



← ALL PAPERS · SUPPORTING MATHEMATICS · V1.0 (WORKING DRAFT)

Constitutive Thresholds as Amplitude Strata of Deformation Complexes

A Unifying Framework for Topological Invariants in Physics

Z. Paz · ORCID 0009-0003-1690-3669 · V1.0 (working draft) · 2026

PURE MATHEMATICS

↓ DOWNLOAD PDF

Constitutive thresholds as amplitude strata of deformation complexes

Z. Paz · ORCID 0009-0003-1690-3669

V1.0 · Working draft, June 2026 · Pure Mathematics

Abstract

We propose that the recurrence of topological invariants in physics is organised by a single mechanism: a physical system carries a *deformation complex* over its space of configurations or moduli, and a *constitutive threshold* is the locus where the cohomology of that complex changes amplitude—where it passes from being concentrated in a single degree (a smooth space) to being spread across several degrees (a singular or stacky locus). We make this precise. A *constitutive datum* is a perfect complex equipped with a self-duality and a Kuranishi map; its *jump stratum* is where the negative-amplitude (isotropy) cohomology jumps. We prove two structural results: that concentration is equivalent to smoothness and that the self-duality forces the isotropy and obstruction jumps to occur in lockstep across the stratum (Theorem); and that the standard invariant of each setting is the transgression of the obstruction across the stratum, with its arithmetic type (integer-quantised, torsion, or signature) fixed by the type of the duality and the Kuranishi map (Theorem). We exhibit three instances. For flat connections on a closed oriented 3-manifold the structure is the (-1) -shifted symplectic moduli of Pantev–Toën–Vaquié–Vezzosi, the threshold is the reducible locus, and the transgression is the torsion factor of the Witten–Reshetikhin–Turaev invariant; this case is rigorous. For topological insulators the structure is the Bloch bundle over the Brillouin torus, the threshold is the gap-closing (Weyl) locus, and the transgression is the Chern number, integer-quantised because the Kuranishi map is winding-valued; the jump relation here is a known theorem. For closed-chain linkages the structure is the symplectic reduction realising the polygon space (Kapovich–Millson, Hausmann–Knutson): loop closure is the $SO(3)$ moment map, the moment-map adjunction makes the duality exact, the threshold is the collinear stratum, and we verify the paired jump and compute the Kuranishi form explicitly at a degenerate configuration, finding a nondegenerate quadratic branch point. We state the central open problem as a *Uniformity Conjecture*—that the self-duality is induced by the closure that generates the threshold—and we bound it: a Bennett $4R$ linkage shows that loop-closure *alone* does not induce a self-duality (its deformation complex is two-term with mismatched ranks), so the operative hypothesis is closure of moment-map or Poincaré type, not closure as such. A categorical backbone realises the flat-connection and linkage instances as specialisations of one shifted-symplectic construction; an explicit spectral test then shows the insulator instance is

not such a specialisation but a topological- K -theory parallel—on the torsion-free torus its Chern invariant is only a rational shadow, and its genuinely K -theoretic (real-symmetry-class) \mathbb{Z}_2 invariants are torsion with no shifted-symplectic counterpart. We are explicit about what does *not* fit the framework—in particular a causal-structure case based on a curvature-rate threshold, whose only invariant is an elementary wall-crossing class and which carries no deformation-complex self-duality.

Introduction

The question

Topological invariants appear across physics with a regularity that the usual explanation does not fully account for. The usual explanation is *robustness*: a topological quantity is stable under perturbation, hence a good observable. This is correct but secondary. It says why an invariant persists once present; it does not say why topology appears in the first place, nor why it appears *where* it does.

We propose a different order of explanation. A physical system is locally governed by a *deformation complex*—the linearisation of whatever equation defines it (flatness, loop-closure, the spectral-projection condition for a gapped insulator), together with the symmetries that act on its solutions. The cohomology of this complex reads off the local structure of the configuration or moduli space: degree-zero cohomology is the space of genuine deformations, negative degrees record symmetries (automorphisms, isotropy), positive degrees record obstructions. A *constitutive threshold* is then not the locus where some quantity changes value, but the locus where this complex *changes amplitude*: where the cohomology stops being concentrated in degree zero and spreads into the symmetry and obstruction degrees. On one side the space is a manifold; on the other it is singular or stacky. The topological invariant is the secondary class that measures the failure to continue the smooth (concentrated) structure across the stratum.

This inverts the standard order: the stratification is primary, the invariant is its shadow. We argue that this inversion is the content behind the slogan that “topology appears where closure generates thresholds”, and we make both halves precise.

Results and structure

Section 2 sets up the abstract object—a *constitutive datum* (Definition 1)—and proves the two structural theorems: amplitude concentration is equivalent to smoothness, and a self-duality on the complex forces the isotropy and obstruction jumps to be paired across the threshold (Theorem 3); the standard invariant is the transgression of the obstruction, its arithmetic type fixed by the duality and Kuranishi types (Theorem 6).

Sections 3–5 give three instances of increasing concreteness and decreasing completeness of the literature. The flat-connection case (§3) is rigorous and is the richest, carrying a genuine (-1) -shifted symplectic structure. The topological-insulator case (§5) is the case in which the transgression principle is already a theorem. The linkage case (§4) is the finite-dimensional one in which the central conjecture is checkable by explicit computation; we set it up correctly as a symplectic reduction and verify the duality, the paired jump, and the Kuranishi signature in matrices.

Section 6 assembles the common dictionary, states the Uniformity Conjecture, and discusses the interpretation (which we keep clearly separate from the mathematics). Section 7 gives the categorical backbone: an abstract self-dual perfect complex, the structural lemma in sharpest form, and the realisation of Instances A and B as specialisations of one Pantev–Toën–Vaquié–Vezzosi construction, with an explicit spectral test showing Instance C is a topological- K -theory parallel rather than such a specialisation. Section 8 is an explicit statement of scope: we record a case (causal structure) that does *not* fit, and why. Section 9 lists the open problems and the next computations in order.

What this paper does and does not claim

The framework theorems of §2 are essentially formal given their hypotheses; their value is in isolating the *weakest* structure that forces the threshold phenomenon (a self-duality of some shift, no more). Instance A is rigorous modulo standard results. Instance C’s jump relation is

a known theorem; the deformation-complex reading is the organising lens. Instance B's moment-map structure is established in the literature, the duality is then exact, and we verify the paired jump and the Kuranishi signature explicitly at a degenerate configuration (§4.4). The Uniformity Conjecture is open; we state it sharply and bound it with an explicit counterexample (the Bennett $4R$ linkage, Remark 11), which fixes its operative hypothesis. We do not claim the framework is universal: §8 records its boundary.

The framework

The object

Definition 1 (Constitutive datum). Let B be a derived stack (in the finite cases, a derived scheme). A *constitutive datum* on B consists of:

1. a *deformation complex* E^\bullet : a perfect complex on B , equal to the tangent complex \mathbb{T}_B or a shift of it, with finite-dimensional cohomology sheaves \mathcal{H}^i of locally constant Euler characteristic;
2. a *self-duality* of shift n : a quasi-isomorphism $D: E^\bullet \xrightarrow{\sim} (E^\bullet)^\vee[n]$, graded-symmetric or graded-antisymmetric, inducing perfect pairings on cohomology

$$\mathcal{H}^i \cong (\mathcal{H}^{-n-i})^*;$$

3. a *Kuranishi map* $\kappa: \mathcal{H}^1 \rightarrow \mathcal{H}^2$ (the quadratic part of the defining equation), compatible with D , presenting B locally as $\kappa^{-1}(0)/\exp(\mathcal{H}^{\leq 0})$.

The cohomology carries the local geometry: $\mathcal{H}^{\leq 0}$ is the infinitesimal isotropy (automorphisms/symmetries), \mathcal{H}^0 (modulo isotropy) the deformations, $\mathcal{H}^{\geq 1}$ the obstructions. The *type* of the datum is the triple

$$(n, \text{symmetry of } D, \text{arithmetic flavor of } \kappa).$$

The point of stating the duality at this level of generality—any shift, either symmetry—is that this, and nothing stronger, is what the threshold phenomenon requires. The familiar refinements (a symplectic form, a Calabi–Yau base, a mapping-stack origin) are case-specific.

