Supplemental Material to “Finite spectral windows with an opening tower in seven-dimensional parents”

Draft, 3 November 2026

This is the supplement of the paper *Finite spectral windows with an opening tower in seven-dimensional parents*. Its parts keep the numbers of the full Supplemental Material, so that a citation of a part in the paper refers to the part of the same name here.

Parts contained. Lemma parts: 24, 28, 29, 30.

The full Supplemental Material, with all parts of the project, exists as `supplemental.pdf` and `supplemental_mobile.pdf` in the same folder as this file. A reference inside a part to a part not listed here points to the full version.

Contents

Part 24 **2**

Part 28 **22**

Part 29 **38**

Part 30 **56**

Part 24

Commutative cylinders, spectral walls, and the two-stage compression

By Astra 6 (gpt-6-astra, effort xhigh), 3 Nov 2026; prompt in prompts/24_commutative_parent.md. Reproduced as delivered, with links made relative. The section "Numerical check" at the end is by Claude Opus 5.5. Refereed by Terra 6.1 (effort max, 3 Nov 2026; report): no error in the displayed equations, and one gap in the general statement on compression maps. The record above the section "Referee amendments" stays as delivered, and those amendments apply to it.

A commutative cylinder admits an exact realization of the **operator window** of lemmas 22 and 23 when the two finite gradings are prescribed through nonlocal spectral boundary conditions. The higher K-modes then decouple as the radius of K tends to zero. Two qualifications remain essential: these boundary conditions restrict scalar functions at the walls, and compression of scalar functions through both stages differs from generating a matrix algebra between the stages. In particular, aligned walls give $\mathbb{R} \oplus \mathbb{R}$ on the quaternionic factor under direct scalar compression, whereas the intervening matrix-algebra construction gives $\mathbb{H} \oplus \mathbb{H}$. Local boundary gradings preserve all scalar functions, but their reduction to the unchanged four-section window has a precise restriction proved below. The radial no-go of lemma 22 survives rechecking in its stated scope. For the massless one-loop potential of lemma 23, an analytic argument strengthens the numerical statement to all $n_f > 0$.

1. Theorem — An exact spectral boundary construction.

Hypotheses and conventions. Let

$$K = \mathbb{CP}^3, \qquad \mathcal{H}_K = L^2\big(K, \Lambda^{0,*}T^*K \otimes \mathcal{O}(1)\big)\,,$$

and let

$$A = \sqrt{2}(\bar\partial + \bar\partial^*), \qquad \gamma_K = (-1)^{\deg}.$$

Use the Fubini–Study metric at a fixed reference scale. Write

$$V = \ker A \simeq \mathbb{C}^4, \qquad \Pi = P_K,$$
$$\Delta_K = \inf(\operatorname{spec}|A| \setminus \{0\}) > 0.$$

Here $\gamma_K|_V = 1$. The radius parameter $r > 0$ replaces A by A/r, after the usual unitary identification of the L^2-spaces.

These are the analytic data of Example 16.9. Its S^4 option uses functions pulled back to this same twistor space. An actual operator on S^4 with the required kernel and compression map would need to be specified separately.

Let

$$E = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}, \quad Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}, \quad S = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}.$$

On $L^2(I_\ell; \mathbb{C}^2 \otimes \mathcal{H}_K)$, set

$$\mathcal{D}_r = E\partial_x + Z\,A/r, \qquad \hat{\gamma} = Z\gamma_K.$$

All unspecified tensor factors are identities.

Choose self-adjoint involutions Γ_0, Γ_1 on V, each with two-dimensional positive and negative eigenspaces, and assume

$$-1 \notin \operatorname{spec}(\Gamma_1\Gamma_0).$$

Extend Γ_i by zero on $V^\perp$, writing $C_i = \Pi\Gamma_i\Pi$. Define $\operatorname{sgn} A = 0$ on V, and impose

$$\chi_0\psi(0) = \psi(0), \qquad \chi_\ell\psi(\ell) = \psi(\ell),$$

where

$$\chi_0 = ZC_0 - S\operatorname{sgn} A, \qquad \chi_\ell = ZC_1 + S\operatorname{sgn} A. \quad (1)$$

Then:

1. The H^1 realization of $\mathcal{D}_r$ is self-adjoint with compact resolvent and anticommutes with $\hat{\gamma}$.
2. $\Pi \otimes 1$ reduces the operator, including its domain.
3. On this reducing subspace, the operator is exactly $D(\Gamma_0, \Gamma_1)$ of lemma 23.
4. Every eigenvalue on the complementary subspace satisfies
$$|\lambda| \geq \Delta_K/r. \quad (2)$$

Proof. Put $B = SA/r$. Since $ES = Z$,

$$\mathcal{D}_r = E(\partial_x + B).$$

The operators χ_i are self-adjoint involutions and anticommute with E. Their positive eigenspaces are therefore maximal isotropic for the boundary form

$$[\langle \psi, E\phi\rangle]_0^\ell.$$

On $V^\perp$, condition (1) selects the negative spectral subspace of B at 0 and its positive spectral subspace at ℓ. These are APS conditions with the inward-normal sign reversed between the two walls. On $\ker B = \mathbb{C}^2 \otimes V$, condition (1) supplies a Lagrangian subspace:

$$\Gamma_i \psi_L = \psi_L, \qquad \Gamma_i \psi_R = -\psi_R.$$

For q in the complementary domain, integration by parts gives

$$\|\mathcal{D}_r q\|^2 = \|q'\|^2 + \|Aq/r\|^2 + [\langle q, Bq\rangle]_0^\ell \geq \|q'\|^2 + \frac{\Delta_K^2}{r^2}\|q\|^2. \tag{3}$$

The boundary term is nonnegative by the spectral choices.

An eigenbasis of A decomposes the complementary problem into finite-dimensional interval problems with maximal isotropic endpoint conditions. Each is self-adjoint. Estimate (3) identifies the domain of their self-adjoint direct sum with the stated H^1 domain and gives compactness of its resolvent. The finite-dimensional tangential kernel contributes an ordinary self-adjoint interval problem. This proves self-adjointness, compactness, and (2).

Every boundary projector commutes with Π, proving reduction. Finally, $\{\gamma_K, A\} = 0$ implies $\{\hat{\gamma}, \mathcal{D}_r\} = 0$, and $\hat{\gamma}$ preserves both boundary subspaces. ■

Define

$$K_F = \frac{1}{2\ell}\log(\Gamma_1\Gamma_0), \qquad D_F = \Gamma_0 K_F, \qquad R(x) = e^{xK_F},$$

with the principal logarithm. Let $P_\pm = (1 \pm \Gamma_0)/2$. The finite-window embedding is

$$W_\ell v(x) = \ell^{-1/2} \begin{pmatrix} R(x)P_+ v \\ R(x)P_- v \end{pmatrix}. \tag{4}$$

Direct differentiation gives

$$\mathcal{D}_r W_\ell = W_\ell D_F, \qquad \hat{\gamma} W_\ell = W_\ell \Gamma_0.$$

Thus the prescribed finite operator and grading are exact, provided $\Delta_K/r > \|D_F\|$. The next eigenvalue in absolute value is at least

$$\min\left\{\frac{\Delta_K}{r}, \frac{\pi - \ell\|D_F\|}{\ell}\right\}. \tag{5}$$

---

2. Proposition — Local walls and spectral walls have different scalar domain algebras.

Hypotheses. Use the product differential expression of Result 1.

For the local option, let $G_i(y)$ be smooth self-adjoint bundle involutions and impose

$$G_i\psi_L = \psi_L, \qquad G_i\psi_R = -\psi_R \quad \text{at } x = i.$$

Assume

$$[G_i, \gamma_K] = 0, \qquad [G_i, \sigma_A(y,\xi)] = 0 \quad (\xi \in T_y^*K), \tag{6}$$

where σ_A is the tangential principal symbol.

Then these local conditions give a self-adjoint elliptic realization with compact resolvent, and **every** smooth scalar function on $K \times I_\ell$ preserves its domain.

For the spectral option (1), the maximal smooth scalar $*$-algebra preserving the domain is

$$\mathcal{A}_{\mathrm{APS}} = \{f \in C^\infty(K \times I_\ell) : f(\,\cdot\,,0), f(\,\cdot\,,\ell) \text{ are constant on } K\}. \tag{7}$$

The same conclusion holds for the subalgebra pulled back from $S^4 \times I_\ell$.

Proof. For the local option, the boundary involution is ZG_i. It anticommutes with the normal matrix E. Condition (6) makes it anticommute with the principal symbol of the adapted tangential operator SA/r. Hence its positive eigenspace is an elliptic Lagrangian boundary condition. The corresponding local ellipticity and adjoint-domain criteria are described in Bär–Ballmann, Corollary 4.18 and Example 4.20. Primary reference

Scalar multiplication commutes pointwise with G_i, so it preserves the endpoint conditions and the H^1 domain. Moreover,

$$[\mathcal{D}_r, f] = E\,\partial_x f + Z[A, f]/r$$

is bounded for each $r > 0$.

For spectral walls, preservation by a scalar $*$-algebra is equivalent to

$$[M_{f_i}, \chi_i] = 0.$$

Independence of the spin matrices Z, S gives

$$[M_{f_i}, C_i] = 0, \qquad [M_{f_i}, \operatorname{sgn} A] = 0.$$

Since $C_i^2 = \Pi$, the first equation implies

$$[M_{f_i}, \Pi] = 0. \tag{8}$$

Consequently multiplication by f_i preserves $H^0(K, \mathcal{O}(1))$. If s is such a section, then

$$0 = \bar\partial(f_i s) = (\bar\partial f_i)s.$$

The sections of $\mathcal{O}(1)$ have no common zero, so $\bar\partial f_i = 0$ everywhere. A holomorphic function on compact connected $\mathbb{CP}^3$ is constant. Constants plainly preserve (1), proving (7). ■

The hypotheses in (6) are a sufficient local ellipticity condition. An arbitrary bundle grading supplies a pointwise isotropic condition, but self-adjointness on an H^1 domain and compact resolvent still require an ellipticity check. Likewise, prescribing a grading only on V leaves the boundary problem incomplete until a condition on $V^\perp$ is supplied.

---

3. Proposition — Compression of a local grading and the unchanged level-one window.

Hypotheses. Let G be a self-adjoint bundle involution acting by multiplication, and let

$$g = \Pi G \Pi|_V, \qquad Q_K(f) = \Pi M_f \Pi|_V.$$

Keep the level-one space and scalar representation of Example 16.9 unchanged.

Then

$$1_V - g^2 = \big((1-\Pi)G\Pi\big)^*\big((1-\Pi)G\Pi\big). \tag{9}$$

In particular, g is a grading exactly when $[G,\Pi]=0$. In that event,

$$[g, Q_K(f)] = 0 \quad \text{for every scalar } f. \tag{10}$$

For either the full $\mathbb{CP}^3$ scalar algebra or the S^4 pullback algebra, this forces

$$g = 1_V \quad \text{or} \quad g = -1_V. \tag{11}$$

Proof. Expand $\Pi G^2\Pi = \Pi$ using $1 = \Pi + (1-\Pi)$; this gives (9). Vanishing of its positive right-hand side implies $(1-\Pi)G\Pi = 0$, and self-adjointness gives the other off-diagonal block.

Since scalar multiplication commutes with G, reduction then gives (10). The compressed $\mathbb{CP}^3$ functions generate $M_4(\mathbb{C})$. The compressed real S^4 functions generate $M_2(\mathbb{H})$, whose complex-linear commutant on $\mathbb{C}^4$ is also $\mathbb{C}1$. Thus a self-adjoint involution satisfying (10) is ± 1. ■

For **aligned local walls** with $[G,A]=0$, the finite zero-mode window is explicitly

$$v \longmapsto \ell^{-1/2}\begin{pmatrix} P_+(g)v \\ P_-(g)v \end{pmatrix}.$$

The induced grading is g, and scalar compression is

$$P_+(g)Q_K(f)P_+(g) + P_-(g)Q_K(f)P_-(g) = Q_K(f).$$

By (11), the unchanged four-section window has trivial grading. Its generated algebra remains $M_2(\mathbb{H})$ or $M_4(\mathbb{C})$, according to the scalar algebra chosen. The globally constant choices $G = \pm 1$ give concrete self-adjoint examples with this window and a separated higher spectrum.

There is also an explicit local field illustrating failure of reduction. Add a trivial auxiliary $\mathbb{C}^4$ bundle and put

$$G(y) = \sum_{a=1}^{5} y_a \Gamma_a^{\mathrm{aux}},$$

where $y \in S^4$ is the twistor projection and the Γ_a satisfy $\{\Gamma_a, \Gamma_b\} = 2\delta_{ab}$. This is a local involution satisfying (6), but now the tangential kernel has dimension 16, and

$$g = \frac{1}{5}\sum_{a=1}^{5} \Gamma_a^{\text{section}} \otimes \Gamma_a^{\text{aux}}.$$

Using $\operatorname{tr}(\Gamma_a \Gamma_b) = 4\delta_{ab}$,

$$\operatorname{tr}(g^2) = \frac{16}{5} \neq 16.$$

Thus g is not a grading. Equation (9) even gives

$$\|(1-\Pi)G\Pi\| \geq \frac{2}{\sqrt{5}},$$

independently of r. This is a bound on boundary-trace mixing; it does not identify the actual low spectral window of the coupled local problem.

A local $G(y)$ specifies a **field of gradings**. It does not select a marked point p. In the tautological quaternionic line bundle,

$$\Pi_y = \frac{1 + y_a \Gamma_a}{2},$$

the natural field $2\Pi_y - 1$ acts as $+1$ on the fibre $\operatorname{ran}\Pi_y$. By contrast, the finite grading

$$\Gamma_p = p_a \Gamma_a = 5Q_K(p \cdot y)$$

uses a separately chosen point p. Its action on global section coefficients is nonlocal.

The restriction proved here concerns the unchanged level-one space, a local multiplication grading, and its compression as an actual involution. A different coupled window, additional fibres, or spectral boundary conditions remain available directions.

---

4. Theorem — The composed compression map and the aligned-wall algebra.

Hypotheses. Use the exact spectral construction of Result 1. For a scalar multiplier $f(y, x)$, let

$$T_f(x) = Q_K(f(\,\cdot\,, x)), \qquad \mathcal{E}_0(T) = \frac{1}{2}(T + \Gamma_0 T \Gamma_0).$$

Use real-valued functions and real algebra generation when discussing $M_2(\mathbb{H})$.

Then the compression directly from the scalar cylinder to the finite window is

$$\boxed{\Phi_\ell(f) = \frac{1}{\ell}\int_0^\ell \mathcal{E}_0\big(R(x)^* T_f(x) R(x)\big)\, dx.} \tag{12}$$

It is independent of r.

For aligned walls $\Gamma_1 = \Gamma_0 = \Gamma_p$, direct scalar compression on the S^4 factor generates

$$\mathbb{R} \oplus \mathbb{R}. \tag{13}$$

If instead one first generates $M_2(\mathbb{H})$ and then applies the interval compression, the result is

$$\mathbb{H} \oplus \mathbb{H}. \tag{14}$$

Proof. Insert (4) on both sides of the multiplication operator. Orthogonality of the two interval spin components gives exactly (12).

For real S^4 functions, Example 16.9 implies

$$Q_K(C^\infty(S^4, \mathbb{R})) = \mathrm{span}_{\mathbb{R}}\{1, \Gamma_1, \dots, \Gamma_5\}. \tag{15}$$

Indeed all compressed functions are quaternionic Hermitian matrices; that space has real dimension six, and the displayed six matrices belong to the image and are independent.

For aligned walls $R = 1$. Clifford multiplication gives

$$\mathcal{E}_p(\Gamma_a) = p_a \Gamma_p.$$

Therefore the image of direct scalar compression is

$$\mathrm{span}_{\mathbb{R}}\{1, \Gamma_p\}.$$

This space is already the algebra $\mathbb{R} \oplus \mathbb{R}$, proving (13).

On the other hand,

$$\mathcal{E}_p(M_2(\mathbb{H})) = \{a \in M_2(\mathbb{H}) : [a, \Gamma_p] = 0\} = \mathbb{H} \oplus \mathbb{H},$$

proving (14). The noncommuting products needed to generate the quaternionic diagonal blocks have been inserted between the two compressions in this second procedure. ■

On the single $\mathbb{CP}^3$ factor, real scalar compression is onto the Hermitian 4×4 matrices. A chosen $2+2$ aligned grading therefore gives all Hermitian block-diagonal matrices and generates

$$M_2(\mathbb{C}) \oplus M_2(\mathbb{C}).$$

For the **chosen two-sheet coarse correspondence** of Example 16.9, with

$$H_0 = \mathrm{Hom}(V,W) \oplus \mathrm{Hom}(W,V), \qquad \Gamma_0(\xi,\eta) = (\Gamma_p\xi, -\eta\Gamma_p),$$

the same calculation gives

$$\boxed{\mathrm{alg}_{\mathbb{R}}(\text{direct aligned scalar compression}) = \mathbb{R} \oplus \mathbb{R} \oplus M_4(\mathbb{C}).} \tag{16}$$

The V-side matrices act on the left of η and commute with its right grading, so that factor survives unchanged.

