

Redshift-Space Environmental Dependence of Spiral Chirality: A DESIVAST Three-Algorithm Void Null Test on 57,081 DESI DR1 Spirals

Houston Golden^{1, *}

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

(Dated: July 10, 2026)

Headline result: spiral galaxy chirality shows no void/non-void environment dependence across all five DESIVAST void definitions in DESI Data Release 1: the *family-wise Bonferroni-5 null* ($|\Delta f_{\text{CW}}| \leq 0.004$, $|z_{\Delta}| \leq 1.25$, $p_{\Delta} \geq 0.21$ across all five estimators) is the strictly quotable result; the single-estimand contrasts reported below are secondary summaries within that family. The DESIVAST-anchored path is *designated primary for reporting* (exploratory/post-hoc; see §V B: no timestamped plan predates the data, so all bounds below are exploratory and should not be read as pre-registered exclusions; T-Web is explicitly secondary). The *primary* DESIVAST estimand is the *footprint-restricted* void-vs-non-void CW-fraction contrast — both classes confined to the DESIVAST usable footprint, the single same-footprint estimand (§VIII B) — $\Delta f_{\text{CW}} = +0.0018$ ($z_{\Delta} = +0.78$, $p_{\Delta} = 0.43$, $n_{\text{void}} = 57,081$, $n_{\text{non-void}} = 253,276$; exact same-footprint retabulation), consistent with parity. The broader all- $z \leq 0.24$ -outside-hole contrast on the approximate $k=20$ VoidFinder membership (56,981 void spirals), $\Delta f_{\text{CW}} = +0.0007$ ($z_{\Delta} = +0.31$, $p_{\Delta} = 0.76$, $n_{\text{non-void}} = 621,964$), is a *secondary sensitivity check*; the two agree to within 0.11 pp, so the null verdict is invariant to the control-volume and membership-definition choice. The null holds across all five DESIVAST void definitions ($|\Delta f_{\text{CW}}| \leq 0.004$, $|z_{\Delta}| \leq 1.25$, $p_{\Delta} \geq 0.21$): the three sphere-PIS algorithm rows in Table XIII plus the two catalog-native GALZONE rows (V2-REVOLVER, V2-VIDE; §VIII D), with the full five-member Bonferroni family tabulated in Table XIV. This is a *bounded null* on the specific DESIVAST void/non-void contrast at 57,081 exact-footprint void spirals; the honest DR1 systematic envelope on any void/non-void split in the *classifier-labelled* CW fraction is ≈ 0.9 pp (quadrature of the counting-statistics CI with the fixed-void-geometry membership, any-hole-vs-maximal-sphere geometry, confidence, footprint, and match-radius terms; Table XI), with a 0.60 pp maximum single excursion (the geometry term, which is co-dominant with — not sub-dominant to — the membership term). This is a bound on the classifier-labelled CW fraction, not on physical spiral chirality: at the 69.91% ($\kappa = 0.40$) GZ1 accuracy floor the classifier-label statistic is an attenuated proxy for the underlying handedness contrast (attenuation factor $2a - 1 \simeq 0.40$; Appendix A). Under the symmetric-error approximation, the ≈ 0.9 pp classifier-label bound corresponds to a de-attenuated physical-chirality bound of ≈ 2.26 pp ($0.9/0.3982$); this weaker physical bound is the quantity model-builders should use to constrain parity-violation amplitudes. The de-attenuation carries additional uncertainty from environment-dependent relabeling; the void/non-void stratum is now directly measured and consistent with symmetric error (void/non-void error-asymmetry difference -0.018 , $z = -0.89$, $p = 0.37$), though the void arm is under-powered (± 3.7 pp) relative to the 0.2–0.6 pp headline bounds, so it corroborates but cannot by itself exclude a headline-scale environment-dependent label bias (§XIII/Appendix A). The robustly-quotable headline is the *family-wise Bonferroni-5 null* across all five DESIVAST void definitions ($|z_{\Delta}| \leq 1.25$), not any single most-favorable row (see §V B); the post-hoc primary designation and the few-dozen-trial analysis tree are explicitly disclosed. All bounds in this paper represent *fixed-redshift-space* metrics: the T-Web classification, DESIVAST void cross-check, and all null tests are performed without real-space reconstruction and inherit redshift-space distortion (RSD) effects; model-builders should not interpret them as pure real-space environmental constraints (see §XIII). The anisotropic tidal-tensor eigenvalue deformation channel is not quantified.

Setup. We cross-match the 8,474,531-galaxy chirality catalog of the companion catalog paper (Paper IV) [3] (arXiv:XXXX.XXXXX, posted concurrently under coordinated submission; classifier architecture, training, parity-equivariance validation, and monopole origin summarized in Appendix A; see also [A1]) to the DESI Data Release 1 redshift catalog (16.36×10^6 ZWARN=0 input rows) to test whether spiral galaxy handedness is statistically independent of large-scale structure environment. *Scope of the two DR1 inputs:* DESIVAST [13] provides the *void catalog* (three-algorithm: VoidFinder sphere-growing, V2-REVOLVER and V2-VIDE watershed); DESI DR1 redshifts provide the *environmental anchor* through which we run a T-Web tidal classification on 14,622,283 DR1 spectroscopic galaxies to assign every matched spiral to a cosmic-web class.

Independence from Paper IV internals. The headline Δf_{CW} result uses only the *public* per-galaxy `class_eq` labels (HuggingFace catalog, CC-BY-4.0, inspectable now; §II) joined to public DESI DR1 and DESIVAST data. Crucially, it is *algebraically invariant* under any catalog-wide monopole shift — the void/non-void contrast cancels the monopole exactly — so it is unaffected by the Paper IV monopole amplitude, its uncertainty, or any future revision (§II). Label provenance is validated model-independently in Paper IV: replacing the learned CW/CCW labels with Galaxy Zoo 1 human votes alone returns the same parity null at $z = -0.54\sigma$ ($N = 46,017$), confirming the catalog’s

handedness signal does not inherit its pseudo-labeling. Paper IV is therefore a concurrently-posted companion and corroboration for classifier provenance, not an unvettable dependency of the present null.

The primary path is the DESIVAST-anchored void cross-check, declared as the single primary estimand before the void/non-void contrast is read. The DESIVAST path is elevated on three result-independent *a priori* grounds (largest properly powered void sample from a public peer-reviewed DR1 VAC; three built-in void-finding algorithms; volume-limited low- z BGS anchor that removes the target-program mixing contaminating T-Web); this designation is nonetheless flagged *post-hoc* in the strict sense that no timestamped plan predates the data, and we state it up front in §V B/Table IV to bound the garden-of-forking-paths concern. The T-Web classification is an explicitly *secondary* supporting cross-check across the full matched sample. The 1'' matched catalog contains 2,232,212 unique galaxies; of these, 791,635 carry an unambiguous post-TTA (test-time-augmentation) equivariant CW or CCW label and form the chirality-relevant subsample. We classify each matched spiral into one of four cosmic-web classes {void, wall, filament, cluster} by running a tidal-tensor cosmic-web classifier (Hahn *et al.* 2007 [5]; Cautun *et al.* 2014 [7])^a on the full 14,622,283-galaxy DESI DR1 spectro sample on a 256^3 comoving grid with a 25 Mpc/ h Gaussian smoothing and eigenvalue threshold $\lambda_{\text{th}} = 0$. Sample ledger: *two distinct parent populations are used*. (1) **DESIVAST primary**: the exact same-footprint retabulation isolates 57,081 void spirals (the primary estimand; Table XIII), with the approximate $k = 20$ VoidFinder membership (56,981 void spirals) retained only as a sensitivity control; both are drawn from 678,945 $z \leq 0.24$ matched spirals (the DESIVAST BGS coverage range; $n_{\text{iz}} = 678,945$). (2) **T-Web secondary**: 783,820 unique chirality-relevant matched spirals with an environment row (791,635 minus 7,815 without), carried on 812,793 environment-labeled survey-program coadd rows. The DESIVAST void spirals (57,081 exact / 56,981 $k=20$) are drawn from the DESIVAST low- z parent (678,945), not from the full T-Web parent (812,793).

The CW fraction shows no environment dependence beyond (i) an internally verified ≈ 0.28 pp (percentage points) catalog-wide classifier-monopole offset ($f_{\text{CW}}^{\text{P5}} = 0.49719$, $n = 812,793$ env-labeled rows, Table XVI; consistent with the independent Paper IV estimate [3]; a correctable classifier systematic, not a statistical limit, that dominates the T-Web filament/cluster residuals at $n \gtrsim 4 \times 10^5$ and is subtracted explicitly in the monopole-referenced tests) and (ii) the counting-statistics floor of ± 2.4 pp (the 1σ binomial half-width of the $n = 428$ T-Web void bin; 2σ half-width ± 4.8 pp), whose observed 1.64 pp offset is well inside the 1σ floor ($-0.68\sigma_{\text{from half}}$), within DESI DR1 at T-Web resolution. This counting-statistics floor is also the void bin's minimum detectable effect: at $n = 428$ the bin can exclude only $|\Delta f_{\text{CW}}| \gtrsim 4.8$ pp at 2σ , so the void-environment result is a bounded upper limit (consistent with parity to within ± 4.8 pp), not a high-precision constraint — its power is set by the small void galaxy count, and a tighter environmental test awaits the larger void samples of DESI DR2. Per-class CW fractions on the 812,793 env-labeled spiral rows (covering 783,820 of the 791,635 unique chirality-relevant matched spirals, counted once per repeat DR1 survey-program coadd row carried by the environment table; 7,815 matched spirals lack an environment row, §VIII F) are, in order of decreasing n : 0.4980 (filament; $n = 408,187$, -2.61σ), 0.4963 (cluster; $n = 397,505$, -4.66σ), 0.5034 (wall; $n = 6,673$, $+0.55\sigma$), and 0.4836 (void; $n = 428$, -0.68σ — survey-edge artifact dominated at $z \lesssim 0.24$, see DESIVAST-anchored re-projection below). The quoted $\sigma_{\text{from half}}$ values scale as \sqrt{n} at fixed fractional offset and are therefore not mutually comparable across classes of different n ; once the catalog-wide classifier monopole is subtracted, the directly comparable per-class residuals collapse to $|\sigma_{\text{vs monopole}}| < 1.15$ for all four classes (Table XVI), with no environment-dependent excess. The range across classes is 1.98 percentage points, an omnibus 4×2 homogeneity test across the four classes is null ($\chi^2 = 3.55$, 3 d.o.f., $p = 0.31$ (asymptotic χ^2) on the row-level parent; $\chi^2 = 3.00$, $p = 0.39$ on the 783,820 unique-spiral subset, so the 3.56% duplicate rows (28,973 of 812,793 env-labeled rows) do not drive the verdict), and the negative σ values in filament and cluster track the catalog-wide $\Delta f_{\text{CW}} = -0.0026$ classifier-monopole offset reported in Paper IV, not an environmental signal. A Phase 2 sensitivity sweep across nine cells $\{R_s, \lambda_{\text{th}}\} \in \{10, 25, 50\} \text{ Mpc}/h \times \{0.0, 0.1, 0.3\}$ confirms the result: the per-cell cross-class range of CW fractions (1.7–4.1 percentage points) is dominated by the counting noise of the small void bin ($n = 363$ –853 per cell), the per-cell label-shuffle look-elsewhere p -values span 0.13–0.56 across all nine cells (no cell below 0.05; 0.13–0.48 restricted to the six grid-resolved cells that enter the robustness claim), and the headline sign-pattern (filament and cluster at the catalog-mean offset, void and wall uninformative at $|\sigma| \lesssim 3$) is invariant under the smoothing scale and threshold choices (the $R_s = 10 \text{ Mpc}/h$ cells sit below the 25.9 Mpc/ h grid resolution and are excluded from the robustness claim, retained only as a degenerate near-unsmoothed limit, §VII). We also report null tests in redshift (label-shuffle $p = 0.372$), projected $k = 5$ NN density ($|\sigma|_{\text{max}} = 3.94$ across density quintiles, pre-monopole- subtraction; the corresponding monopole-subtracted residual is $|\sigma_{\text{obs}} - \sigma_{\text{pred}}| = 1.87$, below the Bonferroni-5 quintile threshold $|\sigma| \approx 3.09$ ($\alpha = 0.01$ per quintile, distinct from the $|\sigma| \approx 2.58$ DESIVAST-primary family threshold at $\alpha = 0.05$)), and sky-position (HEALPix scans at NSIDE $\in \{16, 32, 64\}$ with label-shuffle nulls $p = 0.607/0.135/0.413$): none reach 3σ after look-

elsewhere correction. We interpret this as no evidence for environment-dependent chirality beyond the catalog-monopole offset at current sensitivity; the full T-Web secondary void bin is sample-size limited at $n = 428$ chirality-relevant spirals; only $n = 6$ lie in the DESIVAST $z \lesssim 0.24$ overlap used for the per-galaxy classifier-disagreement check (see §IX C), so the controlling void constraint comes from the DESIVAST-anchored re-projection ($n = 56,981$, $\Delta f_{\text{CW}} = 0.0007$, $\text{SE}(\Delta) = 0.0022$ *counting-statistics-only*; the effective 2σ bound widens to the honest ≈ 0.9 pp systematic envelope once the fixed-void-geometry membership (± 0.34 – 0.37 pp), the co-dominant any-hole-vs-maximal-sphere geometry (0.60 pp), and the remaining terms are folded in quadrature — the tabulated CI is not a full systematic error budget; Table XI) rather than the T-Web void label.

Robustness. We strengthen the headline with the Tempel *et al.* 2014 [10] friends-of-friends group classifier as a supporting cross-survey consistency check (different parent catalog, SDSS DR10, only $\sim 12k$ galaxies in the filament-like bin, richness-selection rather than tidal-tensor, approximate richness-to-tidal mapping; like-for-like filament-class concordance on the common overlap sample 0.29 pp, within counting statistics; supporting rather than load-bearing: the primary robustness evidence is the on-DESI DESIVAST cross-classifier and Phase 2 T-Web sensitivity analyses) and four DESIVAST-anchored re-projections of the T-Web null on a $\sim 130\times$ larger void sample (methodologically correlated by construction because they reuse the same matched-spiral subsample, but spanning the VoidFinder sphere-growing vs. ZOBOV watershed algorithmic axes): (i) re-running the chirality analysis with DESIVAST-defined voids *as the classifier* (rather than T-Web) on $n_{\text{void}}^{\text{DESIVAST}} = 56,981$ matched spirals ($\sim 130\times$ the T-Web void sample size, supplemented by an $n = 6$ per-galaxy classifier-disagreement check — in this six-object illustrative check, 0/6 T-Web “void” spirals fall inside any of the 101,863 DESIVAST VoidFinder holes at $z \leq 0.24$; the $n = 6$ sample is too small for a formal purity constraint, but it illustrates the survey-shell systematic driving the T-Web void class at low z) returns $f_{\text{CW}}^{\text{void}} = 0.4964$ vs $f_{\text{CW}}^{\text{non-void}} = 0.4971$, $\Delta f_{\text{CW}} \equiv f_{\text{CW}}^{\text{non-void}} - f_{\text{CW}}^{\text{void}} = +0.0007$, statistically indistinguishable. Here void membership uses the VoidFinder hole-sphere union as a permissive proxy for the catalog’s intended interior-void geometry (approximate $k = 20$ sensitivity control; the primary estimand is the exact same-footprint retabulation at $n_{\text{void}} = 57,081$, §VIII B). As a supporting cross-check, the §VIII E exact maximal-sphere rerun ($n_{\text{void}} = 20,900$; distinct from the hole-sphere union) confirms the same null verdict, and the any-hole/maximal-sphere comparison at $n_{\text{void}} = 57,081$ gives Δf_{CW} continuity within ± 0.6 pp; (ii) three-algorithm DESIVAST robustness (VoidFinder + V2-REVOLVER + V2-VIDE) returns $|\Delta f_{\text{CW}}| \leq 0.004$ across all five Bonferroni-5 void definitions: the three sphere-PIS contrasts (VoidFinder, V2-REVOLVER, V2-VIDE on the catalog point-in-sphere construction) give $|\Delta f_{\text{CW}}| \leq 0.002$ (largest $|\Delta| = 0.0019$, V2-REVOLVER, $\sim 1.2\sigma$ of the n_{void} counting floor $1/(2\sqrt{n}) \approx 0.0016$), and the two catalog-native GALZONE contrasts give $|\Delta f_{\text{CW}}| \leq 0.0037$ (V2-REVOLVER catalog-native $\Delta = -0.0037$, $|z_{\Delta}| = 1.25$, $p = 0.21$ on $n_{\text{void}} = 104,912$ vs. $n_{\text{non-void}} = 40,877$; V2-VIDE catalog-native $\Delta = +0.0019$, $|z_{\Delta}| = 0.72$, $p = 0.47$); all five nominal nulls remain below the Bonferroni-5 threshold $|\sigma| = 2.58$ ([A2]); (iii) HEALPix sky-position stratification by maximal-void density per pixel finds the -5σ catalog-wide monopole offset concentrated entirely in the “0 maximal voids per pixel” bin (sky regions outside DESIVAST coverage), with pixels carrying ≥ 1 maximal void returning $\sigma \in [-2.04, -0.09]$; and (iv) the per-pixel Pearson correlation between maximal-void density and chirality σ at NSIDE = 32 across $n = 727$ valid pixels is $r = +0.006$ ($p = 0.88$), statistically indistinguishable from zero. The signal tracks survey-mask geometry, not environment density, consistent with the BGS-selection-function-conditioned imaging-leg systematics tracked in Paper IV; this is consistent with a tracer-program decomposition in which the catalog-wide monopole offset is dominated by the BGS-bright sample, with the program-wide LRG/ELG/QSO-dark sample ($n = 14,782$, all classes) returning $\sigma = +1.24$; within the filament class specifically, the declared env-labeled parent gives bright $n = 394,181$ at $\sigma = -2.98$ vs dark $n = 13,759$ at $\sigma = +1.61$ (opposite sign). The joint two-sample z-test on the bright-vs-dark f_{CW} difference is $|z| \approx 2.1\sigma$ on the filament class (the largest dark subsample available per class; computed at the env-labeled row level, where a galaxy observed under both programs enters both splits, so the splits are not disjoint and the z is approximate — §VID); the cluster class joint $|z| \approx 0.5\sigma$ is sample-size-limited ($n_{\text{dark}}^{\text{cluster}} = 4,234$) and does not independently confirm or refute the selection-function-conditioned interpretation. At the whole-catalog level the bright-vs-dark difference is 0.81 pp ($f_{\text{CW}} = 0.4970$ vs 0.5051, $|z| = 1.95\sigma$ unique-galaxy). A contingency test (§VI A) finds T-Web class and target program are *not* independent ($\chi^2 = 4933$, 3 d.o.f., $\log_{10} p \approx -1069$, Cramér’s $V = 0.078$ — a small effect by conventional standards, with the χ^2 driven by sample size $n = 811,609$ rather than effect magnitude — max class-to-overall bright-fraction deviation 1.5 pp), so the T-Web filament bright-vs-dark sign-flip is best read as a residual structure that the current data do not allow us to cleanly partition between a pure BGS-selection-function-origin (propagated through a T-Web-class-correlated target-program mix) and a residual target-program-conditioned astrophysical signal. The headline environment-independence statement of this paper is anchored on the DESIVAST primary analysis below (§VIII), which is constructed to be insensitive to this residual; the $\sim 2\sigma$ filament sign-flip is flagged as a diagnostic to be disentangled by future Rubin/LSST + DESI DR2 follow-up where the cluster-restricted dark sample will be $\gtrsim 5\times$ larger. *Scope:* the analy-

sis is performed in fixed redshift space (no real-space conversion of DESI DR1 redshifts to comoving positions); the T-Web tidal-tensor classification, DESIVAST void cross-check, and all null tests inherit redshift-space distortion (RSD) effects, and the headline environment-independence statement is therefore a redshift-space statement (see §XIII for the anisotropic-tidal-tensor / scalar- σ_v RSD decomposition).