Write $a^i(b) = \dim \mathcal{H}_b^i$, with a_{\min}^i the generic (minimal) value.

The stratification

Definition 2 (Constitutive stratification). The *jump stratum* is

$$\Sigma = \{ b \in B : \dim \mathcal{H}_b^{\leq 0} > (\dim \mathcal{H}^{\leq 0})_{\min} \} = \text{supp} (\mathcal{H}^{\leq 0} / \mathcal{H}_{\text{generic}}^{\leq 0}),$$

the locus where the isotropy jumps. We set $\text{\textbf{E}} = B \setminus \Sigma$ (the smooth, “EXISTS” regime) and $\text{\textbf{H}} = \Sigma$ (the singular, “HAPPENS” regime).

Theorem 3 (Concentration \Leftrightarrow smoothness; duality forces paired jumps). *For a constitutive datum:*

1. On $\text{\textbf{E}}$ the complex is concentrated: its cohomology is supported in degree 0 up to a constant central part, κ has locally constant rank, and B is smooth there—a manifold, or a smooth stack of fixed isotropy type if the central part is nonzero—of dimension a_{\min}^0 .
2. Across Σ the amplitude spreads, and the duality D forces the jumps to be paired:

$$\Delta a^i = \Delta a^{-n-i} \quad \text{for all } i.$$

In particular the isotropy jump is mirrored by an obstruction jump; isotropy and obstruction inflate together, by the self-duality, not by an index count.

3. Σ is closed and nowhere dense; $\text{\textbf{E}}$ is open and dense. Hence one enters $\text{\textbf{H}}$ only by specialisation (raising the isotropy, a closed condition) and is returned to $\text{\textbf{E}}$ by a generic perturbation. The asymmetry is order-theoretic—open/generic versus closed/special—and requires no time orientation.

Proof. (1) and (3) use only upper-semicontinuity of a^i , which is part of Definition 1(i): the locus $\{a^{\leq 0} > a_{\min}^{\leq 0}\}$ is closed, and the maximal-rank locus of a family of complexes is open and dense. On $\backslash\text{texts}\{\mathbf{E}\}$ the ranks of the differentials are locally constant, so $\kappa^{-1}(0)$ is cut out by a constant-rank quadratic and the implicit function theorem gives a manifold of dimension a_{\min}^0 after the constant-dimensional quotient. (2) is the content of the duality. The quasi-isomorphism D induces $\mathcal{H}_b^i \cong (\mathcal{H}_b^{-n-i})^*$ pointwise, hence $a^i(b) = a^{-n-i}(b)$ for all b and all i ; differencing across Σ gives the paired jumps. No symplectic form and no rank–nullity identity is used—only D . \square

Remark 4 (Euler characteristic is weaker than duality). Local constancy of $\chi = \sum (-1)^i a^i$ (Definition 1(i)) gives only the *alternating* relation $\sum_i (-1)^i \Delta a^i = 0$. The *term-by-term* pairing $\Delta a^i = \Delta a^{-n-i}$ is strictly stronger and is exactly what the self-duality D supplies. This distinction is the crux of the framework: many constraint systems have a constant χ (a Maxwell–Calladine-type index) without carrying a genuine self-duality, and only the latter forces the isotropy and obstruction to jump together. We return to this in §4.

The invariant

Definition 5 (Constitutive invariant). The invariant of the datum is the *transgression* of the obstruction across Σ : the secondary class measuring the obstruction to extending, over a neighbourhood of Σ , the trivialisation of $\det(E^\bullet)$ that exists on $\backslash\text{texts}\{\mathbf{E}\}$. Concretely it is the image of the obstruction/Kuranishi data under

$$\det(E^\bullet)|_{\text{near } \Sigma} \longrightarrow H^\bullet(\backslash\text{texts}\{\mathbf{E}\}) \quad \text{or} \quad H^\bullet(B, \backslash\text{texts}\{\mathbf{E}\}),$$

the class whose nonvanishing certifies that the concentrated structure on $\backslash\text{texts}\{\mathbf{E}\}$ cannot be continued across the spread on Σ .

Theorem 6 (Transgression, with type-determined arithmetic). *The constitutive invariant is locally constant on $\backslash\text{texts}\{\mathbf{E}\}$ and jumps across Σ by a quantity computed from κ on the enlarged $\mathcal{H}^1|_\Sigma$. Its arithmetic type is fixed by the type of the datum:*

- a primary* class, the integral class underlying D , present and integer-quantised when D is integral;*

- a secondary* class, the transgression of $\det(E^\bullet, D)$, which is a torsion/half-density when D is symmetric of odd shift, a \mathbb{Z} -valued degree when κ is winding-type, and a rank/signature when κ is a nondegenerate quadratic form.*

In each case the secondary class degenerates precisely on Σ , where $\det(E^\bullet)$ acquires zero modes ($\mathcal{H}^{\neq 0} \neq 0$).

Sketch. Local constancy on $\backslash \text{texts}\mathbf{E}$ is Theorem 3(1): no jump locus lies inside a component of $\backslash \text{texts}\mathbf{E}$. The jump is the contribution of the Kuranishi singularity on the normal model of Σ , whose degree/charge is the codimension- c winding of κ on the normal sphere (for $c = 1$, a sign per wall-crossing, a class in relative H^1 ; for $c = 3$ with κ the local Weyl map, a monopole degree in \mathbb{Z}). Quantisation of the primary class is integrality of D . The flavor of the secondary class is the flavor of κ . \square

The whole thesis is contained in Theorems 3 and 6: *the constitutive threshold is the amplitude-spread locus of the deformation complex; the standard invariant is the transgression of its obstruction; quantisation is winding-type Kuranishi.* The three instances below are this structure realised.

Instance A: flat connections (rigorous)

Let M be a closed oriented 3-manifold and G reductive with a nondegenerate invariant pairing \langle , \rangle on \mathfrak{g} . Consider the derived moduli stack of G -local systems,

$$\mathbf{R}\text{Loc}_G(M) = \mathbf{R}\text{Map}(M_B, BG).$$

The complex and the duality

The tangent complex at a flat connection A is the shifted twisted de Rham complex

$$\mathbb{T}_A \simeq C^\bullet(M; \text{ad}A)[1] : \quad \Omega^0 \xrightarrow{d_A} \Omega^1 \xrightarrow{d_A} \Omega^2 \xrightarrow{d_A} \Omega^3 \quad (\text{degrees } -1, 0, 1, 2),$$

so its cohomology is $H_A^0 = \text{Lie Aut}(A)$ (isotropy), H_A^1 (deformations), H_A^2 (obstructions),

$H_A^3 \cong (H_A^0)^*$. Flatness is precisely the condition $d_A^2 = [F_A, \cdot] = 0$ that closes the differential into a complex.

By Pantev–Toën–Vaquié–Vezzosi [1], mapping out of the Calabi–Yau 3-manifold M into the 2-shifted symplectic BG transgresses to a (-1) -shifted symplectic form

$$\omega = \int_M \langle -, - \rangle \quad \text{on } \mathbf{R}\mathrm{Loc}_G(M), \quad \omega: \mathbb{T}_A \otimes \mathbb{T}_A \rightarrow k[-1], \quad (a, b) \mapsto \int_M \langle a \wedge b \rangle.$$

This is the self-duality of Definition 1 with $n = -1$, and it is exactly Poincaré duality on M^3 : $H_A^i \cong (H_A^{3-i})^*$, the source of the bookkeeping $h^0 = h^3, h^1 = h^2, \chi = 0$. The orientation of M (its fundamental class) is the Calabi–Yau structure; that $\dim M = 3$ is what lands the shift at -1 , the value at which the pairing is symmetric on \mathbb{T} and the theory produces a number.

The stratification

By Goldman [3] the irreducible locus is smooth: there $H_A^0 = \mathfrak{z}(\mathfrak{g})$ and $H_A^2 \cong \mathfrak{z}^*$ are minimal, the complex is concentrated in degree 0, and $\mathbf{R}\mathrm{Loc}_G(M)$ is a smooth (symplectic) stack away from the central $BZ(G)$ gerbe. The jump stratum is the *reducible* locus, where the stabiliser jumps above the centre; by the (-1) -shifted self-duality the obstruction H_A^2 jumps in lockstep (Theorem 3(2)). Thus $\mathbf{E} = \text{irreducibles}$, $\mathbf{H} = \text{reducibles}$.

