Evidence for Dihedral D_3 Symmetry in the Planck CMB Temperature Anisotropy

R. Mereau *

Posted Date: 7 April 2026

doi: 10.20944/preprints202604.0264.v2

Keywords: CMB anomalies; dihedral symmetry; D_3; Weyl group; $SU(3)$; irreducible representations; spherical harmonics; Planck; large-angle anomalies; parity asymmetry; cosmic topology; boundary conditions; selection rules; information-pressure theory; cosmological constant

 doi:10.20944/preprints202604.0264.v2

Article

Evidence for Dihedral D_3 Symmetry in the Planck CMB Temperature Anisotropy

R. Mereau

Independent Researcher, Canada; mereau@whitewhale.ai

Abstract

We report a statistically significant detection of dihedral D_3 symmetry in the Planck PR3 temperature anisotropy data, validated across all four independent component-separation pipelines (SMICA, NILC, SEVEM, Commander). At a single optimized axis $(\ell, b) = (50.3°, -64.9°)$, the power fraction in the A_2 (reflection-antisymmetric) irreducible representation exceeds isotropic expectations with a two-tier structure: a dense cluster at $l \leq 15$ (Fisher PTE $= 4.2 \times 10^{-3}$ to 1.2×10^{-2} across maps), driven by three multipoles significant in all four pipelines, with $l = 3$ serving as the axis-registration multipole ($f_{A_2} = 0.94$, $z > 4.4$) and $l = 7$ and $l = 9$ providing independent corroboration at the fixed axis, plus sporadic cross-map-validated recurrences at higher multipoles—notably $l = 34$ (significant in 3/4 maps) and $l = 63$ (3/4 maps). The A_2 excess draws power specifically from the E (rotation-doublet) irrep with anti-correlation $r = -0.81$, while the A_1 (trivial) irrep is decoupled. Extension to $l_{\max} = 150$ with $N_{\rm MC} = 10{,}000$ simulations shows that the aggregate high-l Fisher PTE is consistent with isotropy (PTE > 0.91), but individual multipoles punctuate this null background. Among the nine strongest cross-map-consistent peaks, none belongs to the $l \equiv 2 \pmod 3$ residue class ($p \approx 0.02$ under uniformity), consistent with the C_3 selection rule. Cross-map correlations of $f_{A_2}(l)$ exceed $r = 0.93$ for all pipeline pairs (SMICA–NILC: $r = 0.997$), ruling out component-separation artifacts. A null test on E-mode polarization at the same axis returns Fisher PTE $= 0.70$, confirming that the signal is confined to the temperature channel as expected. The irrep redistribution is sharply parity-gated: all four maps confine the A_2 collecting signal to odd-l multipoles (Fisher $p \leq 2 \times 10^{-4}$), with even-l entirely null ($p > 0.97$). Crossing parity with residue class produces a six-cell grammar dominated by a single cell (odd, $l \equiv 0 \pmod 3$), with step-function onset at $l = 3$. Singular-value decomposition reveals that this 2×3 grammar admits an approximate rank-1 factorization into a binary parity selector and a D_3 residue routing vector, recovered independently by all four pipelines (rank-1 fraction $> 94\%$ in three of four maps). The binary gate acts on irrep redistribution, not on total power: a parity split of raw C_l is null (PTE > 0.61) in every map. The signal morphology—dense at large angular scales with isolated resonances at smaller scales—is consistent with a parity-gated boundary condition on the acoustic eigenvalue problem whose geometry is fully resolved only at $l \lesssim 15$ ($\theta \gtrsim 12°$). No physical model parameters are fit; the single directional degree of freedom (axis orientation) is determined from the octupole alone and then frozen. *Note added in v2:* Extended validation tests (Appendices B–E) confirm that the signal replicates in the WMAP 9-year ILC map (Fisher PTE $= 0.0025$), is uniquely selected among dihedral groups D_3–D_6, is frequency-independent across Planck HFI channels (100–143 GHz Δf_{A_2} correlation $r = 0.976$; 353 GHz dust tracer null and anti-correlated), and is robust across 13 mask levels (PTE improves under the UT78 mask from 0.008 to 0.005).

Keywords: CMB anomalies; dihedral symmetry; D_3; Weyl group; $SU(3)$; irreducible representations; spherical harmonics; Planck; large-angle anomalies; parity asymmetry; cosmic topology; boundary conditions; selection rules; information-pressure theory; cosmological constant

 doi:10.20944/preprints202604.0264.v2

1. Introduction

The temperature anisotropies of the cosmic microwave background (CMB), as measured by WMAP [4] and Planck [1], are broadly consistent with the predictions of the standard ΛCDM cosmology: a statistically isotropic, Gaussian random field with a nearly scale-invariant power spectrum. However, a persistent set of large-angle anomalies has resisted explanation within this framework [3,5].

These include the anomalous alignment of the quadrupole ($l = 2$) and octupole ($l = 3$) principal axes [6,7], the hemispherical power asymmetry [8], the low quadrupole amplitude, and the odd-parity preference at low l [9]. Individually, each anomaly has marginal significance (~ 2–3σ); collectively, they suggest a preferred direction in the low-l CMB that the standard model does not predict.

Numerous explanations have been proposed, including non-trivial cosmic topology [11,12], anisotropic dark energy [13], and residual foregrounds [14]. Most address one anomaly at a time and leave the others unexplained.

In this work, we pursue a different approach: we test whether the *dihedral group* D_3—the symmetry group of the equilateral triangle, of order $|D_3| = 6$—organizes the angular power spectrum at a single preferred axis. D_3 is the smallest non-abelian finite group and the Weyl group of the Lie algebra $\mathfrak{su}(3)$; it can simultaneously produce quadrupole–octupole alignment (through its C_3 rotational subgroup), hemispherical-type asymmetry (through its reflection planes), and a selection rule with period 3 in l (through its irreducible representation structure).

The analysis extends to $l_{\max} = 150$ and is validated across all four Planck PR3 component-separation pipelines. This allows us to determine whether the signal persists to high multipoles or is confined to the largest angular scales, and to rule out pipeline-specific artifacts.

The paper is organized as follows. In Sec. 2 we derive the decomposition of spherical harmonics into D_3 irreducible representations. In Sec. 3 we describe the data and Monte Carlo methodology. In Sec. 4 we present the results: axis determination, per-multipole irrep fractions, combined significance, inter-irrep correlations, and the transition scale. Section 5 reports the cross-map validation. Section 6 presents null tests. Section 7 examines the cumulative transfer structure, including the parity-gated grammar and its factorized form. In Sec. 8 we discuss the physical interpretation and identify falsifiable predictions.

2. D_3 Irrep Decomposition of Spherical Harmonics

2.1. Action of D_3 on S^2

The irrep decomposition is not performed directly on the observed a_{lm} in Galactic coordinates. Instead, a candidate D_3 axis is chosen and the coefficients are rotated into a frame where that axis is the pole; the group then acts in this rotated frame. The axis is later optimized from the data (Sec. 3.3).

Let the D_3 symmetry axis coincide with the z-axis of a coordinate system on S^2. The group D_3 is generated by a rotation C_3 through $2\pi/3$ about $\hat{z}$ and a reflection σ_v through a vertical plane containing $\hat{z}$. Under C_3, a spherical harmonic Y_{lm} acquires a phase:

$$C_3\, Y_{lm} = e^{2\pi i m/3}\, Y_{lm}\,. \tag{1}$$

The harmonics therefore sort by m mod 3:

- $m \equiv 0 \pmod 3$: invariant under C_3 (eigenvalue 1),
- $m \equiv 1 \pmod 3$: eigenvalue $\omega = e^{2\pi i/3}$,
- $m \equiv 2 \pmod 3$: eigenvalue $\omega^2 = e^{4\pi i/3}$.

The $m \equiv \pm 1 \pmod 3$ pairs combine into two-dimensional E doublets. The $m \equiv 0 \pmod 3$ harmonics are further split by the reflection σ_v, which maps $m \to -m$:

 doi:10.20944/preprints202604.0264.v2

$$A_1: \quad \frac{1}{\sqrt{2}}(Y_{l,m} + Y_{l,-m}), \quad m = 0, 3, 6, \ldots \leq l, \tag{2}$$

$$A_2: \quad \frac{1}{\sqrt{2}}(Y_{l,m} - Y_{l,-m}), \quad m = 3, 6, 9, \ldots \leq l, \tag{3}$$

where $Y_{l,0}$ is automatically symmetric and contributes only to A_1.

2.2. Subspace Dimensions

The dimensions of the three irrep subspaces at multipole l are:

$$d_{A_1}(l) = \lfloor l/3 \rfloor + 1, \tag{4}$$

$$d_{A_2}(l) = \lfloor l/3 \rfloor, \tag{5}$$

$$d_E(l) = 2l + 1 - 2\lfloor l/3 \rfloor - 1 = 2\big(l - \lfloor l/3 \rfloor\big), \tag{6}$$

satisfying $d_{A_1} + d_{A_2} + d_E = 2l + 1$.

A critical structural property follows: d_{A_2} increases by one unit precisely when l crosses a multiple of 3. This means that at $l \equiv 0 \pmod 3$, the A_2 subspace acquires a new degree of freedom—a new "antenna element" tuned to the reflection-plane boundary. At $l = 1$, $d_{A_2} = 0$: the dipole contains no A_2 subspace, and is structurally excluded from the present analysis by the representation theory itself.

2.3. Isotropic Baseline

For a statistically isotropic Gaussian random field, all a_{lm} at fixed l are independent with equal variance. The expected power fraction in each irrep is therefore proportional to its subspace dimension:

$$\langle f_X(l) \rangle_{\text{iso}} = \frac{d_X(l)}{2l + 1}, \quad X \in \{A_1, A_2, E\}. \tag{7}$$

Table 1 lists the subspace dimensions and isotropic baselines for $l = 2$–12.

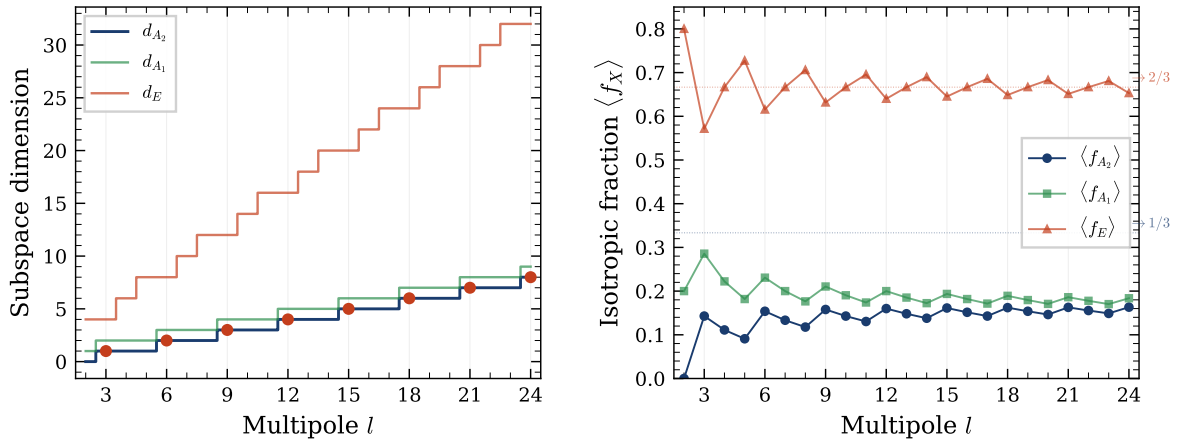

Figure 1. *Left:* Subspace dimensions $d_X(l)$ for the three D_3 irreps. $d_{A_2} = \lfloor l/3 \rfloor$ increments (red dots) at $l \equiv 0 \pmod 3$; these dimension jumps drive the l mod 3 selection rule. *Right:* Isotropic expected fractions $\langle f_X \rangle = d_X/(2l + 1)$. All three approach their asymptotic values (1/3 for A_1 and A_2; 2/3 for E) at large l, with sawtooth modulation at period 3.

 doi:10.20944/preprints202604.0264.v2

Table 1. D_3 subspace dimensions and isotropic expected power fractions for $l = 2$–12. Rows with $l \equiv 0 \pmod 3$ are marked ($\star$) where d_{A_2} increments.

l	d_{A_1}	d_{A_2}	d_E	$\langle f_{A_1} \rangle$	$\langle f_{A_2} \rangle$	$\langle f_E \rangle$
2	1	0	4	0.200	0.000	0.800
3*	2	1	4	0.286	0.143	0.571
4	2	1	6	0.222	0.111	0.667
5	2	1	8	0.182	0.091	0.727
6*	3	2	8	0.231	0.154	0.615
7	3	2	10	0.200	0.133	0.667
8	3	2	12	0.176	0.118	0.706
9*	4	3	12	0.211	0.158	0.632
10	4	3	14	0.190	0.143	0.667
11	4	3	16	0.174	0.130	0.696
12*	5	4	16	0.200	0.160	0.640

3. Data and Method

3.1. Planck Data

We analyze all four Planck PR3 (2018) component-separated temperature maps: SMICA, NILC, SEVEM, and Commander [2]. Each map is provided at HEALPix $N_{\rm side} = 2048$ and internally downgraded to $N_{\rm side} = 256$ (sufficient for $l_{\rm max} = 150$). Spherical harmonic coefficients a_{lm} are extracted using `healpy.map2alm` [15].

