluing the elementary braid and handle surface pieces: their lattice maps are relative boundaries, and the boundary homomorphism commutes with gluing. If Σ is disconnected, choose one proper union of components. The target equation on its marked points is the product-one equation for a same-rank proper selection of eigenvalues at each original puncture and at the final puncture. For an elementary braid toric piece the product of its signs on its two labels is the sign of its transposition. A handle piece has no sign translate; on a union of its covering components its boundary permutation is a commutator and has sign one. Transporting the label sets through the fixed relabelings and multiplying gives the sign of the complete permutation on the chosen union. That permutation is the identity, so this product of signs is one. The defining genericity of V excludes this equation.

For connected Σ, relative homology gives the integral exact sequence

$$0 \longrightarrow H_1(\Sigma;\mathbb{Z}) \longrightarrow \Lambda \xrightarrow{\partial} \mathbb{Z}_0^A \longrightarrow 0, \qquad \mathbb{Z}_0^A = \ker\bigl(\mathbb{Z}^A \xrightarrow{\sum} \mathbb{Z}\bigr).$$

In particular ∂ is surjective as a homomorphism of integral lattices. The scalar subgroup is the entire stabilizer of the gauge torus. This can also be seen directly: its characters on the nonzero coordinates are differences of vertex coordinates, and a spanning tree of the connected graph makes all vertex coordinates equal. A spanning tree generates the primitive type-A root lattice, so there is no additional finite stabilizer.

For completeness, saturation after the quotient can be read from surface homology. Let $\bar{\Sigma}$ be obtained by capping all boundary components. The permutation of a marked surface sends each black point to the next black point along its boundary. Its cycles therefore correspond exactly to the boundary components. Since the total permutation is 1, each boundary component contains exactly one black point a_i. The gauge map ψ sends e_i to the boundary path from a_i to its next black point; in this case that path traverses the entire boundary circle once. Thus $\psi(e_i)$ is its oriented fundamental cycle, up to the common choice of clockwise versus counterclockwise orientation. All boundary circles occur, so the gauge image in $H_1(\Sigma;\mathbb{Z})$ is exactly the image of $H_1(\partial\Sigma;\mathbb{Z})$, not merely a finite-index sublattice. Moreover,

$$H_1(\Sigma;\mathbb{Z})/\operatorname{im} H_1(\partial\Sigma;\mathbb{Z}) \simeq H_1(\bar{\Sigma};\mathbb{Z})$$

is free. Dividing by the effective gauge torus consequently leaves the integral exact sequence

$$0 \longrightarrow H_1(\bar{\Sigma};\mathbb{Z}) \longrightarrow \Lambda/\operatorname{im}\psi \longrightarrow \mathbb{Z}_0^A \longrightarrow 0.$$

Its right-hand term is free, so choose an integral splitting once for each of the finitely many walks. The induced torus map is a product projection. Its pullback by $C_\pi : V \to T_1$ is therefore a product of V with a fixed torus, after undoing the fixed sign translate. The affine factors descend as affine-space bundles of their fixed ranks; their compact-support direct images are given by their canonical traces. Thus a stratum with affine rank a and torus rank b has

$$Rf_{a!}\mathbb{Q} \simeq \mathbb{Q}_V \otimes R\Gamma_c((\mathbb{C}^*)^b,\mathbb{Q})(-a)[-2a]$$

in the derived category with its mixed Hodge structure. This is an isomorphism of relative objects and is preserved by parameter base change.

□

26.3. **Localization in mixed Hodge modules.** Use the actual localization triangles for the closed filtration just constructed. If two terms of a triangle have locally constant ordinary cohomology sheaves, the long exact cohomology sequence expresses each cohomology sheaf of the third as an extension of a kernel and a cokernel of maps of local systems. It is again a local system. In mixed Hodge modules these are smooth objects, equivalently admissible variations shifted by $\dim V$, and are closed under subquotients and extensions. The strict weight filtration therefore gives local systems on every weight grade. In this statement the weight index for a smooth Hodge module is reduced by $\dim V$ to obtain the fibre-weight convention. All boundary

maps are included in the localization triangles; no degeneration of a spectral sequence is asserted or needed.