The transgression

The Chern–Simons functional $CS: \mathbf{R}\mathrm{Loc}_G(M) \rightarrow k/\mathbb{Z}$ is a Lagrangian (isotropic) structure for ω , and its underlying integral class—the generator of $H^4(BG; \mathbb{Z})$ transgressed—is the *quantised level* $k \in \mathbb{Z}$: the primary invariant of Theorem 6. In the large- k stationary-phase expansion of the Witten–Reshetikhin–Turaev invariant [4], each flat connection contributes with weight the Ray–Singer torsion [5]

$$\tau_A = \det'(\mathbb{T}_A, \omega)^{1/2},$$

the regularised determinant of the very complex above. On \mathbf{E} (amplitude in degree 0) this is an honest half-density and the contribution is clean; across Σ the zero

modes $H_A^1 \neq 0$ (and the enlarged stabiliser) force \det' to be reinterpreted as a section of $\det H_A^\bullet$, and the one-loop formula changes form. The torsion is therefore the secondary class of Theorem 6, degenerating precisely on the reducible stratum. Because $\kappa(a) = [a \wedge a]$ is quadratic, the secondary class is torsion-type while the primary class k is the integer.

Theorem 7 (Instance A). *On $\mathbf{RLoc}_G(M^3)$ the tangent complex $C^\bullet(M; \text{ad})[1]$ is self-dual under the (-1) -shifted symplectic form $\omega = \int_M \langle -, - \rangle; \mathbf{E}$ (concentrated amplitude) is the smooth symplectic locus of irreducible local systems and \mathbf{H} (spread amplitude) is the reducible locus, the spread being paired by ω ; the Chern–Simons functional is a Lagrangian structure for ω with quantised primary class $k \in \mathbb{Z}$, and the Ray–Singer torsion is the secondary class whose degeneration is supported on \mathbf{H} .*

Every ingredient is standard ([2, 3] for smoothness and the symplectic pairing; [1] for the shifted-symplectic structure; [4, 5] for the asymptotics). Instance A is thus a theorem, and is the case in which the heaviest available machinery genuinely applies.

Instance B: closed-chain linkages as polygon spaces

This is the finite-dimensional instance, and the one in which the Uniformity Conjecture (§6) is checkable by explicit matrix computation. The setup must be chosen correctly: the loop-closure condition is genuinely a moment map only when the edges are taken as vectors in \mathbb{R}^3 , i.e. as points of symplectic coadjoint orbits. (Treating planar closure as an \mathbb{R}^2 -valued constraint does *not* furnish the moment-map duality, and conflating the two leads to a spurious failure of Theorem 3(2); see Remark 8.)

Closure is a moment map

Take a closed n -bar linkage with edge lengths ℓ_1, \dots, ℓ_n , each edge a vector of fixed length in \mathbb{R}^3 , i.e. a point of the sphere $S_{\ell_i}^2$, a coadjoint orbit of $SU(2)$ and hence symplectic. The product $\prod_i S_{\ell_i}^2$ carries the diagonal $SO(3)$ action (rotating all edges together—the

rigid motion of the closed chain), and loop closure is its moment map:

$$\mu = \sum_{i=1}^n e_i : \prod_i S_{\ell_i}^2 \longrightarrow \mathbb{R}^3 \cong \mathfrak{so}(3)^*.$$

The configuration (shape) space is the symplectic reduction

$$Q = \mu^{-1}(0)/SO(3),$$

the *polygon space* of Kapovich–Millson [6] and Hausmann–Knutson [7]. This is the closure-induced structure: it is forced by the rigid-rotation symmetry of a closed chain, not added by hand.

The complex and the (exact) duality

At $q \in \mu^{-1}(0)$ the deformation complex of the reduction is the three-term Koszul/moment-map complex

$$E^\bullet : \underbrace{\mathfrak{so}(3)}_{\text{deg } -1} \xrightarrow{\rho_q} \underbrace{T_q(\prod S^2)}_{\text{deg } 0} \xrightarrow{d\mu_q} \underbrace{\mathfrak{so}(3)^*}_{\text{deg } +1},$$

where ρ_q is the infinitesimal action (the rigid-rotation vector fields) and $d\mu_q$ the closure Jacobian. Its cohomology is $\mathcal{H}^{-1} = \ker \rho_q$ (the residual stabiliser), $\mathcal{H}^0 = \ker d\mu_q / \text{im } \rho_q$ (reduced shape deformations, the tangent to Q), $\mathcal{H}^+ = \text{coker } d\mu_q$ (obstruction).

The moment-map identity $d\langle \mu, \xi \rangle = \iota_{\rho(\xi)} \omega$ says exactly that $d\mu_q = \rho_q^* \circ \omega^b$, where $\omega^b: T_q \rightarrow T_q^*$ is the (nondegenerate) symplectic isomorphism. Hence

$$\text{coker}(d\mu_q) = \mathfrak{so}(3)^* / \text{im}(\rho_q^* \omega^b) = \mathfrak{so}(3)^* / \text{im}(\rho_q^*) \cong (\ker \rho_q)^*,$$

the cokernel of an adjoint being the dual of the kernel. This is the self-duality of Definition 1 with $n = 0$: $\mathcal{H}^{+1} \cong (\mathcal{H}^{-1})^*$, with \mathcal{H}^0 carrying the reduced symplectic self-pairing $\mathcal{H}^0 \cong (\mathcal{H}^0)^*$. It is *exact*—it holds at every configuration—so $\dim \mathcal{H}^{+1} = \dim \mathcal{H}^{-1}$ identically, and Theorem 3(2) ($\Delta a^{-1} = \Delta a^{+1}$) is forced.

The stratification

A generic spatial polygon has trivial $SO(3)$ stabiliser, so $\ker \rho_q = 0$, hence (by the duality) $\text{coker } d\mu_q = 0$: the reduction is smooth and Q is a symplectic manifold there. The jump stratum is the *collinear* locus, where all edges are parallel: there the stabiliser jumps to $SO(2)$ (rotations about the common axis), $\mathcal{H}^{-1} = \ker \rho_q$ jumps from 0 to $\mathfrak{so}(2) \cong \mathbb{R}$, and the obstruction \mathcal{H}^{+1} jumps with it. By the theorem of [6] the singular points of Q are exactly these collinear configurations. Thus $\text{\textbf{E}} = \text{non-collinear polygons}$ (smooth symplectic shape space), $\text{\textbf{H}} = \text{collinear polygons}$ (stabiliser increase, singular). Since κ is a quadratic form (the second-order obstruction $d^2\mu$), the secondary invariant of Theorem 6 is rank/signature-type, classifying the local branch structure at the degenerate configuration. The Sjamaar–Lerman stratification of singular symplectic reduction [8] is the general statement of this picture.

Remark 8 (A correction, and the role of χ). A naive treatment takes a planar linkage with closure $\Phi(\theta) = \sum_i \ell_i e^{i\theta_i} \in \mathbb{C}$ and the diagonal phase $\rho = (1, \dots, 1)$, and posits the same duality. This fails: the structure group of the would-be moment map is 1-dimensional ($\mathfrak{u}(1)$) while Φ is \mathbb{R}^2 -valued, so the adjunction $d\mu = \rho^* \omega^b$ does not typecheck, and one finds (e.g. at the collinear configuration of a rhombus) that $\text{coker } d\mu$ jumps while $\ker \rho$ does not, an apparent violation of Theorem 3(2). The violation is an artifact of the wrong setup: planar closure as an \mathbb{R}^2 -valued constraint is not the moment map. What survives in that setup is only the Euler characteristic (the Maxwell–Calladine index [9]), consistent with Remark 4: χ is constant, but the term-by-term duality is absent because the duality was never genuinely present. The correct closure-induced structure is the spatial $SO(3)$ moment map above, for which the duality is exact and the paired jump holds.