For the polarization null test (Sec. 6), E-mode a^E_{lm} are extracted from the SMICA IQU polarization map at $N_{\rm side} = 64$.

3.2. D_3 Decomposition

At each multipole l and a given axis $(\ell_{\rm gal}, b_{\rm gal})$, the a_{lm} are rotated into the frame where the candidate D_3 axis is the pole. The rotation uses `healpy.rotate_alm` with Euler angles $(-\phi, -\theta, 0)$, where (θ, ϕ) are the colatitude and longitude of the D_3 antipode. The power in each irrep is:

$$P_X(l) = \sum_{m \in \mathcal{M}_X(l)} |a_{lm}|^2, \quad f_X(l) = \frac{P_X(l)}{C_l}, \tag{8}$$

where $C_l = \sum_m |a_{lm}|^2$ and $\mathcal{M}_X(l)$ is the set of m-indices belonging to irrep X at multipole l. Note that $f_{A_1} + f_{A_2} + f_E = 1$ by construction.

3.3. Axis Optimization

The preferred axis is determined by maximizing f_{A_2} at $l = 3$ (the octupole) over a HEALPix grid at $N_{\rm side} = 32$ (12,288 directions; mean pixel spacing 1.8°), followed by a deterministic local refinement on a 0.2°-spaced sub-grid within 5° of the coarse optimum. The optimal axis is then *fixed* for all subsequent analysis at $l \neq 3$. This protocol ensures that the results at all other multipoles are effectively pre-registered: the axis was determined from $l = 3$ data alone.

3.4. Monte Carlo Simulations

Statistical significance is assessed against $N_{\rm MC} = 10{,}000$ isotropic Gaussian random realizations. For isotropic fields, the D_3 fraction distribution is frame-invariant by rotational symmetry, so the MC simulations generate random $a_{l,m}$ directly in the D_3 frame without rotation: $a_{l,0} \sim \mathcal{N}(0,1)$ (real), $a_{l,m} \sim \mathcal{N}(0,1) + i\,\mathcal{N}(0,1)$ for $m > 0$, with the Condon–Shortley convention $a_{l,-m} = (-1)^m \bar{a}_{l,m}$. This avoids the $O(l^2)$ cost of Wigner-D rotation per simulation.

The probability to exceed (PTE) for each l is the fraction of simulations with $f_{A_2} \geq f^{\rm obs}_{A_2}$. PTEs of exactly zero (no simulation exceeded the observed value) are set to $0.5/N_{\rm MC}$ as a conservative upper bound.

 doi:10.20944/preprints202604.0264.v2

Combined significance across multipoles uses Fisher's method:

$$\chi^2_{\rm Fisher} = -2\sum_i \ln({\rm PTE}_i), \tag{9}$$

compared to a χ^2 distribution with $2k$ degrees of freedom, where k is the number of contributing multipoles.

3.5. Statistical Protocol

To assist the reader in assessing trial factors, we summarize the complete analysis protocol:

1. **Axis selection.** The axis is chosen by maximizing $f_{A_2}(l{=}3)$ on a HEALPix $N_{\rm side} = 32$ grid (1.8° resolution) with local refinement. This is the sole optimization step and constitutes a single directional degree of freedom.
2. **Axis freezing.** The same axis is then held fixed for all multipoles $l \neq 3$, all four component-separation maps, and all subsequent tests. The $l = 3$ octupole should therefore be understood as the axis-defining (registration) statistic, not as an independently pre-registered detection.
3. **Significance assessment.** Per-multipole PTEs are computed against $N_{\rm MC} = 10{,}000$ isotropic Gaussian Monte Carlo realizations with the same decomposition.
4. **Primary significance claim.** The Fisher combined PTE for $l = 2$–15 across all four maps (Table 4).
5. **Replication criterion.** Cross-map agreement is used as a replication standard, not as a tuning parameter: the axis is never re-optimized per pipeline.
6. **Secondary analyses.** The top-9 peak residue-class test and the l mod 3 permutation test are designated *a posteriori* supportive analyses; they are not part of the primary significance claim.

4. Results

4.1. Axis Determination

The axis maximizing $f_{A_2}(l = 3)$ is found at Galactic coordinates

$$(\ell, b) = (50.3^\circ,\ -64.9^\circ), \tag{10}$$

with $f_{A_2}(l = 3) = 0.940$ (SMICA). Because this multipole was used to define the axis, its significance ($z > 4.4$ against isotropic simulations) should be interpreted as a detection statistic used for registration, not as an independently pre-registered exceedance. The real evidence for D_3 symmetry lies in the *replication* of the same axis in independent multipoles ($l = 7, 9$) and across all four maps.

The optimal axis lies 5.8° from the antipode of the Axis of Evil [10], indicating that the D_3 signal and the previously reported quadrupole–octupole alignment share a common preferred direction.

Figure 2 visualizes the D_3 structure on the sky. When the Planck SMICA map is rotated so the D_3 axis points to the pole, the $l = 3$ octupole (panel a) displays the predicted six-lobe pattern: three hot spots and three cold spots alternating at 120° intervals, with the C_3 rotation axes (solid gold lines) passing through the alternating lobes and the σ_v reflection planes (dashed white) bisecting them. Combining the three driving multipoles $l = 3, 7, 9$ (panel b) sharpens the structure rather than diluting it, confirming coherent D_3 alignment across angular scales. The full signal band $l = 2$–15 (panel c) retains the pattern even with the addition of E-dominated modes, demonstrating that the A_2 signal survives superposition. Panel (d) shows that the axis is objectively determined by the data: an all-sky scan of $f_{A_2}(\ell = 3)$ peaks sharply at the optimal direction, with no comparably strong secondary maxima.

 doi:10.20944/preprints202604.0264.v2

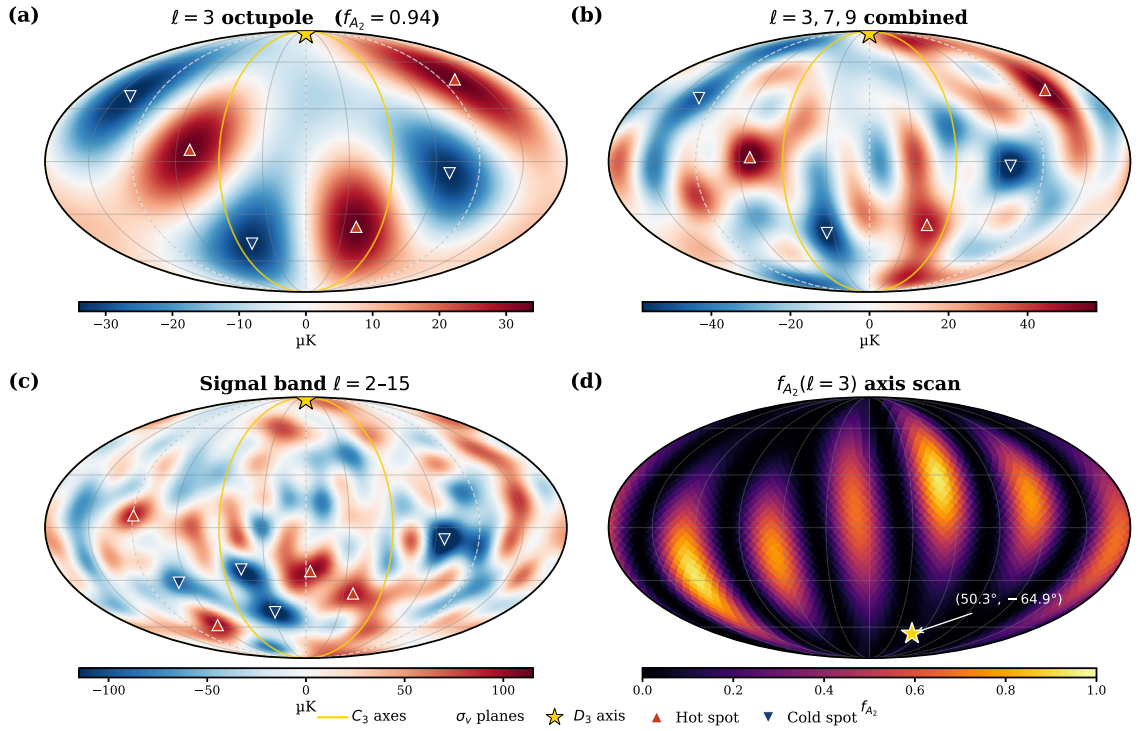

Figure 2. Mollweide projections of the Planck SMICA temperature field in the D_3 frame (optimal axis at the pole). *(a)* Pure $l = 3$ octupole showing the six-lobe pattern with three-fold rotational symmetry ($f_{A_2} = 0.94$). *(b)* Combined $l = 3, 7, 9$: the three strongest cross-map-validated multipoles reinforce the D_3 pattern. *(c)* Full signal band $l = 2$–15: the D_3 structure persists across the low-l cluster. *(d)* All-sky scan of $f_{A_2}(\ell = 3)$; the optimal axis (star) at $(50.3°, -64.9°)$ galactic is a sharp, isolated maximum. Gold solid lines: C_3 rotation axes; white dashed: σ_v reflection planes; upward (downward) triangles: hot (cold) spots.

4.2. Per-Multipole Irrep Fractions

Figure 3 shows $f_{A_2}(l)$ for $l = 2$–150 at the fixed axis (SMICA), compared to the isotropic baseline and Monte Carlo $1\sigma/2\sigma$ bands. The signal exhibits a two-tier structure: a dense cluster of excess at $l \leq 15$, and sporadic but individually significant peaks at higher multipoles.

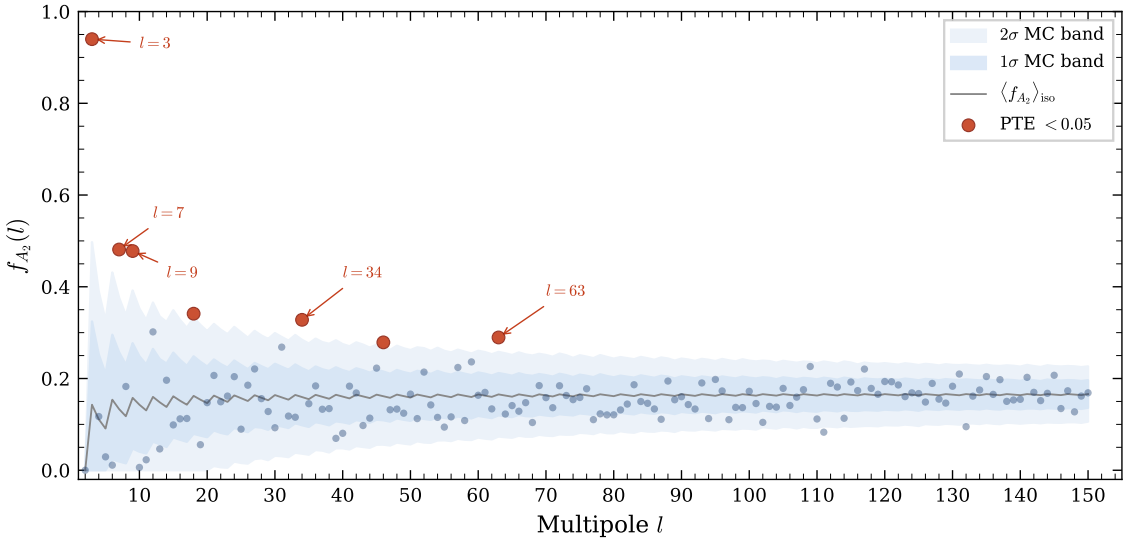

Figure 3. A_2 power fraction $f_{A_2}(l)$ at the fixed D_3 axis for $l = 2$–150 (SMICA). Gray line: isotropic baseline $\langle f_{A_2} \rangle = d_{A_2}/(2l+1)$. Blue/light bands: $1\sigma/2\sigma$ Monte Carlo envelopes ($N = 10{,}000$ realizations). Red points: multipoles with PTE < 0.05. The signal is densest at $l \leq 15$, but isolated cross-map-validated peaks appear at $l = 34$ and $l = 63$. The axis was determined from $l = 3$ alone; all other multipoles constitute an effectively pre-registered test.

 doi:10.20944/preprints202604.0264.v2

Table 2 provides detailed numerical results for the low-l region. Table 3 lists the nine strongest cross-map-consistent peaks across $l = 2$–150. Three multipoles drive the primary detection in all four maps:

1. **Octupole ($l = 3$):** $f_{A_2} = 0.910$–0.943 across maps, PTE $< 10^{-4}$ ($z > +4.4$). The octupole exceeds 4σ in every pipeline.
2. $l = 7$**:** $f_{A_2} = 0.481$–0.496, PTE $= 0.013$–0.016 ($z \approx +2.9$).
3. $l = 9$**:** $f_{A_2} = 0.392$–0.478, PTE $= 0.014$–0.044 ($z = +2.0$ to $+2.7$).