For the affine-space bundle (26.4), the smooth trace is a canonical isomorphism

$$R\epsilon_!\mathbf{Q}_{\widetilde{\mathcal{X}}_{\mathrm{flag}}}(c_n)[2c_n] \xrightarrow{\sim} \mathbf{Q}_{\mathcal{X}_{\mathrm{flag}}}. \tag{26.8}$$

On a trivializing open it is the trace $H_c^{2c_n}(\mathbf{A}^{c_n},\mathbf{Q})(c_n)\to\mathbf{Q}$; all other compact fiber cohomology groups vanish. Complex orientation makes these traces agree on overlaps. The map is defined in mixed Hodge modules, and compact-support base change identifies every parameter pullback with the trace of the pulled-back bundle. Consequently

$$Rf_{\mathrm{flag}!}\mathbf{Q}_{\mathcal{X}_{\mathrm{flag}}} \simeq R\widetilde{f}_!\mathbf{Q}_{\widetilde{\mathcal{X}}_{\mathrm{flag}}}(c_n)[2c_n]. \tag{26.9}$$

The original full-flag family therefore has the same local-constancy and Tate assertions, with this explicit shift and twist.

26.4. **Relative Springer sign extraction.** Here is an explicit relative construction of the sign identification. Let $\mathscr{N}_n$ be the variety of all unipotent matrices, let $\pi_n:\widetilde{\mathscr{N}_n}\to\mathscr{N}_n$ be its Springer flag resolution, and let $i_1:\{I_n\}\hookrightarrow\mathscr{N}_n$. In the usual Springer normalization the universal sign summand is

$$\big(R\pi_{n*}\mathbb{Q}_{\widetilde{\mathscr{N}_n}}\big)^{\mathrm{sgn}} \simeq i_{1*}\mathbb{Q}[-2c_n](-c_n).$$

The morphism is restriction to the identity fibre followed by integration along its flag variety, as in [42, Proposition 8.5]; its Tate twist is the top class of that flag variety. To obtain the action in mixed Hodge modules, start with the finite étale regular-semisimple Grothendieck cover. Its deck action preserves the variation of type $(0,0)$. Smallness identifies its direct image with the intermediate extension, so the action extends in pure Hodge modules. Proper base change to the nilpotent cone gives the restriction-normalized Springer action; translate by $x\mapsto I_n+x$ to the unipotent variety [2, §2.5 and §§3.4–3.5]. Restriction to the identity and integration over the flag variety are morphisms in mixed Hodge modules [48]. Apply this morphism to the image of the genuine idempotent $e_{\mathrm{sgn}}=(n!)^{-1}\sum_{w\in\mathfrak{S}_n}\mathrm{sgn}(w)w$ in that category. Its underlying rational map on this summand is the isomorphism of [42, Proposition 8.5]. Conservativity therefore gives the displayed isomorphism of mixed Hodge modules. These constructions are conjugation-equivariant, giving the same action and map in the equivariant category.

Let $\mathcal{R}_{\mathrm{uni}}$ be the framed representation family over B_μ with the final matrix unipotent. The map $\mathcal{R}_{\mathrm{uni}}\to\mathscr{N}_n$ sends a tuple to that matrix. The framed flag family at $t=1$ is its Cartesian pullback of π_n, and the inverse image of I_n is the original framed family. Proper base change pulls back the displayed universal isomorphism. Since the effective conjugation action is free, the equivariant isomorphism descends to the geometric quotients. Writing π for the resulting flag map and $i:\mathcal{X}_\mu\hookrightarrow\mathcal{X}_{\mathrm{uni}}$ for the identity-monodromy inclusion gives, over the whole parameter base,

$$\big(R\pi_*\mathbb{Q}_{\mathcal{X}_{\mathrm{flag},1}}\big)^{\mathrm{sgn}} \simeq i_*\mathbb{Q}_{\mathcal{X}_\mu}[-2c_n](-c_n).$$