Worked example: the paired jump and its signature

We verify the structure explicitly. At a collinear closed polygon $e_i = \epsilon_i \ell_i \hat{n}$ ($\epsilon_i = \pm 1$, $\sum_i \epsilon_i \ell_i = 0$) one has $\ker \rho_q = \{\omega : \omega \times e_i = 0 \forall i\} = \mathbb{R} \hat{n}$ and $\text{im}(d\mu_q)^\perp = \{\omega : \omega \parallel e_i \forall i\} = \mathbb{R} \hat{n}$, so \mathcal{H}^{-1} and \mathcal{H}^{+1} are each one-dimensional, against 0 generically. The Kuranishi map is the geodesic curvature of the constituent spheres: a tangent step v_i has geodesic acceleration $\ddot{e}_i = -|v_i|^2 \ell_i^{-2} e_i$, whence for

$v \in \ker d\mu_q$,

$$d^2\mu(v, v) = - \sum_i \frac{|v_i|^2}{\ell_i^2} e_i, \quad \kappa(v) = -\frac{1}{2} \left(\sum_i \frac{\epsilon_i}{\ell_i} |v_i|^2 \right) \hat{n} \in \text{coker } d\mu_q,$$

a diagonal quadratic form whose coefficient on each edge is set by the side ϵ_i of the axis on which that edge lies.

Computing rank ρ_q and rank $d\mu_q$ independently (so that the equality $\dim \mathcal{H}^{-1} = \dim \mathcal{H}^{+1}$ is a genuine check of the adjunction, not an identity built in) gives:

CONFIGURATION	$\hat{-1}$	$\hat{0}$	$\hat{+1}$	$=2N-6$
pentagon (1, 1, 1, 1, 2), generic	0	4	0	4
pentagon (1, 1, 1, 1, 2), collinear	1	6	1	4
rhombus (1, 1, 1, 1), generic	0	2	0	2
rhombus (1, 1, 1, 1), collinear	1	4	1	2

In each row $\dim \mathcal{H}^{-1} = \dim \mathcal{H}^{+1}$, and across the threshold the isotropy and obstruction jump together, $\Delta a^{-1} = \Delta a^{+1} = +1$, with χ constant and the deformation space \mathcal{H}^0 absorbing the slack ($\Delta a^0 = +2$): Theorem 3(2) realised in matrices. The Kuranishi form κ on \mathcal{H}^0 is *nondegenerate* in both cases—signature (2, 2) for the rhombus (eigenvalues $\{\pm \frac{1}{2}, \pm \frac{1}{2}\}$) and (4, 2) for the pentagon—so the singular configuration is a quadratic (Morse) cone point at which two smooth shape branches cross transversally, the bifurcation through the collinear locus. The local model is the cone over $\{\kappa = 0\}/SO(2)$, of dimension $2n - 5 - 1 = 2n - 6 = \dim Q$, matching the Sjamaar–Lerman description of the singular reduced space. The rhombus row is instructive: it is the same configuration whose paired jump *appears* to fail in the incorrect planar setup of Remark 8; in the correct spatial $SO(3)$ setup the duality holds exactly, confirming that the apparent failure was an artifact of the wrong arena.

Remark 9 (Status). With §4.4 the moment-map duality of Instance B is verified explicitly: the structure (closure = $SO(3)$ moment map; Q a polygon space) is established [6, 7], the duality $\mathcal{H}^{+1} \cong (\mathcal{H}^{-1})^*$ is an exact consequence of the adjunction, and the paired jump and Kuranishi signature are confirmed in matrices above. The contrast with a *non-polygon* linkage is sharp and instructive: for a Bennett $4R$ the closure is not a moment map, the deformation complex is two-term with mismatched ends, and no self-duality exists—see Remark 11, which uses this contrast to bound the central conjecture (§6.2).

Instance C: topological insulators

The complex and the duality

A gapped Bloch Hamiltonian $H(k)$ over the Brillouin torus T^d has a spectral projector onto the occupied bands,

$$P(k) = \frac{1}{2\pi i} \oint_{\Gamma} (z - H(k))^{-1} dz,$$

a smooth family of rank- m projectors whose image is the Bloch bundle $\mathcal{E} \rightarrow T^d$. The deformation differential is first-order perturbation theory,

$$L_H: \delta H \mapsto \delta P = \sum_{a \in \text{occ}, b \in \text{unocc}} \frac{\langle b | \delta H | a \rangle}{E_a - E_b} |a\rangle \langle b| + \text{h.c.},$$

whose energy denominators are the gap. The relevant duality is Poincaré–Lefschetz duality on T^d (in K -theory for the gapped class), pairing the curvature of \mathcal{E} against the fundamental class. The isotropy is the eigenframe symmetry: $U(1)^N$ generically (each band a line), jumping to a non-abelian $U(k)$ at a band touching where the occupied subspace cannot be separated.

The stratification and the transgression

The jump stratum Σ is the gap-closing locus $\{E_a = E_b\}$, where P is undefined and L_H singular (the denominators vanish) and the isotropy jumps to $U(k)$. Generically Σ is *codimension 3* (von Neumann–Wigner [21]: the local Weyl form $H_{\text{eff}} = \vec{d} \cdot \vec{\sigma}$ degenerates where $\vec{d} = 0$, three conditions), so the threshold is a point carrying a charge. The Kuranishi map is \vec{d} , and it is *winding-valued*: the local invariant is $\deg(\vec{d}/|\vec{d}|): S^2 \rightarrow S^2$, the Berry monopole charge.

The Chern number $C = \frac{1}{2\pi i} \int_{T^2} \text{tr}(P dP \wedge dP)$ is defined on \mathbf{E} , constant within each gapped phase, and *jumps across Σ by the monopole charge*—a standard theorem [22, 23, 24, 25]. This is exactly Theorem 6: the EXISTS-side invariant is the transgression of an obstruction concentrated at the HAPPENS stratum, the accumulated monopole charge of the touchings. Because the Kuranishi map is winding-valued rather than a quadratic form, the transgression is integer-quantised: this is why the topological-insulator invariant is a robust \mathbb{Z} where the linkage invariant (Instance B, quadratic κ) is only signature-valued. Bulk–boundary correspondence is then the statement that an interface between phases of different C must contain a gap-closing: it cannot interpolate the bundle invariant across the boundary without crossing Σ , so the edge mode is the real-space avatar of the HAPPENS locus.

Here the transgression principle is already a theorem; the contribution of the framework is to place the result alongside Instances A and B as the same structure with a winding-type Kuranishi map. That placement is at the level of the abstract pattern (Theorems 3–6) only: an explicit spectral test (§7.4, Remark 14) shows that C does *not* specialise from the PTVV construction that realises A and B, and is related to them by analogy rather than base change.

Synthesis

The dictionary

The three instances are one structure realised at three types:

FLAT CONNECTIONS (A)** **LINKAG	E (B)**	INSULATOR (C)	
deformation diff.	d_A	$d\mu_q$ (and ρ_q)	$L_H : \delta H \mapsto$
isotropy $\mathcal{H}^{\leq 0}$	gauge stabiliser	residual $SO(3)$ stab.	eigenframe $U(1)^N \rightarrow U$
obstruction	$H_A^2 \cong (H_A^1)^*$	$\text{coker } d\mu \cong (\ker \rho)^*$	Berry monop
Kuranishi κ	$[a \wedge a]$ (quad.)	$d^2\mu$ (quad.)	\vec{d} (winding)
duality shift n	-1 (shifted symp.)	0 (reduction)	PD on T^d (K th.)
threshold Σ	reducibles	collinear	gap-closing (codim 3)
invariant	WRT (level + torsion)	mobility / signature	Chern number $\in \mathbb{Z}$
status	theorem	moment-map exact; verified (§4.4) jump = thm	

The two differences across the row are *explained*, not papered over. The arithmetic of the invariant (\mathbb{Z} -quantised vs. torsion vs. signature) is the arithmetic of (D, κ) ; the codimension of Σ (a wall vs. a charged point) is the codimension of the Kuranishi zero locus.

