WORKING PAPER / GEOMETRIC PROBES AND VERIFIED SECTOR REFINEMENT

Helike as a Probe of a Quantum Yang–Mills Model

From a classical geometric direction to residual-controlled quantum states.

Abstract. We map the homogeneous member of Helike to gauge-invariant observables in a published SU(2) Yang–Mills sphere matrix model. Following an all-spin GPU survey, we construct its scalar invariant ring and spin-2 covariant module directly, retaining the same Hamiltonian, domain and signed determinant. Exact rational Gaussian moments assemble both the projected Hamiltonian and the Gram matrix of its unprojected action. This gives the full-operator residual, including numerical eigensolver error, rather than only a small finite-matrix residual. Across two frozen controls and eight cutoffs, the vacuum and first scalar and tensor states jointly meet the shared residual allowance at $N = 20$; all three also satisfy $\epsilon \leq 0.01$ from $N = 22$. At $N = 24$, the two sector matrices have dimensions 169 and 244. All 16 cases and four precision reassemblies pass evidence verification. This clears a measured residual bottleneck in the selected model; full-spectrum and continuum conclusions remain separate tasks.

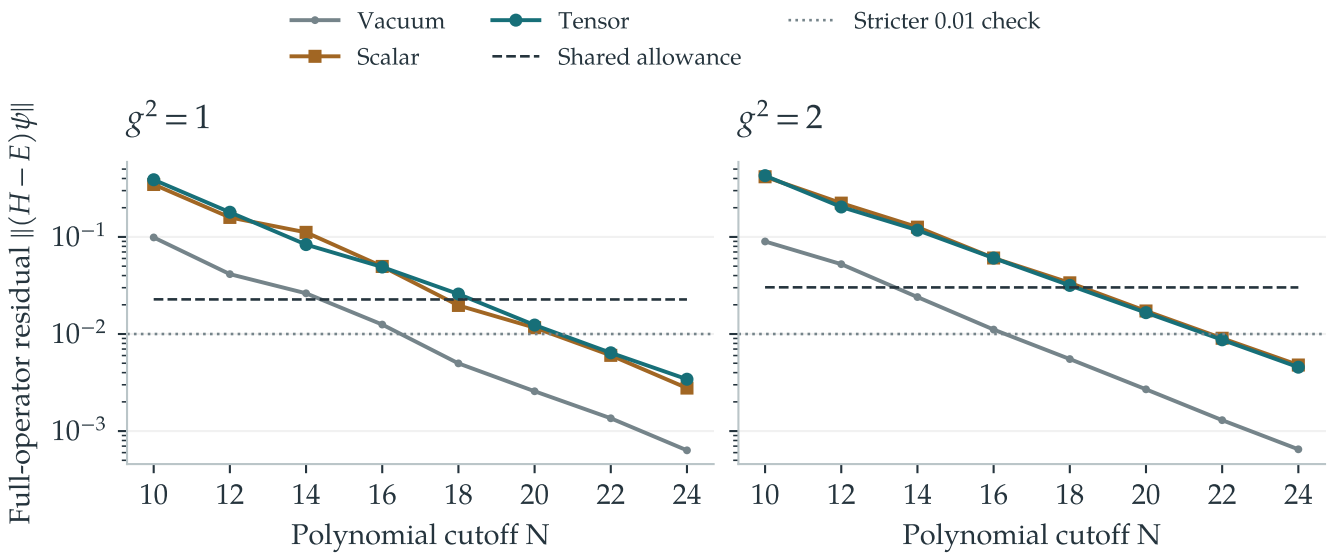

Figure 1: The new result: full-operator residuals at every registered cutoff, with no fitted curve. The dashed allowance is shared by the vacuum and both target sectors; the dotted line is the stronger fixed 0.01 check. Both controls first pass jointly at $N = 20$, and pass the stronger check at $N = 22$. The lines join measured points; they do not prove an asymptotic rate.

Why this matters. The calculation no longer asks a large undifferentiated basis to rediscover the scalar and tensor structure. It builds that structure explicitly, then measures what the full operator still pushes outside it. The improvement is achieved by a mathematically faithful change of representation and larger cutoffs, not a changed physical model or relaxed tolerance.

Contribution. A geometry-derived probe, explicit invariant/covariant operator calculus, and replayable full-residual validation are presented together. Symmetry reduction and variational diagonalization have established precedents; the particular construction and its measured results are the contribution documented here, not a claim that those ingredients were first invented here.

SURVEY / 23 NOVEMBER 2026

The initial geometric response

This section preserves the original all-spin survey. Its response weights were not recomputed in the later sector run. Its convergence verdict remains attached to that experiment, not substituted for the new residual result.

Original survey summary. A gauge-invariant probe derived from Helike's homogeneous geometry is evaluated in a published SU(2) Yang–Mills matrix model on the round three-sphere. An unrestricted colour-singlet calculation within each finite basis supplies the reference spectrum. At the largest registered cutoff, the directional channel places 99.154% of its excited response in the lowest computed level at $g^2 = 2$. At $g^2 = 1$ and 0.5, a separately registered scalar composite reaches the lowest level instead. Fifteen cases pass finite-matrix validation and evidence replay. Cutoff drift and exterior residuals remain too large for a converged model-level conclusion.

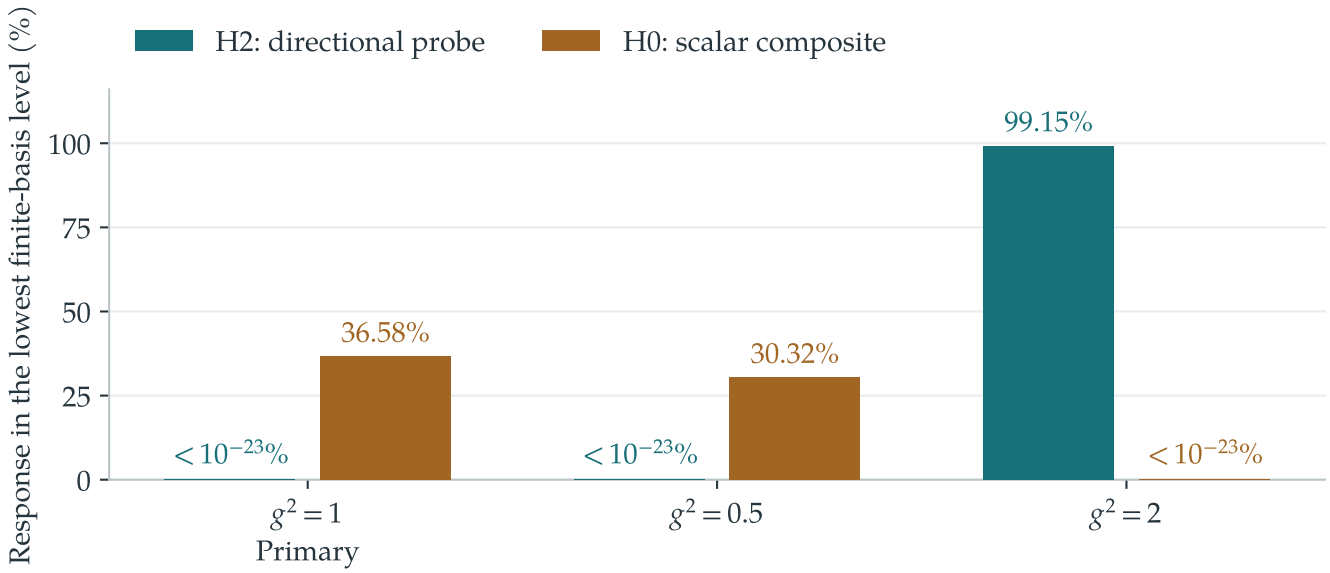

Figure 2: Response carried by the lowest complete excited level at $N = 12$ (1,134 physical basis states). Each channel is normalized by its own finite-basis excited norm after vacuum subtraction. Tiny opposite-channel weights are below 10^{-25} as fractions; they are not asserted to be exact zeros. All three cases remain **UNRESOLVED** under the convergence contract.

What this establishes. The mapped observables produce a reproducible, shape-selective response in an external quantum model. They are compared against the whole physical finite spectrum, not a spectrum restricted to look like the probe.

What this survey left open. The identity and energy of the true lowest excitation had not settled under basis growth. The matched-control efficiency comparison was therefore not run. This paper makes no claim of a continuum mass gap, computational advantage, or a result for the full four-dimensional theory.

