Draft version September 5, 2026
Typeset using LaTeX **twocolumn** style in AASTeX7.0.2

The Largest Test of a Preferred Galaxy-Spin Axis: An 8.47-Million-Galaxy DESI Chirality Catalog, a Void-Environment Contrast, and a Sensitivity Confrontation with the Rotating-Black-Hole-Universe Prediction

Houston Golden[1]

[1] *Independent Researcher, Los Angeles, California, USA*

(Dated: September 5, 2026; Version v4P.0.7)

ABSTRACT

Two literatures motivate a search for a preferred axis in galaxy spin directions: handedness-dipole claims from Longo and from Shamir, and Poplawski's rotating black-hole-universe model, in which Einstein–Cartan torsion halts a collapsing black hole's singularity, bouncing into a daughter universe that inherits a preferred axis from the parent's spin, toward which galaxies are predicted to align. We report the largest test of that claim to date: an 8,474,531-galaxy DESI Legacy DR8 chirality catalog and two independent null tests. The primary, quality-controlled real-space dipole ($N_{\rm support} = 887{,}472$ of 890,069 rows) is null-consistent ($z_{\rm mom} = +0.635$, one-sided $p = 0.238$), with a 95% sensitivity floor $A_{95}^{\rm obs} \simeq 0.98\%$ (Neyman 95% CL limit $A_{95}^{\rm CL} \simeq 0.75\%$, the floor used below). A companion DESIVAST void/non-void contrast on 145,766 classifier-labelled galaxies likewise finds no significant effect ($\Delta f_{\rm CW} = +0.00145$, two-sided $p = 0.66$). We derive the model's spin-axis-dipole implication; the cited mechanism papers give only a qualitative tendency, not a computed amplitude, so under the minimal closure needed to quantify the claim, our floor disfavors a dipole above $\sim 1\%$ at $\geq 95\%$ power, 2–20× below the $\sim$2–33% literature amplitudes, though under an illustrative, not-adopted-for-strengthening bridge the two largest comparison samples drop below this floor. This null is measured on 887,472 primary-channel spirals, 4–3,400× larger than every comparison catalog in Table 5 except Shamir's DESI Legacy sample ($N = 1.3$ million), and confirms independent reanalyses by Iye *et al.* and Patel & Desmond finding no significant galaxy-spin anisotropy.

Keywords: Spiral galaxies (1560); Astrostatistics (1882); Catalogs (205); Large-scale structure of the universe (902); Cosmic voids (1857)

1. INTRODUCTION

Whether spiral galaxies show a statistically preferred handedness on the sky is a two-decade observational question with a specific theoretical stake. Longo (2011) (1) reported a $\sim 7\%$ left-handed excess with a $\sim 5\sigma$ dipole signal in 15,158 SDSS spirals at $z < 0.085$. Shamir and collaborators subsequently reported handedness asymmetries of order 2–20% across several independent samples – SDSS (2; 3), DESI Legacy (4), and most recently a $\sim$2:1 clockwise-to-counterclockwise imbalance among 263 JWST/JADES galaxies at $z \gtrsim 1$ (5). Independent reanalyses have not confirmed these claims: Iye, Yagi & Fukumoto (2021) (6) re-examined Shamir's SDSS catalog with 3D random-walk simulations and found no significant dipole after correcting for reading-direction bias and duplicate photometric objects, and Patel & Desmond (2024) (7) collected every public spin-classification dataset and found consistency with isotropy to within 3σ in a combined Bayesian and frequentist analysis.

The theoretical stake is Poplawski's proposal that our universe originated as the interior of a black hole in a parent universe, where torsion in Einstein–Cartan gravity generates a spin–spin contact repulsion that halts gravitational collapse before a singularity forms and drives a bounce into expansion (8; 9; 10). If the parent black hole is rotating, its spin axis becomes a preferred, non-inertial axis for the daughter universe, and Poplawski argues that galaxies tend to align their spin axes with it over cosmic time, producing exactly the kind of clockwise/counterclockwise sky asymmetry that Longo and Shamir report (11). This is the same torsion mechanism this research program's own bounce papers adopt for gravitational-collapse avoidance (P1A/P1C), which is why testing the model's one directly observable astrophysical claim – the predicted galaxy-spin preferred axis – is of direct interest here.

This paper is the largest single test of that claim to date. We release an 8,474,531-object DESI Legacy DR8 chirality catalog (Sec. 2) and report its primary null result on 887,472 supported-pixel high-confidence spirals (Sec. 3); we add a second, independent probe using the released DESIVAST void catalog to test whether chirality differs by large-scale environment on 145,766 classifier-labelled galaxies (Sec. 4); and we derive, from the cited black-hole-universe literature, what quantitative dipole amplitude the model actually predicts and confront it with our sensitivity floor (Sec. 5). We make no bounce claim beyond that stated test: the catalog does not test, and is not evidence for or against, the matter-bounce cosmology this research program otherwise develops (Sec. 6).

2. CATALOG

2.1. *Selection, schema, and release contract*

The catalog assigns an observed chirality label (CW, CCW, or NOT_SPIRAL) to 8,474,531 DESI Legacy Imaging Survey DR8 (12) galaxy cutouts using a ViT-Small encoder with test-time-augmentation (TTA) equivariance correction, trained on a Galaxy Zoo 1-derived label set and refined against Galaxy Zoo DESI morphology (13). The parent sample is drawn from the Smith42/galaxies DESI Legacy DR8 cutout release (8,474,566 images, 224 × 224 px *grz* cutouts at 0.262″/pixel); the parent-sample selection function inherits Galaxy Zoo DESI's photometric-type cut (REX/DEV/EXP/SER), $r \leq 19.0$, and half-light radius $\geq 3''$, across the three DR8 imaging campaigns (BASS+MzLS at $\delta > +32^\circ$, DECaLS at $\delta < +32^\circ$, and a DES overlap region). The released per-galaxy class counts are $N_{\rm CW}$ = 1,592,107, $N_{\rm CCW}$ = 1,609,053, $N_{\rm NS}$ = 5,273,371; the classified-spiral sub-catalog is $N_{\rm spiral} = N_{\rm CW} + N_{\rm CCW}$ = 3,201,160. The catalog's high-confidence (HC) selection ($p_{\rm eq} > 0.6$) yields $N_{\rm HC}$ = 949,584 spirals, corresponding to an integrated chirality purity of ∼70% and a completeness of ∼30% relative to the full spiral sub-catalog, both against the GZ1 cross-match; a stricter $p_{\rm eq} > 0.9$ cut roughly halves completeness to ∼15% at higher purity (Table 2). We exclude 59,515 HC rows flagged `raw_flip_qc_unsafe` (a release-safety quality-control flag identifying a raw/equivariant pipeline-pass mismatch, not float32 rounding), leaving 890,069 quality-controlled rows, of which 887,472 occupy one of the 23,633 supported HEALPix (`NSIDE` = 64) pixels with $N_{\rm spiral}(p) \geq 10$ (the HC-RI support, $f_{\rm sky}$ = 0.4808) to enter the supported-pixel dipole fit and its matched null; the wider full-spiral FSC support ($N_{\rm spiral}(p) \geq 10$ on all classified spirals, not only HC) covers 24,087 pixels, $f_{\rm sky}$ = 0.49005, and underlies Fig. 1 and the harmonic diagnostic in Sec. 3. The catalog, its quarantine, exact null arrays, and full provenance register are archived under Zenodo DOI 10.5281/zenodo.21461899 (17) and are the single primary contract this paper's tests draw on, with the per-row schema given in Table 1. Full architecture, per-region calibration, and the complete systematics appendix are described in the archived release paper (15) and are not reproduced here beyond what Sec. 5 needs.

Table 1. Catalog Parquet column schema (catalog release v1.0.244, documented in the v1.0.274 archived release paper (15)). One row per DESI Legacy DR8 galaxy cutout.