The central conjecture, and its boundary

The uniform content—Theorems 3 and 6—is proved from Definition 1. What is *not* proved is that the self-duality D is structural rather than observed case by case. The three instances are suggestive, but they share a feature more specific than “a closure condition”: in each, the closure equips the configuration space with a self-dual deformation theory of a recognisable type—a moment map on a symplectic space for A (flatness as the gauge moment map, PTVV) and B (the $SO(3)$ moment map on $\coprod S^2$), and a Poincaré–Lefschetz

duality on a closed manifold for \mathbb{C} (the Brillouin torus). The computation of §4.4 and the counterexample below show that this specificity is necessary, not incidental.

Conjecture 10 (Uniformity, sharpened). *A constitutive threshold carries a canonical self-duality D on its deformation complex—hence (Theorem 3) paired isotropy/obstruction jumps and (Theorem 6) a transgression invariant whose arithmetic is read off the type—when the closure equips the configuration space with a self-dual deformation theory: a moment map for the symmetry acting on a symplectic space (Instances A, B), or a Poincaré/Poincaré–Lefschetz duality on a closed configuration manifold (Instance C). The shift and symmetry of D are then determined by the dimension and orientation data of that structure. The hypothesis is essential: an arbitrary closure constraint does not suffice (Remark 11).*

Remark 11 (The Bennett counterexample: closure alone is not enough). The naive form of the conjecture—“any loop-closure induces D ”—is false, and the Bennett $4R$ linkage is a clean finite-dimensional counterexample. For a closed spatial linkage the deformation complex in screw coordinates is the *two-term* map

$$J: \mathbb{R}^{\#\text{joints}} \longrightarrow \mathfrak{se}(3) \cong \mathbb{R}^6, \quad J = [\$_1 \cdots \$_{\#\text{joints}}],$$

with $\mathcal{H}^0 = \ker J$ the first-order mobility and $\mathcal{H}^1 = \text{coker } J$ the self-stresses (reciprocal wrenches). For a $4R$ loop the ends are \mathbb{R}^4 and \mathbb{R}^6 : a self-duality $E^\bullet \cong (E^\bullet)^\vee[n]$ would require equal end-dimensions, so none exists—the dimension mismatch $6 \neq 4$ is the overconstraint that makes a generic spatial $4R$ rigid. Explicit computation (supplementary script; $\alpha = 60^\circ, \beta = 90^\circ, a = 1, b = \sin \beta / \sin \alpha$) confirms the picture: the Bennett linkage closes over a continuum (genuine 1-DOF finite mobility, κ -obstruction vanishing), and at a closed configuration the joint screws have $\text{rank } J = 3$, giving $(\dim \mathcal{H}^0, \dim \mathcal{H}^1) = (1, 3)$; a generic spatial $4R$ has $\text{rank } J = 4$ and $(0, 2)$. The numerical $\text{rank } J = 3$ has a geometric reason that holds throughout the motion, not merely at the tested configuration: the four Bennett axes lie on a common regulus (a ruling family of a one-sheeted hyperboloid, signature $(2, 1)$) at every configuration [10, 14], and three lines on a regulus force the fourth to be screw-dependent, so $\text{rank } J = 3$ persists for the whole cycle and the mobility is genuinely finite. In both, $\chi = \dim \mathcal{H}^0 - \dim \mathcal{H}^1 = -2$ (the Maxwell–Calladine / Grübler index $M = n - 6$ [9, 10]), so the Bennett threshold *does* show a paired jump $\Delta a^0 = \Delta a^1 = +1$ —but for a two-term complex this is forced by χ

alone (Remark 4), *not* by a self-duality, and indeed $\dim \mathcal{H}^0 \neq \dim \mathcal{H}^1$ at every configuration (they differ by 2). Contrast Instance B’s polygon, where $\dim \mathcal{H}^{-1} = \dim \mathcal{H}^{+1}$ exactly (§4.4): there the duality is genuine. The reciprocal (Klein) form on $\mathfrak{se}(3)$, of signature $(3, 3)$, pairs $\text{im } J$ with $\text{coker } J$ inside $\mathfrak{se}(3)$ but does not make the mismatched complex self-dual. The lesson is exactly the sharpening above: the operative hypothesis is “closure = moment map / Poincaré duality,” not “closure.”

Remark 12 ($\chi = 0$ is necessary but not sufficient). One might hope the dividing line is simply the balanced complex $\chi = 0$, since a self-duality $\mathcal{H}^0 \cong (\mathcal{H}^1)^*$ forces $\dim \mathcal{H}^0 = \dim \mathcal{H}^1$. It is not. For a closed jR loop $\chi = j - 6$, so the balanced case is the $6R$ loop, and there $\dim \mathcal{H}^0 = \dim \mathcal{H}^1$ holds at every configuration (generic $6R$: $(0, 0)$, rigid; a mobile overconstrained $6R$ such as a Sarrus or Bricard linkage: $(1, 1)$). The necessary dimensional balance is thus met. But no self-duality is induced: the only closure-given data are J and the Klein form on $\mathfrak{se}(3)$, the canonical pairing they build, $(m, \sigma) \mapsto \langle Jm, \sigma \rangle_K$, vanishes identically on $\ker J \times \text{coker } J$ because $Jm = 0$, and there is no other closure-canonical bridge from the joint domain to $\mathfrak{se}(3)$. The genuine relation between mobility and self-stress is the *second-order* (Kuranishi / prestress) quadratic form $\sigma \cdot d^2(\text{closure})(m, m)$, not a linear self-duality of the complex; this is precisely the *force-weighted Hessian* $\sum_i s_i \nabla^2 f_i$ restricted to $\ker J$ of the quasi-variant/shaky theory of mechanism rigidity [13, 11, 12, 10], whose definiteness certifies second-order (quasi-variant) rigidity. A concrete family makes this vivid: the regular n -gon linkages of [10] are *hypo-paradoxical*— $\dim \ker J \neq 0$, so the classical count predicts motion, yet $\dim C = 0$ (rigid)—exactly because this Hessian is definite. There $\mathcal{H}^0 = \ker J$ overcounts the actual (zero) mobility and the Kuranishi form does the cutting, which is the general mechanism by which \mathcal{H}^0 is only the first-order tangent. Structurally, a mobile $6R$ has a 1-dimensional (odd) configuration curve, which cannot be symplectic, whereas the polygon space is even-dimensional and symplectic; this is why the polygon carries the moment-map duality and the balanced linkage does not. Hence *the moment-map/symplectic hypothesis of Conjecture 10 is strictly stronger than $\chi = 0$* : balance is necessary, moment-map closure on a symplectic configuration space is the operative sufficient condition. More broadly, the established fact that linkage mobility is genuinely second-order—the Jacobian is insufficient and the Hessian decides [11, 12, 10]—is independent confirmation, from mechanism theory, that general spatial linkages live in the nonlinear Kuranishi layer and

therefore lie outside the linear self-duality of Lemma 13; only the moment-map polygon subclass (§4.4), where closure is first-order and the configuration space symplectic, is self-dual.

This is the precise sense of “topology appears where closure generates thresholds”: a closure of moment-map/Poincaré type produces the self-duality D , and D is what makes the threshold topological in the strong (paired-jump, transgression) sense; closures without that structure—even balanced ones with $\chi = 0$ —produce only the weaker Euler-characteristic bookkeeping.

Interpretation

We record an interpretive reading, kept separate from the mathematics. Theorem 3(2) says that across Σ the isotropy $\mathcal{H}^{\leq 0}$ changes—the automorphisms of the object, hence the equivalence relation by which two configurations count as “the same”, genuinely differs on the two strata, and B ceases to be a manifold and becomes a stack. In the vocabulary of Fine [29] this is an *essential* rather than accidental difference: the spread of the complex is constitutive, not a removable property. In the vocabulary of Wiggins [30], different sortal concepts carry different criteria of identity, and a jump in $\mathcal{H}^{\leq 0}$ is exactly such a change of identity-criterion. This is the one place the philosophy does non-ornamental work: it licenses calling the two strata different *kinds* rather than two regions of one space, by pointing at the precise mathematical object—the jump in the isotropy—rather than at a slogan. We make no stronger metaphysical claim.

Toward a base-change formulation

We now give the categorical backbone promised by the framework: an abstract duality-equipped object, the structural lemma in its sharpest form, and the sense in which the genuine instances are specialisations of a single construction. Two of the three (A, B) are rigorously such specialisations of established constructions of Pantev–Toën–Vaquié–Vezzosi [1]; the third (C) is the parallel statement in a topological setting, which we state

honestly as requiring the broader Calabi–Yau-categorical framework. The point of the section is to show that “EXISTS/HAPPENS = amplitude of a self-dual complex” is not three coincidences but the truncation of one structure, and to locate precisely the part that is a theorem.