01 / EXACT GEOMETRY

The probe preserves a direction, not just a number

The model coordinate $M \in \mathbb{R}^{3\times 3}$ has spatial rows and colour columns. Colour rotations act on the right. Thus $S = MM^T$ is gauge invariant. The classical field comes from the Floor manuscript [7]. The two real quadratures of its homogeneous member are

$$B_c = \frac{1}{\sqrt{2}}\begin{pmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & -1 \end{pmatrix}, \qquad B_s = \frac{1}{\sqrt{2}}\begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix}.$$

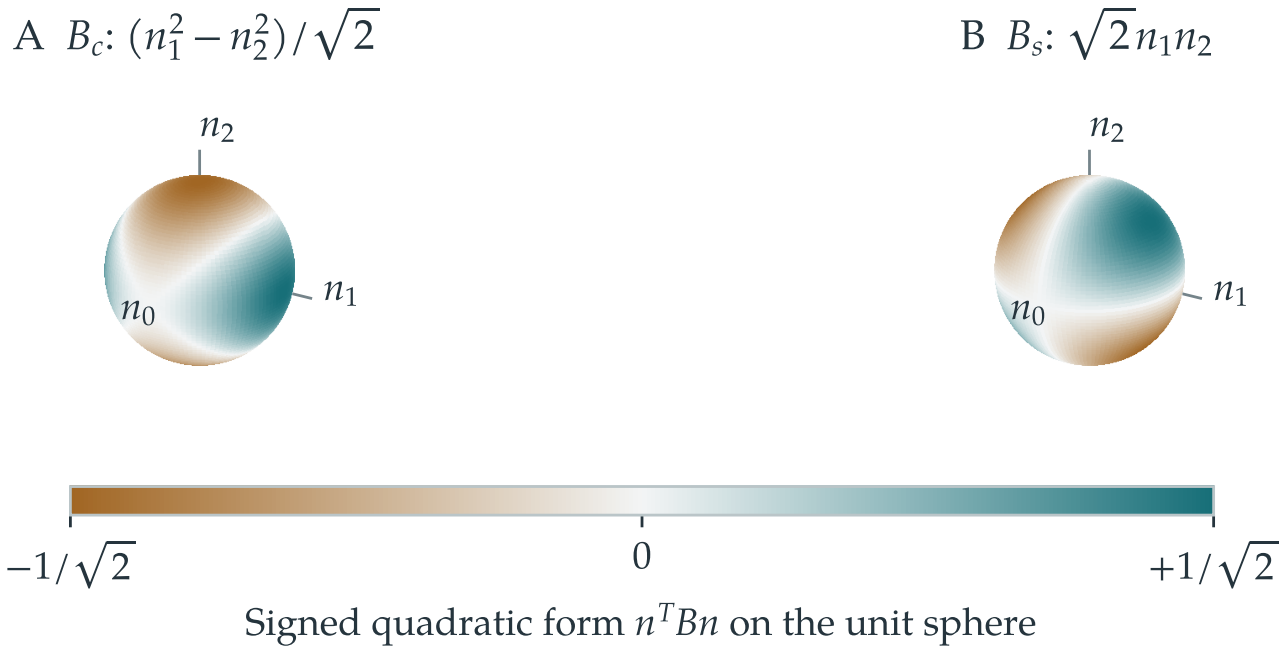

Figure 3: Exact signatures n^TBn on $\|n\| = 1$, with one camera and colour scale. Only the fixed-radius sphere's colour changes. These are declared quadratic forms, not vacuum shapes or quantum eigenfunctions. Coordinates use the registered order $(0, 1, 2)$.

Define $O_a(M) = \mathrm{tr}(B_aMM^T)$. At the mapped classical reference connection $M_* = I/2$,

$$\nabla_M O_a\big|_{M_*} = 2B_aM_* = B_a.$$

This is the exact geometric link: the observable's first response is the original direction. M_* is a classical saddle used for calibration, not the assumed quantum vacuum. The latter is independently solved across the declared domain.

Rotating these tensors spans all five symmetric trace-free spatial directions. Their observables O_A form the **directional spin-2 channel, H2**. A separate scalar channel is their squared strength:

$$O_0 = \sum_{A=1}^{5} O_A^2 = \mathrm{tr}(S^2) - \tfrac{1}{3}(\mathrm{tr}\, S)^2.$$

H0 is a quartic composite, not the original linear eigenfield. From a spin-0 vacuum, H2 couples to spin 2 and H0 to spin 0. Neither observable detects every sector. Only the homogeneous classical member is translated; the Berger floor theorem is not imported as a quantum conclusion.

02 / MEASURED RESPONSE

Where the geometry lands in the spectrum

For the computed vacuum Ω, let Q remove its component and let P_n project onto a complete excited energy level. The measured response is

$$\phi_A = QO_A\Omega, \qquad W_n = \frac{\sum_A \|P_n\phi_A\|^2}{\sum_A \|\phi_A\|^2}.$$

Degenerate levels are scored as whole subspaces, not by choosing one eigenvector. The full complex transition kernels are retained in the evidence.

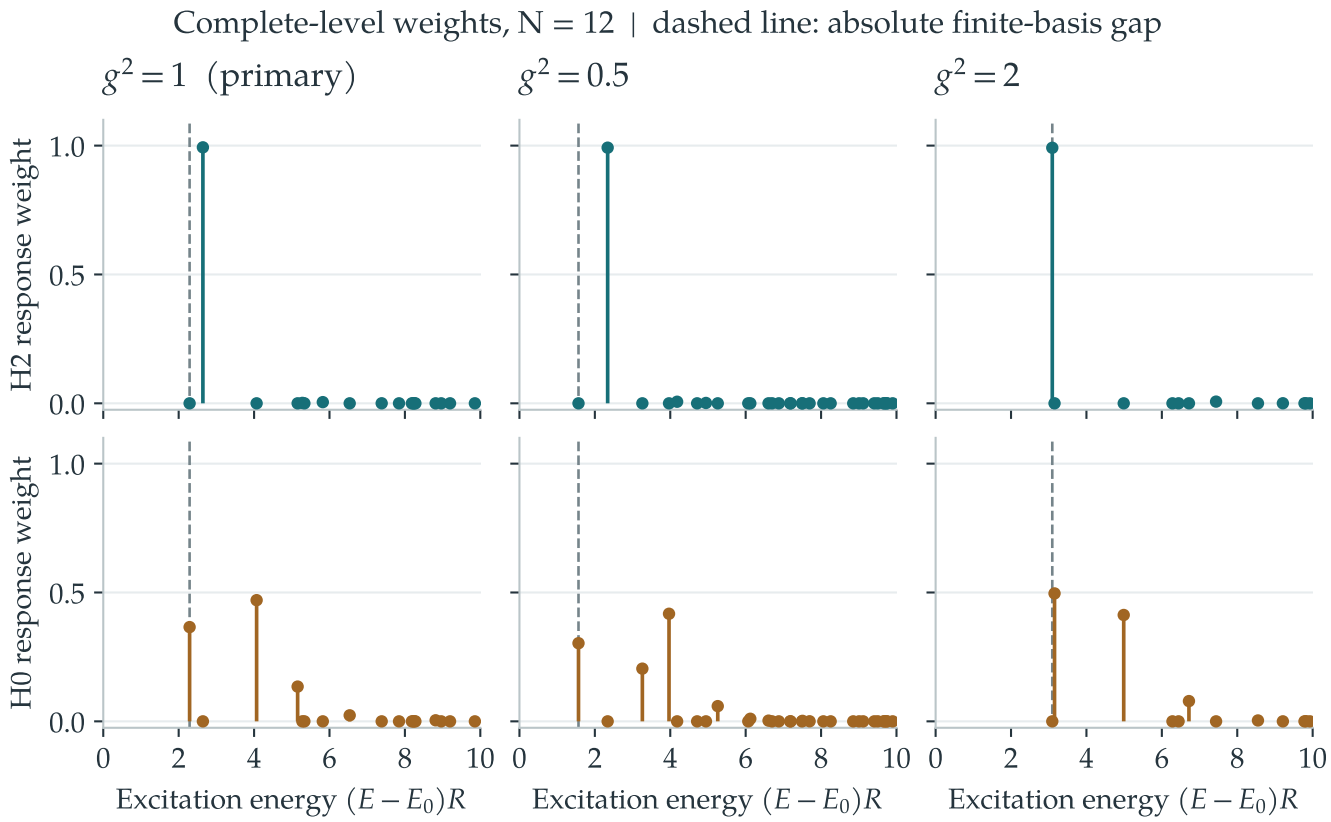

Figure 4: Finite-basis spectral fingerprints at $N = 12$. All complete levels with excitation energy at most $10/R$ are shown; weights are normalized over *all* finite levels, including those beyond the plotted window. H2 (top) is concentrated in its first accessible spin-2 level at all three couplings. H0 (bottom) spreads over several scalar levels. Dashed lines mark the independently computed absolute gap, which need not be the probe's first accessible gap.

g^2	E_0R	$\Delta_{\mathrm{abs}}R$	J_{lowest}	$\Delta_{\mathrm{H2}}R$	$\Delta_{\mathrm{H0}}R$
1	5.433212201	2.289814624	0	2.642429574	2.289814624
0.5	4.891790030	1.566333799	0	2.339593687	1.566333799
2	6.259094501	3.094157392	2	3.094157392	3.153908889

At $g^2 = 1$ and 0.5, the strong directional response is *above* the absolute lowest level. At $g^2 = 2$, it coincides with the lowest five-state level at this cutoff. "Accessible" uses the registered weight threshold 10^{-4}; smaller weights remain in the evidence. Nine displayed decimals describe the finite matrix computation, not nine-digit accuracy for the untruncated model.

03 / REFINEMENT HISTORY

The answer must survive a larger space of states

The reference basis includes every colour-invariant polynomial through degree N, multiplied by a Gaussian. It retains all spatial spin sectors and the determinant-odd states. Increasing N adds possibilities; it does not steer the solver toward Helike.