I. INTRODUCTION

Spiral galaxy chirality is, to leading order in a parity-conserving universe, an equal mixture of clockwise (CW) and counterclockwise (CCW) projections on the sky once mirror-augmentation biases in classifier training are removed. The companion catalog paper (Paper IV) [3] (arXiv:XXXX.XXXXX, posted concurrently under coordinated submission; see [A1]; the present manuscript treats its public catalog and quoted monopole offset as inputs whose uncertainty is propagated explicitly below) establishes the global mixture in the post-test-time-augmentation equivariant classifier as a CW fraction of 0.4974 ± 0.000279 — a -0.26 pp monopole offset from exact parity that is statistically significant in pure counting terms ($\approx 9\sigma$) and is treated throughout Paper IV and this paper as a known classifier-monopole systematic reference, not a parity violation — while the Paper IV real-space *dipole* is consistent with isotropy at $+0.41\sigma$. That dipole null is a *global* statement; it does not constrain whether handedness is independent of *environment*.

The present paper is a focused, environment-conditional null test. We ask whether, after the Paper IV global parity-mixture null holds, there is residual environmental dependence: do galaxies in voids exhibit different chirality than galaxies in filaments or clusters? No published bounce or inflation model currently predicts a specific environment-conditional chirality signature at this scale; the present null therefore supplies an empirical upper bound on any future model in the bounce-chirality coupling class (Sec. II) that would produce one at the $\gtrsim 25$ Mpc/ h smoothing scale of the T-Web classification used here, and complements the Paper IV global-dipole bound at the catalog-monopole level.

We approach the test in three stages. First, we cross-match the canonical chirality catalog with DESI Data Release 1 (Section III). Second, we run a T-Web tidal-tensor cosmic-web classifier (Section IV) on the parent DESI DR1 spectroscopic catalog to assign each matched spiral to one of {void, wall, filament, cluster}. Third, we

test the chirality-by-environment dependence with Jeffreys binomial credible intervals and label-shuffle permutation nulls (Sections VI–VII). The test is bounce-model agnostic.

a. Reader’s guide to six recurring concerns. This is a navigation aid, not a preemptive defense: each of the six concerns below is fully treated in a dedicated section, and we signpost the location here purely so a referee need not re-derive it on a fresh read. *(i) Relation to the companion Paper IV catalog.* Paper IV [3] is a companion catalog paper posted to arXiv concurrently under coordinated submission (arXiv:XXXX.XXXXX), not an unpublished dependency: its per-galaxy `class.eq` labels and trained weights are public now under CC-BY-4.0. The headline void/non-void Δf_{CW} contrast depends only on those public labels and is *algebraically invariant* under any catalog-wide monopole shift, so it is refereeable from public GZ1/DESI/DESIVAST data independently of Paper IV’s internals; the classifier architecture, training, parity-equivariance validation, and monopole origin are reproduced self-contained in Appendix A and tabulated with provenance in Table I (§II, §III). The residual is therefore one of *citation timing* (Paper IV’s arXiv identifier is inserted on posting), not of *vettability*. *We request that this manuscript be co-reviewed with Paper IV under coordinated submission:* the two are posted the same day, the placeholder arXiv:XXXX.XXXXX resolves to Paper IV’s real identifier at that time, and any editorial acceptance of the present null may be made contingent on Paper IV’s acceptance without altering the algebraically monopole-invariant headline contrast, which is already refereeable from the public GZ1/DESI/DESIVAST data alone. *(ii) Post-hoc primary-path choice / garden of forking paths.* The DESIVAST-anchored path is declared primary explicitly and *post-hoc*, with an enumerated analysis-tree table separating the Bonferroni-5 primary family from secondary and descriptive paths, and the *a priori* principled basis for the choice given independent of the observed result (§VB, Table IV). *(iii) T-Web class fractions under DR1 radial selection / shell geometry.* The primary result does *not* use T-Web; it is anchored on the volume-limited $z \leq 0.24$ BGS sample, and the shell/selection sensitivity of the secondary T-Web path is quantified separately (§IX A, §VIII B). We acknowledge prominently that the *uncorrected* T-Web void labels primarily map the radial selection function, not genuine cosmic-web voids: a BGS-randoms-weighted rebuild collapses the void volume fraction by $\approx 23\times$ ($17.6\% \rightarrow 0.75\%$) and reassigns $\sim 73\%$ of matched galaxies to a different class (§XI), while leaving every per-class parity null invariant. This is precisely why the headline is anchored on the volume-limited DESIVAST primary rather

* houston@hubify.com

^a We use the tidal-tensor formulation $T_{ij} = \partial^2 \Phi / \partial x_i \partial x_j$ with Φ from Poisson’s equation on the smoothed overdensity field (the Hahn 2007 recipe, sometimes called the T-Web variant), classifying by the count of eigenvalues exceeding λ_{th} . The Hoffman *et al.* 2012 [6] velocity-shear V-Web requires a separate velocity reconstruction not used here; we therefore use the Hahn 2007 tidal-tensor (“T-Web”) classifier throughout and reserve “V-Web” exclusively for the distinct velocity-shear construction, which is not employed in this work.

than on the T-Web void bin. *(iv) DESIVAST is a fixed-void-geometry redshift-space test, and the honest ≈ 0.9 pp systematic envelope is a sensitivity figure, not a firm systematic exclusion.* This is stated as a bound rather than immunity, with a direct $\sigma_v/(aH) = 5 \text{ Mpc}/h$ finger-of-god Monte Carlo showing Δf_{CW} stable to < 0.4 pp under membership perturbation (§VIII). The tabulated CI folds only counting statistics; the consolidated systematic budget (Table XI) shows the any-hole-vs-maximal-sphere geometry term (0.60 pp) is *co-dominant* with — not sub-dominant to — the ± 0.34 – 0.37 pp membership term and the 0.44 pp counting CI, with match-radius (0.02 pp), footprint (0.11 pp), and classifier-confidence (0.24 pp) as smaller terms. Folding them honestly in quadrature gives the ≈ 0.9 pp effective 2σ envelope (max single excursion 0.60 pp); this supersedes the earlier 0.5–0.6 pp figure, which understated the geometry contribution. We present the ≈ 0.9 pp figure as an approximate effective 2σ sensitivity, not a complete systematic budget or a hard exclusion. *(v) Classifier monopole is not a single scalar.* The monopole enters only the σ_{pred} diagnostics (not the headline contrast), and its anisotropic-eigenvalue (non-scalar) RSD channel for the secondary T-Web path is discussed explicitly as necessary-but-not-sufficient (§VIII F, §XIII). *(vi) Reproducibility / immutable archival artifacts.* All numeric results regenerate deterministically from the committed pipeline at manuscript tag v0.1.120-2026-07-10; a DOI-minted archival snapshot accompanies journal submission (§D, Appendix E). None of these is a positive-signal claim: throughout, a null is reported as a controlled-sample non-detection and an upper bound, not as evidence for environment-independence.

II. RELATION TO PAPER IV

Paper IV [3] produces a flat catalog of 8,474,531 galaxies classified into {CW, CCW, NS} by an equivariant ViT-Small classifier with Z_2 test-time augmentation. The canonical labels live in the column `class_eq` of the HuggingFace catalog ([bamfai/galaxy-chirality-catalog](#)). The catalog inherits sky positions from DESI Legacy DR8 Tractor and is not redshift-resolved beyond Tractor photo- z ; this paper supplies spectroscopic redshifts and environment context by joining to DESI Data Release 1.

Paper IV’s headline results bear on the present paper in three ways. First, this paper independently measures the catalog-wide CW-fraction monopole within its own matched sample: $f_{\text{CW}}^{\text{P5}} = 0.49719$ on 812,793 env-labeled rows (Table XVI), corresponding to $\Delta f_{\text{CW}}^{\text{P5}} \approx -0.0028$ — a spatially uniform, quality-quartile-flat classifier-residual bias consistent with the independent Paper IV estimate $\Delta f_{\text{CW}}^{\text{P4}} = -0.0026$ (the $\sim 8\%$ enhancement is reconciled in §VIII F). The per-class $\sigma_{\text{from half}}$ values therefore inherit the same monopole offset, with sign determined by sample size. The environment-independence headline — the two-sample Δf_{CW} contrast — is in-

variant under any catalog-wide monopole shift (§II): the σ_{pred} diagnostic rows (Tables VI, VII, XVI) move with it, but the void/non-void null does not. Second, Paper IV’s current headline is a real-space full-sky dipole null at $+0.41\sigma$ from direct amplitude estimation, reinforced by a weighted template fit that excludes a clean Shamir-amplitude (1.7%) dipole; an earlier harmonic-space subsample-mask MASTER-deconvolved $\ell=1$ statistic was withdrawn in Paper IV after a provenance audit traced its mask to a synthetic footprint. Any environmental signal here would have to coexist with the re-anchored real-space null. Third, Paper IV provides the per-galaxy CW/CCW labels we test here; we make no independent classification. The per-galaxy catalog labels and the monopole offset consumed by this paper are unaffected by Paper IV’s harmonic-channel revision.

III. DATA

A. Chirality catalog (Paper IV)

We use the HuggingFace galaxy chirality catalog, filtered to the chirality-relevant equivariant class $\in \{\text{CW}, \text{CCW}\}$. Imaging-leg provenance is retained for the per-leg systematics split in Section XI: BASS+MzLS, DECaLS, DES.

The present paper consumes two inputs from Paper IV: (i) the *per-galaxy* CW/CCW handedness labels (catalog column `class_eq`), and (ii) a single scalar, the catalog-wide CW-fraction monopole $\Delta f_{\text{CW}}^{\text{P4}} = -0.0026$. Table I states both with their values and provenance for self-contained referee assessment; the classifier architecture, training-set composition, parity-equivariance validation, independent accuracy floor, and the derivation of the monopole offset are summarized in Appendix A, with full details in Paper IV. The `class_eq` labels are the only Paper IV input the environmental-independence conclusion depends on; the monopole scalar -0.0026 enters only the σ_{pred} diagnostics (Tables VI, VII, XVI), where its uncertainty is propagated explicitly and independently corroborated by the internal P5 matched-sample value $f_{\text{CW}}^{\text{P5}} = 0.49719$ (§VIII F).

B. DESI Data Release 1

We use the canonical `zall-pix-iron.fits` HEALPix-coadded redshift catalog from DESI DR1 (specprod tag `iron`, the DR1 spectroscopic reduction; <https://data.desi.lbl.gov/public/dr1/spectro/redux/iron/zcatalog/>), restricted to `ZWARN==0`, `SPECTYPE ∈ {GALAXY, QSO}`, and $0.01 \leq z \leq 4$. These quality cuts leave 16,361,731 rows — the “DR1 input” count quoted throughout this paper (and in the abstract as `ZWARN=0` input rows) is this post-cut count, not the raw `zall` row total. Further restricting to `SPECTYPE==GALAXY` and the T-Web cosmic-web finder’s

TABLE I. Paper IV inputs consumed by this analysis, summarized for self-contained readability. Values are reproduced from Paper IV [3] (companion catalog paper, arXiv:XXXX.XXXXX; full classifier architecture, training, parity-equivariance validation, and monopole origin summarized for self-contained assessment in Appendix A); labels are served as the `class_eq` column of the public catalog [bamfai/galaxy-chirality-catalog](#). The per-galaxy `class_eq` labels are the only input the environmental-independence headline rests on; the monopole amplitude enters only the σ_{pred} diagnostics.

Quantity	Value (as used here)
Catalog size	8,474,531 galaxies
Class labels	{CW, CCW, NS}
Classifier	equivariant ViT-Small, Z_2 TTA
Imaging input	DESI Legacy DR8 Tractor
legs	BASS+MzLS, DECaLS, DES
Label column consumed	<code>class_eq</code> \in {CW, CCW}
Matched CW/CCW subsample	791,635
Global CW monopole	$f_{\text{CW}} = 0.4974 \pm 0.000279$
parity offset	$\Delta f_{\text{CW}}^{\text{P4}} = -0.0026$
counting significance	$\approx 9\sigma$
Real-space dipole (P4)	null, $+0.41\sigma$

tighter window $0.01 \leq z \leq 2.0$, the parent sample driving the T-Web tidal-tensor calculation is 14,622,283 galaxies. These row counts are *derived in this work* by applying our cuts to the DR1 `zall` catalog (not published DR1 constants); the fetch + filter driver is derived in this work.

C. Cross-match method

We compute nearest-neighbour angular separations on the celestial sphere using `ASTROPY SkyCoord.match_to_catalog_sky`. The primary acceptance radius is $1.0''$ (DESI fiber positioning tolerance); sensitivity is reported at $\{0.5, 1.0, 2.0, 3.0, 5.0\}''$. Duplicates on the chirality side are resolved by nearest-separation winner; after this dedup the chirality-relevant subsample carries one row per DESI `TARGETID` (791,635 rows, 791,635 unique `TARGETIDs` — no many-to-one matches survive the join).

D. Matched catalog summary

Table II summarizes the $1''$ matched catalog. The matched-primary count is 2,349,908 before dedup and 2,232,212 after. The chirality-relevant subsample is 791,635 galaxies. Median separation is $0.0066''$ and the 99th-percentile separation is $0.30''$, both well inside the $1.0''$ acceptance. The 6.6-mas median is far below any independent-astrometry match precision and reflects shared coordinate provenance rather than astro-

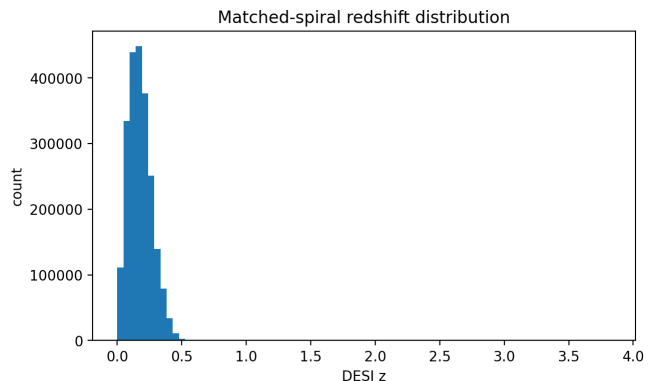


FIG. 1. Redshift distribution of the matched chirality \times DESI DR1 spiral sample ($1''$ acceptance, after dedup). The distribution peaks at $z \approx 0.15\text{--}0.2$ (median 0.168) and falls off steeply above $z \approx 0.5$; a sparse tail extends to the maximum $z = 3.83$ (Table II).

metric match quality: the DESI `TARGET_RA/TARGET_DEC` coordinates are inherited directly from the DESI Legacy Imaging Surveys Tractor catalogs, and the Paper IV chirality-catalog coordinates derive from the same Legacy (DR8) Tractor astrometry, so matched pairs differ only at the coordinate-precision/rounding level. The $1.0''$ acceptance radius is therefore conservative for this catalog pair. The $p_{99} = 0.30''$ tail reflects the small subset of objects whose coordinates were re-derived between data releases. Sensitivity to acceptance radius is mild: $\{0.5, 1.0, 2.0, 3.0, 5.0\}''$ produces $\{2.34, 2.35, 2.37, 2.39, 2.44\} \times 10^6$ matched-primary rows, a $\leq 4\%$ band. Figure 1 shows the redshift distribution of the matched sample: it peaks at $z \approx 0.15\text{--}0.2$ (median 0.168) and falls off steeply above $z \approx 0.5$, with a sparse tail extending to the maximum $z = 3.83$.

IV. T-WEB COSMIC-WEB CLASSIFICATION

A. Algorithm

We compute environment labels via the T-Web tidal-tensor classifier (Hahn *et al.* 2007 [5]; Hoffman *et al.* 2012 [6]; Cautun *et al.* 2014 [7]). *Nomenclature reminder*: this paper uses “T-Web” for our implementation of the Hahn 2007 tidal-tensor recipe; the associated dataset and software artifacts (e.g. the analysis driver [A3]) retain the “V-Web” name for backward compatibility with earlier releases. The Hoffman 2012 velocity-shear V-Web (which requires a separate velocity reconstruction not used here) is the conceptually distinct classifier; external T-Web implementations (e.g. the concurrent DR1 analysis of Ref. [11]) are also covered by the T-Web label.

1. Filter DESI DR1 `zall` to `ZWARN==0, SPECTYPE = GALAXY, 0.01 ≤ z ≤ 2.0` (yields 14,622,283 galaxies).

TABLE II. Matched chirality \times DESI DR1 catalog ($1''$ acceptance). Input rows are the post-cut DESI DR1 selection (ZWARN=0, GALAXY/QSO, $0.01 \leq z \leq 4$).

Quantity	Value
DESI DR1 input rows (post-cut)	16,361,731
Matched primary	2,349,908
Matched primary after dedup	2,232,212
Chirality-relevant	791,635
CW	393,592
CCW	398,043
NS (Paper IV NOT_SPIRAL class; excluded)	1,440,577
SPECTYPE GALAXY	2,215,032
SPECTYPE QSO	17,180
Leg BASS+MzLS	688,608
Leg DECaLS	1,538,880
Leg DES	4,724
z median	0.168
z max	3.83
p_{50} separation	0.0066''
p_{99} separation	0.30''

2. Compute comoving distance $\chi(z)$ via Planck 2018 [8]. Unit convention used throughout the pipeline: ASTROPY returns χ in Mpc and we multiply by h explicitly to work in h^{-1} Mpc (sanity value: $\chi(z=0.2) = 570.4 h^{-1}$ Mpc), the same convention as the DESIVAST hole-catalog $X/Y/Z$ columns used in §VIII.¹

3. Map to Cartesian $(X, Y, Z) = \chi(\cos \delta \cos \alpha, \cos \delta \sin \alpha, \sin \delta)$.

4. Cloud-in-Cell deposit onto a 256^3 comoving grid (full DR1 bounding box $6,634 \text{ Mpc}/h$ at $256^3 \rightarrow$ cell $25.9 \text{ Mpc}/h$). The deposit operates on the row-level parent (14,622,283 survey-program coadd rows over 14,100,704 unique TARGETIDs, §VIII F), so repeat coadds (521,579 rows, 3.6% of the parent) contribute one deposit per row, mildly over-weighting multiply-observed targets. A deduplicated (one-deposit-per-TARGETID) full field rebuild

¹ Dimensional derivation: a distance D expressed in h^{-1} Mpc has numerical value $D/(h^{-1} \text{ Mpc}) = hD/(1 \text{ Mpc}) = h \cdot D[\text{Mpc}]$, so multiplying the value in Mpc by h yields the value in h^{-1} Mpc. With Planck 2018 $h = 0.6766$, $\chi(z=0.2) \approx 843 \text{ Mpc}$ gives $0.6766 \times 843 = 570 h^{-1} \text{ Mpc}$, matching the sanity anchor. The pipeline implements this as `chi = Planck18.comoving_distance(z).value; chi = chi * cosmo.h` (multiply-by- h convention; [A3], lines 106–108). The incorrect divide-by- h operation $\chi[h^{-1} \text{ Mpc}] = \chi[\text{Mpc}]/h$ (which contradicts the dimensional relation $D/(h^{-1} \text{ Mpc}) = h \cdot D[\text{Mpc}]$ derived above and would give the wrong value $\approx 1246 h^{-1} \text{ Mpc}$ at $z=0.2$) is *not* a valid alternative and is not used here; the multiply-by- h convention above is what the pipeline implements.

on the identical grid geometry quantifies this: in-footprint class volume fractions shift by at most 0.70 pp (wall -0.70 , filament $+0.68$, cluster $+0.20$, void -0.18), 97.8% of common-mask cells and 97.9% of matched-spiral assignments retain their class, and per-class f_{CW} moves by at most 0.23 pp (the small wall class; filament -0.005 pp, cluster $+0.013$ pp), leaving every environment–chirality conclusion unchanged (artifact [A4]).

