hexagon:2611.00427v2

A PROOF OF THE HAUSEL–RODRIGUEZ-VILLEGAS CONJECTURE

QIYUE TANG

Abstract. We prove the Hausel–Rodriguez-Villegas conjecture on the mixed Hodge polynomials of primitive-twisted GL_n character varieties.

Contents

1. Introduction 2
1.1. Background and previous results 3
1.2. Proof of the formula 5
1.3. Notation and coefficients 5
1.4. Moduli and the two support categories 6
1.5. Established inputs 7
2. Critical models and their orientations 7
2.1. The framed critical model 7
2.2. The full deformation complex and its trace 9
2.3. Closed forms and the determinant square 12
2.4. The fixed positive quotient 13
2.5. Support gluing and the derivative 14
2.6. The comparison over the full parameter ring 16
3. Monodromic coefficients and localization 18
3.1. An explicit coefficient projector 18
3.2. The integral localization map 18
3.3. One-sided constructibility on finite-type windows 20
4. Harder–Narasimhan restrictions 20
4.1. The framed heart and its numerical bounds 20
4.2. Complete coherent flags 22
4.3. The closed-stratum counit 23
4.4. Vacuum removal and finite wall transport 24
4.5. The comparison arrow and its endpoints 26
5. The Hecke relation on the comparison arrow 29
5.1. Directed bounded presentations 29
5.2. Choices over the whole marked curve 30
5.3. The two derived words 31
5.4. A proper carrier and its supported class 32
5.5. Lifting the class to the entire arrow 33
5.6. Independence of the auxiliary presentation 35
5.7. Finite-band vanishing 36

Date: November 13, 2026.

2020 *Mathematics Subject Classification.* 14D20, 14C30, 14N35, 14F45.

Key words and phrases. Character varieties, mixed Hodge structures, Higgs bundles, stable pairs, perverse filtration.

6. The positive source identity 37
6.1. The unframed input and its hypotheses 37
6.2. Finite spectral addition and reduced spectral curves 39
6.3. The Hodge calculation at a nodal stratum 41
6.4. Purity and reconstruction of the source 43
7. The scalar critical primitive 44
7.1. The critical scheme on arbitrary ordinary bases 44
7.2. The normal Hessian and its determinant root 45
7.3. Pushforward and the scalar shift 46
8. Scalar packets and the global physical identity 46
8.1. Detection of the comparison error 47
8.2. The physical product over coprime packets 47
8.3. Stratified descent and induction 50
8.4. Restriction to the original nil coefficient 51
9. The fixed loci and their cohomology 52
9.1. The fixed functor 53
9.2. Closed attracting filtrations and punctual purity 55
9.3. Indices and the diagonal hook identity 56
9.4. Moving support and finite symmetric actions 58
9.5. Global purity and proof of the relative formula 59
9.6. Properness of the positive spectral map 60
9.7. The whole nil locus 61
10. Hodge recovery and the character-variety formula 62
10.1. The identity in full Hodge classes 62
10.2. Recovering perverse-graded Hodge structures 63
10.3. The hook expression 63
10.4. Proof of Theorem 1.1 64
10.5. Ranks one and two 65
References 65

1. Introduction

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

$$\mathcal{M}_{g,n} = \left\{ (A_1, B_1, \dots, A_g, B_g) \in \mathrm{GL}_n(\mathbf{C})^{2g} : \prod_{j=1}^{g} [A_j, B_j] = \zeta I_n \right\} /\!\!/ \mathrm{GL}_n(\mathbf{C}). \tag{1.1}$$

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 $2N_n$. Its cohomology is of Hodge–Tate type [16, Corollary 4.1.11]. We write W for Deligne's weight filtration, use ordinary cohomology, and set

$$H(\mathcal{M}_{g,n}; q, t) = \sum_{D,k} \dim \mathrm{Gr}^W_{2k} H^D(\mathcal{M}_{g,n}, \mathbf{Q})\, q^k t^D. \tag{1.2}$$

For a box b of a partition λ, let a_b and l_b be its arm and leg lengths, counting boxes to its right and below, respectively, and put $h_b = a_b + l_b + 1$. With Q a formal rank variable, define

$$\mathcal{H}_\lambda(z,w) = \prod_{b\in\lambda} \frac{(z^{2a_b+1} - w^{2l_b+1})^{2g}}{(z^{2a_b+2} - w^{2l_b})(z^{2a_b} - w^{2l_b+2})}, \tag{1.3}$$

$$\Omega(Q;z,w) = \sum_\lambda Q^{|\lambda|}\mathcal{H}_\lambda(z,w), \tag{1.4}$$

$$\overline{H}_n(z,w) = (z^2-1)(1-w^2)[Q^n]\operatorname{Log}\Omega(Q;z,w). \tag{1.5}$$

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

$$\operatorname{Log} F = \sum_{d\geq 1} \frac{\mu(d)}{d} \log F(Q^d; z^d, w^d). \tag{1.6}$$

Theorem 1.1. *For every $g \geq 2$ and every $n \geq 1$,*

$$H(\mathcal{M}_{g,n}; q, t) = (t\sqrt{q})^{2N_n}\overline{H}_n\left(\sqrt{q}, -\frac{1}{t\sqrt{q}}\right). \tag{1.7}$$

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

Corollary 1.2. *For every $g \geq 0$ and every $n \geq 1$, the mixed Hodge polynomial of $\mathcal{M}_{g,n}$ is given by* (1.7).

Proof. Theorem 1.1 treats $g \geq 2$. If $g = 0$, the empty commutator product is I_n, so $\mathcal{M}_{0,1}$ is a point and $\mathcal{M}_{0,n}$ is empty for $n > 1$. The genus-zero identity [16, Theorem 4.3.1] gives

$$\operatorname{Log}\Omega(Q;z,w) = \frac{Q}{(z^2-1)(1-w^2)}.$$

Thus $\overline{H}_1 = 1$ and $\overline{H}_n = 0$ for $n > 1$. Since $N_1 = 0$, these are exactly the required mixed Hodge polynomials; the zero expression remains zero when $N_n < 0$.

If $g = 1$, one has $\mathcal{M}_{1,n} \cong (\mathbf{C}^*)^2$ [16, Theorem 2.2.17]. The genus-one identity [5, Theorem 1.0.2], specialized at $u = 1$, is

$$\operatorname{Log}\Omega(Q;z,w) = \frac{(z-w)^2 Q}{(z^2-1)(1-w^2)(1-Q)}.$$

Hence $\overline{H}_n(z,w) = (z-w)^2$ for all $n \geq 1$. Here $N_n = 1$, and substitution gives

$$(t\sqrt{q})^2\left(\sqrt{q} + \frac{1}{t\sqrt{q}}\right)^2 = (1+qt)^2 = H((\mathbf{C}^*)^2; q, t).$$

□

1.1. **Background and previous results.** Character varieties and Higgs moduli provide two algebraic structures on the spaces arising from nonabelian Hodge theory. Hitchin introduced the moduli of solutions of the self-duality equations on a Riemann surface and its integrable system in [19, 20]. Simpson developed the correspondence between Higgs bundles and local systems in [50]. For $g \geq 2$, in the primitive-twisted case the Betti space (1.1) is diffeomorphic to the moduli of stable K_C-Higgs bundles of rank n and degree d, where $\gcd(n,d) = 1$ and $\zeta = \exp(2\pi i d/n)$. The Hitchin map sends a Higgs bundle to the coefficients of its characteristic polynomial. It is proper, and its generic fiber is a Picard variety of fixed degree on a smooth spectral curve. Its

perverse Leray filtration provides a second grading on Higgs cohomology. On the Betti side, Deligne's weight filtration records the mixed Hodge structure of the affine character variety.

Hausel and Thaddeus related mirror symmetry for the SL_n and PGL_n Hitchin systems to stringy Hodge invariants [18]. The mixed Hodge partition formula was announced in Hausel's 2004 survey in the equivalent PGL_n normalization [13, Conjecture 5.1]. Hausel and Rodriguez-Villegas developed it in their 2006 preprint, published in 2008 [16, Conjecture 4.2.1]. Their starting point was arithmetic: the character table of $\mathrm{GL}_n(\mathbb{F}_q)$ gives point counts, and Katz's theorem in the appendix converts polynomial counts into compactly supported E-polynomials. The resulting partition sum suggested the bigraded mixed Hodge formula. The two gradings in (1.2) retain both cohomological degree and weight. With compactly supported E-polynomials, Poincaré duality gives

$$E_c(\mathcal{M}_{g,n};q) = q^{2N_n} H(\mathcal{M}_{g,n};q^{-1},-1).$$

Their paper also established the rank-two formula in arbitrary genus and related the predicted symmetries to curious hard Lefschetz and quiver invariants [16, Theorem 1.1.3 and Sections 4–5]. Hausel, Letellier and Rodriguez-Villegas subsequently formulated a Macdonald-polynomial extension for generic punctured character varieties and proved its arithmetic and quiver-theoretic specializations [14].

The small-genus cases have explicit partition identities. For genus zero the primitive-twisted variety is a point in rank one and is empty in higher rank; the corresponding identity is [16, Theorem 4.3.1]. For genus one it is $(\mathbb{C}^*)^2$ in every rank. Carlsson and Rodriguez-Villegas proved the required identity using vertex operators [5, Theorem 1.0.2]; it also follows from the Macdonald-polynomial Nekrasov–Okounkov formula of Rains and Warnaar [45, Theorem 1.3]. In the notation above, $\overline{H}_n(z,w) = (z-w)^2$ for $g = 1$ and every $n \geq 1$. For any genus, the rank-one variety is $(\mathbb{C}^*)^{2g}$ and its mixed Hodge polynomial is $(1+qt)^{2g}$. Theorem 1.1 treats all ranks in genus at least two for the primitive central twist in (1.1).

Several general specializations were proved through arithmetic and symmetric functions. Schiffmann counted indecomposable vector bundles and stable Higgs bundles over finite fields [48]. Mozgovoy and Schiffmann obtained formulas for twisted Higgs bundles and type-A quiver bundles [41]. Mellit proved the predicted Poincaré polynomial for the unpunctured coprime moduli [34], and established the punctured Poincaré formula through Macdonald polynomials and affine Springer fibers [35]. His integrality theorem for the Hausel–Letellier–Villegas kernels controls the coefficients of the symmetric-function series [33]. His toric stratifications establish curious hard Lefschetz for the character varieties [36]. The distinction among the counting, integrality, Poincaré, and full mixed Hodge statements is recorded in [37, Conjecture 4.5 and Theorems 4.6–4.8].

The relation between the two geometric gradings was formulated by de Cataldo, Hausel and Migliorini as the $P = W$ conjecture. They proved it in rank two for the groups GL_2, SL_2, and PGL_2 [8]. For coprime GL_n moduli in arbitrary rank, the theorem was proved by Maulik and Shen [31] and by Hausel, Mellit, Minets and Schiffmann [15]. Maulik, Shen and Yin obtained a further proof through perverse filtrations and Fourier transforms [32]. With the normalized perverse filtration, $P = W$ identifies P_k with $W_{2k} = W_{2k+1}$ under nonabelian Hodge theory. The remaining enumerative task is to calculate the dimensions in each cohomological degree and each perverse degree of the Hitchin system.

The connection with curve counting predates these proofs of $P = W$. Chuang, Diaconescu and Pan expressed the HRV partition sum as a refined Gopakumar–Vafa expansion for a local Calabi–Yau curve and identified its BPS multiplicities with the primitive multiplicities in the relative Lefschetz decomposition of perverse graded Higgs cohomology [7]. Their construction used a proposed refined stable-pair formula together with $P = W$. Chuang, Diaconescu, Donagi

and Pantev extended this approach to parabolic refined invariants and Macdonald polynomials [6]. These works supply the enumerative origin of the local-curve series used here. Monavari described the fixed stable-pair loci by double nested Hilbert schemes and proved their Euler-characteristic hook product [39]. He subsequently computed the K-theoretic local Donaldson–Thomas series in arbitrary genus and degree and proved the local K-theoretic DT/PT correspondence [40]. Here we calculate the cohomological stable-pair series in mixed Hodge modules, retaining the orientation maps, finite symmetric actions, and cohomological grading.

1.2. **Proof of the formula.** We compute cohomological stable-pair invariants on $X_1 = \mathrm{Tot}_C(K_C \oplus \mathcal{O}_C)$ in two ways. A positive quotient $K_C \oplus \mathcal{O}_C \twoheadrightarrow L$ relates X_1 to the split threefold $X_0 = \mathrm{Tot}_C(L \oplus K_C L^{-1})$. Both are described by critical charts on the derived stack of surface pairs on $\mathrm{Tot}_C L$. Section 2 constructs these charts and compares their closed forms and orientations. Sections 4–5 construct the comparison morphism over the positive Hitchin base. Elementary modifications act on this morphism and its cofiber. Their commutator is multiplication by the rank. After passing to the quotient by contributions with at least two smaller positive ranks, the cofiber is supported in a finite interval of Euler characteristics. The commutator forces it to vanish.

Section 6 computes the positive source using the Hitchin support theorem, the smooth-stack integrality theorem, and the relative Hilbert schemes of reduced spectral curves. A critical calculation identifies its primitive contribution with the individual perverse direct-image coefficients of the canonical Hitchin map. Disjoint spectral packets and induction on rank then give the global source identity. Its proof uses the comparison morphism and its orientation constructed in the preceding sections.

Section 9 computes the same stable-pair series by two torus actions. The fixed loci are Quot schemes of the dualizing modules of monomial thickenings of C. The full attracting filtrations give purity in every cohomological degree. Their indices and finite permutation actions express the series through symmetric powers of C and the arms, legs, and hooks of partitions. Throughout these calculations, $H^1(C, \mathbf{Q})$ remains a Hodge structure. In Section 10, purity separates the cohomological degrees in the equality of Hodge classes, and $P = W$ identifies the resulting dimensions with the coefficients of (1.2). Taking the weight character yields (1.7).

1.3. **Notation and coefficients.** For the proof of Theorem 1.1, 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

$$q_0 = (a, b) : \mathcal{O}_C \oplus K_C \twoheadrightarrow L, \qquad A = K_C L^{-1}. \tag{1.8}$$

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

$$X_0 = \mathrm{Tot}_C(L \oplus A), \qquad X_1 = \mathrm{Tot}_C(K_C \oplus \mathcal{O}_C) \tag{1.9}$$

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

$$B_r = \ell r^2, \qquad a_r = -hr - \frac{\ell r(r-1)}{2}, \qquad \alpha_r = \frac{a_r}{r}, \qquad c_r = \alpha_r - \frac{1}{2r}, \qquad N_r = hr^2 + 1. \tag{1.10}$$

The symbol a_r in (1.10) is unrelated to an arm length a_b.

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

$$\mathrm{add} : \prod_\nu \mathcal{A}_{r_\nu}^D \longrightarrow \mathcal{A}_r^D, \qquad \sum_\nu r_\nu = r. \tag{1.11}$$

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 actual 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 [10]. Set

$$\mathbf{T} = \mathbf{Q}(-1), \qquad \mathbb{L} = \mathbf{T}[-2], \qquad \vartheta^2 = \mathbf{T}, \qquad \Lambda = \vartheta[-1]. \tag{1.12}$$

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 9. 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 (1.12). The construction and all orientation maps used in this paper are given in Section 2.

1.4. **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

$$(E, \Theta, s), \qquad \Theta : E \longrightarrow E \otimes L, \quad s : \mathcal{O}_C \longrightarrow E. \tag{1.13}$$

Here E is a vector bundle. The additional fields, the relation variable, and the coherent realization of its framing are described in Section 2. 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

$$g_q(E, \Theta) = \frac{1}{2}\langle \alpha_q, \mathrm{Tr}(\Theta^2)\rangle. \tag{1.14}$$

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 perverse Hodge modules 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 generated by the perverse constituents of finite-addition convolutions of the framed and unframed coefficients occurring in proper flag correspondences, with at least two positive-rank factors, all of rank less

than r. We include their indicated shifts, Tate twists, and source vanishing-cycle images. In particular $\mathscr{I}_1 = 0$. For either category $\mathscr{J}$ let $D_{\mathscr{J}}$ denote the complexes all of whose perverse cohomology objects 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.

1.5. **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 [30, Theorems 0.1 and 0.4]. The smooth-stack integrality theorem of [4, Theorem 10.2.7 and Lemma 10.2.9], with its commutative cotangent orientation, applies to the fixed-slope positive Higgs stack. Section 6 verifies its hypotheses and fixes the scalar stabilizer parameter.

The integral localization theorem of [27, Corollary 7.19] supplies the morphism for filtered oriented 3-Calabi–Yau objects. We use its chart-level monodromic Hodge realization together with the specialization exchange of [12]. The smooth descent and the comparison of orientations are included below. The conductor transversality criterion of [38, 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 [17, Corollary 1.3.2]. The $P = W$ theorem for coprime GL_n Higgs moduli is [31, Theorem 0.2]. The rational tautological classes have the Hodge types stated in [49]; 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.