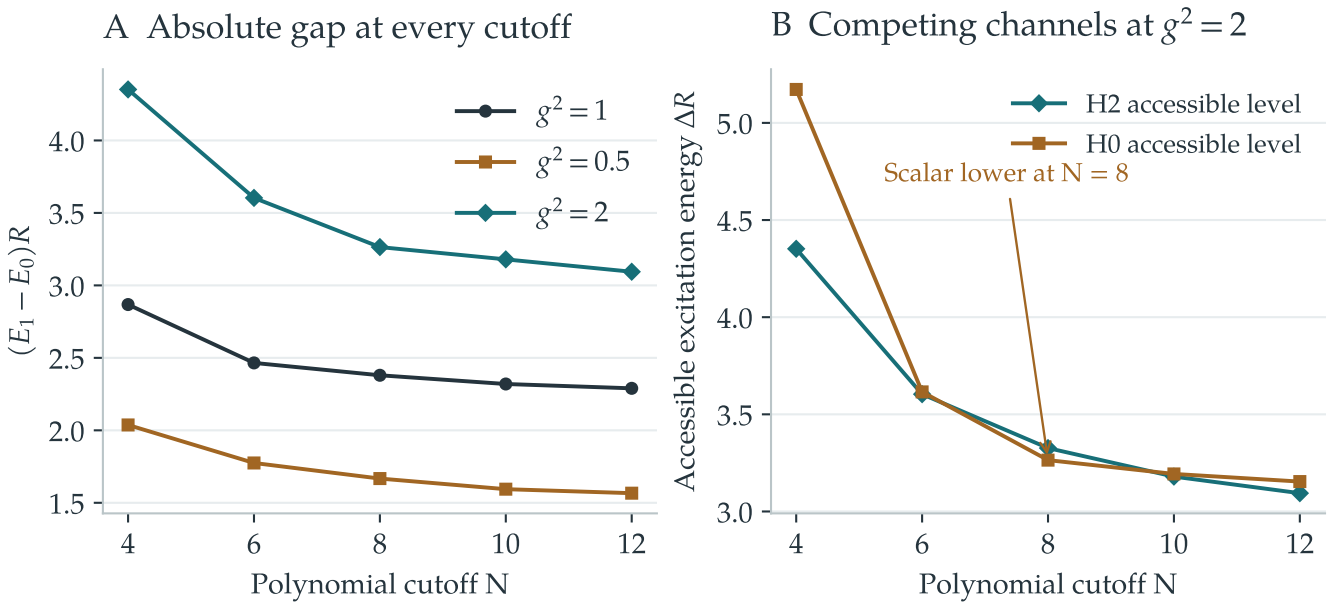

Figure 5: Registered cutoff history. (A) Absolute gaps at all three couplings. (B) The first robustly accessible scalar and directional levels at the strongest coupling. At $N = 8$ the scalar is lower, whereas the directional level is lower at $N = 10$ and 12. Lines connect computed cutoffs only; they are neither interpolated eigenstates nor an extrapolation to infinite cutoff.

N	Dimension	ΔR, $g^2 = 1$	ΔR, $g^2 = 0.5$	ΔR, $g^2 = 2$	J, $g^2 = 2$
4	29	2.867846746	2.037058771	4.351929591	2
6	91	2.465167225	1.774780654	3.603696384	2
8	238	2.379861209	1.667228247	3.264328103	0
10	546	2.319649058	1.593988390	3.179460361	2
12	1,134	2.289814624	1.566333799	3.094157392	2

The strongest-looking result has a necessary qualification. The 99.154% directional weight at $N = 12$ cannot be detached from the change in level ordering at $N = 8$. This is unresolved truncation dependence, not evidence of a physical phase transition.

The registration requires two successive increments for which E_0, E_1 and each quoted gap change by no more than $10^{-3}\max(1, |\text{quantity}|)$, and channel weights by no more than 0.01. The survey does not meet these rules. A difference of variational energy estimates is also not automatically an upper or lower bound on the true gap.

Boundary of this run. $N = 12$ was the only authorized extension beyond $4, 6, 8, 10$. The survey stopped there. No coupling, probe, Hamiltonian or acceptance threshold was changed after inspecting the results.

04 / ACCEPTANCE

Accurate finite matrices; unresolved model convergence

Finite-matrix agreement answers whether the included problem was solved accurately. Exterior checks ask whether too much of the full operator's action still falls outside that included space. These are different questions.

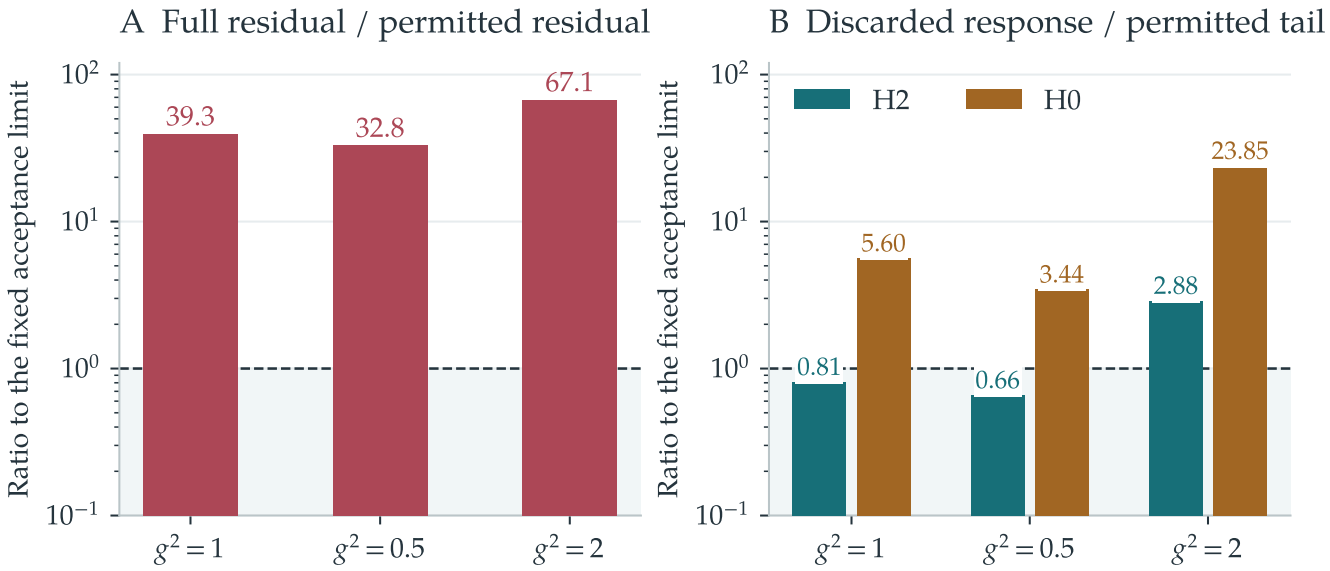

Figure 6: Final-cutoff diagnostics divided by their fixed acceptance limits. Ratios at or below the dashed line pass that individual test. (A) Worst full Hamiltonian residual among the vacuum, lowest level and first accessible probe levels, divided by $0.01\,\max(1, \Delta_{\mathrm{abs}})$. (B) Discarded squared probe response divided by 0.001. The shaded region is the passing range. Every coupling fails the residual test; the scalar tail fails in all three.

Verified in this experiment	Not established by this experiment
20 focused tests; 15 exact symbolic identities	A cutoff-converged model gap
8 mechanism-removal checks rejected defects	The lowest level's eventual ordering
Independent physical-space construction at $N = 4$	A Helike efficiency advantage
CPU/GPU spectra and low-level projectors agree	A cross-model universality result
Matrix-exponential and spectral kernels agree	A four-dimensional continuum mass gap
All 15 cases replay from stored evidence	All HP00–HP10 research gates closed

The maximum scaled CPU/GPU spectral difference was 1.161×10^{-15}; the maximum low-level projector difference was 1.646×10^{-12}, both below the registered 10^{-10} limit. This high numerical precision does not erase the exterior residuals shown above.

Original survey verdict: UNRESOLVED. The random-seed and scalar-trace utility comparisons remain `NOT_RUN` because their convergence prerequisite failed. There is no hidden failed comparator from which an advantage is inferred. A new approximation strategy or a larger run requires a separate dated amendment. The sector refinement reported later follows the 24 November protocol; it does not erase this history.

05 / METHODS AND EVIDENCE

A published Hamiltonian and a replayable calculation

The external model is the pure SU(2) sphere reduction in Ref. [1], with the radius convention checked against Ref. [2]. On the registered regular Cartesian colour-singlet domain, setting $q = M/g$ gives

$$H = \frac{1}{R}\left[-\frac{1}{2}\Delta_q + \frac{1}{2}\operatorname{tr}(qq^T) - 3g\det q + \frac{g^2}{4}\left\{(\operatorname{tr}(qq^T))^2 - \operatorname{tr}((qq^T)^2)\right\}\right].$$

The runs use $R = \hbar = 1$, positive g, no added mass or fitted energy offset, and no extra boundary extension. The two classical wells remain in the Cartesian domain. The signed cubic determinant is not replaced by its absolute value.

Basis and independent reference. The physical polynomial basis is generated by the six entries of qq^T and $\det q$, with determinant powers zero or one. The special orthogonal invariant theorem and $(\det q)^2 = \det(qq^T)$ justify this spanning construction [3]. Exact Wick ordering is a triangular change of this basis, not a dropped direction. The overlap matrix is retained in the generalized eigenproblem. An independent $N = 4$ oscillator construction imposes the three Gauss generators and assembles the Hamiltonian before projection. Direct quartic projection is used for H0; squaring already projected quadratic probes would give a different operator.

Computation. Double-precision CUDA solves and matrix exponentials ran on an NVIDIA GeForce RTX 5070 Laptop GPU, with independent CPU comparisons. CUDA requests could not fall back to CPU. The worker survey took 72.817 seconds including construction checks, excluding the separate evidence replay. Peak sampled host memory was 2.4165 GiB; peak PyTorch GPU allocation was 157.992 MiB and reservation 226 MiB. These device figures exclude driver/context allocations and other applications. They are resource measurements, not a speedup benchmark.

Provenance. Run `20261124T021903Z-c326a3c5` retains matrices, normalized eigenvectors, metric, angular Casimir, observables, complete-level kernels, cutoffs, residuals and source hashes. The source snapshot is included, because the computation predates the publishing commit. Evidence replay returned `VERIFIED`, 15 cases, `scientific_acceptance: false`. The base source revision is: `ed5d3cf5209bb8cecd6d97c4f0df3065af90363e`. Per-file hashes identify the actual experimental implementation. No server endpoint was changed or deployed.