Pushing to B_μ by compact supports yields

$$\big(Rf_{\mathrm{flag},1!}\mathbb{Q}\big)^{\mathrm{sgn}} \simeq Rf_!\mathbb{Q}[-2c_n](-c_n).$$

Thus the comparison is an isomorphism of relative direct images, not only an isomorphism independently chosen at each point. The sign idempotent is horizontal and weight preserving. Together with the affine trace and the preceding localization argument, this proves the claimed local constancy when all the original eigenvalues vary.

In particular, on each fiber the normalization is

$$H_c^i(X_{\mathrm{flag},(b,1)})^{\mathrm{sgn}} \simeq H_c^{i-2c_n}(\mathcal{X}_{\mu,b})(-c_n). \tag{26.10}$$

The relative sign idempotent is

$$e_{\mathrm{sgn}} = \frac{1}{n!} \sum_{\sigma \in S_n} \mathrm{sgn}(\sigma)\, \sigma.$$

Its image in each weight-graded local system is a local system. Undoing the displayed shift and twist gives (26.2), with Tate fibers.

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 relatively open in 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.

The convention is $0 \le \alpha_{p,a} < 1$: weight zero is allowed, weight one is represented by zero, and an auxiliary identity marking has its weight fixed at zero. For the existential choice in Theorem 27.1, choose the generic unitary point also outside the loci $\xi_{p,a} = 1$ for every unfixed coordinate that is nonconstant on its determinant component. This is possible by the same finite proper-character avoidance argument. With at least two unfixed coordinates every coordinate is nonconstant: a coordinate character cannot be proportional to the determinant exponent vector, which has positive entries in all those coordinates. If there is only one unfixed coordinate, the determinant component is a point, and the zero-dimensional case is treated directly.

The chosen point consequently has $0 < \alpha_{p,a} < 1$ in every nonconstant coordinate. A sufficiently small logarithm neighborhood then gives a nonempty relatively open subset of the rational age hyperplane H_S inside the weight box. No branch discontinuity is crossed. The additional chamber assertion is understood for a chamber relatively open in H_S after fixed auxiliary coordinates are omitted. It does not assert avoidance of a hyperplane containing an imposed boundary face merely because that hyperplane does not contain all of H_S. If a separate application prescribes such a face, the noncontainment hypothesis must refer to the affine span of that face. The finite active-slope wall argument and rational density apply to the relative open chamber exactly as stated. In the zero-dimensional case a hyperplane not containing H_S misses it, so no further perturbation is needed.

$$H_S = \{\alpha : \mu \cdot \alpha = S\}. \tag{27.4}$$

Its component is fixed by $\zeta = e^{-2\pi iS/d}$, 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 14.5. 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 25.3, 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

$$[R^n] \,\mathrm{Coeff}_{\boldsymbol{\mu}} \,\mathrm{Log}\, Z_{\mathrm{nil}} = \frac{x\mathbf{T}^{-N}}{(1-x)(1-x\mathbf{T})} \sum_i (-x)^i A_i. \tag{28.3}$$

This is Theorem 14.10, 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 and h is proper. Apply the shared weight and duality calculation of Subsection 12.2 to $(M, \mathcal{A}, N, P_i)$. In particular $H_c^j(\mathcal{A}, P_i)$ is pure of weight $2N + i + j$. If

$$F_{\boldsymbol{\mu}}(x) = x^{-1}\mathbf{T}^N (1-x)(1-x\mathbf{T})[R^n] \,\mathrm{Coeff}_{\boldsymbol{\mu}} \,\mathrm{Log}\, Z_{\mathrm{nil}},$$

then (28.3) yields

$$[H_c^j(\mathcal{A}, P_i)] = (-1)^{i+j} \,\mathrm{Gr}^W_{2N+i+j}\big([x^i] F_{\boldsymbol{\mu}}(x)\big). \tag{28.4}$$