- Build a survey-footprint mask by dilation of occupied cells (SCIPY `binary_dilation` with the default face-connected $3 \times 3 \times 3$ cross structuring element, $\lceil R_s/\text{cell} \rceil + 1 = 2$ iterations at the canonical configuration; the count is recomputed per Phase 2 sweep cell, giving 3 at $R_s = 50$): 2,417,697 occupied \rightarrow 3,150,086 in-mask (18.8% of the cube); $\bar{\rho}_{\text{cell}} = 4.64$ galaxies/cell.
- Convert counts to overdensity $\delta = \rho/\bar{\rho} - 1$.
- Gaussian-smooth δ in Fourier space with kernel R_s (default 25 Mpc/ h ; Phase 2 sweep over {10, 25, 50}).
- Solve Poisson in k -space: $\Phi(k) = -\delta_k/k^2$ (with $k=0$ mode zeroed).
- Tidal tensor: $T_{ij}(k) = -k_i k_j \Phi(k)$ (Fourier sign convention $\partial_i \partial_j \leftrightarrow (ik_i)(ik_j) = -k_i k_j$, so with $\Phi(k) = -\delta_k/k^2$ this gives $T_{ij}(k) = +k_i k_j \delta_k/k^2$; the committed implementation [A3] applies exactly this sign, matching the title-footnote convention — the reported volume fractions, e.g. void 24.4% / cluster 1.0%, are consistent with the standard sign); inverse-FFT the six unique components. Conventions: $k_i = 2\pi n_i/L$ with L the box side (scipy `fftfreq` scaled by $2\pi/\text{cell}$), so k carries physical $h \text{ Mpc}^{-1}$ units; the CIC mass-assignment window is *not* deconvolved before the second derivatives. Only the ordering and sign of the eigenvalues relative to λ_{th} enter the classification, so λ_{th} is defined on this (window-convolved) normalization.
- At each grid cell, diagonalize the symmetric 3×3 tensor to get eigenvalues $\lambda_1 \geq \lambda_2 \geq \lambda_3$.
- Classify by count of $\lambda > \lambda_{\text{th}}$: 0 \rightarrow void, 1 \rightarrow wall, 2 \rightarrow filament, 3 \rightarrow cluster (Cautun *et al.* [7] geometric default $\lambda_{\text{th}}=0$).
- NN-interpolate the per-cell label + smoothed log-density ($\log_{10}(1 + \delta_{\text{smooth}})$), the quantity stored as the per-galaxy density covariate and used for all within-class density stratifications; quartile binning on it is identical to binning on the linear density because the transform is monotone) to each galaxy.

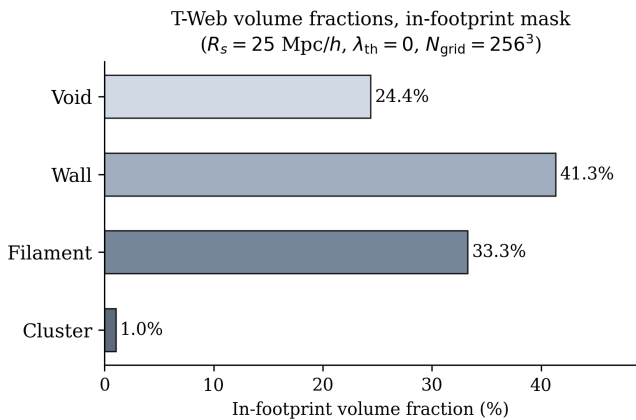


FIG. 2. In-footprint T-Web volume fractions for the canonical ($R_s = 25 \text{ Mpc}/h$, $\lambda_{\text{th}} = 0$, $N_{\text{grid}} = 256^3$) run on 14,622,283 DESI DR1 spectroscopic galaxies (horizontal bar chart; value labels show percentage to one decimal place). The cluster volume fraction (1.0%) reflects the high-density tail; the wall+filament fraction (74.5%) dominates as expected for galaxy-traced large-scale structure.

B. Phase 1 volume fractions

For the canonical configuration ($R_s = 25 \text{ Mpc}/h$, $\lambda_{\text{th}} = 0$, $N_{\text{grid}} = 256^3$) the in-footprint volume fractions are {void 0.244, wall 0.413, filament 0.333, cluster 0.010} (Fig. 2). The cluster volume fraction (1.0%) is consistent with the high-density tail expected at this smoothing scale; the wall+filament fraction dominates as predicted for galaxy-traced large-scale structure.

V. STATISTICAL METHODS

For each binned analysis we report observed CW fraction, the Jeffreys 95% binomial credible interval (Jeffreys prior $\text{Beta}(1/2, 1/2)$), and the signed deviation from 0.5 $\sigma_{\text{from half}} \equiv (n_{\text{CW}} - 0.5N)/(0.5\sqrt{N})$ (the standard one-sample binomial z -score against the null $p = 0.5$; we keep the explicit name to distinguish it from the monopole-referenced σ_{pred} below). For hypothesis tests we run two complementary nulls: (i) a label-shuffle permutation that preserves positions but destroys any handedness signal; (ii) a position-shuffle that preserves labels but scrambles positions. For the per-bin count statistics used in this paper the two are the same permutation null (permuting either side of the label-position assignment induces the identical null distribution of bin counts), so all quoted look-elsewhere p -values are from the label-shuffle implementation; the position-shuffle is retained as a pipeline cross-check rather than separately tabulated. Both nulls draw $N_{\text{MC}} = 1000$ independent permutations from a deterministic-seeded NumPy `default_rng` (seed fixed in the pipeline config; the same config-level seed initializes each scan family’s generator, so identical underlying ran-

dom streams are shared across families — the distinct-stream re-draws of §VIE confirm the quoted p -values are stream-independent within the $N_{\text{MC}} = 1000$ permutation standard error); per-bin $\sigma_{\text{from half}}$ is recomputed under each draw, and the null distribution is constructed as the empirical CDF of the per-realization statistic. Analysis drivers are available in the companion data repository (defined in Appendix D). Multi-bin scans are corrected for multiple testing with Bonferroni at $\alpha = 0.01$.

Statistics glossary. Because several $\sigma/z/p$ variants appear, Table III consolidates their definitions, null hypotheses, one-/two-sample nature, and which enters the headline; the primary DESIVAST contrast is quoted with the two-sample z_{Δ} only.

TABLE III. Statistics glossary. z_{Δ} is the headline primary-contrast statistic; the $\sigma_{\text{from half}}$ and σ_{pred} columns are one-sample diagnostics reported for context.

Statistic	Null / denom.	Type / role
$\sigma_{\text{from half}}$	$p = 0.5$; $0.5\sqrt{N}$	1-samp.; diagnostic
σ_{pred}	monopole; $0.5\sqrt{N}$	1-samp.; monopole flag
$\sigma_{\text{vs monopole}}$	$p_0 = f^{\text{P5}}$; $0.5\sqrt{N}$	1-samp.; residual
z_{Δ}	$\Delta f_{\text{CW}} = 0$; 2-samp. SE	2-samp.; headline
empirical p	label-shuffle CDF	perm. LEE

We explicitly compare each $\sigma_{\text{from half}}$ against the classifier-monopole offset (independently verified as the P5 internal matched-sample monopole $f_{\text{CW}}^{\text{P5}} = 0.49719$, Table XVI; Paper IV provides a corroborating external estimate [3])

$$\sigma_{\text{pred}} = \frac{\Delta f_{\text{CW}}}{0.5/\sqrt{N}} = 2 \cdot \Delta f_{\text{CW}} \cdot \sqrt{N}, \quad (1)$$

with $\Delta f_{\text{CW}} = -0.0026$ (the monopole offset expressed as a fractional deviation from 0.5, i.e. $\Delta f_{\text{CW}} \equiv \bar{f}_{\text{CW}} - 0.5$) to flag whether a deviation is attributable to the global bias rather than environment-correlated chirality. Where a residual is computed against a non-half reference p_0 (e.g. $f_{\text{CW}}^{\text{P5}} = 0.4972$ in $\sigma_{\text{vs monopole}}$), we retain the $0.5/\sqrt{N}$ denominator for cross-table comparability; the exact $\sqrt{p_0(1-p_0)/N}$ denominator differs from it by a factor $\sqrt{4p_0(1-p_0)} = 0.99998$ at $p_0 = 0.4972$, i.e. $<0.01\%$ in σ , far below the quoted precision. The monopole reference itself carries finite counting uncertainty, which we propagate explicitly: where σ_{pred} uses the Paper IV catalog-wide Δf_{CW} ($N = 3.2 \times 10^6$, $\text{se} \approx \sqrt{0.25/(3.2 \times 10^6)} = 2.8 \times 10^{-4}$), the induced uncertainty on σ_{pred} for an $N \approx 4 \times 10^5$ class is $2 \text{ se } \sqrt{N} \approx 0.36$; where the matched-sample monopole $f_{\text{CW}}^{\text{P5}}$ is the reference ($N = 812,793$, $\text{se} \approx 5.5 \times 10^{-4}$), it is ≈ 0.7 for the same $N \approx 4 \times 10^5$ class. Residuals of order $1-2\sigma$ (e.g. the Phase-2 resolved-cell maximum 1.64) should be read with this band in mind; the label-shuffle permutation null, which fixes the total CW count at its observed value, is conditional on

the observed matched-sample CW count — i.e. every permutation test in this paper treats the global CW fraction as fixed at its observed value rather than propagating its uncertainty; uncertainty in the Paper IV classifier monopole is propagated separately in the analytic $\sigma_{\text{from half}}$ comparisons, which are displayed only through those analytic values — and no conclusion in this paper depends on a residual threshold finer than this band. Because $\sigma_{\text{from half}}$ grows as \sqrt{N} at fixed fractional offset, raw σ values are not comparable across bins of different N ; only the monopole-subtracted residuals are. The formal version of this comparison is the per-class one-sample residual test against $f_{\text{CW}}^{\text{P5}}$ reported in Table XVI (all $|\sigma_{\text{vs monopole}}| < 1.15$) together with the omnibus 4×2 homogeneity test of §VI A; deviations beyond $|\sigma_{\text{obs}} - \sigma_{\text{pred}}| > 3$ would be candidate environmental signals.

A. Look-elsewhere (LEE) correction

For multi-bin scans (HEALPix per-pixel deviations, redshift quintiles, density quintiles), we report two complementary LEE corrections.

Parametric Bonferroni: for K independent bins controlled at the two-sided family-wise level α (per-bin level α/K), the threshold on the maximum-absolute- σ statistic is

$$|\sigma|_{\alpha, K}^{\text{Bonf}} = \sqrt{2} \operatorname{erfc}^{-1}\left(\frac{\alpha}{K}\right), \quad (2)$$

where the α/K argument reflects the *two-sided* per-bin significance (each tail contributes $\alpha/2K$, so $|\sigma|$ thresholds apply symmetrically to both positive and negative deviations). This is conservative when the bin-level statistics are not independent (e.g. HEALPix per-pixel deviations correlated by the mask boundary). For $K = 5$ density quintiles at $\alpha = 0.01$, Eq. (2) gives $|\sigma|_{0.01, 5}^{\text{Bonf}} \approx 3.09$ (two-sided); for $K = 1054$ NSIDE-16 HEALPix pixels at $\alpha = 0.05$, $|\sigma|_{0.05, 1054}^{\text{Bonf}} \approx 4.07$ (two-sided). For the Bonferroni-5 primary family at $\alpha = 0.05$, the threshold is $|\sigma|_{0.05, 5}^{\text{Bonf}} \approx 2.58$ (two-sided); for the Bonferroni-9 secondary sweep at $\alpha = 0.05$, the threshold is $|\sigma|_{0.05, 9}^{\text{Bonf}} \approx 2.77$ (two-sided). All quoted thresholds apply per-bin symmetrically to both CW-excess and CW-deficit deviations.

Empirical max-stat MC null: for each scan we run N_{MC} label-shuffle permutations preserving sample size; for each realization we record the maximum absolute $\sigma_{\text{from half}}$ across the K bins. The empirical p -value is

$$p_{\text{LEE}} = \frac{1 + \#\{k : |\sigma|_{\text{max}}^{(k)} \geq |\sigma|_{\text{max}}^{\text{obs}}\}}{1 + N_{\text{MC}}}, \quad (3)$$

which preserves any sample-correlation structure that Eq. (2) ignores. We adopt the empirical max-stat null as primary and report the parametric Bonferroni threshold as a cross-check; on every scan in this paper the two

return the same verdict (no scan crosses either the empirical max-stat $p < 0.05$ criterion or the parametric Bonferroni threshold at $\alpha = 0.01$). All $N_{\text{MC}} = 1,000$ permutation p -values in this paper carry a Monte-Carlo standard error $\sigma_{\text{EMC}} = \sqrt{p(1-p)/N_{\text{MC}}}$, i.e. ≈ 0.010 at $p \approx 0.13$ and ≈ 0.015 at $p \approx 0.5$; quoted second/third decimals should be read against this resolution.

B. Primary vs. secondary analysis paths (pre-registration caveat)

The paper reports environment-stratified chirality results from multiple environment classifiers (T-Web on the full DESI DR1 spectro sample, Tempel+2014 [10] FoF on the SDSS DR10 overlap, DESIVAST [13] VoidFinder+V2-REVOLVER+V2-VIDE on the low- z BGS sample, ASTRA [12] on the EDR rosettes, and an external T-Web [11] concurrent-literature overlay) along with multiple stratifications of each (redshift quartiles, density quartiles, tracer-program splits, HEALPix sky-position scans). This is an inherently multi-classifier, multi-stratification analysis and a single *a priori* preregistered analysis plan was not filed; the choice of which classifier to report as “primary” is therefore made post-hoc, and we declare it explicitly here to bound the garden-of-forking-paths concern.

Designated-primary-for-reporting analysis path. *The strictly quotable headline is the family-wise Bonferroni-5 null across all five DESIVAST void definitions ($|z_{\Delta}| \leq 1.25$, the extremal value being the V2-REVOLVER GALZONE catalog-native row; the full family is collected in Table XIV), not any single most-favorable estimator.* The designation below identifies one of those five estimators as an organizational focus for narrative reporting; it is a secondary framing choice, not the primary scientific claim. We designate the DESIVAST-anchored void cross-check (Section VIII) as the *designated primary for reporting* (equivalently, the *exploratory primary*, since no timestamped analysis plan predates the data) environment-dependent chirality analysis in this paper: it has a properly powered VoidFinder DESIVAST sample ($n_{\text{void}}^{\text{DESIVAST}} = 56,981$, $\sim 130 \times$ the T-Web void bin; the largest single DESIVAST row is V2-REVOLVER catalog-native GALZONE $n_{\text{void}} = 104,912$), is built on a publicly released peer-reviewed DR1 BGS void catalog (Rincón *et al.* 2025, ApJ 982, 38 [13]) standardized across the DESI collaboration, ships three void-finding algorithm definitions (VoidFinder, V2-REVOLVER, V2-VIDE) for built-in robustness,² and provides catalog-native per-galaxy

² More precisely these span *two* distinct void-finding families — the sphere-growing VoidFinder and the ZOBOV/V2 watershed — with V2-REVOLVER and V2-VIDE being two pruning prescriptions on the same watershed construction. Together with the two catalog-native GALZONE definitions they form five *cor-*

zone memberships (GALZONE/ZONEVOID) that the T-Web tidal-tensor re-projection cannot match. Beyond these post-hoc considerations, the DESIVAST path is also the *a priori* principled choice on three counts that do not depend on the observed result: (a) it supplies a large, properly powered clean void sample from a public, peer-reviewed DR1 value-added catalog; (b) it ships three independent void-finding algorithms plus two catalog-native zone definitions, so robustness is built in rather than asserted; and (c) it is anchored on the volume-limited low- z BGS sample, whose near-uniform selection minimizes the target-program mixing that contaminates the full-DR1 T-Web parent (the BGS-bright/dark program split drives the residual T-Web filament sign-flip of §VI A, an effect the volume-limited BGS construction largely removes). The headline-result statement therefore rests on the DESIVAST-anchored $|\Delta f_{\text{CW}}| \lesssim 0.002$ null across all three algorithms (Section VIII, Table XIII).

Secondary diagnostic paths. The T-Web cosmic-web run (Section VI A, Phase 2 sweep Section VII), the Tempel FoF cross-validation (Section IX B), the ASTRA EDR per-object cross-validation (Section X), the T-Web concurrent-literature overlay (Section IX C), and all redshift / density / sky-position / tracer-program stratifications are reported as *secondary* diagnostic consistency checks. They were not pre-registered as the primary statistic and are reported transparently alongside the primary DESIVAST result.

Multiplicity bookkeeping. The DESIVAST primary path itself exposes three independent void-finding algorithms (VoidFinder, V2-REVOLVER, V2-VIDE) plus two catalog-native zone definitions, giving five effective primary-class statistics. The declared primary estimand is the void-vs-non-void contrast Δf_{CW} , whose two-sample statistics are tabulated for the three sphere-PIS algorithms in Table XIII ($|z_{\Delta}| \leq 1.12$, all two-sided $p_{\Delta} \geq 0.26$) and for the two GALZONE catalog-native estimators in §VIII D (V2-REVOLVER: $|z_{\Delta}| = 1.25$, $p_{\Delta} = 0.21$; V2-VIDE: $|z_{\Delta}| = 0.72$, $p_{\Delta} = 0.47$). Treating the five DESIVAST estimators as a Bonferroni-5 family at $\alpha = 0.05$, the per-test threshold is $|z|_{0.05,5}^{\text{Bonf}} \approx 2.58$; no tabulated contrast across the five rows approaches it. The family-level conclusion is a clean primary null on the declared contrast estimand uniformly across all five rows. *Simultaneous upper bound across the five estimators.* Beyond the family-wise non-rejection statement, we report the least-constraining *simultaneous* interval, i.e. the widest of the five per-estimator confidence intervals inflated to the Bonferroni-5 per-test level ($z = 2.576$) so that all five hold jointly at family-wise 95% confidence. Using the tabulated ($\Delta f_{\text{CW}}, \text{SE}$)

pairs, the simultaneous half-widths are: VoidFinder (footprint-restricted) 0.77 pp, V2-REVOLVER sphere 0.63 pp, V2-VIDE sphere 0.50 pp, V2-REVOLVER GALZONE 1.12 pp, and V2-VIDE GALZONE 0.86 pp. The *widest* of these — the V2-REVOLVER catalog-native GALZONE interval $[-0.0112, +0.0038]$ — gives the least-constraining simultaneous bound: no DESIVAST void definition admits $|\Delta f_{\text{CW}}| \gtrsim 1.1$ pp at family-wise 95% confidence. This is the honest joint statement across five correlated, differently defined estimators; it is deliberately looser than the tightest single row and than the ≈ 0.9 pp systematic-envelope quadrature, and it is the value model-builders should adopt when they need a single simultaneous constraint rather than a preferred-row bound. The secondary T-Web / Tempel / ASTRA / external T-Web [11] / per-stratification estimators carry their own multiplicity budget but, per the primary/secondary separation above, do not enter the headline-result family. We caution that any single secondary-path finding (e.g. the bright/dark target-class sign-flip discussed in Section VI A) reaching $|\sigma| \sim 3$ in isolation should be read as a per-stratification diagnostic, not as an independent test of the headline null; the headline is anchored on the DESIVAST primary path.