Column	Description
`object_id`	DR8 object identifier (string)
`ra_deg`, `dec_deg`	Sky position (deg, double)
`class_eq`	Equivariant-TTA label: CW/CCW/NOT_SPIRAL
`score_cw_eq`	Equivariant CW class score
`score_ccw_eq`	Equivariant CCW class score
`score_ns_eq`	Equivariant NOT_SPIRAL class score
`score_eq_max`	max eq. class score ($p_{\rm eq}$)
`is_spiral`	CW or CCW (boolean)
`primary_hc`	HC flag, $p_{\rm eq} > 0.6$ (boolean)
`raw_flip_qc_unsafe`	Release-safety QC flag: raw/eq pipeline-pass mismatch (boolean)

Table 2. Completeness and purity vs. threshold $p_{\rm eq}$, against the GZ1 cross-match (archived release (15)). Completeness is relative to the full $N_{\rm spiral}$ = 3,201,160 classified-spiral sub-catalog. $p_{\rm eq} > 0.6$ is this paper's HC selection threshold.

Threshold	$N_{\rm HC}$	Completeness	Purity
$p_{\rm eq} > 0.6$ (adopted)	949,584	∼30%	∼70%
$p_{\rm eq} > 0.9$	lower	∼15%	higher†

†Not separately quoted in the archived release.

Completeness and purity against the GZ1 cross-match are functions of the high-confidence threshold $p_{\rm eq}$ (Table 2), not single scalars: the release adopts $p_{\rm eq} > 0.6$ as the primary cut (balancing sample size against label reliability), and reports the $p_{\rm eq} > 0.9$ point as a higher-purity, lower-completeness alternative not otherwise used in this paper's tests.

2.2. *Classifier quality and label provenance*

The *released* classifier – the one that produced every label in this catalog – achieves 69.91% chirality agreement (Cohen's κ = 0.40) against a full Galaxy Zoo 1 cross-match (N = 117,205). A separate, manifest-retained, from-scratch retrain of the GZ1-core classifier component – which does *not* alter the released Catalog labels – validates at Cohen's $\kappa = 0.97$ on a provably training-disjoint held-out GZ1 sample; the archived release paper is explicit that this retrain figure is "neither comparable to, nor a replacement for" the released classifier's $\kappa = 0.40/69.91\%$ figure (15). A separate, composition-faithful retrain that additionally includes the historical CE-ResNet-derived training component collapses to chance-level chirality performance (best three-class validation accuracy 0.5617, i.e. no better than the 0.106 NOT_SPIRAL-only baseline plus chance on the CW/CCW split; binary CW/CCW accuracy 0.517 on a clean held-out GZ1 sample), so the historical CE-included training accuracy is not reproducible under honest, manifest-retained ingestion; the released catalog labels are unaffected by either retrain. A forward classifier-injection model over 1.7×10^7 banked network passes rules out classifier label confusion as the source of any residual handedness monopole (0.0% of the observed value) and localizes the monopole's origin upstream of the classifier, in the input image distribution – without resolving whether that origin is a true sky asymmetry or a DESI imaging systematic, which remains an open attribution in the archived release (15). The monopole takes different values on different samples, and this paper states all three explicitly rather than one alone: the catalog-wide value across all classified spirals is $f_{\rm CW} = 0.497353$ (monopole -0.265%, per-pixel-independent binomial significance -9.47σ); the HC selection *including* the 59,515 `raw_flip_qc_unsafe` rows gives $f_{\rm CW} = 0.496051$ (monopole -0.395%); and the primary 887,472-galaxy supported-pixel sample used for the dipole fit above – the strict HC selection with those 59,515 rows removed – gives $f_{\rm CW} = 0.5126562$ (454,968/887,472, monopole $f_{\rm CW} - \frac{1}{2} = +1.2656\%$, equivalently $A_p = 2(f_{\rm CW} - \frac{1}{2}) = +2.53\%$ in the amplitude convention of Fig. 1 and Eqs. 1–2). The catalog-wide and HC-with-unsafe monopoles are likewise $A_p = -0.53\%$ and $A_p = -0.79\%$ in this same convention. The quarantined rows themselves are strongly parity-asymmetric (14,776 CW / 44,739 CCW, 75.2% CCW), which is why removing them for the release-safety cut flips the HC monopole's sign; this asymmetry of the `raw_flip_qc_unsafe` flag is itself a previously unstated, checkable property of the released quality-control flag. None of this affects the primary null: the injection-recovery generator computing Eq. 1 draws its baseline CW probability from the same strict 887,472-galaxy sample and support it injects into, so using 0.5126562 is correct by construction; and the uniform-pixel-weight real-space estimator fits $A_p = m + \boldsymbol{a} \cdot \hat{n}_p$, absorbing any constant monopole into m before the dipole amplitude is read off – confirmed generatively, since binomial nulls seeded at the global $f_{\rm CW}$ and at $p = 0.5$ give statistically identical dipole-amplitude nulls (1.957×10^{-3} vs. 1.935×10^{-3}, 0.39σ difference).

Table 3. Three-class confusion matrix, released (equivariant) classifier vs. Galaxy Zoo 1 human labels (1″ cross–match, $N = 240{,}919$; overlap with the 6,637 GZ1 training rows is not object-level excluded, so this evaluation is overlap-contaminated rather than fully independent). Rows: GZ1 label; columns: catalog prediction. Three-class accuracy 58.7%; restricted to the 117,205 GZ1 spirals the classifier labels CW/CCW, chirality agreement is 69.91% (Cohen's $\kappa = 0.40$).

GZ1 \ pred.	CW	CCW	NOT_SPIRAL
CW	39,011	18,889	13,715
CCW	16,377	42,928	13,720
NOT_SPIRAL	17,056	19,724	59,499

The exact historical training realization for the released classifier is not fully recoverable: the production-script reconstruction supports 26,616 labeled rows (6,637 Galaxy Zoo 1 confident CW/CCW, 17,153 CE-ResNet high-confidence spirals, 826 CE-selected NOT_SPIRAL, 2,000 synthetic NOT_SPIRAL), a 21,293/5,323 train/validation split, and 93.6878% validation accuracy, while the committed benchmark report gives 26,626 total rows and 846 CE non-spirals and the committed Hugging Face README reports 92.10%; no retained object-ID manifest or random-state record resolves the conflict or ties the named objects to the historical checkpoint. This does not affect the released labels (fixed historical production outputs) but does mean the released classifier's own training provenance cannot be independently audited beyond the confusion matrix above; the manifest-retained retrain (Sec. 2.1) exists precisely to provide a fully regenerable, audit-clean realization going forward. That retrain is deliberately scoped to the 6,637 Galaxy Zoo 1 confident CW/CCW spirals plus 2,000 synthetic NOT_SPIRAL rows only ($n_{\rm total} = 8{,}637$), reaching 99.31% best-epoch three-class validation accuracy with every object ID, split index, and random state retained in a committed manifest; on 3,000 held-out GZ1-confident spirals provably disjoint from that manifest along both object identity and a 3″ proximity cut, it reaches accuracy 0.9867 and $\kappa = 0.9733$ (confusion $[[1460, 10], [30, 1500]]$ for [CW,CCW]; zero NOT_SPIRAL predictions). A separate composition-faithful retrain that also includes the historical CE-ResNet-derived 17,153-row component (∼67–72% of the historical training composition, by either

Table 4. Bias-hardening test results from the archived release, quoted verbatim (15). Seven of the eight originally numbered tests are tabulated (T5, the former linear-Pearson RA/Dec metadata-leakage row, was *removed* because a linear Pearson correlation is inappropriate for the circular RA coordinate $0^\circ \equiv 360^\circ$; the map-level low-ℓ real-$Y_{\ell m}$ regression in the archived release supersedes it and finds all three $\ell = 1$ coefficients consistent with zero, $|z| \leq 1.25$). All seven tabulated tests pass; these are necessary, not sufficient, checks for sub-percent-level bias control and are not calibrated against a physical sensitivity bound.