The abstract object and the structural lemma

Let B be a derived stack and $E \in \text{Perf}(B)$ a perfect complex with a *self-duality of shift n* : a quasi-isomorphism $D: E \xrightarrow{\sim} E^\vee[n]$, graded-symmetric or graded-antisymmetric. The case of interest is $E = \mathbb{T}_B$ with D underlying an n -shifted symplectic structure $\omega \in \mathcal{A}^{2,\text{cl}}(B, n)$ in the sense of [1]; then $D = \omega^\flat$ is the induced $\mathbb{T}_B \simeq \mathbb{L}_B[n]$. The amplitude stratification (Definition 2) is the rank stratification of E : the loci $\{b : \dim \mathcal{H}^i(E_b) \geq r\}$ are closed (upper-semicontinuity of fibrewise cohomology rank for a perfect complex), so $\text{amplitude}(E)$ (concentrated amplitude) is open and $\Sigma = \text{amplitude}(E)$ closed.

Lemma 13 (Self-duality forces paired ranks). *For a perfect complex E on B with a shift- n self-duality $D: E \simeq E^\vee[n]$, the fibrewise cohomology ranks satisfy*

$$\dim \mathcal{H}^i(E_b) = \dim \mathcal{H}^{-n-i}(E_b) \quad \text{for all } b \in B, i \in \mathbb{Z},$$

and hence across any stratum Δ $\dim \mathcal{H}^i = \Delta \dim \mathcal{H}^{-n-i}$.

Proof. Fix b and work over the residue field. The quasi-isomorphism D gives $\mathcal{H}^i(E_b) \cong \mathcal{H}^i(E_b^\vee[n]) = \mathcal{H}^{i+n}(E_b^\vee)$. Over a field $\text{Hom}(-, k)$ is exact, so $\mathcal{H}^k(E_b^\vee) \cong (\mathcal{H}^{-k}(E_b))^*$; with $k = i + n$ this is $(\mathcal{H}^{-n-i}(E_b))^*$. Thus $\mathcal{H}^i(E_b) \cong (\mathcal{H}^{-n-i}(E_b))^*$, and dimensions agree. Differencing across a stratum gives the paired jump. \square

This is the precise content of Theorem 3(2), and it isolates exactly what the polygon has and the Bennett/6R linkage lacks. The polygon complex is genuinely self-dual ($n = 0$), so $\dim \mathcal{H}^{-1} = \dim \mathcal{H}^{+1}$ (§4.4); the linkage two-term complex $\mathbb{R}^j \rightarrow \mathfrak{st}(3)$ carries no self-duality (mismatched ends, Remark 11), so the equality $\dim \mathcal{H}^0 = \dim \mathcal{H}^1$ that holds at $\chi = 0$ is *not* an instance of Lemma 13—it is the unrelated coincidence of a two-term Euler

characteristic (Remark 12). Lemma 13 is therefore the formal discriminator between the two: the paired jump is a self-duality phenomenon iff it is the equality of \mathcal{H}^i with \mathcal{H}^{-n-i} , not merely of the two ends of a short complex.

The universal construction

The self-dualities of the genuine instances are not postulated; they are produced by a single mechanism—transgression of a pairing on a target against an orientation of a source. The two relevant theorems of [1] are:

(*Mapping stacks.*) If X is \mathcal{O} -compact and d -oriented (a fundamental class $\int_X: \mathbf{R}\Gamma(X, \mathcal{O}_X) \rightarrow k[-d]$) and Y is m -shifted symplectic, then $\mathbf{Map}(X, Y)$ is $(m - d)$ -shifted symplectic, with form $\int_X \mathrm{ev}^* \omega_Y$.

(*Lagrangian intersection.*) If $L_1, L_2 \rightarrow Z$ are Lagrangian morphisms to an m -shifted symplectic Z , then $L_1 \times_Z^h L_2$ is $(m - 1)$ -shifted symplectic.

Both are instances of “orientation \times pairing \rightarrow self-duality,” and the shift of the resulting self-duality is (target pairing degree) minus (source orientation dimension). The deepest common source is a *Calabi–Yau structure*: a d -CY structure on a smooth proper dg-category \mathcal{C} induces a $(2 - d)$ -shifted symplectic structure on its moduli of objects $\mathcal{M}_{\mathcal{C}}$ [18], of which the mapping-stack and Lagrangian-intersection statements are special and relative cases.

Instances A and B as specialisations (rigorous)

A (flat connections). Take $X = M_{\mathbf{B}}$, the Betti stack of a closed oriented 3-manifold (so $d = 3$), and $Y = BG$, which is 2-shifted symplectic from the invariant pairing $\langle \cdot, \cdot \rangle$ on \mathfrak{g} [1]. Then

$$\mathbf{R}\mathrm{Loc}_G(M) = \mathbf{Map}(M_{\mathbf{B}}, BG) \quad \text{is} \quad (2 - 3) = -1\text{-shifted symplectic,}$$

and the resulting D is exactly the Poincaré-duality pairing $\int_M \langle -, - \rangle$ of §3. This is verbatim a PTWV example; the 3-CY structure on $\mathrm{Loc}(M)$ is the orientation of M .

B (polygon spaces). The shifted cotangent $T^*[1]BG = [\mathfrak{g}^*/G]$ is 1-shifted symplectic. For $M = \prod S^2_{\ell_i}$ symplectic with the $SO(3)$ moment map μ , the descended map $[\mu]: [M/G] \rightarrow [\mathfrak{g}^*/G]$ is a Lagrangian morphism, and so is the zero section $BG \rightarrow [\mathfrak{g}^*/G]$ (this is the derived content of “ μ is a moment map”; see [16, 17]). Hence the polygon space is the Lagrangian intersection

$$Q = [\mu^{-1}(0)/G] = [M/G] \times_{[\mathfrak{g}^*/G]}^h BG \quad \text{is} \quad (1 - 1) = 0\text{-shifted symplectic,}$$

recovering the $n = 0$ self-duality $\mathcal{H}^{+1} \cong (\mathcal{H}^{-1})^*$ of §4.4. Thus A and B are the same construction with (source, target) = (M^3, BG) and (reduction of $\prod S^2, [\mathfrak{g}^*/G]$) respectively; the difference -1 versus 0 is exactly the difference in the input dimensions.

Instance C does not specialise from the construction (a spectral test)

For C the source is the Brillouin torus T^d and the duality is *Poincaré–Lefschetz duality in topological K-theory*: T^d is framed, hence KU -oriented [20], $KU^0(T^d) \cong KU_0(T^d)$ by cap with the fundamental class, and the Chern number is the pushforward

$\int_{T^d}: KU^0(T^d) \rightarrow \mathbb{Z}$. This has the same *shape* as §7.2—transgression against the fundamental class of an oriented source—and the moduli of gapped Hamiltonians is, like A, a mapping space (into the classifying space of the symmetry class, e.g.

$\mathbb{Z} \times BU \simeq \Omega^\infty KU$ for class A). One might therefore hope to subsume C by a topological (spectral) analogue of §7.2—Calabi–Yau structures on KU -linear categories transgressing a K -orientation. An explicit test shows the resemblance does not upgrade to a base-change specialisation, and the obstruction is structural rather than a missing construction.