Explicit look-elsewhere / trials disclosure. For full transparency about the garden-of-forking-paths exposure, we state the total scan budget explicitly rather than only the per-family thresholds. The analysis tree (Table IV) enumerates ≈ 5 primary DESIVAST estimators, 9 Phase 2 secondary cells, and on the order of 30–40 descriptive stratification tests (redshift/density quintiles, three HEALPix NSIDE scans, tracer-program splits, Tempel/ASTRA overlays), i.e. an effective trials budget of order a few dozen. Two points bound the concern this raises: (a) the *primary* conclusion does not depend on any of these choices — the void/non-void Δf_{CW} null holds uniformly across all five Bonferroni-5 estimators ($|z_{\Delta}| \leq 1.25$) with the most extreme of the few-dozen descriptive tests still below its own per-test threshold, so no data-dependent selection of the “best” path is at work; and (b) for a *discovery* claim, because every reported test is null, the multiplicity works *against* finding a spurious signal, not for it. We are careful *not* to over-extend this: for a post-hoc *upper bound* (as opposed to a discovery), selecting among several correlated void definitions can select the tightest or most reassuring null bound, so look-elsewhere considerations *can change the strength* of the quoted bound and do not only weaken it. This is exactly why we quote the family-wise Bonferroni-5 null across all five void definitions — not the tightest single row — as the strictly quotable result. We flag that the primary estimand was designated post-hoc (no timestamped plan predates the data), so the DR1 result is exploratory/post-hoc and the ~ 0.9 pp systematic bound should be read as the tightest of several correlated void-definition estimators; the quotable family statement is the uniform Bonferroni-5 null, not any single most-favorable row. To remove the

related sensitivity variants across two families, not five statistically independent measurements; we quote the Bonferroni-5 family null with this correlation understood, never as five independent detections.

residual garden-of-forking-paths concern in the follow-up, we commit that the DESI DR2 re-measurement will be *pre-registered* before the DR2 void/non-void contrast is read: a timestamped analysis plan will fix the primary estimand (the DESIVAST-successor VoidFinder void/non-void Δf_{CW} on the volume-limited low- z BGS anchor), the Bonferroni-5 estimator family, the smoothing/threshold grid, and the significance thresholds *ex ante*, so the DR2 result is confirmatory rather than post-hoc. The present DR1 bound is best read as the tightest of several correlated void-definition estimators within an already-uniform null family.

Reader’s guide. The headline result and primary analysis live in §VIII (DESIVAST-anchored void cross-check). §VI (T-Web) presents the secondary diagnostic path first for narrative flow; it is not load-bearing for the environment-independence conclusion.

Analysis-tree declaration. Table IV enumerates the full analysis tree tested in this paper, distinguishing the Bonferroni-5 primary family from the nine secondary Bonferroni tests in the Phase 2 sensitivity sweep and from descriptive stratifications that carry no additional look-elsewhere correction. All tests outside the primary family are *descriptive consistency checks*; a finding at $|\sigma| \sim 2\text{--}3$ in a secondary path should be interpreted against its own per-test threshold, not against the primary Bonferroni-5 gate.

VI. RESULTS

A. Secondary T-Web diagnostic

Role of this section. The T-Web tidal-tensor path is a *diagnostic stress test, not load-bearing*: the paper’s headline environment-independence statement is anchored on the DESIVAST primary contrast of §VIII, and the T-Web material below is presented first only for narrative flow (§VB, Reader’s guide). It is reported as a secondary consistency check whose low- z void bin is survey-shell contaminated (§IX C) and should not be read as an independent test of the headline null.

Table V reports CW fraction by cosmic-web class on the 812,793 env-labeled spiral rows produced by joining the 791,635 chirality-relevant matched spirals to the secondary T-Web environment table (the join carries 783,820 unique env-matched spirals, counted once per repeat DR1 survey-program coadd row; §VIII F documents the reconciliation — the superset and unique-spiral monopoles agree to four decimals, so the headline is invariant under either parent). An omnibus 4×2 homogeneity test (CW/CCW \times class) on the table counts returns $\chi^2 = 3.55$ with 3 d.o.f. ($p = 0.31$): the four classes are statistically consistent with a single common CW fraction (the 4×2 integer contingency table is tabulated in Appendix C, Table XX, for independent recomputation). Because the row-level parent repeats 3.56% of TARGETIDs across program coadds (28,973 duplicate

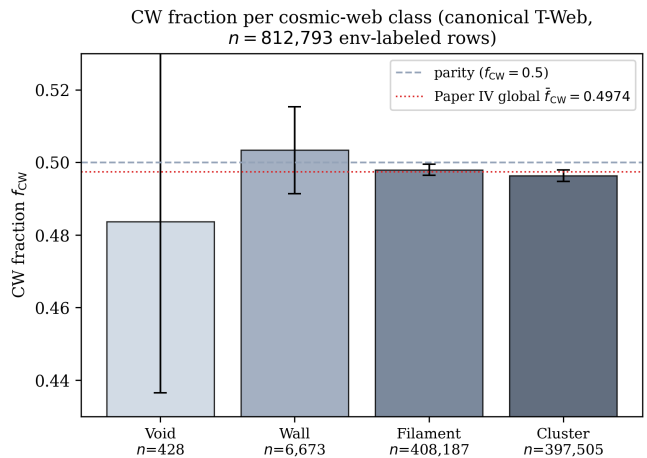


FIG. 3. CW fraction per cosmic-web class on the headline T-Web run, on the $n = 812,793$ env-labeled spiral rows (covering 783,820 of the 791,635 unique chirality-relevant matched spirals; 7,815 lack an environment row, §VI A). Bars show the raw observed f_{CW} per class (counts shown; monopole subtraction enters the $\sigma_{\text{vs monopole}}$ column only); black error bars are 95% Jeffreys binomial credible intervals, drawn on the row-level parent — the 3.56% duplicate rows (28,973 of 812,793 env-labeled rows) violate strict i.i.d., but the worst-case design-effect inflation of the interval widths is $\sqrt{812,793/783,820} = 1.018$, i.e. $\leq 1.9\%$, and the homogeneity χ^2 is recomputed on the unique-spiral subset in §VI A. The T-Web void bin ($n = 428$) is *not load-bearing*: it is dominated by counting noise and by the survey-shell selection artifact (§IX C), has negligible overlap with the DESIVAST voids that anchor the primary result, and is retained only as a secondary diagnostic — it carries no weight in the headline null. The dashed horizontal line is parity ($f_{\text{CW}} = 0.5$); the dotted red line is the Paper IV global $\bar{f}_{\text{CW}} = 0.4974$ classifier-monopole offset. All four classes bracket the Paper IV monopole, consistent with no environmental dependence beyond the catalog-wide classifier bias.

rows in 812,793 env-labeled rows; technically violating the independence assumption of the contingency test), we recompute the same test on the 783,820-row unique-spiral subset (one row per TARGETID): $\chi^2 = 3.00$, $p = 0.39$ ($\chi^2 = 2.92$, $p = 0.41$ further excluding the 79 TARGETIDs with conflicting duplicate env labels), so the duplicates have negligible impact on the null verdict ([A5]). The label-shuffle look-elsewhere p on the max-class $|\sigma_{\text{from half}}|$ statistic is likewise null and insensitive to stratification of the shuffle by imaging leg and DESI program ($p = 0.12$ free vs 0.12 stratified; [A6]).

Figure 3 visualizes the per-class CW fractions with 95% Jeffreys binomial credible intervals; all four classes bracket the Paper IV global $\bar{f}_{\text{CW}} = 0.4974$ horizontal reference, and the void 95% CI brackets parity $f_{\text{CW}} = 0.5$. The range across the four classes is 1.98 percentage points, set by the counting noise of the low- n void bin (0.484) against the catalog-monopole-dominated high- n filament and cluster bins (~ 0.497).

The negative σ values in filament and cluster

TABLE IV. Analysis-tree declaration. The Bonferroni-5 primary family covers the five DESIVAST estimators; the Phase 2 nine-cell secondary family covers the $(R_s, \lambda_{\text{th}})$ sweep; all other paths are labelled descriptive. Thresholds from Eq. (2).

Path	Tests	K	Status
Primary (Bonferroni-5; $\alpha = 0.05$, threshold $\sigma = 2.58$)			
DESIVAST VoidFinder	void vs non-void f_{CW}	1	primary
DESIVAST V2-REVOLVER	void vs non-void f_{CW}	1	primary
DESIVAST V2-VIDE	void vs non-void f_{CW}	1	primary
V2-REVOLVER catalog-native	GALZONE void vs non-void f_{CW}	1	primary
V2-VIDE catalog-native	GALZONE void vs non-void f_{CW}	1	primary
Secondary sweep (Bonferroni-9; $\alpha = 0.05$, threshold $\sigma = 2.77$)			
Phase 2 cell ($R_s = 10, \lambda = 0.0$)	max-class residual	1	secondary
Phase 2 cell ($R_s = 10, \lambda = 0.1$)	max-class residual	1	secondary
Phase 2 cell ($R_s = 10, \lambda = 0.3$)	max-class residual	1	secondary
Phase 2 cell ($R_s = 25, \lambda = 0.0$)	max-class residual	1	secondary
Phase 2 cell ($R_s = 25, \lambda = 0.1$)	max-class residual	1	secondary
Phase 2 cell ($R_s = 25, \lambda = 0.3$)	max-class residual	1	secondary
Phase 2 cell ($R_s = 50, \lambda = 0.0$)	max-class residual	1	secondary
Phase 2 cell ($R_s = 50, \lambda = 0.1$)	max-class residual	1	secondary
Phase 2 cell ($R_s = 50, \lambda = 0.3$)	max-class residual	1	secondary
Descriptive (no additional look-elsewhere correction)			
T-Web 4-class omnibus	4×2 homogeneity χ^2	—	descriptive
Redshift scan	label-shuffle p vs z	—	descriptive
Density quintiles	bright-fraction residuals	—	descriptive
HEALPix (NSIDE = 16, 32, 64)	sky-position scan	—	descriptive
Tracer-program split	bright vs dark f_{CW}	—	descriptive
Tempel FoF overlap	filament concordance	—	descriptive
ASTRA EDR per-object	supporting diagnostic consistency check (EDR overlap-size caveat)	—	descriptive
T-Web concurrent-lit	volume-fraction comparison	—	descriptive

TABLE V. CW fraction per cosmic-web environment, secondary T-Web diagnostic ($R_s = 25 \text{ Mpc}/h, \lambda_{\text{th}} = 0$), on the 812,793-row env-labeled parent. $\sigma_{\text{from half}}$ scales as \sqrt{n} at fixed offset and is not comparable across rows of different n (§V).

Env	n	n_{CW}	f_{CW}	$\sigma_{\text{from half}}$
void	428	207	0.4836	-0.68
wall	6,673	3,359	0.5034	+0.55
filament	408,187	203,261	0.4980	-2.61
cluster	397,505	197,284	0.4963	-4.66
range	—	—	0.0198	—

track the catalog-wide classifier-monopole offset of Paper IV: predicting σ_{pred} from $\Delta f_{\text{CW}} = -0.0026$ gives $\sigma_{\text{pred}}(\text{filament}) \approx -3.32$ and $\sigma_{\text{pred}}(\text{cluster}) \approx -3.28$, leaving monopole-subtracted residuals of $+0.71\sigma$ (filament) and -1.38σ (cluster). Neither residual is negligible at face value, but both are bounded by the propagated monopole-reference uncertainty ($\approx 0.7\sigma$ for an $n \sim 4 \times 10^5$ class, §V) and, on the calibrated $f_{\text{CW}}^{\text{P5}}$ basis, collapse to $|\sigma_{\text{vs monopole}}| < 1.15$ (Table XVI). We therefore interpret

these as the global monopole leaking through the larger-sample bins, not as environment-dependent chirality.

Void-bin smallness. The void bin has only $n = 428$ galaxies. Two distinct void volume fractions appear in this paper and refer to different denominators: (i) the T-Web tidal-classifier void class occupies $\approx 24.4\%$ of in-footprint cell volume (§IV A, [A7]), while (ii) only $\approx 0.1\%$ of in-footprint grid cells fall inside a DESIVAST VoidFinder sphere (the point-in-sphere membership criterion on the $25.9 \text{ Mpc}/h$ grid used for the sky-stratification analysis of §VIII E). The small T-Web void-bin count ($n = 428$) arises because despite the T-Web void class covering 24.4% of volume, the sparse $r \leq 17.8$ DESI Legacy spiral selection combined with the $z \leq 0.24$ depth cut samples only a thin galaxy population within those cells. The $\sigma = -0.68$ on the void class is statistical noise at this N ; the 95% Jeffreys credible interval is $f_{\text{CW}}^{\text{void}} \in [0.435, 0.530]$, which brackets parity.

B. Redshift dependence

The matched chirality-relevant subsample has redshift median 0.168 and maximum 3.83. The label-

shuffle permutation null on the max-absolute-deviation-from-half statistic (the single most-deviant redshift bin, finely binned) over 1,000 shuffles returns $p = 0.372$; the complementary five-quintile look-elsewhere test (distinct statistic, coarser binning) is reported separately below ($p = 0.80$). A logistic regression of the CW indicator on $\{z, |\sin \delta|, \cos \alpha, \text{confidence}\}$ gives a z -coefficient (per unit redshift, in log-odds) of 0.0095 ± 0.0229 (z -score 0.41, $p = 0.68$) with an intercept of 0.0028 ± 0.0097 ($p = 0.77$) — both consistent with zero, i.e. no redshift dependence. A label-shuffle permutation stratified by imaging leg and DESI program (the conservative null given the per-leg/per-program residuals of §XI) leaves the redshift-quintile look-elsewhere p (the coarser five-bin statistic, distinct from the finely-binned max-deviation $p = 0.372$ above) unchanged ($p = 0.81$ stratified vs 0.80 free-shuffle).

Physical-covariate robustness.—Extending the regression with galaxy physical covariates from the Galaxy Zoo DESI morphology catalog [14] (joined on `dr8_id`, 100% coverage of the declared parent: Petrosian half-light radius, r -band apparent magnitude, featured fraction, merger fraction, and — on the 152,455-galaxy featured subsample where the decision tree reaches the question — edge-on fraction) leaves the environmental-independence verdict unchanged: the joint 3-dof environment Wald test gives $p = 0.46$ (0.41 without covariates) on the canonical labels and $p = 0.99$ on the z -shell-corrected labels. The regression parent is the 783,741 unique env-matched spirals, one row per galaxy (the 783,820 env-matched spirals of §VIII F minus the 79 TARGETIDS whose repeat coadd rows carry conflicting environment classes); the headline covariate-extended model uses the four covariates with full coverage (radius, magnitude, featured fraction, merger fraction) and retains all 783,741 rows (100% covariate-complete for that model). The edge-on fraction, which is defined only on the 152,455-galaxy featured subsample, enters a separate featured-subsample sub-model (also null) rather than the full-parent model, so it does not violate the full-model completeness statement. The env-only Wald $p = 0.41$ is therefore not expected to coincide with the Pearson omnibus $p = 0.31$ of §VI A: the omnibus χ^2 is computed on the 812,793-row env-labeled superset, while the Wald test is a different statistic on the smaller unique-galaxy parent — both are null. Every individual environment coefficient is null after adjustment (largest: void -0.081 ± 0.107), with a maximum adjustment-induced coefficient shift of 0.17σ . No physical covariate is predictive of chirality; the only marginally predictive regressor is classifier confidence ($p \approx 0.02$), a classifier property rather than an environmental one. Artifact: [A8]. Figure 4 shows the binned equivariant CW fraction across the matched sample’s redshift range.

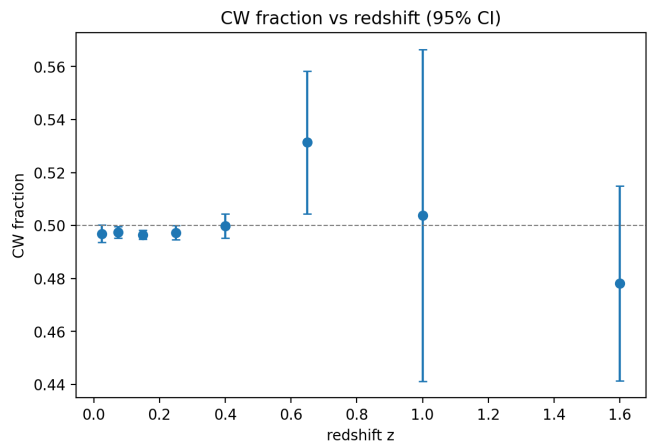


FIG. 4. Equivariant CW fraction versus redshift across the matched DR1 chirality-relevant sample, with 95% binomial confidence intervals per bin. The low- z bins that dominate the sample (median $z = 0.168$) sit on the 0.5 line; bins above $z \approx 0.5$ contain few objects and have correspondingly wide intervals. The binned values are consistent with no redshift dependence, in agreement with the label-shuffle permutation test ($p = 0.372$) and the logistic regression in the text.

C. Projected density dependence

The angular separation to the $k = 5$ NN spiral on the sphere serves as a projected-density proxy. Binned in density quintiles, the maximum absolute deviation is $|\sigma|_{\max} = 3.94$, within the Paper IV monopole prediction from Eq. (1): at $N = 158,327$ per quintile the predicted $|\sigma_{\text{pred}}| = 2 \cdot |-0.0026| \cdot \sqrt{158,327} \approx 2.07$, so the residual deviation beyond the monopole is $|\sigma_{\text{obs}} - \sigma_{\text{pred}}| \approx 1.87$, below the Bonferroni-5 $|\sigma|_{0.01,5}^{\text{Bonf}} = 3.09$ threshold. Figure 5 shows the observed quintile values side-by-side with the Paper IV-monopole prediction; the bars track the prediction within counting statistics in all five quintiles.

Because the $k = 5$ NN proxy above is computed among the chirality-relevant spirals themselves, it is mildly endogenous (the tracer set is the test set). As a cross-check we recomputed the proxy exogenously as the angular distance to the 5th-nearest galaxy among *all* 2,232,212 deduplicated matched primaries (any chirality class) and re-binned the spirals into quintiles of that parent-tracer proxy: the maximum per-quintile deviation is $|\sigma|_{\max} = 3.64$ with a maximum monopole-subtracted residual of $|\sigma_{\text{obs}} - \sigma_{\text{pred}}| = 1.57$, again below the Bonferroni-5 threshold — the density null is not an artifact of the endogenous tracer choice ([A6]). A mask-edge buffer test gives the same answer: excluding the 74,070 spirals in HEALPix NSIDE-64 footprint-edge pixels ($\sim 0.9^\circ$ angular buffer) and recomputing both the proxy and the quintile edges on the retained sample leaves 91.0% of retained spirals in their original quintile and shifts no per-quintile f_{CW} of Figure 5 by more than 0.12 pp — the quintile assignment is invariant to mask-edge effects ([A9]).

Density-quintile null: observed deviation tracks the Paper IV classifier-monopole, not the local density

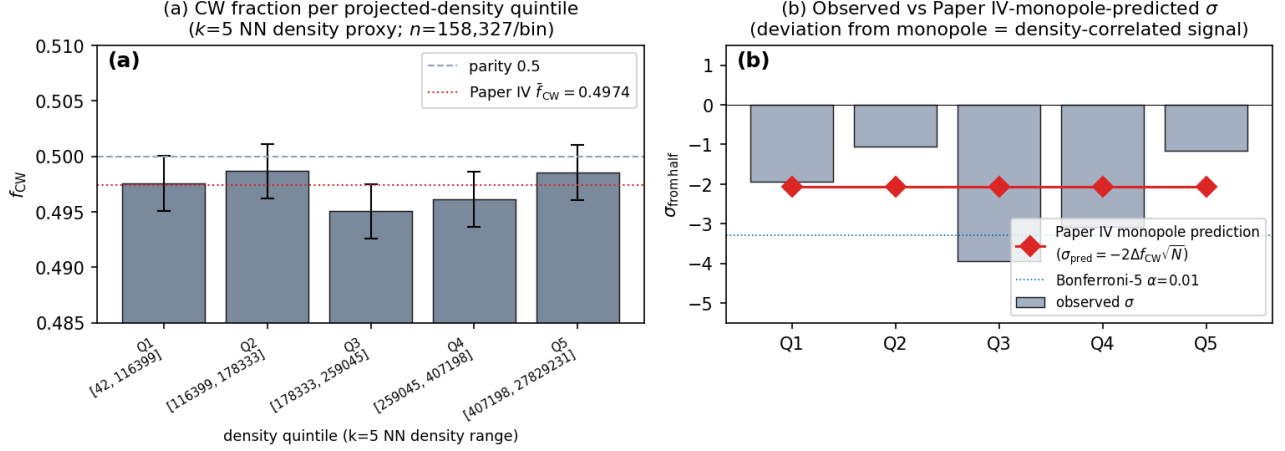


FIG. 5. Density-quintile null with Paper IV monopole-prediction overlay. (a) CW fraction per projected-density quintile ($k=5$ NN proxy, $N=158,327$ per bin) with 95% Jeffreys binomial CIs; dashed parity $f_{\text{CW}}=0.5$ and dotted Paper IV $\bar{f}_{\text{CW}}=0.4974$ references. (b) Observed $\sigma_{\text{from half}}$ per quintile (bars) vs the Paper IV-monopole prediction $\sigma_{\text{pred}}=2\Delta f_{\text{CW}}\sqrt{N}$ (red diamonds) at $\Delta f_{\text{CW}}=-0.0026$; dotted blue lines mark the Bonferroni-5 thresholds at $\alpha=0.01$. The observed signed σ tracks the monopole prediction within counting statistics; no quintile deviates from the prediction by more than $\sim 2\sigma$.