2. 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.

2.1. **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

$$0 \longrightarrow K_C L^{-1} \longrightarrow N \xrightarrow{q} L \longrightarrow 0. \tag{2.1}$$

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

$$g_q(E,\Theta,s)=\frac{1}{2}\,\mathrm{Tr}_C\Big(\alpha_q\,\mathrm{Tr}(\Theta^2)\Big). \tag{2.2}$$

Here $\mathrm{Tr}_C : H^1(C,K_C)\to\mathbf{C}$ is Serre trace. Formula (2.2) defines a function on the derived stack as well: use a Čech or Thom–Whitney representative of α_q and the universal trace. Its derivative in a bundle and field direction (b,v) is $\mathrm{Tr}_C(\alpha_q\,\mathrm{Tr}(\Theta v))$, since the additional term $\mathrm{Tr}(\Theta[b,\Theta])$ vanishes.

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

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

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

Proposition 2.1. *The boundary-framed physical enhancement is naturally equivalent, on the indicated locus, to* $\mathrm{Crit}(\mathscr{P}_L,g_q)$*. The equivalence is defined over connective Artin differential graded algebras, and respects maps of objects, their homotopies, direct sums, and finite exact diagrams.*

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

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

The projective bundle formulas give

$$R\pi_*H=V^\vee,\qquad R\pi_*H^{-1}=0,\qquad R\pi_*H^{-2}=K_S[-1]. \tag{2.3}$$

The relative diagonal has resolution

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

For a framed perfect object I, define

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

Applying this diagonal resolution and then tensoring by H recovers

$$I\simeq\mathrm{fib}(\pi^*U\longrightarrow\pi^*G\,H). \tag{2.4}$$

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

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

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

$$d\varphi+se=c_G,\qquad de=0. \tag{2.5}$$

Conversely these equations construct the upper triangle, its map to the lower triangle, and the fiber (2.4). Both composites are the unit and counit of the displayed diagonal kernel. This also proves the assertion on mapping spaces and higher homotopies, after every derived base change.

The tangent of $\mathscr{P}_L$ at a pair is $\mathrm{RHom}_S(J_G, G)$. Surface duality with proper support identifies its shifted cotangent with $\mathrm{RHom}_S(GK_S^{-1}, J_G[1])$. Consequently (2.5) describes the derived zero stack of the cotangent section $(c_G, 0)$. In the monic resolution

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

multiplication by y is chain homotopic to multiplication by Θ. The torsor cocycle $\alpha_q y$ therefore represents the covector $\mathrm{Tr}_C(\alpha_q \,\mathrm{Tr}(\Theta\, d\Theta))$, including bundle directions by the trace identity preceding the proposition. It is dg_q. In a local splitting of the quotient, $X : E \to E \otimes A$ is the transverse field, and $\varepsilon : E \to K_C$ is the dual framing field. Write ψ for the cotangent field dual to Θ, and $\rho_{\mathrm{cot}}, \rho_{\mathrm{phys}}$ for the relation generators in the cotangent and physical models. The cone convention in (2.5) gives

$$\psi = -X, \qquad e = -\varepsilon, \qquad \rho_{\mathrm{cot}} = -\rho_{\mathrm{phys}}. \tag{2.6}$$

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

For the split bundle $L \oplus K_C L^{-1}$, the function is zero. For $N = K_C \oplus \mathcal{O}_C$ and the fixed positive quotient q_0, it is the function (2.2). These two uses have the same cone convention and the same dictionary (2.6).

2.2. **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

$$\begin{aligned} \alpha &= (z_0 X - z_2,\ -(z_0\Theta - z_1),\ z_0^2 \varepsilon), \\ \beta &= (z_0\Theta - z_1,\ z_0 X - z_2,\ s), \qquad \gamma = -z_0^2 \rho : M^{-1} \longrightarrow M^1. \end{aligned} \tag{2.7}$$

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

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

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

The source and target types are $g, a, b, w : V_i \to V_j$, $r : R_i \to R_j$, $u : R_i \to V_j$, and $v : V_i \to R_j$, with the normal twists in (2.7). Its differential is

$$\begin{aligned} D(g, r) &= (g\Theta_i - \Theta_j g,\ gX_i - X_j g,\ gs_i - s_j r,\ r\varepsilon_i - \varepsilon_j g) - g\rho_i + \rho_j g, \\ D(a, b, u, v) &= aX_i - X_j a + \Theta_j b - b\Theta_i + u\varepsilon_i + s_j v, \qquad Dw = 0. \end{aligned} \tag{2.8}$$

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

The map to monad morphisms sends (g,r) to $\mathrm{diag}(g,g,g,r,g)$, a field to minus the derivative of (2.7), and w to the lower-to-top entry $-z_0^2w$. Denote it by κ. Its field entries, in the order lower, middle, frame, top, are

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

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

$$(a_2,b_2,u_2,v_2)(a_1,b_1,u_1,v_1)=w(-a_2b_1+b_2a_1-u_2v_1). \tag{2.9}$$

Degree-zero products and their actions are ordinary rectangular composition. Equations (2.8) and (2.9) follow by multiplication of the displayed matrices. Thus κ is a morphism of differential graded categories.

Lemma 2.2. *The morphism κ is a quasi-isomorphism on every universal chart and after arbitrary derived base change. The relative boundary-framing complex is its subcomplex $r=0$.*

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

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

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

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

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

Lemma 2.3. *The proper relative Serre trace on $\mathcal{C}$ is the cochain pairing*

$$q_{\mathcal{C}}(g,w)=q_{\mathcal{C}}(w,g)=\mathrm{Tr}(gw),$$
$$q_{\mathcal{C}}((a,b,u,v),(a',b',u',v'))=-\mathrm{Tr}(ab'-ba')-\mathrm{Tr}(v'u)+\mathrm{Tr}(vu'). \tag{2.10}$$

The last two traces are taken on the respective framing spaces. After the shift $T=\mathcal{C}[1]$, it agrees under (2.6) *with the canonical critical pairing.*

Proof. Let π be the projection from the relative projective plane to its curve-and-bundle chart, and let π_D be its boundary restriction. Put $G_j=R\pi_*\,\mathrm{RHom}(M,M(-jD))$. The second-order

framing gives $G_2 \simeq G_1$, because its boundary cone is $\mathrm{End}(R) \otimes R\pi_{D*}\mathcal{O}_D(-D) = 0$. The pairing in question is the specific product

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

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

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

The contractions away from the point are

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

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

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

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

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

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 $\mathrm{str}(c_u v_1) = -1/(z_0 xy)$. For the gauge class the top component of its lift is $c_g = -h_x h_y / z_0^2$; 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 [42, Proposition 5.2.3]. Tensoring these calculations with arbitrary rectangular Hom spaces gives exactly (2.10).

There is also a direct differential check. If P is minus the field pairing in (2.10), 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 (2.6) therefore matches the gauge–relation block as well. □

2.3. **Closed forms and the determinant square.** We specify the closed lift of the trace just calculated. In the convention of [44, Theorem 2.12], the leading component of Ch_2 is $\frac{1}{2}\operatorname{Tr}(\mathrm{mult})$. We therefore use

$$\omega_{\mathrm{Perf}}^{\mathrm{Tr}} = 2\mathrm{Ch}_2. \tag{2.11}$$

On Y, the given volume has $K_Y = \mathcal{O}_Y(-3D)$. The framed mapping construction of [51, Theorem 3.3], with extra divisor D, applies: the lift of the specified framing from D to $2D$ is étale because $R\Gamma(D, \mathcal{O}_D(-D))[1] = 0$. Its framing isotropic path is zero, since the ideals of D and $2D$ multiply to zero on $3D$. The other path is the boundary triangle $K_Y \to \mathcal{O}_Y \to \mathcal{O}_{3D}$. This constructs a closed (-1)-shifted form whose leading pairing is precisely the proper product in Lemma 2.3. The factor in (2.11) does not change the coefficient $1/2$ of (2.2): the latter is its quadratic primitive coefficient.

Let $\mathrm{pr} : \mathcal{C}_{ij} \to \mathcal{B}_{ij}$ be the projection to the surface framing complex, forgetting b, v, w. It is a strict functor. Its kernel has b, v in degree one and w in degree two, and has square zero by (2.9). The pairing

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

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

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

Here k lies in the kernel and h, b are homogeneous composable entries. In generator degree two they are rectangular trace identities, using (2.9); all other cases vanish. Thus, after curve sections and Serre duality, there is a natural paired triangle

$$B_{ji}^{\vee}[-3] \longrightarrow C_{ij} \longrightarrow B_{ij}, \qquad B_{ij} = R\Gamma(C, \mathcal{B}_{ij}), \quad C_{ij} = R\Gamma(C, \mathcal{C}_{ij}). \tag{2.12}$$

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

The leading pairing comparison of Lemma 2.3 extends to the full closed forms. Both forms have weight two for normal dilation. If δ is their difference and h_0 is its leading homotopy, then $\delta - Dh_0 \in F^3\mathrm{DR}$. For the Euler field $\mathcal{E}$, Cartan's formula gives

$$H_0 = h_0 + \tfrac{1}{2}\iota_{\mathcal{E}}(\delta - Dh_0) \in F^2\mathrm{DR}, \qquad DH_0 = \delta. \tag{2.13}$$

This uses the full de Rham mixed complex. In a formal presentation with invariant ideal J, the filtration $\prod_p J^{\max(m-p,0)}\Omega^p$ makes the operation continuous. Construct (2.13) on the equivariant formal diagram, with its transition maps, before restricting to a completion at a moving point.

For a diagonal object put $B = B_{ii}$. The half line is $\det B$; its square is given by (2.12), duality, and the canonical critical factor $(-1)^{\mathrm{vdim}\,\mathscr{P}_L}$ in [27, Example 3.20]. This defines the orientation used throughout the paper. The bimodule structure and the determinant double-sequence identity [27, Lemma 2.4] give its direct-sum and finite-flag compatibilities.

Proposition 2.4. *Let Ψ be the equivalence of Proposition 2.1, or one of the supported equivalences constructed below, and let M_A, M_B be the orientation lines of its source and target. Its specified half-line comparison $\gamma : M_A \to \Psi^* M_B$ satisfies*

$$o_A = \det(d\Psi^*)\, o_B\, \gamma^{\otimes 2}, \qquad d\Psi^* : \Psi^*\mathbb{L}_B \longrightarrow \mathbb{L}_A. \tag{2.14}$$

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

Proof. First consider (2.1). Write its paired triangle as $K \to C \to B$, with $h : K \xrightarrow{\sim} B^{\vee}[-3]$. The cotangent triangle is

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

Identify its last object with $B[2]$ by the inverse of $h^\vee[-1]$, using the same shifted double-dual convention as in [27, 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 2.3, is

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

Consequently these are equivalent diagrams of fiber sequences. Taking their graded determinants, and inserting the same critical factor on both sides, proves (2.14). This keeps the direction of the cotangent derivative. For example, scaling h by c changes the raw square by $c^{\chi(B)}$, and negating the whole dual half changes its derivative determinant by $(-1)^{\chi(B)}$; these are the same factor. Dualizing and shifting inverts the full determinant map.

For a finite exact diagram, apply this proof to the paired bimodule triangles (2.12) on every entry of its double Hom grid. Naturality of duality and determinant double-sequence coherence give the asserted square for the resulting half-line map. The supported case and the identification of its full map with the derivative of the specified Ψ are spelled out in Lemma 2.7. □

2.4. **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 2.5. *The two orientation roots differ by $\mathcal{C}_{\mathrm{res}}^{r^2}$, where $\mathcal{C}_{\mathrm{res}}$ is one fixed quadratic line depending only on the two quotients. Choosing its quadratic trivialization once gives a square-preserving comparison, compatible with direct sums and flags. Its factors cancel canonically between a finite modification and its same-rank baseline.*

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

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

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

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

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 $\mathrm{End}(E)$ 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. □

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

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

Transport the fixed positive-quotient orientation by Lemma 2.5. Both quotient degrees are even, and their kernels are generically transverse: otherwise the component $\mathcal{O}_C \to L$ of the positive quotient would vanish, forcing a surjection $K_C \to L$ contrary to $\deg L > \deg K_C$. The fixed quadratic factor has exponent r^2 and cancels between a modification and its same-rank baseline. All Hom, section, support, and duality maps below are equivariant for the torus H defined next. Their determinant grids are taken weight by weight: weight zero supplies the orientation of the fixed enhancement, and the moving weights supply the localization index.

Let $Z_\lambda \subset \mathrm{Tot}_C(K_C \oplus \mathcal{O}_C)$ 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 $\omega_{Z_\lambda} = \mathrm{Hom}_C(\mathcal{O}_{Z_\lambda}, K_C)$.

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

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

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

Lemma 2.6. *On the selected support branch there is a formal equivalence of the specified H-fixed physical enhancements whose truncation is BA^{-1} on every ordinary Artin base. It is constructed on all connective Artin bases by restriction and gluing with fixed determinant.*

Proof. For a rank-one perfect object define

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

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

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

$$\big(\mathrm{Perf}(Y) \times_{\mathrm{Pic}(Y)} \{\det I_0\}\big) \underset{\mathrm{Perf}(U) \times_{\mathrm{Pic}(U)} \{\det I_0|_U\}}{\times} \{I_0|_U\}. \tag{2.15}$$

It has deformation complex $R\Gamma_Z \mathscr{A}_I$. An étale map on the selected finite support is universally injective and proper there, hence a closed immersion onto that support. Delete the other sheets over it. This gives an elementary distinguished square with the open complement. Descent for Perf and Pic gives an equivalence of (2.15) on the two sides. Its inverse is the Mayer–Vietoris homotopy fiber of the difference of restriction maps; the same fiber on mapping complexes gives full faithfulness and the higher cells [52, Tag 08GS]. This construction uses the open complement in the derived ambient.

For the monomial baseline, its transverse ideal has scalar endomorphisms, local Ext^1 equal to the Hilbert tangent representation, and no higher local Ext, by the two-variable Hilbert–Burch resolution. The two tangent characters at a box are $a^{h_\flat}$ and $a^{-h_\flat}$; 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 2.7. *For the equivalence in Lemma 2.6, the algebraic half-line comparison defined by the four supported Hom and section maps has square* (2.14) *for the derivative of that same equivalence.*

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

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

Superscripts $H, +, -$ denote the invariant, positive-weight, and negative-weight parts for H. The map is precomposition by the section. For the paired triangle (2.12), the positive-weight part of the physical cotangent complex has determinant

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

The localization orientation $M(G) \otimes \det(\mathbb{L}^{>0}_{\mathrm{phys}})^{-1}$ therefore identifies with $M_H(G)$, using the graded regrouping maps of [27, Section 6.5]. The double Hom grid is H-equivariant before taking determinants. Its full determinant identity is

$$\frac{M(G)}{M(G_0)} = \det \mathrm{RHom}(T,G_0) \otimes \det \mathrm{RHom}(T,T) \otimes \det \mathrm{RHom}(G_0,T) \otimes (\det R\Gamma T)^{-1}. \tag{2.16}$$

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

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

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

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

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

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

$$R\mathcal{H}om_Y(J, I_0), \qquad \mathscr{A}_{I_0} \oplus R\mathcal{H}om_Y(J, J), \qquad R\mathcal{H}om_Y(I_0, J). \tag{2.17}$$

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

After supported sections and H-invariants, the baseline term in (2.17) is acyclic and the other terms are perfect with finite support, so their determinants are defined. Restriction and gluing preserve the displayed filtration. Consequently its determinant map is the determinant of the derivative of the functor in Lemma 2.6. To compare with the proper Hom grid, use the natural maps from supported to proper sections on its surviving finite-support terms. The possible trace-projector correction on $\mathrm{RHom}(J, J)$ is $1 - uv$, where $u(1) = \mathrm{id}_J$ and $v = \operatorname{str}_J$. Since $vu = \mathrm{rk}(J) = 0$, the identity $\det(1 - uv) = \det(1 - vu)$ gives determinant one, even on nonreduced bases.

The paired triangle (2.12), applied to the same exact arrow diagram, has surface halves

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

Its full half determinant is (2.16); its weight-zero half determinant is the same formula with invariants in every Hom and section term. Applying the proof of Proposition 2.4 to this grid identifies the square of γ with the derivative determinant just computed. The three full surface Hom Euler pairings are zero, and the section term has Euler characteristic $d = \operatorname{length} T$. 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 2.5. □

2.6. **The comparison over the full parameter ring.** The preceding construction also compares the closed forms. On full endomorphism complexes, stable Koszul local cohomology gives the supported product. The boundary section is invertible on the support. Denote the maps from supported sections to $R\Gamma\,\mathrm{RHom}(I, I(-jD))$ by J_j. They satisfy $J_i(a)J_j(b) = J_{i+j}(ab)$. Restricting the inputs to the trace kernel only after taking these products gives the proper pairing of Lemma 2.3. The selected étale maps preserve the volume. Their canonical bundle maps and duality counits therefore compare this pairing on the full moving family, including directions that move the support. The residue change-of-generators, localization, and base-change identities in [42, §§4,10] apply to these counits. Locally a monic characteristic polynomial in the curve coordinate, together with high powers of the two normal coordinates, gives a bounded flat Koszul model. Lifting its coefficients to a connective presentation gives the same construction after derived base change. Different lifts compute the same support functor and its already defined duality counit.