Remark 14 (The spectral test: C is analogy, not specialisation). Two facts decide it (supplementary script). (i) *The target is not shifted-symplectic.* The PTVV mapping-stack theorem produces a symplectic form by transgressing a symplectic form on the target BG ; the TI target $\Omega^\infty KU$ is a spectrum, not a shifted-symplectic derived algebraic stack, and the TI transgression produces a K -theory class, not a symplectic structure. The constructions rhyme, but their target data—a symplectic form versus a K -orientation—are different kinds of object, and only the former is PTVV. (ii) *On the torus, class A is only a rational shadow.* $KU^*(T^d)$ is torsion-free (the Atiyah–Hirzebruch spectral sequence collapses on $\Lambda^*\mathbb{Z}^d$), so the Chern character is rationally injective with full-rank image and

the class-A invariant is determined by its Chern numbers. The apparent agreement of C with an algebraic story is therefore the degenerate torsion-free case and is *no evidence* for a genuinely spectral self-duality. The genuinely K -theoretic content lives in the *real* symmetry classes: an explicit class-All (time-reversal) model has total Chern number 0 in *both* its trivial and nontrivial phases, while a \mathbb{Z}_2 invariant—a KR/KO 2-torsion class, here the parity of a partial Chern number [26, 27, 28]—distinguishes them. That invariant is invisible to the Chern character and arises from an antiunitary symmetry with no counterpart in complex shifted-symplectic geometry. Hence C realises the abstract pattern of Theorems 3–6 but is *not* a specialisation of the PTVV construction of A and B; a spectral or Real (KR/KO) K -theoretic unification, if one exists, would be a genuinely different framework, not a base change of §7.

The transgression invariant, abstractly

At this level the transgression invariant of Definition 5 is the obstruction-theoretic class of the determinant line $\det E$ with the self-dual (orientation) structure that D induces. On \mathbf{E} , where E is concentrated, $\det E$ has a canonical self-dual trivialisation; its failure to extend across Σ , where $\mathcal{H}^{\neq 0} \neq 0$, is the secondary class. For (-1) -shifted symplectic structures this self-dual trivialisation is precisely *orientation data* in the sense of Kontsevich–Soibelman [19], and its anomaly is the source of the Ray–Singer torsion factor in A. The arithmetic type of the secondary class is, as in Theorem 6, fixed by the duality and Kuranishi types: torsion (A, symmetric D , quadratic κ), signature (B, quadratic κ), \mathbb{Z} -degree (C class A, winding κ). A single secondary class covering all three would have to bridge the k -linear determinant/orientation theory of A and B and the topological- K -theory setting of C, which Remark 14 shows are not related by base change; in the real symmetry classes the secondary class is KR/KO 2-torsion, outside the k -linear theory entirely. So the determinant/orientation account is the right general statement for A and B, and C’s is its analogue in a different coefficient world, not a specialisation of it.

Status of this section

Theorem (rigorous): Lemma 13; the realisations of A as $\mathbf{Map}(M_B, BG)$ and of B as a Lagrangian intersection in $[\mathfrak{g}^*/G]$, hence both as specialisations of the PTVV transgression/intersection construction with their self-dualities determined by source

dimension and target pairing [1, 16, 17]. *Shown negative:* C does *not* specialise from this construction (Remark 14)—its target is a spectrum rather than a shifted-symplectic stack, on the torsion-free torus its class-A invariant is only the rational shadow, and its genuinely K -theoretic (real-class) invariants are KR/KO torsion with no shifted-symplectic counterpart. So the honest backbone is: A and B are one PTVV construction, rigorously; C realises the abstract pattern of Theorems 3–6 but is related to A and B by the analogy “oriented transgression of a pairing,” not by base change. The single mechanism unifies two of the three; the third rhymes in a different coefficient world.

Scope: what does not fit

The framework is not universal, and it is worth stating its boundary explicitly, because the distinction between “a threshold with a deformation-complex self-duality” and “a regular value of a scalar functional” is exactly what the framework turns on.

Consider a proposed causal-structure instance: X a space of globally hyperbolic Lorentzian metrics (modulo diffeomorphism), \mathcal{D} a smooth real functional (e.g. a curvature rate), and $\Sigma = \{\mathcal{D} = c\}$ a level set at a regular value c . This is a genuine stratification— Σ separates X into sub- and super-threshold regions—but it is *not* an instance of Definition 1. There is no deformation complex with a self-duality at the threshold; the threshold is the regular level set of a scalar, and the only invariant it carries is the wall-crossing class of a co-oriented hypersurface. That class is the Thom class of a codimension-1 submanifold, hence lives in relative $H^1(X, X_{<c}) \cong H^0(\Sigma)$ —it is the “which side of the wall” class, of degree 1, evaluated on *paths* by their intersection number with Σ . It is therefore not a degree-2 gerbe, not integer-quantised in any nontrivial sense, and carries none of the structure (isotropy jump, paired obstruction, transgression) of Theorems 3–6. A numerical normalisation attached to such a class (a fixed real period) is a coefficient, not a complex and not a duality, and does not change its degree or supply quantisation. We therefore exclude this case from the unification. It would enter only if a genuine deformation complex with a self-duality were exhibited at the metric threshold—which we do not have.

Remark 15 (Compatibility with relative cohomology and the $4\pi^2$ gerbe). This exclusion applies to *absolute* cohomology: the threshold $\Sigma = \{\mathcal{D} = c\}$ carries a degree-1 class in $H^1(\Sigma; \mathbb{Z})$, not a degree-2 gerbe in $H^2(X; \mathbb{Z})$. This is compatible with the treatment in [31, 32], where the same class appears as a degree-2 Thom class $\tau(\Sigma_X) \in H^2(X, X_{\text{sub}}; \mathbb{Z})$ via the Thom isomorphism $H^2(X, X \setminus \Sigma; \mathbb{Z}) \cong H^1(\Sigma; \mathbb{Z})$. The two descriptions are the same class in different languages: absolute degree-1 on Σ and relative degree-2 on the pair (X, X_{sub}) . The real period $4\pi^2$ of this relative class is determined not by the level-set structure of \mathcal{D} but by the Compton-Hopf identification (§9(4); [31] §7): the winding torus $T_\gamma^2 \subset \mathbb{P}\mathbb{N}$ has symplectic area $\int_{T_\gamma^2} \omega_R \wedge \omega_A = 4\pi^2$, and this class is pulled back to the configuration space via the diffeomorphism $\varphi: T_\gamma^2 \rightarrow T_{\text{winding}}^2$. The \mathbb{R} -gerbe of [32] lives on T_γ^2 ; it arrives on X_{sup} via φ^* , not from the level-set structure of \mathcal{D} alone.

This exclusion is part of the content: the framework identifies which thresholds are “topological” in the operative sense (those carrying a self-dual deformation complex) and separates them from thresholds that are merely separations of a space by a scalar. We note that this exclusion applies to the curvature-rate case specifically; the causal *contact* threshold $y \prec x$ —a different mathematical object carrying a Čech cocycle and an integer winding class rather than a level-set class—is not excluded by this argument, and whether it assembles into a perfect complex with self-duality remains an open question recorded in §9(4).

Open problems and next computations

In order of priority and tractability:

1. *The boundary of the conjecture (settled in the linkage family)*. The closed-chain polygon case is verified (§4.4); the general spatial linkage is understood: loop-closure alone does not induce a self-duality (Bennett $4R$, Remark 11), and the natural weaker guess fails too— $\chi = 0$ (the balanced $6R$ loop) supplies the necessary dimensional balance but not the duality (Remark 12). So within mechanisms the moment-map/symplectic hypothesis is exactly the right one and is strictly stronger than $\chi = 0$. The open

direction is now categorical rather than computational: an intrinsic characterisation of which closures are of moment-map/Poincaré type—e.g. whether a (-2) -shifted symplectic structure [15] ever upgrades an overconstrained mechanism, which we doubt given the odd-dimensional configuration curves involved.

2. *Instance C and Real K -theory (resolved here; new direction)*. The spectral test (Remark 14) shows C does not specialise from the PTVV construction: its target is a spectrum, not a shifted-symplectic stack. The open direction this opens is not “fit C into PTVV” but its opposite—whether a *Real* (KR/KO) K -theoretic analogue of the amplitude/transgression framework governs the symmetry-class \mathbb{Z}_2 invariants [26, 27, 28], which the present (k -linear, complex) framework cannot see. That would be a parallel theory in a different coefficient world, sharing only the “oriented transgression” shape.

3. *The uniform theorem (derived form): the honest two of three*. Section 7 delivers the backbone: Lemma 13 (self-duality \Rightarrow paired ranks) is the abstract Theorem 3(2), and Instances A and B are realised as specialisations of one construction— $\mathbf{Map}(M_B, BG)$ and a Lagrangian intersection in $[\mathfrak{g}^*/G]$, both PTVV [1, 16, 17]. Instance C is *not* subsumed (Remark 14); it realises the abstract pattern but in topological K -theory, related by analogy only. So the uniform theorem is genuinely a theorem for the two k -linear instances; whether a single framework can also contain the spectral instance C is the deep open question, and the evidence here is that it would require leaving k -linear shifted-symplectic geometry altogether.