Reuse. Figures are supplied as PDF, SVG and PNG, with editable LaTeX captions and receipt-generated tables. Separate CSV and JSON exports preserve the original survey and the new sector controls without merging their scope.

References

[1] A. P. Balachandran, S. Vaidya and A. R. de Queiroz, *A Matrix Model for QCD*, arXiv:1412.7900v1, doi:10.1142/S0217732315500807.

[2] N. Acharyya and A. P. Balachandran, *Matrix Model of QCD: Edge Localized Glue Balls and Phase Transitions*, arXiv:1702.06430v2, Eq. 2.14.

[3] H. Kraft and C. Procesi, *Classical Invariant Theory*, Theorem 10.2.

06 / SCALAR CONSTRUCTION

Put the operator in the geometry's own coordinates

Let $x = \sqrt{\alpha}\, q$ and, on this page onward, set

$$G = xx^T, \qquad T = \operatorname{tr} G, \qquad P = \operatorname{tr} G^2, \qquad d = \det x.$$

Spatial rows and colour columns match the probe convention. The sector source uses the transposed convention, $x_{\text{code}} = x^T$ and $G = x_{\text{code}}^T x_{\text{code}}$; traces, determinant, Laplacian and the spatial tensor are unchanged. This is an index convention, not an extra physical map.

The complete scalar polynomial space at a cutoff. Colour- and spatial-rotation invariants are generated by T, P, d. The elementary symmetric functions of the squared singular values are $T, (T^2 - P)/2, d^2$. Keeping d, rather than replacing it by $|d|$, retains the orientation information to which the cubic interaction couples. Thus

$$\mathcal{V}_N^{(0)} = \operatorname{span}\left\{T^a P^b d^c e^{-T/2} : a, b, c \geq 0, \quad 2a + 4b + 3c \leq N\right\}.$$

This is a complete scalar polynomial sector at degree N, not a complete finite-dimensional description of the infinite Hilbert space. Rotation averaging preserves degree, so it is the scalar part of the original cutoff space.

The inherited geometry supplies the differential operator. With $u = (T, P, d)$, the Cartesian Laplacian acts on $F(u)$ as

$$\Delta_x F = \sum_{i,j} \Gamma_{ij} \partial_i \partial_j F + 18 \partial_T F + 28 T \partial_P F, \quad \Gamma = \begin{pmatrix} 4T & 8P & 6d \\ 8P & -8T^3 + 24TP + 48d^2 & 4Td \\ 6d & 4Td & (T^2 - P)/2 \end{pmatrix}.$$

Here $\Gamma_{ij} = \nabla_x u_i \cdot \nabla_x u_j$ is the induced co-metric of the invariant coordinates. It is not a fitted covariance matrix. It can degenerate on singular orbits; the calculation retains the original Cartesian measure and regular domain rather than declaring a flat measure in (T, P, d).

For a homogeneous monomial m of degree $D = 2a + 4b + 3c$, the polynomial action is

$$e^{T/2} H(m e^{-T/2}) = \alpha \left[-\tfrac{1}{2}\Delta_x m + (D + \tfrac{9}{2})m - \tfrac{T}{2} m\right] + \frac{T}{2\alpha} m - \frac{3g}{\alpha^{3/2}} dm + \frac{g^2}{4\alpha^2}(T^2 - P)m.$$

The action stays in the untruncated invariant ring but can exceed degree N. Those higher-degree terms are retained when measuring the residual.

Exact integration, not sampled integration. Write $\langle F \rangle = \pi^{-9/2} \int F(x) e^{-T}\, d^9 x$. Integration by parts gives

$$\langle \nabla F \cdot \nabla u \rangle + \langle F \Delta u \rangle = \langle F \nabla T \cdot \nabla u \rangle.$$

Using $u = T, P, d$ yields terminating rational recurrences for every polynomial moment. For example $\langle T^a \rangle = (2a + 7)\langle T^{a-1} \rangle / 2$, $\langle T \rangle = 9/2$, and $\langle d^2 \rangle = 3/4$. No Monte Carlo error or integration grid enters these moments.

07 / TENSOR CONSTRUCTION

Retain the turning shape, not five duplicate solves

The Helike-derived quadratic observable pairs a trace-free spatial tensor with G. Its natural spin-2 carrier is therefore the trace-free part of G. That first carrier alone is not the full sector. Define

$$A_1 = \operatorname{dev} G, \qquad A_2 = \operatorname{dev}(G^2), \qquad \operatorname{dev} X = X - \tfrac{1}{3}\operatorname{tr}(X) I.$$

For a generic diagonal G, its rotation stabilizer forces a symmetric equivariant tensor to be diagonal in the same frame. Polynomial covariance and Cayley–Hamilton reduce its generators to I, G, G^2; removing the trace leaves A_1, A_2. Scalar factors belong to $\mathbb{R}[T, P, d]$. The cutoff module is

$$\mathcal{V}_N^{(2)} = \operatorname{span}\{m A_k e^{-T/2} : k = 1, 2,\ \deg m + 2k \le N\}.$$

It retains every tensor polynomial in this sector, not just the probe's two initial quadratures. The following exact identities close its kinetic action:

	A_1	A_2
$\nabla T \cdot \nabla$	$4A_1$	$8A_2$
$\nabla P \cdot \nabla$	$8A_2$	$16(TA_2 - e_2 A_1)$
$\nabla d \cdot \nabla$	0	$4dA_1$
Δ_x	0	$22A_1$

$e_2 = (T^2 - P)/2.$

The product rule $\Delta(mA_k) = (\Delta m)A_k + 2\sum_u (\partial_u m)(\nabla u \cdot \nabla A_k) + m\Delta A_k$ then reduces every matrix entry to the same scalar moments. For instance,

$$A_1 : A_1 = P - T^2/3, \qquad A_1 : A_2 = -T^3/2 + 7TP/6 + 3d^2.$$

Frobenius contraction sums five components. Rotational invariance gives each member one fifth of the Gram entries; the common factor cancels in normalized energies and residuals. One matrix thus represents complete five-state multiplets.

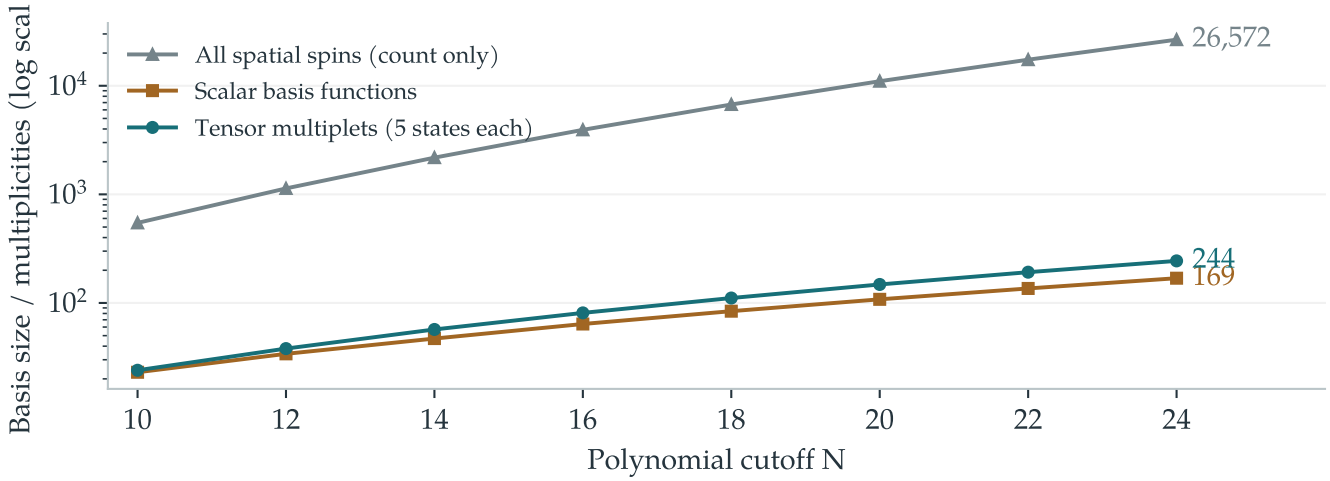

Figure 7: Exact basis counts, not runtime measurements. At $N = 24$, the all-spin colour-singlet polynomial space has 26,572 states; the scalar solve uses 169 functions and the tensor solve 244 multiplets. Other spins are not included in these two solves. The gray curve counts an unexecuted all-spin extension; it is not a measured performance baseline.

The counts follow $h_0(z) = [(1 - z^2)(1 - z^3)(1 - z^4)]^{-1}$ and $h_2(z) = (z^2 + z^4)h_0(z)$, summed through N. The all-spin count is $\binom{\lfloor N/2 \rfloor + 6}{6} + \binom{\lfloor (N-3)/2 \rfloor + 6}{6}$. At matching cutoffs and widths, sector energies and residuals agree with all six saved nine-variable control calculations. The smaller representation does not artificially lower the error at the same cutoff; it makes further refinement practical.

08 / FULL-OPERATOR ERROR

Measure the part the finite matrix cannot see

A small residual inside a diagonalized finite matrix does not test the discarded states. We instead compute the norm of the original Hamiltonian acting on the approximate wavefunction, including its component outside the trial space.

For basis functions ϕ_i, assemble three Gram matrices:

$$S_{ij} = \langle \phi_i, \phi_j \rangle, \qquad H_{ij} = \langle \phi_i, H\phi_j \rangle, \qquad K_{ij} = \langle H\phi_i, H\phi_j \rangle.$$

Crucially, $H\phi_j$ in K is the unprojected polynomial action. Substituting the square of the finite Hamiltonian would remove precisely the error being tested. Let $S = LL^T$ and orthonormalize to

$$\widetilde{H} = L^{-1}HL^{-T}, \quad \widetilde{K} = L^{-1}KL^{-T}, \quad X = \widetilde{K} - \widetilde{H}^2.$$

If U maps coefficient vectors isometrically into the Hilbert space and $\Pi = UU^*$, then

$$X = U^*H(I - \Pi)HU \geq 0.$$

This is the Gram matrix of the escaped operator action, not an estimated tail based on differences between successive eigenvalues.