Normal dilation gives both closed forms weight two. Equation (2.13) consequently constructs their full closed homotopy on the common parameter diagram. The following argument upgrades this statement to the particular ordinary maps used in coefficient descent.

Proposition 2.8. *On a neighborhood of the punctual locus, the maps* A, B *identify the algebraic d-critical sections and the oriented vanishing-cycle coefficients. The comparison holds over the entire smooth moving-support and coordinate parameter space.*

Proof. The ordinary map A is étale. Its unique derived étale lift relative to the specified physical enhancement is identified with the lift in Lemma 2.6: their ordinary maps agree and their lifting spaces agree on all Artin bases. On a selected completion this makes the A-side the fixed physical formal problem extended over $R = \mathbf{C}[[s_1, \dots, s_e]]$. Choose one minimal critical model at its reference point and extend scalars:

$$\mathcal{A}_R = R[[x_1, \dots, x_m]][e_1, \dots, e_m], \quad |e_j| = -1, \quad de_j = I_j(x) \in (x)^2. \tag{2.18}$$

Put $I = (I_1, \dots, I_m) \subset R[[x_1, \dots, x_m]]$. The existence of such a model is the formal critical-chart statement of [3]; it is chosen once over $\mathbf{C}$.

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

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

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

Give $x, e, d_{\mathrm{DR}}x, d_{\mathrm{DR}}e$ Euler degree one and R degree zero. On positive Euler degree v, $\iota_{\mathcal{E}}/v$ contracts the augmented de Rham differential. The internal differential raises that degree, by (2.18). The perturbation series

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

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

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

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

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

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

$$F = -\sum_{i,j} M_{ij} I_i I_j \in I^2 \quad \text{in } R[[x]]. \tag{2.19}$$

This is a single identity over the full parameter ring.

For original charts with acyclic Hessian variables y, put $J = (I, y)$. A class modulo J^2 has the form $f_0(x) + \sum y_a f_a(x)$. If its differential vanishes modulo J, its dy_a coefficients give $f_a \in I$. Thus (2.19) detects equality in the original charts, with their Hessian factors retained. The product parameter directions contribute the identical factor K_{par} to both half lines. The maps are over the parameter base, so their derivative matrices have an identity parameter block and introduce no additional determinant.

The two d-critical sections and the half-line map of Lemma 2.7 existed algebraically before completion. Faithful flatness of local completion detects equality of their sections of $\mathcal{O}/J^2$, and then equality of the two existing square maps on the reduced base. The identification of the critical cotangent determinant with the d-critical canonical line is [27, §§3.5–3.6], with the same Hessian and critical factor used above. Equation (2.14) is therefore the oriented d-critical comparison, proving the assertion about vanishing-cycle coefficients. □

3. 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 [10, §2.1].

3.1. **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

$$J_V F = \mathrm{cofib}(\pi^* R\pi_* F \to F) \tag{3.1}$$

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

$$J_V F \simeq R(\mathrm{id}_V \times +)_*(F \boxtimes j_!\mathbf{Q}_{\mathbf{G}_m}[1]). \tag{3.2}$$

The first map of the resulting triangle is exactly the counit in (3.1). Equation (3.2) is the perverse-exact projector of [10, §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 [57]. The nilpotent and semisimple monodromy operators remain in this construction, as do the Hodge and weight filtrations.

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

$$\mathrm{Filt}(X) = \mathrm{Map}([\mathbf{A}^1/\mathbf{G}_m], X), \qquad \mathrm{Grad}(X) = \mathrm{Map}(B\mathbf{G}_m, X). \tag{3.3}$$

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

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

Distinguish the functors

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

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

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

as in [27, Definition 7.21]. For the surface-half model (2.12), it is $\chi(B^-) - \chi(B^+)$.

Proposition 3.1. *For the oriented (-1)-shifted symplectic stacks used here, the integral map is an isomorphism in monodromic mixed Hodge modules*

$$\zeta_X : \Phi_{\mathrm{Grad}(X)}\Lambda^{-I_\mu} \xrightarrow{\sim} R\,\mathrm{gr}_*\,\mathrm{ev}^!\,\Phi_X. \tag{3.4}$$

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

$$(p^+)_!(\eta^+)^*\Phi_X \simeq \Phi_{X^\mu}\Lambda^{I_\mu}. \tag{3.5}$$

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 (2.12) and Proposition 2.4. These are the stack hypotheses of [27, 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 [12, Theorem 3.7, Proposition 6.2]. Its staircase inverse follows from the adjunction triangle identities. Tensor it with the finite-torsor map of [27, (6.75)]. The normalization on a moving hyperbolic quadratic block is the compact Thom map of the same repelling axis as in [27, Theorem 7.9]. Switching that axis would introduce $(-1)^r$ in rank r. The scalar is checked on the normalized quadratic point, where the constructed maps lie in the perverse heart. All other maps are transformations in the Hodge six-functor system.

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

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

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

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

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

The indices obey $I_U = I_X + d_+ - d_-$, so both the shift and Tate factor of the local map descend. The common charts and their fixed overlaps are [27, Proposition 5.27, Corollary 5.28]. On every level p_m of their smooth nerve, the two ordinary pullbacks are the same invertible shift and twist of perverse heart objects. Hence their mapping spaces have $\pi_r = \mathrm{Hom}(A, B[-r]) = 0$ for $r > 0$. The smooth comparison squares, which are purity and Beck–Chevalley squares, give equal maps on double overlaps. Their compositions give the higher equalities. The discrete mapping spaces then descend the isomorphism and its inverse through the full nerve. Degeneracy compatibility is obtained by pullback of these identities, without applying smooth purity to a nonsmooth diagonal.

For two ordered cuts, write their block weights as a_i, b_i in the same order. Then $(a_i - a_j)(b_i - b_j) \geq 0$ on every cross-Hom block. This checks matching on the full cotangent complex, including relation and obstruction terms. Finitely many ordered cuts admit a common multi-torus chart by the construction of [27, 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 [27, Lemma

7.4]. Thus the product and matching-refinement constructions in [27, Propositions 7.10–7.12, Lemmas 7.14–7.15] are the same Hodge transformations, with determinant associators from (2.12). This proves (3.4). Using the inverse action changes the index sign and gives (3.5). □

Lemma 3.2. *Suppose a filtered correspondence $G \xleftarrow{q} F \xrightarrow{p} U$ occurring in* (3.4) *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 actual counit. The closed/open localization triangle proves the cofiber assertion. No properness of h_U is needed. □

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

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

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

4. 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 2.

4.1. **The framed heart and its numerical bounds.** Let $X = \mathrm{Tot}_C(N)$ be either threefold in (1.9), 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

$$\mathscr{A}_X = \langle \mathcal{O}_Y, \mathrm{Coh}_{c,\leq 1}(X)[-1] \rangle_{\mathrm{ex}}. \tag{4.1}$$

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 $\mu(J) \leq \mu(Q)$ 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 [55, 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 [55, Proposition 2.12].

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

$$R\pi_*\mathcal{O}_Y(-D_\infty) = 0. \tag{4.2}$$

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

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

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 $s \leq e-1$, $\sum r_j = e$, and

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

Every line summand of $\mathrm{Sym}^j(N^\vee)$ has degree at least $-j\ell$ in either row. A rank-r_j quotient has degree at least $-j\ell r_j$, and saturation increases degree. Therefore

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

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

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

$$0 \longrightarrow \pi_*\mathcal{O}_Z \longrightarrow E_I \longrightarrow \pi_*Q \longrightarrow 0. \tag{4.3}$$

It follows by applying $R\pi_*(-D_\infty)[1]$ to the truncation triangle and using (4.2). 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 (4.3) now gives the lower bound $\min(\alpha_e, t) = \alpha_e$ when $t > 0$. □

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

4.2. **Complete coherent flags.** The diagonal resolution of Section 2 realizes framing rank at most one by coherent diagrams on S with row

$$0 \longrightarrow \mathcal{O}_S \longrightarrow U \longrightarrow FK_S^{-1} \longrightarrow 0, \tag{4.4}$$

where F denotes the surface projection of the physical sheaf. The diagrams also include the compatible global field arrows. The nonsplit torsor changes the extra arrow but not this exact row. Exact sequences and flat families are termwise exact; in the split case this is also the construction of [56, Theorem 4.5.7 and Corollary 4.5.9].

Lemma 4.2. *For fixed graded Hilbert polynomials and a fixed position of the framing, compatible complete coherent flags in* (4.4) *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 4.3. *For fixed r, χ, c, a finite list of coherent receivers suffices for the Harder–Narasimhan decomposition in a sufficiently large positive chamber. A single chamber parameter can be chosen after the complete receivers, including their torsion quotients, have been formed.*

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

$$\mu(E') \le \frac{\chi - \alpha(r-f)}{f}, \qquad b = \max_{1\le f\le r} \frac{\chi - \alpha(r-f)}{f}. \tag{4.5}$$

For $t > \max(b, 0)$ the first factor is framed; an earlier unframed factor would violate (4.5). Write its class as $(1, e, \chi_0)$ and the successive unframed classes as $(0, r_j, \eta_j)$. They satisfy

$$a_e \le \chi_0 \le \chi - \alpha(r-e), \qquad \eta_j > cr_j, \qquad \sum_j r_j = r - e, \qquad \chi_0 + \sum_j \eta_j = \chi. \tag{4.6}$$

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 (4.4) form finite-type families, being represented by bounded perfect relative Hom complexes. Form the projective receivers of Lemma 4.2 for the finite list (4.6).

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_{fl}.

Choose a generic T above b, b_{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 (4.3) and Lemma 4.1. The remaining $H^1(I)$ is zero-dimensional, and $\mathrm{Ext}^2(Q, \mathcal{O}_X) = 0$ on the smooth threefold lifts the truncation triangle to an ordinary pair. Absence of zero-dimensional subobjects makes its sheaf pure. Conversely the pure subobject slopes of a bounded stable-pair family are uniformly bounded. Thus T is a stable-pair chamber simultaneously for every receiver under consideration. □

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

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

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

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

$$i_\tau : S_\tau \hookrightarrow U_\tau. \tag{4.7}$$

This also gives the unique relative flag over nonreduced bases.

Use U_τ also for the derived open of the physical window with this classical truncation. Choose ordered integral weights for this flag type, with weight zero on its unique framed graded piece. Let G_τ be the open product of its derived good graded-factor stacks. Using (3.3), define

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

The framing at infinity stays fixed in this mapping-stack problem. Proposition 2.1, applied to universal finite exact diagrams, identifies its truncation with the coherent flag rows and compatible fields constructed above. Every extension of the good factors lies in the cutoff window by Lemma 4.1. Its decreasing semistable flag is its Harder–Narasimhan flag, so it lies in the current open U_τ. Conversely a filtered object with these good graded terms has that coherent flag. These are open conditions on the common exact-diagram problem, determined by geometric fibers; hence the identification holds over derived test bases. The Rees construction realizes the chosen weights, and its classical evaluation is $t_0(p) = i_\tau$ in (4.7).

Let K_{G_τ} denote the product DT coefficient with the moving weight normalization in integral localization. The oriented localization map of Proposition 3.1, equivalently [27, Corollary 7.19] with its monodromic Hodge construction, is

$$\zeta : K_{G_\tau} \xrightarrow{\sim} Rq_* p^! \Phi_{U_\tau}. \tag{4.8}$$

The good product orientation is the one obtained from the same physical orientation; its moving determinant and all refinement maps are those of [27, Proposition 8.34].

Lemma 4.4. *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*

$$Rh_{U_\tau,*} i_{\tau,*} i_\tau^! \Phi_{U_\tau} \longrightarrow Rh_{U_\tau,*} \Phi_{U_\tau}. \tag{4.9}$$

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 (4.9). The localization triangle for the complementary open j has cofiber $Rh_{U_\tau,*} Rj_* j^* \Phi_{U_\tau}$. Products, refinements, and restrictions use these same units and counits, so their comparison cells are retained. □

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

4.4. **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 4.5. *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*

$$\rho_\chi(c_1, c_2) : W^\times_{r,\chi}(c_1) \longrightarrow W^\times_{r,\chi}(c_2), \qquad c_1 < c_2 \le c_r. \tag{4.10}$$

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

Proof. First use the single sufficiently large T chosen in Lemma 4.3. 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 $a - b = -\chi$, so the vacuum costalk is the unframed coefficient with factor $\mathbb{L}^{\chi/2}$.

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

$$\operatorname{cofib}(W^\times_{r,\chi}(c) \to R_{r,\chi}) \in D_{\mathscr{I}_r}. \tag{4.11}$$

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

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

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

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

Thus torsion is absent and all slopes equal λ.

The framed factor is wall-semistable as well. In $0 \to J \to I \to Q[-1] \to 0$, a high-slope subobject of J is one of I, and a low-slope unframed quotient of J gives, after extension by Q, a low-slope quotient of I. For the opposite order, a high-slope subobject of J lifts to an extension by Q of high slope, and a low-slope quotient is already a quotient of I. Conversely, intersections and images of an arbitrary subobject in these diagrams show, by weighted averages, that every extension of good wall factors is wall-semistable. An adjacent unstable framed factor splits off a positive-rank factor of slope λ and has smaller framed rank. Ordering by that rank, the proper receivers and Lemma 4.4 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 (4.11) through the finitely many walls between t and T.

Two cutoffs $c_1 < c_2 \le c_r$ contain the same stable-pair open, by Lemma 4.1. The restrictions have their ordinary composite law. An octahedral triangle applied to their common stable-pair restriction proves (4.10). For the full source, the rank-r unframed term is allowed in $\mathscr{S}_r$; every other term is in that category by smaller-rank induction and finite addition. The same proof gives its stated restriction property. □

For later induction, note that

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

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

4.5. **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 (5.16) and (5.17). The comparison arrow is

$$\Delta_\chi(c) : \phi_{g_{\mathcal{A}}} Rh_* \Omega_{\mathscr{P}} \xrightarrow{e_h} Rh_* \phi_g \Omega_{\mathscr{P}} \xrightarrow{Rh_* \delta_{\mathscr{P}}} R(h\pi_g)_* \Phi_{\mathrm{Crit}(\mathscr{P},g)}. \tag{4.12}$$

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

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

All restriction maps in Proposition 4.5 act on this same arrow and its cofiber.

Lemma 4.6 (The split framed reflection)**.** *Let $\mathscr{X}_{r,\chi}^t$ be the full t-semistable framed stack on $X_0 = \mathrm{Tot}_S(K_S)$, with its specified boundary framing. The zero section is retained. For finite t, derived duality induces an isomorphism over $\mathcal{A}_r^L$*

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

It identifies the oriented Donaldson–Thomas coefficients of Section 2. Write $h_{r,\chi}^t : \mathscr{X}_{r,\chi}^t \to \mathcal{A}_r^L$ for the Hitchin map. If

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

then the normalized identity is

$$\mathcal{B}_{r,\chi}^t \simeq \mathcal{B}_{r,-\chi}^{-t} \otimes \mathbb{L}^\chi. \tag{4.13}$$

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

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

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

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

$$\mathfrak{P}_{r,\chi}^\dagger = \mathrm{Spec}_{\mathfrak{M}_{r,\chi}} \mathrm{Sym}(\mathcal{E}), \qquad \mathfrak{P}_{r,\chi}^\sharp = \mathrm{Spec}_{\mathfrak{M}_{r,\chi}} \mathrm{Sym}(\mathcal{E}^\vee[1]). \tag{4.14}$$

Their test-base descriptions are, respectively,

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

These descriptions include the mapping spaces. Relative duality for the proper support supplies the second description and commutes with derived base change. One may form these stacks on the projective surface $\overline{S} = \mathbf{P}_C(\mathcal{O}_C \oplus L)$ and then restrict to sheaves whose support misses its infinity divisor. Thus they are the corresponding opens of the pair and dual-pair stacks of [56, §4.4.1, Lemmas 4.4.1–4.4.2].

Write

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

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

$$Rp_* D_S \mathcal{F} \simeq \mathcal{C}^\vee[-1]. \tag{4.15}$$

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

The remaining step exchanges the two cotangent presentations over the same $\mathfrak{M}_{r,\chi}$. The physical pure stack is the classical truncation of $T^*[-1]\mathfrak{P}^\dagger_{r,\chi}$ by Proposition 2.1. The same coherent diagram has a second projection to $\mathfrak{P}^\sharp_{r,\chi}$: it retains the extension class of

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

Over this dual pair, its other datum is the map $U \to F$. The tangent triangle and surface duality identify the latter datum with the shifted cotangent fiber. Thus both projections describe the same classical physical moduli functor, including families over nonreduced bases. This is also the construction of [56, Proposition 4.4.3 and §4.5.6].