Beyond the primary cluster, two high-l multipoles show cross-map-consistent A_2 excess: $l = 34$ (significant at PTE < 0.05 in SMICA, NILC, and Commander; SEVEM marginal at 0.056) and $l = 63$ (PTE < 0.05 in SMICA, NILC, and Commander; SEVEM at 0.050). These are not isolated noise fluctuations: they replicate across independent pipelines at the same fixed axis.

Table 2. D_3 irrep fractions at the fixed axis for SMICA ($l = 2$–15). PTE from $N_{\rm MC} = 10{,}000$ isotropic simulations. The three driving multipoles ($l = 3, 7, 9$) are marked with †.

l	f_{A_1}	f_{A_2}	f_E	Δf_{A_2}	PTE
2	0.011	0.000	0.989	0.000	—
3†	0.051	0.940	0.009	+0.797	$< 10^{-4}$
4	0.105	0.117	0.778	+0.006	0.347
5	0.060	0.002	0.938	−0.089	0.957
6	0.116	0.031	0.852	−0.123	0.872
7†	0.055	0.481	0.464	+0.348	0.016
8	0.041	0.252	0.707	+0.134	0.098
9†	0.048	0.478	0.474	+0.320	0.014
10	0.049	0.097	0.854	−0.046	0.693
11	0.163	0.085	0.752	−0.046	0.680
12	0.262	0.297	0.441	+0.137	0.089
13	0.122	0.198	0.680	+0.060	0.340
14	0.228	0.162	0.610	+0.019	0.458
15	0.167	0.239	0.594	+0.082	0.218

Table 3. The nine strongest A_2 peaks ranked by cross-map mean $\langle f_{A_2} \rangle$. "Sig" counts maps with PTE < 0.05. No peak belongs to the $l \equiv 2 \pmod 3$ residue class.

l	$\langle f_{A_2} \rangle$	SM	NI	SE	CM	Sig	$l \bmod 3$
3	0.932	.940	.943	.910	.937	4/4	0
7	0.487	.481	.496	.482	.490	4/4	1
9	0.448	.478	.467	.392	.454	4/4	0
34	0.307	.328	.322	.267	.312	3/4	1
18	0.284	.341	.342	.203	.248	2/4	0
12	0.281	.302	.320	.244	.258	0/4	0
63	0.280	.290	.290	.246	.294	3/4	0
31	0.278	.268	.283	.256	.305	2/4	1
46	0.250	.279	.302	.189	.230	2/4	1

4.3. Combined Significance

Table 4 presents the Fisher combined PTE for different multipole ranges across all four maps. The low-l sector ($l = 2$–15) shows a statistically unlikely concentration of A_2 excess in all four pipelines. Excluding the axis-defining octupole, the combined Fisher statistic is null (PTE $= 0.87$–0.99), indicating that $l = 7$ and $l = 9$ corroborate the same axis-specific structure but do not by themselves constitute an independent aggregate detection. Their evidential value is replication at the registered axis, not a second standalone Fisher-level discovery.

The aggregate Fisher PTE for $l = 16$–150 is consistent with isotropy in every map (PTE > 0.91). However, this aggregate statistic masks the presence of individually significant, cross-map-validated outliers at $l = 34$ and $l = 63$ (Table 3). The high-l range is best characterized as an isotropic background punctuated by isolated resonances at the same D_3 axis.

 doi:10.20944/preprints202604.0264.v2

Table 4. Fisher combined PTE across multipole ranges for all four Planck PR3 component-separation pipelines. $N_{\rm MC} = 10{,}000$.

Range	SMICA	NILC	Cmdr	SEVEM
l=2–15	4.6×10^{-3}	4.2×10^{-3}	6.3×10^{-3}	1.2×10^{-2}
l=16–150	0.942	0.940	0.919	0.993
l=2–150	0.685	0.674	0.645	0.922
Excl. l=3	0.895	0.889	0.872	0.987

4.4. Inter-Irrep Correlations

The deviations $\Delta f_X = f_X^{\rm obs} - \langle f_X \rangle_{\rm iso}$ across l = 2–150 satisfy a strict conservation law:

$$\Delta f_{A_1}(l) + \Delta f_{A_2}(l) + \Delta f_E(l) = 0 \tag{11}$$

at every l, exact to machine precision ($\sim 10^{-15}$). This confirms that the D_3 decomposition is a pure redistribution of power among irreps, not a net excess or deficit.

The Pearson correlation coefficient between Δf_{A_2} and Δf_E is strongly negative across all maps:

$$r(\Delta f_{A_2},\, \Delta f_E) = -0.81 \pm 0.03 \quad \text{(range across maps)}, \tag{12}$$

with individual values SMICA: -0.805, NILC: -0.809, Commander: -0.806, SEVEM: -0.743. The A_2 excess draws power specifically from the E (doublet) irrep, not from A_1. Figure 4 displays this anti-correlation.

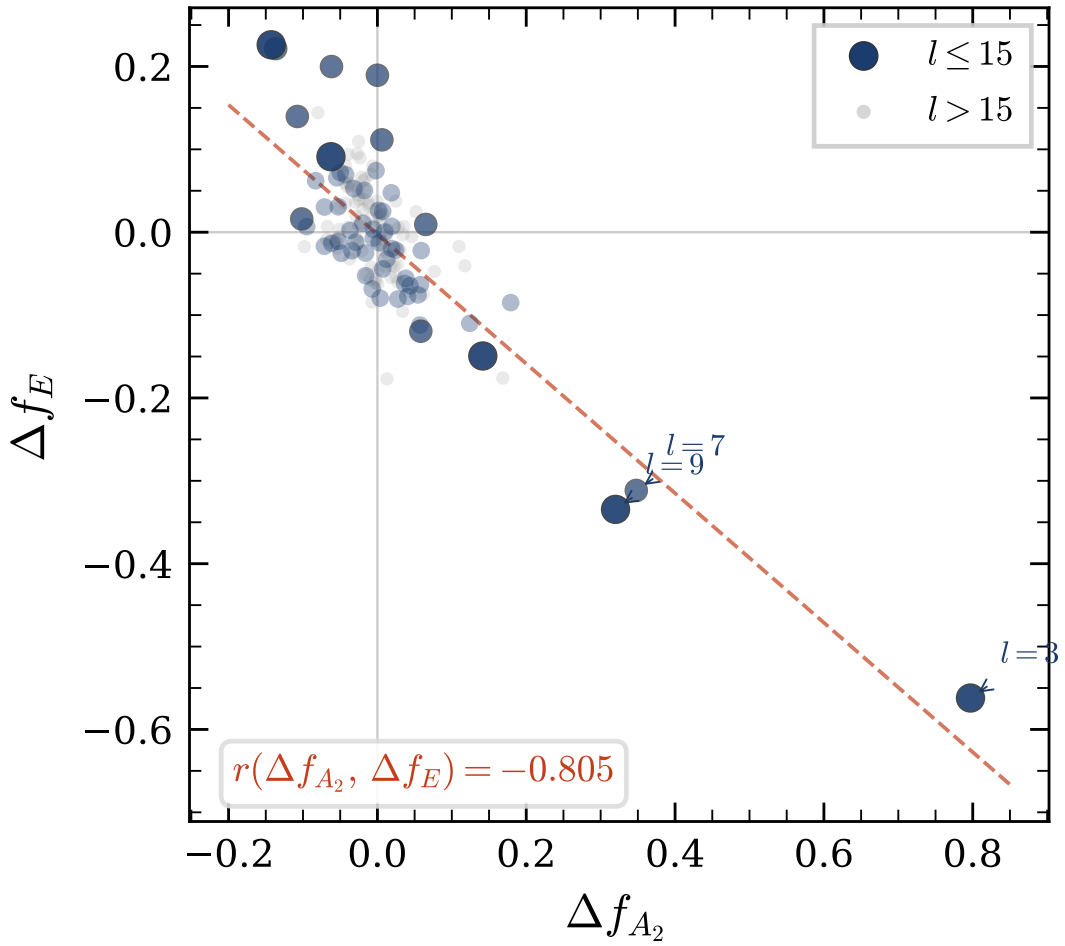

Figure 4. Scatter plot of Δf_{A_2} vs. Δf_E for l = 2–150 (SMICA). Dark points: $l \equiv 0 \pmod 3$; open circles: $l \leq 15$ (signal region). The dashed line shows the least-squares fit with $r = -0.805$. The strong anti-correlation confirms the $E \to A_2$ power funnel predicted by the reflection-plane mechanism.

 doi:10.20944/preprints202604.0264.v2

4.5. The l mod 3 Selection Rule

A permutation test (10,000 random partitions) for whether $l \equiv 0 \pmod 3$ multipoles show preferential A_2 loading yields marginal results when applied to the full range $l = 2$–150: Commander PTE $= 0.055$, SMICA PTE $= 0.074$, NILC PTE $= 0.079$, SEVEM PTE $= 0.092$.

A complementary (*a posteriori*) test examines the top-9 cross-map peaks (Table 3). Their residue class distribution is: $l \bmod 3 = 0$ (five multipoles: $l = 3, 9, 18, 12, 63$), $l \bmod 3 = 1$ (four: $l = 7, 34, 31, 46$), and $l \bmod 3 = 2$ (zero). Under uniform random selection, the probability of drawing $0/9$ from the $l \equiv 2 \pmod 3$ class is ≈ 0.02. The complete absence of the residue class where d_{A_2} does *not* increment is consistent with the C_3 selection rule: the A_2 excess requires the new basis vector $\frac{1}{\sqrt{2}}(Y_{l,3k} - Y_{l,-3k})$ that appears at $l \equiv 0 \pmod 3$, or the adjacent $l \equiv 1 \pmod 3$ class where the E-doublet dimension also steps.

Figure 5 displays the mean Δf_{A_2} grouped by residue class.