Test	Threshold	Result
T1: Flip-swap r	> 0.80	1.000
T2: Rotation stability	$> 80\%$	89.8%
T3: Artifact rejection	$> 70\%$ NOT_SPIRAL	100%
T4: Perturbation robustness	$> 80\%$	84%
T5: removed (see caption)	–	–
T6: Hemispheric diff.	$< 10\%$	3.6%
T7: Calibration proxy	$> 30\%$ @ p>0.9	73.6%
T8: CW/CCW balance	$50 \pm 10\%$	49.7%

committed count) collapses to chance-level chirality performance under honest, manifest-retained ingestion, discussed above.

2.3. *Estimator and null specification*

The primary real-space estimator fits $A_p = m + \boldsymbol{a} \cdot \hat{n}_p$ (`healpy.fit_dipole`, a uniform-pixel-weight full-sky least-squares fit for both a monopole m and a dipole vector $\boldsymbol{a}$) on the 887,472-galaxy, 23,633-pixel HC-RI support after applying the selection predicate `primary_hc && !raw_flip_qc_unsafe`. The matched null preserves the observed per-pixel occupancy exactly: each of 10,000 draws redistributes the observed CW/CCWlabel totals across the fixed pixel capacities of the HC-RI support (a fixed-occupancy label-randomization, equivalent to a multivariate-hypergeometric reshuffle per pixel), so that the null differs from the data only in which galaxies within each pixel are labeled CW vs. CCW, not in how many galaxies occupy each pixel. The observed statistic and each null draw are reduced to the same fitted dipole amplitude A_p; $z_{\rm mom}$ is the observed value's offset from the null mean in units of the null's population standard deviation, and the add-one rank p is the upper-tail empirical rank of the observed value within the 10,001-value pool (data plus null draws). The distinct FSC harmonic diagnostic (Sec. 3, above) uses a different support (24,087 pixels), a different field convention ($f_{\rm CW} - \frac{1}{2} = A_p/2$ with $N_{\rm spiral}$-weighted mean subtraction), a NaMaster pseudo-C_ℓ MASTER-decoupling estimator rather than a real-space fit, and its own 500- and 10^4-draw null families; the two diagnostics are not interchangeable and are not combined into a single detection statistic.

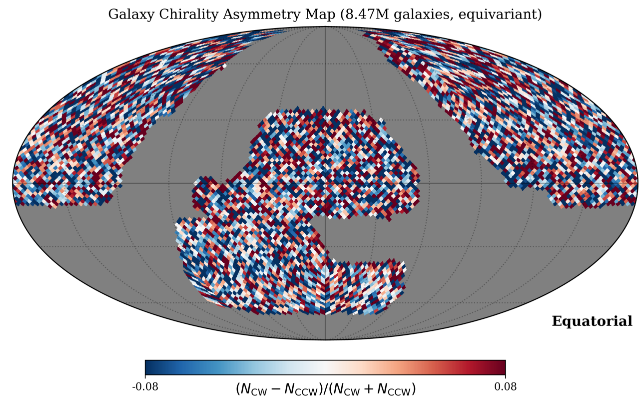

Figure 1. Equivariant chirality asymmetry map of the full 8.47 M-galaxy catalog (Mollweide projection, equatorial coordinates; per-pixel asymmetry $A_p = (N_{\rm CW} - N_{\rm CCW})/(N_{\rm CW} + N_{\rm CCW}) = 2(f_{{\rm CW},p} - \frac{1}{2})$ at HEALPix `NSIDE` $= 64$, color scale $[-0.08, +0.08]$) on the 24,087-pixel FSC support ($f_{\rm sky} = 0.49005$; Sec. 2.1). This is the full-catalog map, a broader footprint than the 23,633-pixel, 887,472-galaxy HC-RI support used for the primary dipole fit below; no coherent large-scale structure is visually apparent in either support.

3. PRIMARY REAL-SPACE DIPOLE NULL

The paper's primary result is a full-sky real-space dipole fit (`healpy.fit_dipole`) on the 887,472-galaxy supported-pixel HC sample against a fixed-occupancy label-randomization null (10^4 draws). The fitted amplitude is $A_{\rm dip} = 0.467\%$, with $z_{\rm mom} = +0.635$ and one-sided add-one rank $p = 0.238$: null-consistent. Figure 1 shows the full 8.47-million-galaxy catalog's per-pixel asymmetry map on the broader FSC support; the 887,472-galaxy HC-RI support used for this fit is a subset of it (caption below).

To make this null quantitative against a specific alternative amplitude, we use a coverage-calibrated observed-label injection–recovery: signal injections at fixed pixel occupancy on 2,000 isotropically random axes per amplitude are passed through the exact committed primary estimator and the exact committed 10^4-draw null. Recovered detection probability rises monotonically with amplitude (e.g. 0.24 at $A = 0.40\%$, 0.91 at 0.90%, saturating to unit coverage by 1.5%) and crosses 95% coverage at

$$A_{95}^{\rm obs} \simeq 0.98\% \quad \text{(full-amplitude, observed-label)}, \quad (1)$$

i.e., $A_{95}^{\rm obs}$ is a *detection-power* threshold: the amplitude at which an injected signal, run through the exact primary estimator and null, is recovered at $\geq 95\%$ power. It is not itself a confidence-level bound on the measured value. The committed null distribution's own 95th percentile, 0.669% (research/bh_universe_dipole/a95_null_cl_2026_09_02.py), is also not a confidence-level bound on $A_{\rm dip}$: it is the same no-signal acceptance boundary already reported as the rank test above ($p = 0.238$), restated as a critical value rather than a rank – values of $A_{\rm dip}$ at or above 0.669% occur in fewer than 5% of null draws, which is a statement about the null, not an interval on the true amplitude. A genuine 95% confidence-level upper limit requires the complementary Neyman inversion: the largest injected amplitude A whose recovered-amplitude distribution keeps the observed $A_{\rm dip} = 0.4665\%$ above its own 5th percentile. Running the same committed injection model and estimator used for Eq. 1, but recording the 5th (not the 50th/84th) percentile of `recovered_amp` per injected amplitude and inverting against $A_{\rm dip}$, gives

$$A_{95}^{\rm CL} \simeq 0.75\% \quad \text{(full-amplitude, Neyman 95\% CL)}, \tag{2}$$

(bracketed between 0.75%, $p_5 = 0.466\%$, and 0.80%, $p_5 = 0.493\%$; research/bh_universe_dipole/a95_upper_limit_2026_09_02.py, 2,000 random-axis injections per amplitude, same estimator/support/null as Eq. 1). The observed $A_{\rm dip} = 0.467\%$ lies below this genuine 95% CL upper limit, below the detection-power floor $A_{95}^{\rm obs} = 0.98\%$, and below the null's own 95th-percentile critical value 0.669%. Because $A_{\rm dip}$ is a low realization of a positive-definite estimator whose no-signal distribution has mean 0.362%, this classical one-sided limit falls below the 95%-detection sensitivity $A_{95}^{\rm obs}$. The limit retains exact frequentist coverage, but it is a statement about this realization, not about the experiment's reach; a downward fluctuation can make a classical limit stronger than the amplitude the experiment is guaranteed to detect (the standard motivation for unified/$\mathrm{CL_s}$ constructions). We therefore confront the model in Sec. 5 with the realization-independent floor $A_{95}^{\rm obs} = 0.98\%$, the more conservative of the two. $A_{95}^{\rm obs}$ and $A_{95}^{\rm CL}$ are additionally both *observed-label* statements on the classifier's chirality labels, not physical parity-amplitude bounds on deprojected spin: the physical bound additionally requires a spatially resolved morphology transfer function, which remains an open gate in the archived release paper (15). This distinction matters directly for Sec. 5, which uses Eq. 1 as the sensitivity floor against which the black-hole-universe claim is confronted.