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

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

The resulting identification of classical truncations is the natural one of [56, Lemma 4.6.13]; it is independent of the displayed presentation. It also identifies the d-critical sections, because both sections are represented by this same function w.

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

$$\mathscr{K}_{\mathfrak{M}} \otimes \det \mathcal{E}. \tag{4.16}$$

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

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

for the two presentations. If $a = \mathrm{rk}\, \mathcal{E}^{-1}$ and $b = \mathrm{rk}\, \mathcal{E}^0$, their virtual dimensions are $\mathrm{vdim}\, \mathfrak{M} + (b-a)$ and $\mathrm{vdim}\, \mathfrak{M} - (b-a)$. Hence their canonical critical square signs agree. The determinant comparison in (4.16) therefore preserves the orientation squares. An acyclic refinement adds the same equation-variable hyperbolic block on both sides; determinant contraction gives the same root map and square map. Smooth atlas changes contribute the same atlas determinant on both sides. These comparisons descend to the global determinant isomorphism. The cotangent lift of D_S carries the canonical cotangent orientation by its cotangent derivative. Together with Proposition 2.4, this identifies the specified physical orientations. Functoriality of the oriented vanishing-cycle construction now identifies the coefficients.

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

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

With the usual shifted-dual signs the preceding reflection is

$$(E, \Theta, X, s, \varepsilon) \longmapsto (E^\vee K_C, \Theta^{\mathsf{t}}, X^{\mathsf{t}}, \varepsilon^{\mathsf{t}}, -s^{\mathsf{t}}). \tag{4.17}$$

The transformed moment expression is the negative transpose of the old one. Its square is identified with the identity by the gauge transformation $-\operatorname{id}_E$. Equivalently this is the duality $R\mathcal{H}om_Y(I, \mathcal{O}_Y)$ with its induced boundary framing; the coherent-diagram comparison is [56, Lemma 4.5.14].

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

Reflection exchanges these pure exact sequences, reversing arrows, and sends their unframed Euler slopes μ to $-\mu$. Thus the subobject inequality $\mu \le t$ becomes the quotient inequality $-t \le -\mu$, and conversely. This proves the stability equivalence for every finite t. Since stability defines open substacks, the all-base isomorphism above restricts to the asserted isomorphism of chamber moduli functors. The same classical chamber duality is stated in [56, Lemma 4.4.6].

Since Θ is transposed without a sign, reflection is over the identity of $\mathcal{A}_r^L$. Ordinary direct image carries the coefficient isomorphism to an isomorphism over this base. Finally $\Lambda^\chi = \Lambda^{-\chi} \otimes \mathbb{L}^\chi$, which proves (4.13).

It remains to justify the cutoff assertion. The coherent projection sequence (4.3) applies to every finite-t semistable object. Its cyclic term has minimum Euler slope at least α_r. Every positive-rank HN quotient of the other term has slope at least t. For $t = \pm\epsilon$ and $0 < \epsilon < 1/r$, a finite such slope with denominator at most r is nonnegative whenever it is at least $-\epsilon$. Purity, the cyclic bound, and extension closure therefore give $\mu_{\min} \ge \alpha_r > c_r$. This proves containment for both complete intrinsic chamber stacks. □

Lemma 4.7. *The core surface window is empty for $\chi < a_r$. There is a constant M_r, independent of perverse degree, such that $K_\chi(c_r) \in D_{\mathscr{I}_r}$ for $\chi > rM_r$. For the ordinary source modulo $\mathscr{S}_r$, the upper vanishing endpoint is instead $\chi > -a_r$.*

Proof. Since ℓ is even, α_r is an integer. A rational slope of denominator at most r below α_r is at most $\alpha_r - 1/r$. Thus $\mu_{\min} > c_r$ is equivalent to $\mu_{\min} \ge \alpha_r$. If $\chi < a_r$, the average slope is at most $\alpha_r - 1/r < c_r$, so the whole core is empty. The same lower bound applies to both vertices of (4.12).

A rank-r semistable L-Higgs bundle has

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

Consequently $H^1(C, E) = 0$ when $\chi/r > h + (r-1)\ell$. Increase this bound to the projective framing threshold. Consider a semistable physical object in the immediate chamber $t = \chi/r + \epsilon$, $0 < \epsilon < 1/r^2$. If its projected Higgs bundle E were unstable, its highest Higgs Harder–Narasimhan factor H would satisfy $\mu(H) > \chi/r$. The denominators are at most r, so $\mu(H) > t$. The preceding bound, applied to the semistable bundle H, gives $H^1(C, H) = 0$. Serre duality then gives $\varepsilon|_H = 0$ for the dual framing $\varepsilon : E \to K_C$. The moment relation makes the off-diagonal transverse field $H \to (E/H)A$ a Higgs morphism. In the nonsplit row the local transverse lifts differ by a multiple of Θ, so their off-diagonal blocks agree. Negativity of A and the decreasing Higgs slopes force this morphism to vanish. Thus H is an unframed physical subobject of slope greater than t, a contradiction. Consequently E is semistable, $H^1(C, E) = 0$, and $\varepsilon = 0$. If $s = 0$, the vacuum subobject and its full unframed quotient would give $t \le \chi/r$, a contradiction. Thus $s \ne 0$. The same negative-A Hom vanishing preserves every equal-slope subobject. The remaining framed stability condition is the Joyce–Song condition: the nonzero section is contained in no proper

Higgs subbundle of the same slope. Denote this pair space by $\mathscr{P}^{\mathrm{JS}}$. The ordinary pair stack is therefore the smooth stable framed open, projective over the positive semistable Hitchin moduli; the projective framing construction is described in [43, Definition 2.11 and (2.37)]. The relative framing map is smooth because $H^1(E) = 0$, and the positive Higgs stack is smooth. The framed stability kills stabilizers. Thus its Hitchin map is proper and its two critical presentations are $T^*[-1]\mathscr{P}^{\mathrm{JS}}$ and $\mathrm{Crit}(\mathscr{P}^{\mathrm{JS}}, g)$ with the orientations of Section 2. On this open δ and the proper exchange e_h are isomorphisms. Proposition 4.5 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 4.5. 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 6.2; 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 4.6 places both immediate chambers in the core and identifies their normalized direct images by (4.13). If $\chi > -a_r$, then $-\chi < a_r$, so the whole positive immediate chamber in class $-\chi$ is empty by the lower endpoint already proved. Consequently $\mathcal{B}^{-\epsilon}_{r,\chi} = 0$, and the full ordinary source belongs to $D_{\mathscr{S}_r}$. This proves the upper source endpoint. □

5. The Hecke relation on the comparison arrow

The next construction proves a degree-zero Weyl relation on the whole arrow (4.12), including its square homotopy. It will also supply the ordinary-source relation used to prove support before any Hilbert-scheme calculation.

5.1. **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

$$X_\chi = \mathrm{Tot}(V_\Theta \oplus V_s \oplus R_\Theta^\vee \oplus R_s^\vee), \qquad e = mr(r+1), \quad d_\chi = \chi + B_r + 2e. \tag{5.1}$$

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

$$F_\chi = \langle \mathrm{res}\,\Theta, \xi\rangle + \langle \mathrm{res}\, s, \eta\rangle + \tfrac{1}{2}\beta_D(\mathrm{Tr}\,\Theta^2). \tag{5.2}$$

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 (4.12) also take the nonzero meromorphic-section open; on S_χ it is exactly the nonzero holomorphic-section condition.

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

$$\mu_{\min}(E_+) - 1 \le \mu_{\min}(E_-) \le \mu_{\min}(E_+). \tag{5.3}$$

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 (5.3) 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.

5.2. **Choices over the whole marked curve.** The auxiliary divisor can be chosen in a family over all of $C_x \times C_y$, including the diagonal. On the relative curve C_t put

$$D_{\mathrm{mark}} = \Gamma_x + \Gamma_y, \qquad K = \omega_{C_t/(C_x \times C_y)}(D_{\mathrm{mark}}) \otimes \mathcal{O}_{C_x \times C_y}(-\Delta_C).$$

In an étale coordinate the expression

$$\frac{(x-y)\,dt}{(t-x)(t-y)} \tag{5.4}$$

trivializes K along the full Cartier divisor D_{mark}. Here $x - y$ is the local basis of the invertible ideal of Δ_C. Under coordinate change f, write $f[u, v] = (f(u) - f(v))/(u - v)$ for the divided difference, extended to $f[u, u] = f'(u)$ on the diagonal. The ratio is $f[x, y]f'(t)/(f[t, x]f[t, y])$, equal to one modulo both graph ideals. Their intersection in the universal smooth three-variable ring is their product; the identity therefore holds on the doubled divisor as well. Swapping the marks changes (5.4) by -1.

A large even power K^{2k} consequently has a symmetric prescribed unit on D_{mark}. Sections extending that unit form an affine-vector-bundle torsor. Their zero divisors are finite flat and disjoint from both graphs, over arbitrary base rings. Extensions β_D form a second affine torsor. Denote the resulting smooth parameter space by $T \to C_x \times C_y$. Full frames along D are handled separately by their action groupoid.

In a local lattice chart the modification is $D_x = \mathrm{diag}(t - x, 1, \dots, 1)$. 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.

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

$$R_D = p^*\mathcal{O}_C(-D), \qquad R\,\mathrm{Hom}_{S_D}(R_D, F) = R\,\mathrm{Hom}_C(\mathcal{O}_C(-D), p_*F). \tag{5.5}$$

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

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

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

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

$$\mathfrak{H}_{\chi,D} = \mathfrak{H}^{\mathrm{der}}_{\chi,D} \overset{\mathbf{R}}{\underset{\mathfrak{P}_D}{\times}} P_D. \tag{5.6}$$

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 (5.6) is

$$\mathbb{L}_{\mathfrak{P}_D}|_{\mathfrak{H}_{\chi,D}} \longrightarrow \mathbb{L}_{\mathfrak{H}^{\mathrm{der}}_{\chi,D}}|_{\mathfrak{H}_{\chi,D}} \oplus \mathbb{L}_{P_D}|_{\mathfrak{H}_{\chi,D}} \longrightarrow \mathbb{L}_{\mathfrak{H}_{\chi,D}}. \tag{5.7}$$

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

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

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

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

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

Contracting the $r-1$ unit entries of D_x leaves the Koszul resolution of $(t-x, w-\lambda)$. Replacing a local extension $\lambda(t)$ by $\lambda(x)$ is a column operation since their difference is divisible by $t-x$. Under this contraction, the remaining arrow obstruction is its evaluation modulo $t-x$, the top obstruction class of the exterior point. Intrinsically this is the isomorphism

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

It is the derivative of c_Q, identified by the Koszul contraction. Since $P_D \to \mathfrak{P}_D$ induces isomorphisms in tangent degrees $-1, 0$, the tangent fiber of (5.7) 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

$$W_+ = \mathfrak{H}_{\chi,x} \overset{\mathbf{R}}{\underset{X_{\chi+1,T}}{\times}} \mathfrak{H}_{\chi,y}, \qquad W_- = \mathfrak{H}_{\chi-1,y} \overset{\mathbf{R}}{\underset{X_{\chi-1,T}}{\times}} \mathfrak{H}_{\chi-1,x}. \tag{5.8}$$

They are the upper and lower diagrams

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

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

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

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

The stable pushout recovers U. Finite limits and colimits supply the inverse functors and their unit and counit on every connective base, including mapping spaces and higher homotopies. These identifications commute with the maps to $\mathfrak{P}_D \times \mathfrak{P}_D$ after transposing the labels. They consequently commute with the base change from $P_D \times P_D$ in (5.6).

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

$$\operatorname{Hom}(\ell_x, \ell_y) \otimes \left[\mathcal{O} \xrightarrow{(u,v)^t} \mathcal{O}^2 \xrightarrow{(-v,u)} \mathcal{O}\right], \quad u = y - x, \quad v = \lambda_y - \lambda_x. \tag{5.9}$$

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

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

$$H^0(C, \mathcal{H}om(E, F)(-D)) = 0 \tag{5.10}$$

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 (5.10); 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

$$a^* F_{\mathrm{initial}} = b^* F_{\mathrm{final}}, \qquad a^{-1} S_{\mathrm{initial}} = b^{-1} S_{\mathrm{final}} \tag{5.11}$$

as functions and closed schemes.

For ordinary Hodge Borel–Moore theory use the mapping spectrum

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

Each derived word is a derived base change of the regular diagonal of the smooth framed middle endpoint. Both are quasi-smooth of virtual dimension q. The fundamental class is the coordinate-class specialization, localization boundary, and inverse Thom map; its composition and base-change paths are those of [23, §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 [11, Sections 3.2.19–3.2.21]. Swapping the two odd coordinate classes contributes -1, and swapping the two odd localization boundaries contributes -1. They cancel. This fixes the binary path, including its sign.

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

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

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

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

$$(y - x) s_j(y) = 0. \tag{5.12}$$

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

$$[W_+] - [W_-] = r[\mathfrak{D}] \quad \text{in } \pi_0 \,\mathrm{BM}_q(Z), \tag{5.13}$$

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

5.5. **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

$$\begin{aligned} \mathrm{BM}_q(Z) &\simeq \mathrm{Map}(a^* \omega_X, b^! \omega_Y \mathbb{L}^{q-d}) \\ &\simeq \mathrm{Map}(b_* a^* \omega_X, \omega_Y \mathbb{L}^{q-d}). \end{aligned} \tag{5.14}$$

Pushing a class properly inserts the pullback unit followed by the proper counit. Thus (5.14) 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 (5.11), A_0 is supported on S_Y. The mapping spectrum of squares to $\Gamma_Y K_Y \to K_Y$ is

$$\operatorname{Map}(A_0, \Gamma_Y K_Y) \underset{\operatorname{Map}(A_0, K_Y)}{\overset{h}{\times}} \operatorname{Map}(A_1, K_Y). \tag{5.15}$$

The first map to the base is an equivalence, since Γ_Y is the coreflector to supported objects. Consequently (5.15) projects equivalently to $\operatorname{Map}(A_1, K_Y)$. Its inverse produces the supported map and its square homotopy with contractible choice. It is additive and preserves the binary composition paths.

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

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

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

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

$$\phi_F(i_*i^!\mathscr{D}_\chi) \longrightarrow \phi_F \mathscr{D}_\chi. \tag{5.16}$$

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

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

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

$$Rp_*\phi_{p^*g}i^!\mathscr{D}_\chi = \phi_g\omega_{t_0\mathscr{P}}\mathbb{L}^{\chi-e} = \phi_g\Omega_{\mathscr{P}}\mathbb{L}^{\chi/2}\mathbb{L}^{-d_0/2}. \tag{5.17}$$

Its target is $R\pi_{g,*}\Phi_{\mathrm{Crit}(\mathscr{P},g)}\mathbb{L}^{\chi/2}\mathbb{L}^{-d_0/2}$. Thus (5.16) is the local δ in (4.12), with the same common factor at both vertices. Closed base change identifies its supported correspondence with the ordinary $\Omega_{\mathscr{P}}$ correspondence after this vector-bundle pushdown.

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

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

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

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

its mate identities give

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

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

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

and includes the two square homotopies. Composing the horizontal arrows is an exact functor on enhanced diagram categories. It takes (5.13) to a relation on precisely Δ, and then on its cofiber.

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

For one hyperbolic pair write $c = z_2 + \mathrm{i}z_1$, $\zeta = z_2 - \mathrm{i}z_1$. Then $c\zeta = z_1^2 + z_2^2$. Projection of the positive real thimble to ζ preserves orientation, whereas projection to c reverses it. Thus the raw support-counit scalar for a rank-k block in the compact dimensional-reduction convention is $(-1)^k$. Equivalently this is $\binom{2k}{2} \equiv k \pmod 2$ in the duality normalization of [24, Lemma 3.4]. The grouped orientation form is

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

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

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

Changing bounded bundle opens is handled by supported open excision of (5.16). 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 5.1. *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

$$F_\chi^c E_\chi^{c-1} - E_{\chi-1}^c F_{\chi-1}^{c-1} = r\rho_\chi(c-2, c). \tag{5.18}$$

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

$$\mathcal{A}_\chi(c) = Rh_*\omega_{t_0 \mathscr{P}_{r,\chi}^{\mathrm{all}}(c)} \mathbb{L}^{\chi + B_r/2}, \tag{5.19}$$

where the zero section is allowed.