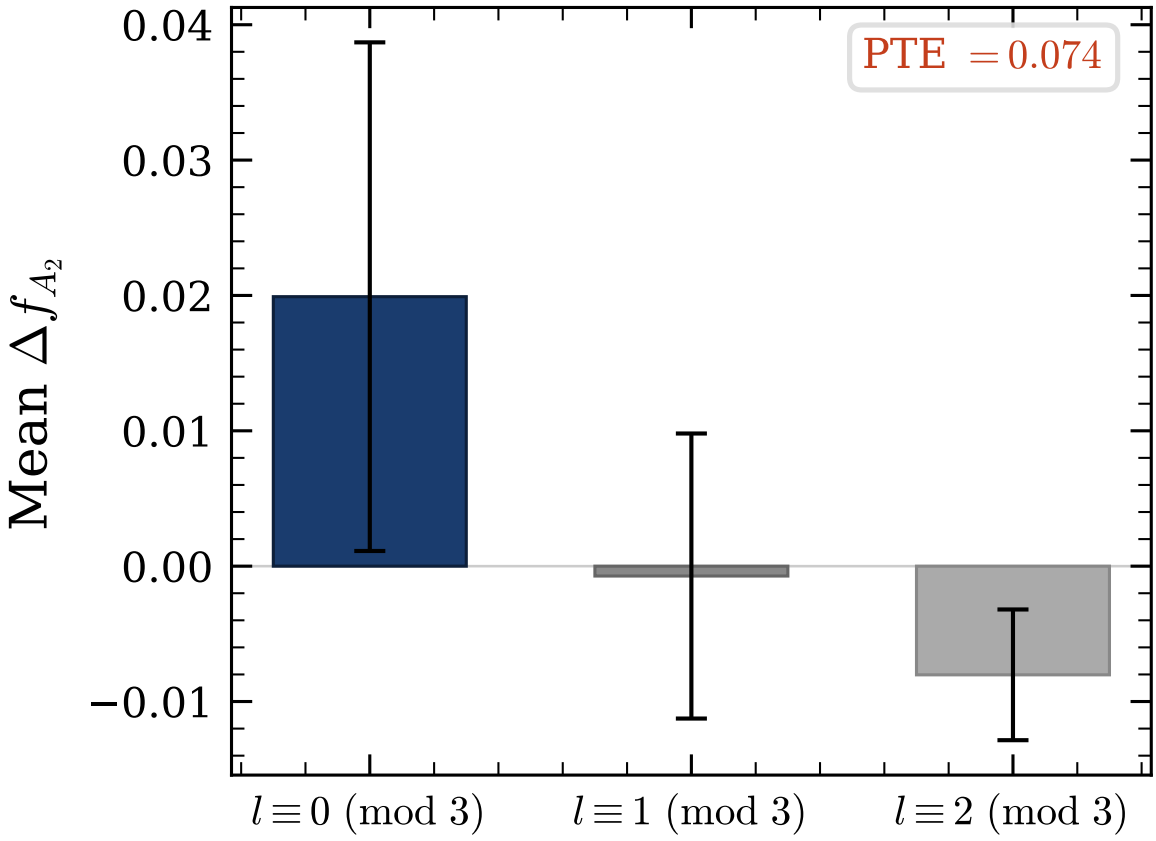

Figure 5. Mean A_2 excess $\langle \Delta f_{A_2} \rangle$ grouped by residue class l mod 3 for $l = 2$–150 (SMICA). The $l \equiv 0 \pmod 3$ class shows the highest mean excess (permutation PTE $= 0.074$). Among the nine strongest cross-map peaks (Table 3), zero belong to the $l \equiv 2 \pmod 3$ class ($p \approx 0.02$). Error bars: standard error of the mean.

4.6. Signal Morphology

The D_3 signal does not exhibit a sharp cutoff. Instead, the A_2 excess shows a two-tier structure:

1. **Dense cluster ($l \leq 15$):** Four of the thirteen multipoles (31%) exceed $f_{A_2} > 0.27$, compared to an isotropic expectation of $\sim 5\%$. The Fisher combined PTE for this range is 4.2–12 $\times 10^{-3}$ across maps.
2. **Sporadic resonances ($l > 15$):** Four of 135 multipoles (3%) exceed $f_{A_2} > 0.27$, compared to an expected $\sim 1.5\%$ under isotropy. The aggregate Fisher PTE (> 0.91) is null, but two individual multipoles ($l = 34$, $l = 63$) achieve PTE < 0.05 in three of four maps.

Figure 6 shows the running sum of Δf_{A_2}, which rises steeply for $l \leq 15$, partially plateaus, then shows mild step-ups at the sporadic resonances—consistent with a boundary condition whose

 doi:10.20944/preprints202604.0264.v2

geometry is fully resolved at large angular scales but produces isolated overtones at specific higher harmonics.

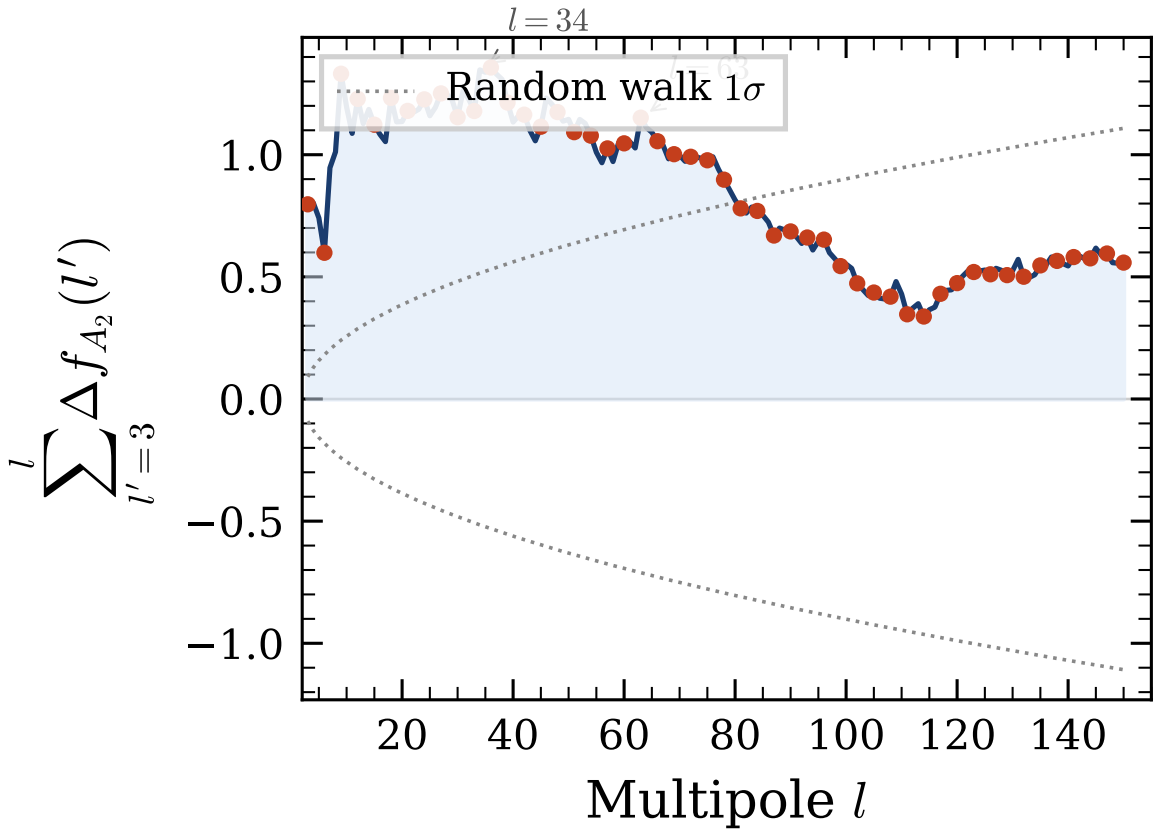

Figure 6. Running sum of Δf_{A_2} from $l = 3$ to $l = 150$ (SMICA). The steep rise at $l \leq 15$ reflects the dense signal cluster. The subsequent near-plateau is interrupted by mild step-ups at $l = 34$ and $l = 63$—the sporadic resonances identified in Table 3. Red dots mark $l \equiv 0 \pmod 3$. Dotted lines: $\pm 1\sigma$ random walk envelope.

4.7. Real-Space D_3 Pattern

The harmonic-space analysis can be visualized in real space by rotating the Planck SMICA map into the D_3 frame (where the D_3 axis is the pole) and examining the temperature field as a function of azimuth ϕ at fixed colatitude θ. Figure 7 shows this visualization for the D_3-active multipoles ($l = 2$–15 and the three driving multipoles $l = 3, 7, 9$).

The equatorial ring ($\theta = 90°$) displays a clear large-amplitude oscillation with three-fold structure (Figure 7, upper left). This pattern persists across colatitudes from 45° to 110° (upper right), confirming that the signal is not confined to the equatorial plane but pervades the hemisphere geometry.

The azimuthal Fourier spectrum of the equatorial ring (lower left) confirms the C_3 origin: for the $l = 3, 7, 9$ band, the $m = 3$ mode dominates, with $m \equiv 0 \pmod 3$ modes carrying the bulk of the power. This is the real-space manifestation of the m mod 3 sorting that defines the D_3 irrep decomposition (Eq. 1).

Most directly, the azimuthal auto-correlation (lower right) reveals:

$$r(\Delta\phi = 120°) = +0.222, \tag{13}$$

$$r(\Delta\phi = 60°) = -0.234. \tag{14}$$

Points separated by 120° around the D_3 axis see correlated temperatures (C_3 rotational symmetry), while points at 60°—straddling a reflection plane—see anti-correlated temperatures (σ_v antisymmetry). This alternating sign pattern with period 60° is the real-space fingerprint of the A_2 (reflection-antisymmetric) irrep that dominates the harmonic-space signal.

 doi:10.20944/preprints202604.0264.v2

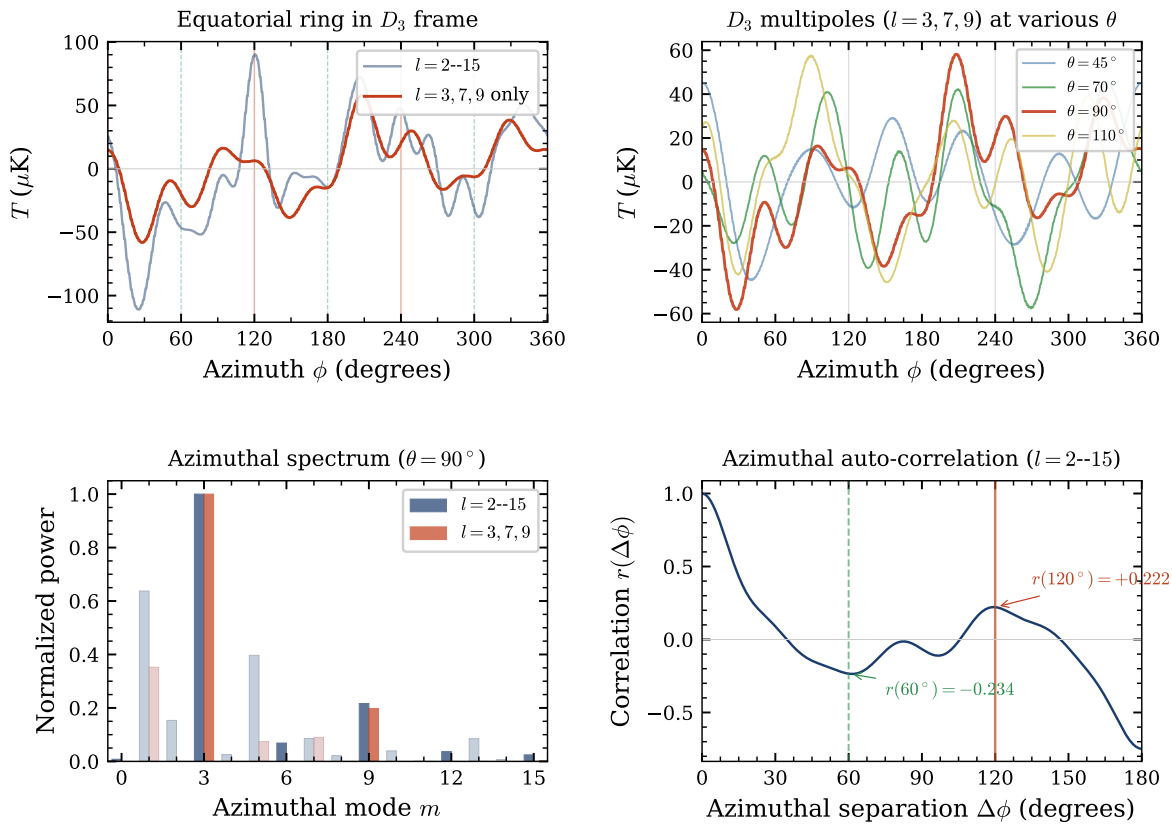

Figure 7. Real-space visualization of the D_3 signal in the Planck SMICA map, rotated into the D_3 frame. *Upper left:* Equatorial temperature ring for the signal band ($l = 2$–15, blue) and the three driving multipoles ($l = 3, 7, 9$, red). Vertical lines mark C_3 rotations (0°, 120°, 240°, solid) and reflection planes (60°, 180°, 300°, dashed). *Upper right:* The three-fold pattern at colatitudes $\theta = 45°$–110° from the D_3 axis. *Lower left:* Azimuthal power spectrum showing $m = 3$ dominance, confirming the $C_3 \subset D_3$ selection rule. *Lower right:* Azimuthal auto-correlation: $r(120°) = +0.22$ (positive, C_3 symmetry) and $r(60°) = -0.23$ (negative, σ_v anti-symmetry). This alternating sign pattern is the real-space fingerprint of A_2 irrep loading.