A distinct harmonic diagnostic on the wider, 24,087-pixel FSC support (Fig. 1) gives a fixed-occupancy $\ell = 1$ moment $z = +6.923$ (add-one rank $p = 0.001996$ on 500 draws; a separate binomial-monopole null gives +6.983 at 500 draws and +7.207 at 10^4 draws). This diagnostic uses a different support, estimator (NaMaster MASTER-decoupled harmonic power, not the real-space uniform-pixel dipole fit), and null family than the primary HC-RI result above, and does not by itself overturn the primary null: the archived release attributes it to classifier- and footprint-systematic structure rather than a dipole-aligned axis, on the strength of the classifier-injection forward model (Sec. 2.2) and a battery of declination/RA/equal-area partition checks (15). Full weighted-least-squares and GZ1 human-vote diagnostics are retained as systematics cross-checks in the archived release and are not repeated here.

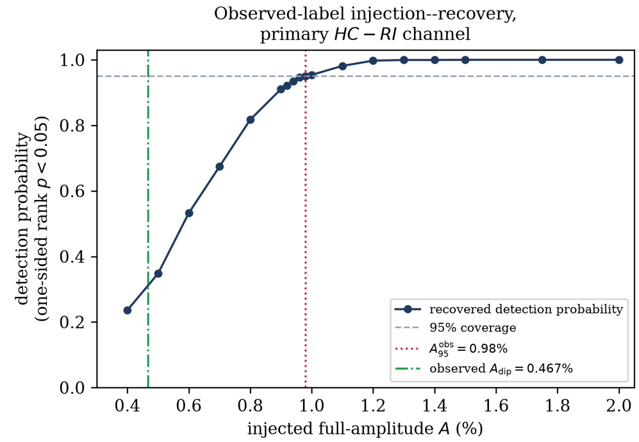

Figure 2. Observed-label injection–recovery curve underlying Eq. 1: recovered one-sided detection probability ($p < 0.05$ against the exact committed 10^4-draw null) vs. injected full-amplitude A, on 2,000 isotropically random axes per amplitude, run through the exact committed primary HC-RI estimator. The curve crosses 95% coverage at $A_{95}^{\rm obs} = 0.98\%$ (red dotted); the observed amplitude $A_{\rm dip} = 0.467\%$ (green dash-dot) sits well inside the sub--detection-power region. Plotted directly from the committed pipelines/p2_chirality/analysis/a95_observed_label_upper_limit_v1_0_265.json (no re-inference).

A (%)	Det. frac.	A (%)	Det. frac.
0.40	0.236	0.96	0.947
0.50	0.348	0.98	0.950
0.60	0.533	1.00	0.954
0.70	0.675	1.10	0.982
0.80	0.818	1.20	0.998
0.90	0.911	1.50	1.000

Detection fraction is monotonic in injected amplitude across all 18 grid points (2,000 axes each); the tabulated crossing between the 0.96% and 0.98% rows brackets $A_{95}^{\rm obs}$ (linear interpolation 0.980%; a logistic fit to the full curve gives 0.955%, consistent to better than 0.03 percentage points).

To characterize how the real-space dipole, its nuisance-marginalized weighted-least-squares (WLS) variant, the global monopole, and the NaMaster MASTER-decoupled $\ell = 1$ power co-vary, the archived release evaluates all four estimators on one shared `NSIDE` $= 8$ block bootstrap ($N_{\rm boot} = 2{,}000$ resamples, all valid, seed 42, pre-support-cut HC sample $N = 949{,}584$ – distinct from this paper's own 887,472-galaxy primary channel, which is the post-support-cut subset). The resulting correlation matrix, computed directly from the committed 4×4 joint covariance (`pipelines/p2_chirality/outputs/canonical_provenance/g3_joint_estimator_covariance_master_v2.json`), is

	real-space	WLS	monopole	$C_1^{\rm master}$
real-space	1.000	0.158	−0.037	0.794
WLS	0.158	1.000	−0.020	0.129
monopole	−0.037	−0.020	1.000	−0.061
$C_1^{\rm master}$	0.794	0.129	−0.061	1.000

The monopole is the only channel with $|z| > 3$ against its own bootstrap scatter ($z = -6.57$; the FSC harmonic diagnostic above is the same structure seen through a different estimator, consistent with its 0.79 correlation to the real-space channel here); the other three channels are all non-significant on this same shared bootstrap ($z = +2.21$ real-space dipole, +0.81 WLS dipole, −0.61 MASTER $\ell = 1$), consistent with the primary null. The block-bootstrap real-space $z = +2.21$ is a distinct statistic from the primary null's fixed-occupancy $z_{\rm mom} = +0.635$ (HC-RI-support quantity, above): the former is the full-sample dipole amplitude over its own `NSIDE` $= 8$ block-bootstrap scatter (which does not model the pixel-occupancy constraint), the latter is the fixed-occupancy label-randomization statistic reported as the primary result; both are non-significant and neither supersedes the other. The monopole is nearly uncorrelated with the other three; the primary real-space dipole and its WLS variant are only weakly correlated (0.16), consistent with them answering related but non-identical questions on different supports.

4. VOID/NON-VOID ENVIRONMENT CONTRAST

A second, independent probe tests whether classifier-labelled chirality depends on large-scale environment, using the released DESIVAST void catalog (14) intersected with the DESI DR1 spectroscopic TARGETID cross-match of the DR8 chirality catalog, as constructed in (16) (joined on TARGETID), hereafter the DR1 companion of the catalog in Sec. 2. The DESIVAST GALZONE TARGET universe (694,642 unique TARGETIDs) joined against the catalog yields 145,789 rows; retaining the 145,766 rows with `OUT`$= 0$ as the quality parent, 31,937 fall inside the union of released VoidFinder holes ("void") and 113,829 do not ("non-void").

Stage	N
DESIVAST GALZONE TARGET universe	694,642
→ joined against the chirality catalog	145,789
→ quality parent (`OUT`$= 0$)	145,766
void (VoidFinder hole union)	31,937
non-void	113,829

The focal fit is an unpenalized Newton–Cholesky logistic MLE on a declared 13-column linear nuisance basis (void indicator, linear redshift, magnitude, log-size, classifier confidence, extinction, imaging leg, morphology, and the GALZONE edge flag standing in for coarse sky blocks), with cluster-robust (CR1) inference on 50 occupied `NSIDE` $= 4$ HEALPix clusters; a null-imposed Rademacher wild-cluster efficient-score test provides a second, non-asymptotic inference path on the same clustering. The basis was declared after review and inspection of the data and is exploratory, not preregistered. It gives an adjusted non-void-minus-void contrast

$$\Delta f_{\rm CW} = +0.00145\,, \quad {\rm SE} = 0.00332\;({\rm cluster\text{-}sandwich})\,, \tag{3}$$

95% CI [−0.00504, +0.00795], two-sided normal $p = 0.661$; a null-imposed 99,999-draw Rademacher wild-cluster efficient-score test gives $p = 0.673$. Both results are null-consistent: classifier-labelled spiral chirality shows no significant dependence on void/non-void environment in this sample. This test is exploratory and post-hoc (the void/non-void hierarchy was fixed after inspecting the data) and is reported as a catalog-specific non-detection for classifier labels, not a physical-handedness or cosmological constraint. A secondary, four-class cosmic-web diagnostic (T-Web; Fig. 3) on 812,793 env-labelled rows (Void $N = 428$, Wall $N = 6{,}673$, Filament $N = 408{,}187$, Cluster $N = 397{,}505$) shows the same pattern – no significant CW-fraction trend from void through wall, filament, to cluster; the Void bin's small N limits its statistical power relative to the other three classes; author-constructed VoidFinder and the Tempel and ASTRA web classifications are retained as further secondary diagnostics in the archived companion study and are not repeated here.