Proof. At a fixed mark creation has target $\Sigma_{\chi+1}\mathbb{L}^{-1}$; removal already has degree zero. Indeed the ambient dimensions in (5.1) and (5.14), together with the common factor $\mathbb{L}^{-d_0/2}$ in (5.17), 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 (5.18). 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 (5.12) 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 (5.19) 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. □

5.7. **Finite-band vanishing.**

Lemma 5.2. *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 5.3. *For every r, χ and $c \le c_r$, the cofiber of the original comparison arrow* (4.12) *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 4.5 makes every cutoff restriction an isomorphism. Use its unique inverse to transport the operators in Proposition 5.1 to

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

Restriction naturality makes their relation $[F, E] = r\,\mathrm{id}$. Lemma 4.7 leaves a finite characteristic band, independent of j. Lemma 5.2 gives $V_\chi = 0$ for all χ. All cutoffs follow by their actual 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 5.4. *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 6.2. 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 4.7 leave the band $[a_r, -a_r]$. Apply the ordinary-source relation of Proposition 5.1 and Lemma 5.2 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 4.6 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 4.5. 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 6.2. Finite addition and the Serre property place every closed term in $D_{\mathscr{I}_r}$. The triangles therefore give support for these chamber and wall coefficients, completing the induction. Dimensional reduction identifies the ordinary source (5.19) with the split physical coefficient. Taking a sufficiently large positive chamber gives the stable-pair assertion. □

Corollary 5.5. *Let $R_{0,r,\chi}$ and $R_{1,r,\chi}$ be the normalized stable-pair direct images from X_0 and X_1. Then*

$$[R_{1,r,\chi}] - [\phi_{g_{\mathcal{A}}} R_{0,r,\chi}] \in K_0(\mathscr{I}_r). \tag{5.20}$$

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 5.3 and Proposition 4.5. They form a zigzag of isomorphisms in each perverse quotient. The stable-pair endpoints are bounded proper direct images by Proposition 9.5. Taking their Grothendieck classes and applying the Serre localization sequence gives (5.20). □

6. 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

$$R^L_{r,\chi} = R(h_{0,r,\chi})_*(\Phi_{P_{r,\chi}(X_0)} \otimes \Lambda^\chi). \tag{6.1}$$

Dimensional reduction and the ordinary-source argument of Theorem 5.4 put every perverse cohomology object of (6.1) 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

$$P_i^L = {}^p\mathcal{H}^i R(h_L)_* \mathbf{Q}_{M_L(r,\chi_*)}[B_r + 1]. \tag{6.2}$$

This Hodge module is independent of χ_* by [30, Theorem 0.4]. Define

$$b_r = \frac{\ell r(r+1)}{2} - rh, \qquad \gamma_r = 1 + rh + \frac{\ell r(r-1)}{2}. \tag{6.3}$$

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 6.1. *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,*

$$1 + \sum_{r\geq 1,\chi} Q^r x^\chi [R^L_{r,\chi}] = \mathrm{Exp}_{\mathrm{add}}\left(\sum_{r\geq 1} Q^r \sum_i \sum_{u,v\geq 0} x^{i+1+u+v} [P_i^L[-i-1-2v] \otimes \mathbf{T}^{v-B_r/2}] \right). \tag{6.4}$$

Only u, v are constrained to be nonnegative; i runs through the finite perverse range in (6.2)*. The exponential includes the derived symmetric-group actions.*

6.1. **The unframed input and its hypotheses.** We begin with the unframed support statement used in Section 4.

Proposition 6.2. *For any fixed Euler slope, every individual perverse direct-image coefficient of the split semistable unframed moduli stack has strict supports among the spectral-addition images Σ_ν. The rank-r primitive is the positive coarse intersection complex, whose proper Hitchin direct image is independent of the Euler characteristic as an object of $D^b\,\mathrm{MHM}(\mathcal{A}_r^L)$.*

Proof. Fix an Euler slope μ and let

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

We apply the smooth linear-stack integrality theorem [4, 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

$$\chi_S(E,F) = -\ell\,\mathrm{rk}(E)\,\mathrm{rk}(F), \qquad \mathrm{Ext}^2_S(E,F) = \mathrm{Hom}_S(F, E\otimes A)^\vee = 0. \tag{6.5}$$

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 [4].

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 [4, §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 [30, 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 [4, 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 $\mathrm{IC}_Y^{\mathrm{vir}} = \mathrm{IC}_Y \otimes \vartheta^{-\dim Y}$, the theorem gives the cohomological induction isomorphism

$$\mathrm{Sym}_{\oplus}\left(\bigoplus_{\eta\neq 0} \mathrm{IC}^{\mathrm{vir}}_{M_\eta} \otimes H^*(B\mathbf{G}_m)_{\mathrm{vir}}\right) \xrightarrow{\sim} Rp_*\,\mathrm{IC}^{\mathrm{vir}}_{\mathfrak{N}_\mu}, \quad H^*(B\mathbf{G}_m)_{\mathrm{vir}} = \Lambda \otimes \bigoplus_{v\geq 0} \mathbf{T}^v[-2v]. \tag{6.6}$$

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 (6.5) makes the ordinary transverse field zero; the derived statement follows from the physical model in Section 2. Its half/dual-half triangle identifies the fixed orientation with the cotangent orientation. Numerical symmetry and [4, 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 [4, §§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 [30, 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 (6.6) 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 $\mathscr{S}_r$. □

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

Lemma 6.3. *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*

$$e_\nu = \ell \sum_{i<j} r_i r_j. \tag{6.7}$$

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 (6.3) 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 6.4. *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 6.3, 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 6.5. *On U_r^{red} the split stable-pair problem is the relative Hilbert scheme of length*

$$m = \chi + \gamma_r - 1. \tag{6.8}$$

This Hilbert scheme has smooth total space of dimension $b_r + m = B_r + \chi$. Its normalized coefficient in (6.1) *is $\mathbf{Q}[B_r] \otimes \mathbf{T}^{-B_r/2}$.*

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

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

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

$$R(\pi_T p_T)_* R\mathcal{H}om_{S_T}(G, G \otimes p_T^* A) \simeq R\pi_{T,*}[\mathcal{E}nd(E) \otimes A \xrightarrow{[\Theta,-]} \mathcal{E}nd(E) \otimes A \otimes L]. \tag{6.9}$$

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 (6.8).

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

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

Here $K_S = p^*A$ is negative on each component. The monic equation gives $\mathcal{N} = p^* L^r|_D$ and $H^0(D, \mathcal{N}) = \bigoplus_{j=1}^r H^0(C, L^j)$. Thus the Hitchin tangent directions surject onto the conductor quotient. They are transverse to the equigeneric tangent space. The all-length criterion [38, Theorem 4.16(3)], valid also for reducible locally planar curves, proves smoothness of every

relative Hilbert scheme on this open. At a nodal curve it gives independent node-smoothing parameters.

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

6.3. **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

$$\epsilon_I = \det(E_I)^\vee \otimes \det(P_I). \tag{6.10}$$

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

$$\mathrm{Gr}^W H_{\lim} = H^1(G, \mathbf{Q}) \oplus V \oplus H_1(G, \mathbf{Q}) \otimes \mathbf{T}. \tag{6.11}$$

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

$$\det(E_I)^\vee \otimes \mathbf{T}^p \otimes \wedge^{j-p}(V \oplus A_{G\setminus I}). \tag{6.12}$$

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

$$\sum_j (-y)^j [\mathrm{IC}_0(\wedge^j H_{\mathrm{sm}})|_{S_\nu}] = \wedge_{-y} V\, K_G(y), \quad K_G(y) = \sum_{\substack{I \subset J \\ \pi_0(G\setminus I) = \pi_0(G)}} (y\mathbf{T})^{|I|} \epsilon_I \wedge_{-y} A_{G\setminus I}. \tag{6.13}$$

These are identities of Hodge classes with their deck actions.

Proof. Normalization and nearby cycles give (6.11) 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 (6.12) 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 (6.13) 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 [46, §§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 [58, Theorems 4.3, 4.5 and Proposition 4.8] retains $F, W, N, \mathrm{can}, \mathrm{Var}$ and the coordinate-order compatibility. On an overlap $t'_e = u_e t_{\sigma(e)}$, logarithmic coordinate change acts by exponentials of N_e. Each logarithm lowers the untwisted weight by two, so its action on the weight graded is the identity apart from the recorded graph permutation and the transition of V. The same holds on every graded image by strictness. Thus the displayed Hodge-graded terms descend and determine the Grothendieck class of the limiting mixed variation. Finite labeling descent is kept equivariant throughout. □

Lemma 6.7. *Let π_m be the relative length-m Hilbert scheme over S_ν and set $\mathcal{Z}_G(y) = \sum_{m\ge 0} y^m[R(\pi_m)_*\mathbf{Q}]$. Then*

$$\mathcal{Z}_G(y) = \wedge_{-y} V \sum_{I\subset J} (y\mathbf{T})^{|I|}\epsilon_I \frac{\wedge_{-y} A_{G\setminus I}}{\wedge_{-y} B_{G\setminus I}}. \tag{6.14}$$

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

$$\mathcal{Z}_G(y) = \wedge_{-y} V \sum_{\tau} (y\mathbf{T})^{|J_\tau|}\epsilon_{J_\tau} \frac{K_{G\setminus J_\tau}(y)}{\wedge_{-y} B_{G\setminus J_\tau}}. \tag{6.15}$$

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

$$P(y)(1-y)(1-\epsilon y) = 1 - \epsilon y + \mathbf{T}y^2. \tag{6.16}$$

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

$$(1-\epsilon y^d)(1-\epsilon\mathbf{T}^d y^d) + \epsilon\mathbf{T}^d y^d = 1 - \epsilon y^d + \mathbf{T}^d y^{2d}. \tag{6.17}$$

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 (6.16). 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 (6.14). Formula (6.17) 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 (6.13) yields (6.15). This argument includes loops, multiple edges, and disconnected normalizations. □

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

Proof of Theorem 6.1. On U_r^{red}, Lemma 6.5 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 6.4. 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 6.5, the length-m Hilbert direct image of $\mathbf{Q}[b_r+m]$ has perverse-degree-i piece pure of weight b_r+m+i. On a strict support of codimension e, its generic variation has weight and unshifted degree

$$\mathrm{wt} = m+i+e, \qquad \text{shift } [-m-e-i]. \tag{6.18}$$

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 6.6. Subtract them in (6.15). 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 (6.18). 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 $\nu = (r_1, \ldots, r_s)$, (6.7) gives

$$B_r - 2e_\nu = \sum_j B_{r_j}, \qquad \gamma_r - 1 - e_\nu = \sum_j (\gamma_{r_j}-1). \tag{6.19}$$

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 (6.10); the remaining action is the ordinary derived symmetric action, with its Koszul signs. The supportwise computation is therefore the right-hand side of (6.4) on U_r^{red}.

Both sides have perverse constituents in $\mathscr{S}_r$: on the left by Section 4, on the right by finite addition and Lemma 6.4. 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. □

7. 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

$$P_i^K = {}^p\mathcal{H}^i R(h_K)_* \mathbf{Q}_{M_K}[2N_r]. \tag{7.1}$$

They include every strict support of the canonical Hitchin map.

Proposition 7.1. *There is a closed immersion*

$$f_r : \mathcal{A}_r^{K_C} \times \mathbf{A}^1 \longrightarrow \mathcal{A}_r^L, \qquad (\theta, \lambda) \longmapsto b\theta + a\lambda \tag{7.2}$$

on characteristic coefficients. For the potential $g_{\mathcal{A}}$ in (1.14)*, one has*

$$\phi_{g_{\mathcal{A}}} P_i^L \cong (f_r)_*(P_i^K \boxtimes \mathbf{Q}_{\mathbf{A}^1}[1]) \otimes \mathbf{T}^{(B_r - 2N_r)/2}. \tag{7.3}$$

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

7.1. **The critical scheme on arbitrary ordinary bases.** The derived critical-stack identification [26, Theorem 5.6 and Proposition 5.7], restricted to the semistable locus as in [25, Proposition 3.6], uses the function

$$g(E, \Phi) = \frac{1}{2}\langle \alpha_q, \mathrm{Tr}(\Phi^2) \rangle = g_{\mathcal{A}} \circ h_L. \tag{7.4}$$

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

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

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

$$\mathcal{C}_T = R\pi_*[\mathcal{E}nd(E) \xrightarrow{[\theta, -]} \mathcal{E}nd(E) \otimes K_C] \tag{7.5}$$

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 (7.5) commutes with derived ordinary base change, so the trace-free H^0 remains zero over every ordinary base. Consequently

$$z = \lambda \,\mathrm{id}_E, \qquad \lambda = \frac{1}{r}\,\mathrm{Tr}(z) \in \Gamma(T, \mathcal{O}_T). \tag{7.6}$$

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 [52, Tags 0A1G, 0G99 and 0656]. Conversely, a stable canonical Higgs family and a scalar λ give the required commuting family. These constructions are inverse on objects and morphisms and commute with base change.

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

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

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

$$Z := t_0\,\mathrm{Crit}(g : M_L \to \mathbf{A}^1) \cong M_K \times \mathbf{A}^1. \tag{7.7}$$

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 (7.7) identifies the gradient ideal. The extra $\mathbf{A}^1$ in that equation parametrizes the commuting scalar λ.

7.2. **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 (7.7) has dimension $2N_r + 1$ in the smooth space M_L of dimension $B_r + 1$. Its normal bundle $\mathcal{N}$ has even rank

$$c = B_r - 2N_r = (\ell - 2h)r^2 - 2. \tag{7.8}$$

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,

$$\begin{aligned} \det \mathrm{RHom}_S(\widetilde{E}, \widetilde{E}) &= \det R\Gamma\mathcal{E}nd(E) \otimes \left(\det R\Gamma(\mathcal{E}nd(E) \otimes L)\right)^{-1} \\ &= \bigotimes_{p\in D} \det(\mathcal{E}nd(E_p) \otimes L_p)^{-1}. \end{aligned} \tag{7.9}$$

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, (7.9) 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 [25, 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 an actual trivialization

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

Consequently

$$\det \mathrm{Hess}(g) = u_Z \delta^{\otimes 2}, \qquad u_Z \in \Gamma(Z, \mathcal{O}_Z)^\times. \tag{7.10}$$

The units in (7.10) 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 [9, §2.2]. These facts also exclude vertical components: smooth irreducible components are open and closed, and their images under a proper flat morphism are closed and open in the connected base. Each would meet the generic fiber, whose connectedness permits only one. Hence M_K is integral. Its finite Stein algebra is generically the base algebra and lies in its function field; normality of the affine Hitchin base gives

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

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

$$\phi_g \mathbf{Q}_{M_L}[B_r + 1] \cong j_* \mathbf{Q}_Z[2N_r + 1] \otimes \mathbf{T}^{c/2}, \qquad j : Z \hookrightarrow M_L. \tag{7.11}$$

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

7.3. **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 (1.8) 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 (7.2) is a closed immersion over every ordinary base.

Proper vanishing-cycle exchange for h_L, followed by (7.11), gives (7.3) 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 7.1. In particular, the positive primitive in (6.4), after vanishing cycles and restriction to $\lambda = 0$, becomes

$$x^{i+1+u+v} P_i^K[-i - 2v] \otimes \mathbf{T}^{v-N_r}. \tag{7.12}$$

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

8. Scalar packets and the global physical identity

Write

$$R^K_{r,\chi} = R(h_{1,r,\chi})_*(\Phi_{P_{r,\chi}(X_1)} \otimes \Lambda^\chi), \tag{8.1}$$

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 9.5. Corollary 5.5 gives

$$[R^K_{r,\chi}] - [\phi_{g_{\mathcal{A}}} R^L_{r,\chi}] \in K_0(\mathscr{I}_r). \tag{8.2}$$

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

$$\begin{aligned} &1 + \sum_{r,\chi} Q^r x^\chi [C_{r,\chi}] \\ &\quad = \mathrm{Exp}_{\mathrm{add}} \left(\sum_{r \geq 1} Q^r \sum_i \sum_{u,v \geq 0} x^{i+1+u+v} \Big[(f_r)_* (P_i^K \boxtimes \mathbf{Q}_{\mathbf{A}^1}[1]) \right. \\ &\qquad\qquad \left. [-i - 1 - 2v] \otimes \mathbf{T}^{v-N_r} \Big] \right). \end{aligned} \tag{8.3}$$

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

8.1. **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 8.1. *Every simple support in $\mathscr{I}_r$ meets U_r, and*

$$K_0(\mathscr{I}_r) \longrightarrow K_0(\mathrm{MMHM}(U_r)) \tag{8.4}$$

is injective.