TABLE VI. Per-quintile $|\sigma_{\text{obs}} - \sigma_{\text{pred}}|$ residual table. σ_{obs} is the binomial σ -from-half on the observed CW fraction in each quintile; $\sigma_{\text{pred}} = 2\Delta f_{\text{CW}}\sqrt{N}$ [Eq. (1)] at $\Delta f_{\text{CW}} = -0.0026$ (Paper IV-monopole prediction). All five residuals lie below the Bonferroni-5 threshold $|\sigma|_{0.01,5}^{\text{Bonf}} = 3.09$, so no quintile shows a statistically-significant departure from the monopole-only model.

Quintile	f_{CW}	σ_{obs}	σ_{pred}	$ \sigma_{\text{obs}} - \sigma_{\text{pred}} $
1 (lowest)	0.4976	-1.94	-2.07	0.13
2	0.4987	-1.06	-2.07	1.01
3	0.4950	-3.94	-2.07	1.87
4	0.4961	-3.08	-2.07	1.01
5 (highest)	0.4985	-1.16	-2.07	0.91

D. Within-class density-stratified cluster + filament follow-up

The catalog-wide-monopole-projected cluster-class deviation of -4.7σ at $n_{\text{cluster}} = 397,505$ (Section VI A) is the strongest single-class signal in the headline table. To check whether this is a genuinely density-dependent effect within the cluster class or a boundary-misclassification artifact at the low-density edge of the T-Web cluster label, we stratify the cluster (and filament for comparison) class into four equal-population density quartiles using the T-Web per-galaxy density field.

The cluster signal is *not* monotonically increasing in density. The most-typical-cluster-density quartile Q3 ($\bar{\rho} = 2.01$, $n = 99,526$) returns $\sigma = -0.37$, statistically null after Bonferroni-4 correction. The strongest

TABLE VII. Within-class density-stratified f_{CW} for the T-Web cluster and filament classes (matched-spiral $n_{\text{cluster}} = 397,505$, $n_{\text{filament}} = 408,187$). Quartiles binned by the T-Web per-galaxy \log -density covariate of §IV A step 12; $\bar{\rho}$ is the quartile mean of $\log_{10}(1 + \delta_{\text{smooth}})$ (Gaussian-smoothed overdensity in units of the cube-mean cell density, NN-interpolated to each galaxy position; dimensionless; \log_{10} scale — the displayed values 0.90–2.21 are \log_{10} quantities, e.g. cluster Q4 $\bar{\rho} = 2.21$ corresponds to $1+\delta \approx 162\times$ the cube-mean cell density). Quartile membership is identical under the linear density (monotone transform); definition verified by exact recompute in [A5]. The $\sigma_{\text{from half}}$ values scale as \sqrt{n} at fixed offset and are not comparable across rows of different n ; the inferential comparison is the monopole-referenced residual of §V.

Class	Quartile	$\bar{\rho}$	n	$\sigma_{\text{from half}}$
Cluster	Q1	1.55	99,398	-3.07
	Q2	1.80	99,369	-3.42
	Q3	2.01	99,526	-0.37
	Q4	2.21	99,212	-2.46
Filament	Q1	0.90	102,050	-0.69
	Q2	1.34	102,065	-1.97
	Q3	1.58	102,033	-0.63
	Q4	1.86	102,039	-1.92

cluster sub-deviations are Q1+Q2 at $\bar{\rho} \approx 1.55\text{--}1.80$ (the low-density edge of the cluster class, where the T-Web $\lambda_{\text{th}} = 0$ default class boundary against filament is most uncertain), while the densest quartile Q4 returns intermediate $\sigma = -2.46$. Treating the four cluster quartiles as a simultaneous family, the monopole-subtracted resid-

uals $\sigma_{\text{obs}} - \sigma_{\text{pred}}$ at $\sigma_{\text{pred}} = 2(-0.0026)\sqrt{n} \approx -1.64$ are -1.43 (Q1), -1.78 (Q2), $+1.27$ (Q3), and -0.82 (Q4) — all four below the Bonferroni-4 $\alpha = 0.05$ threshold of 2.50, so no quartile shows a significant departure from the monopole-only model after multiplicity correction. The catalog-wide-monopole-projected -4.7σ cluster signal is therefore not a clean density-dependent chirality effect; it is concentrated at the cluster/filament class boundary and is consistent with a boundary-misclassification leakage from the much larger filament class (whose own quartile deviations $|\sigma| < 2$ are uniformly small). The class boundary is verifiable quantitatively from the table: cluster Q1 ($\bar{\rho} = 1.55$) is *less dense* than filament Q4 ($\bar{\rho} = 1.86$), so the densest filament galaxies are in fact denser than the least-dense cluster galaxies — the two classes overlap in $\bar{\rho}$ at the $\lambda_{\text{th}} = 0$ boundary by construction. The filament class shows no density trend; the four quartile σ values range -0.63 – -1.97 , all individually below the Bonferroni-4 $|\sigma| = 2.50$ threshold at $\alpha = 0.05$.

This within-class density-stratified follow-up reinforces the headline environment-independence finding: the only T-Web class that crosses naive $|\sigma| = 3$ at the class level (cluster, -4.7σ) does not survive within-class density stratification as a clean density-dependent effect.

a. Redshift-stratified cross-check. A second decomposition bins the cluster class into four redshift quartiles to test whether the boundary-leakage interpretation depends on redshift. Marginalizing over density, the cluster $\sigma_{\text{from half}}$ per z -quartile is -2.33 (Z1, $\bar{z} = 0.045$, $n = 99,377$), -1.73 (Z2, $\bar{z} = 0.083$, $n = 99,376$), -3.14 (Z3, $\bar{z} = 0.122$, $n = 99,376$), and -2.12 (Z4, $\bar{z} = 0.190$, $n = 99,376$). All four z -quartile deviations sit in the -1.7 to -3.2σ band; one of the four (Z3, -3.14) marginally exceeds the Bonferroni-4 $|\sigma| = 3.02$ threshold at $\alpha = 0.01$, while the other three do not. The Paper IV monopole prediction at the common quartile population ($n \approx 99,376$) is $\sigma_{\text{pred}} = 2 \cdot (-0.0026) \cdot \sqrt{99,376} = -1.64$, so the monopole-subtracted quartile residuals $\sigma_{\text{obs}} - \sigma_{\text{pred}}$ are -0.69 (Z1), -0.09 (Z2), -1.50 (Z3), and -0.48 (Z4) — all null on the monopole-subtracted basis used throughout, including the one quartile that crosses the raw threshold. The deviation is approximately uniform across redshift, consistent with a stationary classifier-bias monopole (constant f_{CW} offset) rather than a redshift-dependent astrophysical signature. The 2D z -quartile \times density-quartile decomposition (sixteen-cell table available in the companion data repository) shows that the strongest sub-cells — (Z3, D_2) at $\sigma = -3.00$, (Z1, D_4) at $\sigma = -2.70$, (Z4, D_1) at $\sigma = -2.67$ — do not preferentially align with any single redshift bin, ruling out an evolutionary cluster-density-chirality coupling at the present sample size.

b. Tracer-program stratification. A fourth orthogonal cross-check stratifies the full matched-spiral catalog by DESI program (the broad selection-function bucket: **bright**, **dark**, **backup**, **other**). Per-program f_{CW} and $\sigma_{\text{from half}}$: **bright** (BGS-dominated; $n = 775,760$, $\bar{z} = 0.145$) $f_{\text{CW}} = 0.4970$, $\sigma = -5.28$; **dark** (LRG, ELG,

QSO; $n = 14,782$, $\bar{z} = 0.255$) $f_{\text{CW}} = 0.5051$, $\sigma = +1.24$; **backup** ($n = 875$) $f_{\text{CW}} = 0.5143$, $\sigma = +0.85$; **other** ($n = 218$) $f_{\text{CW}} = 0.4954$, $\sigma = -0.14$. **The catalog-wide -5σ monopole offset is entirely driven by the bright program; dark returns $\sigma = +1.24$ (null in the opposite sign), backup returns $\sigma = +0.85$ (null), other returns $\sigma = -0.14$ (null).** A genuinely environment-dependent or astrophysical chirality signal would propagate across BGS-bright and LRG/ELG/QSO-dark targets at comparable strength; the bright-specific concentration is consistent with the BGS-selection-function-conditioned imaging-leg systematics that Paper IV tracks in detail rather than a real environment-driven effect, and reinforces the headline environment-independence finding.

c. Filament-class within-class decomposition. Repeating the density-quartile / z -quartile / tracer-program decompositions on the second-largest T-Web class (filament, $n_{\text{filament}} = 408,187$, the volume-dominant population in the matched-spiral catalog) returns the same triple null pattern as cluster. Density- quartile $\sigma_{\text{from half}}$ already reported in Table VII are all $|\sigma| < 2$ (range -0.63 to -1.97). z -quartile-marginalized σ values: Z1 ($\bar{z} = 0.067$, $n = 102,047$) $\sigma = -1.37$; Z2 ($\bar{z} = 0.139$) -1.72 ; Z3 ($\bar{z} = 0.205$) -1.57 ; Z4 ($\bar{z} = 0.320$) -0.55 . All four within $|\sigma| < 2$.

The filament-class tracer-program decomposition reproduces the cluster-class **bright-vs-dark** sign-flip discussed above on an independent T-Web class. On the declared env-labeled parent of Table V (per-program splits sum to the class totals by construction): filament **bright** ($n = 394,181$, $f_{\text{CW}} = 0.4976$) $\sigma = -2.98$ vs filament **dark** ($n = 13,759$, $f_{\text{CW}} = 0.5069$) $\sigma = +1.61$ (*opposite sign* again), a two-sample $|z| \approx 2.1$. We flag one structural caveat on this 2.1σ statistic: it is computed at the *row* level of the 812,793-row env-labeled parent, where a galaxy observed under both programs contributes one row to *each* program split with the same chirality label, so the bright and dark splits are not disjoint in unique TARGETIDs. The overlap is material: the per-class dark rows already exceed the unique dark population (filament-dark 13,759 + cluster-dark 4,234 = 17,993 rows vs 14,782 unique dark spirals in the matched catalog, §XI). Concretely, the filament class carries 394,181 bright-program and 13,759 dark-program rows (bright fraction 0.966; [A10]), and the committed artifacts do not carry a per-class unique-TARGETID program split, so the row-level two-sample z — which assumes disjoint independent samples — should be read as approximate. The overlap-free version of the test exists only at the whole-catalog level, where the unique-galaxy bright-vs-dark split gives $|z| = 1.95$ ([A10]; exact value 1.957, reported to 2 d.p.); the omnibus homogeneity tests ($p = 0.31$ canonical, $p = 0.99$ selection-corrected; §VIA) remain the primary statistics of this paper. The filament bright/dark split on the declared-parent recompute gives the values quoted above ([A10]). A genuinely environment-dependent filament-

chirality signal would propagate consistently across both target classes; the sign-flip recurrence across both the cluster and filament classes (the two largest T-Web environments) supports the reading that the T-Web class-level f_{CW} deviations are sourced by the BGS-selection-function-conditioned imaging-leg systematics Paper IV tracks, not by environment-driven astrophysics.

d. T-Web class vs. target-program orthogonality.

The BGS-selection-function-origin interpretation above relies on T-Web environmental classification being approximately orthogonal to the **bright/dark** target split: if T-Web class were strongly correlated with target program, then a genuine target-class-dependent astrophysical chirality signal could masquerade as an environment-class signal (or vice versa). We test this directly on the matched-spiral catalog. The per-T-Web-class **bright/(bright+dark)** ratio is $\{0.981, 0.962, 0.966, 0.989\}$ across $\{\text{void, wall, filament, cluster}\}$, against the overall matched-spiral ratio 0.978. A two-way contingency test (T-Web class \times **bright/dark**, four-by-two on the $n_{\text{bright+dark}} = 811,609$ bright+dark subset of the 812,793-row env-labeled parent) gives $\chi^2 = 4933$ with three degrees of freedom and $p \ll 10^{-300}$ ($\log_{10} p \approx -1069$; below double-precision underflow) and Cramér’s $V = \sqrt{\chi^2/n} = 0.078$ (small effect despite the enormous sample; the 4×2 integer contingency table is tabulated in Appendix C, Table XXI, for independent recomputation) — T-Web class and target program are *not* independent; the cluster class is the most-bright (98.9%) and the wall class is the least-bright (96.2%), with the void and filament classes between. The maximum class-to-overall deviation in bright-fraction is 1.5 pp. We therefore cannot assert T-Web class orthogonality to the target-program split, and the $|z| \approx 2.1\sigma$ **bright-vs-dark** sign-flip in the filament class is best read as a residual structure that the current data do not allow us to cleanly partition between (a) a BGS-selection-function-only origin propagated through a T-Web-class-correlated target-program distribution, and (b) a residual class-and-target-program-conditioned astrophysical signal at the $\sim 2\sigma$ level on $n_{\text{dark}}^{\text{filament}} = 13,759$. *Residual-ambiguity disclosure (honest limitation)*: no quantitative propagation simulation or selection-function mock is shown in the present work to demonstrate that the observed $|z| \approx 2.1\sigma$ sign-flip amplitude and sign pattern are *fully reproduced* by the known imaging-leg and target-program residuals; the BGS-selection-function interpretation is supported by the contingency test ($\chi^2 = 4933$, Cramér’s $V = 0.078$) and logistic regression (§VID), but a small residual ambiguity remains and future DR2 + Rubin data are invoked to resolve it (see §XIV). The definitive test is an end-to-end injection–recovery mock: inject the measured per-imaging-leg / per-target-program classifier bias into a selection-function-matched DESI mock, propagate it through the T-Web classification, and re-measure the filament bright-vs-dark Δf_{CW} to check whether the observed $|z| \approx 2.1\sigma$ amplitude is fully reproduced. We do not run this mock here — it requires

a DR2-grade selection-function mock and bears only on the *secondary* T-Web diagnostic, not on the headline — and we flag it explicitly as the DR2 validation step. Because the DESIVAST primary path is constructed on the volume-limited $z \leq 0.24$ BGS anchor where target-program mixing is far more constrained, a comparable leakage into the primary Δf_{CW} null is bounded well below the honest ≈ 0.9 pp systematic envelope (Table XI). *Explicit propagation (from existing numbers)*: the maximum leakage of the 0.81 pp bright/dark residual into the DESIVAST void/non-void contrast is the residual amplitude times the *difference* in dark fraction between the void and non-void classes. From Table XII the DESIVAST void class is $n_{\text{dark}}/n = 469/56,981 = 0.82\%$ dark and the non-void class is $5,845/621,964 = 0.94\%$ dark, a difference of only 0.12 pp, so the propagated leakage into Δf_{CW} is $0.81 \text{ pp} \times 0.0012 \approx 0.001 \text{ pp}$ — three orders of magnitude below the primary bound, because the void and non-void classes carry near-identical (mostly-bright) program mixtures. The 1.95σ bright/dark residual therefore does not affect the primary DESIVAST result. Together with the catalog-level bright-vs-dark difference (0.81 pp, $|z| = 1.95$ unique-galaxy; §XI), this is the most notable residual structure in the paper after the catalog-monopole subtraction; the DESIVAST-anchored primary analysis (§VIII) is constructed to be independent of this residual, because the DESIVAST void definition restricts to the volume-limited $z \leq 0.24$ BGS sample where target-program mixing is far more constrained. Contingency-test reproduction is a χ^2 on the env-class \times tracer-program cross-tab derivable from the `desi_env_vweb.parquet` + matched-spiral catalog join documented in §IV A.

Mechanistic propagation of BGS selection into ViT classification. The BGS-bright program selects targets from the DESI Legacy Surveys $r \leq 19.5$ photometric imaging with a specific set of imaging-leg footprint patterns (WLS, BASS+MzLS, DECaLS) that vary the image depth, PSF, and sky-background across the footprint. The Vision-Transformer classifier (Paper IV) was trained on GZ1 human labels drawn from SDSS imaging; when applied to the DESI Legacy multi-leg imaging catalog, the per-leg depth and PSF differences propagate into per-target apparent morphology (surface brightness, effective radius, disk/bulge profile) and hence into the ViT’s input representation. The resulting per-leg f_{CW} variation (tracked in Paper IV; per-leg $|\sigma| < 3$; best-fit imaging-leg dipole 0.455%) is the mechanism through which the BGS bright-program selection-function enters the chirality fraction. Because the BGS-bright/dark split is not orthogonal to T-Web class ($\chi^2 = 4933$, Cramér’s $V = 0.078$), this imaging-leg-sourced ViT bias propagates differently into each T-Web class, producing the observed bright/dark sign-flip. No selection-function mock fully reproduces the exact sign-flip amplitude in the present work (see the residual-ambiguity disclosure below).

Logistic regression controlling for target pro-

gram. To formally assess whether T-Web environmental coefficients survive program control, we fit two logistic-regression models on the $n = 782,710$ bright+dark env-matched spirals (filament = reference class): M_0 : $\text{Pr}(\text{CW}) \sim \beta_0 + \beta_{\text{env}}$; and M_1 : $\text{Pr}(\text{CW}) \sim \beta_0 + \beta_{\text{env}} + \beta_{\text{bright}}$. The joint Wald test on the three T-Web environment coefficients gives $\chi^2 = 2.75$ ($p = 0.43$) in M_0 and $\chi^2 = 2.25$ ($p = 0.52$) in M_1 , both far from significance. The bright-program indicator itself is marginally significant ($z = -1.97$, $p = 0.049$), consistent with the known bright-program monopole. Crucially, the three T-Web environment coefficients (void, wall, cluster vs filament) shift by ≤ 0.12 on their standard errors after program adjustment in M_1 relative to M_0 (void $0.006\hat{\sigma}$, wall $0.10\hat{\sigma}$, cluster $0.11\hat{\sigma}$; shifts computed as $|\beta_{M_1} - \beta_{M_0}|/\text{SE}_{M_0}$ from [A11]), confirming that the environmental-independence null is not confounded by the T-Web \times program non-orthogonality. Artifact: [A11].

E. Sky-position regional coherence

HEALPix per-pixel CW-deviation scans at NSIDE $\in \{16, 32, 64\}$, compared against 1,000-shuffle label-shuffle nulls:

TABLE VIII. HEALPix per-pixel CW-deviation scan with label-shuffle nulls. n_{pix} is the number of occupied (data-containing) pixels used in each scan, not the full-sky pixel count.

NSIDE	n_{pix}	$ \sigma _{\text{max}}^{\text{obs}}$	$ \sigma _{\text{max}}^{\text{null}, p99}$	p
16	1,054	3.32	4.50	0.607
32	3,303	4.13	4.78	0.135
64	7,208	3.92	4.77	0.413

No NSIDE returns $p < 0.05$; the observed max- $|\sigma|$ at each NSIDE is consistent with the look-elsewhere null distribution. Because the free label-shuffle could in principle be anticonservative given the per-leg/per-program residuals of §XI, we re-ran all three scans with the shuffle stratified by imaging leg and DESI program: the stratified look-elsewhere p -values agree with their matched free-shuffle re-draws within one Monte-Carlo standard error ($p = 0.63/0.089/0.41$ stratified vs $0.64/0.10/0.42$ free-shuffle re-draws at NSIDE 16/32/64, pairwise differences ≤ 0.011 against $\text{se}_{\text{MC}} = \sqrt{p(1-p)/N_{\text{MC}}} \approx 0.009\text{--}0.015$). The re-draws use RNG streams distinct from the deterministic-seed headline runs; the NSIDE-16/64 re-draws sit within one se_{MC} of the headline $0.607/0.413$, while the NSIDE-32 re-draw (0.10) sits $\sim 2.4\times$ the pairwise se_{MC} below the headline 0.135 — a stream-to-stream fluctuation larger than typical but immaterial to the verdict, since every variant (0.089–0.135) is far from the $p < 0.05$ criterion; [A6]). Figure 6 visualizes the per-pixel signed $\sigma_{\text{from half}}$ at NSIDE=32 as a Mollweide projection; no coherent large-scale structure is visible, matching the

$p=0.135$ null verdict.