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 25.3,

$$\beta_{D,k} = \dim \mathrm{Gr}^W_{2k} H^D(\mathcal{M}_{\boldsymbol{\mu}}, \mathbf{Q}), \qquad H(\mathcal{M}_{\boldsymbol{\mu}}; Q, u) = \sum_{D,k} \beta_{D,k} Q^k u^D. \tag{28.7}$$

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). As in (12.17), work coefficientwise in rank and alphabet degree in the rational-function field $\mathbb{Q}(x,s)$. 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). The box cancellation in (12.18), with $h = g-1$, cancels the baseline in (28.2) partition by partition. Thus the resulting series is the rank-marked version of $\Omega_{g,m}$ in (2.7).

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

$$\mathbb{H}_{\boldsymbol{\mu}}(z,w) = (zw)^{-2N} \sum_i (-z^{-2})^i E_c(P_i; zw). \tag{28.9}$$

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

$$\mathbb{H}_{\boldsymbol{\mu}}(z,w) = \sum_{D,k} \beta_{D,k} (-1)^D z^{2k-D} w^{d_{\boldsymbol{\mu}}-D}. \tag{28.10}$$

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,

$$(u\sqrt{Q})^d (-1)^D (\sqrt{Q})^{2k-D} \left(-\frac{1}{u\sqrt{Q}}\right)^{d-D} = Q^k u^D.$$

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

$$\begin{aligned} H_c(\mathcal{M}_{\boldsymbol{\mu}}; Q, u) &= Q^d u^{2d} H(\mathcal{M}_{\boldsymbol{\mu}}; Q^{-1}, u^{-1}) \\ &= (u\sqrt{Q})^d \mathbb{H}_{\boldsymbol{\mu}}(Q^{-1/2}, -u\sqrt{Q}). \end{aligned}$$

Transposing every summation partition in (2.7) 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 $D \leq d$ and $2k \geq D$ whenever $\beta_{D,k} \neq 0$. 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 (2.8) 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] P. N. Achar, A. Henderson, D. Juteau and S. Riche, *Weyl group actions on the Springer sheaf*, arXiv:1304.2642v2 (2013).

[3] K. Behrend and B. Noohi, *Uniformization of Deligne–Mumford curves*, arXiv:math/0504309 (2005), Proposition 5.7.

[4] A. A. Beilinson, J. Bernstein and P. Deligne, *Faisceaux pervers*, Astérisque **100** (1982), 5–171; Numdam digital edition.

[5] O. Biquard and P. Boalch, *Wild non-abelian Hodge theory on curves*, Compos. Math. **140** (2004), no. 1, 179–204.

[6] 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.

[7] J. Briançon, M. Granger and J.-P. Speder, *Sur le schéma de Hilbert d'une courbe plane*, Ann. Sci. Éc. Norm. Supér. (4) **14** (1981), no. 1, 1–25; Numdam digital edition.

[8] T. Braden, *Hyperbolic localization of intersection cohomology*, Transform. Groups **8** (2003), no. 3, 209–216; arXiv:math/0202251v5.

[9] C. Brav, V. Bussi and D. Joyce, *A 'Darboux theorem' for derived schemes with shifted symplectic structure*, arXiv:1305.6302v4 (2018).

[10] C. Brav, V. Bussi, D. Dupont, D. Joyce and B. Szendrői, *Symmetries and stabilization for sheaves of vanishing cycles*, J. Singul. **11** (2015), 85–151; doi:10.5427/jsing.2015.11e.

[11] C. Bu, B. Davison, A. Ibáñez Núñez, T. Kinjo and T. Pădurariu, *Cohomology of symmetric stacks*, arXiv:2502.04253v2 (2025).

[12] 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.

[13] 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.

[14] 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.

[15] 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.