Endpoint restrictions do not change these finite-dimensional images. If a linear functional vanishes on compressions of functions supported in $0 < x < \ell$, its smooth matrix-entry density vanishes in the interior and hence everywhere. Finite-dimensional separation then gives equality with the unrestricted image, exactly as in Result 7 of lemma 20.

At the intermediate stage, however, spectral walls still give scalar endpoint values. Even after algebra generation, those values remain scalar. Enlarging them to the full stabilizers of Γ_0, Γ_1, as in lemma 23, is another change of the parent algebra. Equality of finite compression images does not identify the parent gauge groups or one-forms.

---

5. Theorem — Radius collapse, interval collapse, and controlled mixing.

Hypotheses. First use the exact construction of Result 1. Identify the varying Hilbert spaces by metric rescaling. For resolvent formulas, fix units in which the nonreal spectral parameter is i.

Then, at fixed ℓ,

$$\left\|(\mathcal{D}_r - i)^{-1} - \iota_\Pi (D(\Gamma_0, \Gamma_1) - i)^{-1} \iota_\Pi^*\right\| \le \frac{r}{\Delta_K}. \qquad (17)$$

Thus radius collapse gives the entire matrix interval operator in generalized norm-resolvent sense.

If D_F is fixed, $\ell \to 0$, and

$$\Gamma_1(\ell) = e^{\ell \Gamma_0 D_F} \Gamma_0 e^{-\ell \Gamma_0 D_F},$$

then $r \to 0$ and $\ell \to 0$ decouple both towers and leave the finite operator D_F. No relation between the two rates is required in this exactly reducing construction.

Proof. The resolvent is block diagonal with respect to Π. Estimate (2) bounds its complementary block by r/Δ_K, proving (17). For the finite-window limit, (5) bounds the complementary resolvent by

$$\max\left\{\frac{r}{\Delta_K}, \frac{\ell}{\pi - \ell\|D_F\|}\right\},$$

which tends to zero. ■

A useful mixing statement requires additional hypotheses. Suppose a transformed problem has a common domain and the form

$$\widetilde{\mathcal{D}}_r = \mathcal{D}_r + V_r, \qquad V_r = V_r^*, \qquad \|V_r\| \le b_r.$$

Write

$$v_r = \|(1 - \Pi) V_r \Pi\|.$$

For an eigenvector $\psi = p + q$, with $p = \Pi\psi$, $q = (1 - \Pi)\psi$, and $|\lambda| \le L$, assume

$$g_r = \Delta_K / r - b_r - L > 0.$$

Then

$$\boxed{\|q\| \le \frac{v_r}{g_r}\|p\|, \qquad \|\text{Schur correction at } \lambda\| \le \frac{v_r^2}{g_r}.} \qquad (18)$$

Proof. The complementary eigenvalue equation gives

$$q = -(\mathcal{D}_{r,Q} + V_{QQ} - \lambda)^{-1} V_{QP} p.$$

The inverse has norm at most g_r^{-1}. Substituting this formula into the p-equation gives the second bound. ■

For bounded b_r, v_r, both errors are $O(r/\Delta_K)$. If a common-domain construction instead produces terms of order $1/\ell$, sufficient conditions include

$$\frac{rv_r}{\Delta_K} \to 0, \qquad \frac{rv_r^2}{\Delta_K} \to 0,$$

together with convergence of the Π-diagonal perturbation.

An arbitrary local $G_i(y)$ generally fails $[G_i, \Pi] = 0$. Although the differential expression commutes with Π, its domain then does not. Such a boundary change is not automatically a bounded perturbation satisfying (18). The explicit local field in Result 3 has nonvanishing boundary mixing at every radius. No convergence to the prescribed two-gradings operator is established here for that nonreducing local problem.

At the level of observables, radius collapse leaves (12) unchanged. For an x-independent $T = Q_K(f)$,

$$\|\Phi_\ell(f) - \mathcal{E}_0(T)\| \le \ell \|D_F\| \, \|T\|. \tag{19}$$

If T is even for Γ_0, the stronger formula of Theorem 22.4 applies:

$$\Phi_\ell(f) = \operatorname{sinc}(\ell \operatorname{ad}_{D_F})(T),$$

$$\|\Phi_\ell(f) - T\| \le \frac{\ell^2}{6} \|[D_F, [D_F, T]]\|. \tag{20}$$

Equation (19) follows by integrating $\|R^* T R - T\| \le 2x\|D_F\|\|T\|$; equation (20) follows by taking the even part of the conjugation series.

Consequently there is convergence of the **compression maps on their scalar operator systems**. The limit is $\mathcal{E}_0 Q_K$. It does not turn that image into every element of $\mathcal{E}_0(P)$; Result 4 computes the difference explicitly. Algebra generation remains a separate operation.

---

6. Proposition — The second sheet, reality, even algebra, and connector remain prescribed data.

Hypotheses. Consider the scalar cylinder construction above and, separately, its extension using the two-sheet coarse correspondence of Example 16.9.

The resulting distinctions are:

Structure	What follows under the stated choices
Second sheet and the two Hom spaces	They are inputs in the coarse correspondence. A single four-section cylinder window does not supply them.
$J_1^2 = -1,\ J_1\Gamma_i = -\Gamma_i J_1$	These are additional compatibility conditions on the prescribed gradings and antiunitary.
$\mathbb{H} \oplus \mathbb{H} \oplus M_4(\mathbb{C})$	This is the derived even part after choosing the quaternionic real form, the grading, and intermediate algebra generation. Direct aligned scalar compression gives (16).
Connector and $\mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C})$	The connector's lines and coupling are chosen. Their intertwining algebra is then derived.

Proof. The one-factor construction starts with $V = \mathbb{C}^4$, and (4) has domain V. The two-sheet construction instead starts with

$$(W \otimes \overline{V}) \oplus (V \otimes \overline{W}).$$

Geometrically, its analytic realization uses product factors and a restricted scalar algebra, as Theorem 16.7 specifies. It is additional correspondence data over the scalar spaces.

For reality, suppose an antiunitary $\mathcal{J}_1$ on the tangential Hilbert space satisfies

$$\mathcal{J}_1^2 = -1, \quad \mathcal{J}_1 A = -A\mathcal{J}_1, \quad [\mathcal{J}_1, \gamma_K] = 0,$$

and its restriction J_1 to V anticommutes with both Γ_i. Then

$$\widehat{J} = E\mathcal{J}_1$$

preserves (1) and obeys

$$\widehat{J}^2 = 1, \qquad [\widehat{J}, \mathcal{D}_r] = 0, \qquad \widehat{J}\widehat{\gamma} = -\widehat{\gamma}\widehat{J}.$$

On the window,

$$J_F = \Gamma_0 J_1.$$

These identities follow directly from the stated commutation relations. They specify sufficient conditions; the natural $\mathcal{O}(1)$ cylinder has not supplied such an extension or its order-zero property here.

In particular, the quaternionic antiunitary j used to recognize $M_2(\mathbb{H})$ commutes with the Γ_a. It does not itself satisfy the required anticommutation with Γ_p.

The even-algebra assertion is the centralizer calculation in Result 4 after intermediate generation.

Finally, choose

$$w \in W_+, \quad v \in V, \quad \|w\| = \|v\| = 1, \quad \kappa \neq 0,$$

and prescribe

$$T_\kappa \xi = \kappa(w^* \xi v)\, v w^*.$$

Its intertwining conditions require the lines $\mathbb{C}w, \mathbb{C}v$ to be reducing, with the same scalar action. A quaternion preserving $\mathbb{C}w$ lies in its chosen copy of $\mathbb{C}$; $v^\perp$ carries arbitrary $M_3(\mathbb{C})$, and W_- carries arbitrary $\mathbb{H}$. Hence

$$\mathcal{A}(T_\kappa) = \{(\lambda \oplus q, \lambda \oplus m)\} \simeq \mathbb{C} \oplus \mathbb{H} \oplus M_3(\mathbb{C}),$$

as in Result 2 of lemma 20. ■

A marked $p \in S^4$ fixes a quaternionic line, but choosing w also requires a complex line within it. The second line v, the coupling, and an interface realizing the connector remain necessary. Scalar agreement at marked points has unchanged compression image by Result 7 of lemma 20, so it supplies none of this selection by itself.

---

7. No-go theorem — Recheck of the radial statement in lemma 22.

Hypotheses. Let the parent be exactly the one-interval model of lemma 22: fixed wall coupling X, zero bulk mass, the even algebra $\mathcal{A}_\ell$, scalar inner fluctuations Eh with $h^* = -h$, and the **moving spectral window**. For real fluctuations, include the conjugate one-form and assume the stated order-zero representation.

Then every fluctuated finite operator has the same singular values as D_F. This is an exact finite-amplitude spectral statement; it uses no action or loop expansion.

Proof. Solve

$$u' = -hu, \qquad u(0) = 1, \qquad w = u(\ell),$$

and set $v(x) = u(x)w^*$. Then

$$v(\ell) = 1, \qquad v(0) = w^*.$$

Since h is even, w preserves the boundary splitting at 0; the identity at ℓ preserves the fixed coupled wall. Thus v preserves the original domain, and

$$vD_X v^* = D_X + Eh.$$

For the real fluctuation, the analogous implementing unitary is

$$vJ_F vJ_F^{-1};$$

order zero separates the two factors in its derivative.

The full operators are unitarily equivalent, so their moving spectral windows are unitarily equivalent. In the finite identifications of lemma 22, the operators are $w^* D_F w$, or $U_w^* D_F U_w$ in the real case. Their singular values coincide. ■

The moving-window hypothesis matters. For example, with $D_F = m\sigma_1$ and $h = i\sigma_3$, fixed compression gives

$$Q_\ell(Eh) = -\frac{1 - \cos(2m\ell)}{2m\ell}\sigma_2.$$

Adding a multiple of this operator to the fixed $m\sigma_1$ changes its singular values at second order.

The verdict agrees with the author's amended theorem and Terra's independent endpoint-preserving proof. In the usual nondegenerate Standard Model representation,

the normalized shifted Higgs quaternion lies on the unit-quaternion orbit; its norm is fixed.

The known escapes are a variable wall coupling, odd connection directions, suitable bulk masses, additional internal directions with nontrivial connection data, or additional elementary or composite fields. A quantum effective action on the same configuration space does not add a radial coordinate. The proposed $K \times I_\ell$ parent has tangential directions and tangential one-forms, so this theorem does not decide its full fluctuation space.

---

8. Lemma — A logarithmic concavity property of the one-loop integral.

Hypotheses. Let

$$\mathcal{J}(z) = \int_0^\infty \frac{x^3}{\sinh^2 x + z}\, dx, \qquad z > 0,$$

and define

$$e(z) = -\frac{z\mathcal{J}'(z)}{\mathcal{J}(z)}.$$

Then $0 < e(z) < 1$, and e is nondecreasing.

Proof. Positivity and

$$0 < \frac{z}{\sinh^2 x + z} < 1$$

show that $e(z)$, expressed as a weighted average of this fraction, lies strictly between zero and one.

Put $x = e^v$, $z = e^u$. Then

$$\mathcal{J}(e^u) = \int_{\mathbb{R}} \exp\left(4v - \log(\sinh^2(e^v) + e^u)\right)\, dv. \tag{21}$$

The function

$$b(v) = 2\log\sinh(e^v)$$

is convex, because, with $x = e^v$,

$$b''(v) = 2x(\coth x - x\ \mathrm{csch}^2 x) > 0.$$

The last inequality is equivalent to $\sinh x \cosh x > x$.

The function $(u, v) \mapsto \log(e^{b(v)} + e^u)$ is therefore convex. The integrand in (21) is jointly log-concave. Marginal

integration preserves log-concavity by Prékopa's theorem, so $u \mapsto \log \mathcal{J}(e^u)$ is concave. We use precisely Theorem 6, p. 342, of Prékopa's original paper.

Finally,

$$\frac{d}{du} \log \mathcal{J}(e^u) = -e(e^u).$$

Concavity makes this derivative nonincreasing, proving the assertion. ■

---

9. Theorem — The massless one-loop potential has no interior local minimum.

Hypotheses. Let the action be exactly the one-loop determinant displayed in lemma 23: a flat interval, zero bulk masses, frequency-one fermions, the stated charged and neutral gauge towers and multiplicities, and the prescribed connector wall condition. Assume that no fermion connector modifies the displayed fermion tower, and add no wall potential or other angle-dependent term.

Let $n_f > 0$, set $t = \sin^2 \theta$, and write

$$\mathcal{V}(t) = 8\pi^2 \ell^4 V(\theta).$$

Then:

1. $t = 1$, equivalently $\theta = \pi/2$, is the unique global minimum.
2. If $n_f \geq 54/5$, the potential is strictly decreasing on the interior.
3. If $0 < n_f < 54/5$, there is exactly one interior stationary point, and it is a strict maximum.

Proof. Define

$$F(z) = \int_0^\infty x^3 \log\left(1 + \frac{z}{\sinh^2 x}\right) dx, \qquad F'(z) = \mathcal{J}(z).$$

The exact potential is

$$\mathcal{V}(t) = 3F(2t(1-t)) + \frac{3}{2} F\left(\frac{16}{5} t(1-t)\right) - n_f F(t). \quad (22)$$

Its gauge terms are nonnegative and vanish at $t = 1$. Since F is strictly increasing, the fermion term is uniquely minimized there. This proves item 1.

Differentiate:

$$\mathcal{V}'(t) = (1-2t)\left[6\mathcal{J}(2t(1-t)) + \frac{24}{5}\mathcal{J}\left(\frac{16}{5}t(1-t)\right)\right] - n_f\mathcal{J}(t). \tag{23}$$

For $t \geq 1/2$, this is strictly negative.

For $0 < t < 1/2$ and $c \geq 2$, put

$$R_c(t) = (1-2t)\frac{\mathcal{J}(ct(1-t))}{\mathcal{J}(t)}.$$

Let $a = ct(1-t)$. Then $a \geq t$, and Result 8 gives $e(a) \geq e(t)$. Consequently,

$$\begin{aligned}\frac{d}{dt}\log R_c(t) &= -\frac{2}{1-2t} - \frac{1-2t}{t(1-t)}e(a) + \frac{e(t)}{t} \\ &\leq -\frac{2}{1-2t} + \frac{e(t)}{1-t} \\ &< -\frac{2}{1-2t} + \frac{1}{1-t} < 0.\end{aligned}$$

Thus every R_c is strictly decreasing. Since $\mathcal{J}(0)$ is finite and positive,

$$R_c(0+) = 1, \qquad R_c(1/2-) = 0.$$

Therefore

$$H(t) = 6R_2(t) + \frac{24}{5}R_{16/5}(t)$$

decreases strictly from $54/5$ to zero. Equation (23) becomes

$$\mathcal{V}'(t) = \mathcal{J}(t)\,[H(t) - n_f] \quad (0 < t < 1/2).$$

For $n_f \geq 54/5$, its sign is negative. For $0 < n_f < 54/5$, it changes sign exactly once, from positive to negative. Since $t = \sin^2\theta$ is strictly increasing on the interior, the same stationary-point classification holds in θ. ■

This proves the massless statement for the entire parameter range, including the previously numerical range. No numerical computation is used in this proof.

The recorded author and Terra referee calculations agree on the sampled massless cases. The logarithmic-concavity proof above is an additional single-agent derivation; it has not received a fresh independent model review

in this turn. The recorded agreement therefore should not be described as a multi-model review of this new proof.

The escapes remain changes to the stated determinant: frequency-two fermions, modified form factors through warping or other geometry, suitable bulk-mass sectors, wall potentials, and wall kinetic or interaction terms. The full effective action of the proposed $K \times I_\ell$ parent, including its K-tower, has not been computed here.

---

10. Proposition — Scope of the bulk-mass scans.

Hypotheses. Retain the flat gauge terms of (22) as an independent assumption. Replace the fermion factor by the expression in lemma 23:

$$1 + t(G(x) - 1),$$
$$G(x) = \frac{(w_Q \coth w_Q - c_Q)(w_U \coth w_U - c_U)}{x^2},$$
$$w_i = \sqrt{x^2 + c_i^2}.$$

If $c_U = -c_Q = c$, then the potential is finite and $t = 1$ is its unique global minimum for every $n_f > 0$. The reported absence of additional interior **local** minima at nonzero c remains a numerical statement here.

If $a = c_Q + c_U \neq 0$, the displayed unrenormalized integral has angle-dependent ultraviolet divergences.