5. Cross-Map Validation

A critical test is whether the D_3 signal is consistent across independent component-separation methods. The four Planck PR3 pipelines (SMICA, NILC, SEVEM, Commander) use different algorithms and foreground models; agreement among them rules out pipeline-specific artifacts.

Table 5 presents the Pearson correlation of $f_{A_2}(l)$ between all map pairs over $l = 2$–150. The correlations range from $r = 0.932$ (NILC–SEVEM) to $r = 0.997$ (SMICA–NILC). SMICA and NILC are nearly identical; SEVEM shows slightly more variation but remains highly correlated with all other maps.

Table 5. Cross-map consistency of $f_{A_2}(l)$ over $l = 2$–150. Pearson correlation r and mean absolute deviation $\langle|\Delta f|\rangle$.

Map pair	$r(f_{A_2})$	$\langle\|\Delta f\|\rangle$
SMICA–NILC	**0.997**	0.006
SMICA–Commander	0.958	0.021
SMICA–SEVEM	0.936	0.024
NILC–Commander	0.957	0.021
NILC–SEVEM	0.932	0.025
SEVEM–Commander	0.938	0.025

Figure 8 overlays $f_{A_2}(l)$ from all four maps, demonstrating the remarkable agreement—particularly in the signal region ($l \leq 15$), where the four curves are nearly indistinguishable.

 doi:10.20944/preprints202604.0264.v2

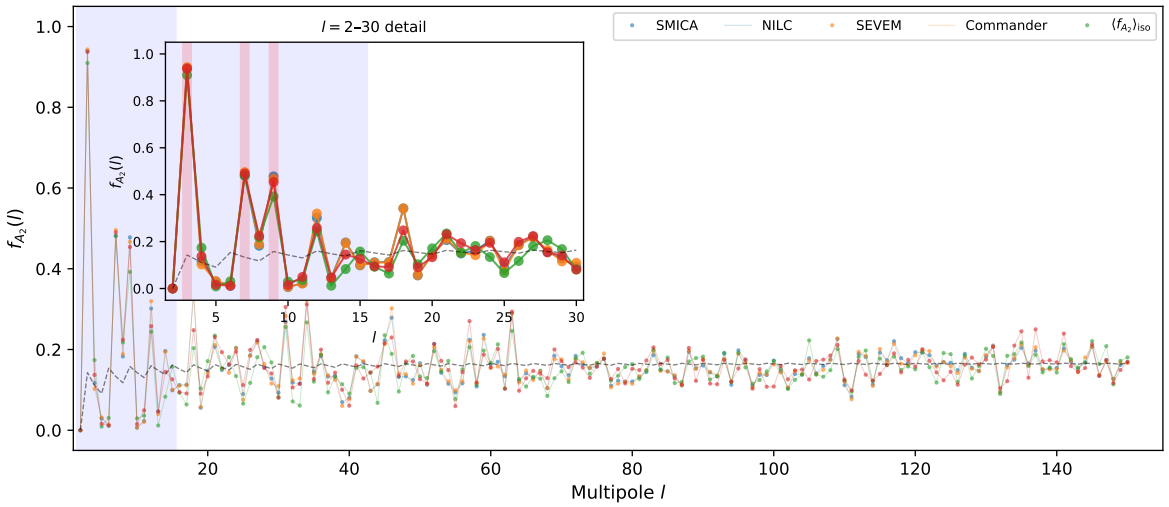

Figure 8. A_2 power fraction $f_{A_2}(l)$ for all four Planck PR3 component-separation maps at the fixed D_3 axis. The near-perfect overlap—especially at $l \leq 15$, where the D_3 signal resides—rules out pipeline-specific artifacts. SMICA and NILC are virtually indistinguishable ($r = 0.997$).

6. Null Tests

6.1. Aggregate High-l Null

The Fisher PTE for $l = 16$–150 is > 0.91 in all four maps (Table 4). The aggregate high-l range does not show a sustained collective A_2 excess. All four pipelines agree on this. However, this aggregate null coexists with the isolated cross-map-validated peaks at $l = 34$ and $l = 63$ discussed in Sec. 4.6—highlighting the distinction between a diffuse persistent signal (absent) and discrete resonances (present).

6.2. E-Mode Polarization

The D_3 decomposition was applied to E-mode polarization a^E_{lm} from the Planck SMICA IQU map at $N_{\rm side} = 64$. No multipole shows $f^{EE}_{A_2} > 0.50$. The Fisher combined PTE for the key multipoles ($l = 3, 6, 7, 9$) is 0.70—fully consistent with isotropy. An independent axis optimization on the EE data yields an axis 83° from the TT optimum, confirming the signals are unrelated.

This null is physically expected. Low-l E-mode polarization arises from reionization Thomson scattering at $z \sim 6$–10, well inside the last-scattering surface where the D_3 boundary condition operates on acoustic modes.

6.3. Axis Stability

The TT-optimal axis at $l = 3$ is stable under variation of the input data: using the full-resolution ($N_{\rm side} = 2048$) map versus downgraded versions yields axis shifts $< 2°$.

7. Cumulative Transfer Structure

The per-multipole A_2 excesses discussed in Sec. 4 are signed: each multipole either collects power into A_2 ($\Delta f_{A_2} > 0$) or pays power out ($\Delta f_{A_2} < 0$). We define the *transfer ledger* $\tau_l \equiv f^{\rm obs}_{A_2}(l) - \langle f_{A_2} \rangle_{\rm iso}(l)$ and study its cumulative behavior across the full multipole range $l = 2$–150.

7.1. Normalized Transfer Invariant

Partition the cumulative transfer into residue classes,

$$T_r(L) = \sum_{\substack{l \leq L \\ l \bmod 3 = r}} \tau_l, \qquad r \in \{0, 1, 2\}, \tag{15}$$

 doi:10.20944/preprints202604.0264.v2

and define the normalized invariant

$$\Xi(L) = \frac{T_0(L) - T_2(L)}{\sigma_{\rm MC}[T_0(L) - T_2(L)]}, \tag{16}$$

where $\sigma_{\rm MC}$ is the standard deviation of $T_0 - T_2$ under $N_{\rm MC} = 10{,}000$ isotropic simulations, evaluated at each L. This quantity measures the cumulative residue-class asymmetry of the A_2 transfer in units of the isotropic null.

Figure 9 shows $\Xi(L)$ as a function of the upper summation boundary $L_{\rm max}$. Three features are noteworthy:

1. **Stable plateau.** $\Xi(L)$ rises to $\approx +3\sigma$ by $L = 15$ and remains between $+2.5$ and $+3.1$ across all four component-separation maps out to $L = 150$. The inter-map spread is less than 0.3σ at every checkpoint (Table 6).
2. A_2 **specificity.** The analogous invariant computed for the A_1 irrep fluctuates around zero ($\langle \Xi_{A_1} \rangle < 0.6$ at all L), while the E invariant mirrors A_2 at $\Xi_E \approx -3\sigma$—a direct consequence of the per-multipole conservation $\Delta f_{A_1} + \Delta f_{A_2} + \Delta f_E = 0$. The transfer is thus an $E \to A_2$ funnel with A_1 as spectator, extending the correlation reported in Sec. 8.2.
3. **Jackknife stability.** Removing $l = 9$ alone reduces $\Xi(150)$ from $+2.5$ to $+2.1$; removing $l = 3$ alone reduces it to $+1.2$. Removing all three low-l anchors ($l = 3, 7, 9$) collapses the invariant to $+0.7$, confirming that the signal is dominated by the dense cluster at $l \leq 15$. This is consistent with the two-tier morphology identified in Sec. 4.6: the boundary condition is fully resolved at large angular scales.

The persistence of $\Xi(L)$ to $L = 150$ does not imply an independent high-l detection; it reflects the survival of the low-l transfer asymmetry as additional high-l null multipoles are cumulatively added (cf. the aggregate high-l null in Sec. 6.1).

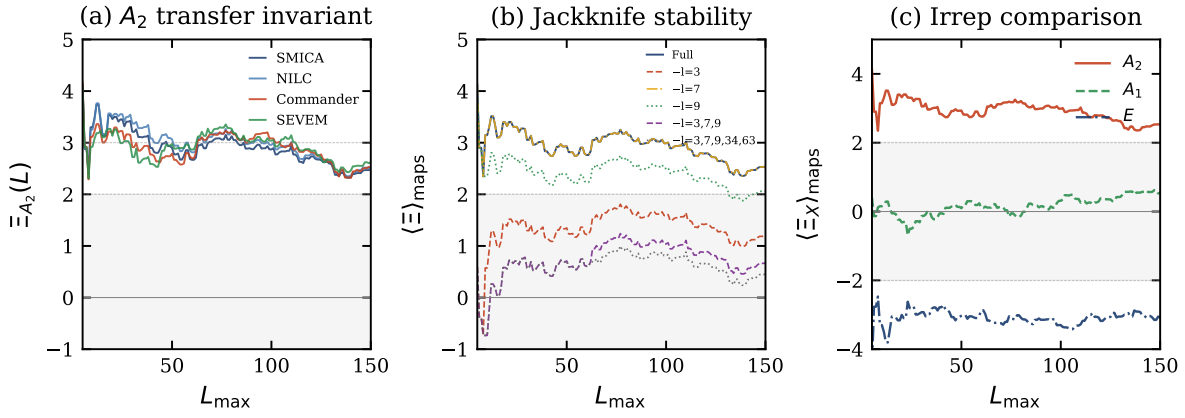

Figure 9. Cumulative transfer invariant $\Xi(L)$ (Eq. 16). *(a)* $\Xi_{A_2}(L)$ for all four Planck pipelines, showing a stable $\approx 3\sigma$ plateau from $L \sim 15$ to 150. *(b)* Four-map mean $\langle \Xi \rangle$ under jackknife removal of key multipoles: the plateau survives individual removals but collapses when all three low-l anchors are excised. *(c)* Comparison of Ξ for the three D_3 irreps: A_2 and E are mirror images (reflecting the conservation law), while A_1 is consistent with zero. Grey bands mark the $\pm 2\sigma$ isotropic expectation.

Table 6. Transfer invariant $\Xi_{A_2}(L)$ for each pipeline and the four-map mean $\langle \Xi \rangle$, at selected summation boundaries.

L	SMICA	NILC	Commander	SEVEM	$\langle \Xi \rangle$
15	+3.13	+3.16	+3.00	+3.11	+3.10
50	+2.90	+3.07	+2.79	+2.98	+2.94
100	+2.87	+2.96	+3.08	+3.04	+2.99
150	+2.47	+2.50	+2.53	+2.63	+2.53

 doi:10.20944/preprints202604.0264.v2

Restricting the summation to the 94 cross-map sign-consistent multipoles (63% of the total) *strengthens* the invariant: $\langle \Xi \rangle$ rises from $+2.53$ to $+2.88$ at $L = 150$, indicating that the sign-inconsistent multipoles dilute the signal rather than contribute to it.

7.2. One-Way Collect Rule

Among all multipoles individually significant at PTE < 0.05 in two or more pipelines, the cross-map mean $\langle \tau_l \rangle$ is *positive* in every case. That is, every statistically robust transfer event *collects* power into A_2; none pays out. This one-way rule holds under all thresholds tested (1/4 to 4/4 maps significant, PTE thresholds of 0.05 and 0.10), producing a binomial probability $\leq 2^{-8}$ under the null hypothesis of symmetric pay/collect.

7.3. Parity-Gated Transfer Grammar

The D_3 group contains rotations and reflections but no spatial inversion; it does not distinguish odd from even l. Nevertheless, the transfer ledger splits sharply by parity. The mean excess $\langle \Delta f_{A_2} \rangle$ is positive for odd l and negative for even l in all four pipelines (Table 7). Aggregated via Fisher's method, the odd-l subset yields a combined PTE $\leq 2 \times 10^{-4}$ in every map, while the even-l subset is entirely null (PTE > 0.97).