Orthogonal residual decomposition. For the returned numerical eigenpair (E, v), let $\psi = Uv/\|v\|$. Since the retained and discarded components are orthogonal,

$$(H - E)Uv = (I - \Pi)HUv + U(\widetilde{H}v - Ev),$$
$$\epsilon^2 = \|(H - E)\psi\|^2 = \frac{v^T X v + \|\widetilde{H}v - Ev\|^2}{v^T v}.$$

The second term is retained even when the finite eigensolve appears accurate. The identity is exact for the mathematical matrices. Its implementation is mixed-precision floating-point arithmetic, not an interval certificate.

Precision policy. Moments are rational. Assembly, Cholesky whitening, subtraction in X, and the final residual contractions use 60 decimal digits. NumPy diagonalization uses float64. Each top-cutoff sector is freshly reassembled at 100 digits; the six reported energies and residuals agree at stored output precision. This does not imply 100-digit eigenvalue accuracy. Gram-conditioning checks and a -10^{-12} exterior-eigenvalue refusal guard the computation. Only tiny negative contractions within an explicitly reported precision-scaled roundoff allowance are floored to zero; materially negative values are refused.

What the norm means. For a normalized vector in a self-adjoint operator's domain, $\mathrm{dist}(E, \sigma(H)) \leq \epsilon$. This is a statement about proximity to some spectral value, not its index. It does not by itself identify the vacuum, prove the lowest excited state, certify a Temple denominator, or bound a gap by subtracting two Ritz estimates. These familiar residual principles are used here to test the geometry-derived calculation, not claimed as new theorems.

09 / MEASURED ADVANCE

The local residual bottleneck is cleared

The frozen controls are $(g^2, \alpha) = (1, 1.7)$ and $(2, 2.3)$, with $N = 10, 12, 14, 16, 18, 20, 22, 24$. No held-out coupling was run. At each cutoff the same allowance is applied to the vacuum and both first sector excitations:

$$a_N = 0.01 \max\{1, \min(E_s - E_0, E_t - E_0)\}.$$

This is explicitly a two-sector allowance, not a verified global gap. Both controls first satisfy all three tests at $N = 20$. From $N = 22$, all three also meet $\epsilon \leq 0.01$, independently of that sector-gap estimate.

N	ϵ_s, $g^2 = 1$	ϵ_t, $g^2 = 1$	ϵ_s, $g^2 = 2$	ϵ_t, $g^2 = 2$
10	0.347055	0.388349	0.416759	0.427875
12	0.158880	0.179696	0.223010	0.203615
14	0.111476	0.083442	0.126115	0.117383
16	0.049674	0.048997	0.060669	0.060551
18	0.019698	0.025808	0.033673	0.031669
20	0.011626	0.012368	0.017271	0.016597
22	0.006026	0.006402	0.009042	0.008717
24	0.002774	0.003422	0.004793	0.004552

The table gives first scalar and tensor residuals. At $N = 24$, the vacuum residuals are 0.000633454 and 0.000649068. Some higher reported states have larger residuals: the next tensor's is 0.015631922 at $g^2 = 1$ and 0.036781034 at $g^2 = 2$. The target-state result is not promoted to all levels.

Why this is a stronger result than stable-looking energies. At $g^2 = 1$, the scalar residual falls from 0.347055 at $N = 10$ to 0.002774 at $N = 24$; at $g^2 = 2$ it falls from 0.416759 to 0.004793. These norms test the original Hamiltonian acting on the represented states. They do not infer accuracy merely because two successive energy estimates look similar. The tensor and vacuum undergo the same check, rather than being assigned the scalar's error estimate.

What was held fixed. The coupling, Gaussian width within each control, domain, Hamiltonian and acceptance rule remain fixed while the nested polynomial space grows. The sector representation removes the need to assemble unrelated spin blocks; it does not remove the higher-degree action of the operator. That action is precisely what the exterior Gram measures. Exact sector reduction preserves the target problem, and refinement reduces its observed residual.

Evidence checks address different failure modes.

- Low-cutoff agreement with the nine-variable implementation checks the sector construction against an independent implementation.
- Exact-moment and Cartesian derivative controls check the operator tables.
- Free-oscillator and nested-cutoff tests check normalization and variational behavior.
- Four top-cutoff reassemblies check sensitivity to assembly precision.
- Saved-matrix replay checks every reported energy, residual and verdict against its retained matrices; it is not fresh physical assembly.

Together these support the stated numerical result. They do not replace certified spectral separation or the unrun higher-spin comparison.

09 / BRANCHES AND EVIDENCE

Follow the scalar and tensor branches separately

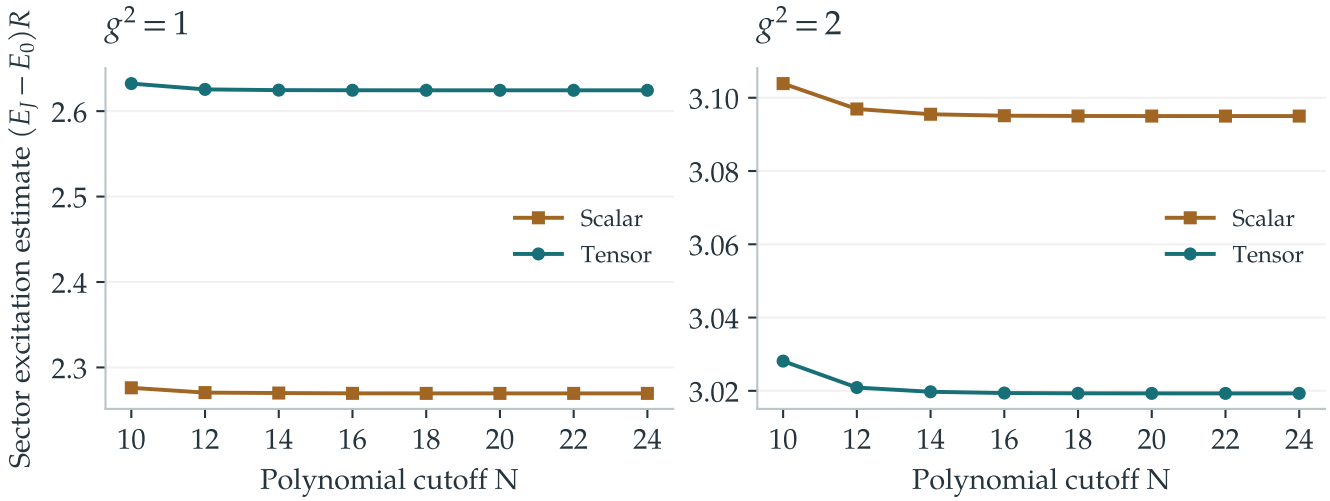

Figure 8: The two sector excitation estimates, followed separately. The scalar Ritz excitation lies lower at $g^2 = 1$ and the tensor at $g^2 = 2$. These are two controls, not a measured crossing curve in coupling. No error bars are drawn because no certified level-specific enclosures were established.

g^2	E_0	E_s	E_t	$E_s - E_0$	$E_t - E_0$
1	5.430381421	7.699987972	8.054765705	2.269606552	2.624384285
2	6.248017521	9.343017843	9.267315993	3.095000322	3.019298471

At $N = 24$, $E_t - E_s$ is +0.354777733 and −0.075701850 respectively. The conditional Temple ordering survives the registered halved-separation test, but the separation is still assumed from Ritz data. The unconditional ordering therefore remains `ORDER_UNRESOLVED`. Displayed digits reproduce the stored calculation, not a claim of nine-digit physical accuracy.

Keep the probe map and normalization explicit. With $M = gq = gx/\sqrt{\alpha}$ and $G = xx^T$, the original observables are

$$O_A = \frac{g^2}{\alpha} B_A : A_1, \qquad O_0 = \frac{g^4}{\alpha^2} A_1 : A_1 = \frac{g^4}{\alpha^2}(P - T^2/3).$$

Thus the geometric direction and its scalar strength belong to exactly the modules being refined. This algebraic bridge does not supply a new transition weight: recomputing $QO_A\Omega$ and its discarded tail at higher cutoffs is still required. The present improvement is in the candidate quantum states and their full-operator residuals.

Execution and evidence. All 16 cases, four 100-digit reassemblies and 36 saved matrix archives pass replay. The run's 69 focused tests and 19 rejected mechanism removals are retained. Runtime was 828.343 seconds (13.8 minutes), with 288.9 MiB peak sampled worker memory, CPU only and one BLAS thread. This is a recorded resource result, not a matched CPU/GPU speedup comparison.

10 / POTENTIAL GEOMETRY

The same operator exposes its wells and an exact floor

In this section $T_q = \mathrm{tr}(qq^T)$ and $d_q = \det q$ refer to the physical coordinate q, not the width-scaled coordinate x. Define

$$W(q) = \tfrac{1}{2}T_q - g d_q, \qquad \nabla W = q - g \operatorname{cof} q.$$

The cofactor identities $q : \operatorname{cof} q = 3d_q$ and $\|\operatorname{cof} q\|^2 = (T_q^2 - \mathrm{tr}((qq^T)^2))/2$ give the potential exactly:

$$\tfrac{1}{2}\|\nabla W\|^2 = \tfrac{1}{2}T_q - 3g d_q + \tfrac{g^2}{4}\{T_q^2 - \mathrm{tr}((qq^T)^2)\} = V(q).$$

Since $\Delta d_q = 0$ and $\Delta T_q = 18$, $\Delta W = 9$. On the registered regular Cartesian domain, with $\mathcal{A} = \nabla + \nabla W$,

$$H = \tfrac{1}{2}\mathcal{A}^*\mathcal{A} + \tfrac{9}{2} \;\geq\; \tfrac{9}{2}.$$

This is a floor relative to the declared energy origin. It is *not* the vacuum-to-excitation gap. Saturation would require $\psi \propto e^{-W}$, which is not normalizable for $g > 0$: the cubic term produces open directions of growth. Nonattainment alone does not prove a strictly separated spectral infimum. Every computed ground energy satisfies the non-strict floor.