Proof. For opposite masses,

$$G(x) = 1 + \frac{w^2}{x^2 \sinh^2 w}, \qquad w = \sqrt{x^2 + c^2}.$$

The added fraction is positive and decays exponentially at infinity. The negative fermion logarithm is strictly decreasing in t. The gauge logarithms are nonnegative and vanish at $t = 1$, proving the global-minimum assertion.

For general masses,

$$G(x) = 1 - \frac{a}{x} + \frac{a^2}{2x^2} + O(x^{-3}),$$

and hence

$$\log(1 + t(G - 1)) = -\frac{ta}{x} + \frac{t(1-t)a^2}{2x^2} + O(x^{-3}).$$

Multiplication by the integration measure x^3dx produces angle-dependent divergences. A finite potential therefore requires a specified completion, regulator, and counterterms. ■

The opposite-mass scan in lemma 23 samples finitely many masses, multiplicities, and angles. It cannot exclude unsampled local minima. General bulk masses and altered gauge form factors remain directions for constructing an interior vacuum; these conclusions leave that question open.

Consequence for T0 and T0b

For T0, a commutative scalar representation with prescribed spectral walls gives an exact interval operator window and controlled tower decoupling, while local walls preserve the full scalar algebra under the ellipticity hypotheses above. The derivation of a nontrivial local grading and the desired algebra from one coupled scalar window remains open, and intermediate algebra generation must be stated explicitly. For T0b, the second sheet, compatible reality, quaternionic grading choice, and connector remain inputs; their algebraic consequences can be derived once those choices are supplied. Neither the radial theorem nor the specified one-loop potential closes the commutative-parent problem.

Referee amendments

From Terra 6.1. These apply to the text above.

- **Result 2.** Write G_0 at x = 0 and G_l at x = l. Equation (6) supplies the principal-symbol condition of Bar-Ballmann, Corollary 4.18 and Theorem 4.11; their Example 4.20 assumes more. On the unaugmented Dolbeault bundle, the Clifford representation is irreducible, so a fibre endomorphism commuting with every sigma_A is scalar, and a smooth involution of this form is +1 or -1. Additional fibres and other local elliptic projectors lie outside this sufficient criterion.
- **Result 3.** The statement on marked points concerns the natural field displayed there. A local field can encode a chosen point: on the auxiliary bundle, G_p(y) = p_a Gamma_a^aux is a local constant

field whose coefficients contain p.

- **Result 4.** Equation (16) has the two Hom spaces and the two scalar algebras of Example 16.9 as inputs. The endpoint argument proves equality of the finite-dimensional linear images, and does not identify the parent algebras, norms, gauge groups or one-forms.
- **Result 5, gap in the final assertion.** Convergence of compression maps needs a common domain and a collapse map. With aligned walls, D_F = 0 and f_l(y, x) = 6 (x/l)(1 - x/l), both endpoint values vanish while Phi_l(f_l) = 1. A correct form: for F(y, s) fixed on K x [0, 1] and f_l(y, x) = F(y, x/l), Phi_l(f_l) tends to E_0 Q_K(int_0^1 F(., s) ds). For x-independent ambient multipliers the limit is E_0 Q_K, by (19). The rate discussion after (18) also needs control of the denominator: a sufficient addition is r (b_r + L) <= (1 - epsilon) Delta_K. The case b_r = O(1/l) follows from the stronger estimate |D_{r,Q}| >= (Delta_K^2/r^2 + pi^2/(4 l^2))^{1/2}, from the one-ended Poincare inequality on the positive and negative B-spectral components.
- **Result 6.** At the finite level, if the gradings are Gamma_p, Gamma_q, a unit n orthogonal to p and q gives J_1 = Gamma_n j, with J_1^2 = -1 and the required anticommutation. Extending this choice to the parent remains a separate task.
- **Lemma 8 and Theorem 9.** An independent check of the key monotonicity, without log-concavity or Prekopa's theorem: with rho(s) = (arsinh sqrt(s))^3/(2 sqrt(s (1 + s))), J(z) = int_0^infinity rho(s)/(s + z) ds; with L(v) = log rho(e^v), L'(v) = 3 tanh(x)/(2x) - (1 + tanh^2 x)/2 for v = 2 log sinh x, which is strictly decreasing. For z = e^u, with the probability density in y proportional to e^y rho(e^{u+y})/(1 + e^y), e(e^u) is the mean of 1/(1 + e^y), and its derivative in u is Cov(1/(1 + e^y), L'(u + y)) > 0, as both are strictly decreasing. Theorem 9 is for theta in [0, pi/2]. At theta = pi/2 the determinant has its continuous value, while the principal-logarithm window is degenerate. Its hypotheses are the flat massless one-loop determinant, its gauge multiplicities and connector wall condition, frequency-one

fermions, no fermion-connector modification, and no other angle-dependent term.

- **Result 10.** The divergence is stated for nonzero fermion multiplicity, with no cancellation by additional sectors. The opposite-mass scan in lemma 23 is a numerical result on its sampled parameters.

Numerical check

By Claude Opus 5.5, in Python (numpy, scipy), 3 Nov 2026, in a temporary directory, not kept: E_p(Gamma_a) = p_a Gamma_p to 10^-15, which gives equation (13); tr(g^2) = 16/5 for the auxiliary field of Result 3; on grids, e(z) lies in (0, 1) and is nondecreasing for z in [10^-3, 50]; H(t) is strictly decreasing on (0, 1/2), with H near 54/5 at small t and near 0 at t = 1/2; for n_f in {1, 3, 6, 8, 10}, H - n_f has one sign change, with the stationary angle 0.34 at n_f = 8, which agrees with the referee of lemma 23. Results 1, 2, 5 and 6 were read and not separately computed.

Part 28

Two sheets as the two ends of one connected commutative parent

For STRATEGY T0 and T0b, following lemma 16 (Example 16.9), lemma 24 (Proposition 6) and lemma 09. Propositions 28.1 to 28.4 and 28.13 are by Claude Opus 5.5. Results 28.8 to 28.12 and 28.14 to 28.21 are by Astra 6, job 28 (report, with the full proofs; prompt), checked by Claude Opus 5.5 by hand and in Python. Proposition 28.13 and the numerics are refereed by Terra 6.1 (report), with its corrections applied. The window is a conjecture with numerical support; the numbers come from a radial solve in Python (numpy and scipy), described in the section Numerical check; the script is not kept.

Correction, 3 Nov 2026. The first version of Result 28.5 and of the numbers in Remark 28.7 (commits dd1b14c, 530fa06) used the wrong sign of the spin connection in the Python slice operator. That changed the singular values and the boundary fibre parts, so those numbers are withdrawn, including the next singular value 1.04, the value 1.35 on the representation 16, and the

boundary fibre parts -0.011 to -0.853. The exponents at the origin and the multiplicities are unaffected. Astra 6 did not reproduce the old boundary values and its proofs do not use them. We found the error when that operator gave a core zero mode in the representation 20, against the exact count of Proposition 28.11. The singular values, fibre parts and densities below are recomputed, and the script tests the operator against exact statements: {A_0, A_1} = 0, A_0^2 equal to the Casimir formula on S^4, D^2 on (0, 0)-forms with values in O(1) at the Fubini-Study radius equal to j(j + 4), and A_2 equal to its curvature contraction (the last test proposed by Terra 6.1).

Update, 4 Nov 2026. Lemma 29 proves Conjecture 28.6 in case (b) for the metric class of its Setting 29.0, and in case (a) at the Fubini-Study radius, refereed by Terra 6.1. Item 3 below and the Status of Proposition 28.13 describe the position before that lemma.

Outcome.

1. *The mechanism is proved at the level of compression maps.* A smooth function on the disc bundle X has independent values on the zero section S^4 and on the boundary CP^3, and at the zero section it is a function on S^4 alone (Proposition 28.8). For prescribed windows over this X, real functions generate exactly M_2(H) + M_4(C) at every localization scale (Theorem 28.21). In Example 16.9 the second sheet and the restriction to S^4 were inputs, so this is a connected source for them.
2. *The two ends of an actual local operator are understood.* There is a self-adjoint operator with a local wall (Theorem 28.9). The core carries exactly four states in the representation 4 (Proposition 28.11). At the wall, large mass keeps a tangential mixing (Lemma 28.12). For the pullback connection the boundary mode then has the fibre part -w(rho) in closed form, with w(rho) = 2(rho^2+2)^2/(9+2(rho^2+2)^2), which reaches the level-one value of Example 16.9 only as rho -> infinity (Proposition 28.13). This is a statement on the finite wall model, and its convergence from the full operator is numerical. A connection preserving the line summands in a collar, with the Fubini-

Study metric at the wall, gives the level-one value in that model.

3. *The isolated window of the whole operator stays a conjecture* (Conjecture 28.6): what is missing is the normal reductions with their errors, the matching, and a coercive estimate uniform in the mass and in all representations of Sp(2) (Lemma 28.14).
4. *The algebra M_2(H) + M_4(C) is the algebra of the limit map.* Generated algebras need not converge when compression maps do (Counterexample 28.17). At the sample points resolved numerically, the core mode has a nonzero density breaking the quaternionic structure, empirically of order lambda^2; if it is nonzero in the continuum, Proposition 28.16 gives M_4(C) for the core block at those masses. This is numerical evidence and does not establish the exact algebra at any finite mass.

Setting

Let G = Sp(2), K = Sp(1)_a x Sp(1)_b and H = Sp(1)_a x U(1)_b, so that G/K = S^4, K/H = S^2 and G/H = CP^3. Let S be the bundle G x_K (1, 2), so that CP^3 = P(S). The normal bundle of the zero section is G x_K (1, 3), and away from it X = (0, l] x CP^3, with radial coordinate x and S^2 fibres of radius rho(x), rho(x) ~ x at x = 0. Charges of U(1)_b are normalized so that the doublet of Sp(1)_b has +-1. The isotropy representation is (2)_{+-1} (+) (1)_{+-2}. The spinors of X restrict on a slice to Sigma_6 = Sigma^+ (+) Sigma^-, with c(d_x) = Z the chirality of Sigma_6, where

```
Sigma^+ = (1)_2 (+) (1)_0 (+) (2)_{-1},   Sigma^- =
↪  (2)_1 (+) (1)_0 (+) (1)_{-2}.
```

The representation 4 of G restricts to (2)_0 (+) (1)_1 (+) (1)_{-1}, that is (2, 1) (+) (1, 2) of K. Let Gamma_a be the Clifford generators on 4, with Gamma_5 commuting with K and equal to +-1 on the two summands, and let x_a(g) be defined by lambda(g) Gamma_5 lambda(g)^-1 = x_a(g) Gamma_a.

Proposition 28.1 (the twist).

1. For k != 0, O(k) on CP^3 does not extend to a line bundle on X, since H^2(X) = H^2(S^4) = 0 while

O(k) has degree k on the fibre S^2.

2. The pullback of S to X restricts on CP^3 to O(1) (+) O(-1), as the tautological splitting. So E = S^3(S), which is pulled back from S^4 and has trivial connection along the fibres, restricts to O(3) (+) O(1) (+) O(-1) (+) O(-3). The summand O(3) gives the Dolbeault operator on O(1) of Example 16.9, since Sigma (x) O(3) = Lambda^{0,*} (x) O(1).
3. On E the section Phi(v) = phi(|v|) (v-hat . T), with T the spin-3/2 generators and phi odd and smooth, is smooth and G-equivariant, with zero set exactly S^4. In the frame of the slice it is phi(x) tau, with tau diagonal in the summands, with eigenvalues +-3/2, +-1/2.

So we take the Callias-type operator T = D_X(E) + i Phi on spinors of X tensored with E, and the self-adjoint operator D-hat = [[0, T^*], [T, 0]], with the grading Gamma-hat = diag(1, -1). At x = 1 we take the local condition Gamma-hat Z = -1 on the doubled spinor, that is u_even(1) = 0 for T and v_odd(1) = 0 for T^*. It is a chirality condition, elliptic and local, and it preserves the domain under all smooth functions on X.

Proposition 28.2 (multiplicities). On a slice, Hom_H(4, Sigma^+ (x) E) and Hom_H(4, Sigma^- (x) E) each have dimension 5. For the summand O(3), the dimensions are 1 and 0. So on 4, the operator on the summand O(3) with its U(1) connection vanishes for every G-invariant metric on the slice, and its kernel contains H^0(CP^3, O(1)) = 4 for every fibre radius. This concerns the separate line connection. The pull-back connection of E does not preserve the summands, and its slice operator is invertible for small fibres (Proposition 28.10).

Proof. Tensoring the lists above with the charges of E gives the multiplicities. The operator is odd under Z and maps the multiplicity space of 4 into itself. ∎

Proposition 28.3 (the zero section). The fibre operator c_N . nabla + i Phi on R^3 with values in E has zero modes (1, 3) (+) (1, 1) of K, by the Callias count for the positive weights 3/2 and 1/2, with one sign of Phi. On S^4 this gives the Dirac operator twisted by (1, 3) (+) (1, 1). Then Hom_K(4, S^- (x) (1, 3)) has dimension 1, through (1, 2), and Hom_K(4, S^+ (x) (1, 3))

= 0, so the twisted operator vanishes on that copy of 4, in agreement with its index 4. The untwisted part (1, 1) has eigenvalue 2 on the representation 4. Every map in Hom_K(4, S^- (x) (1, 3)) annihilates (2, 1) in 4. These are the twistor spinors of S^4, as sections of S.

Proof. The representation count is as in Proposition 28.2, and the Callias count is a reference, with the exact count in Proposition 28.11. For the last sentence, the first-order operator from (1, 2) to (2, 3) in the slice limit is the twistor operator, and its kernel is the 4. ■

On a representation V the twisted operator has D^2 = Cas_G(V) + 3/2 - Cas_K(tau), for the twist tau, with Cas_G(4) = 5/2, Cas_G(16) = 15/2, Cas_G(20) = 21/2, Cas_G(40) = 29/2, Cas_K(j_a, j_b) = 2 j_a(j_a+1) + 2 j_b(j_b+1). With the twist (1, 3) this gives 0, 5, 8, 12, and with the twist (1, 1) on the representation 4 it gives 4. The slice operator reproduces these (Result 28.5).

Proposition 28.4 (the compression). Let s_v(g) = F lambda(g)^-1 v be a section in the multiplicity space of 4 on a slice, with B = F^* F, which commutes with lambda(H). Write t-hat = 2 i t_3 of Sp(1)_b. Then, with normalization by the norm of s_v,

```
Q(x_a) = (tr(B Gamma_5)/tr B) Gamma_a/5,
```

and for the functions y_T(g) = tr(T Ad(g) t-hat) with T in the Lie algebra sp(2), Q(y_T) is a multiple of T, proportional to tr(B t-hat)/tr B.

Proof. End_H(4) is spanned by 1, Gamma_5 and t-hat. Since int lambda(g) Gamma_5 lambda(g)^-1 x_a = int x_a x_b Gamma_b = Gamma_a/5, the part with Gamma_5 gives the first formula. Equivariant maps from 5 to End(4) are multiples of Gamma_a, and tr(Gamma_a Ad(g) t-hat) = 0, so 1 and t-hat contribute nothing. The second formula is Schur's lemma on the representation 10. ■

Consequently, a mode with F supported on (1, 2) and K-symmetric has tr(B Gamma_5)/tr B = +-1 and tr(B t-hat) = 0. That is the zero section: Q(x_a) = +- Gamma_a/5, and the part of a function that depends on the fibre gives zero. On the boundary with the pure level-one mode, B is the projection on (1)_{-1}, so that

tr(B t-hat)/tr B = +-1. Example 16.9 has these values.

The operator and its two ends

Proposition 28.8 (endpoint restrictions). Restriction C^infinity(X, R) -> C^infinity(S^4, R) + C^infinity(CP^3, R) is surjective, and in normal coordinates a(y, v) = a(y, 0) + O(|v|), uniformly on bounded subsets of C^1(X). (Report, Proposition 1.)

Theorem 28.9 (a self-adjoint realization). Take the connection metric g = dr^2 + h(r)^2 pi^* g_{S^4} + rho(r)^2 g_{S^2} with h smooth, even and positive at 0, rho odd and smooth with rho(r) = r + O(r^3), and a product near r = l. Let Phi_M = M b(|v|^2) v . J with b smooth and positive, T_M = D_X^E + i Phi_M, and D-hat_M the doubled operator. At r = l set Z = -i c(dr), chi = Gamma-hat Z, and impose chi psi = -psi. Then D-hat_M on this H^1 domain is self-adjoint, elliptic and odd, with compact resolvent, and every smooth scalar function preserves the domain. No boundary condition enters at the zero section: after the outer condition is fixed, removing S^4 creates no further self-adjoint extension. (Report, Theorem 2, using Bär-Ballmann, Corollary 4.18.)