Proof. The lift of a in the extension (1.8) 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 6.4. 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 (8.4). The exclusion of the single rank-r unframed factor from $\mathscr{I}_r$ is essential at precisely this step. □

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

Lemma 8.2 (Excision for separated scalar packets)**.** *Work on an ordinary coprime factorization chart for the characteristic polynomial of z, and fix an allocation of the Euler characteristic to its factors. The Bézout decomposition of a stable pair is induced by an equivalence of the formal physical enhancements on all connective Artin bases. Under this same equivalence the closed Calabi–Yau form is the sum of the forms of its packets, with the original trace normalization $2\,\mathrm{Ch}_2$. Consequently the ordinary decomposition map preserves the algebraic d-critical sections.*

Proof. Use the projective compactification Y and its boundary framing from Section 2. 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

$$\mathfrak{D}(Y, D_\infty) = \mathrm{fib}_{(\mathcal{O}_Y, \mathcal{O}_{D_\infty}, \mathrm{id})} \left(\mathrm{Perf}(Y) \longrightarrow \mathrm{Pic}(Y) \underset{\mathrm{Pic}(D_\infty)}{\overset{h}{\times}} \mathrm{Perf}(D_\infty) \right), \tag{8.5}$$

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

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

For the closed-form comparison use the strict trace kernel $\mathscr{A}_I$ from Lemma 2.6. On U_0 the supertrace is the identity of $\mathcal{O}_{U_0}$, so $\mathscr{A}_I|_{U_0} = 0$. Excision therefore gives, compatibly with the derivative of the preceding gluing functor,

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

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

Equivalently, the canonical identifications with $\mathcal{O}$ off the supports give relative perfect-object classes. The category of perfect objects supported on the disjoint union is the product of the categories supported on the Z_i. We use the natural spectrum-valued Chern character and its additivity homotopies [21, Theorem 6.5], together with localization of the mixed complex [22, Sections 1.5, 1.12 and 2.4]. Taking the relative fibers in the localization diagram, and then the finite direct sum over the disjoint supports, gives the homotopy

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

in the relative negative-cyclic complex. The relative class is defined by its trivialization on the open complement, and the cyclic words above represent its supported trace. The constant object $\mathcal{O}_Y$ has no positive-degree form in the moduli directions. The fixed boundary and volume

transgression of $2\,\mathrm{Ch}_2$ is linear. Hence the very gluing equivalence constructed above pulls back the closed form to the sum of the packet forms, with all its higher Hodge components.

The passage from a (-1)-shifted symplectic enhancement to its d-critical section is functorial for these equivalences; in critical charts it is the canonical section determined by the closed form, as in [27, §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 8.3. *Over a base carrying a coprime factorization of the characteristic polynomial of z, the physical stable-pair coefficient factors into its scalar-packet coefficients with one common framing object. With the normalization Λ^{χ}, this factorization is compatible with refinement and with the ordinary derived permutation action. It holds over all ordinary bases of that factorization chart.*

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

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

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

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

$$D_{ij} := \mathrm{fib}(A_{ij}) = R\pi_{X,*}R\mathcal{H}om_{X_1\times T}(F_i, F_j) \tag{8.6}$$

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 2 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 2.4 is therefore

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

The cross summand of the cotangent complex is

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

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

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

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

$$o_1 \boxtimes o_2 = \det(d\Psi^*) \circ \Psi^* o \circ \gamma_{12}^{\otimes 2}, \qquad d\Psi^* : \Psi^* \mathbb{L}_{\mathfrak{P}} \longrightarrow \mathbb{L}_{\mathfrak{P}_1 \times \mathfrak{P}_2}. \tag{8.7}$$

Lemma 8.2 identifies the original algebraic d-critical sections for this same gluing map. The paired determinant comparison above is its derivative comparison, with the algebraic root γ_{12} and square (8.7). Consequently it identifies the original oriented vanishing-cycle coefficients.

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

Finally, Riemann–Roch gives

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

The opposite cross contractions have even parity. The raw oriented DT interchange has the extra character $(-1)^{\kappa_i \kappa_j}$. For raw degrees p, q, after the normalization its exponent is $pq + \kappa_i\kappa_j + \chi_i q + \chi_j p$. The ordinary derived exchange of the shifted degrees has exponent $(p + \chi_i)(q + \chi_j)$. These agree modulo two. The fixed positive-quotient orientation comparison is the same paired determinant comparison in Section 2; its opposite residue blocks have even length and contribute no additional permutation character. Hence the normalized packet factorization has exactly the derived symmetric action. □

8.3. **Stratified descent and induction.**

Theorem 8.4. *For every rank and Euler characteristic,*

$$[R_{r,\chi}^K] = [C_{r,\chi}]. \tag{8.8}$$

Thus (8.3) *is the full monodromic Hodge class identity for the original physical stable-pair coefficients on* X_1.

Proof. First apply vanishing cycles to (6.4). 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 7.1 now identifies the result with (8.3). In particular, (8.2) becomes

$$[R_{r,\chi}^K] - [C_{r,\chi}] \in K_0(\mathscr{I}_r). \tag{8.9}$$

Suppose (8.8) is known in all smaller positive ranks. Stratify U_r by the multiplicities of its distinct evaluated roots. For a multiplicity partition $\nu = (r_1, \dots, r_s)$, $s \geq 2$, denote the stratum by $\mathscr{R}_\nu$ and its ordered-root cover by $p_\nu : \widetilde{\mathscr{R}}_\nu \to \mathscr{R}_\nu$. This is a finite étale torsor under the group G_ν permuting roots of equal multiplicity. These covers are used separately on the root-multiplicity strata.

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

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

Spectral addition with the ordered roots defines a finite map $v_\nu : W_\nu \to \widetilde{\mathscr{R}}_\nu$. It is the closed coefficient-matching locus in a base change of the finite addition map; its image can be proper. Proper base change and Lemma 8.3 give the G_ν-equivariant expression

$$p_\nu^*(R^K_{r,\chi}|_{\mathscr{R}_\nu}) \cong (v_\nu)_* \left(\bigoplus_{\sum_j \chi_j = \chi} \boxtimes_j (R^K_{r_j,\chi_j}|_{T_{r_j}}) \Big|_{W_\nu} \right). \tag{8.10}$$

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

The coefficients $C_{r,\chi}$ have the same expression. Each of their primitive factors lies on one scalar by (7.3); grouping its symmetric powers according to distinct scalar roots gives (8.10) with C in place of R^K. This uses the induction identity for finite symmetric projectors and keeps the within-root stabilizers. Equalities within a repeated-root packet are used on reductions, which determine the Hodge modules.

The induction hypothesis is an equality of full Grothendieck classes, and this is sufficient here. The tensor power of a short exact sequence has a finite equivariant filtration whose graded terms are induced tensor products of the subobject and quotient. Degree shifts carry their sign representations. Thus tensor and symmetric powers preserve the lower-rank equalities with the full G_ν action. Rational finite invariants are exact, so these equalities descend through p_ν. Finite closed/open localization combines the root strata and gives equality on all of U_r. By (8.9) and Lemma 8.1, 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 (8.10) is bounded below and, with fixed sum, bounded above. There are finitely many partitions and perverse indices, completing the coefficientwise argument. □

8.4. **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,

$$K_{r,\chi} = \Phi_{P_{r,\chi}(X_1)}|_{N_{r,\chi}} \otimes \Lambda^\chi. \tag{8.11}$$

The map $h_{K,\mathrm{nil}}$ records the canonical field θ. On reductions, $h_1|_{N_{r,\chi}} = f_r|_{\lambda=0} \circ h_{K,\mathrm{nil}}$; simultaneous triangularization of commuting θ, z proves this equality on geometric points, hence on the reduced schemes. Hodge modules and their six operations are unchanged by nilpotent thickenings. Properness in Proposition 9.5 gives the required proper base change.

Corollary 8.5. *Put $R^{\mathrm{nil}}_{r,\chi} = R(h_{K,\mathrm{nil}})_* K_{r,\chi}$. In the coefficientwise completed Grothendieck group of monodromic Hodge modules on the canonical Hitchin bases,*

$$1 + \sum_{r,\chi} Q^r x^\chi [R^{\mathrm{nil}}_{r,\chi}] = \mathrm{Exp}_{\mathrm{add}} \left(\sum_{r\geq 1} Q^r \sum_i \sum_{u,v\geq 0} x^{i+1+u+v} [P_i^K[-i-2v] \otimes \mathbf{T}^{v-N_r}] \right). \tag{8.12}$$

After compact integration, if

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

then

$$[Q^r] \operatorname{Log} Z_{\mathrm{nil}}(Q,x) = \frac{x\mathbf{T}^{-N_r}}{(1-x)(1-x\mathbf{T})} \sum_i (-x)^i [R\Gamma_c(\mathcal{A}_r^{K_C}, P_i^K)]. \tag{8.13}$$

Once the ordinary nil cohomology theorem 9.2 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 (8.3) by proper base change, use (7.12), and identify its closed image with the canonical Hitchin base. This gives (8.12) with the ambient restriction (8.11).

Compact integration commutes with these external products, finite proper addition, and finite projectors. The plethystic logarithm extracts the primitive. Summing $u, v \geq 0$ gives $x/((1-x)(1-x\mathbf{T}))$ and the shift $[-i-2v]$ gives $(-1)^i$, proving (8.13) in $K_0(\mathrm{MMHS})$. The scalar primitives are ordinary by Proposition 7.1; the nil cohomology is ordinary by Theorem 9.2. 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 3. 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 [10, §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. □

9. 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

$$H : \ (\theta, z) \longmapsto (u\theta, u^{-1}z). \tag{9.1}$$

Retain the characteristic-nil locus $N_{r,\chi}$ and its ambient coefficient $K_{r,\chi}$ from (8.11).

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

$$\omega_\lambda = \mathrm{Hom}_C(\mathcal{O}_{Z_\lambda}, K_C), \qquad D_{\lambda,d} = \mathrm{Quot}^H_{Z_\lambda}(\omega_\lambda, d). \tag{9.2}$$

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

$$L_{\lambda,d} = \bigoplus_a Rf_{d,a,*}(\Phi_{d,a} \otimes \Lambda^a[d]). \tag{9.3}$$

Here the enhancement and $\Phi_{d,a}$ are inherited from the fixed problem for physical stable pairs. The identification with a Quot scheme in (9.2) concerns its truncation.

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

$$d(\mathbf{r}) = \sum_b h_b r_b, \quad R(\mathbf{r}) = \sum_b r_b, \quad A(\mathbf{r}) = \sum_b (a_b+1) r_b, \quad L(\mathbf{r}) = \sum_b l_b r_b, \tag{9.4}$$
$$e_b = l_b - a_b - 1, \qquad E(\mathbf{r}) = \sum_b e_b r_b, \qquad C_{\mathbf{r}} = \prod_{b\in\lambda} \operatorname{Sym}^{r_b} C.$$

The map

$$\mu_{\mathbf{r}} : C_{\mathbf{r}} \longrightarrow \operatorname{Sym}^{d(\mathbf{r})} C, \qquad (D_b)_b \longmapsto \sum_b h_b D_b \tag{9.5}$$

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

Theorem 9.1. *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*

$${}^p\mathcal{H}^k L_{\lambda,d} \simeq \bigoplus_{\substack{d(\mathbf{r})=d \\ R(\mathbf{r})-2A(\mathbf{r})=k}} (\mu_{\mathbf{r}})_* \mathbf{Q}_{C_{\mathbf{r}}}[R(\mathbf{r})] \otimes \vartheta^{E(\mathbf{r})}. \tag{9.6}$$

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

$$\sum_{d\geq 0} v^d [L_{\lambda,d}] = \operatorname{Sym}_{\mathrm{add}} \left(\sum_{b\in\lambda} v^{h_b} (\delta_{h_b})_* \mathbf{Q}_C \otimes \vartheta^{e_b}[2(a_b+1)] \right), \tag{9.7}$$

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

Theorem 9.2. *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*

$$H_c^j(N_{r,\chi}, K_{r,\chi}) \simeq \bigoplus_{\substack{\lambda\vdash r \\ d_\lambda \geq 0}} \bigoplus_{d(\mathbf{r})=d_\lambda} H^{j-2L(\mathbf{r})}(C_{\mathbf{r}}, \mathbf{Q})(-L(\mathbf{r})). \tag{9.8}$$

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

9.1. **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 $\operatorname{Hilb}^r(\mathbf{A}^2)^H$. Each difference-weight summand of $\mathbf{C}[\theta, z]$ is a rank-one module over $\mathbf{C}[\theta z]$. An H-homogeneous ideal of finite colength is supported at the origin, so each such summand is a power submodule in θz. This gives the monomial ideals indexed by partitions. Let t_1, t_2 be the geometric tangent characters of (z, θ). At a partition ideal the two tangent characters attached to a box are

$$t_1^{-l_b} t_2^{a_b+1}, \qquad t_1^{l_b+1} t_2^{-a_b}. \tag{9.9}$$

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

$$0 \longrightarrow \mathcal{O}_{Z_\lambda} \longrightarrow F \longrightarrow Q \longrightarrow 0. \tag{9.10}$$

The quotient is finite and flat over the base, by the fiberwise injection criterion applied to the two curve bundles. Curve duality sends (9.10) 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 (9.10). These operations commute with base change. They identify the fixed moduli functor with (9.2) on all ordinary Artin bases, with the physical enhancement retained as specified above.

Use $\mathbf{P}^1$ as a reference curve and restrict the support of the defect to its affine coordinate t. The second volume-preserving torus is

$$R: \ (t, \theta, z) \longmapsto (vt, v^{-1}\theta, z). \tag{9.11}$$

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 9.3. *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

$$t^{m_{ij}} e_{ij} \longmapsto \sum_k c_{ij,k} t^{m_{ij}-k} e_{i+k,j+k}, \qquad 0 \le m_{ij} - k < m_{i+k,j+k}. \tag{9.12}$$

For $k \le 0$ the displayed inequality contradicts monotonicity. For $k > 0$ apply θ. Every target term remains nonzero after this operation: its new box lies in λ and its height is at least m_{ij}. Thus the coefficient propagates to column $j-1$. At column zero the source has zero θ-image whereas every target still survives; the coefficient is zero. Induction on j proves $\mathrm{Hom}(K, M/K)^{H,R} = 0$.

A common fixed quotient is supported at $d[0]$ and is killed by t^d, by Cayley–Hamilton. It lies in a projective finite-module Quot fiber. The residual torus has a limit there, which is monomial. A nonconstant residual orbit would have positive-dimensional closure through this limit, contradicting the vanishing of its fixed tangent. This proves the classification and reducedness.

Differentiating the Artin-family identification above identifies the invariant tangent with $\mathrm{Ext}^1(I,I)_0^{H,R}$. Stability kills the trace-free Ext^0, and proper-support Calabi–Yau duality kills the invariant Ext^2 and Ext^3. The derived fixed tangent is therefore acyclic. For $d = 0$ the same assertion follows before taking curve cohomology: the trace-free local self-Ext complex of the partition ideal has only the nonzero hook characters in (9.9). Its normalized oriented coefficient is the unit. The orientation comparisons of Section 2 give the unit for every common fixed point. □

9.2. **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 2, it is

$$p^+_!(i^+)^*\Phi_Y \simeq \Phi_{Y^0} \otimes \Lambda^{I_\mu}. \tag{9.13}$$

Here I_μ is the index of Proposition 3.1, 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 [27]. 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 (9.13). The chart maps, orientation torsors, and their descent in MMHM were fixed in Section 2.

Lemma 9.4. *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

$$([v_w], v) \longmapsto ([t^{w-w_i} v_w], t \cdot v). \tag{9.14}$$

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 (9.14). Separatedness gives uniqueness. The intersections of the fixed locus with the distinct projective weight blocks are open and closed. This proves the assertion. □

For the Quot problem on $\mathbf{A}^1$, let $p(t)$ be the characteristic polynomial of the support. A quotient is killed by $p(t)$ and is therefore a quotient of the vector bundle $M/p(t)M$. Stability under t, θ, z and the characteristic-polynomial condition are closed incidence equations in a quotient Grassmannian. Consequently its support map to $\mathrm{Sym}^d \mathbf{A}^1$ is projective. The monomial basis supplies a linearization; the base weights are $1, \ldots, d$. Lemma 9.4 and Lemma 9.3 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

$$H^j_c(D_{\lambda,d,a}(\mathbf{A}^1), \Phi_{d,a}) \text{ is a sum of copies of } \vartheta^j. \tag{9.15}$$

The untwisted coefficient $\Phi_{d,a}$ is monodromically Verdier self-dual. Properness of the support map and the positive-base contraction adjunction [2, 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 (9.15) makes it a sum of ϑ^j.