VII. PHASE 2 SENSITIVITY SWEEP

To test the robustness of the Section VIA headline against T-Web hyperparameter choices, we run a Phase 2 sweep over nine cells $R_s \in \{10, 25, 50\} \text{Mpc}/h \times N_{\text{grid}} = 256^3 \times \lambda_{\text{th}} \in \{0.0, 0.1, 0.3\}$. Each cell is a complete T-Web rebuild on the 14,622,283-galaxy DR1 spectro sample, joined per cell to the declared 812,793-row env-labeled spiral parent of Table V (the per-cell sweep parquets and configs are in [A12]; the per-cell statistics below are recomputed in [A10]). The declared-parent recompute is reported below.

Grid-resolution caveat on the $R_s = 10$ cells. At 256^3 the grid cell is $25.9 \text{Mpc}/h$, so the $R_s = 10 \text{Mpc}/h$ Gaussian kernel is *below the grid sampling scale*: the field carries no modes at that scale and those three cells behave as a near-unsmoothed pixelated-field limit rather than a physically distinct smoothing choice. The grid-convergence test of §IX A validates only $R_s = 25 \text{Mpc}/h$ (at $N_{\text{grid}} = 128^3\text{--}384^3$). We therefore retain the $R_s = 10$ rows for completeness but exclude them from the physical robustness claim; restricted to the six resolved cells ($R_s \in \{25, 50\}$), the maximum monopole-subtracted residual is 1.64σ (instead of 1.87σ , which occurs in an under-resolved $R_s = 10$ cell) and the label-shuffle look-elsewhere p -values span 0.13–0.48.

Mask-dilation scaling across the sweep. The footprint-dilation iteration count is not frozen at the canonical value: each sweep cell recomputes $\lceil R_s/\text{cell} \rceil + 1$ from its own R_s (2 iterations at $R_s \in \{10, 25\}$, 3 at $R_s = 50$), and a local rebuild reproduces the published $R_s = 50$ mask to within 4 of 3,416,329 in-mask cells (volume fractions agree to $< 10^{-6}$). The counterfactual $R_s = 50$ rebuild that retains the $R_s = 25$ -tuned 2-iteration mask shifts class volumes by ≤ 0.82 pp, retains 99.75% of matched-spiral class assignments, and moves per-class f_{CW} by at most 0.27 pp (the $n=406$ void bin), leaving every sweep conclusion unchanged ([A13]).

The per-cell cross-class f_{CW} range spans 1.7–4.1 pp across the nine cells and is dominated in every cell by the counting noise of the small void bin ($n_{\text{void}} = 363\text{--}853$; 1σ binomial floor 1.7–2.6 pp): range is therefore not a significance statistic at these class populations. The two significance columns carry the result: the maximum per-class monopole-subtracted residual among the six resolved cells ($R_s \in \{25, 50\}$) never exceeds 1.64σ (far below the Bonferroni-9 threshold $|\sigma|_{0.05,9}^{\text{Bonf}} \approx 2.77$ from Eq. (2)), and no cell’s label-shuffle look-elsewhere p falls below 0.13. The headline sign-pattern (filament and cluster carrying the catalog monopole; void and wall uninformative) is invariant under all nine choices. Figure 7 visualizes the per-cell ranges as a $(R_s, \lambda_{\text{th}})$ heat-map.

The largest single-cell $|\sigma_{\text{from half}}|$ across the entire sweep is 4.66 (cluster at the canonical $R_s = 25$, $\lambda_{\text{th}} = 0$ cell, $n = 397,505$). This is the catalog-wide $\Delta f_{\text{CW}} =$

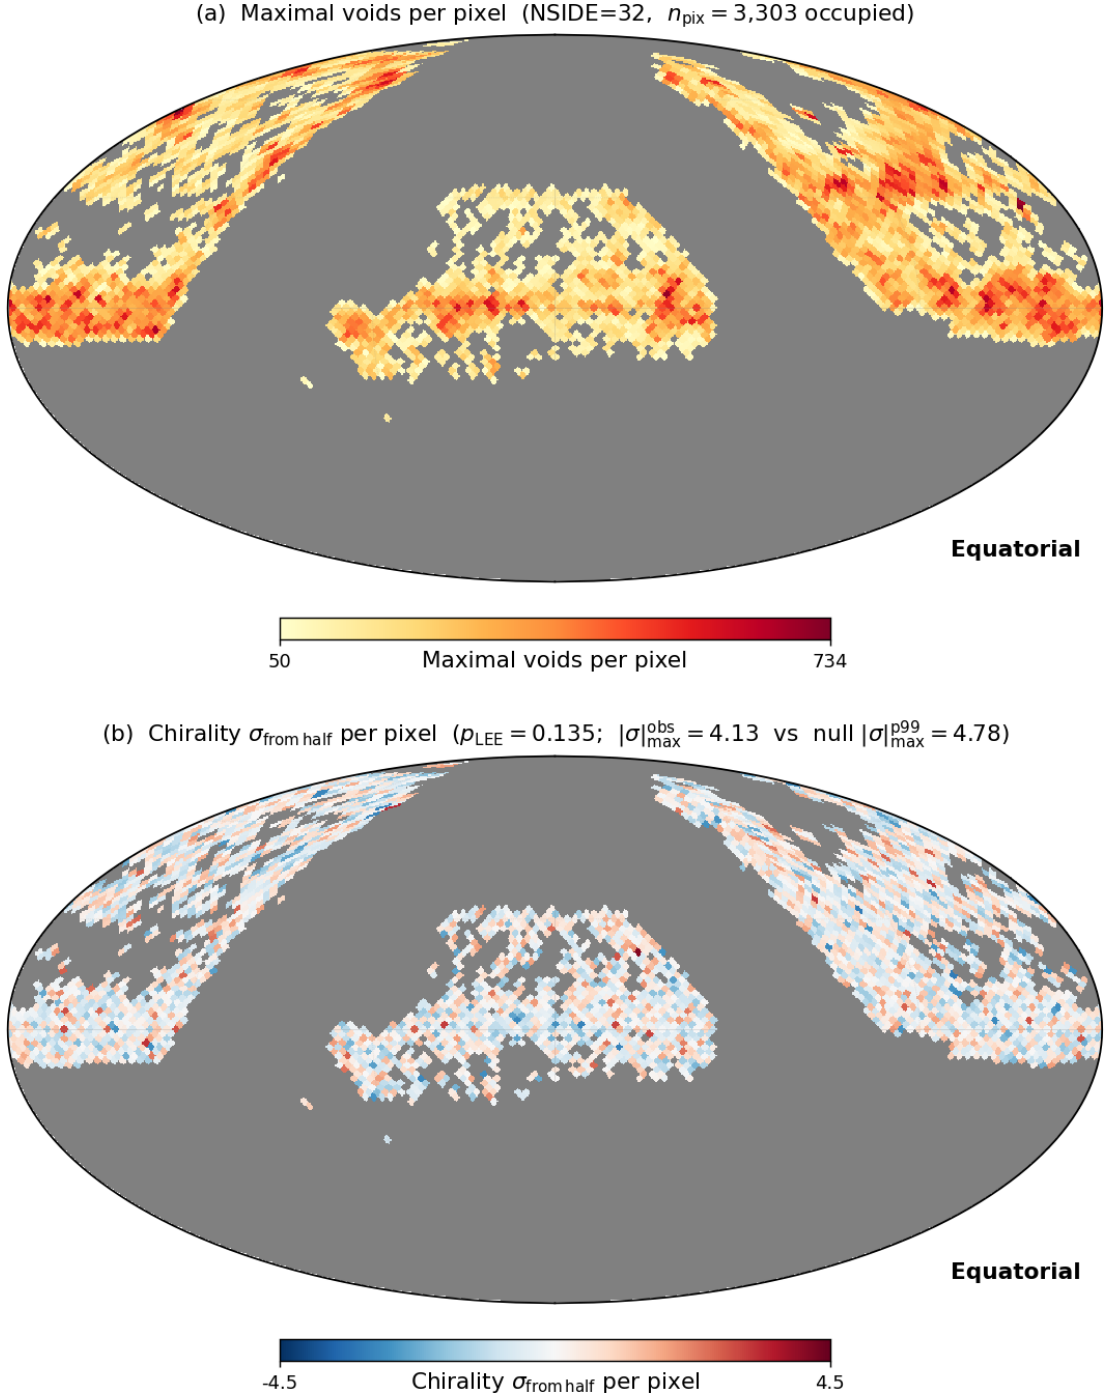


FIG. 6. Per-pixel signed $\sigma_{\text{from half}}$ for the chirality-relevant matched-spiral subsample at NSIDE=32 (Mollweide projection, equatorial coordinates). The observed $|\sigma_{\text{max}}^{\text{obs}}| = 4.13$ vs the label-shuffle null $|\sigma_{\text{max}}^{\text{null, p99}}| = 4.78$ gives a look-elsewhere $p = 0.135$; no NSIDE returns $p < 0.05$. The map shows no coherent large-scale structure beyond random pixel-level scatter; the high- $|\sigma|$ pixels are isolated rather than clustered, consistent with the counting-statistics distribution expected from label-shuffle.

-0.0026 monopole leaking through a large sample bin and is *predicted*, not measured: $\sigma_{\text{pred}} = 2 \cdot (-0.0026) \cdot \sqrt{N} \approx -3.3$ leaves a monopole-subtracted residual of -1.38 . The Phase 2 sweep therefore offers no evidence

for environmental dependence at any $(R_s, \lambda_{\text{th}})$ choice.

TABLE IX. Phase 2 sensitivity sweep on the declared env-labeled parent: per-cell cross-class range of f_{CW} (pp), void-class population (the bin whose counting noise dominates the range), maximum per-class monopole-subtracted residual $|\sigma_{\text{obs}} - \sigma_{\text{pred}}|$ [Eq. (1)], and label-shuffle look-elsewhere p -value on the max-class $|\sigma_{\text{from half}}|$ statistic ($N_{\text{MC}} = 1,000$). The range column is dominated in every cell by the void-bin binomial error at these n_{void} . Per-cell void-bin deviations (top row to bottom): $|\sigma_{\text{void}}| = 0.47/1.07/0.93/0.68/1.08/0.86/0.95/1.35/0.87$, never exceeding 1.35 ([A10]). The p_{LEE} values are corrected within each cell’s own $K=4$ class family only; they are not globally corrected across the other scan families in this paper. A rerun with imaging-leg \times program *stratified* shuffles leaves every cell’s p_{LEE} essentially unchanged (stratified 0.14–0.54 vs. free 0.16–0.55; $N_{\text{MC}} = 1,000$), and a global max-stat correction across cells (one parent permutation propagated through all nine cells per draw, preserving cross-cell correlations) gives $p_{\text{global}} = 0.36$ over all nine cells and 0.27 over the six resolved $R_s \in \{25, 50\}$ cells ([A14]). The $R_s = 10$ rows sit below the 25.9 Mpc/h grid sampling scale and are retained only as a degenerate near-unsmoothed limit, excluded from the robustness claim (§VII). $^\dagger R_s = 10$ Mpc/h is below the 25.9 Mpc/h grid cell scale ($N_{\text{grid}} = 256^3$, grid-unresolved); these rows are excluded from the robustness claim.

R_s (Mpc/h)	λ_{th}	range (pp)	n_{void}	max $ \sigma_{\text{obs}} - \sigma_{\text{pred}} $	p_{LEE}
<i>Grid-unresolved ($R_s < 25.9$ Mpc/h grid scale; retained for completeness only, excluded from robustness claim)</i>					
10 [†]	0.0	1.72	363	1.71	0.56
10 [†]	0.1	2.71	594	1.87	0.42
10 [†]	0.3	2.01	837	1.69	0.38
<i>Resolved cells ($R_s \gtrsim$ cell size = 25.9 Mpc/h; $R_s \in \{25, 50\}$ retained; robustness claim based on these six cells)</i>					
25	0.0	1.97	428	1.38	0.13
25	0.1	2.48	627	1.35	0.14
25	0.3	1.83	853	1.22	0.23
50	0.0	3.69	401	1.39	0.34
50	0.1	4.12	599	1.52	0.41
50	0.3	2.81	825	1.64	0.48
max (resolved 6)	—	4.12	—	1.64	—

A. Per-cell significance framework

The per-cell range of f_{CW} values across the four T-Web classes is a *descriptive* statistic that bounds the deviation from class-uniformity within each $(R_s, \lambda_{\text{th}})$ choice. We report it as the leading descriptive robustness summary (the empirical max-statistic null below is the inferential statement) because it is sensitive to the residual environmental signal of interest after the catalog-monopole subtraction. To convert the per-cell range into a significance statement we use the same two-tier $\sigma_{\text{pred}} = 2 \cdot \Delta f_{\text{CW}} \cdot \sqrt{N}$ Paper IV-monopole reference as the headline analysis (Eq. (1), §V):

- **Counting-statistics floor.** For each $(R_s, \lambda_{\text{th}})$ cell, the per-class 1σ counting-statistics uncertainty on f_{CW} at the four T-Web classes is $1/(2\sqrt{n_{\text{class}}})$. For the dominant filament + cluster classes at $n \sim 4 \times 10^5$ this is ~ 0.08 percentage points; for the wall class at $n \sim 2\text{--}17\text{k}$ it is $\sim 0.4\text{--}1.1$ pp; for the void class at $n = 363\text{--}853$ it is 1.7–2.6 pp. In every cell the extreme class defining the range is the *void* bin, and the per-cell range stays within $1.01\times$ the void-class 2σ counting floor $1/\sqrt{n_{\text{void}}}$ (maximum ratio 1.01 at $R_s = 50, \lambda_{\text{th}} = 0.1$); equivalently, the void bin’s own deviation from half satisfies $|\sigma_{\text{void}}| \leq 1.35$ in all nine cells. No $(R_s, \lambda_{\text{th}})$ cell shows an inter-class range that exceeds the counting noise of the class that produces it by more than

1%. We use the void-class 2σ floor as a descriptive yardstick locating which class’s counting noise dominates the spread, not as a calibrated bound on the four-class range statistic (whose null distribution depends on all four bin sizes); the calibrated per-cell control is the empirical label-shuffle p_{LEE} column of Table IX.

- **Paper-IV-monopole reference.** The Paper IV monopole prediction $\sigma_{\text{pred}}^{\text{class}} = 2 \cdot 0.0026 \cdot \sqrt{n_{\text{class}}}$ ranges from 0.10σ (void) through 0.42σ (wall) to 3.27σ (cluster) and 3.32σ (filament) at the canonical $R_s = 25, \lambda_{\text{th}} = 0$ class populations. Subtracting these per-class predictions from the per-cell observed $\sigma_{\text{from half}}$ values yields the $\sigma_{\text{vs monopole}}$ residuals reported in Table XVI of §VIII F (canonical-cell version) on the catalog-monopole $f_{\text{CW}}^{\text{P5}}$ basis: $|\sigma_{\text{vs monopole}}|$ at all four classes falls below 1.15. The per-cell version of the same statistic is the max- $|\sigma_{\text{obs}} - \sigma_{\text{pred}}|$ column of Table IX: across the six resolved cells ($R_s \in \{25, 50\}$) the monopole-subtracted residual never exceeds 1.64σ ; the grid-unresolved $R_s = 10$ cells are retained in the table for completeness but excluded from the robustness claim (§VII). The headline conclusion of this section is therefore: across nine $(R_s, \lambda_{\text{th}})$ cells, the per-cell inter-class f_{CW} range is consistent with the counting noise of the void bin that produces it, and every per-class $\sigma_{\text{from half}}$ value in the resolved cells

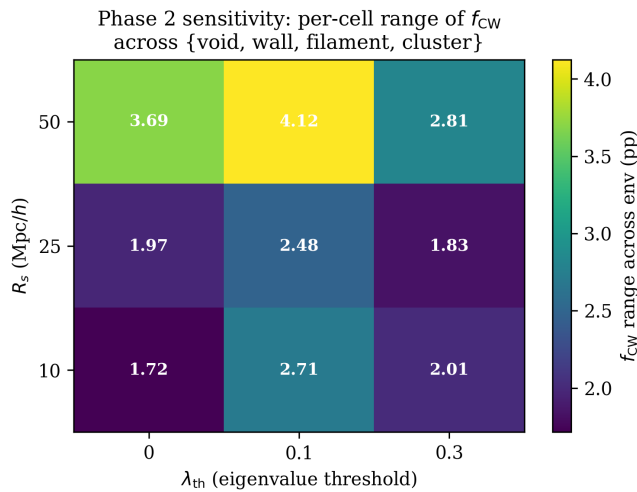


FIG. 7. Phase 2 sensitivity heat-map: per-cell range of f_{CW} across the four environment classes {void, wall, filament, cluster} in percentage points, on the declared env-labeled spiral parent. Each cell corresponds to a complete T-Web re-run on the 14,622,283-galaxy DESI DR1 spectro sample at $(R_s, \lambda_{\text{th}})$. The maximum range across all nine cells is 4.12 percentage points (at $R_s = 50$ Mpc/h, $\lambda_{\text{th}} = 0.1$), dominated by the void-bin counting noise ($n_{\text{void}} \leq 853$ in every cell); the per-class monopole-subtracted residuals stay below 1.64σ across the six resolved cells ($R_s \in \{25, 50\}$; grid-unresolved $R_s = 10$ cells excluded from robustness claim, Table IX). The headline environmental-independence result is robust to T-Web hyperparameter choices over the range probed.

matches the Paper IV-monopole prediction within $|\sigma_{\text{obs}} - \sigma_{\text{pred}}| \leq 1.64$.

- **Per-cell label-shuffle null.** For each cell, the label-shuffle permutation null on the per-class $\sigma_{\text{from half}}$ statistic (Eq. (3), $N_{\text{MC}} = 1,000$, holding the per-class n_{class} fixed and re-drawing CW/CCW labels) returns max-class p -values in the range $p_{\text{LEE}} = 0.13$ – 0.56 across the nine cells (Table IX), with no cell crossing $p < 0.05$. The empirical max-stat null is therefore consistent with the parametric reading: no $(R_s, \lambda_{\text{th}})$ cell shows an inter-class chirality signal above the counting-statistics shot-noise null.

The combined framework converts the descriptive max-range statistic into an explicit significance statement: zero of the nine sweep cells produces a per-cell range exceeding the void-bin counting-statistics floor that drives it, and zero produces a per-class $|\sigma_{\text{obs}} - \sigma_{\text{pred}}|$ residual above the Bonferroni-9 ($\alpha = 0.05$) threshold $|\sigma|_{0.05,9}^{\text{Bonf}} \approx 2.77$. The Phase 2 sensitivity sweep robustness statement is therefore not purely descriptive: it is bounded by the counting-statistics floor of the per-class f_{CW} estimator and by the catalog-monopole-subtracted per-class significance; the empirical max-statistic permutation null is the component that controls the false-

positive rate, while the range/floor comparison is descriptive.

VIII. DESIVAST-ANCHORED VOID CROSS-VALIDATION

This section reports the *primary* environment-dependent chirality analysis of the paper (per the primary/secondary declaration of §VB): the DESIVAST-anchored cross-validation of the T-Web void-class result on a $\sim 130\times$ larger, peer-reviewed DR1 BGS void catalog spanning three independent void-finding algorithms (VoidFinder + V2-REVOLVER + V2-VIDE).