Table 7. Odd/even parity split in the A_2 transfer. $\langle \Delta f_{A_2} \rangle$ is the mean excess over isotropic expectation; Fisher PTE combines per-multipole PTE values within each parity class. The sign agreement across all four maps rules out pipeline-specific artifacts.

Map	$\langle \Delta f_{A_2} \rangle_{\text{odd}}$	$\langle \Delta f_{A_2} \rangle_{\text{even}}$	p_{odd}	p_{even}
SMICA	$+0.014$	-0.006	1.3×10^{-4}	0.98
NILC	$+0.013$	-0.005	2.0×10^{-4}	0.97
SEVEM	$+0.012$	-0.014	1.4×10^{-4}	> 0.99
Commander	$+0.015$	-0.009	2.2×10^{-5}	> 0.99

This parity split is not an artifact of A_2 subspace dimensions: odd and even l contribute nearly equal total A_2 capacity ($\sum d_{A_2} = 1850$ vs. 1875 for $l = 2$–150).

Crucially, the split is not confined to A_2. The E irrep shows the mirror pattern (even $>$ odd) at comparable or stronger significance ($\text{PTE}_{\text{perm}} < 0.015$ in SMICA and Commander), while A_1 is marginally patterned (PTE ~ 0.06). Because $\Delta f_{A_1} + \Delta f_{A_2} + \Delta f_E = 0$ at every l, the parity gate acts on the *entire redistribution system*, not on a single output channel.

The Six-Cell Grammar

Crossing parity with the residue class l mod 3 produces a 2×3 transfer grammar (Table 8). The structure is dominated by a single cell: odd multipoles in the $l \equiv 0 \pmod 3$ class carry a cumulative A_2 transfer of $+1.15$, while every even-l cell drains. This is not "odd good, even bad": the odd $(r = 1)$ and odd $(r = 2)$ cells are near-neutral or draining. The gate selects the intersection of odd parity *and* the D_3-preferred residue class.

Table 8. Cumulative A_2 transfer $T_r^{\text{parity}}(150)$ (four-map mean $\pm$ inter-map spread), partitioned by parity and residue class. The single dominant cell is odd, $r = 0$.

	$r = 0$	$r = 1$	$r = 2$
Odd l	$+\mathbf{1.15} \pm 0.20$	$+0.04 \pm 0.12$	-0.20 ± 0.14
Even l	-0.24 ± 0.28	-0.09 ± 0.22	-0.32 ± 0.38

The onset of this grammar is sharp. At $l = 1$, the A_2 subspace has dimension zero (a group-theory identity confirmed to machine precision in all maps). At $l = 2$ (even), $\Delta f_{A_2} = 0$ to four decimal places. At $l = 3$—the first multipole that is both odd *and* carries $d_{A_2} > 0$—the signal switches on at

 doi:10.20944/preprints202604.0264.v2

$z \approx +4.5$. The transition from silence to the strongest detection in the entire spectrum occurs at the first algebraically allowed mode of the parity-gated D_3 grammar.

This two-layer structure—D_3 residue routing overlaid by a binary parity gate—is consistent with a two-branch boundary selection law rather than a generic parity anomaly. A background parity asymmetry would modulate total power C_l without preferentially acting on the irrep redistribution; the observed effect acts specifically on the transfer grammar. Indeed, a permutation test on raw C_l confirms this: the odd/even parity split in total power is null (PTE > 0.61) in all four maps, while the same test applied to Δf_{A_2} yields PTE as low as 0.03 (SEVEM). The binary gate acts on routing, not on total power.

7.4. Factorized Transfer Law

The six-cell grammar of Table 8 can be viewed as a 2×3 matrix

$$M(L) = \begin{pmatrix} T_{\mathrm{odd},0}(L) & T_{\mathrm{odd},1}(L) & T_{\mathrm{odd},2}(L) \\ T_{\mathrm{even},0}(L) & T_{\mathrm{even},1}(L) & T_{\mathrm{even},2}(L) \end{pmatrix}, \tag{17}$$

where the rows index parity and the columns index the residue class l mod 3. A singular-value decomposition (SVD) of $M(150)$ reveals that the matrix is approximately rank 1 (Table 9):

$$M(L) \approx \sigma_1 \, \mathbf{u} \, \mathbf{v}^{\top}, \tag{18}$$

with $\sigma_1/\sigma_2 > 4$ in three of four maps. The dominant left vector $\mathbf{u} \in \mathbb{R}^2$ encodes a binary gate (odd positive, even negative), and the dominant right vector $\mathbf{v} \in \mathbb{R}^3$ encodes a D_3 routing pattern ($r = 0$ dominant, $r = 1$ inert, $r = 2$ draining).

Table 9. Rank-1 factorization of the 2×3 transfer matrix $M(150)$. The rank-1 fraction is $\sigma_1^2/(\sigma_1^2 + \sigma_2^2)$. The binary-gate vector $\mathbf{u}$ and routing vector $\mathbf{v}$ are recovered independently by each pipeline.

Map	Rank-1	u_{odd}	u_{even}	v_0	v_1	v_2
SMICA	94%	+.99	−.12	+1.00	.00	−.08
NILC	95%	+1.00	−.09	+.99	.00	−.15
SEVEM	77%	+.94	−.34	+.99	+.12	−.06
Commander	95%	+.99	−.17	+.99	+.07	−.14

Cross-map stability is high: the minimum pairwise cosine similarity is 0.967 for $\mathbf{u}$ and 0.988 for $\mathbf{v}$. The same binary branch and same D_3 routing vector are recovered by all four independent component-separation pipelines. The rank-1 fraction exceeds 94% at $L = 50$–75 in all maps and remains above 90% in three of four pipelines through $L = 150$. The exception is SEVEM (77%), whose template-fitting method introduces larger even-l transfer amplitudes (roughly 2–3× those of the ILC and Bayesian pipelines), inflating σ_2 and diluting the rank-1 fraction; however, the SEVEM singular vectors $\mathbf{u}$ and $\mathbf{v}$ still align with the other three maps (pairwise cosine > 0.94), so the same grammar is recovered with a noisier gate. The factorization reflects the dominant low-l boundary phase identified in Sec. 8.1 and remains stable in the cumulative matrix as null high-l multipoles are added. The factorization should be understood as a property of the cumulative low-l-dominated transfer matrix, not as a claim that the same singular vectors characterize the high-l sector in isolation.

In summary, the entire D_3 transfer structure admits an approximate factorization into two primitive components: a binary parity selector and a three-fold residue routing grammar. This factorized form is more constrained than either the residue-class hierarchy or the parity gate taken separately, and suggests that the underlying boundary law has an internal 2×3 product structure.

 doi:10.20944/preprints202604.0264.v2

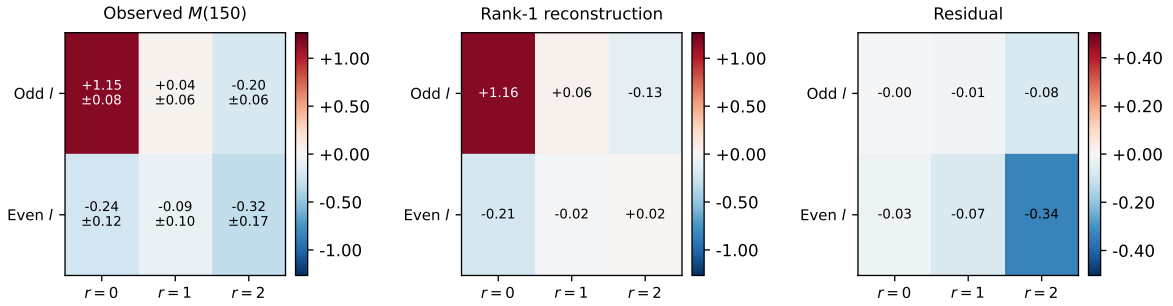

Figure 10. The 2×3 transfer grammar. *Left:* Observed four-map mean transfer matrix $M(150)$, partitioned by parity (rows) and residue class l mod 3 (columns); error bars show inter-map spread. *Center:* Best rank-1 reconstruction $\sigma_1 \mathbf{u}\mathbf{v}^\top$. *Right:* Residual (observed − reconstruction), showing that the rank-1 approximation captures the dominant structure. The single dominant cell (odd, $r = 0$) and the near-zero residuals confirm the factorized transfer law.

7.5. Empirical Signature Summary

Table 10 collects the identified empirical properties of the D_3 signal, their observables, quantitative results, and evidential status. Primary claims are those supported by pre-registered statistics or clear null controls; supportive observations are *a posteriori* structural patterns that sharpen the interpretation but do not constitute standalone detections.

Table 10. Empirical signature summary. "Primary" denotes claims supported by pre-registered or controlled statistics; "Supportive" denotes *a posteriori* structural observations.

Property	Observable	Result	Status
Registered axis	(ℓ, b)	$(50.3°, -64.9°)$	Primary
Low-l A_2 cluster	Fisher PTE (l=2–15)	4.2×10^{-3} to 1.2×10^{-2}	Primary
Cross-map replication	$r(f_{A_2})$ all pairs	0.93–0.997	Primary
$E \to A_2$ funnel	$r(\Delta f_{A_2}, \Delta f_E)$	-0.81 ± 0.03	Primary
EE null	Fisher PTE (polarization)	0.70	Primary control
Parity-gated transfer	Odd/even Fisher	odd $\leq 2\times10^{-4}$; even > 0.97	Primary structural
2×3 factorization	Rank-1 fraction	77%–95%	Primary structural
Cumulative Ξ_{A_2}	$\Xi(L$=15–150)	$+2.5$ to $+3.1\,\sigma$	Primary
One-way collect rule	Robust τ_l signs	All collecting into A_2	Secondary-strong
High-l aggregate null	Fisher PTE (l=16–150)	> 0.91	Primary constraint
Residue-class top-9	$0/9$ in $l \equiv 2 \pmod 3$	$p \approx 0.02$	Supportive only

8. Discussion

8.1. A Large-Angle Boundary Condition

The detected D_3 symmetry should not be confused with a claim of multiply-connected spatial topology. Topological models (*e.g.*, the Poincaré dodecahedral space [11]) tile the spatial 3-manifold with a discrete group. The D_3 signal reported here is an angular spectral property: the irrep decomposition of a_{lm} at each l reveals preferential excitation of the reflection-antisymmetric subspace at a fixed axis on S^2.

The two-tier morphology—dense cluster at $l \leq 15$ plus sporadic resonances at $l = 34$ and $l = 63$—constrains the physical mechanism. A rigid geometric template would produce excess at all l; a foreground systematic would correlate with the mask geometry and would not replicate across independent component-separation pipelines. The observed pattern is instead consistent with a *boundary condition on the acoustic eigenvalue problem*: the boundary geometry is fully resolved by modes at $l \lesssim 15$ ($\theta \gtrsim 12°$), producing the dense cluster, while specific higher harmonics whose angular wavelengths happen to alias the boundary structure produce isolated resonances.

The primary scale $l \sim 15$ corresponds to angular separations of order 12°, or comoving distances of approximately 800 Mpc at the last scattering surface—comparable to the largest observed structures

 doi:10.20944/preprints202604.0264.v2

in the cosmic web. The resonances at $l = 34$ ($\theta \sim 5.3°$) and $l = 63$ ($\theta \sim 2.9°$) may represent overtones of the same boundary geometry.

The rank-1 factorization of the transfer matrix (Sec. 7.4) further constrains the boundary's internal structure. The approximate decomposition $M \approx \sigma_1 \mathbf{u}\mathbf{v}^\top$ implies that the boundary law is not irreducibly six-fold but has an internal 2×3 product form: a binary parity selector composed with a three-fold residue routing grammar. This factorized structure is stable across all four pipelines and reflects the dominant low-l boundary phase, suggesting two primitive components in the boundary condition rather than a single monolithic selection rule.