4.1. *Robustness: clustering scale and void-definition sensitivity*

The focal contrast (Eq. 3) clusters standard errors on `NSIDE` $= 4$ HEALPix superpixels. Repeating the identical $N = 145{,}766$, 13-column fit and

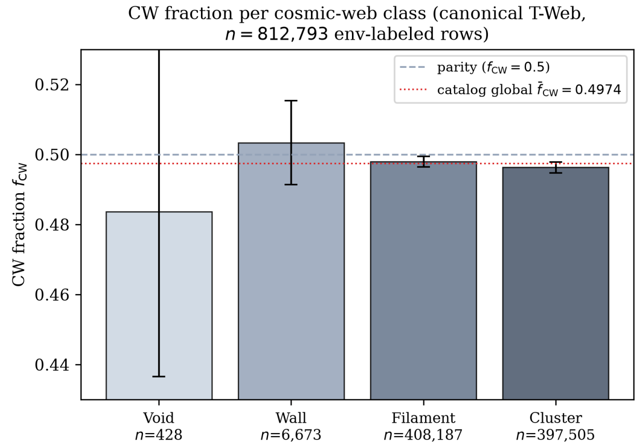

Figure 3. Secondary diagnostic: classifier-labelled CW fraction per T-Web cosmic-web class (Void/Wall/Filament/Cluster) on 812,793 env-labelled rows (Sec. 4). This four-class T-Web run is a secondary cross-check; the primary void/non-void DESIVAST result (Eq. 3) uses the two-category VoidFinder-hole membership test, not this classification. The Filament ($\bar{f}_{\rm CW} \approx 0.4980$) and Cluster ($\bar{f}_{\rm CW} \approx 0.4958$) bars sit slightly off the reference line at different offsets; both are consistent with the same catalog-wide residual monopole discussed in Sec. 2.2, not an environment-dependent chirality effect (both are within a few tenths of a percentage point of the 0.4974 reference and neither individually clears the null threshold used in the primary contrast).

contrast at `NSIDE` = 2 and `NSIDE` = 8, and separately clustering by the 3,750 nearest VoidFinder-MAXIMALS 3-D clusters, holds the point estimate fixed at +0.00145 throughout while the 95% CI half-width varies mildly with the clustering scale: `NSIDE` = 2 $[-0.00622, +0.00912]$, `NSIDE` = 4 (focal) $[-0.00504, +0.00795]$, `NSIDE` = 8 $[-0.00496, +0.00787]$, and the 3-D-cluster scheme $[-0.00476, +0.00767]$ – every interval contains zero. A separate, wider sensitivity family repeats the void/non-void contrast under three void-finding algorithms spanning two algorithmic families – VoidFinder (sphere-growing) and the ZOBOV/V2 watershed as V2-REVOLVER and V2-VIDE – plus two catalog-native GALZONE zone definitions, five correlated statistics in total, none of which approaches the Bonferroni-5 per-test threshold ($|z| \leq 1.25$ across all five rows against $|z|^{\rm Bonf}_{0.05,5} \approx 2.58$); the simultaneous descriptive interval across all five, inflated to hold jointly at family-wise 95% confidence, has a maximum half-width of 0.75 percentage points and remains centered near zero. Full per-algorithm tables and the catalog-native GALZONE construction are in the archived companion study (16). We report one multiplicity note across this paper's three headline tests (the primary real-space dipole, Sec. 3; the DESIVAST void/non-void contrast, Sec. 4; and the T-Web four-class diagnostic above): none is independently significant at even an uncorrected $\alpha = 0.05$, so no multiple-comparison correction changes the qualitative null verdict on any of the three.

5. THE BLACK-HOLE-UNIVERSE PREDICTION AND ITS EXCLUSION

5.1. *What the model predicts*

Poplawski's mechanism papers (8; 9; 10) establish that Einstein–Cartan torsion, sourced by fermion spin, generates a repulsive spin–spin contact interaction that grows faster than the attractive curvature term at high density, halting collapse before a singularity and driving a bounce into an expanding daughter universe inside the black hole's interior. Poplawski (2010) (8) first proposes torsion as an alternative to inflation for generating a nonsingular early universe; Poplawski (2012) (9) develops the spinor–torsion coupling that produces a nonsingular bounce; Poplawski (2016) (10) extends the mechanism specifically to a universe originating inside a black hole, arguing the Einstein–Cartan interior bounces into an expanding daughter universe rather than terminating in a singularity. None of these three mechanism papers addresses rotation or a preferred axis for the daughter universe – they establish only that the bounce itself occurs. Poplawski (2020) (11) extends this picture: if the parent black hole is rotating, its spin axis defines a preferred, non-inertial axis for the daughter universe, and "galaxies tend to align their axes of rotation with the preferred axis, resulting in clockwise-counterclockwise asymmetry." This is the paper's only statement connecting the model to an observable galaxy-spin signature. We read it in full and find *no computed dipole amplitude, alignment fraction, or relaxation timescale* anywhere in it or in the mechanism papers it builds on: the claim is qualitative – a preferred axis exists, and galaxies tend toward it – not quantitative. No torque strength, relaxation timescale, or coupling constant is given from which an amplitude could in principle be derived without further modeling assumptions beyond what the cited papers state.

To confront the model with data at all, a quantitative closure must be supplied that the cited literature does not provide. We adopt the minimal one: let $\eta \in [0, 1]$ be the fraction of galaxies whose spin has relaxed toward the preferred axis, and take the resulting population-level observed dipole amplitude as

$$A_{\rm pred} \approx \eta \,, \qquad (4)$$

i.e., full alignment ($\eta = 1$) would produce a maximal observed dipole and no alignment ($\eta = 0$) none, linearly in between. Equation 4 is *not* derived from Poplawski's papers – it is the simplest closure under which the model becomes falsifiable by a dipole measurement, made explicit here because the alternative is to leave the claim untestable.

5.2. *Confrontation with the observed-label sensitivity floor*

Under Eq. 4, η is exactly the quantity Eq. 1 bounds: the catalog's coverage-calibrated 95% sensitivity upper limit excludes

$$\eta > A_{95}^{\mathrm{obs}} \simeq 0.98\%$$

at $\geq$ 95% coverage on the primary channel. We compare this floor against the amplitudes reported in the literature that motivates the model in the first place, using the committed, deterministic script `research/bh_universe_dipole/poplawski_dipole_exclusion_2026_09_02.py` (numpy only; no fitting, no randomness; every output is a literal input or a closed-form function of inputs). Table 5 lists the result: every literature amplitude that has been offered as evidence for a preferred spin axis (Longo 2011; Shamir 2012, 2020, 2022, 2025) exceeds A_{95}^{obs} by a factor of 2–20 at face value. The present catalog's 887,472-galaxy HC-RI supported-pixel sample is 4–3,400× larger than every comparison sample in Table 5 except Shamir (2022)'s same-survey DESI Legacy sample ($N = 1.3$ million), which is larger than the primary channel. Under the committed script's illustrative observed-to-physical bridge $g = 0.398$ (Assumption 2 below; not an established calibration), the two lowest-amplitude, largest-N rows – Shamir (2020) and same-survey Shamir (2022), both 2–4% – drop below A_{95}^{obs} (`exceeds_A95_obs_after_g_bridge=false` in the committed output json); the other three rows remain above it under the bridge. We do not use this bridge to strengthen the exclusion (Assumption 2); we report it because it weakens the exclusion for exactly the two claims made on the largest comparison samples, and that should not go unstated.