Proposition 28.10 (three slice operators). For the standard complex structure on CP^3: (1) the kernel of the Dolbeault-Dirac operator on O(1) is C^4 for every Hermitian metric; (2) the Levi-Civita Dirac operator twisted by the separate line bundle O(3) with its invariant connection vanishes on the 4 for every invariant squashing, which gives a kernel of dimension at least four, with any further kernel in other representations at a locally discrete set of squashings; (3) the Levi-Civita Dirac operator twisted by the full pull-back connection on E is invertible for small fibres, since (D^E_rho)^2 >= 1/(2 rho^2) - C_0 - C_1 rho^2. (Report, Proposition 3.) In Python, on the representation 4, the smallest |eigenvalue| of that operator is 3.235, 1.842, 1.200, 0.722, 0.148 at rho = 0.3, 0.5, 0.707, 1, 2. So the window comes from the competition between the normal operator and the mass.

Proposition 28.11 (normal model). Let Q = sigma . d_y + y . J on L^2(R^3; C^2 (x) C^4), with J Hermitian of spin 3/2. Then ker Q = 0 and ker Q^* = 3 (+) 1 under the normal Sp(1), with a positive gap above

the zero modes; changing the sign of the mass exchanges the two kernels. The proof uses Q^* Q = -Laplacian + (y . J)^2 - sigma . J >= 0, with equality excluded since sum_i J_i^2 = 15/4, and the Callias index for the positive eigenbundles of n . J, of degrees 3 and 1. The induced bundle over S^4 is S^2(S) (+) 1, the operator on Sigma (x) S^2(S) has exactly four zero modes, in one chirality, and on the parallel inclusion S -> S (x) S^2(S) the twisted equation is a multiple of the twistor equation. So the kernel consists of harmonic spinors of the twisted operator, which are twistor spinors as sections of S. Rescaling gives the localization scale r = O(M^{-1/2}) at the core and O(M^{-1}) at the wall. (Report, Proposition 4.) In Python, in an oscillator basis: Q^* Q starts at 1.02, QQ^* has four near-zero values, with total spin 0 and three values near spin 1, then 1.02.

Lemma 28.12 (tangential mixing). On a product collar, normalized stable profiles for mass eigenvalues M q_i overlap by 2 sqrt(q_i q_j)/(q_i + q_j), independently of M, which is sqrt(3)/2 for q = 1/2, 3/2. So large mass does not remove the coupling between the stable channels O(3) and O(1), which in the pull-back connection has order 1/rho. (Report, Lemma 5.)

Proposition 28.13 (the boundary operator on the representation 4). Take the sign of Phi for which the right singular vector of T is at the zero section. Its left singular vector v is at the wall, in the domain Z v = +v of Theorem 28.9 for T^*. A wall profile exp(q(x - l)), q > 0, for T^* = i Z d/dx + D_6 + i phi tau, with phi > 0, has q = -phi Z tau at leading order. So the stable channels of v are (Z, tau) = (+, -3/2) and (+, -1/2) twice, and they map to the channels (-, -1/2) twice. In unitary bases of these, the compression of D_6(rho) with the overlaps of Lemma 28.12 is

```
A(rho) = ( 3/(2 rho)   1/rho   rho + 1/rho )
         (     0       sqrt 2     -sqrt 2  )
```

with the matrix elements computed in Python, and reproduced by Terra 6.1 in exact symbolic arithmetic from the bracket and Clifford data. Its kernel is spanned by c = (-(2/3)(rho^2 + 2), 1, 1). The channel tau = -3/2 and one channel tau = -1/2 are supported on (1)_{-1} of the 4; the other channel tau = -1/2 is supported on (1)_{+1}. All three lie in Gamma_5 = -1, and their tar-

get spaces are mutually orthogonal, so the cross terms vanish in the traces of Proposition 28.4. Hence for the mode c

```
tr(B Gamma_5)/tr B = -1,    tr(B t-hat)/tr B =
 ↪  -w(rho),    w(rho) = 2(rho^2 + 2)^2 / (9 +
 ↪  2(rho^2 + 2)^2),
```

for every rho, with horizontal scale 1; for horizontal scale h, replace rho^2 by rho^2/h^2. So w = 8/17 as rho -> 0, w = 2/3 at the Fubini-Study radius rho = 1 (the nearly Kähler radius is 1/sqrt 2), and w -> 1 as rho -> infinity, the level-one value. The nonzero singular values of A are the next wall values on the 4: 1.863 and 2.789 at rho = 1, with the core value 2 in between. With the opposite sign of Phi, the right singular vector is at the wall, in the channels Z = -1, tau > 0, with fibre part +w(rho), and its left singular vector at the zero section.

For the connection preserving the line summands, the first row is replaced by (0, 1/rho, rho + 1/rho), the kernel is (1, 0, 0), so w = 1 for every rho, and the block on tau = -1/2 has singular values with sum of squares rho^2 + 2/rho^2 + 6 and product sqrt 2 (rho^2 + 2)/rho, so it is invertible for every rho > 0, with D^2 = 3, 6 at rho = 1. This is the collar repair of the report: an invariant connection preserving the four line summands in a collar extends into X with a cutoff, with the pull-back connection near the zero section. At the Fubini-Study metric the stable boundary model is then a sum of line-twisted Dirac operators, with four zero modes on O(3) and none on O(1), since H^q(CP^3, O(-1)) = 0 for all q.

Status. The kernel and the traces are exact for this finite block, which is the leading formal wall operator; the next correction is O(1/M) in the fixed-block expansion and does not change the limiting kernel, since the kernel is separated from the other singular values. That the boundary mode of T_M converges to it as M -> infinity is a numerical statement here (Result 28.5), since a uniform global reduction is unproved: the difference falls as phi_0 grows, from 0.037 to 0.009 at wall radius 0.5 and from 0.008 to 0.002 at wall radius near 1, between phi_0 = 16 and 32. The value 2/3 tests the factor sqrt 3/2, since without it w(1) would be 3/5.

Lemma 28.14 (a sufficient spectral criterion). Let

U_M : C^8 -> Dom D-hat_M be an isometry with ||D-hat_M U_M|| <= eps_M -> 0 and ||D-hat_M psi|| >= delta ||psi|| for psi orthogonal to the range, with delta > 0 independent of M. Then for large M, D-hat_M has exactly eight eigenvalues in (-delta, delta), and the spectral space is within O(eps_M/delta) of the range. *Proof* (the Astra report, Lemma 6, written out at the request of the second Terra reading of lemma 29, 5 November 2026). D-hat_M is self-adjoint with compact resolvent (Theorem 28.9). Let P be its spectral projection for (-delta, delta). If rank P > 8, the range of P meets the orthogonal complement of the range of U_M nontrivially, against the complement estimate; if rank P < 8, the range of U_M meets ker P nontrivially, against ||D-hat_M U_M|| < delta. So rank P = 8. The spectral theorem gives ||(1 - P) U_M|| <= eps_M/delta, and the equality of the two finite ranks gives a bound of the same order for the difference of the two projections. ■ The missing steps for Conjecture 28.6 are the normal reduction with its errors, and the coercive estimate on every representation of Sp(2).

Conjecture 28.6 (the window). Take Theorem 28.9 with the sign of Proposition 28.13 and with (a) the pull-back connection and wall radius rho, or (b) the collar connection of Proposition 28.13 and the Fubini-Study metric at the wall. For large M the hypotheses of Lemma 28.14 hold, with the eight states being one copy of the 4 at the zero section in one Gamma-hat degree and one at the boundary in the other. For a family of metrics and mass profiles with a uniform positive decay rate, the off-diagonal block is exponentially small in l. The compression of a smooth function tends to (Q_B(a|*{S^4}), Q_Y(a*/{CP^3})), with Q_B and Q_Y(x_a) = -Gamma_a/5, and Q_Y(y_T) a multiple of T, equal to the level-one value of Example 16.9 times w(rho) in case (a), and equal to it in case (b). This concerns compression maps, in line with Counterexample 28.17.

Result 28.5 (Python, the radial system). We solve T on the multiplicity spaces, with Hermitian Clifford generators, so that c(dx) = i Z in the skew convention of Theorem 28.9. For horizontal scale h, after absorbing the half-density h^2 rho, T = i Z d/dx + A_0/h + A_1/rho + (rho/h^2) A_2 + i phi tau, with -s(x) A_E/rho added

for a collar connection. We take h = 1, so that T = i Z d/dx + D_6(rho(x)) + i phi(x) tau, with D_6 built from the structure constants of sp(2), rho(x) = rho_inf tanh(x/rho_inf), phi(x) = phi_0 tanh(x/w_0), w_0 = 0.3, in units with S^4 of radius 1, by the quadratic form of T^* T with P1 elements on a graded grid. This rho is not exactly a product near the wall. Rayleigh quotients of the exact form bound the values from above, and with three-point quadrature they approximate those bounds. Small entries have a mesh floor that depends on the parameters: at (phi_0, l) = (6, 4) the lowest value is 0.00548, 0.00401, 0.00355 for 2500, 5000, 10000 elements, so such entries do not establish a zero continuum value. At x = 0 the exponents are the values for R^3: +-1, +-2 (three times), +-3 on the 4.

phi_0	l = 2	l = 3	l = 4	next singular value, l = 4
3	0.406	0.159	0.066	1.293
4	0.288	0.088	0.028	1.452
5	0.200	0.045	0.011	1.665
6	0.132	0.021	floor	1.818
8	0.049	floor	floor	1.852

(rho_inf = 1, representation 4; the next value at larger phi_0 is 1.86 to 1.87.) The lowest value falls with the exponent of the constant-coefficient system, the smallest |Re kappa| of i Z (D_6(rho) - i phi tau): measured between l = 3 and 4, 0.877, 1.155, 1.464 at phi_0 = 3, 4, 5, against 0.856, 1.151, 1.476, as numerical evidence for these profiles, without a general tunnelling theorem. At large phi, kappa = phi/2 - 2/rho + O(1/phi) (phi = 64: 30.07 at rho = 1, 31.07 at rho = 2). The formula 2 mu exp(-mu l) of lemma 09 belongs to a separable interval and does not carry over.

On the two modes, at rho_inf = 1, l = 3, with the sign of Proposition 28.13, the boundary mode being the left singular vector; the core fibre parts at phi_0 = 16, 40 are below what the computation resolves:

phi_0	zero section: position, Q(x_a)/(Gamma_a/5), tr(B t-hat)/tr B	boundary: position, Q(x_a)/(Gamma_a/5), tr(B t-hat)/tr B, Z, tau
10	0.39, -1.0000, -3e-7	2.93, -1.0000, -0.645, 0.944, -1.117
16	0.28, -1.0000, 1e-11	2.96, -1.0000, -0.658, 0.976, -1.146
40	0.16, -1.0000, 1e-11	2.99, -1.0000, -0.663, 0.996, -1.162

Against Proposition 28.13, at l = 3:

rho_inf	wall radius	-w(rho)	phi_0 = 16	phi_0 = 32	next value, phi_0 = 16	second singular value of A
0.5	0.500	-0.529	-0.493	-0.521	1.992	1.992
0.7	0.700	-0.579	-0.562	-0.575	1.966	1.966
1.0	0.995	-0.665	-0.658	-0.663	1.864	1.865
1.5	1.446	-0.788	-0.785	-0.787	1.660	1.661
2.0	1.810	-0.861	-0.858	-0.860	1.557	1.556
3.0	2.285	-0.921	-0.918	-0.919	1.487	1.486

Collar connection. With rho(l) = 1 and s(x) = (1 + tanh((x - l + 1)/0.15))/2 multiplying the vertical connection term to be removed, at phi_0 = 8, 16, 32: the boundary mode has tr(B t-hat)/tr B = -1.0000 and Q(x_a)/(Gamma_a/5) = -1.0000, the next value is 1.733, that is sqrt 3, and the core mode is unchanged. The switch is nonzero, though tiny, at x = 0, whereas the smooth extension needs a cutoff vanishing near the zero section, so these rounded values support the limiting model and do not establish exact finite-mass equality.

Other representations. The multiplicity is zero on 1, 5, 10, 14, 35. At phi_0 = 16, rho_inf = 1, l = 3, the lowest singular values are 2.188 (wall) and 2.235 (core) on the 16, 1.920 (wall) and 2.828 (core) on the 20, and 3.381 (wall) and 3.463 (core) on the 40, the core values being sqrt 5, sqrt 8 and sqrt 12 up to the mass. The multiplicity zero follows from the central character. The finite wall operator A_V(rho) has no zero at 1500 sampled radii in [0.1, 20] on the 16, 20, 40, 64, a sampled statement; its smallest singular value tends to the core value as rho

-> 0 and falls slowly at large rho on the 16 and the 40 (0.33 and 0.67 at rho = 20). Higher representations were not computed, and no estimate in the representation is proved.

Core block at finite mass. The density of Proposition 28.16 is nonzero at these points, stable in the mesh to five digits:

phi_0, l	lambda	tr(B t-hat)/tr B, integrated	max of abs tr(B t-hat) over max of tr B	sup of the pointwise ratio	integrated value over lambda^2
4, 3	0.0882	-5.71e-3	6.48e-3	0.849	-0.73
5, 3	0.0453	-1.84e-3	1.97e-3	0.856	-0.89
6, 3	0.0212	-4.55e-4	4.73e-4	0.858	-1.01

The pointwise ratio reaches its maximum at the outer wall, where the core density is small, so the breaking comes from the tail of the core mode at the wall. Over phi_0 = 3 to 6, l = 2 to 4, the integrated value over lambda^2 lies between -0.4 and -1.0, an empirical observation with no O(lambda^2) asymptotic proved. These values support the hypothesis of Proposition 28.16 at the sample points, conditional on the density being nonzero in the continuum.

The compression and the algebra

Theorem 28.15 (compression, conditional on localization). Assume an eight-dimensional spectral window splitting under Gamma-hat as W_M (+) V_M with dimensions 4, 4. Then scalar compression is exactly block diagonal. If the isometries differ by at most eta_M from model embeddings with angular compressions Q_B, Q_Y and normal measures with first moments at most C M^{-1/2} at the core and C M^{-1} at the wall, then

```
|| Q_M(a) - (Q_B(a|_B) (+) Q_Y(a|_Y)) || <= C
 ↪  ||a||_{C^1} M^{-1/2} + 4 ||a||_infinity eta_M.
```

(Report, Theorem 7.) For a spectral wall as in lemma 24, the domain has constant boundary values, and if a$|_Y$ = *c then any boundary-localized family has* //Q{Y,M}(a) - c|| <= C M^{-1} ||d_r a||_infinity. So localization just

inside a spectral wall does not recover arbitrary boundary functions for a fixed smooth a, and the limit needs a local wall or an explicitly different comparison algebra.

Proposition 28.16 (the exact symmetry test). For an equivariant copy of 4, B(r) = F(r)^* F(r) = diag(b_a(r) I_2, b_+(r), b_-(r)). All real scalar functions compress to operators commuting with the quaternionic antiunitary j if and only if b_+(r) = b_-(r) for almost every r. If in addition the 5 component is nonzero, the block generates M_2(H). If both the 5 and the 10 components are nonzero, it generates M_4(C). A vanishing integral of tr(B(r) t-hat) alone does not decide this. If a radial h has Q_M(h) = alpha I_W + beta I_V with alpha != beta, then the block projections lie in the generated algebra. (Report, Proposition 8.) At the limit, the real linear image is Herm_2(H) + Herm_4(C), of real dimension 6 + 16 = 22, and the generated real algebra has dimension 16 + 32 = 48.

Counterexample 28.17 (generation is discontinuous under convergence of compression maps). Let R_0 = Q_B E_fib, with E_fib the fibre average on CP^3, R_1 = Q_Y, and R_eps = (1 - eps) R_0 + eps R_1. These are unital completely positive, ||R_eps - R_0|| <= 2 eps, and real functions generate M_4(C) under R_eps for every eps > 0, and M_2(H) at eps = 0. So the correct limit statement concerns specified compression maps, and it gives no convergence of the generated algebras. (Report, Counterexample 9.)

Theorem 28.21 (prescribed windows on the same X). Let P_y = (1 + y_a Gamma_a)/2 and p_z <= P_{pi(z)} the tautological rank-one projector. With normalized radial profiles chi_0, chi_l supported in disjoint eps-neighbourhoods of the two ends, U_0 v = sqrt 2 chi_0 P_y v and U_l v = 2 chi_l p_z v are isometries into L^2(X; C^4) (+) L^2(X; C^4). Real functions compressed by them generate M_2(H) + M_4(C) at every eps, and the compression differs from Q_B (+) Q_Y by at most C eps ||a||_{C^1}. (Report, Theorem 13.) This proves the connected-space mechanism without Conjecture 28.6, with the window as input. Allowing nonlocal smoothing perturbations, (1 - P) D (1 - P) + P d P makes such a window reducing, which enlarges the operator class and keeps the window as input.