The wells are geometric sets. Solving $q = g \operatorname{cof} q$ gives $q = 0$ and $q = O/g$, $O \in SO(3)$. A nonzero singular matrix cannot solve this equation: its cofactor has a different rank. For an invertible solution, multiplication by q^T gives $q^T q = g \det(q) I$, forcing the stated orientation-preserving orbit. Thus the second well is not just another point in nine coordinates.

An explicit connecting path, not just an endpoint difference. For fixed $O \in SO(3)$, put

$$q(t) = \frac{a(t)}{g} O, \qquad a(t) = \frac{1}{1 + e^{-(t - t_0)}}.$$

Then $\dot{a} = a(1 - a)$ and $\dot{q} = \nabla W(q)$, so the Euclidean action square-completes as

$$\begin{aligned} S_E &= \int \left(\tfrac{1}{2}\|\dot{q}\|^2 + \tfrac{1}{2}\|\nabla W\|^2\right) dt \\ &= \tfrac{1}{2}\int \|\dot{q} - \nabla W\|^2 dt + W(q_+) - W(q_-) \\ &= \frac{1}{2g^2} \quad \text{on this path.} \end{aligned}$$

This provides a checkable connection between the signed determinant, well geometry and the declared model. Superpotential/instanton structure in Yang–Mills matrix mechanics has precedents [6]; this calculation is not a claim of priority for that structure. It does not determine the quantum excitation gap from the action alone.

11 / FLUX BACKGROUND

Hold Helike's field, then ask which excitation is lowest

Update, 25 November 2026. Pages 11–14 are new. They follow the same Helike channel around a held flux background, with rigorous bounds on the physical operator, a proved discrete-spectrum theorem and a Kato–Temple argument.

The question. The edge paper's classical theorem says Helike's eigenvalue is helicity squared minus the flux along her circle. Its quantum counterpart asks: around a background that carries flux along the long circle, does the lowest vibration switch to Helike's channel where that mechanism says it should? We hold the background with external sources and study the fluctuations around it.

The model, unchanged in form. With $q \in \mathbb{R}^{3\times3}$ (rows are frame indices, columns colour), $\omega = (\sqrt{t}, 1/\sqrt{t}, 1/\sqrt{t})$ and sources $\mathrm{src} = (a, b, b)$,

$$H_{\mathrm{src}} = -\tfrac{1}{2}\Delta_q + \tfrac{1}{2}\|\omega \circ q - g\,\mathrm{cof}\,q\|^2 - \tfrac{1}{2}\textstyle\sum_k \mathrm{src}_k\|q_k\|^2.$$

The sources $a = (1-s)(t-2s)$ and $b = (1-s)((1-s)/t - s)$ hold the Maurer–Cartan-line background $q_{\mathrm{bg}} = \mathrm{diag}(2s/\sqrt{t}, 2s, 2s)/(2g)$ stationary. Its harmonic frequencies split by locked charge: Helike's doublet at charge ±2, a competitor at charge 0, a charge-±1 pair and a second charge-0 mode. At the controls ($s = 1.15$, $t = 3$) they are 1.0901 (×2), 1.1067, 2.4297 (×2) and 2.6777.

The window is the stability condition. Expanding the square shows that the quadratic part of H_{src} has coefficients

$$\omega_0^2 - a = s(2 + t - 2s), \qquad \omega_1^2 - b = \frac{s\,[(2+t) - s(1+t)]}{t}.$$

The second is positive exactly when $s < (t+2)/(t+1)$, the registered window. Figure 9 shows it against every registered point.

The registered grid is held. The full flux-edge grid runs at $g^2 = 0.03$–0.07. Everything below is at the two weak-coupling controls, $g^2 = 0.004$ and 0.002, where every bound can be checked. The grid is released only through a recorded method amendment, which is not yet written.



Figure 9: Why the window is the stability condition. The quadratic coefficient $\omega_1^2 - b$ (teal) is positive exactly for $s < (t+2)/(t+1)$ (shaded); $\omega_0^2 - a$ is far larger throughout. Dots are the registered grid points, the square is the controls. Closed forms; nothing is fitted.

12 / A DOMAIN THAT RESPECTS THE GEOMETRY

Two sheets, one fold, and a basis that stays inside

The domain was the first problem. A local Fock continuation of the reduced operator turned out not to be bounded below: it reached matrices that are not Gram matrices. The repair keeps every trial point positive definite. The Gram matrix is written $G = LL^T$ with L lower triangular, logarithmic diagonal and free off-diagonal entries. The chart's exact volume and co-metric are carried in full, and trial functions are $\psi = \chi/\sqrt{v}$.

The interface is a fold, not a wall. The positive and negative determinant sheets meet at rank two. There the original Cartesian operator imposes matched values and a vanishing conormal sum: an orientation double cover. It imposes no Dirichlet, Neumann or Robin wall. The signed cubic distinguishes the two sheets, $V(w) - V(-w) = -2g(\Sigma\omega)r_0r_1w$.

Cutoff quasimodes avoid the fold. Each Ritz state is multiplied by a smooth cutoff that equals one where the state lives and vanishes well before rank loss. The result is a smooth, compactly supported function of q. It lies in the domain of the physical operator whatever happens at the interface, so every bound below is a bound for the physical operator.

Two independent solvers agree. The Python reference forms dense moment matrices. The Rust solver applies H matrix-free with exact Kronecker moments: each Laurent term factorises into one-dimensional moments, and the operator acts on the occupation simplex without ever forming a matrix. From 924 to 3,003 trial functions, the identified gaps change by at most 6.7×10^{-4}, under the registered 0.003. The two implementations agree on every gap to about 10^{-13} (Figure 10, lines).

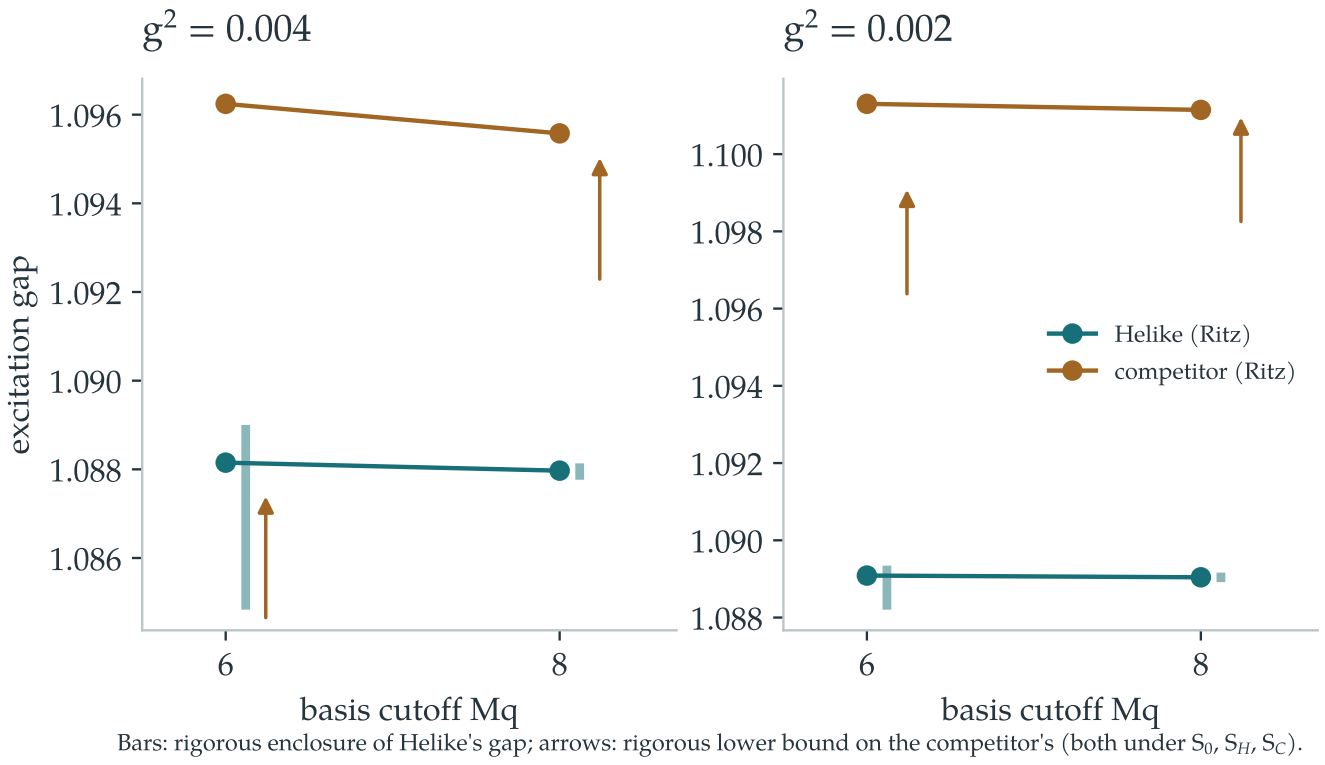

Figure 10: Basis refinement and rigorous bounds together. Lines: Ritz gaps of Helike (teal) and the competitor (gold) at $M_q = 6$ and 8. Bars: rigorous enclosure of Helike's exact gap. Arrows: rigorous lower bound on the competitor's exact gap. Bounds are conditional on the stated separation hypotheses; the Ritz values lie inside them.