9.3. **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

$$\mathcal{B}_\lambda = \{\lambda_i - i - 1 : i \geq 0\}, \quad \sigma(c) = \begin{cases} 1 & c \in \mathcal{B}_\lambda, \\ -1 & c \notin \mathcal{B}_\lambda, \end{cases} \quad d_c = \sum_{j-i=c} m_{ij}. \tag{9.16}$$

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

$$\begin{aligned} v_m = \mathsf{Z}_m - \kappa\overline{\mathsf{Z}_m} \\ - (1-t_1)(1-t_2)\big(-t_3\mathsf{Z}_\lambda\overline{\mathsf{Z}_m} + \overline{\mathsf{Z}_\lambda}\mathsf{Z}_m + (1-t_3)\mathsf{Z}_m\overline{\mathsf{Z}_m}\big). \end{aligned} \tag{9.17}$$

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, $-1-\operatorname{sgn}(e-1)$, $\operatorname{sgn}(e)+\operatorname{sgn}(e-1)$, and $1-\operatorname{sgn}(e)$. For a legal height increase at (i,j), put $\alpha=m_{ij}+1$, $c=j-i$, and $q=t_1^{-i}t_2^{-j}t_3^{\alpha}$. Replacing Z_m by Z_m+q in (9.17) gives

$$\begin{aligned}\Delta v &= q-\kappa\overline{q}\\ &\quad -(1-t_1)(1-t_2)\big[-t_3\mathsf{Z}_\lambda\overline{q}+\overline{\mathsf{Z}_\lambda}q+(1-t_3)(q\overline{\mathsf{Z}_m}+\mathsf{Z}_m\overline{q}+1)\big].\end{aligned}$$

Pair opposite weights in the half-sign evaluation. For the box (l,k) put $D=k-l-c$ and $e=\alpha+j-k$. Its baseline contributes $-A(D,e)$, and its occupied levels contribute

$$\sum_{\beta=1}^{m_{lk}}\big(-A(D,e-\beta)+A(D,e-\beta+1)\big)=A(D,e)-A(D,e-m_{lk}).$$

The term $q-\kappa\overline{q}$ contributes $\operatorname{sgn}(-c,\alpha+j)$, and the constant term $-P$ contributes 1. Adding these contributions gives

$$\Delta I=\operatorname{sgn}(-c,\alpha+j)+1-\sum_{(l,k)\in\lambda}A(k-l-c,\alpha+j-k-m_{lk}). \tag{9.18}$$

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 (9.18) is

$$1+2(L_{c+1}-L_c)+2\mathbf{1}_{c=0}. \tag{9.19}$$

Counting row endpoints gives $L_c-L_{c+1}=\mathbf{1}_{c\in\mathcal{B}_\lambda}-\mathbf{1}_{c<0}$. Since $\alpha+j>0$, (9.18) reduces to $\Delta I=\sigma(c)$.

Restricting (9.17) to H kills its quadratic term in Z_m. A legal increment has difference $F_c(u)-F_c(u^{-1})$, with

$$F_c(u)=u^{-c}\big(1-(1-u^{-1})(1-u)\overline{\mathsf{Z}_\lambda}\big).$$

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

$$\begin{aligned}\Delta I_H &= \operatorname{sgn}(-c)+L_{c-1}-L_{c+1}=\tfrac{1}{2}\big(\sigma(c-1)+\sigma(c)\big),\\ I_H &= \tfrac{1}{2}\sum_c d_c\big(\sigma(c-1)+\sigma(c)\big),\quad I_R=\tfrac{1}{2}\sum_c d_c\big(\sigma(c)-\sigma(c-1)\big).\end{aligned} \tag{9.20}$$

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 (9.3), the punctual exponent and degree are

$$J=I_H-I_R=\sum_c d_c\sigma(c-1),\qquad \text{degree}=J-d. \tag{9.21}$$

We use the diagonal-trace rim-hook identity [54, 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

$$\sum_{v=c-l}^{c+a}\sigma(v-1)=l-a-1=e_b,\qquad e_b-h_b=-2(a_b+1). \tag{9.22}$$

Equations (9.21)–(9.22) and the already proved punctual purity give the punctual graded cohomology object as the symmetric object on the generators

$$v^{h_b}\vartheta^{e_b}[2(a_b+1)], \qquad b\in\lambda. \tag{9.23}$$

The intrinsic weight $k+d$ separates degree k before any global symmetric-power operation is taken.

9.4. **Moving support and finite symmetric actions.** The punctual calculation must descend with its Hodge structure and its finite permutation actions. Let x, y be two étale coordinates on an open set $U\subset C$. On a neighborhood V of the diagonal in $U_t\times U_p$, discard the other sheets of the two étale equivalence relations. Then

$$X=x(t)-x(p), \qquad Y=y(t)-y(p)=Xv(t,p),$$

where v is a unit. Over $B=U_p\times\mathbf{A}^1_s$ form

$$W=\operatorname{Spec}\big(\mathcal{O}(V)[s,\tau]/(X-s\tau)\big), \qquad \upsilon=\tau v(t,p). \tag{9.24}$$

It has two étale maps to $\mathbf{A}^1_\tau\times B$ and $\mathbf{A}^1_\upsilon\times B$. At $s=1$ they are the original coordinate maps; at $s=0$ their ratio is the linear change $\upsilon=(dy/dx)(p)\tau$. Cotangent lifting preserves the relative threefold volume and the transverse monomial module. For a fixed d, the support-cycle fiber d times the distinguished section is the constant punctual Quot scheme over B. This holds on nonreduced bases: formal étaleness and Cayley–Hamilton identify the quotients killed by the dth power of the section coordinate.

The supported fixed-determinant gluing construction in Proposition 2.8 applies to these two maps. On the completed parameter ring $R=\mathbf{C}[[p-p_0,s-s_0]]$ its reference branch is constant. The full-parameter calculation there gives equality of the two potentials modulo the square of the common critical ideal, for these same maps. Its algebraic four-Hom line map has square equal to the determinant of the same cotangent derivative. Faithful flatness of completion therefore identifies the existing algebraic oriented critical structures over B. Tensor the line map with id_{K_B} to retain the smooth-base factor.

Proper pushforward of the punctual comparison gives a map between pullbacks of one punctual cohomology complex over $U\times\mathbf{A}^1$. Pullback from U is fully faithful, by the affine-line cohomology adjunction. The two endpoint maps are equal. At $s=0$ this is the linear coordinate action. The physical orientation map gives an action over $\mathbf{G}_m$, and its induced maps on cohomology objects are morphisms of constant variations. They equal the identity at 1, and hence everywhere on the connected group. Thus all moving coordinate transitions are identity on the punctual Hodge variations.

On an ordered configuration space the finite defect decomposes canonically by support idempotents, also over nilpotent bases: $T=\bigoplus_i T_i$. Use a single baseline G_0 and the triangles $G_0\to G_i\to T_i$. Put

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

Use the equivariant surface and transported orientation of Lemma 2.7. The following grid is H-equivariant before taking determinants. Its weight-zero part compares the fixed orientations; its moving parts supply the index and permutation factors below. The supported Hom and section triangles give the line map

$$M(G)M(G_0)^{-1}\simeq\bigotimes_i\big(M(G_i)M(G_0)^{-1}\big). \tag{9.25}$$

The cross complexes $\mathrm{RHom}(T_i,T_j)$ vanish canonically for $i\neq j$; the volume trace splits by the same idempotents. Taking weight zero in this grid before determinants gives the oriented fixed-coefficient comparison, with the determinant functor's regrouping maps. The merge morphism

is étale on the labeled branch, since the idempotents lift uniquely over infinitesimal extensions. When two reference curves have baselines of different parity, divide first by their baseline point coefficients. The quadratic-point calibration in [24, Lemma 3.4] and [27, Section 7.8] is a rational unit map of the associated orientation torsors. The compact Thom and product-thimble maps have scalar $+1$. This fixes the cancellation in (9.25). The ordered-configuration variations are consequently constant; their cohomology still retains the topology of the configuration space.

We check the symmetry used in their descent. In the J-model of monodromic Hodge modules, let $r(x) = x^2$, let $\iota(x) = -x$ be its deck involution, and put

$$\vartheta = e_- r_* \mathbf{Q}_{\mathbf{A}^1}[1], \qquad e_- = (1 - \iota)/2.$$

Its global cohomology vanishes. Its kth convolution power is the full-sign summand of $(\sum x_i^2)_* \mathbf{Q}_{\mathbf{A}^k}[k]$. A smooth quadratic fiber retracts equivariantly onto S^{k-1}; its vanishing cohomology has determinant character. A coordinate permutation has its sign on this cohomology, and the shifted [1] factors have the same Koszul sign. The product is $+1$. These objects lie in the perverse heart, where rational realization is faithful. This proves that ϑ is even and Λ is odd, as stipulated in (1.12).

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

$$m_0 + m_+ + m_- = -d, \qquad a = m_- - m_+, \qquad \kappa = m_0 + 2m_+, \qquad \kappa + a = -d. \tag{9.26}$$

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

$$\kappa_i \kappa_j + (d_i - a_i)(d_j - a_j) = 2d_i d_j + 2a_i a_j.$$

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 (9.23). 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 [47, Theorem 5.3].

9.5. **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 $\sum d_i p_i$, proper base change and (9.25) 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:

$$\mathrm{wt}(A_a) = 0, \qquad \mathrm{wt}(L_{\lambda,d}) = d. \tag{9.27}$$

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 (9.7) has the same description: it assigns box occupancies at each support point, with total weighted multiplicity $\sum h_b r_b$. Taking finite rational invariants is exact. Equivariant tensor-power filtrations of short exact sequences show that passage to full K_0 classes preserves these projectors. For a finite group quotient of C^s, the invariant direct image of its constant object is the constant object of the symmetric product, including at diagonals: the unit followed by the projector is an isomorphism on the finite orbit fibers. Therefore finite closed/open localization on the base gives

$$[L_{\lambda,d}] = \left[\bigoplus_{d(\mathbf{r})=d} (\mu_{\mathbf{r}})_* \mathbf{Q}_{C_{\mathbf{r}}} \otimes \vartheta^{E(\mathbf{r})}[2A(\mathbf{r})] \right]. \tag{9.28}$$

Each box on the right has weight $e_b + 2(a_b+1) = h_b$. Hence both sides of (9.28) 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 9.1.

Twisting by the symmetric charge ϑ^{-d} changes the box generator to

$$v^{h_b}(\delta_{h_b})_* \mathbf{Q}_C(a_b+1)[2(a_b+1)]. \tag{9.29}$$

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.

9.6. **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 9.5. *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* (9.1) *and the positive grading on $\mathcal{A}_r^L$, and admits an H-linearized ample line bundle.*

Proof. The closed immersions f_s of (7.2) and the finite spectral addition of Lemma 6.3 show that, for each rank partition, the map

$$\prod_i (\mathcal{A}_{r_i}^{K_C} \times \mathbf{A}^1) \longrightarrow \mathcal{A}_r^L \tag{9.30}$$

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 (9.30) parameterize all compact effective cycles of rank r, with their multiplicities. Given a DVR family in $\mathcal{A}_r^L$ and a generic stable pair, after a finite extension its generic cycle lifts to one of these finite maps. Its lift extends over the integral-closure DVR and has a limiting cycle entirely in X. The projective stable-pair moduli on $Y = \mathbf{P}_C(\mathcal{O}_C \oplus K_C \oplus \mathcal{O}_C)$ extend the generic pair. Separatedness of its projective cycle identifies the limit with the cycle just constructed. Purity of the limiting sheaf prevents an additional isolated support component at infinity. The extended pair is therefore supported in X. Boundary avoidance descends through the finite extension, proving properness.

For $X_0 = \mathrm{Tot}_C(L \oplus A)$, the pullback of A to the normalization of any compact irreducible curve has negative degree. Its transverse coordinate is consequently zero. Every compact cycle lies in $\mathrm{Tot}_C(L)$ and is a finite spectral divisor. The same DVR argument, now with the ordinary positive spectral coefficients as the cycle parameters, proves properness for this row as well.

At a geometric nil point, the commuting θ, z are simultaneously triangularizable over the curve's generic field after extension, and the diagonal entries of z vanish. Hence

$$\mathrm{char}(az + b\theta) = \mathrm{char}(b\theta), \quad \mathrm{char}(au^{-1}z + bu\theta) = \mathrm{char}(bu\theta). \tag{9.31}$$

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 [53]. 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. □

9.7. **The whole nil locus.** Proposition 9.5 and Lemma 9.4 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 (9.13) computes the layers with exactly the coefficient (8.11).

The baseline Euler characteristic and H-index are

$$\chi_0(\lambda) = \chi\left(\bigoplus_{(i,j)\in\lambda} K_C^{-j}\right) = -h \sum_{(i,j)\in\lambda} (2j+1) = -h\|\lambda\|^2, \tag{9.32}$$

$$I_0(\lambda) = \sum_{b\in\lambda} \chi(K_C^{a_b+1}) = h\sum_{b\in\lambda}(2a_b+1) = h\|\lambda\|^2 = -\chi_0(\lambda). \tag{9.33}$$

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

$$\begin{aligned} R\Gamma(D_{\lambda,d,a}, \Phi_{d,a} \otimes \Lambda^{\chi+I_0+a}) &= R\Gamma(D_{\lambda,d,a}, \Phi_{d,a} \otimes \Lambda^{d+a}), \\ \bigoplus_a R\Gamma(D_{\lambda,d,a}, \Phi_{d,a} \otimes \Lambda^{d+a}) &= R\Gamma(\mathrm{Sym}^d C, \vartheta^d[-2d]L_{\lambda,d}). \end{aligned} \tag{9.34}$$

The latter is pure of complex weight zero by (9.27). The symmetric charge changes each generator of (9.7) to

$$\vartheta^{h_b+e_b}[2(a_b+1) - 2h_b] = \mathbf{T}^{l_b}[-2l_b]. \tag{9.35}$$

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 (9.34), 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 (9.8), including its degree and auxiliary monodromy assertions, and proves Theorem 9.2.

Remark 9.6. The perverse degree in (9.6) 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 10. Theorem 9.2 supplies the ordinary Hodge cohomology and its purity in each degree.

10. Hodge recovery and the character-variety formula

We combine the two calculations before taking any numerical realization. All formal variables in this section are even. For a bounded Hodge complex K, its class is $[K] = \sum_j (-1)^j [H^j(K)]$. Let

$$\zeta_C(y) = \sum_{m\geq 0} [H^*(\mathrm{Sym}^m C)]\, y^m = \frac{\bigwedge_{-y} H^1(C)}{(1-y)(1-\mathbf{T}y)}. \tag{10.1}$$

The last identity is the symmetric-product formula for a curve [28]; it is an identity of Hodge classes, with the derived symmetric action on $H^1(C)$ giving the exterior numerator.

10.1. **The identity in full Hodge classes.** Define the series of the original restricted physical coefficients by

$$Z_{\mathrm{nil}}(Q,x) = 1 + \sum_{r\geq 1}\sum_{\chi} Q^r x^{\chi} [R\Gamma_c(N_{r,\chi}, K_{r,\chi})]. \tag{10.2}$$

Theorem 9.2 gives

$$Z_{\mathrm{nil}}(Q,x) = \sum_{\lambda} Q^{|\lambda|} x^{-h\|\lambda\|^2} \prod_{b\in\lambda} \zeta_C(x^{h_b}\mathbf{T}^{l_b}). \tag{10.3}$$

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 (1.1). 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

$$P_i^K = {}^p\mathcal{H}^i Rh_{K,*}\mathbf{Q}_{M_K}[2N_r], \qquad A_i = [R\Gamma_c(\mathcal{A}_r^{K_C}, P_i^K)]. \tag{10.4}$$

The scalar primitive of Section 7, restricted to the characteristic-nil locus, is

$$x^{i+1+u+v} P_i^K[-i-2v] \otimes \mathbf{T}^{v-N_r}, \qquad u,v \geq 0. \tag{10.5}$$

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 (10.5) give

$$[Q^r]\,\mathrm{Log}\, Z_{\mathrm{nil}}(Q,x) = \frac{x\mathbf{T}^{-N_r}}{(1-x)(1-x\mathbf{T})} \sum_i (-x)^i A_i. \tag{10.6}$$

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 (10.6) in $K_0(\mathrm{MHS})$.

Equivalently, define

$$F_r(x) = x^{-1}\mathbf{T}^{N_r}(1-x)(1-x\mathbf{T})[Q^r]\,\mathrm{Log}\, Z_{\mathrm{nil}}(Q,x). \tag{10.7}$$

Then

$$F_r(x) = \sum_i (-x)^i [R\Gamma_c(\mathcal{A}_r^{K_C}, P_i^K)]. \tag{10.8}$$

The variable x records Euler characteristic; the source primitive identifies its exponent i with the Hitchin perverse coefficient in (10.8).

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 $r = \sum r_i$,

$$\sum_i a_{r_i} - a_r = \ell \sum_{i<j} r_i r_j \geq 0. \tag{10.9}$$

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 (10.1) 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 (10.3)–(10.8) hold.