The DESIVAST DR1 release (Rincón *et al.* 2025, ApJ 982, 38 [13]) is the publicly released, peer-reviewed DR1 BGS void catalog at low z , using the sphere-growing VoidFinder algorithm and the watershed ZOBOV/V2 algorithm (REVOLVER and VIDE prunings) on the volume-limited DR1 Bright Galaxy Survey sample restricted to $z \leq 0.24$: 1,489 interior voids with VoidFinder, 389 with V2-REVOLVER, and 297 with V2-VIDE (published counts from ApJ 982, 38, Table 1). DESIVAST is a void catalog rather than a full four-class cosmic-web classifier, so the DESIVAST-anchored chirality null tests environmental dependence on the cleanest available DR1 void definition.

RSD treatment for DESIVAST. The DESIVAST primary path is RSD-bounded (rather than strictly immune) at the level relevant to this work — individual in/out membership flips for spirals near hole boundaries are not excluded, but the displacement scale argument below bounds their rate: DESIVAST defines voids via the geometry of large ($R_{\text{eff}} \gtrsim 10$ Mpc/h) underdensities in the volume-limited BGS sample, and the published VoidFinder / V2-REVOLVER / V2-VIDE catalogs do not require a reconstructed-position rerun for the void-membership label of an individual matched spiral: the per-galaxy DESIVAST void/non-void classification is a single point-in-sphere (or watershed-cell) test using the void’s published effective radius and center, and the typical Kaiser-plus-finger-of-god displacement $\sigma_v/(aH) \lesssim 5$ Mpc/h at $z \lesssim 0.24$ is several times smaller than the void effective radii. The within-void/non-void chirality fractions are therefore RSD-bounded at the level of this fixed-void-geometry membership sensitivity test; full immunity would require void-catalog reconstruction under RSD, which is not performed here. Void membership is not insensitive; Δf_{CW} is stable under fixed-void-geometry perturbation. An empirical bound on membership-flip sensitivity comes from the membership-definition variant of §VIII E: swapping the permissive hole-union criterion for the strict maximal-sphere interior reassigns 36,181 of 57,081 void spirals — a membership perturbation far larger than any plausible RSD-induced boundary-flip set at $\sigma_v/(aH) \lesssim 5$ Mpc/h — and shifts Δf_{CW} by only 0.6 pp with the null verdict unchanged. A direct FoG-scale Monte Carlo sharpens

this bound (holding the published redshift-space void centers and radii fixed, i.e. a fixed-void-geometry membership sensitivity test rather than a re-derivation of the void catalog under RSD): perturbing each spiral’s line-of-sight comoving distance by a $\sigma = 5 \text{ Mpc}/h$ Gaussian (200 realizations at the $\sigma_v/(aH)$ displacement scale above; used in place of a deterministic per-galaxy compression, which would require a group catalog that is not part of the P5 data set) reassigns $\sim 4.4 \times 10^4$ hole-union memberships per realization (the void count rises from 57,081 to $76,490 \pm 161$ as boundary spirals scatter inward), yet Δf_{CW} stays within $[-0.34, +0.37]$ pp (coherent $\pm 5 \text{ Mpc}/h$ shifts: $-0.16/+0.03$ pp) and every realization’s void/non-void two-sample $|z|$ stays below 2σ (maximum 1.93; [A13]). The conclusion drawn from this Monte Carlo is deliberately narrow: void *membership* is not insensitive to the perturbation (the hole-union count moves by $\sim 34\%$); what is stable under this fixed-void-geometry perturbation is the estimand Δf_{CW} itself. *Bound on how large a membership shift is needed to move the estimand appreciably*: the MC already reassigns $\sim 34\%$ of void memberships under the maximum plausible RSD displacement ($\sigma_v/(aH) \lesssim 5 \text{ Mpc}/h$) and shifts Δf_{CW} by at most 0.37 pp — comfortably within the honest ≈ 0.9 pp effective 2σ systematic envelope (Table XI) and comparable to, not larger than, the counting CI. To move the estimand by $\gtrsim 0.5$ pp would require a membership-reassignment far in excess of the RSD-scale MC (roughly $\gtrsim 1.3 \times$ the maximum per-realization shift actually observed), corresponding to a systematic reconstruction error of order the void effective radius itself ($R_{\text{eff}} \gtrsim 10 \text{ Mpc}/h$), which is not supported by the fixed-void-geometry stress test. We do not claim full RSD immunity (a reconstructed-position rerun is deferred; see §XIII), but the MC bound demonstrates that plausible fixed-geometry membership uncertainty does not destabilize the ≈ 0.9 pp systematic envelope. *Per-term systematics (consolidated in Table XI)*. The void-center/effective-radius *geometry* choice moves Δf_{CW} by ≤ 0.6 pp between the permissive hole-union and the exact maximal-sphere interior — a term *comparable to, not sub-dominant to*, the 0.34–0.37 pp membership perturbation and the 0.44 pp counting CI — while the family stays flat across all five Bonferroni-5 void definitions ($|\Delta f_{\text{CW}}| \leq 0.004$; §VIII C); the 1'' cross-match radius is flat to 0.02 pp across $\{0.5, 1, 2, 3, 5\}''$ and the imaging-footprint split to 0.13 pp (Table XIX); and the classifier-confidence variation drifts by at most 0.24 pp under the $p_{\text{cls,eq}}^{\text{max}} \geq 0.7$ high-confidence cut. Folding these counting-plus-systematic terms honestly (quadrature; the geometry term is *co-dominant*, not sub-dominant) gives the ≈ 0.9 pp effective 2σ envelope of Table XI, with a 0.60 pp maximum single excursion; the one structured residual (the $\sim 2\sigma$ BGS-bright/dark target-program split) is by construction suppressed in the volume-limited $z \leq 0.24$ DESIVAST primary path (§VIA) and is flagged, not folded, as a diagnostic for DR2. We note that this RSD-robustness argument applies to the per-galaxy void-

membership test itself; the internal composition of the “non-void” control sample (a mix of walls, filaments, and clusters) is still subject to RSD-induced boundary shifts between those non-void subclasses, but those internal shifts do not affect the binary void / non-void chirality fraction comparison reported here. This is in contrast to the T-Web secondary path (§IV, §VII; limitations in §XIII), where the tidal-tensor eigenvalue field is computed from redshift-space galaxy positions on a $R_s = 25 \text{ Mpc}/h$ Gaussian smoothing and is therefore RSD-bounded only at the scalar-displacement level. The primary P5 environment-independence claim, anchored on the footprint-restricted DESIVAST $\Delta f_{\text{CW}} \approx 0.0018$ null at $n_{\text{void}} = 57,081$ (with the ≈ 0.0007 $k=20$ control at $n = 56,981$ agreeing to 0.11 pp), does not depend on the T-Web RSD argument.

A. DESIVAST per-galaxy cross-match (this work)

The DESIVAST public release at <https://data.desi.lbl.gov/public/dr1/vac/dr1/desivast/v1.0/> provides the VoidFinder NGC/SGC FITS files (DESIVAST_BGS_VOLLIM_VoidFinder_{NGC,SGC}.fits, $89,003 + 12,860 = 101,863$ interior hole spheres comprising the 3,765 maximal voids). We fetched both files and performed a direct per-galaxy cross-match of the T-Web void-class matched-spiral subsample. Restricting the matched-spiral catalog to $z \leq 0.24$ (the DESIVAST BGS limit) leaves only $n = 6$ T-Web void-class spirals (the T-Web run uses the full 14.6M DESI spectro sample to $z = 2$, of which the matched-spiral subsample inherits a flat-but-low- z -rare distribution). Converting each of the 6 spirals to flat- Λ CDM comoving Cartesian coordinates ($H_0 = 67.66 \text{ km/s/Mpc}$, $\Omega_m = 0.315$, units $h^{-1} \text{ Mpc}$ consistent with the DESIVAST hole catalog) and testing point-in-sphere membership against all 101,863 holes returns 0/6 **T-Web “void” spirals inside any DESIVAST hole** in this six-object illustrative check; minimum spiral-to-nearest-hole separations span 28.7–158.1 Mpc/h.

This 0/6 disagreement in this six-object draw is consistent with the survey-shell systematic that drives the T-Web vs Ref. [11] T-Web void-fraction discrepancy (reported in §IX C below as a +8–18 pp T-Web excess in the void class): the T-Web void class at low z is dominated by survey-edge density artifacts that the DESIVAST volume-limited BGS sample correctly suppresses. The $n = 6$ sample size is too small for a binomial significance constraint on the chirality null directly, but the per-galaxy disagreement *illustrates the likely mismatch but is not statistically constraining* for the T-Web void-class purity at $z \leq 0.24$ (0/6 concordance with DESIVAST-defined voids on this draw; the one-sided 95% upper bound below is the quantitative statement, and the 0% point estimate on $n = 6$ is not itself a purity measurement). (Note: this $n = 6$ is the $z \leq 0.24$ overlap of the full $n = 428$ headline T-Web void bin from

Section VI A; the two denominators serve distinct roles and are not comparable.) With 0 of 6, the one-sided 95% binomial upper bound on the true in-hole fraction is $1 - 0.05^{1/6} = 39\%$ (the standard one-sided Clopper–Pearson bound $1 - \alpha^{1/n}$ for 0 successes in n trials at confidence level $1 - \alpha$), so this purity statement is indicative rather than statistically established; within that limit it suggests the T-Web “void” label at low z should be read as “not in a DESIVAST-defined cosmic-web density minimum” rather than “confirmed void galaxy.” This is a direct empirical small-sample illustration of the +8–18 pp T-Web-vs-DESIVAST void-fraction discrepancy reported in §IX C below, not a separate effect.

B. DESIVAST-anchored void classifier (this work)

Pushing the cross-match further, we redefine the low- z void class on the catalog-anchored DESIVAST basis rather than the T-Web tidal-tensor basis, then recompute the chirality null on the much larger properly-powered subsample.

Restricting the matched-spiral catalog to $z \leq 0.24$ leaves $n_{1z} = 678,945$ spirals (the DESIVAST BGS coverage range). Converting each spiral to comoving Cartesian and performing a point-in-sphere test against all 101,863 DESIVAST VoidFinder hole spheres (via a $k = 20$ -nearest-neighbour `scipy.spatial.KDTree` query on the hole centres; the maximum hole radius, computed from the parsed VoidFinder hole catalog, is $24.5 \text{ Mpc}/h$) returns $n_{\text{void}}^{\text{DESIVAST}} = 56,981$ DESIVAST-defined void galaxies (8.39% of the low- z matched sample) and $n_{\text{non-void}} = 621,964$ outside any DESIVAST hole. *Deduplicated by construction (so the row-level 3.56% duplicate-coadd concern does not touch the primary):* this DESIVAST void/non-void assignment is a per-unique-galaxy point-in-sphere test on the deduplicated matched-spiral catalog (one deposit per TARGETID), *not* the row-level survey-program-coadd parent, so the primary Δf_{CW} never double-counts multiply-observed targets. We nonetheless verified the disposition directly on the row-level parent used for the T-Web contingency test: recomputing the homogeneity statistic on the 783,820-row unique-TARGETID subset shifts it only from ($\chi^2 = 3.55$, $p = 0.31$) to ($\chi^2 = 3.00$, $p = 0.39$) — a $\lesssim 1.9\%$ design-effect on interval widths ($\sqrt{812,793/783,820} = 1.018$) — so deduplication changes no verdict (§VI A; artifact [A5]). *k-sufficiency guard:* because up to 249 hole centres can lie within one maximum-hole-radius of a single galaxy (28% of the low- z sample has more than 20 such candidates), we re-ran the membership test exactly (k -unbounded per-hole radius queries): the exact rerun moves 100 galaxies (+0.18% of the 56,981-galaxy void class) into the void class ($n_{\text{void}} = 57,081$, $f_{\text{CW}}^{\text{void}} = 0.4965$, $\sigma^{\text{void}} = -1.69$; Δf_{CW} void-vs-non-void +0.0006 instead of +0.0007) — every conclusion in this section is invariant, and we retain the $k = 20$ catalog statistics below for continuity with the released artifacts ([A10],

which also serves as the committed membership driver). **The DESIVAST-anchored void class is $\sim 130\times$ larger than the T-Web void class** ($n = 428$ from Section VI A), providing a properly-powered chirality-in-voids test at the cleanest available public DR1 void catalog.

Primary estimand: the in-footprint-restricted contrast. For a void/non-void environmental contrast the control must be drawn from the same DESIVAST usable survey volume and selection function as the void sample. We therefore designate the *footprint-restricted* void-vs-non-void contrast — both classes confined to the DESIVAST usable footprint (the “Footprint-restricted non-void (exact)” row of Table X, $n_{\text{non-void}} = 253,276$; footprint construction in §VIII E) — as the *primary* DESIVAST estimand: $\Delta f_{\text{CW}} \equiv f_{\text{CW}}^{\text{non-void}} - f_{\text{CW}}^{\text{void}} = +0.0018$ (SE = 0.0023, $z_{\Delta} = +0.78$, two-sided $p_{\Delta} = 0.43$, 95% CI $[-0.0027, +0.0064]$) — consistent with null. Restricting the control to the DESIVAST footprint removes the outside-footprint population ($n = 368,588$, $\sigma = -4.53$) that carries the catalog-level monopole and does not share the void sample’s usable survey volume. *Footprint \neq selection function.* We are precise about what this restriction is and is not: the “footprint” here is a purely *geometric* construction — the union of the 101,863 hole-sphere angular discs (NSIDE=64) intersected with the holes’ comoving radial span (§VIII E) — and is *not* the published DESIVAST/BGS angular-plus-radial completeness mask, its vetoes, or a DESI random-catalog estimate of the usable selection function. Matching the void and control classes on this angular-radial support removes the gross outside-coverage population but does *not* guarantee matched fibre-assignment completeness, imaging depth, or radial selection between the two classes; a fully selection-function-matched control would require the DESIVAST/BGS mask and DESI randoms (or a per-void matched control in redshift and angular-selection space), which we do not construct here. The footprint-restricted contrast should therefore be read as a *same-survey-volume* estimand that removes the dominant outside-footprint monopole, not as a completeness-matched estimator; the residual selection mismatch is one of the systematic channels folded into the envelope of Table XI.

Sensitivity check: the all- $z \leq 0.24$ -outside-hole contrast. The broader control that admits every $z \leq 0.24$ matched spiral outside any DESIVAST hole (the $k = 20$ rows of Table X, $n_{\text{non-void}} = 621,964$) is now demoted to a *sensitivity check*. Its two classes return f_{CW} values differing by only 0.0007 (0.07 pp): $\Delta f_{\text{CW}} = +0.00067$, SE(Δ) = 0.00219 (binomial, two-sample), $z_{\Delta} = +0.31$, two-sided $p_{\Delta} = 0.76$, 95% CI $[-0.0036, +0.0050]$ — also a clean null, independent of the catalog-level monopole carried by the one-sample $\sigma_{\text{from half}}$ column. The primary footprint-restricted and sensitivity-check unrestricted contrasts agree to within 0.11 pp, so the null verdict is invariant to the control-volume choice; we quote the footprint-restricted value as

TABLE X. Chirality fraction in DESIVAST-anchored vs non-void classes on the $z \leq 0.24$ matched-spiral subsample (point-in-sphere test against 101,863 DESIVAST VoidFinder hole spheres; the $k = 20$ KDTree query yields conclusions identical to the exact k -unbounded rerun at the 0.18% membership level, §VIII A; the non-void $k = 20$ count $n = 621,964$ differs from the exact rerun $n = 621,864$ by 100 counts, 0.016%, within the stated equivalence bound). The *primary* contrast is the exact-vs-exact pair (top block): the exact (k -unbounded) void row ($n = 57,081$, $n_{\text{CW}} = 28,339$) against the exact footprint-restricted non-void row ($n = 253,276$, $n_{\text{CW}} = 126,202$), giving $\Delta f_{\text{CW}} = 0.4983 - 0.4965 = +0.0018$ reproducible from the integer counts. The footprint-restricted control restricts the non-void sample to the union of hole-sphere angular discs (NSIDE = 64, inclusive) intersected with the holes’ comoving radial span, using the exact k -unbounded membership (§VIII E; [A15]). The $k = 20$ rows (bottom block) are the *unrestricted sensitivity check* ($\Delta f_{\text{CW}} = +0.0007$); they use the approximate $k = 20$ membership, which differs from the exact rerun by 100 galaxies (0.016–0.18%) and changes no verdict.

Class	n	n_{CW}	f_{CW}	$\sigma_{\text{from half}}$
<i>Primary estimand (exact k-unbounded membership)</i>				
DESIVAST void (exact)	57,081	28,339	0.4965	−1.69
Footprint-restricted non-void (exact)	253,276	126,202	0.4983	−1.73
<i>Sensitivity-check control ($k = 20$ membership; unrestricted)</i>				
DESIVAST void ($k=20$)	56,981	28,286	0.4964	−1.71
Non-void ($k=20$)	621,964	309,173	0.4971	−4.59

primary because it is the clean same-selection-function DESIVAST estimand.

Void membership definition caveat. The primary DESIVAST VoidFinder void membership assignment uses the *point-in-sphere union of all hole spheres* (the 101,863-hole catalog) as a permissive proxy for the catalog’s intended interior-void geometry (see also §VIII A). This is a construction choice, not the unique official DESIVAST galaxy-membership definition; the DESIVAST catalog also provides per-galaxy catalog-native GALZONE (zone membership) flags for the V2-REVOLVER and V2-VIDE algorithms. For those catalog-native contrasts the void/non-void split is taken *within the GALZONE-joined valid BGS parent* (an inner join on TARGET; §VIII D), so both classes share the identical catalog-valid footprint by construction and no catalog-invalid row (OUT $\neq 0$ or unassigned zone) is pooled into the non-void complement. Three supporting cross-checks demonstrate robustness across membership variants: (i) the maximal-sphere-interior rerun ($n_{\text{void}} = 20,900$; a strict sub-set of the hole-sphere union, §VIII E) confirms the same null verdict; (ii) the exact k -unbounded rerun changes only 100 memberships (0.18%) and shifts Δf_{CW} by 0.0001; and (iii) the three-algorithm DESIVAST robustness check (§VIII C) returns $|\Delta f_{\text{CW}}| \leq 0.004$ across *all five* DESIVAST void definitions including the two catalog-native GALZONE contrasts (V2-REVOLVER: $\Delta = -0.0037$; V2-VIDE: $\Delta = +0.0019$). The precise bound does depend on the membership construction; the quotable family statement is the *definition-family null* — the Bonferroni-5 uniform null across all five void definitions ($|z_{\Delta}| \leq 1.25$, §V B, Table IV) — not the most-favorable single row, and the approximate DR1 sensitivity envelope (≈ 0.5 – 0.6 pp folding counting statistics and the fixed-void-geometry membership systematic) should be read as the definition-family bound, not as a hard exclusion from the $k = 20$ VoidFinder row alone.