13 / RIGOROUS BOUNDS

Every number an enclosure, every error accounted

What is enclosed. For each saved state, and with outward-rounded arithmetic throughout:

- the full-operator residual $r = \|(H - E)\psi\|$, from an exact finite exponential-polynomial moment sum, with no quadrature;
- the charge defect $c = \|(L - m)\psi\|$, through the exact Laurent velocity of the rotation, and from it $q \le \|P_m \eta \psi\|$;
- the cutoff tail t and leakage $\ell = \|[H, \eta]\psi\|$, by Gaussian majorants integrated to infinity in 128-bit ball arithmetic;
- the linear moment $s = \langle \psi, (H - E)\psi \rangle$, by the same moment sum, checked against an independent polarization of two residual enclosures.

The projected state then satisfies $\|(H - E)\phi\| \le R = (r + \ell)/q$.

The leakage is negligible, and now we know it. An absolute 10^{-10} line on the leakage was never derived. It was replaced, before any further run, by the rule the certificate actually needs: $\ell \le 10^{-3}\, r_{\mathrm{lower}}$. All sixteen states pass. The worst ratio is 3.1×10^{-5}, as predicted in advance (Figure 11).

One defect found and fixed on the way. The first gate evaluation refused at its endpoint-pairing check. A JSON parser that is not correctly rounded had read three Arb upper bounds one ulp low. Parsing was made exact, and endpoints are now computed upward from the exact dyadics. No bound changed.

g^2	Mq	state	r (lower)	$c \le$	$\ell \le$	$\ell/r \le$	$R \le$
0.004	6	ground	0.015147	1.59e-03	2.9e-09	1.9e-07	0.015147
0.004	6	Helike	0.048725	8.23e-03	4.9e-09	1.0e-07	0.048727
0.004	6	competitor	0.080123	9.70e-03	2.9e-09	3.6e-08	0.080127
0.004	8	ground	0.006468	5.75e-04	2.4e-08	3.7e-06	0.006469
0.004	8	Helike	0.008821	1.18e-03	2.7e-07	3.1e-05	0.008822
0.004	8	competitor	0.017696	1.73e-03	1.7e-07	9.7e-06	0.017697
0.002	6	ground	0.005394	5.60e-04	1.4e-27	2.5e-25	0.005395
0.002	6	Helike	0.023894	3.93e-03	1.7e-27	7.0e-26	0.023894
0.002	6	competitor	0.039019	4.57e-03	9.8e-28	2.5e-26	0.039019
0.002	8	ground	0.002191	1.90e-04	1.1e-26	5.1e-24	0.002192
0.002	8	Helike	0.002809	3.64e-04	1.8e-25	6.4e-23	0.002810
0.002	8	competitor	0.005525	5.15e-04	1.1e-25	2.1e-23	0.005527

Evidence identity. Pages 11–14 read four hash-pinned receipts: the M_q $6 \to 8$ convergence run, the interval sector bounds, the relative leakage gate and the Kato–Temple ordering. The builder refuses any of them that promotes a certificate, releases the grid or reports a reversed order. The Rust interval code and its receipts are on `feat/helike-lab`.

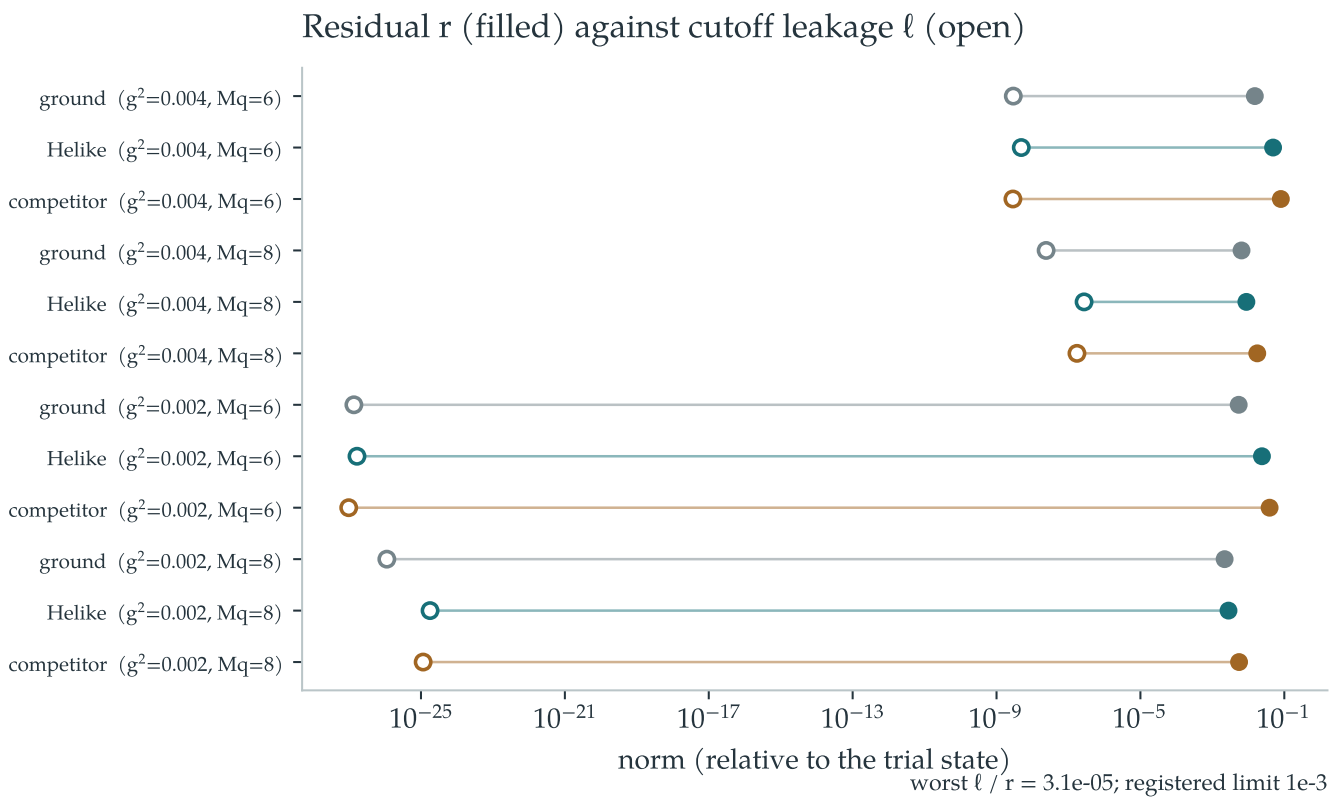

Figure 11: Where the error lives. Filled: rigorous lower bound on the full-operator residual. Open: rigorous upper bound on the cutoff leakage. The leakage is at least four orders of magnitude below the residual at every state, and far smaller at the weaker coupling.

14 / THE ORDERING

Helike sits below the competitor at weak coupling

Theorem A (discrete spectrum). On the window $t > 1$, $1 < s < (t+2)/(t+1)$, and for every $g > 0$,

$$V(q) \;\geq\; \tfrac{\kappa}{2}\|q\|^2 - \frac{S^4}{2g^2}, \qquad \kappa = \min_k(\omega_k^2 - \mathrm{src}_k) > 0, \quad S = \sum \omega.$$

The proof pairs the cubic term against one quartic term by AM–GM in singular values. Hence H has compact resolvent and discrete spectrum in every colour-singlet charge sector.

Three lemmas. (1) The Rayleigh quotient of the trimmed, projected state lies within $(s \pm e)/[q^2, 1]$ of the saved centre, with $e = 2tr + \ell + c(r + \ell)$. (2) Kato–Temple bounds the largest eigenvalue below any threshold ρ from below. (3) The ground and competitor Weinstein windows are disjoint, so two charge-0 eigenvalues lie below $\rho_C = E_0 + 2\omega_H - \omega_{\min}/4$.

One stated hypothesis, S_C. In charge 0, at most two eigenvalues lie below ρ_C. The trial spectrum puts the third level more than 0.25 above it. Helike's side needs no hypothesis: her level is bounded above by min–max.

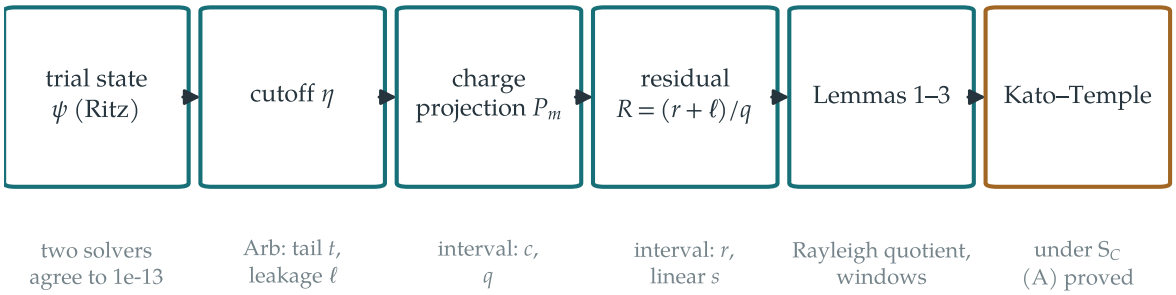

Figure 12: The certificate chain, from a trial state to a verdict about the physical operator. Each box is computed in outward-rounded arithmetic or proved; the last step is the only one that uses a hypothesis.

g^2	Mq	Ritz gaps H / C	$U_H - E_0$	$L_C - E_0$	margin	verdict
0.004	6	1.088151 / 1.096244	1.088602	1.087547	-0.001055	UNRESOLVED
0.004	8	1.087970 / 1.095578	1.087982	1.095162	+0.007179	ORDERED
0.002	6	1.089087 / 1.101302	1.089193	1.099236	+0.010042	ORDERED
0.002	8	1.089042 / 1.101146	1.089044	1.101105	+0.012061	ORDERED

Result. At three of the four controls the exact charge-±2 ground level lies strictly below the exact second charge-0 level, conditional on S_C. At $M_q = 6$, $g^2 = 0.004$, the rule is UNRESOLVED: a residual of 0.08 cannot separate levels 0.008 apart. Every verdict and margin matches a prediction recorded before computation.