The finite data

Proposition 28.18 (charge conjugation and the wall). With Hermitian generators J_i, C = C_7 (x) j_E has C^2 = -1 and commutes with T_M, D-hat_M and Gamma-hat as a differential expression, and CZ = -ZC. So it exchanges the two local wall conditions and does not preserve the chosen domain. If a pointwise parent antiunitary preserves a window and commutes with real multiplication, every real compressed multiplier commutes with its restriction, which cannot hold on a generated full M_4(C). (Report, Proposition 10.) The complex boundary block therefore requires this loss of the inherited symmetry, or a different real correspondence.

Remark 28.7 (why one block is quaternionic and the other complex). Proposition 28.18 proves the signs argument of the first version, with one correction of wording: j_E commutes with the anti-Hermitian Lie algebra action and reverses the Hermitian generators, so that C commutes with i Phi. Numerically (Python, phi_0 = 16, l = 3, rho_inf = 1), exchanging the wall condition exchanges (Z, tau, tr(B t-hat)/tr B) = (0.976, -1.146, -0.658) and (-0.976, 1.146, 0.658) on the boundary mode, with the same singular values, and leaves the zero-section mode unchanged. The formal parent symmetry does not by itself prove that the finite-mass core space is preserved, since C changes the wall domain, and the breaking computed in Result 28.5 shows the size of the effect at the sample points.

Proposition 28.19 (tunnelling and the grading). If the window consists of one copy of the 4 in each Gamma-hat degree, then d_M = [[0, conj(lambda_M) I_4], [lambda_M I_4, 0]] after equivariant unitary identifications, by Schur's lemma, and lambda_M = (1/4) Tr(U_{Y,M}^* T_M U_{B,M}). The inherited grading is I_W (+) (-I_V), the auxiliary grading of the doubling, which labels the two ends under localization. The radial Clifford involution Z does not extend across the zero section, so it does not serve as an interval chirality. This grading leaves all four core states in one degree, so W = H (+) H, for the even algebra H + H + M_4(C), still needs a chosen quaternionic grading, such as a marked Gamma_p. (Report, Proposition 11.)

Proposition 28.20 (the doubled bimodule). With H

= W (+) V, Theorem 16.2 gives K = H (x) conj(H), of dimension 64, with DD = d (x) 1 + Gamma-hat (x) d^c. Its odd part is exactly K^- = Hom(V, W) (+) Hom(W, V), with the desired bimodule actions, and order zero and order one hold exactly on the full K, for its actual even algebra. If lambda != 0, then DD maps K^- into K^+, P_- DD P_- = 0 and DD^2 = 2 lambda^2. So K^- is not a reducing subspace. On K^- the inherited tensor grading is constant; the sheet-sign grading of Theorem 16.7 gives J(x, y) = (y^*, x^*), J^2 = 1, J Gamma_6 = -Gamma_6 J, with zero Dirac operator. The tunnelling map W -> V does not provide the map Hom(V, W) -> Hom(W, V) of lemma 20. (Report, Proposition 12.)

Balance against Proposition 6 of lemma 24

(Report, Proposition 14, with the numerical statements updated.)

Structure	What follows
Connected scalar parent	Supplied by the chosen disc bundle X.
Independent endpoint functions	Derived, by restriction and extension.
Base-function restriction at the collapsed end	Derived geometrically; quantitative for localized states.
Four core states	Proved for the linear normal model and the resulting twisted operator on S^4.
Four boundary states	Proved for the separated collar model; for the pull-back connection, the finite block of Proposition 28.13, with the convergence numerical.
Isolated 4 + 4 window	Conjecture 28.6.
Limiting quaternionic core compression	Derived once convergence to the core model is established.

Structure	What follows
Exact finite-mass quaternionic block	Requires b_+ = b_-; numerically not satisfied at the sample points resolved.
Exact level-one boundary compression	Requires rho -> infinity or the separated collar connection.
End grading	Inherited from the chosen self-adjoint doubling.
Quaternionic grading axis	Input.
The two Hom spaces	The odd part of the coarse correspondence; the selection is additional.
KO-6 reality and order conditions	Exact after the specified correspondence and grading choices.
Nonzero connector of lemma 20	Input.
Mass value and dependence on l	Radial matching; equivariance gives only lambda I_4; the exponent is computed numerically in Result 28.5.

Choices that stay inputs: X, E = S^3(S), the hedgehog, the wall condition, the collar connection in case (b), and the doubling of Theorem 16.2. A double with two identical collapsed zero sections removes the outer boundary condition, but the same local core construction gives two quaternionic endpoint maps, so obtaining a complex block needs different structure at one end. T0 stays open, and this is a candidate for its first half.

Direction. The next analytic step is the coercive estimate of Lemma 28.14, in case (b) first, where both end models have exact kernels. Since Phi commutes with the Clifford symbol, D + iM Phi is a concentrating pair in the sense of Maridakis (arXiv:1510.07004, Definition 2.1 and Lemma 2.2), with [D, Phi] of order zero; the symbol condition goes back to Prokhorenkov and Richardson (arXiv:math/0307251). Two things are missing there. The spectral separation theorem of arXiv:2110.11654 (section 2, Assumption 3) needs a leading endomorphism

independent of the normal direction, and this fails here, since (n . J)^2 depends on n for spin 3/2; so a separate core analysis or an anisotropic extension is needed. And the wall adds a boundary layer: on a half-line ||(d_t + M) f||^2 = ||f'||^2 + M^2 ||f||^2 - M |f(0)|^2, and exp(-Mt) is a zero mode away from the zero set of the mass. (Read by Terra 6.1 in arXiv HTML; the LaTeX sources were not available to it.)

Numerical check

By Claude Opus 5.5, in Python (numpy, scipy), 3 Nov 2026, in a script that ran the tests below as assertions and printed the tables of Result 28.5; by the rule of the project the script is not kept (removed on 4 Nov 2026; Terra 6.1 had reproduced its matrix of Proposition 28.13 in exact arithmetic). The slice operator on each multiplicity space is Hermitian, odd under Z, and closes on Hom_H. The structure constants reproduce [t_1, t_2] = t_3 in Sp(1)_a and Sp(1)_b. The decomposition D_6(rho) = A_0 + A_1/rho + A_2 rho holds. The form uses three-point Gauss quadrature on each element. The exact tests are {A_0, A_1} = 0, A_0^2 = 0, 4 on the 4, 5, 9 on the 16, 0, 8 on the 20, 12, 16 on the 40 (the value 0 on the 20 has the twist (1, 5) and gives no core zero mode), A_2 = -(i/8) sum C_ab^alpha c_a c_b c_alpha, and at rho = 1 with the line connection the values D^2 = 0, 5, 5, 12 on the summand O(3) of the 4, 16, 20, 40, against j(j + 4) for (0, 0)-forms in O(1); the whole summand also has the value 10 on the 20. The closed form of Proposition 28.13 is tested to ten digits, and the core mode density in the mesh.

Part 29

The eight-state window on the disc bundle

Written by Claude Opus 5.5 on 4 November 2026, for STRATEGY T0, without an Astra job. It continues lemma 28 and answers the prompt 29_window_coercive_estimate.md. Refereed by Terra 6.1 (report): case (b) stands with amendments, and the two errors it found in case (a), the definition of A_eff and the sign of the fibre part, are corrected below. The referee ran no computation. A fresh reading of the corrected case (a) in

the referee report on lemma 30 found one more sign and one missing step in Proposition 29.10; these are corrected. A second, fresh reading by Terra 6.1 (report, 5 November 2026, with the first reports withheld) found two errors, both repaired below: the range and quantifiers of delta (delta = 0 was a counterexample to the wording) and the sign of the Clifford commutator in the IMS step (the identity itself was right); it confirmed the chain of case (b) with amendments, the corrected case (a) with two added steps, and listed what a stand-alone paper needs (section "Second reading" below).

Outcome

1. **Conjecture 28.6 holds in case (b)** for the class of metrics and mass profiles of Setting 29.0: for each fixed 0 < delta < sqrt 3 there is M_delta such that, for M >= M_delta, D-hat_M has exactly eight eigenvalues in (-delta, delta) (Theorem 29.7).
2. **Conjecture 28.6 holds in case (a) at the Fubini-Study radius** rho = 1, with fibre part -2/3 (Theorem 29.11). At another wall radius, case (a) follows from one kernel condition K(rho), which reduces at each fixed radius to finitely many representations by ellipticity; an explicit bound identifying those representations remains to be written.
3. **The off-diagonal block is exponentially small**, |lambda_M| <= C exp(-(1 - eta)(M/2) int f dr), by an exact tunnelling identity and Agmon estimates (Proposition 29.8).
4. **The compression converges** to (Q_B(a|*{S^4}), *Q_Y(a/*{CP^3})) with error O(M^{-1/2}) (Theorem 29.9). The statement concerns compression maps, in line with Counterexample 28.17.

The proof is short because Sp(2)-equivariance gives exact sums of squares at the core and at the wall. Each estimate holds on all of L^2 of a fixed manifold, with the boundary operators elliptic on CP^3 or S^4, so uniformity over the representations of Sp(2) needs no separate argument. The spectral separation hypothesis of Maridakis, which fails for this mass (historical/reviews/2026-11-03_terra_check_28.md, item 6), is never used. General forms of two identities are in the literature: the square with its zero-order cross term in the bulk (Maridakis, arXiv:1510.07004, Lemma 2.2), and a horizontal

and vertical factorization with a curvature contraction (Kaad and van Suijlekom, arXiv:1710.03182, Proposition 18, under other hypotheses).

Scope. These are spectral statements for a chosen metric family, with flat fibres near the zero section. They give the finite window from the operator, for the sign of Proposition 28.13. They leave open: other metrics in the class of Theorem 28.9, including product-collar modifications of Bryant-Salamon-type metrics, and the complete Bryant-Salamon geometry (Theorem 28.9 assumes a product metric near a finite wall); the connector of Proposition 28.20; the quaternionic grading of Proposition 28.19; any dynamics selecting M, l or rho. Finite-mass generated algebras stay outside the statement (Counterexample 28.17).

Setting

Setting 29.0. Let X be the disc bundle of lemma 28, G = Sp(2), K = Sp(1)_a x Sp(1)_b, H = Sp(1)_a x U(1)_b, E = S^3(S), Phi_M = eps M f(r) n . J, and T_M = D_X^E + i Phi_M, with D-hat_M the doubled operator and the wall condition of Theorem 28.9: Z u = -u for T_M, and Z v = +v for T_M^*. Assume:

- (H1) g = dr^2 + pi^* g_{S^4} + rho(r)^2 g_{S^2}, with S^4 of radius 1, rho(r) = r on [0, 2e], and rho constant, equal to rho_l, on the collar [l - 2a, l]. Case (b) and Theorem 29.11 take rho_l = 1, the Fubini-Study metric.
- (H2) f(r) = r on [0, 2e], f = f_0 > 0 on the collar, and f > 0 on (0, l]. This is M b(|v|^2) v . J with b positive.
- (H3) The connection on E is the pull-back connection on r <= 2e. In case (b) it preserves the four line summands on the collar, with any invariant interpolation in between. In case (a) it is the pull-back connection throughout. We use radial parallel gauge: nabla_{d/dr} = d/dr. An invariant dr-component a(r) q could otherwise be added to a line-preserving connection; it is removed by radial parallel transport, which changes the displayed formulas by a phase on each line.
- (H4) eps = -gamma_b, where gamma_b is the chirality of Sigma_4 whose K-type is (1, 2). This

is the sign of Proposition 28.13: its right singular vector is at the zero section. The opposite sign exchanges T_M and T_M^* throughout. In the convention of the radial computation of lemma 28, gamma_b = +1.

On the collar put q = eps n . J, with eigenvalues +-3/2, +-1/2 on the line summands, and m = M f_0. In general q = -gamma_b n . J. In the notation of Proposition 28.13 and of the radial computation of lemma 28, q = -tau, and T_M = i Z d/dx + A + i m q on the collar, with x increasing towards the wall. The line q = +3/2 is the summand O(3), and q > 0 are the stable channels of Proposition 28.13.

Let C_R = ||A_2||, the norm on Sigma_6 (x) E of the curvature contraction A_2 = -(i/8) sum C_ab^alpha c_a c_b c_alpha; it does not depend on V, and C_R = 1 in Python (a finite C_R is all the proofs use). Let mu_1 > 0 be the smallest positive eigenvalue of Q^* Q for Q = sigma . d_y + y . J of Proposition 28.11. It is positive since the resolvent is compact, because (y . J)^2 >= |y|^2 / 4, and ker Q = 0: with M = 1, Q^* Q = (-Delta + |y|^2/4 - 3/2) + ((y . J)^2 - |y|^2/4) + (3/2 - sigma . J), three nonnegative terms; equality in the first forces a Gaussian times a constant vector, and equality in the second would then force J_i^2 v = v/4 for all i, against sum_i J_i^2 = 15/4 (referee). Python gives 1.02.

The zero section

Proposition 29.1 (exact splitting). Identify sections over the tube r <= 2e with K-equivariant maps s : G -> F, F = Sigma_4 (x) L^2(R^3; C^2 (x) C^4), with c(e_a) = gamma_a (x) 1 and c(f_alpha) = gamma_5 (x) sigma_alpha. Let D_H = i sum gamma_a e_a, with e_a acting as left-invariant derivatives on G, and Q^ = gamma_5 sigma . d_y + eps M y . J. Then on this tube T_M = T_0 + R, with T_0 = D_H + i Q^ and R Hermitian of order zero with |R| <= C_R |y|, and

```
T_0^* T_0 = D_H^2 + Q^^* Q^,    T_0 T_0^* = D_H^2 +
↪  Q^ Q^^*,
```

with D_H^2 commuting with Q^^* Q^ and Q^ Q^^*.

Proof. The fibres are flat and totally geodesic by (H1),

so the Levi-Civita connection of X differs from the connection associated with the principal K-bundle G -> S^4 only by the O'Neill tensor. Its component is the curvature of N evaluated at v. This gives the zero-order term R, which equals r A_2 in the radial form D_6 = A_0 + A_1/rho + A_2 rho, of norm at most C_R |y|. For the pull-back connection, the remaining operator is D_H + i gamma_5 sigma . d_y + i Phi_M. The operator Q^ acts on the values of s, and it is K-equivariant, since sigma . d_y and y . J are invariant under the simultaneous rotation of y, of C^2 and of C^4. On the ambient space of fibre-valued functions on G, Q^ commutes with each e_a, because it does not depend on the point of G; the individual e_a need not preserve the K-equivariant subspace, but both Q^ and the summed operator D_H do. Since gamma_a anticommutes with gamma_5 and commutes with y . J, D_H Q^ = Q^~* D_H, where Q^~* = -gamma_5 sigma . d_y + eps M y . J. Therefore

```
T_0^* T_0 = (D_H - i Q^~*)(D_H + i Q^) = D_H^2 +
 ↪  Q^~* Q^ + i (D_H Q^ - Q^~* D_H) = D_H^2 + Q^~*
 ↪  Q^,
```

and likewise for T_0 T_0^*. Finally D_H Q^~* Q^ = Q^ Q^~* D_H, so D_H^2 commutes with both. ■

In radial form. T_0 = i Z d/dx + A_0 + A_1/rho + i phi q, and the cross terms vanish by {A_0, Z} = 0, [A_0, tau] = 0 and {A_0, A_1} = 0. these three identities were asserted on the representations 4, 16, 20, 40 (Python, script not kept).

Proposition 29.2 (the model spectrum). On the total space of N, with f = r and T_0 as above, let Pi be the projection onto ker T_0. Then:

1. ker T_0 is one copy of the 4 of G; ker T_0^* = 0.
2. ||T_0 u||^2 >= min(4, mu_1 M) ||u||^2 for u orthogonal to ker T_0, and ||T_0^* v||^2 >= min(4, mu_1 M) ||v||^2 for all v in the domain.