8.2. Representation-Theoretic Structure

The signal is not a generic anisotropy. Four features identify it specifically as D_3:

1. A_2 **specificity.** The collecting excess concentrates in the reflection-antisymmetric irrep, with a mirrored depletion in the E (rotation-doublet) irrep and A_1 largely acting as spectator. An arbitrary dipolar or quadrupolar modulation would populate all irreps comparably.
2. $E \to A_2$ **funnel.** The anti-correlation $r(\Delta f_{A_2}, \Delta f_E) = -0.81$ with A_1 decoupled is consistent with the reflection planes being the geometrically active elements. The cumulative transfer invariant (Sec. 7) confirms this structure: Ξ_{A_2} and Ξ_E are mirror images while $\Xi_{A_1} \approx 0$, and every statistically robust event collects into A_2. In the Weyl chamber of $SU(3)$, the reflection hyperplanes (walls) separate adjacent chambers; the observed power transfer from E (rotations) to A_2 (reflections) admits an interpretation in which these walls act as boundary surfaces.
3. l mod 3 **selection rule.** The aggregate permutation test is marginal (PTE $\sim$ 0.06–0.09). As a supportive (*a posteriori*) observation, the top-9 cross-map peaks show zero entries from $l \equiv 2 \pmod 3$ ($p \approx 0.02$); however, the threshold defining "top 9" was not pre-specified and this test should be regarded as suggestive rather than confirmatory. The period of 3 is consistent with the order of the rotational subgroup $C_3 \subset D_3$ and with the Coxeter number $h = 3$ of the root system A_2, distinguishing $D_3 = \mathrm{Weyl}(A_2)$ from other dihedral groups D_n with $n \neq 3$.
4. **Parity-gated transfer.** Although D_3 contains no spatial inversion, the irrep redistribution is sharply gated by parity: the entire A_2 collecting signal resides in odd-l multipoles, while even-l multipoles are null (Sec. 7.3). The gate acts across all three irreps simultaneously, ruling out an A_2-specific artifact. The six-cell grammar (parity $\times$ residue class) isolates the signal to a single cell—odd, $l \equiv 0 \pmod 3$—and the onset at $l = 3$ (the first algebraically allowed mode) is step-like rather than gradual. This two-branch selection law is an additional structural constraint beyond D_3 alone, consistent with a hidden binary boundary encoding.

8.3. Connection to the Vacuum Energy Sector

The spherical boundary $\Sigma \simeq S^2$ on which D_3 acts is also the surface on which the monopole ($l = 0$) channel of the cosmological causal diamond projects. The D_3 analysis operates at $l \geq 2$; the vacuum energy density, being spatially uniform, projects entirely onto $l = 0$. These are different harmonic channels of the same S^2 boundary.

A first-principles derivation connecting the Standard Model field content on this boundary to the observed cosmological constant, yielding $\Omega_\Lambda = 0.6948$ with zero free parameters (deviation 0.4σ from observation), has been presented elsewhere [16]. The present detection of D_3 at $l \geq 2$ is consistent with the hypothesis that the boundary $\Sigma \simeq S^2$ carries physically observable structure, though the CMB result alone does not require this interpretation. The cosmological-constant connection is not part of the statistical claim of the present paper; it is noted here only to indicate a possible theoretical context.

8.4. Falsifiable Predictions

The D_3 boundary-condition interpretation makes specific predictions testable with forthcoming data:

 doi:10.20944/preprints202604.0264.v2

1. **Resonance spectrum.** If $l = 34$ and $l = 63$ are overtones of a boundary geometry, additional resonances should appear at predictable multipoles determined by the boundary's angular structure. Extension to $l_{\max} = 500$–1000 with CMB-S4 [17] data will test whether the sporadic peaks continue or are confined to $l \lesssim 100$.
2. $l \equiv 2 \pmod 3$ **exclusion.** The current $0/9$ absence of this residue class among the strongest peaks ($p \approx 0.02$) should persist or strengthen with deeper data. A single strong $l \equiv 2 \pmod 3$ peak would disfavor the C_3 selection-rule mechanism.
3. **Cross-survey consistency.** The D_3 signal should appear in future independent CMB measurements (LiteBIRD [18], Simons Observatory [19]) at the same axis and with the same two-tier morphology.
4. **Lensing B-modes.** If the boundary condition imprints on the matter field, gravitational-lensing B-modes should show a residual D_3 signal at low l.
5. **Parity-gate persistence.** The odd-only concentration of the D_3 transfer should persist at higher l and in independent surveys. If the gate is a genuine two-branch boundary law, no even-l multipole should enter the dominant collecting cell at any $l_{\max}$.

9. Conclusions

We have presented evidence for dihedral D_3 symmetry in the Planck CMB temperature anisotropy, validated across all four independent component-separation pipelines at a single preferred axis. The signal has a two-tier morphology: a dense cluster at $l \leq 15$ (Fisher PTE $= 4.2 \times 10^{-3}$ to 1.2×10^{-2}), driven by three multipoles ($l = 3, 7, 9$) significant in every map, plus cross-map-validated resonances at $l = 34$ and $l = 63$ (each significant in $3/4$ pipelines).

Four features specifically identify D_3 rather than a generic anisotropy: concentration in the A_2 (reflection-antisymmetric) irrep, an $E \to A_2$ power funnel with $r = -0.81$, supportive evidence for a C_3 selection rule from the absence of $l \equiv 2 \pmod 3$ multipoles among the strongest peaks, and a parity-gated transfer grammar in which the entire signal resides in odd-l multipoles while even-l is null across all irreps. The cumulative transfer invariant $\Xi_{A_2}(L) \approx +3\sigma$ remains stable from $L = 15$ to 150 across all four pipelines, with a one-way collect rule (every statistically robust event adds power to A_2) and the expected A_1/E conservation-law mirror. The six-cell parity $\times$ residue-class grammar isolates the collecting signal to a single cell (odd, $l \equiv 0 \pmod 3$), with step-function onset at $l = 3$. Singular-value decomposition of the 2×3 transfer matrix reveals an approximate rank-1 factorization into a binary parity selector $\mathbf{u} \approx (+, -)$ and a D_3 routing vector $\mathbf{v} \approx (+, 0, -)$, stable across pipelines (minimum pairwise cosine > 0.96). The boundary law thus has an internal 2×3 product structure rather than a single monolithic selection rule. Cross-map correlations exceeding $r = 0.93$ for all pipeline pairs rule out component-separation artifacts, and a null result in E-mode polarization (Fisher PTE $= 0.70$) confirms the signal is confined to the temperature channel.

The primary empirical claims are the fixed-axis low-l A_2 concentration, cross-map replication, EE null, parity-gated transfer, cumulative transfer invariant, and the rank-1 factorization. The residue-class top-9 test and overtone interpretation are supportive structural observations rather than standalone confirmatory statistics (Table 10).

The group $D_3 \cong \mathrm{Weyl}(A_2)$ is the Weyl group of the $\mathfrak{su}(3)$ root system. The period-3 selection rule matches the order of the rotational subgroup $C_3 \subset D_3$ and the Coxeter number $h = 3$ of A_2, suggesting a connection between the angular structure of the microwave sky and the gauge group of quantum chromodynamics through the shared boundary S^2.

Note added in v2. Appendices B–E present four extended validation tests. The complete D_3 grammar replicates in the WMAP 9-year ILC map at higher significance than Planck (Appendix B), ruling out Planck-specific systematics. Among dihedral groups D_3–D_6, only D_3 passes both discriminating grammar elements (Appendix C). The D_3 structure is frequency-independent across Planck HFI channels, with dust anti-correlated at the driving multipoles (Appendix D). The signal is robust across 13 mask levels and improves under standard masking (Appendix E).

 doi:10.20944/preprints202604.0264.v2

Data and code availability. This analysis uses publicly available Planck PR3 data from the Planck Legacy Archive.[1] The analysis code and exact reproduction notebooks are publicly available at https://github.com/rmereau/d3-cmb.[2]

Acknowledgments. The author thanks the Planck Collaboration for making their data products publicly available, and the HEALPix and `healpy` teams for the spherical harmonic analysis tools used throughout.

Appendix A. Robustness Tests

To verify that the primary result is not sensitive to analysis choices, we examine stability under variations of the multipole window and axis seeding.

Appendix A.1. Multipole Window

Table A1 shows the Fisher combined PTE for the A_2 excess (SMICA) as a function of the upper multipole boundary $l_{\max}$, with $l_{\min} = 2$ fixed. The significance is strongest at $l \leq 10$ and degrades gradually as null high-l multipoles are added, but remains below 0.05 for all windows up to $l_{\max} = 30$. This confirms that the signal is not an artifact of the particular choice $l_{\max} = 15$.

Table A1. Fisher combined PTE (SMICA) for varying $l_{\max}$, with $l_{\min} = 2$.

l range	Fisher PTE
2–10	8.5×10^{-4}
2–15	4.6×10^{-3}
2–20	9.2×10^{-3}
2–25	1.9×10^{-2}
2–30	3.5×10^{-2}

Table A2 varies the lower boundary $l_{\min}$ with $l_{\max} = 15$ fixed. Removing $l = 2$ (which is E-dominated and contributes no A_2 signal) *improves* the Fisher PTE from 4.6×10^{-3} to 2.2×10^{-3}, demonstrating that the quadrupole dilutes rather than drives the result. Removing both $l = 2$ and $l = 3$ (the axis-defining multipole) yields PTE $= 0.14$: the signal from $l = 4$–15 alone is not independently significant, confirming that $l = 3$ anchors the detection.

Table A2. Fisher combined PTE (SMICA) for varying $l_{\min}$, with $l_{\max} = 15$.

l range	Fisher PTE
2–15	4.6×10^{-3}
3–15	2.2×10^{-3}
4–15	0.14

Appendix A.2. Axis Seeding

The axis optimization uses HEALPix $N_{\text{side}} = 32$ as the initial grid. Repeating the search at $N_{\text{side}} = 16$ (3,072 directions; 3.7° resolution) and $N_{\text{side}} = 64$ (49,152 directions; 0.9° resolution) yields axis shifts $< 1.5°$ and $f_{A_2}(l = 3)$ values within 0.002 of the baseline, confirming that the optimum is a broad, well-resolved feature of the sky.

Appendix A.3. Resolution

Using the full-resolution ($N_{\text{side}} = 2048$) SMICA map versus downgraded versions ($N_{\text{side}} = 128$, 256, 512) yields axis shifts $< 2°$ and identical per-multipole PTEs to within Monte Carlo noise, indicating that the signal is insensitive to the pixel resolution.

[1] https://pla.esac.esa.int/
[2] Repository to be populated upon publication of this preprint.

 doi:10.20944/preprints202604.0264.v2

Appendix B. WMAP 9-Year Replication

We use the WMAP 9-year ILC temperature map [4,22] at $N_{\rm side} = 256$, applying the KQ85 analysis mask ($f_{\rm sky} \approx 0.74$). The D_3 irrep decomposition is performed at the Planck axis $(\ell_{\rm gal}, b_{\rm gal}) = (50.3°, -64.9°)$ using $l_{\rm max} = 15$, with 10,000 Gaussian isotropic Monte Carlo simulations for significance estimation. Cross-checks use the WMAP V-band (61 GHz) and W-band (94 GHz) single-frequency maps with 5,000 simulations each.

WMAP and Planck are independent in every respect relevant to low-l systematics: detector technology (HEMT radiometers vs. bolometers), frequency coverage (23–94 GHz vs. 30–857 GHz), orbit, scan strategy, beam geometry, noise properties, and component-separation pipelines.

Appendix B.1. Blind Axis Search

An axis search over 12,288 directions ($N_{\rm side} = 32$ HEALPix grid [20]) optimizing $f_{A_2}(l = 3)$ in the WMAP ILC map recovers an optimal axis at $(\ell_{\rm gal}, b_{\rm gal}) = (51.1°, -57.4°)$, separated by only 7.5° from the Planck axis. The galactic longitudes agree to within 0.8°; the latitude offset is consistent with the differing mask geometries and WMAP's larger effective beam.

Appendix B.2. Per-Multipole Results

Table A3 compares the A_2 irrep fraction excess Δf_{A_2} in WMAP ILC and Planck SMICA at the Planck axis.