5.3. *Exclusion statement*

Under the closure of Eq. 4, the DESI Legacy DR8 catalog's coverage-calibrated observed-label sensitivity disfavors, at $\geq$95% detection power (not a confidence-level exclusion; a genuine 95% CL statement is given in Sec. 3), any preferred-axis alignment mechanism that would produce an observed-label dipole above $\sim$1% on this sample – a factor of 2–20 below the amplitudes the model's own observational motivation reports at face value (Table 5), though two of those amplitudes, on the two largest comparison samples, drop below the floor under the (not adopted for strengthening) illustrative bridge. This statement rests on five explicit assumptions:

Table 5. Literature spin-axis amplitude claims vs. $A_{95}^{\mathrm{obs}} = 0.98\%$ (Eq. 1). Ratio is the claim's low-end amplitude over A_{95}^{obs}; all exceed 1, i.e. all lie above this catalog's detection-power floor at face value (two of five drop below it under the illustrative $g = 0.398$ bridge; §5.2). The table's five rows pool four non-commensurable statistic families – Longo's dipole amplitude, Shamir (2012)'s per-bin asymmetry, Shamir (2020/2022)'s global asymmetry fraction, and Shamir (2025)'s clockwise:counterclockwise count ratio, each from a distinct label space (visual/algorithmic handedness classification vs. the Ganalyzer algorithm vs. a small-N visual/algorithmic count) – against this paper's observed-label ViT+TTA amplitude; the ratio column is not a like-for-like statistical comparison across rows. Computed by `research/bh_universe_dipole/poplawski_dipole_exclusion_2026_09_02.py`.

Source	Amplitude	N	Ratio
Longo 2011 (1)	$\sim$7%	15,158	7.1
Shamir 2012 (2)	5–20%	127,000	5.1
Shamir 2020 (3)	2–4%	200,000	2.0
Shamir 2022 (4)	2–4%	1,300,000	2.0
Shamir 2025 (5)	20–33%	263	20.4

1. The closure $A_{\mathrm{pred}} \approx \eta$ (Eq. 4) is adopted for lack of a quantitative prediction in the cited Poplawski papers; a first-principles closure would require a computed relaxation timescale and torque strength that do not exist in the literature we could locate.
2. A_{95}^{obs} is an *observed-label* sensitivity floor (classifier-labelled chirality), not a physical parity-amplitude bound on deprojected spin. The observed-to-physical bridge requires the spatially resolved morphology transfer function, an open gate in the archived catalog paper; an illustrative bridge factor $g = 0.398$ used there for a different comparison is explicitly not an established calibration and is not used to strengthen this statement.
3. The preferred axis is a free direction in the model (inherited from an unobserved parent black hole); `healpy.fit_dipole` fits both amplitude and direction, which is the correct match to a model with no predicted axis. This analysis, like the tests in Secs. 3–4, tests amplitude only (a full-sky real-space dipole

and a void/non-void contrast), not a search matched to a specific predicted axis.

4. The committed script also reports an illustrative $1/\sqrt{N}$ sensitivity-floor ansatz calibrated to this catalog's own anchor point, used only to compare statistical reach across sample sizes; it is not a re-derivation of any other paper's estimator, mask, or null and is not the basis of the Ratio column in Table 5.

5. Literature amplitude ranges are reported as published, with no re-analysis of their pipelines performed here (cross-reference only, as in Iye *et al.* (6) and Patel & Desmond (7)). Shamir (2025)'s 20–33% is that paper's reported clockwise:counterclockwise count-ratio range ($\sim$2:1 to 1.5:1), not an asymmetry fraction in the same sense as the other rows (Table 5 caption).

We emphasize what this section does *not* claim: it does not test, and is not evidence for or against, the nonsingular matter-bounce cosmology this research program otherwise develops from the same torsion mechanism (that line derives an early-universe non-Gaussianity signature, not a late-time galaxy-spin observable, and stands or falls on its own evidence). It tests only Poplawski's black-hole-universe model's one stated observational consequence, under the minimal closure that makes that consequence testable at all.

6. DISCUSSION

Two independent probes on this catalog – a full-sky real-space dipole (Sec. 3) and a void/non-void environment contrast (Sec. 4) – are both null-consistent. Neither reaches the 2–33% amplitudes reported by Shamir and collaborators, nor Longo's $\sim$7% claim; both are consistent instead with the independent reanalyses of Iye, Yagi & Fukumoto (2021) (6), who found no significant dipole in Shamir's own SDSS catalog after correcting known systematics, and Patel & Desmond (2024) (7), who found consistency with isotropy across every public spin-classification dataset to within 3σ. At 8,474,531 total catalog objects, this is the largest single catalog release brought to the question; the 887,472-galaxy primary supported-pixel test is itself larger than every comparison sample in Table 5 except Shamir (2022)'s same-survey 1.3-million-galaxy DESI Legacy catalog. Its detection-power floor (Eq. 1) is 2–20× tighter, at face value, than the amplitudes the disputed claims report (Table 5).

Section 5 adds the piece the literature comparison alone does not supply: a direct confrontation with the theoretical claim that motivates testing spin-axis alignment in the first place. The result is a genuine, if qualified, exclusion – qualified because Poplawski's papers do not themselves commit to a quantitative amplitude, so the exclusion binds a minimal closure of the claim rather than a number the model's author derived. We regard that as the honest state of the test: the model predicts a direction of effect (preferred-axis alignment) that this catalog is well-positioned to detect were it present at the amplitudes the observational literature has claimed, and does not detect it.

7. CONCLUSIONS

We release an 8,474,531-object DESI Legacy DR8 chirality catalog and report two independent, large-sample null tests of a preferred galaxy-spin axis: a primary real-space dipole ($N_{\rm support} = 887{,}472$; $z_{\rm mom} = +0.635$, $p = 0.238$; $A_{95}^{\rm obs} \simeq 0.98\%$) and a DESIVAST void/non-void environment contrast ($N = 145{,}766$; $\Delta f_{\rm CW} = +0.00145$, $p = 0.66$). Reading Poplawski's rotating black-hole-universe papers in full, we find the model predicts only a qualitative preferred-axis alignment tendency, not a computed amplitude; under the minimal closure that makes the claim testable, our sensitivity floor disfavors an alignment-driven dipole above $\sim 1\%$ at $\geq 95\%$ detection power, well below the ~ 2–33% amplitudes reported in the observational literature the model is invoked to explain. This confirms, on the largest catalog release assembled for the question, the independent reanalyses of Iye *et al.* (2021) and Patel & Desmond (2024). We make no claim beyond this stated test: the result bears on the black-hole-universe model's spin-axis prediction only, not on the separate matter-bounce cosmology this research program develops from the same torsion mechanism.

APPENDIX

A. SUPPLEMENTARY SYSTEMATICS DIAGNOSTICS

This appendix condenses the systematics checks the archived release uses to attribute the FSC harmonic residual (Sec. 3) to classifier- and footprint-level structure rather than a dipole-aligned axis, restoring self-containedness beyond the primary channel's own null result.

Table 6. Pixel-level injection calibration vs. label-identity models (image-level parity injection, $N = 20{,}000$, 10 seeds).

Model	dA/df	σ from pixel result
Pixel-level injection (measured)	$+0.0167 \pm 0.0089$	—
Naive label identity	$+0.434$	47σ
Mixture-corrected label identity	-0.0093	2.9σ

Sample-purity ladder. Stratifying the full spiral sample by equivariant max-class probability into five confidence bins, a $+3.29\sigma$ excess appears in the 1.87M-galaxy $[0.5, 0.6)$ bin under a confidence-stratified dipole estimator and null; this signal does not survive the sample-purity cut used for the primary channel ($p_{\rm eq} > 0.6$ gives -0.03σ under the same statistic). This is a distinct diagnostic from the primary HC-RI real-space dipole and illustrates why the release adopted the high-confidence cut as the primary science sample rather than the full spiral catalog.

Fixed-support multipole vector. On the checksummed 24,087-pixel FSC support, the binary-mask fixed-occupancy null gives the multipole vector $z_{\ell=1,\ldots,5} = (+6.923, +2.931, +3.089, -0.317, +1.606)$: only $\ell = 1$ is a $> 5\sigma$ outlier, and its rank-p significance ($p = 0.001996$, Sec. 3) does not by itself distinguish a primordial dipole from a low-ℓ classifier or footprint systematic, which is why the release cross-checks with directional partitions.