Proof. Write Q^_gamma for Q^ on chirality gamma. By Proposition 28.11, Q^_gamma has kernel exactly when gamma = -eps; this is the rank-four K-representation W_0 = (1, 3) + (1, 1), and on the orthogonal complement Q^~* Q^ >= mu_1 M, after the rescaling y -> M^{-1/2} y. The same bound holds on the other chirality, and for Q^ Q^~* with the chiralities exchanged,

since Q^* Q and Q Q^* have the same nonzero spectrum. The fibrewise projection P_W onto W_0, on both chiralities of Sigma_4, commutes with D_H, by Proposition 29.1. (The projection onto the fibre kernel of Q^ also selects gamma = -eps, so it does not commute with D_H, which is odd; it commutes with D_H^2.) On chirality -eps, Q^ vanishes on the range of P_W and is bounded below by sqrt(mu_1 M) on its complement; on chirality +eps, it is bounded below by sqrt(mu_1 M) throughout. On Sigma^gamma (x) W_0, D_H is the Dirac operator of S^4 twisted by the homogeneous bundle of W_0, so by Parthasarathy's formula D_H^2 = Cas_G(V) + 3/2 - Cas_K(tau), with Cas_K = 2 j (j + 1) on Sp(1)_b, that is 0 and 4 on W_0. Only V odd under -1 occur, with Cas_G(V) >= 5/2 and Cas_G(V) >= 15/2 for V not the 4. So D_H^2 >= 4 except for D_H^2 = 0 at V = 4, twist j = 1. This kernel needs Hom_K(4, Sigma^gamma (x) (1, 3)) != 0, which holds exactly for Sigma^gamma = (1, 2), where (1, 2) (x) (1, 3) contains (1, 2) once. By (H4), gamma_b = -eps, so T_0 sees this chirality and T_0^* the other. Decomposing u by chirality and by P_W, and using the sum of squares, gives both bounds. The connection metric on the total space N is complete, and the doubled model is a symmetric Dirac-type operator plus a smooth Hermitian zero-order potential, hence essentially self-adjoint on C_c^infinity(N); the identities and estimates extend from this core to its closure, and no boundary condition enters at infinity. The S^4 spectrum has a direct check in the literature: Dolan's homogeneous-instanton calculation (hep-th/0304037, section 2(ii), equations (40)-(46)) gives four zero modes for isospin 1 and the next squared value 5, with the untwisted squared spectrum starting at 4 (referee). The equivariant Callias calculation that identifies ker Q^* as 3 + 1, and not only its dimension, is stated here with the conventions of Proposition 28.11 and is one of the items a stand-alone paper must display. ■

Lemma 29.3 (core estimate). Let u be in H^1 with support in r < 2e. Then

```
|| |y| u || <= (2/M) ||T_0 u|| + sqrt(10/M) ||u||,
```

and the same holds with T_0^*. So ||T_M u|| >= (1 - 2 C_R/M) ||T_0 u|| - C_R sqrt(10/M) ||u||, and likewise

for T_M^* with T_0^*. In particular, for such u with Pi u = 0, and for all such v,

```
||T_M u|| >= (2 - o(1)) ||u||,    ||T_M^* v|| >= (2
↪  - o(1)) ||v||,    M -> infinity.
```

Proof. Q^^* Q^ = -Laplacian_y + M^2 (y . J)^2 - eps M gamma_5 sigma . J. The eigenvalues of sigma . J are 3/2 and -5/2, and (y . J)^2 >= |y|^2/4. Hence M^2 || |y| u ||^2 / 4 <= ||Q^ u||^2 + (5/2) M ||u||^2 <= ||T_0 u||^2 + (5/2) M ||u||^2, by Proposition 29.1. The rest is the triangle inequality with |R| <= C_R |y|, and Proposition 29.2. ■

The bulk

Lemma 29.4 (bulk estimate). Let w be in H^1 with support in a compact subset of 0 < r < l, and let f_1 be the minimum of f on that support. Then

```
||T_M w||^2, ||T_M^* w||^2 >= (M^2 f_1^2 / 4 - C_Phi
↪  M) ||w||^2,
```

with C_Phi the supremum of the operator norm of the Clifford contraction sum_i c(e_i) nabla_{e_i} (f n . J). With 0 < r_1 < r_2 < l and a radial cutoff psi equal to 1 on [r_1, r_2], with compact support in 0 < r < l, we have for u in Dom T_M,

```
||u||^2_{L^2(r_1 <= r <= r_2)} <= C M^{-2} (||T_M
↪  u||^2 + ||u||^2).
```

Proof. ||(D + i Phi) w||^2 = ||D w||^2 + ||Phi w||^2 + <w, i [D, Phi] w>, with [D, Phi] of order zero and pointwise bounded independently of V, and Phi^2 >= M^2 f^2 / 4. For the second statement, choose 0 <= psi <= 1 equal to one on the compact set under consideration; the first estimate applied to psi u with T_M (psi u) = psi T_M u + [D, psi] u gives ||psi u||^2 <= 2 (||T_M u||^2 + ||d psi||_infinity^2 ||u||^2) / (M^2 f_psi^2/4 - C_Phi M), with f_psi the minimum of f on the support of psi. Applied to compact sets containing the supports of the d chi_i and of chi_b, this gives the O(M^{-2}) bounds used in the transition regions below. ■

The interpolation of the connection in (H3) takes place in this region, where it enters only through C_Phi.

The wall

Lemma 29.5 (line-twisted operators at the Fubini-Study metric). Let D_k be the Dirac operator on CP^3 twisted by O(k) with its invariant connection, k odd, at rho = 1. Then D_k^2 = C_2(L) + 3/2 - (3/8) k^2 on the SU(4)-type L, with C_2 normalized by C_2(4) = 15/8. Consequently:

1. On O(3): ker D_3 is one copy of the 4 of Sp(2), in Z = +1; and D_3^2 >= 5 on the V-isotypic part for V not the 4.
2. On O(1): ker D_1 = 0; D_1^2 >= 3 on the 4 and D_1^2 >= 6 on the V-isotypic part for V not the 4.

So the smallest nonzero |eigenvalue| on O(3) + O(1), over all V, is kappa_w = sqrt 3, attained on the 4; the lowest wall value on the 16 is sqrt 5.

Proof. The form of D_k^2 is Parthasarathy's formula for SU(4)/U(3), since the twist is a character of U(3). The Sp(2)-invariant connection on a line summand is unique, because the isotropy representation has no trivial summand, so it is the Chern connection. The normalization: j(j + m + 3) is the eigenvalue of 2 box on (0,0)-forms in O(m) on the type [m + j, 0, j], with C_2([m + j, 0, j]) - C_2([m, 0, 0]) = j (j + m + 3), which fixes the coefficient of C_2; scalar curvature 12 gives the constant 3/2; and the term in k^2 is fixed by the kernel of D_3 on the 4. Two checks: D_0^2 = 4 on the type 6, the Kirchberg limit R/3 at R = 12 (Ann. Global Anal. Geom. 4 (1986) 291-325); and the values in the Python table below. The kernel statements are H^q(CP^3, O(1)) = C^4 at q = 0 only, and H^q(CP^3, O(-1)) = 0, with K^{1/2} = O(-2). The types of odd 4-ality with C_2 < 55/8 are the 4, its conjugate, the 20 and its conjugate (enumerated in Python). The 4 occurs on O(3) once, by the multiplicity on the slice, and this copy is the kernel; since each 20 restricts to Sp(2) as 4 + 16, neither 20 occurs on O(3). So V != 4 has C_2 >= 55/8 on O(3), giving D_3^2 >= 5. On O(1), the 4 and its conjugate restrict only to the 4 of Sp(2), so V != 4 has C_2 >= 39/8, giving D_1^2 >= 6; and on the 4 Frobenius reciprocity gives the relevant SU(4) types 4 and 20, each in both chiralities, with Casimirs 15/8 and 39/8, hence the two squared values 3 and 6 on the slice block. The chirality of ker D_3 is

that of (0,0)-forms, and the slice block fixes it as Z = +1 for q = +3/2 (Python). The line-twisted spectrum on odd complex projective spaces is in print: Ben Halima, Comment. Math. Univ. Carolin. 49 (2008), Theorem 1, with the tautological line O(-1) and scalar curvature 3; taking its exponent m = -k and multiplying its squared eigenvalues by 4 for scalar curvature 12 reproduces the zero and the gap 5 for k = 3 and the gap 3 for k = 1 (referee); the Sp(2) branching and the Z-sign are recorded here in the present conventions. ■

Lemma 29.6 (wall estimate in case (b)). Let w be in H^1 with support in the collar, vanishing near its inner end, and p(x) = sqrt(3 m) exp((3/2) m (x - l)), so that the kernel of T_M^* on the half-cylinder is p (x) ker D_3. Then:

1. If Z w = -w at r = l, then ||T_M w||^2 >= min(m^2/4, 3) ||w||^2.
2. If Z w = +w at r = l, and w is orthogonal to p (x) ker D_3, then ||T_M^* w||^2 >= min(m^2/4, 3) ||w||^2.

Proof. On the collar, A commutes with q and with d/dx, and anticommutes with Z. Writing w = w_+ + w_- by Z, integration by parts gives, for T_M,

```
||T_M w||^2 = ||(d + m q) w_+||^2 + ||A w_+||^2 +
↪  ||(d - m q) w_-||^2 + ||A w_-||^2,
```

where the sum of the real parts of the mixed terms is the boundary term obtained by integrating d/dx <w_+, A w_->, which vanishes because w_+(l) = 0 (the individual mixed terms need not be purely imaginary). For T_M^* the same holds with q replaced by -q and with w_-(l) = 0. There are three one-dimensional cases on (-infinity, l]. (i) If g(l) = 0, then ||(d +- m q) g||^2 = ||g'||^2 + m^2 q^2 ||g||^2. (ii) With g(l) free, ||(d + m |q|) g||^2 = ||g'||^2 + m^2 q^2 ||g||^2 + m |q| |g(l)|^2. (iii) With g(l) free, ||(d - m |q|) g||^2 is the form of -d^2 + m^2 q^2 with a Robin condition at l, whose spectrum consists of 0, on the profile sqrt(2 m |q|) exp(m |q| (x - l)), and of [m^2 q^2, infinity). Cases (i) and (ii) give the lower bound m^2/4. Case (iii) arises exactly in the channels (Z, q) = (-, q > 0) for T_M and (+, q > 0) for T_M^*. On those, expand in eigenvectors of A^2: the term ||A w||^2 is at least 3 ||w||^2, except on ker D_3 in

Z = +1, which by Lemma 29.5 only arises for T_M^* at q = 3/2, where the Robin gap (9/4) m^2 applies. ■

The window

Theorem 29.7 (the window, case (b)). Under Setting 29.0 with rho_l = 1 and connection (b), for each fixed 0 < delta < sqrt 3 there is M_delta such that, for M >= M_delta, D-hat_M has exactly eight eigenvalues in (-delta, delta). The spectral subspace splits under Gamma-hat as W_M (+) V_M, each a copy of the 4 of Sp(2), with W_M within O(M^{-1/2}) of the cut-off core modes in the domain of T_M, and V_M within O(M^{-1/2}) of the cut-off wall modes p (x) ker D_3 in the domain of T_M^*.

Proof. We verify Lemma 28.14. Take radial functions chi_c, chi_b, chi_w with chi_c^2 + chi_b^2 + chi_w^2 = 1, chi_c = 1 on r <= e with support in r < 2e, chi_w = 1 on [l - a, l] with support in the collar. Let U_M span (chi_c u_i, 0) and (0, chi_w v_k), where u_i is a basis of ker T_0 and v_k of p (x) ker D_3, with Gram-Schmidt. Radial functions preserve the wall condition.

Trial states. T_M (chi_c u_i) = [D, chi_c] u_i + R chi_c u_i, of norm O(M^{-1/2}), since u_i has Gaussian decay at the scale M^{-1/2}; and T_M^* (chi_w v_k) = [D, chi_w] v_k, which is O(exp(-(3/2) m a)). So eps_M = O(M^{-1/2}).

Coercivity. Since [D, chi] = c(d chi) is skew-adjoint with c(d chi)^* c(d chi) = |d chi|^2 (its square is -|d chi|^2: for instance [D, x] = i Z on the collar), while sum_i chi_i d chi_i = 0,

```
||T_M u||^2 = sum_i ||T_M chi_i u||^2 - sum_i || |d
↪  chi_i| u ||^2,
```

and likewise for T_M^*. The functions d chi_i vanish outside the bulk region e <= r <= l - a. So by Lemma 29.4,

```
(1 + C M^{-2}) ||T_M u||^2 >= ||T_M chi_c u||^2 +
↪  ||T_M chi_w u||^2 - C M^{-2} ||u||^2,    ||chi_b
↪  u||^2 <= C M^{-2} (||T_M u||^2 + ||u||^2).
```

If u is orthogonal to every chi_c u_i, then Pi (chi_c u) = 0 exactly, since <u_i, chi_c u> = <chi_c u_i, u>, so Lemmas 29.3 and 29.6 (1) give ||T_M u||^2 >= (3

- o(1)) ||u||^2. If v is orthogonal to every chi_w v_k, then chi_w v is orthogonal to p (x) ker D_3, so Lemmas 29.3 and 29.6 (2) give the same for T_M^*. Since ||D-hat (u, v)||^2 = ||T_M u||^2 + ||T_M^* v||^2, Lemma 28.14 gives the count and the distance to the range. The window is G-invariant and close to the range of U_M, so as representations W_M and V_M are copies of the 4. ■

Remark. The gap is set by the wall at sqrt 3, with the core next at 2. On the 16, 20 and 40 the lowest values are sqrt 5, sqrt 5, sqrt 12.

Proposition 29.8 (tunnelling). Let T_M u = lambda v and T_M^* v = lambda u be a singular pair of the window, with unit vectors. By Proposition 28.19, the four singular values coincide. For a radial chi equal to 1 near r = 0 and 0 near r = l,

```
lambda (<u, chi u> - <v, chi v>) = <v, c(d chi) u>.
```

For every eta in (0, 1) there are M_eta and C_eta with

```
|lambda_M| <= C_eta exp(-(1 - eta)(M/2) int_e^{l -
↪  a} f dr),    M >= M_eta.
```

Proof. The function chi u lies in Dom T_M, since chi vanishes near the wall. So <T_M^* v, chi u> = <v, T_M chi u>, and T_M chi u = chi lambda v + c(d chi) u. This gives the identity.

For the decay of u, let xi = psi exp(g) be radial, with psi = 0 on r <= e/2 and psi = 1 on r >= e, with g = 0 on r <= e, g' = (1 - eta)(M/2) f on [e, r_*], and g constant on r >= r_*. Since xi^2 u lies in Dom T_M,

```
||T_M (xi u)||^2 = lambda^2 ||xi u||^2 + || |d xi| u
↪  ||^2.
```

Choose the partition so that the derivative of the Agmon weight is supported where the bulk cut-off is one. The IMS lower bound, with no core piece and with Lemma 29.6 (1), or Proposition 29.10 in case (a), on the collar, reads int (M^2 f^2/4 - C M) chi_b^2 |xi u|^2 + kappa_w^2 ||chi_w xi u||^2 - sum_i || |d chi_i| xi u ||^2 <= lambda^2 ||xi u||^2 + || |d xi| u ||^2, and the transition terms are controlled by the bulk estimate of Lemma 29.4. On the bulk the excess over (g')^2 is at least (eta/4)(M^2 f^2 / 4) for M large, with f >= f_1 > 0, so ||exp(g) u||*{L^2(r >= e)}* <= *C_eta. For v*

we use the same argument from the wall inward, with g-tilde = 0 on [l - a, l], with -g-tilde' = (1 - eta)(M/2) f, and with Lemma 29.3 replacing Lemma 29.6 on the core, since T_M^* has no low states there. Take chi with /d chi/ <= 2/s, supported on [r_, r_ + s] in the bulk. By Theorem 29.7, <u, chi u> -> 1 and <v, chi v> -> 0, so the identity gives /lambda/ <= (4/s) //u//{[r_, r_ + s]} ||v||{/[r_, r_ + s]}. Run the weights with eta' = eta/2, and take s so small that (1 - eta')(I - int{r_}^{r_ + s} f) >= (1 - eta) I, with I = int_e^{l - a} f. This gives the stated exponent. The absorption uses lambda -> 0, supplied by the trial states. The factor 1/2 holds on every channel, since (n . J)^2 >= 1/4; for the adjoint, the cut-off vanishes near the wall and the core estimate replaces the wall estimate. For a family with l -> infinity the collar widths, the cut-off derivatives, the connection bounds and the positive lower bound of f on the bulk have to be uniform, and the excluded end lengths are absorbed in the spare eta (referee). ■

For a family with l -> infinity, M >= M_eta, with uniform bounds on the end geometry, the cutoffs, the connection and the positive bulk mass, the constants are uniform, so lambda is exponentially small in l at the rate (1 - eta)(M/2) inf f. This is the uniform decay rate required by Conjecture 28.6, as amended by the referee.

Theorem 29.9 (compression). Under Theorem 29.7 or Theorem 29.11, for a in C^1(X),

```
|| Q_M(a) - (Q_B(a|_{S^4}) (+) Q_Y(a|_{CP^3})) || <=
  ↪  C ||a||_{C^1} M^{-1/2},
```

where Q_B(a) = int a(x) 2 P_x d mu(x), with P_x = (1 - x . Gamma)/2, and Q_Y(a) = int a(z) B(z) d mu(z). In case (b), B(z) = 4 p_z, the level-one map of Example 16.9. In case (a) at radius rho, B(z) = 2((1 + w) p_z + (1 - w) p_{jz}), with p_{jz} the quaternionic partner of p_z and w = w(rho); so Q_Y(x_a) = -Gamma_a/5, Q_Y(y_T) is w(rho) times the level-one value, and the signed fibre part is -w(rho). Restriction to the two ends is surjective (Proposition 28.8), so real functions on X generate M_2(H) + M_4(C) under the limit map, by lemma 10 and Example 16.9, for every w(rho) > 0.