4. *The causal contact threshold: excluded on the wrong grounds, genuinely open*.

Section 8 excludes a causal-structure instance on the grounds that a curvature-rate threshold $\{\mathcal{D} = c\}$ is a regular level set of a scalar, carrying only a wall-crossing class of degree 1 and no deformation-complex self-duality. That exclusion is correct for the curvature-rate case. It does *not* apply to the causal contact threshold $y \prec x$, which is a different mathematical object: it carries an explicit Čech cocycle

$f_{01} = \langle \lambda, \alpha \rangle^{-1} \cdot (Z^A Z_A)^{-1} \in H^1(\mathbb{P}T^+, \mathcal{O}(-2))$, an integer winding class in $H^1(T_\gamma^2, \mathbb{Z})$, and a cup-product pairing $\int_{T_\gamma^2} \omega_R \wedge \omega_A = 4\pi^2$ whose topological origin is proved in [31]. These are the ingredients of the deformation-complex framework.

What has *not* been established is whether they assemble into a perfect complex over a

space of causal configurations with a self-duality D in the sense of Definition 1—whether the retarded Green function cocycle sits in a $\mathbb{P}T^+/\mathbb{P}T^-$ -induced self-dual complex on $\mathbb{P}N$ (the null twistors), analogously to Instance B where the polygon complex is induced by the two Lagrangians meeting at their intersection. This is genuinely open and has not been attempted. A positive resolution would establish Instance D and reopen the scope of §8 accordingly.

Acknowledgements

The formalisation in this paper—the deformation-complex framework of §2, the dictionary of §6, and the corrections recorded in Remarks 4 and 8—was developed and stress-tested in extended dialogue with an AI assistant. All claims, and responsibility for any errors, remain the author’s. The framework was arrived at by repeatedly demanding that each claim close on its own equations; several earlier formulations (including a residue/restriction theorem and a relative-cohomology “gerbe” identification, and the planar-linkage setup of Remark 8) were discarded in that process, and the present scope—three instances, an explicit exclusion, and a bounded conjecture—reflects what survived. The matrix computations of §4.4 (the polygon paired jump and Kuranishi signature), Remark 11 (the Bennett mobility, screw ranks, and self-duality test), and Remark 12 (the balanced $6R$ test) are reproduced by three short, self-contained supplementary scripts.

References

1. T. Pantev, B. Toën, M. Vaquié, G. Vezzosi, *Shifted symplectic structures*, Publ. Math. IHÉS **117** (2013) 271–328.
2. M. F. Atiyah, R. Bott, *The Yang–Mills equations over Riemann surfaces*, Philos. Trans. Roy. Soc. London Ser. A **308** (1983) 523–615.

3. W. M. Goldman, *The symplectic nature of fundamental groups of surfaces*, Adv. Math. **54** (1984) 200–225.
4. E. Witten, *Quantum field theory and the Jones polynomial*, Comm. Math. Phys. **121** (1989) 351–399.
5. D. B. Ray, I. M. Singer, *R-torsion and the Laplacian on Riemannian manifolds*, Adv. Math. **7** (1971) 145–210.
6. M. Kapovich, J. J. Millson, *The symplectic geometry of polygons in Euclidean space*, J. Differential Geom. **44** (1996) 479–513.
7. J.-C. Hausmann, A. Knutson, *The cohomology ring of polygon spaces*, Ann. Inst. Fourier (Grenoble) **48** (1998) 281–321.
8. R. Sjamaar, E. Lerman, *Stratified symplectic spaces and reduction*, Ann. of Math. (2) **134** (1991) 375–422.
9. C. R. Calladine, *Buckminster Fuller's "tensegrity" structures and Clerk Maxwell's rules for the construction of stiff frames*, Internat. J. Solids Structures **14** (1978) 161–172.
10. N. Shvalb, O. Medina, *The geometry of infinitesimal mobility of closed-loop linkages*, Mech. Mach. Theory **220** (2026) 106332.
11. E. N. Kuznetsov, *Underconstrained Structural Systems*, Springer, 2012.
12. Z. Li, A. Müller, *Mechanism singularities and shakiness from an algebraic viewpoint*, Mech. Mach. Theory **191** (2024) 105510.
13. R. Connelly, *Generic global rigidity*, Discrete Comput. Geom. **33** (2005) 549–563.
14. J. E. Baker, *The Bennett linkage and its associated quadric surfaces*, Mech. Mach. Theory **23** (1988) 147–156.
15. D. Borisov, D. Joyce, *Virtual fundamental classes for moduli spaces of sheaves on Calabi–Yau four-folds*, Geom. Topol. **21** (2017) 3231–3311.
16. D. Calaque, *Lagrangian structures on mapping stacks and semi-classical TFTs*, in *Stacks and Categories in Geometry, Topology, and Algebra*, Contemp. Math. **643** (2015) 1–23.
17. P. Safronov, *Quasi-Hamiltonian reduction via classical Chern–Simons theory*, Adv. Math. **287** (2016) 733–773.

18. C. Brav, T. Dyckerhoff, *Relative Calabi–Yau structures*, *Compos. Math.* **155** (2019) 372–412.
19. M. Kontsevich, Y. Soibelman, *Stability structures, motivic Donaldson–Thomas invariants and cluster transformations*, arXiv:0811.2435 (2008).
20. M. F. Atiyah, R. Bott, A. Shapiro, *Clifford modules*, *Topology* **3** (1964) 3–38.
21. J. von Neumann, E. Wigner, *Über das Verhalten von Eigenwerten bei adiabatischen Prozessen*, *Phys. Z.* **30** (1929) 467–470.
22. B. Simon, *Holonomy, the quantum adiabatic theorem, and Berry's phase*, *Phys. Rev. Lett.* **51** (1983) 2167–2170.
23. D. J. Thouless, M. Kohmoto, M. P. Nightingale, M. den Nijs, *Quantized Hall conductance in a two-dimensional periodic potential*, *Phys. Rev. Lett.* **49** (1982) 405–408.
24. X.-L. Qi, S.-C. Zhang, *Topological insulators and superconductors*, *Rev. Mod. Phys.* **83** (2011) 1057–1110.
25. M. Z. Hasan, C. L. Kane, *Colloquium: Topological insulators*, *Rev. Mod. Phys.* **82** (2010) 3045–3067.
26. C. L. Kane, E. J. Mele, *\mathbb{Z}_2 topological order and the quantum spin Hall effect*, *Phys. Rev. Lett.* **95** (2005) 146802.
27. A. Kitaev, *Periodic table for topological insulators and superconductors*, *AIP Conf. Proc.* **1134** (2009) 22–30.
28. D. S. Freed, G. W. Moore, *Twisted equivariant matter*, *Ann. Henri Poincaré* **14** (2013) 1927–2023.
29. K. Fine, *Essence and modality*, *Philos. Perspectives* **8** (1994) 1–16.
30. D. Wiggins, *Sameness and Substance Renewed*, Cambridge University Press, 2001.
31. Z. Paz, *Topological closure of the null cone and the $4\pi^2$ causal pairing*, STF Framework Paper, V6.3 (2026).
32. Z. Paz, *Topological Invariants as Signatures of Constitutive Thresholds: From Closed-Loop Linkages to Topological Insulators and Chern-Simons Theory*, STF Framework Paper, V10 (2026).

CITATION

```
@article{paz2026constitutive,  
  author = {Paz, Z.},  
  title = {Constitutive Thresholds as Amplitude Strata of Deformation  
Complexes},  
  year = {2026},  
  version = {V1.0},  
  url = {https://existshappens.com/papers/constitutive-thresholds/}  
}}
```

[← Topological Invariants](#)

[All Papers](#)

[Topological Closure →](#)

EXISTS | HAPPENS The Selective Transient Field Framework

[Framework](#) [Papers](#) [Predictions](#) [News](#) [ATLAS](#) [About](#)

© 2026 Z. Paz · existshappens.com