Not established. S_C itself; anything at the registered grid couplings; the location of the switch; any continuum statement.

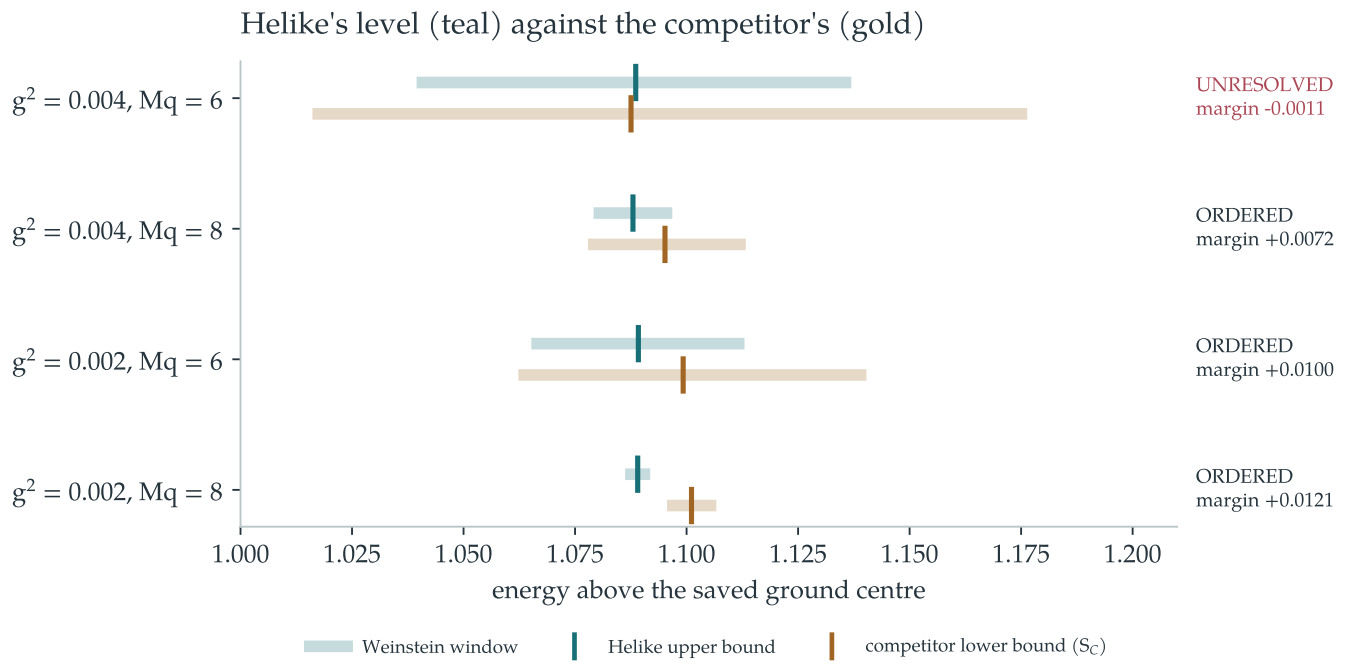

Figure 13: The ordering, control by control. Shaded: Weinstein windows, each containing a true eigenvalue. Ticks: Helike's upper bound and the competitor's Kato–Temple lower bound under S_C. ORDERED means the gold tick lies right of the teal one.

15 / CONTRIBUTION AND CONTEXT

What is distinctive, and what is inherited

The research question is geometric. Instead of starting from a low eigenvalue and naming it, we start from a specified classical field direction. Its quadratic observable has the exact first variation $\nabla O_a(M_*) = B_a$. Its rotational orbit defines a tensor channel; its squared strength defines a separate scalar channel. The reference Hamiltonian is not fitted to either. The all-spin survey tests where those observables land, and the sector construction makes their relevant spectral branches accessible to tighter full-operator error checks.

Established ingredients deserve explicit attribution. Pavel [4] uses symmetric gauge variables, elementary symmetric polynomials, harmonic-oscillator trial states and spin-resolved variational calculations in spatially constant SU(2) Yang–Mills mechanics. That is a close methodological precedent, though not the same sphere Hamiltonian with its quadratic and signed cubic terms. Acharyya et al. [5] compute low-lying spectra variationally in an SU(3) sphere matrix model. The model and its boundary-domain issues are inherited from [1,2]. Rayleigh–Ritz reduction, Gaussian integration, representation theory and residual estimates are not inventions of this study.

The construction documented here. The contribution is the explicit chain from the Helike direction through gauge-invariant probes to the $\mathbb{R}[T, P, d]$ calculus and two-generator spin-2 module, with exact moments for both H and its unprojected action, mixed-precision residual accounting, independent low-cutoff implementation checks, and source-pinned cutoff evidence. It clears a practical residual obstacle on the frozen controls while preserving the original physical problem. This is more specific than claiming that geometrical methods themselves are new.

Priority and utility are separate questions. The sources checked on 24 November establish substantial precedents for the ingredients. They do not settle priority for this exact combination and probe map. We therefore make no "first-ever" claim. The measured result stands without one. Nor is a smaller sector matrix by itself a benchmark against other solvers: the all-spin calculation answers a broader question.

The next scientific boundary. The new run does not recompute probe-response weights or their discarded tails, compare high-cutoff complete-level geometry, exclude every higher-spin competitor, or certify spectral separation. Its polynomial refinement concerns the selected matrix model; it does not restore the spatial field modes omitted in the sphere reduction. A held-out crossing experiment and the full continuum problem remain separate. The former 99.154% response is still an $N = 12$ observation, not an $N = 24$ result.

Additional references

[4] H.-P. Pavel, *SU(2) Yang–Mills quantum mechanics of spatially constant fields*, Phys. Lett. B 648 (2007), 97–106, arXiv:hep-th/0701283.

[5] N. Acharyya, A. P. Balachandran, M. Pandey, S. Sanyal and S. Vaidya, *Glueball Spectra from a Matrix Model of Pure Yang–Mills Theory*, Int. J. Mod. Phys. A 33 (2018), 1850073, arXiv:1606.08711v3.

[6] N. Acharyya, M. Pandey and S. Vaidya, *Axial Anomaly in SU(N) Yang–Mills Matrix Models*, Phys. Rev. Lett. 127 (2021), 092002.

[7] *The Floor Theorem for the Yang–Mills Hessian at the Berger Field: A Floor Law Across the Berger Spheres*, version 3.3, November 2026, repository manuscript.

16 / NEXT EXPERIMENTAL LAYER

Use GIGI to connect evidence across models

Plan, not a deployed capability. The model supplies observations. GIGI can organize those observations into versioned bundles, compare compatible geometric signatures and expose which features survive refinement. The goal is to follow a physical pattern through changing approximations, not merely line up similarly sized gap numbers.

Evidence layer	What must be retained
Model contract	Hamiltonian, domain, gauge group, geometry, units, boundary conditions and normalization.
Run and cutoff	Coupling, basis/mesh size, source and array hashes, resource data, status and all failed gates.
Complete level	Energy, multiplicity, symmetry sector, projector identity and residual; not an arbitrary eigenvector label.
Probe response	Full transition kernel, normalized spectral weights, vacuum-subtracted norm and discarded response.
Comparison map	The explicit bridge between models or cutoffs, its preserved quantities, errors and validity domain.

First deliverable: two evidence families, kept distinct. Import the original 15 all-spin cases and the 16 new sector cases with their immutable array references into isolated GIGI bundles. Recover every recorded gap, residual and available weight exactly. Follow the scalar and tensor branches separately, so the $N = 8$ ordering change cannot be hidden by sorting each run and calling its first row the same state.

Second deliverable: a measured geometric through-line. Within a common configuration space, embed successive cutoff subspaces into a common metric and compare complete-level projectors. Across genuinely different models, require a declared map or compare a common set of physical observables. The shared signature can include dimensionless energies, spin labels, spectral weight distributions and correlation kernels at fixed dimensionless times. An apparent match must survive normalization, uncertainty and adversarial controls.

Third deliverable: refinement before universality. Test convergence within each model separately. Then test whether the retained patterns agree across compatible models, on held-out cases. Raising this polynomial cutoff approaches the selected matrix model; it does not restore the spatial field modes omitted by that model. A continuum claim requires its own refinement parameters and limits.

The detailed plan is linked from the accompanying README. It separates current source-backed GIGI operations from the missing adapters and scientific gates. No graph spectral gap, PCA variance, or generic database curvature is relabelled as a Yang–Mills excitation energy.

Success criterion. GIGI should let a researcher ask: *Is this the same geometric excitation surviving a change of approximation, or only a similar number?* The answer must identify the retained pattern, its comparison map and the evidence that could falsify the match. It may also correctly return *not yet comparable*.

New evidence identity. Run `controls-abd81f97` uses source `abd81f97`; the full revision, hashes, source snapshots and replay instructions are linked in the README. Figures are regenerated from receipts without a new physical run. Verification is `PASS`; scientific acceptance remains `UNRESOLVED`. Typesetting this paper deploys nothing to GIGI.