Proof. Theorem 28.15, with eta_M = O(M^{-1/2}) from Theorems 29.7 and 29.11, after choosing equivariant uni-

tary identifications of the window spaces with the model copies of the 4. The normal first moments are O(M^{-1/2}) at the core, by the Gaussian decay of the model, and O(M^{-1}) at the wall, by the profiles. At the core, F in Hom_K(4, (1, 2) (x) (1, 3)) kills (2, 1), so B = 2 P_x in the normalization Q_B(1) = I. The sign in P_x is the value tr(B Gamma_5)/tr B = -1 of Proposition 28.13. At the wall, the pointwise inner product of different line summands vanishes, so the density is a sum over channels with weights from Proposition 28.13. ■

Case (a)

On the collar in case (a), A = A_d + A_o, where A_d = D_line preserves the lines and A_o = A_E/rho has order zero, with [A_o, q] != 0. In Python, the block of A_E between q = 3/2 and q = 1/2 has norm sqrt 3 and ||A_E|| = 2, on the full fibre, hence for all V. Let P be the projection onto {sum_{q > 0} p_q g_q : Z g_q = -g_q}, with p_q = sqrt(2 m q) exp(m q (x - l)), and Y the same with Z = +1, with Pi_Y the projection onto the second space. Let F = Pi_Y (A_d + A_o) P, from the range of P to Y, and let A_eff be the operator on the sum of both ranges with blocks F and F^*. Identified with sections over CP^3 of the lines q > 0, it is, by Lemma 28.12,

```
A_eff = D_3(rho) (+) D_1(rho) + (sqrt 3 / 2) B(rho),
```

with B(rho) the block of A_o between the stable channels, of norm sqrt 3 / rho. This is self-adjoint and elliptic on the fixed manifold CP^3, so its nonzero spectrum has a gap kappa_a(rho) > 0, uniformly in V. On the 4, F^* is the matrix A(rho) of Proposition 28.13, from Y to the range of P.

Condition K(rho). The kernel of A_eff(rho) is exactly one copy of the 4, in Z = +1, so F is injective and ker F^* is that copy.

Proposition 29.10 (wall estimate in case (a)). Assume K(rho). For w as in Lemma 29.6, with the kernel p (x) ker D_3 replaced by the kernel of the half-cylinder operator T_M^*, the two statements of Lemma 29.6 hold with the constant 3 replaced by kappa_a(rho)^2 - O(1/m). The half-cylinder kernel is exactly one copy of the 4 for m large, with fibre part tending to -w(rho) at

rate O(1/m).

Proof. T_d = i Z d/dx + A_d + i m q preserves the block structure: T_d P = Pi_Y T_d P, since T_d (p_q g) = p_q A_d g. Also T_d maps the complement of P to the complement of Y, because for y in Y the function T_d^* y = A_d y lies in P. On the complement of P, the sum of squares of Lemma 29.6, with the channels with Robin kernel removed, gives the bound m/2 both for T_d : P^perp -> Y^perp and for its adjoint, so this block is bijective; for m > 2 ||A_o|| its bounded perturbation B_22 = (1 - Pi_Y)(T_d + A_o)(1 - P) is bijective with ||B_22^{-1}|| <= (m/2 - ||A_o||)^{-1}, which is the surjectivity the Schur complement F^* - B_21^* (B_22^*)^{-1} B_12^* needs, with B_12 = Pi_Y A_o (1 - P) and B_21 = (1 - Pi_Y) A_o P (referee). So with w = w_1 + w_2 and w_1 = P w,

```
||T_M w|| >= ||(1 - Pi_Y) T_M w|| >= (m/2) ||w_2|| -
 ↪  ||A_o|| ||w||,
||T_M w|| >= ||Pi_Y T_M w|| >= ||F w_1|| - ||A_o||
 ↪  ||w_2||.
```

Eliminating ||w_2||, one gets (1 + 2 ||A_o|| / m) ||T_M w|| >= ||F w_1|| - (2 ||A_o||^2 / m) ||w||. Under K(rho), F is bounded below by kappa_a(rho), and F^* is bounded below on the complement of its kernel. The same elimination holds for T_M^*, with P and Y and F and F^* exchanged. On the half-cylinder T_M is thus injective with closed range, so T_M^* is surjective; its kernel is the kernel of the Schur complement of the block T_d + (1 - Pi_Y) A_o (1 - P), whose inverse has norm at most (m/2 - ||A_o||)^-1. So this Schur complement differs from F^* by at most ||A_o||^2 / (m/2 - ||A_o||), and the graph of the Schur solution gives ||P_{ker T_M^*} - P_{ker F^*}|| = O(1/m). This projection estimate turns orthogonality to the exact half-cylinder kernel into orthogonality to ker F^* up to O(1/m), which gives the adjoint bound. Off the 4, this is invertible. On the 4 the block is a 2 x 3 matrix of rank 2, so its kernel has dimension exactly 1, within O(1/m) of (-(2/3)(rho^2 + 2), 1, 1). ■

Theorem 29.11 (the window in case (a)). Under Setting 29.0 with the pull-back connection, K(1) holds, and for each fixed 0 < delta < min(2, kappa_a(1)) there is M_delta such that for M >= M_delta the conclu-

sion of Theorem 29.7 holds at the Fubini-Study radius, with the wall modes replaced by the kernel of Proposition 29.10, and with fibre part tending to -w(1) = -2/3. Here kappa_a(1) >= sqrt 5 - 3/2 off the 4, and it equals 1.863 on the 4, the squared values there being (45 +- 3 sqrt 33)/8. At another wall radius the same conclusion holds under K(rho).

Proof. By Lemma 29.5, on V != 4 we have |D_3 (+) D_1| >= sqrt 5, while ||(sqrt 3/2) B(1)|| = 3/2. So A_eff is invertible there, with lower bound sqrt 5 - 3/2 > 0. On the 4, F^* is A(1), with kernel of dimension one in Z = +1, the vector (-2, 1, 1), and nonzero singular values 1.863 and 2.789, by Proposition 28.13. So K(1) holds. The proof of Theorem 29.7 then applies with Proposition 29.10 in place of Lemma 29.6. The trial wall states are cut-offs of the exact half-cylinder kernel: on the 4-isotypic component the adjoint kernel equation is v' = (m Z q + i Z A_4) v, with A_4 a fixed finite Hermitian matrix anticommuting with Z, so the coefficient is Hermitian, every L^2 solution on the negative half-cylinder lies for large m in its positive spectral subspace, whose eigenvalues are at least m/2 - ||A_4||, and its mass beyond distance a is O(exp(-(m/2 - ||A_4||) a)); the cut-off defect is exponentially small. The finite matrix A(rho) of Proposition 28.13 is used here as given; its derivation from the geometric data, with bases and signs, is not in the live text and is an obligation for a stand-alone paper (referee). ■

Other radii. For K(rho), only finitely many V can carry a kernel, those with |D_3(rho) (+) D_1(rho)| <= 3/(2 rho), by ellipticity. In Python, on the 16, 20, 40 and for rho in [0.2, 4], the compressed block has smallest singular value at least 1.30. A check of K(rho) on all relevant V at a given radius needs a lower bound for D_line(rho) on all V, which we have not written down away from rho = 1.

Remarks

Remark 29.12 (the index). The summands of E on CP^3 have indices 4, 0, 0, -4, of total zero, as equivariant cobordism invariance requires: E extends over X (Witten 1983, preprint pp. 43-44). The window has one copy of the 4 in each Gamma-hat degree, so its Sp(2)-index is 4 -

$4 = 0$. So index theory allows lambda != 0, and the eight states come from the estimates above. The sign in (H4) is consistent with this: with the opposite sign, the end-model modes at the core belong to T_M^* and those on O(-3) at the wall to T_M (Proposition 28.13), so the two ends never lie in the same Gamma-hat degree. At finite mass the full operator can have nonzero lambda, so these states need not lie in exact kernels. The relative sign of the two ends enters through the chirality of ker D_3. It follows from the slice weights, which contain the 4 once in Z = +1 and not in Z = -1, and the sign of the core is the chirality statement behind (H4). If the O(3) zero mode lay in Z = -1 with q > 0, it would be a wall zero mode of T_M, in conflict with Lemma 29.6 (1). As a check in Python, the mean position of the modes is 0.51 at the core for T_M and 2.96 at the wall for the adjoint problem, with l = 3.

Remark 29.13 (physical units). All statements are in shape units. Under the dilation g -> R^2 g, with D -> D/R and the potential divided by R, the operator is unitarily equivalent to D-hat_M / R. So the light mass is lambda_M / R, the gap delta / R, and their ratio is unchanged, while the compression maps stay the same; this observation is from the audit of 3 November 2026 by Astra 6, section 5.2, conditional there on the estimates proved here. A family with R -> infinity, with lambda_M / delta small, closes the physical gap and keeps the eight states separated. Whether a dynamics selects such a family is open. The limiting geometry of such a family is a separate question: nothing here identifies it with a complete Bryant-Salamon manifold.

Remark 29.14 (other metrics). In the class of Theorem 28.9, T_M - T_0 also contains first-order terms with coefficients O(r^2), and M (b(r^2) - b(0)) y . J. One route shrinks the core piece to r < M^{-1/2 + s}, where Lemma 29.4 controls the complement since M^2 r^2 dominates. We have not carried this out, and the horizontal first-order terms need a bound in terms of D_H that takes account of the unbounded curvature terms.

Second reading (GPT-6.1 Terra, 5 November 2026)

The second reading (report) found no counterexample to the corrected Theorem 29.7 with 0 < delta < sqrt 3, confirmed the core-wall channel assignment and the relative sign of Remark 29.12, and confirmed the corrected case (a) (the overlap factor sqrt 3/2, the Schur complement, the fibre sign -w(rho), the kernel (-2, 1, 1) with squared singular values (45 +- 3 sqrt 33)/8) against the finite matrix of Proposition 28.13. Its verdict on a standalone mathematics paper on case (b): not ready in the present form. What such a paper must display, in the referee's order: the spinor and curvature conventions of Proposition 29.1; a self-contained equivariant Callias kernel calculation identifying ker Q^* as 3 + 1; the domains and essential self-adjointness on N; the SU(4)-to-Sp(2) eligibility and chirality calculation of Lemma 29.5; the repaired IMS and weighted estimates with their cut-offs, and the uniformity in l for the family of lemma 30; the derivation of the finite matrix A(rho) of Proposition 28.13 from the geometry; and a literature comparison crediting the published end-model spectra and localization methods.

Prior work identified by the referee, with the passages it checked (the identifiers are in paper/LITERATURE.md, section L): the line-twisted CP^3 spectra of Lemma 29.5 are in Ben Halima (2008), Theorem 1; the S^4 instanton spectrum of Proposition 29.2 is in Dolan, hep-th/0304037; the zero-order cross-term identity of Lemma 29.4 is Maridakis, arXiv:1510.07004, Lemma 2.2, and arXiv:2110.11654 develops normal-model localization under a direction-independence assumption that (n . J)^2 violates here; Kaad and van Suijlekom, arXiv:1710.03182, Proposition 18, give a horizontal-vertical factorization with a curvature contraction under other fibration hypotheses; Bismut and Lebeau (1991), sections VIII-IX, equations (9.155)-(9.156), give model-space projection and spectral-counting arguments of the same kind; Le Peutrec (Witten Laplacian with Neumann condition), Parker, arXiv:2307.00694, Theorem 1.1, and Schulz-Baldes and Stoiber, arXiv:2203.15014, sections 3 and 5, are the methodological ancestors of the boundary localization and the IMS argument; Kottke, arXiv:0909.5661, section 5, is the families Callias framework behind Propo-

sition 28.11. The passages checked do not state the combined eight-state theorem with this S^4 core, CP^3 wall, anisotropic spin-3/2 mass and local boundary condition; this is a finding about the checked records, within a search that did not exhaust forward citations or review databases.

Numerical check

Python (numpy, scipy), scripts not kept by the rule of the project. The exact inputs were asserted: [A, q] = 0 and {A, Z} = 0 at rho = 1 with the line-preserving connection, the values of Lemma 29.5 against the formula, the norms used above, and the identities of Proposition 29.1. The direct solve is a test of the prediction: the limit values below were fixed by Proposition 29.2 and Lemma 29.5 before the run, and a failure would have refuted the argument. Setting: rho(x) = x for x <= 0.4, rho(x) = 1 for x >= 1.2, phi = M rho, l = 3, with the connection switched on [0.6, 1.2], and P1 elements at 2500 cells. Lowest singular values of T_M:

V	M = 8	M = 16	M = 32	predicted limit
4	small, 1.7321, 1.9861, 2.4495	small, 1.7321, 1.9963, 2.4495	small, 1.7323, 1.9991, 2.4497	0, sqrt 3, 2, sqrt 6
16	2.2245, 2.2361, 2.4495	2.2324, 2.2367, 2.4495	2.2351, 2.2409, 2.4497	sqrt 5 (core), sqrt 5 (wall), sqrt 6
20	2.2361, 2.8148, 2.8284	2.2367, 2.8248, 2.8284	2.2409, 2.8275, 2.8286	sqrt 5 (wall), sqrt 8 (core), sqrt 8 (wall)
40	3.4461, 3.4641, 3.4641	3.4585, 3.4641, 3.4645	3.4626, 3.4642, 3.4672	sqrt 12 three times

The wall values have a mesh error growing with M, of size 0.005 at M = 32 on the 16 and 20. The small value on the 4 is uninformative in this table: it halves when the mesh doubles, so it is a mesh floor, and Proposition

29.8 is proved without it. The Python wall values of the adjoint problem round to -1.00000 for the fibre part and the Gamma_5 part at M = 8, 16, 32, consistent with the limiting values of Theorem 29.9 in case (b); the theorem establishes the convergence, and the printed decimals do not establish exact equality at finite M. The scripts behind the table are not kept, by the rule of the project, so the table has not been reproduced by a referee.

Part 30

The massless limit as a joint limit: a finite window and an opening tower

Written by Claude Opus 5.5 on 4 November 2026, for STRATEGY T14. Theorem 30.1, Corollary 30.2, Proposition 30.3, Corollary 30.4 and Proposition 30.6 were proposed in an audit of commit b2ddf98 by Astra 6 (assessment in the review); the proofs below were redone against the source lemmas. Theorem 30.5 and Proposition 30.7 are new here and rest on lemma 29, whose referee is queued. Refereed by Terra 6.1, together with the note (report); its corrections are applied and have not been read again, since the Codex quota is spent.

Outcome

The claim that W and Z become massless as a dimension opens needs a family. Along a general path to M_W = M_Z = 0 nothing forces a tower to open: lemma 23 has theta -> 0 at fixed l, with the tower spacing fixed. The families below have the hierarchy

```
m_weak << E_eps << Delta_KK -> 0,
```

so a moving low-energy window keeps a fixed finite number of states. The hierarchy alone does not make the count at fixed energy diverge: the signed spectrum 0, +-eps^alpha, +-1, +-2, … has the hierarchy with a count bounded at fixed energy below 1. So the divergence at fixed positive energy is a separate property, proved below for each family.

1. On the exact interval this holds with an internal dimension of one (Theorem 30.1), also with localized end modes (Proposition 30.3).

2. On the cylinder of lemma 24 and on the disc bundle of lemma 29 it holds with seven internal dimensions. On the disc bundle the window comes from the operator, with one copy of the 4 at each end, and the count at fixed energy has seven-dimensional Weyl growth (Theorem 30.5).
3. On the disc bundle the core block of the window is quaternionic up to a defect that is O(M^{-1/2}) in general, and exponentially small when the connection preserves the quaternionic structure of E, for all bounded real observables (Propositions 30.6, 30.7).

Scope. These are statements about operators, compression maps and spectra along families given by hand. No dynamics selects them (Remark 30.8). The vector masses of Corollary 30.2 belong to the five-dimensional Yang-Mills toy of lemma 23. On the disc bundle only the fermion window is treated, so the relation between lambda and M_W is unproved there.

The interval

Theorem 30.1 (window and tower on the interval). Let (H_F, D_, *J_F, Gamma_F) satisfy the hypotheses of Theorem 22.1, with the algebra A_F represented by pi as in Theorem 22.4, r = dim H_F, and for 0 < alpha < 1,* $l_ > 0$ and small eps > 0 let D_F(eps) = eps D_, *l_eps =* $l_$ eps^{-alpha}, with D_eps the interval operator D_X of Theorem 22.1 for these data. Let Delta_eps be the smallest element of spec |D_eps| outside spec |eps D_*|. Then for eps small:

1. W_eps^* D_eps W_eps = eps D_*, with J_F and Gamma_F exact.
2. Delta_eps = (pi / l_) *eps^alpha - eps //D_||.*
3. ||Q_eps(a) - pi(a)|| <= (2/3) l_^2 *//D_||^2* eps^{2(1 - alpha)} ||a||.
4. For alpha < beta < 1, E_* > 0 and E_eps = E_* eps^beta, the count N_eps(E_eps) = #{|D_eps| <= E_eps} equals r.
5. For fixed E > 0, N_eps(E) ~ (2 r l_* E / pi) eps^{-alpha}.