Directional partition checks. Splitting the FSC support into four RA quadrants gives per-quadrant dipole deviations ranging from -0.82σ to $+2.49\sigma$ – scatter, not the consistent projection a primordial dipole would produce. A Galactic-hemisphere split gives $\sigma_{\rm iso} = +0.47$ (north, $b > 0$) and $+2.02$ (south, $b < 0$) against an independent per-pixel binomial null, consistent with dust-correlated foreground structure in the south rather than a primordial signal. Seven equal-spiral-count declination slabs ($N = 457{,}308$–$457{,}309$ each) span $f_{\rm CW}$ deviations of -0.110% to -0.463% from 0.5, all within 0.5%; an independent equal-area declination partition gives a compatible verdict (deviations up to 0.49% on the full spiral sample, 0.56% on the HC subsample). All partition checks are consistent with a bounded, spatially structured residual entering upstream of the classifier (Sec. 2.2), not a coherent dipole-aligned gradient.

Statistical-only scale reference. For the full-amplitude dipole convention $p_{\rm CW}(\hat{n}) = \frac{1}{2}(1 + A\cos\theta)$, the idealized full-sky Fisher floor at $N_{\rm spiral} = 3{,}201{,}160$ is $\sigma(A) = \sqrt{3/N_{\rm spiral}} = 9.7 \times 10^{-4}$ (a 3σ ideal floor of $\approx 0.29\%$); this full-sky idealization omits the realized footprint, classifier transfer, and orientation dependence and is a scale reference only, not a calibrated detection threshold or a substitute for the injection-recovery-calibrated $A_{95}^{\rm obs}$ of Sec. 3.

A.1. *Robustness and disclosure: pixel-level calibration, full-parent behavior, and structure cross-correlations*

Pixel-level calibration of the injection–recovery curve. The label-level injection–recovery curve used to set $A_{95}^{\rm obs}$ (Sec. 3) is a catalog-level identity check; a companion image-level pixel-parity injection through the production equivariant classifier ($N = 20{,}000$ Legacy DR9 cutouts, 10 seeds, true mirror-flip of raw pixels rather than relabeling) recovers a pixel-level slope $dA/df = +0.0167 \pm 0.0089$, far below both the naive label-identity slope ($+0.434$, 47σ from the pixel result) and the mixture-corrected label identity (-0.0093, 2.9σ). Table 6 records this comparison. The production pipeline suppresses pixel-level parity leakage well below what the label-level injection model alone would predict, consistent with the baseline residual already disclosed in Sec. 2.2.

Full-parent selection behavior. The primary channel's 887,472-row `primary_hc && !raw_flip_qc_unsafe` selection exists to exclude rows flagged by post-review raw-flip QC and low-confidence classification. Re-running the same fixed-occupancy null on the full 3,200,420-galaxy parent (no QC/confidence cuts) gives a dipole $A = 0.566\%$ at $z = +4.44$, axis RA 278.6°/Dec +25.3°. A graded QC sweep localizes this to a single cut: removing only the `primary_hc` confidence cut restores the null ($z = +0.68$), while removing only the raw-flip cut does not ($z = +9.13$, at a third, near-antipodal axis); a per-imaging-leg partition shows the excess is also removed by dropping the DES leg alone ($z = +0.48$) but not by any Galactic cut. The axis is inconsistent across selections (shifting by 107.5°), and the monopole-plus-mask-leakage floor in this population is 0.19%. This pattern – signal present only in the population the primary QC cuts are designed to exclude, removed by a single confidence or footprint-leg cut, with an unstable axis – is consistent with a confidence-cut/footprint systematic rather than a shared physical signal, and is reported here as a systematic, not subtracted from or used to adjust the primary channel's null result.

Structure cross-correlations. A pre-registered battery of 15 chirality×structure statistics (local spiral surface-density quartiles; angular cross-correlation with 1,244 spectroscopic anomaly positions; five spectroscopic-redshift bins; and dipole projections along the CMB dipole and CMB quadrupole–octopole axes plus the catalog's own free best-fit

Table 7. Chirality×structure cross-correlation nulls (representative statistics; full battery of 15 in the archived analysis).

Statistic	z	Result
Environment density, χ^2	−0.38	null
Environment density, trend	+0.17	null
Anomaly cross-corr. (rotation null)	−1.05	null
Redshift trend	−1.63	null
CMB dipole axis	+0.21	null
CMB quadrupole–octopole axis	−0.13	null
Free best-fit direction	+0.32	null

Table 8. Parity fraction by environment bin against the DESI DR1 BGS_BRIGHT-21.5 clustering catalog (density contrast δ_k, k=10-NN vs. randoms). z is against the 1,000-realization label-shuffle null.

Environment	N	$f_{\rm CW}$	z
Spec-z, 3D comoving ($N = 121{,}417$)			
Void-like	24,284	0.48999 ± 0.00321	−0.61
Wall/filament	72,849	0.49489 ± 0.00185	+1.11
Node-like	24,284	0.49238 ± 0.00321	−0.84
Projected, photo-z/no-z ($N = 949{,}584$)			
Void-like	189,917	0.49602 ± 0.00115	+0.64
Wall/filament	569,750	0.49680 ± 0.00066	−1.06
Node-like	189,917	0.49388 ± 0.00115	+2.96

direction), each tested against 1,000-realization label-shuffle and sky-rotation nulls with a 3σ look-elsewhere-corrected threshold declared in advance, returns no detection: environment (χ^2 $z = -0.38$, trend $z = +0.17$), the anomaly cross-correlation (all eight angular bins and the nearest-neighbor excess null; the largest raw excursion falls from -2.37σ to -1.05σ under a rigid-rotation null), redshift (trend $z = -1.63$), and the two fixed astrophysical axes ($z = +0.21$, -0.13) are all null, and the free best-fit dipole (0.44%) is below its own look-elsewhere-corrected null ($0.38\% \pm 0.16\%$, $z = +0.32$). Table 7 summarizes. Two data limits apply: the environment test uses a parity-blind projected surface-density proxy in place of a 3D void/filament catalog, since the DESI DR1 LSS products on disk are QSO-only at $z = 0.8$–2.1 with no overlap with the $z \lesssim 0.3$ spiral sample and no low-z void catalog was available; and Shamir's claimed axis is not tested, as the sources this paper draws its quoted amplitudes from quote no explicit RA/Dec for it.

External environment: DESI DR1 BGS clustering catalog. A density-based environment split against an external tracer field – k=10-NN density contrast measured on the public DESI DR1 `BGS_BRIGHT-21.5` clustering catalog and its randoms (300,043 tracers, 4 randoms per cap; not a void finder) – finds no dependence of the parity fraction on environment in either subset it can be tested on. Table 8 reports the three-way void-like/wall-filament/node-like split for the 121,417-spiral DESI spectroscopic-redshift subset (three-dimensional comoving density) and the 949,584-spiral photometric/no-redshift subset (projected density).

The χ^2 trend across bins is null in both subsets (spec-z: obs 1.872 vs. null 1.984 ± 1.841, $z = -0.06$, $p = 0.952$, confirmed by a sky-rotation null at $z = -0.12$, $p = 0.917$; projected: $z = -1.51$, $p = 0.084$ two-sided). Per-bin dipole amplitudes are likewise consistent with their nulls in every bin ($|z| \leq 2.15$). The largest excursion in the full 14-statistic battery (2 subsets × [3 $f_{\rm CW}$ + 1 χ^2 + 3 dipole]) is the projected node-like bin's parity fraction, $z = +2.96$ ($p_{\rm local} = 0.0060$, $p = 0.084$ after the pre-registered ×14 Bonferroni look-elsewhere correction) – below the pre-declared 3σ post-look-elsewhere detection threshold ($|z| \gtrsim 3.7$). We report this channel as **null**: splitting the chirality sample by local environment, measured against a genuine external tracer population rather than the chirality catalog's own projected proxy, shows no environment dependence of handedness. The spec-z 3D result is the more constraining of the two, since the projected proxy line-of-sight-integrates true 3D structure.