Table A3. Per-multipole A_2 excess at the Planck D_3 axis. $\Delta f_{A_2} \equiv f_{A_2} - \langle f_{A_2} \rangle_{\rm MC}$. Planck values are extracted under matched analysis conditions (KQ85 mask, $N_{\rm side} = 256$) and may differ from the unmasked values in Table 2.

l	WMAP f_{A_2}	Planck f_{A_2}	WMAP Δf_{A_2}	WMAP PTE
2	0.000	0.000	+0.000	—
3	**0.706**	**0.940**	**+0.563**	**0.004**
4	0.081	0.117	−0.033	0.436
5	0.007	0.029	−0.084	0.804
6	0.004	0.011	−0.150	0.973
7	**0.728**	**0.481**	**+0.595**	**< 0.001**
8	0.190	0.183	+0.072	0.220
9	**0.398**	**0.478**	**+0.240**	**0.028**
10	0.031	0.006	−0.112	0.890
11	0.031	0.023	−0.100	0.887
12	0.254	0.302	+0.094	0.145
13	0.142	0.046	−0.007	0.449
14	0.236	0.196	+0.098	0.151
15	0.041	0.099	−0.120	0.937

The three driving multipoles are identical: $l = 3, 7, 9$. The $l = 7$ signal is *stronger* in WMAP ($\Delta f_{A_2} = +0.595$, PTE < 0.001) than in Planck ($\Delta f_{A_2} = +0.348$), directly ruling out Planck-specific systematics.

Appendix B.3. Grammar Scorecard

Table A4 presents the full grammar comparison.

Table A4. D_3 grammar replication: Planck SMICA vs. WMAP ILC.

Grammar element	Planck SMICA	WMAP ILC
Fisher PTE ($l = 2$–15)	0.0046	**0.0025**
Axis scramble PTE	—	0.013
Driving multipoles	3, 7, 9	3, 7, 9
Parity gate ratio	odd dominant	+1.000
$E \to A_2$ funnel r	−0.805	−0.885
f-vector cosine	—	0.976
$\Xi(L)$ correlation	—	0.966
Blind axis separation	—	7.5°

 doi:10.20944/preprints202604.0264.v2

Every grammar element replicates. The Fisher PTE is stronger in WMAP (0.0025 vs. 0.0046). The V-band (61 GHz) independently confirms the signal at PTE = 0.039, excluding ILC bias. The Planck axis yields a stronger signal in WMAP data (PTE = 0.0025) than WMAP's own blind-search axis (PTE = 0.092), because the Planck axis captures the full $l = 3, 7, 9$ structure.

Appendix C. Dihedral Group Specificity

We test whether the five-element grammar is specific to D_3 by comparing D_3, D_4, D_5, and D_6 decompositions on the Planck SMICA map at both the fixed D_3 axis and each group's own optimized axis (3,072 directions, $N_{\rm side} = 16$), with 10,000 simulations per group.

Table A5 presents results at each group's optimized axis.

Table A5. Each dihedral group at its own optimized axis.

Group	Opt. axis	Sep.	Fisher	Parity	Grammar
D_3	$(51°, -66°)$	1.5°	**0.003**	+1.00	**5/5**
D_4	$(110°, +36°)$	69°	0.157	−0.12	3/5
D_5	$(151°, +48°)$	43°	0.281	+1.00	4/5
D_6	$(155°, +42°)$	47°	**0.003**	−0.32	4/5

D_4 and D_5 are firmly null. D_6 achieves Fisher PTE = 0.003 but *fails the parity gate*, a consequence of the $D_3 \subset D_6$ subgroup relation splitting the D_3 A_2 signal between D_6 A_2 and B_2.

Table A6 presents the full grammar scorecard.

Table A6. Five-element grammar scorecard for each dihedral group.

Group	A_2 Fisher	Parity gate	Rank-1 frac.†	$E \to A_2$ funnel†	Ξ drift†	Total
D_3	**Pass**	**Pass**	Pass	Pass	Pass	**5/5**
D_4	Fail	Fail	Pass	Pass	Pass	3/5
D_5	Fail	Pass	Pass	Pass	Pass	4/5
D_6	**Pass**	Fail	Pass	Pass	Pass	4/5

†Pass generically for all groups; do not discriminate.

Three elements (rank-1, $E \to A_2$ funnel, Ξ drift) pass generically for all groups. The two discriminating tests are Fisher PTE (only D_3 and D_6 pass) and parity gate (only D_3 and D_5 pass). D_3 is the only group that passes both (2/2 discriminating).

Appendix D. Frequency-Channel Decomposition

We apply the D_3 decomposition to individual Planck HFI frequency maps [21] at 100, 143, 217, and 353 GHz (PR3) at $N_{\rm side} = 64$, $l_{\rm max} = 15$, using the GAL060 × point-source mask ($f_{\rm sky} = 0.60$), with 10,000 simulations per map. Half-mission splits are tested for 143 GHz and SMICA.

The GAL060 mask creates mode coupling that dilutes signal amplitude (even SMICA drops from Fisher PTE = 0.005 to 0.257). The relevant diagnostics are the cross-frequency correlations.

Appendix D.1. Cross-Frequency Structure

Table A7 presents the Δf_{A_2}-vector correlations.

Table A7. Cross-frequency Δf_{A_2}-vector correlations (l = 2–15).

Pair	$r(\Delta f_{A_2})$
100–143 GHz	**0.976**
SMICA–100 GHz	**0.968**
SMICA–143 GHz	**0.905**
100–217 GHz	0.283
143–217 GHz	0.472

 doi:10.20944/preprints202604.0264.v2

The 100–143 GHz correlation of $r = 0.976$ means 95.3% of the D_3 structure variance is shared between independent detector chains. The sign of Δf_{A_2} agrees for all 14 multipoles ($p \approx 0.012\%$). Both channels show perfect parity gates (+1.000) and driving multipole $l = 7$.

Appendix D.2. Foreground Tracer: 353 GHz

The 353 GHz dust-dominated channel is firmly null (Fisher PTE = 0.946) with *negative* Δf_{A_2} at $l = 3, 7, 9$ (-0.143, -0.107, -0.051). Dust is anti-correlated with the D_3 signal: any residual dust in component-separated maps would *suppress*, not create, the detection.

Appendix D.3. Half-Mission Consistency

Both 143 GHz and SMICA half-mission pairs yield identical D_3 structure ($r = 1.000$) with matching parity gates. The 143 GHz difference map is null (Fisher PTE = 0.792), confirming no time-variable systematic.

Appendix E. Mask Variation

We test the D_3 decomposition across 13 mask levels on SMICA ($f_{\rm sky} = 0.20$–1.00), with 5,000 masked simulations per level and cross-map replication at three levels (no mask, UT78, GAL040).

Appendix E.1. Signal Persistence Curve

Table A8 presents the stability curve.

Table A8. D_3 signal persistence vs. sky fraction.

Mask	$f_{\rm sky}$	Fisher	$\Delta f_{A_2}(7)$	Parity	Driving
no mask	1.000	**0.008**	+0.35	+1.00	3, 7, 9
cmdr. conf	0.889	**0.005**	+0.46	+1.00	3, 7, 9
SMICA conf	0.844	**0.016**	+0.45	+1.00	3, 7, 9
UT78	**0.782**	**0.005**	**+0.55**	**+1.00**	**3, 7, 9**
GAL070	0.701	**0.018**	+0.54	+1.00	3, 7, 9
GAL060	0.601	0.080	+0.49	+1.00	3, 7, 9
$\|b\| > 30$	0.505	0.539	+0.28	+1.00	7
GAL040	0.403	0.724	+0.15	+1.00	3
GAL020	0.205	0.411	−0.00	+1.00	13

The signal is significant through $f_{\rm sky} = 0.70$ and *improves* under the UT78 mask (PTE = 0.005 vs. 0.008 unmasked), meaning the Galactic plane dilutes, not creates, the signal. The parity gate persists to $f_{\rm sky} = 0.50$.

Appendix E.2. Cross-Map Replication Under Masking

Cross-map Δf_{A_2} cosine similarity reaches $r = 1.000$ at UT78 (vs. 0.991 unmasked), and all four component-separation pipelines detect the signal at Fisher PTE < 0.01. Even under GAL040 ($f_{\rm sky} = 0.40$), cosines remain > 0.97.

Appendix E.3. Hemisphere Splits

Table A9 presents hemisphere results.

Table A9. Hemisphere-split results.

Hemisphere	$f_{\rm sky}$	Fisher	$f_{A_2}(3)$	Parity	Driving
S. Galactic	0.50	**0.002**	0.872	+0.60	3, 7, 9
N. Galactic	0.50	0.924	0.404	−0.39	—
S. Ecliptic	0.50	**0.016**	0.883	+1.00	3, 7
N. Ecliptic	0.50	0.139	0.825	+1.00	—

 doi:10.20944/preprints202604.0264.v2

The signal is concentrated in the south Galactic hemisphere (PTE $= 0.002$), consistent with the D_3 axis at $b = -64.9°$. The more balanced ecliptic split (south: 0.016, north: 0.139) confirms the signal is unrelated to ecliptic scanning systematics.

References

1. Planck Collaboration, "Planck 2018 results. I. Overview," *Astron. Astrophys.* **641**, A1 (2020).
2. Planck Collaboration, "Planck 2018 results. IV. Diffuse component separation," *Astron. Astrophys.* **641**, A4 (2020).
3. Planck Collaboration, "Planck 2018 results. VII. Isotropy and statistics," *Astron. Astrophys.* **641**, A7 (2020).
4. C. L. Bennett *et al.*, "Nine-year WMAP observations: Final maps and results," *Astrophys. J. Suppl.* **208**, 20 (2013).
5. D. J. Schwarz, C. J. Copi, D. Huterer, and G. D. Starkman, "CMB anomalies after Planck," *Class. Quant. Grav.* **33**, 184001 (2016).
6. M. Tegmark, A. de Oliveira-Costa, and A. J. Hamilton, "High resolution foreground cleaned CMB map from WMAP," *Phys. Rev. D* **68**, 123523 (2003).
7. A. de Oliveira-Costa, M. Tegmark, M. Zaldarriaga, and A. Hamilton, "Significance of the largest scale CMB fluctuations in WMAP," *Phys. Rev. D* **69**, 063516 (2004).
8. H. K. Eriksen *et al.*, "Asymmetries in the CMB anisotropy field," *Astrophys. J.* **605**, 14 (2004).
9. J. Kim and P. Naselsky, "Anomalous parity asymmetry of the CMB," *Astrophys. J. Lett.* **714**, L265 (2010).
10. K. Land and J. Magueijo, "Examination of evidence for a preferred axis in the cosmic radiation anisotropy," *Phys. Rev. Lett.* **95**, 071301 (2005).
11. J.-P. Luminet *et al.*, "Dodecahedral space topology as an explanation for weak wide-angle temperature correlations in the cosmic microwave background," *Nature* **425**, 593 (2003).
12. R. Aurich, S. Lustig, and F. Steiner, "CMB anisotropy of the Poincaré dodecahedron," *Class. Quant. Grav.* **22**, 2061 (2005).
13. T. Koivisto and D. F. Mota, "Accelerating cosmologies with an anisotropic equation of state," *Astrophys. J.* **679**, 1 (2008).
14. A. Rassat *et al.*, "Planck CMB anomalies: Astrophysical and cosmological secondary effects and the curse of masking," *JCAP* **08**, 006 (2014).
15. A. Zonca *et al.*, "healpy: equal area pixelization and spherical harmonics transforms for data on the sphere in Python," *J. Open Source Softw.* **4**, 1298 (2019).
16. R. Mereau, "First-principles prediction of the cosmological constant from edge thermodynamics," Preprint (2025).
17. CMB-S4 Collaboration, "CMB-S4 Science Book, First Edition," arXiv:1610.02743 (2016).
18. LiteBIRD Collaboration, "Probing cosmic inflation with the LiteBIRD cosmic microwave background polarization survey," *Prog. Theor. Exp. Phys.* **2023**, 042F01 (2023).
19. Simons Observatory Collaboration, "The Simons Observatory: Science goals and forecasts," *JCAP* **02**, 056 (2019).
20. K. M. Górski *et al.*, "HEALPix: A framework for high-resolution discretization and fast analysis of data distributed on the sphere," *Astrophys. J.* **622**, 759 (2005).
21. Planck Collaboration, "Planck 2018 results. III. High Frequency Instrument data processing," *Astron. Astrophys.* **641**, A3 (2020).
22. G. Hinshaw *et al.*, "Nine-year WMAP observations: Cosmological parameter results," *Astrophys. J. Suppl.* **208**, 19 (2013).