Proof. The condition l_eps ||D_F(eps)|| = l_* ||D_// *eps^{1 - alpha} < pi/2 holds for small eps. Theorem 22.1 gives item 1, and the next level is (pi - theta_max)/l*

= pi/l - ||D_F||, which is item 2. Theorem 22.4 (2) gives item 3. For item 4, eps ||D_|| « eps^beta « eps^alpha. For item 5, the signed spectrum is (n pi +- theta_k)/l_eps, which is r arithmetic progressions of spacing pi / l_eps, counted with multiplicity also when singular values coincide or vanish, each with (2 E l_eps / pi)(1 + o(1)) points in [-E, E]. ■

Theorem 22.1 starts from supplied finite data and a matrix-valued parent algebra, and the wall data change with eps by hand. A fixed Majorana block does not scale as eps D_*, so applying the theorem to an electroweak Dirac block needs a reducing sector or a treatment of the heavy neutral sector.

Corollary 30.2 (vector poles in the toy of lemma 23). For the universal channel of lemma 23, with m_eps = eps m_* and theta_eps = l_eps m_eps -> 0,

```
M_W = arcsin(sin(2 theta_eps)/sqrt 2) / l_eps,
 ↪  M_Z = arcsin(2 sin(2 theta_eps)/sqrt 5) / l_eps,
```

so M_W^2 = 2 m_eps^2 (1 + O(theta_eps^2)) and M_Z^2 = (16/5) m_eps^2 (1 + O(theta_eps^2)). The vector poles follow the hierarchy. The constants 2 and 16/5 belong to this five-dimensional action. The small-mass pole coefficients of the fully massive finite spectral action are 3 and 24/5 (lemma 18, Proposition 14, and lemma 23), for the family with both members of every doublet at the same Dirac mass and no Majorana mass. Lemma 25 connects the two pairs at the level of the leading coefficient: its M^2 = ((3 - d_eff(s))/2) |||[D_F, T]||^2 / ||T||^2, with s = l Lambda, is the leading Taylor quotient of the heat-cutoff action for a generator constant in the transported frame, so the coefficient is 3 and 24/5 at d_eff = 0, and 2 and 16/5 at d_eff = 1. This does not identify the spectral-action pole problem with the Yang-Mills toy, and a joint-limit use needs a remainder uniform in s, since lemma 25 expands at fixed s. Along Lambda = E_eps we have s = l_* E_* eps^{beta - alpha} -> 0, the finite value, and at fixed Lambda > 0 we have s -> infinity, the interval value. Varying the cutoff of the action is an extra prescription: an observational cutoff alone does not change the action.

Proposition 30.3 (a localized family). Take M = -mu and the mixed conditions of lemma 9, with mu_* >

0, 0 < alpha < 1, mu_eps = mu_* eps^alpha and the length l_eps^loc = c_eps / mu_eps, where c_eps = (1 - alpha) log(1/eps) + c_0. This length differs from the l_eps of Theorem 30.1. Let s = kappa l be the root of s = c tanh s, and m = mu sech s. Then

```
m_eps ~ 2 mu_* exp(-c_0) eps,    Delta_eps ~ mu_*
 ↪  eps^alpha,    (1/kappa) / l_eps = 1/s -> 0.
```

With the left profile L(x) proportional to exp(-kappa x) - r exp(kappa x), r = exp(-2 kappa l), the following identities are exact:

```
r = (m / (mu + kappa))^2,    mu - kappa = m^2 / (mu
 ↪  + kappa).
```

Proof. Lemma 9.5 gives m = 2 mu exp(-c)(1 + o(1)) and kappa/mu -> 1. The first oscillatory root has pi < k_1 l < 3 pi/2 and lambda_1^2 = mu^2 + k_1^2, so Delta/mu -> 1. From mu = kappa coth s and m = kappa / sinh s we get mu + kappa = kappa exp(s) / sinh s, hence m / (mu + kappa) = exp(-s). The second identity is m^2 = mu^2 - kappa^2. ■

The localization width and the length both grow, with a shrinking ratio. The finite two-point distance 1/m and the length l are both divergent, with (1/m)/l ~ exp(c)/(2c), so a metric convergence statement would need its own comparison maps.

Corollary 30.4 (the cylinder of lemma 24). Use the operator D_r = E d_x + Z A / r on the interval times CP^3, with A as in lemma 24, finite data on V = ker A = C^4 with gradings of type (2, 2) as in lemma 24, so that r = 4, the scaling of Theorem 30.1, and r_eps = r_* eps^{-gamma}, 0 < gamma < 1. Other values of r need an enlarged tangential bundle, which we do not construct. Then for max(alpha, gamma) < beta < 1, the cutoff E_* eps^beta isolates the finite window, since the next level is at least min((Delta_K / r_) *eps^gamma, (pi / l_)* eps^alpha - eps ||D_*||), by lemma 24, (5). At fixed E > 0 the count is proportional to l_eps r_eps^6 E^7, the seven-dimensional Weyl growth, by Dirichlet and Neumann bracketing of the squared product operator as l_eps and r_eps grow. For gamma = alpha the geometry expands homothetically in all seven directions.

The window here is prescribed by the spectral walls of lemma 24.

The disc bundle

Theorem 30.5 (joint limit on the disc bundle). Fix a shape as in Setting 29.0, with connection (b), or with connection (a) at rho = 1. Let delta > 0 be as in Theorem 29.7 or 29.11, G_0 = int_e^{1 - a} f dr, 0 < alpha < beta < 1, E_* > 0, eta in (0, 1), and

```
M_eps = (2 (1 - alpha) / ((1 - eta) G_0))
 ↪ log(1/eps),    R_eps = eps^{-alpha},    D_eps =
 ↪ D-hat_{M_eps} / R_eps.
```

By the dilation of Remark 29.13, which is exact under the canonical spinor identification and the L^2 unitary factor R^{-7/2}, D_eps is the operator of the metric R_eps^2 g with potential (M_eps / R_eps) f(s / R_eps) n . J in the physical radial coordinate s = R_eps r, of size M_eps / R_eps -> 0. For eps small:

1. D_eps has exactly eight eigenvalues in (-delta eps^alpha, delta eps^alpha), and they are +- m_eps with multiplicity 4 (or zero with multiplicity 8), where m_eps <= C eps. The next level is at most C log(1/eps) eps^alpha, so it tends to zero.
2. For E_eps = E_* eps^beta the count N_eps(E_eps) = #{|D_eps| <= E_eps} equals 8.
3. The compression converges to the endpoint maps, with error C ||a||_{C^1} (log 1/eps)^{-1/2} in shape coordinates.
4. For fixed E > 0, N_eps(E) ~ C_X E^7 eps^{-7 alpha}, with C_X = 64 omega_7 Vol_g(X) / (2 pi)^7 the Weyl constant of the elliptic operator D-hat_0 on X, where omega_7 is the volume of the unit ball and 64 = 8 x 4 x 2 counts the spinors, E and the doubling.

Proof. Item 1: Theorem 29.7 or 29.11 applies since M_eps -> infinity, Proposition 28.19 gives the multiplicity, and Proposition 29.8 gives lambda <= C exp(-(1 - eta)(M_eps/2) G_0) = C eps^{1 - alpha}, hence m_eps = lambda / R_eps <= C eps. The bound on the next level: the shape-unit value is at most s_9(D-hat_0) + c M_eps, with the comparison below. Item 2: eps « eps^beta « eps^alpha. Item 3: Theorem 29.9, with M_eps^{-1/2} of the stated order. Item 4: D-hat_M = D-hat_0 + M V on the same domain, with V bounded by

c = (3/2) sup f, so by min-max N(D-hat_0, Lambda - c M) <= N(D-hat_M, Lambda) <= N(D-hat_0, Lambda + c M). Take Lambda = E R_eps, so that M_eps / Lambda -> 0, and use the Weyl law for the fixed self-adjoint elliptic operator D-hat_0. ■

So the eight finite states stay separated at the moving scale, and at fixed physical energy the spectrum has the growth of a seven-dimensional internal space, for eleven dimensions in M^4 x X. Item 1 is an upper bound on m_eps, since a lower bound on lambda is unproved. A family with l_eps growing at fixed M has the same form, with the rate of Proposition 29.8. The constants M_eta, C_eta of Proposition 29.8 are fixed once the shape and eta are fixed, so no uniformity in a varying shape is used. The limit at fixed energy resolves a neighbourhood whose shape is fixed and whose size grows, so the pointed limit is flat at a fixed interior point and a half-space at a fixed boundary point. Nothing here identifies it with a complete Bryant-Salamon manifold, or with a solution of field equations.

The quaternionic defect

For the 4-copy of the window at the core, let B(r) = F(r)^* F(r) = diag(b_a I_2, b_+, b_-) as in Proposition 28.16, let j be the quaternionic antiunitary commuting with Sp(2), and P_H(A) = (A + j A j^-1)/2. Put

```
eta_t = int |b_+(r) - b_-(r)| dr / int tr B(r) dr.
```

Proposition 30.6 (all bounded real observables). For the normalized compression Q of this mode, ||Q(f) - P_H Q(f)|| <= 2 eta_t for every real f with ||f||_infinity <= 1. If X_1, ..., X_k are such compressions, then ||X_1 ... X_k - P_H(X_1 ... X_k)|| <= 2 k eta_t.

Proof. With angular Haar measure of total mass 1, Schur's lemma gives the angular average (tr B / 4) I, so Q(f) = (4 / int tr B) int int f(r, g) lambda(g)^-1 B(r) lambda(g) dg dr. Since j commutes with lambda(g), Q(f) - P_H Q(f) is the same integral with B - P_H B. Here j exchanges b_+ and b_- and fixes the block b_a, so ||B - P_H B|| = |b_+ - b_-|/2. The triangle inequality gives the first bound. For the second, write X-tilde = j X j^-1, so that X - P_H X = (X - X-tilde)/2 and ||X_i - X-tilde_i|| <= 4 eta_t. Telescoping, X_1 ... X_k - X-

tilde_1 … X-tilde_k is a sum of k terms, each a product of contractions with one factor X_i - X-tilde_i, so its norm is at most 4 k eta_t; dividing by 2 gives 2 k eta_t. ■

So bounded expressions become quaternionic when eta_t -> 0. An unrestricted generated algebra may use coefficients of size 1/eta_t and retain the other directions, as in Counterexample 28.17.

Proposition 30.7 (the size of the defect). Under Theorem 29.7 or 29.11, eta_t = O(M^{-1/2}). Assume moreover that the quaternionic structure j_E of E is parallel on all of X. This holds for the pull-back connection and for the interpolation (1 - s(r)) nabla^pull + s(r) nabla^line with real s, and it can fail for other invariant interpolations allowed by (H3). Then for every eta in (0, 1) there are M_eta and C_eta with

```
eta_t <= C_eta exp(-(1 - eta) M int_{2e}^{l - a} f
 ↪  dr),    M >= M_eta.
```

This is the exponent of the bound on lambda in Proposition 29.8 doubled, with the integral starting at 2e. A bound in terms of lambda itself would need a lower bound on lambda, which is open.

Proof. For the first bound, fix a normalized equivariant model F_0 at the core with C F_0 = F_0 for the local real structure below, so that j B_0 j^-1 = B_0. Theorem 29.7 gives ||u - chi_c F_0|| = O(M^{-1/2}) in L^2, and |b_+ - b_-| <= (|F| + |C F|) |F - C F| gives the bound, by the Cauchy-Schwarz inequality.

For the second, j_E is parallel and j_E (n . J) j_E^-1 = -n . J, so C = C_7 (x) j_E commutes with D_X^E and with i Phi_M, hence with T_M as a differential expression, with C^2 = -1, and it exchanges the wall conditions (Proposition 28.18). For the line-preserving connection, j_E P_q j_E^-1 = P_{-q} gives j_E nabla^line j_E^-1 = nabla^line, and real affine interpolation keeps this. On Hom_G(4, -) put (C F) = C F j^-1, which is antilinear with square +1. It commutes with D_H, with the full fibre term i Q^ and with R separately, since each part of the Clifford structure transforms in the same way, so it acts on the one-dimensional space Hom_G(4, ker T_0) as a real structure. Fix F_0 there with C F_0 = F_0, so that B_0 is j-invariant. Let u = F be the mode with

T_M u = lambda v, T_M^* v = lambda u, with the phase for which <chi_c F_0, u> is real, with chi_c as in Theorem 29.7. Put zeta = u - C u and nu = v - C v. In the interior,

```
T_M zeta = lambda nu,    Pi (chi_c zeta) = 0,
```

where the second holds since <chi_c F_0, C u> is the conjugate of <chi_c F_0, u>. Take the weights of the proof of Proposition 29.8 starting at r = 2e and meeting at r_* = l - a: g = 0 on r <= 2e, g' = (1 - eta)(M/2) f on [2e, l - a], g constant beyond; and g-tilde = G - g on r <= l - a, g-tilde = 0 beyond, with G = g(l - a). So g-tilde is constant on r <= 2e. That proof gives ||exp(g) u|| <= C_eta on X, and ||exp(g-tilde) v|| <= C_eta on r <= l - a/2, and lambda <= C_eta exp(-G). The same bounds hold for C u and C v, since C is pointwise isometric.

Let xi = psi-tilde exp(g-tilde), with psi-tilde = 1 on r <= l - a and 0 on r >= l - a/2. Then ||T_M (xi zeta)|| <= lambda ||xi nu|| + || |d xi| zeta ||. Since xi is constant on the support of chi_c, Pi (chi_c xi zeta) = 0, so Lemmas 29.3 and 29.4 and the partition of Theorem 29.7 bound ||T_M (xi zeta)|| from below by the bulk term with (M/2) f and by (2 - o(1)) on the core. On the bulk |d xi| = (1 - eta)(M/2) f xi, which is absorbed with the factor eta to spare, leaving

```
||xi zeta|| <= C (lambda + ||u||_{L^2(l - a <= r <=
↪  l - a/2)}) <= C exp(-G).
```

Finally B - j B j^-1 = F^* F - (C F)^* (C F), so |b_+ - b_-| <= (|F| + |C F|) |zeta| pointwise in r. On r <= l - a, the Cauchy-Schwarz inequality with the weight exp(g) gives

```
int (|F| + |C F|) |zeta| <= 2 ||exp(g) u|| ||exp(-g)
↪  zeta|| = 2 ||exp(g) u|| exp(-G) ||xi zeta|| <= C
↪  exp(-2 G),
```

since exp(-g) = exp(g-tilde - G). On r >= l - a, both |F|^2 and |C F|^2 integrate to at most C exp(-2 G). ■

The samples in section 6.1 of the audit (Python, in the pull-back connection, rows reproduced by Claude Sonnet 5.5) have eta_t / lambda^2 between 0.5 and 0.9. They vary the length at fixed mass, with profiles outside Setting 29.0, so they are a qualitative comparison and do not test this statement as M -> infinity. Under the added hy-

pothesis the defect comes from the wall, which breaks C: for an operator with the same interior and no wall, on a complete manifold, Hom_G(4, ker T) carries a real structure, and if it is one-dimensional the core block is exactly quaternionic.

Remark 30.8 (what is missing, and the direction). The next steps are, in order:

1. The vector sector on the family of Theorem 30.5: the kinetic matrix and mass Hessian of the vector fluctuations, H_eps u = M^2 K_eps u, and then the full momentum-dependent kernel. The aim is M_W, M_Z ~ c lambda / R with finite nonzero couplings. The tunnelling eigenvalue alone does not establish this.
2. The gauge normalization. For a bulk Yang-Mills field with fixed microscopic coupling and constant profile, g_4^-2 = Vol(X) / g_11^2, so plain dilation gives g_4 proportional to R^{-7/2}, which must be separated from restoration at finite coupling. Gauge fields from isometries have Killing overlap integrals instead.
3. The action in the radius and the electroweak order parameter. A toy V = a r^-p + b eps^2 r^q with a, b, p, q > 0 has its unique minimum at r proportional to eps^{-2/(p+q)}, so alpha = 2/(p+q) lies in (0, 1) if p + q > 2. The repository has not derived such a potential, and adding it by hand would derive nothing.
4. The connector of Proposition 28.20, and the origin of the quaternionic grading of Proposition 28.19.

Numerical check

No statement here needs a computation. The audit tabulates the localized family in Python (mu_* = 1, alpha = 1/2, c_0 = 2), and Claude Sonnet 5.5 recomputed the rows; for eps = 10^-2, c = 4.3026 and m = 2.710 x 10^-3, against the asymptotic 2.707 x 10^-3.