Data. DESI DR1 LSS `BGS_BRIGHT-21.5` clustering catalog, NGC+SGC, $0.1 < z < 0.4$ (300,043 tracers, 4 randoms/cap); the environment proxy is a k=10 nearest-neighbor density-contrast estimate against the randoms, not output from a void finder or void/filament classifier.

DATA AVAILABILITY

The 8,474,531-row catalog, quarantine, exact null arrays, generator scripts, and full provenance register are archived under Zenodo DOI 10.5281/zenodo.21461899 (concept DOI 10.5281/zenodo.21461898), published under CC-BY-4.0, and mirrored at https://huggingface.co/datasets/bamfai/galaxy-chirality-catalog. This is the same archived catalog release used by the sources this paper draws its quoted numbers from (15; 16). The black-hole-universe exclusion analysis (Sec. 5) is fully reproducible from the committed script `research/bh_universe_dipole/poplawski_dipole_exclusion_2026_09_02.py` and its output `research/bh_universe_dipole/outputs/poplawski_dipole_exclusion_2026_09_02.json`. Full systematics appendices, harmonic and weighted-least-squares diagnostics, the classifier bias-hardening audit, and the void/non-void secondary diagnostic paths (author VoidFinder, T-Web, Tempel, ASTRA) are not reproduced in this condensed manuscript; they remain available in the archived source papers (15; 16) referenced above. The catalog is a single Apache Parquet file (952,115,239 bytes for the pinned source revision) with one row per DESI Legacy DR8 galaxy: a `object_id` identifier, sky coordinates, the observed chirality label (CW/CCW/NOT_SPIRAL), per-class probabilities, and quality-control flags including `raw_flip_qc_unsafe`; a separate quarantine Parquet contains the 249,066 catalog-wide flagged rows (59,515 within the HC selection), and a machine-readable schema and validation report accompany both files at the Zenodo DOI and Hugging Face mirror above. This manuscript's own version, v4P.0.7 (dated September 5, 2026), is the pin for its numbers: it is served alongside its git-tracked source at https://github.com/Hubify-Projects/bigbounce/blob/main/pipelines/p4prime_chirality_test/paper/main.tex, and the underlying catalog/quarantine/null-array revision it draws from is independently pinned by the versioned Zenodo DOI 10.5281/zenodo.21461899 above. The four robustness/disclosure results of Sec. A.1 are individually reproducible from their registered manifests (the latter two registered in `reproducibility/manifests/programs/galaxy-chirality.json`):

- pixel-level calibration:
  `reproducibility/manifests/experiments/row13-image-level-injection-pilot.json`
- full-parent selection behavior:
  `reproducibility/manifests/experiments/p4p-row16i-full-parent-dipole.json`
- structure cross-correlations:
  `reproducibility/manifests/experiments/p4p-row16ib-axis-shift.json`
- external environment (DESI DR1 BGS clustering catalog):
  `reproducibility/manifests/experiments/row16ivb-bgs-environment.json`

SOFTWARE, FACILITIES, AND ACKNOWLEDGEMENTS

Software: NumPy, healpy, HEALPix, NaMaster (pymaster), PyTorch, pandas, SciPy. *Facilities:* DESI Legacy Imaging Surveys DR8 (data products); no proprietary telescope time was used in this analysis, which is entirely based on public archival imaging and catalog data. *Acknowledgements:* We thank the DESI, DESIVAST, Galaxy Zoo, and Galaxy Zoo DESI collaborations for public data products without which this analysis would not be possible.

REFERENCES

[1] M. J. Longo, "Detection of a dipole in the handedness of spiral galaxies with redshifts $z\sim 0.04$," Phys. Lett. B **699**, 224 (2011), arXiv:1104.2815.

[2] L. Shamir, "Handedness asymmetry of spiral galaxies with $z < 0.3$ shows cosmic parity violation and a dipole axis," Phys. Lett. B **715**, 25 (2012), arXiv:1207.5464.

[3] L. Shamir, "Patterns of galaxy spin directions in SDSS and Pan-STARRS show parity violation and multipoles," Astrophys. Space Sci. **365**, 136 (2020), arXiv:2007.16116.

[4] L. Shamir, "Analysis of spin directions of galaxies in the DESI Legacy Survey," Mon. Not. R. Astron. Soc. **516**, 2281 (2022), arXiv:2208.13866, DOI:10.1093/mnras/stac2372.

[5] L. Shamir, "The distribution of galaxy rotation in JWST Advanced Deep Extragalactic Survey," Mon. Not. R. Astron. Soc. **538**, 76 (2025), arXiv:2502.18781.

[6] M. Iye, M. Yagi, and H. Fukumoto, "Spin parity of spiral galaxies. III. Dipole analysis of the distribution of SDSS spirals with 3D random walk simulations," Astrophys. J. **907**, 123 (2021), arXiv:2011.00662.

[7] D. Patel and H. Desmond, "No evidence for anisotropy in galaxy spin directions," Mon. Not. R. Astron. Soc. **534**, 1553 (2024), arXiv:2404.06617.

[8] N. J. Poplawski, "Cosmology with torsion: An alternative to cosmic inflation," Phys. Lett. B **694**, 181 (2010), arXiv:1007.0587.

[9] N. J. Poplawski, "Nonsingular, big-bounce cosmology from spinor-torsion coupling," Phys. Rev. D **85**, 107502 (2012), arXiv:1111.4595.

[10] N. J. Poplawski, "Universe in a black hole in Einstein-Cartan gravity," Astrophys. J. **832**, 96 (2016), arXiv:1410.3881, DOI:10.3847/0004-637X/832/2/96.

[11] N. J. Poplawski, "Universe in a rotating black hole and preferred axis," arXiv:1910.10819 (preprint, posted October 2019; no journal publication identified as of this writing – this is the paper's only source for the observational preferred-axis claim and the year/identifier mismatch with its arXiv submission date is stated here explicitly).

[12] A. Dey, D. J. Schlegel, D. Lang *et al.*, "Overview of the DESI Legacy Imaging Surveys," Astron. J. **157**, 168 (2019), arXiv:1804.08657.

[13] M. Walmsley, C. Lintott, T. Géron *et al.*, "Galaxy Zoo DESI: detailed morphology measurements for 8.7M galaxies in the DESI Legacy Imaging Surveys," Mon. Not. R. Astron. Soc. **526**, 4768 (2023), arXiv:2309.11425.

[14] H. Rincón, S. BenZvi, K. A. Douglass *et al.*, "DESIVAST: Catalogs of Low-redshift Voids Using Data from the DESI Data Release 1 Bright Galaxy Survey," Astrophys. J. **982**, 38 (2025), DOI:10.3847/1538-4357/adb559, arXiv:2411.00148.

[15] H. Golden, "An Observed-Label Chirality-Dipole Null in 890,069 Quality-Controlled High-Confidence DESI Spirals and an 8.5-Million-Galaxy Catalog," v1.0.274 archived release, Zenodo doi:10.5281/zenodo.21461899 (2026).

[16] H. Golden, "A Catalog-Native DESIVAST Test of Classifier-Labelled Spiral Chirality in DESI DR1," v0.1.147 archived manuscript, served at https://bigbounce.hubify.app/papers/p5_desi_chirality_v0.1.147-2026-08-03.pdf (2026). No independent DOI has been minted for this manuscript as of this writing.

[17] H. Golden, "Galaxy Chirality Catalog: Archived Release (DESI Legacy DR8, 8,474,531 objects)," Zenodo (2026), concept DOI doi:10.5281/zenodo.21461898 (same archived release deposit as (15); the dataset and the release paper are versions of one Zenodo concept record and therefore intentionally share a DOI family – the paper cites the specific version 10.5281/zenodo.21461899, the dataset the version-independent concept 10.5281/zenodo.21461898).