[16] B. Davison and S. Meinhardt, *Cohomological Donaldson–Thomas theory of a quiver with potential and quantum enveloping algebras*, arXiv:1601.02479v5 (2020).

[17] 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).
[18] F. Déglise, F. Jin and A. A. Khan, *Fundamental classes in motivic homotopy theory*, arXiv:1805.05920v3 (2021).
[19] P. Descombes, *Hyperbolic localization in Donaldson–Thomas theory*, arXiv:2506.22400v1 (2025).
[20] A. Gitlin and D. Keating, *A vertex model for supersymmetric LLT polynomials*, arXiv:2110.10273v1 [math.CO], 2021.
[21] J. Haglund, M. Haiman, and N. Loehr, *A combinatorial formula for Macdonald polynomials*, arXiv:math/0409538v1 [math.CO], 2004.
[22] 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.
[23] T. Hausel, A. Mellit, A. Minets and O. Schiffmann, $P = W$ *via* $\mathcal{H}_2$, arXiv:2209.05429v3 (2026).
[24] T. Hausel and F. Rodriguez-Villegas, *Cohomology of large semiprojective hyperkähler varieties*, Astérisque **370** (2015), 113–156; arXiv:1309.4914v1.
[25] 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.
[26] 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.
[27] A. A. Khan, *Intersection theory on derived stacks*, arXiv:1909.01332v2 (2026).
[28] A. A. Khan, *Algebraic K-theory of quasi-smooth blow-ups and cdh descent*, Annales Henri Lebesgue **3** (2020), 1091–1116, doi:10.5802/ahl.55, arXiv:1810.12858. Theorem numbering and proof locations used here are those of the author version dated 2020-02-01, `https://www.preschema.com/papers/kblow.pdf`.
[29] T. Kinjo, H. Park and P. Safronov, *Cohomological Hall algebras for* 3*-Calabi–Yau categories*, arXiv:2406.12838v2 (2026).
[30] T. Kinjo, *Dimensional reduction in cohomological Donaldson–Thomas theory*, Compos. Math. **158** (2022), 123–167; arXiv:2102.01568v2.
[31] 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.
[32] 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.
[33] J. C. Lee and S. Lee, *Generators for the cohomology of the moduli space of irregular parabolic Higgs bundles*, arXiv:2402.05380v2 (2024).
[34] F. Nironi, *Moduli spaces of semistable sheaves on projective Deligne–Mumford stacks*, arXiv:0811.1949v2, Proposition 4.20 and Section 6.
[35] 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.
[36] L. Maxim, M. Saito and J. Schürmann, *Symmetric products of mixed Hodge modules*, arXiv:1008.5345v3 (2010).
[37] I. G. Macdonald, *Symmetric products of an algebraic curve*, Topology **1** (1962), no. 4, 319–343, doi:10.1016/0040-9383(62)90019-8.
[38] 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.
[39] 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.
[40] D. Maulik and J. Shen, *The* $P = W$ *conjecture for* GL_n, Ann. of Math. **200** (2024), no. 2, 529–556; arXiv:2209.02568v2.
[41] 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.
[42] A. Mellit, *Toric stratifications of character varieties*, Publ. Math. Inst. Hautes Études Sci. **142** (2025), 153–240, doi:10.1007/s10240-025-00158-0.
[43] 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.
[44] H. Nakajima, *More lectures on Hilbert schemes of points on surfaces*, in *Development of Moduli Theory—Kyoto 2013*, Adv. Stud. Pure Math. **69** (2016), 173–205; arXiv:1401.6782v1, Proposition 3.19.
[45] S. Nayak and P. Sastry, *Grothendieck duality and transitivity II: traces and residues via Verdier's isomorphism*, arXiv:1903.01783v2 (2019).
[46] T. Pădurariu and Y. Toda, *Topological K-theory of quasi-BPS categories for Higgs bundles*, arXiv:2409.10800v1 (2024).
[47] T. Pantev, B. Toën, M. Vaquié and G. Vezzosi, *Shifted symplectic structures*, arXiv:1111.3209v4 (2013).