10.2. **Recovering perverse-graded Hodge structures.** The moduli space M_K is smooth semiprojective: properness of h_K and positive scaling of its base give a limit of every scaling orbit, and the fixed locus lies in the projective nil fiber. Thus $H^D(M_K)$ is pure of weight D by [17, Corollary 1.3.2]. Smooth duality gives the same purity for $H^D_c(M_K)$. This is independent of the stable-pair purity proved in Section 9. The proper decomposition theorem in Hodge modules makes $H^j_c(\mathcal{A}_r^{K_C}, P_i^K)$ a direct summand of $H^{j+2N_r+i}_c(M_K)$. It is therefore pure of weight $2N_r + i + j$. Distinct j have distinct weights, so (10.8) implies

$$[H^j_c(\mathcal{A}_r^{K_C}, P_i^K)] = (-1)^{i+j}\, \mathrm{Gr}^W_{2N_r+i+j}([x^i]F_r(x)). \tag{10.10}$$

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 (10.10) recovers the pure Hodge structure $H^j_c(\mathcal{A}_r^{K_C}, P_i^K)$ for every i, j.

Normalize the perverse Leray filtration by $Rh_{K,*}\mathbf{Q}[N_r]$, so its index k corresponds to the perverse object $P^K_{k-N_r}$. The ordinary piece in degree D is

$$V_{D,k} = H^{D-N_r-k}(\mathcal{A}_r^{K_C}, P^K_{k-N_r}). \tag{10.11}$$

Duality gives $\mathbb{D}P_i^K = P^K_{-i}(2N_r)$, whence

$$V_{D,k} \simeq H^{N_r+k-D}_c(\mathcal{A}_r^{K_C}, P^K_{N_r-k})^\vee \otimes \mathbf{T}^{2N_r}. \tag{10.12}$$

In particular the compact indices and weight corresponding to (D, k) are

$$i_c = N_r - k, \qquad j_c = N_r + k - D, \qquad 2N_r + i_c + j_c = 4N_r - D. \tag{10.13}$$

Write $\beta_{D,k} = \dim V_{D,k}$. Equations (10.3), (10.10), and (10.12) determine these dimensions as well as the underlying pure Hodge structures.

10.3. **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 $2g$,

$$\zeta_C(y; s) = \frac{(1 - sy)^{2g}}{(1 - y)(1 - s^2 y)}. \tag{10.14}$$

Write $Z_{\mathrm{nil}}(Q, x; s)$ for the weight character of $Z_{\mathrm{nil}}(Q, x)$. Formula (10.3) gives

$$Z_{\mathrm{nil}}(Q, x; s) = \sum_\lambda Q^{|\lambda|} x^{-h\|\lambda\|^2} \prod_{b\in\lambda} \zeta_C(x^{h_b} s^{2l_b}; s). \tag{10.15}$$

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 (10.6) says

$$[Q^r] \operatorname{Log} Z_{\mathrm{nil}}(Q, x; s) = \frac{x s^{-2N_r}}{(1 - x)(1 - xs^2)} \sum_i (-x)^i E_c(P_i^K; s). \tag{10.16}$$

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

$$x = z^{-2}, \qquad s = zw. \tag{10.17}$$

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

$$z^{-2h(2a+1)}\frac{(z^{2a+1}-w^{2l+1})^{2g}}{(z^{2a+2}-w^{2l})(z^{2a}-w^{2l+2})}. \tag{10.18}$$

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 $Z_{\mathrm{nil}}(Q, z^{-2}; zw) = \Omega(Q; z, w)$ from (1.4). Multiplying (10.16) by $(z^2-1)(1-w^2)$ gives

$$\overline{H}_r(z,w) = (zw)^{-2N_r}\sum_i(-z^{-2})^i E_c(P_i^K; zw). \tag{10.19}$$

The signs are additive inverses in the lambda ring. The Adams operations fix -1, so $\psi_d(-s) = -s^d$.

The compact indices in (10.13) have total sign $(-1)^{i_c+j_c} = (-1)^D$ and weight $4N_r - D$. Substitution in (10.19) gives the finite formula

$$\overline{H}_r(z,w) = \sum_{D,k}\beta_{D,k}(-1)^D z^{2k-D}w^{2N_r-D}. \tag{10.20}$$

10.4. **Proof of Theorem 1.1.** Under nonabelian Hodge theory, the $P = W$ theorem [31, Theorem 0.2] identifies the normalized filtration with

$$P_kH^D(M_K, \mathbf{Q}) = W_{2k}H^D(\mathcal{M}_{g,r}, \mathbf{Q}) = W_{2k+1}H^D(\mathcal{M}_{g,r}, \mathbf{Q}). \tag{10.21}$$

This equality of filtrations under the cohomological identification implies $\dim \mathrm{Gr}^W_{2k} H^D(\mathcal{M}_{g,r}) = \beta_{D,k}$. The Hodge–Tate assertion in [16] now gives $H(\mathcal{M}_{g,r}; q, t) = \sum_{D,k}\beta_{D,k}q^kt^D$. For each summand of (10.20),

$$(t\sqrt{q})^{2N_r}(-1)^D(\sqrt{q})^{2k-D}\left(-\frac{1}{t\sqrt{q}}\right)^{2N_r-D} = q^kt^D.$$

This proves Theorem 1.1.

The argument also determines the multiplicities in the rational Tate decomposition of [49, 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 [29, 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

$$H^D(\mathcal{M}_{g,r}, \mathbf{Q}) \simeq \bigoplus_k \mathbf{Q}(-k)^{\oplus\beta_{D,k}}. \tag{10.22}$$

Formula (10.10) supplies the multiplicities in this splitting.

Corollary 10.1. *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 $D > 2N_r$ [1, Section 4.1]. The ordinary weight bound for smooth varieties [16, Theorem 2.1.2(5)] gives $2k \geq D$ whenever $\beta_{D,k} \neq 0$. Both exponents in (10.20) 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 (10.20) and preserves the parity of D. □

10.5. **Ranks one and two.** In rank one put $m = \chi + g - 1$. For $m \geq 0$, a stable pair is $\mathcal{O}_C \to \mathcal{O}_C(D)$ for an effective divisor D of degree m, together with a field in $H^0(C, K_C)$ and a scalar. Hence

$$P_{1,\chi}(X_1) = \mathbf{A}^{g+1} \times \operatorname{Sym}^m C, \qquad N_{1,\chi} = \mathbf{A}^g \times \operatorname{Sym}^m C.$$

The smooth physical coefficient is $\mathbf{Q} \otimes \Lambda^{-(g+1+m)}$. Its normalization by Λ^χ is $\mathbf{T}^{-g}[2g]$. Compact integration over $\mathbf{A}^g$ contributes $\mathbf{T}^g[-2g]$, so

$$R\Gamma_c(N_{1,\chi}, K_{1,\chi}) \simeq R\Gamma(\operatorname{Sym}^m C, \mathbf{Q}).$$

Consequently

$$[Q]Z_{\mathrm{nil}} = x^{-h}\zeta_C(x), \qquad F_1(x) = x^{-g}\mathbf{T}^g \bigwedge_{-x} H^1(C),$$

and the final substitution gives $(1+qt)^{2g}$.

For rank two, let $Z_j = [Q^j]Z_{\mathrm{nil}}$. The partitions (2) and (1, 1) give

$$Z_1 = x^{-h}\zeta_C(x), \qquad Z_2 = \zeta_C(x)\left(x^{-4h}\zeta_C(x^2) + x^{-2h}\zeta_C(x^2\mathbf{T})\right).$$

Expanding (1.6) through rank two yields

$$[Q^2]\operatorname{Log} Z_{\mathrm{nil}} = Z_2 - \frac{1}{2}Z_1^2 - \frac{1}{2}\psi_2(Z_1).$$

Under the weight character,

$$\psi_2(\zeta_C(x;s)) = \frac{(1-s^2x^2)^{2g}}{(1-x^2)(1-s^4x^2)}.$$

The two partition terms and the two subtractions therefore give

$$\begin{aligned}\frac{H(\mathcal{M}_{g,2};q,t)}{(1+qt)^{2g}} = {} & \frac{(1+q^2t^3)^{2g}}{(q^2t^2-1)(q^2t^4-1)} + \frac{q^{2g-2}t^{4g-4}(1+q^2t)^{2g}}{(q^2-1)(q^2t^2-1)} \\ & - \frac{q^{2g-2}t^{4g-4}(1+qt)^{2g}}{2(qt^2-1)(q-1)} - \frac{q^{2g-2}t^{4g-4}(qt-1)^{2g}}{2(q+1)(qt^2+1)}.\end{aligned}$$

This is the rank-two formula of [16, Theorem 1.1.3]. At $g = 2$ it reduces to

$$\begin{aligned}H(\mathcal{M}_{2,2};q,t) = (1+qt)^4(1 & + q^2t^2 + 4q^2t^3 + q^2t^4 + q^4t^4 \\ & + 4q^4t^5 + q^4t^6 + q^6t^6).\end{aligned}$$

References

[1] A. A. Beilinson, J. Bernstein and P. Deligne, *Faisceaux pervers*, Astérisque **100** (1982), 5–171; `https://archive.numdam.org/item/AST_1982__100__1_0/`.

[2] T. Braden, *Hyperbolic localization of intersection cohomology*, Transform. Groups **8** (2003), no. 3, 209–216; arXiv:math/0202251v5.

[3] C. Brav, V. Bussi and D. Joyce, *A 'Darboux theorem' for derived schemes with shifted symplectic structure*, arXiv:1305.6302v4 (2018).

[4] C. Bu, B. Davison, A. Ibáñez Núñez, T. Kinjo and T. Pădurariu, *Cohomology of symmetric stacks*, arXiv:2502.04253v2 (2025).

[5] E. Carlsson and F. Rodriguez-Villegas, *Vertex operators and character varieties*, Adv. Math. **330** (2018), 38–60; arXiv:1603.09267v1.

[6] 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.

[7] 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.

[8] M. A. A. de Cataldo, T. Hausel and L. Migliorini, *Topology of Hitchin systems and Hodge theory of character varieties: the case A_1*, Ann. of Math. **175** (2012), no. 3, 1329–1407, doi:10.4007/annals.2012.175.3.7.

[9] 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).

[10] B. Davison and S. Meinhardt, *Cohomological Donaldson–Thomas theory of a quiver with potential and quantum enveloping algebras*, arXiv:1601.02479v5 (2020).

[11] F. Déglise, F. Jin and A. A. Khan, *Fundamental classes in motivic homotopy theory*, arXiv:1805.05920v3 (2021).

[12] P. Descombes, *Hyperbolic localization in Donaldson–Thomas theory*, arXiv:2506.22400v1 (2025).

[13] T. Hausel, *Mirror symmetry and Langlands duality in the non-Abelian Hodge theory of a curve*, in F. Bogomolov and Y. Tschinkel (eds.), *Geometric Methods in Algebra and Number Theory*, Progr. Math. **235**, Birkhäuser, Boston, 2005, 193–217; arXiv:math/0406380v2.

[14] 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.

[15] T. Hausel, A. Mellit, A. Minets and O. Schiffmann, *$P = W$ via $\mathcal{H}_2$*, arXiv:2209.05429v3 (2026).

[16] T. Hausel and F. Rodriguez-Villegas, *Mixed Hodge polynomials of character varieties*, with an appendix by N. M. Katz, Invent. Math. **174** (2008), 555–624; arXiv:math/0612668v3.

[17] T. Hausel and F. Rodriguez-Villegas, *Cohomology of large semiprojective hyperkähler varieties*, arXiv:1309.4914v1 (2013).

[18] T. Hausel and M. Thaddeus, *Mirror symmetry, Langlands duality, and the Hitchin system*, Invent. Math. **153** (2003), no. 1, 197–229, doi:10.1007/s00222-003-0286-7.

[19] N. J. Hitchin, *The self-duality equations on a Riemann surface*, Proc. London Math. Soc. (3) **55** (1987), no. 1, 59–126, doi:10.1112/plms/s3-55.1.59.

[20] N. J. Hitchin, *Stable bundles and integrable systems*, Duke Math. J. **54** (1987), no. 1, 91–114.

[21] 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.

[22] 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.

[23] A. A. Khan, *Intersection theory on derived stacks*, arXiv:1909.01332v2 (2026).

[24] T. Kinjo, *Dimensional reduction in cohomological Donaldson–Thomas theory*, Compos. Math. **158** (2022), 123–167; arXiv:2102.01568v2.

[25] 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.

[26] 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.

[27] T. Kinjo, H. Park and P. Safronov, *Cohomological Hall algebras for 3-Calabi–Yau categories*, arXiv:2406.12838v2 (2026).

[28] I. G. Macdonald, *Symmetric products of an algebraic curve*, Topology **1** (1962), no. 4, 319–343, doi:10.1016/0040-9383(62)90019-8.

[29] 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.

[30] 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.

[31] D. Maulik and J. Shen, *The $P = W$ conjecture for GL_n*, Ann. of Math. **200** (2024), no. 2, 529–556; arXiv:2209.02568v2.

[32] D. Maulik, J. Shen and Q. Yin, *Perverse filtrations and Fourier transforms*, Acta Math. **234** (2025), 1–69; arXiv:2308.13160v3.

[33] A. Mellit, *Integrality of Hausel–Letellier–Villegas kernels*, Duke Math. J. **167** (2018), no. 17, 3171–3205; arXiv:1605.01299v2.

[34] A. Mellit, *Poincaré polynomials of moduli spaces of Higgs bundles and character varieties (no punctures)*, Invent. Math. **221** (2020), 301–327; arXiv:1707.04214.

[35] 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.
[36] A. Mellit, *Toric stratifications of character varieties*, Publ. Math. Inst. Hautes Études Sci. **142** (2025), 153–240, doi:10.1007/s10240-025-00158-0.
[37] A. Mellit, *Cohomology rings of character varieties*, arXiv:2507.12454v1 (2025).
[38] 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.
[39] S. Monavari, *Double nested Hilbert schemes and the local stable pairs theory of curves*, Compos. Math. **158** (2022), no. 9, 1799–1849, doi:10.1112/S0010437X22007606.
[40] S. Monavari, *The refined local Donaldson–Thomas theory of curves*, Geom. Topol. **30** (2026), no. 9, 3423–3484; arXiv:2506.14359v3.
[41] S. Mozgovoy and O. Schiffmann, *Counting Higgs bundles and type A quiver bundles*, Compos. Math. **156** (2020), 744–769; arXiv:1705.04849.
[42] S. Nayak and P. Sastry, *Grothendieck duality and transitivity II: traces and residues via Verdier's isomorphism*, arXiv:1903.01783v2 (2019).
[43] T. Pădurariu and Y. Toda, *Topological K-theory of quasi-BPS categories for Higgs bundles*, arXiv:2409.10800v1 (2024).
[44] T. Pantev, B. Toën, M. Vaquié and G. Vezzosi, *Shifted symplectic structures*, arXiv:1111.3209v4 (2013).
[45] E. M. Rains and S. O. Warnaar, *A Nekrasov–Okounkov formula for Macdonald polynomials*, J. Algebraic Combin. **48** (2018), 1–30; arXiv:1606.04613v3.
[46] M. Saito, *Mixed Hodge modules*, Publ. Res. Inst. Math. Sci. **26** (1990), no. 2, 221–333, doi:10.2977/prims/1195171082.
[47] M. Saito, *Thom–Sebastiani theorem for Hodge modules*, arXiv:2605.05782v1 (2026).
[48] O. Schiffmann, *Indecomposable vector bundles and stable Higgs bundles over smooth projective curves*, Ann. of Math. **183** (2016), no. 1, 297–362, doi:10.4007/annals.2016.183.1.6.
[49] V. Shende, *The weights of the tautological classes of character varieties*, arXiv:1411.4975v2 (2015).
[50] C. T. Simpson, *Higgs bundles and local systems*, Publ. Math. Inst. Hautes Études Sci. **75** (1992), 5–95, doi:10.1007/BF02699491.
[51] T. Spaide, *Shifted symplectic and Poisson structures on spaces of framed maps*, arXiv:1607.03807v1 (2016).
[52] The Stacks Project Authors, *The Stacks Project*, `https://stacks.math.columbia.edu`; Tags 08GS, 0A1G, 0G99 and 0656.
[53] J. Stoppa and R. P. Thomas, *Hilbert schemes and stable pairs: GIT and derived category wall crossings*, arXiv:0903.1444v3 (2009).
[54] R. Sulzgruber, *Inserting rim-hooks into reverse plane partitions*, arXiv:1710.09695v2 (2018).
[55] Y. Toda, *Curve counting theories via stable objects I. DT/PT correspondence*, arXiv:0902.4371v3 (2009).
[56] Y. Toda, *Categorical Donaldson–Thomas theory for local surfaces*, arXiv:1907.09076v4 (2021).
[57] S. Tubach, *Mixed Hodge modules on stacks*, Forum Math. Sigma **13** (2025), e175; arXiv:2407.02256v3.
[58] J. Wildeshaus, *Modules de Hodge mixtes à croisements normaux*, C. R. Math. Acad. Sci. Paris **337** (2003), no. 7, 467–472, doi:10.1016/j.crma.2003.08.001.

Zhili College, Tsinghua University, Beijing 100084, China
Email address: `tangqy24@mails.tsinghua.edu.cn`