[48] M. Saito, *Mixed Hodge modules*, Publ. Res. Inst. Math. Sci. **26** (1990), no. 2, 221–333, doi:10.2977/prims/1195171082.
[49] M. Saito, *Thom–Sebastiani theorem for Hodge modules*, arXiv:2605.05782v1 (2026).
[50] T. Schürg, B. Toën and G. Vezzosi, *Derived algebraic geometry, determinants of perfect complexes, and applications to obstruction theories for maps and complexes*, author version, April 2013, Section 3.1.
[51] V. Shende, *The weights of the tautological classes of character varieties*, Int. Math. Res. Not. IMRN (2017), no. 22, 6832–6840; arXiv:1411.4975v2.
[52] C. T. Simpson, *Harmonic bundles on noncompact curves*, J. Amer. Math. Soc. **3** (1990), no. 3, 713–770.
[53] C. T. Simpson, *Moduli of representations of the fundamental group of a smooth projective variety. I*, Publ. Math. IHÉS **79** (1994), 47–129, publisher's PDF.
[54] C. T. Simpson, *Moduli of representations of the fundamental group of a smooth projective variety. II*, Publ. Math. IHÉS **80** (1994), 5–79, publisher's PDF.
[55] C. Simpson, *Local systems on proper algebraic V-manifolds*, Pure Appl. Math. Q. **7** (2011), no. 4; arXiv:1010.3363v1, Theorem 9.7 and Proposition 9.9.
[56] T. Spaide, *Shifted symplectic and Poisson structures on spaces of framed maps*, arXiv:1607.03807v1 (2016).
[57] The Stacks Project Authors, *The Stacks Project*, stacks.math.columbia.edu; Tags 08GW, 08H9, 08GG, 08HB, 0A1G, 0G99 and 0656.
[58] J. Stoppa and R. P. Thomas, *Hilbert schemes and stable pairs: GIT and derived category wall crossings*, arXiv:0903.1444v3 (2009).
[59] R. Sulzgruber, *Inserting rim-hooks into reverse plane partitions*, arXiv:1710.09695v2 (2018).
[60] B. Toën and M. Vaquié, *Moduli of objects in dg-categories*, author version, April 2007, Lemma 3.1.
[61] Y. Toda, *Categorical Donaldson–Thomas theory for local surfaces*, arXiv:1907.09076v4 (2021).
[62] Y. Toda, *Curve counting theories via stable objects I. DT/PT correspondence*, arXiv:0902.4371v3 (2009).
[63] S. Tubach, *Mixed Hodge modules on stacks*, Forum Math. Sigma **13** (2025), e175; arXiv:2407.02256v3.
[64] 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.
[65] K. Yoshioka, *Twisted stability and Fourier–Mukai transform*, arXiv:math/0106118v1, Definition 5.1 and Theorem 7.1 (2001).
[66] S. Monavari, E. Pavia and A. T. Ricolfi, *Derived hyperquot schemes*, arXiv:2409.16858v2 (2026).
[67] J. Hall and D. Rydh, *Perfect complexes on algebraic stacks*, arXiv:1405.1887v3 (2017).
[68] D. Ben-Zvi, J. Francis and D. Nadler, *Integral transforms and Drinfeld centers in derived algebraic geometry*, arXiv:0805.0157v5 (2010).
[69] F. Nironi, *Grothendieck duality for Deligne–Mumford stacks*, arXiv:0811.1955v2 (2009).
[70] D. Calaque, *Lagrangian structures on mapping stacks and semi-classical TFTs*, arXiv:1306.3235v2 (2014).
[71] M. Haiman, *Macdonald polynomials and geometry*, in *New Perspectives in Algebraic Combinatorics*, MSRI Publications **38** (1999), 207–254; author manuscript dated March 14, 1999.

Zhili College, Tsinghua University, Beijing 100084, China
Email address: tangqy24@mails.tsinghua.edu.cn