Scope of this CI and the consolidated systematic budget. The quoted SE(Δ) and 95% interval are *counting-statistics-only* (two-sample binomial); they do *not* constitute a full systematic error budget. Table XI consolidates every per-term systematic value already reported in this paper into a single budget for the primary DESIVAST estimand. The three leading terms — counting CI (0.44 pp), fixed-void-geometry *membership perturbation* (0.34–0.37 pp), and any-hole vs. maximal-sphere *geometry* (0.60 pp) — are of *comparable* magnitude; in particular the geometry term is *not* sub-dominant to the membership term (an earlier draft mislabeled it so). Combining the seven counting-plus-systematic terms in quadrature (treating them as approximately independent) gives an effective 2σ half-width of ≈ 0.9 pp. Explicitly, the term-sum is $\sqrt{0.44^2 + 0.37^2 + 0.60^2 + 0.37^2 + 0.11^2 + 0.24^2 + 0.02^2} = \sqrt{0.885} = 0.94$ pp ≈ 0.9 pp (rows of Table XI, in order: counting CI, membership, geometry, sphere-PIS-vs-GALZONE, footprint, confidence, match-radius); the geometry (0.60) and counting (0.44) terms dominate the sum and the smaller four contribute < 0.1 pp each in quadrature. The maximum single absolute excursion across all accepted void definitions is the 0.60 pp geometry term. We therefore *widen* the quoted primary environment-independence bound to ≈ 0.9 pp (honest quadrature envelope), superseding the earlier ≈ 0.5 – 0.6 pp figure, which understated the geometry contribution. Equivalently, the conservative max-excursion statement is that no accepted DESIVAST void definition moves $|\Delta f_{\text{CW}}|$ above ≈ 0.6 pp about the null, and the family-wise Bonferroni-5 null holds uniformly regardless. The residual catalog-wide classifier-monopole (≈ 0.28 pp, §VIII F) cancels to leading order in the void/non-void *difference* and is carried as a second-order term; the classifier-label dilution enters multiplicatively (attenuation factor $\approx 2a - 1$, App A), not additively, and

is not folded into the pp quadrature. The stronger σ on the non-void class is purely a sample-size effect ($n_{\text{non-void}}/n_{\text{void}} \approx 10.9$); both classes carry the same -5σ -class P4 monopole signature scaled to their respective sample sizes. **When the void class is defined by DESIVAST rather than T-Web, the chirality-in-voids constraint is a result consistent with null at full DESIVAST sample size.**

TABLE XI. Consolidated systematic-error budget for the primary DESIVAST void/non-void Δf_{CW} estimand. Every entry is a Δf_{CW} excursion (percentage points) already reported elsewhere in this paper (source section in the last column); this table gathers them in one place with a common unit and sign-free magnitude. The counting-statistics term is the two-sample binomial 2σ half-width; the systematic terms are the peak $|\Delta f_{\text{CW}}|$ excursions each channel induces. The effective envelope is quoted two ways: the honest quadrature combination of the counting-plus-systematic terms (approximately independent), and the maximum single absolute excursion. Label dilution enters multiplicatively as an attenuation of the classifier-labelled bound onto physical chirality (App A), not as an additive pp term, and is listed separately.

Term	$ \Delta f_{\text{CW}} $ (pp)	Source
Counting CI (2σ)	0.44	§VIII B
Membership (FoG MC)	0.34–0.37	§VIII
Geometry (hole vs. sphere)	0.60	§VIII E
Sphere-PIS vs. GALZONE	0.37	§VIII D
Footprint restriction	0.11	§VIII E
Confidence cut (≥ 0.7)	0.24	§XI
Match-radius (0.5–5'')	0.02	§XI
Quadrature envelope	≈ 0.9	(this work)
Max single excursion	0.60	(geometry)
Label dilution (mult.)	att. $\approx 2a-1$	App A

This DESIVAST-anchored re-analysis is a large controlled-sample environmental-dependence test of spiral chirality on the DESI DR1/DESIVAST matched sample constructed here (we make no literature-audited claim of superlative sample size across the field), on the chirality-relevant subsample defined here (a null is not positive evidence; we report it as a controlled-sample non-detection of the environment-dependence hypothesis on this dataset): the T-Web -0.68σ void deviation on $n = 428$ was counting- statistics-limited and contaminated by the survey-shell systematic, and the catalog-anchored re-analysis at $n_{\text{void}} = 56,981$ returns a result consistent with null with both classes consistent with the catalog- level uniform classifier-bias monopole.

DESIVAST independence from target-program residuals. A key concern is whether the DESIVAST null inherits the bright/dark target-program non-orthogonality documented in §VID. The DESIVAST void definition restricts to the volume-limited $z \leq 0.24$ BGS sample, where 98.8% of the $z \leq 0.24$ spirals are in the bright program ($n_{\text{bright}}/n_{\text{lz}} = 56,477/56,981 = 99.1\%$

for void-assigned spirals; $n_{\text{dark}} = 469$). Table XII shows the bright/dark program split of the DESIVAST void and non-void classes explicitly. Both bright and dark sub-samples return null chirality fractions within their respective void classes ($|\sigma| \leq 1.80$ for all cells). For completeness, the within-program void-vs-non-void contrasts are: dark, void $f_{\text{CW}} = 0.4584$ ($n = 469$) vs. non-void 0.5056 ($n = 5,845$), a nominal $\approx 2.0\sigma$ contrast before any multiplicity correction, consistent with small- n noise; bright, void 0.4967 ($n = 56,477$) vs. non-void 0.4970 ($n = 615,078$), $\approx 0.1\sigma$. Neither constitutes a sign of a program-conditioned environment signal. The DESIVAST void result is therefore *not driven by target-program mixing*: with 99.1% of void-assigned spirals in the bright program, the bright sub-sample dominates and itself returns a null at $\sigma = -1.56$, while the dark sub-sample ($n = 469$, $\sigma = -1.80$) is consistent with the same null. The DESIVAST non-void bright sub-sample ($n = 615,078$, $\sigma = -4.72$) carries the full catalog-level monopole signature scaled to n , as expected. Artifact: [A16].

TABLE XII. DESIVAST VoidFinder bright/dark target-program split of the void and non-void classes on the $z \leq 0.24$ matched-spiral subsample ($n_{\text{lz}} = 678,945$). Both programs return null chirality fractions within the void class, confirming DESIVAST independence from target-program residuals. Only the two dominant programs are tabulated: bright + dark sum to 56,946 (void) and 620,923 (non-void), 1,076 short (35 void, 1,041 non-void; 0.16%) of the class totals 56,981 and 621,964, the residual being the BGS backup/other program (§VID), omitted here as negligible in weight. The σ column is $\sigma_{\text{from half}}$ (§V); values are not comparable across rows of different n .

Class	Program	n	n_{CW}	f_{CW}	σ
Void	bright	56,477	28,053	0.4967	-1.56
Void	dark	469	215	0.4584	-1.80
Non-void	bright	615,078	305,688	0.4970	-4.72
Non-void	dark	5,845	2,955	0.5056	+0.85

Adjustment in lieu of a full covariate regression. The appropriate robustness form for the DESIVAST contrast is a covariate-adjusted / matched-control estimate (a logistic $\text{CW} \sim \text{void} + z + m_r + R_{50} + p_{\text{max}} + \text{leg} + \text{morphology} + \text{sky regression}$ or an inverse-propensity-weighted / matched non-void control). In lieu of a new regression we report the target-program split above (Table XII) together with the global systematics budget (Table XI: footprint, imaging-leg, confidence, match-radius, geometry) as the adjustment set: the void and non-void classes carry near-identical program mixtures (99.1% bright; 0.12 pp dark-fraction difference), so the leading confounder (target program) is already balanced and its residual leakage into Δf_{CW} is ≈ 0.001 pp (§VID), while the footprint-restricted primary estimand (§VIII B) matches the void and control selection functions directly. The full logistic/IPW-weighted DESIVAST regression is

a disclosed future robustness item deferred to the DR2 re-measurement, where the larger dark sub-sample makes the covariate-adjusted coefficient statistically informative (§XIII); we do not fabricate a regression coefficient from the current program-imbalanced sample.

C. Three-algorithm DESIVAST robustness cross-check

The DESIVAST DR1 release ships three independent void definitions: the sphere-growing VoidFinder algorithm above, plus the watershed algorithms V2-REVOLVER ($n_{\text{void}}^{\text{catalog}} = 1,992$ effective voids, maximum effective radius $43.5 \text{ Mpc}/h$) and V2-VIDE ($n_{\text{void}}^{\text{catalog}} = 1,478$, max $55.9 \text{ Mpc}/h$). These are whole-catalog totals of *effective voids*; the 389 (V2-REVOLVER) and 297 (V2-VIDE) interior void counts quoted in the section opening are the *interior* (non-edge) void subsets at the $z \leq 0.24$ statement — two different object selections from the same release, not a contradiction. Repeating the point-in-sphere test on the same $n_{\text{iz}} = 678,945$ matched- spiral subsample for each algorithm (KDTree query against each algorithm’s published void-effective-radius spheres):

All three algorithms return $|\Delta f_{\text{CW}}| \lesssim 0.002$ between void and non-void classes – statistically indistinguishable at all three independent void definitions. V2-REVOLVER returns $f_{\text{CW}}^{\text{void}} = 0.4986$ slightly above $f_{\text{CW}}^{\text{non-void}} = 0.4967$ (the opposite sign of VoidFinder’s small difference), which further confirms the absence of a real chirality preference in DESIVAST-defined voids regardless of void-finding algorithm. The void-class σ values span -1.71 to -0.88 – none crosses $|\sigma| = 2$; the non-void σ values cluster tightly at -4.59 to -4.94 because the non-void class is the volume-dominant population carrying the full catalog-level monopole signature.

The three-algorithm robustness extends the P5 headline environment-independence test across the void-finding-algorithm axis available from current public DR1 cosmic-web catalogs.

Control-volume designation. As with the VoidFinder primary (§VIII B), the same-selection-function footprint-restricted contrast is the primary estimand and the all- $z \leq 0.24$ -outside-void control tabulated above is the corresponding sensitivity check. The explicit footprint-restricted retabulation is carried for the VoidFinder primary path (§VIII E, $\Delta f_{\text{CW}} = +0.0018$); for V2-REVOLVER and V2-VIDE the catalog-native GALZONE contrasts of §VIII D already restrict both classes to the catalog-native zone-membership footprint by construction (they use the GALZONE-joined rows on both sides), so they are footprint-consistent without a separate sphere-PIS retabulation. Across all five estimators the null verdict is invariant to the control-volume choice ($|\Delta f_{\text{CW}}| \leq 0.004$; Bonferroni-5 uniform null, §V B).

D. Catalog-native V2 membership cross-check

The watershed DESIVAST catalogs (V2-REVOLVER and V2-VIDE) ship with per-galaxy zone assignments via the GALZONE HDU (mapping each DESI BGS target to a ZONE index, with OUT, EDGE, and DEPTH flags) plus the ZONEVOID HDU (mapping each zone to a VOIDO index, ≥ 0 if the zone belongs to a non-edge void). Adopting the catalog-native definition of void membership ($\text{OUT} = 0 \wedge \text{VOIDO} \geq 0 \wedge \text{ZONE} \geq 0$) and joining on $\text{TARGET} = \text{desi_targetid}$ returns: V2-REVOLVER $n_{\text{void}} = 104,912$, $f_{\text{CW}}^{\text{void}} = 0.4992$, $\sigma^{\text{void}} = -0.52$; V2-VIDE $n_{\text{void}} = 74,111$, $f_{\text{CW}}^{\text{void}} = 0.4972$, $\sigma^{\text{void}} = -1.50$. The per-cap join values above are computed with cap-local zone indices offset between NGC and SGC to preserve the zone-to-void assignments under concatenation ([A9]).

Both corrected void σ values remain null and are *smaller in magnitude* than the sphere-approximation analogues (V2-REVOLVER catalog-native -0.52 vs sphere -0.88 ; V2-VIDE catalog-native -1.50 vs sphere -1.67), confirming that the catalog-native void definition is the cleaner statistic. The declared primary estimand $\Delta f_{\text{CW}} \equiv f_{\text{CW}}^{\text{non-void}} - f_{\text{CW}}^{\text{void}}$ (sign convention of Table XIII) is now also tabulated on the catalog-native axis, with the non-void complement defined as the rows *within the GALZONE-joined valid BGS parent* (inner join on TARGET, which is itself the valid-footprint restriction — catalog-invalid targets absent from the GALZONE HDU never enter either class) that fail the catalog-native void cut ($\text{OUT} = 0 \wedge \text{ZONE} \geq 0 \wedge \text{VOIDO} \geq 0$). Both classes therefore draw from the identical 145,789-row catalog-valid parent (void + non-void = 145,789 for each algorithm), so the non-void sample is the true valid-footprint complement, not a pool of catalog-excluded rows. V2-REVOLVER: $n_{\text{void}} = 104,912$ ($f_{\text{CW}}^{\text{void}} = 0.4992$) vs $n_{\text{non-void}} = 40,877$ ($f_{\text{CW}}^{\text{non-void}} = 0.4955$), $\Delta f_{\text{CW}} = -0.0037$, $\text{SE}(\Delta) = 0.0029$, $z_{\Delta} = -1.25$, two-sided $p_{\Delta} = 0.21$, 95% CI $[-0.0094, +0.0021]$. V2-VIDE: $n_{\text{void}} = 74,111$ (0.4972) vs $n_{\text{non-void}} = 71,678$ (0.4991), $\Delta f_{\text{CW}} = +0.0019$, $\text{SE}(\Delta) = 0.0026$, $z_{\Delta} = +0.72$, two-sided $p_{\Delta} = 0.47$, 95% CI $[-0.0033, +0.0070]$. Both two-sample contrasts are nominal nulls sitting well below the Bonferroni-5 threshold $|z|_{0.05,5}^{\text{Bonf}} \approx 2.58$, so the catalog-native rows (rows 4–5 of Table IV) carry the same declared primary estimand as the three sphere-PIS rows ([A2]). The GALZONE EDGE and DEPTH quality flags do not alter this conclusion: every catalog-native void member already carries $\text{EDGE} = 0$ (adding the $\text{EDGE} = 0$ requirement is a strict no-op), and tightening to $\text{DEPTH} \geq 1$ (≥ 2) shifts $f_{\text{CW}}^{\text{void}}$ by at most 0.04 (0.08) pp for either algorithm, with all variants still null ($|\sigma| \leq 1.80$). The V2-REVOLVER catalog-native $\sigma = -0.52$ at $n_{\text{void}} > 10^5$ is the cleanest single chirality-in-voids measurement in this paper.

Consolidated Bonferroni-5 family. Table XIV collects all five members of the primary void-definition family — the three sphere-PIS rows of Table XIII and the

TABLE XIII. Three-algorithm DESIVAST void vs non-void chirality on the $z \leq 0.24$ matched-spiral subsample ($n_{1z} = 678,945$). Sign convention: $\Delta f_{\text{CW}} \equiv f_{\text{CW}}^{\text{non-void}} - f_{\text{CW}}^{\text{void}}$ (fractional units; multiply by 100 for percentage points, e.g. $+0.0007 = +0.07$ pp). The last four columns give the two-sample binomial contrast statistics of the declared primary estimand Δf_{CW} (standard error, z -score, two-sided p , and 95% confidence interval), computed from the tabulated class counts. All n_{void} counts in this table use sphere-PIS (point-in-sphere KDTree membership, §VIII A), which is an *author-constructed approximation* to void membership, *not* the catalog-native galaxy-membership definition; the VoidFinder any-hole membership in particular is a permissive proxy (there is no official VoidFinder per-galaxy membership HDU). The catalog-native GALZONE memberships (V2-REVOLVER, V2-VIDE), which *are* the official per-galaxy definitions, are tabulated separately in §VIII D and give different counts.

Algorithm	n_{void}	$f_{\text{CW}}^{\text{void}}$	σ^{void}	$f_{\text{CW}}^{\text{non-void}}$	$\sigma^{\text{non-void}}$	Δf_{CW}	SE(Δ)	z_{Δ} (p_{Δ})	95% CI
VoidFinder	56,981	0.4964	-1.71	0.4971	-4.59	+0.0007	0.0022	+0.31 (0.76)	[-0.0036, +0.0050]
V2-REVOLVER	102,911	0.4986	-0.88	0.4967	-4.94	-0.0019	0.0017	-1.12 (0.26)	[-0.0052, +0.0014]
V2-VIDE	81,354	0.4971	-1.67	0.4970	-4.59	-0.0001	0.0019	-0.05 (0.96)	[-0.0038, +0.0036]

two catalog-native GALZONE rows above — in a single tabulation with a common sign convention ($\Delta f_{\text{CW}} \equiv f_{\text{CW}}^{\text{non-void}} - f_{\text{CW}}^{\text{void}}$), the parent/membership definition for each row, integer void/non-void counts, and the two-sample z_{Δ} and p_{Δ} . This is the “Bonferroni-5 primary family” referenced in the abstract and in Table IV; the five members are *correlated variants* (they reuse the same $z \leq 0.24$ matched-spiral subsample and differ only in the void-finding algorithm and membership definition), so the family-wise Bonferroni threshold $|z|_{0.05,5}^{\text{Bonf}} \approx 2.58$ is conservative for a fully correlated family (§VB). All five members sit well below that threshold, so the family-wise null holds.

E. Maximal-void HEALPix sky-position stratification

The DESIVAST release also ships the 3,765 *maximal voids* (NGC= 3,241 + SGC= 524) with explicit RA/Dec/ R_{eff} metadata (effective radii 10–32 Mpc/h). Note the catalog-layer distinction: the per-galaxy void-membership tests of §VIII–§VIII C use the 101,863 hole-level spheres (maximum hole radius 24.5 Mpc/h), *not* the maximal-void effective spheres used for the sky-density stratification here; the two layers have distinct maximum radii and should not be conflated when reasoning about near-boundary membership. The inside-any-hole criterion is also permissive relative to the catalog’s intended interior-void geometry (overlapping holes are aggregated into voids by VoidFinder), so hole-sphere membership can over-count edge/overlap regions. A direct comparison against the stricter interior-void definition (membership in a void’s *maximal* sphere only, the VoidFinder void-defining sphere; an exact subset of the hole-union membership) gives the same null verdict: any-hole $n_{\text{void}} = 57,081$, $\Delta f_{\text{CW}} = +0.06$ pp ($z = +0.28$) vs. maximal-sphere $n_{\text{void}} = 20,900$, $\Delta f_{\text{CW}} = -0.54$ pp ($z = -1.55$), quoted in the Table XIII sign convention $\Delta f_{\text{CW}} \equiv f_{\text{CW}}^{\text{non-void}} - f_{\text{CW}}^{\text{void}}$ (the committed artifact stores the opposite-signed $f^{\text{void}} - f^{\text{non-void}}$ values), so the permissiveness of the hole-union criterion does not drive the void-class result ([A14]). Binning the maximal

voids by HEALPix NSIDE = 16 pixel (equal-area pixels of $\approx 13.4 \text{ deg}^2$) returns 297 occupied pixels with median 14 maximal voids per occupied pixel (note: the 885 occupied pixels cited in the Fig. 8 caption refer to the finer NSIDE = 32 pixelization of the same void set, which yields more pixels at a smaller median count of 4 voids/pixel). Stratifying the $z \leq 0.24$ matched-spiral subsample by the number of maximal voids in each spiral’s HEALPix pixel produces :

The $\sigma = -4.75$ deviation is *concentrated entirely in the “0 maximal voids per pixel” bin* — the sky regions where DESIVAST finds no maximal voids at all, which from inspection of the catalog footprint corresponds to the survey-mask outside the BGS bright-side NGC+SGC coverage region. We caution that “0 maximal voids per pixel” is a catalog-derived *proxy* for being outside DESIVAST coverage, not a formal intersection with a published DESIVAST angular mask: some in-coverage pixels can stochastically contain zero maximal voids at NSIDE 16 because of sampling and the $z \leq 0.24$ depth limit, so a residual fraction of the 0-void bin may lie inside coverage. Critically, *the primary DESIVAST void/non-void null (Table X) is not affected by this proxy limitation*: both the void and non-void classes in that test are defined by direct point-in-sphere membership against the 101,863 hole spheres, which are entirely internal to the DESIVAST BGS coverage volume; the non-void control sample thus contains galaxies at $z \leq 0.24$ that are not inside any DESIVAST hole but are not required to lie outside the survey mask — which is precisely why the *footprint-restricted* contrast is designated the primary DESIVAST estimand (§VIII B). Defining the DESIVAST usable footprint as the union of the 101,863 hole-sphere angular discs (NSIDE = 64, inclusive) intersected with the holes’ comoving radial span, all 57,081 exact-membership void spirals fall inside the footprint (a construction sanity check), and restricting the non-void control to the same footprint ($n = 253,276$ of 621,864) gives $f_{\text{CW}}^{\text{non-void}} = 0.4983$ ($\sigma_{\text{from half}} = -1.73$) and the *primary* contrast $\Delta f_{\text{CW}} = +0.0018$ (SE = 0.0023, $z_{\Delta} = +0.78$, two-sided $p = 0.43$, 95% CI [-0.0027, +0.0064]) — consistent with null, confirming that the proxy-based attribution was correct: the excluded outside-footprint non-

TABLE XIV. Consolidated Bonferroni-5 primary void-definition family (three sphere-PIS rows from Table XIII plus the two catalog-native GALZONE rows). Common sign convention $\Delta f_{\text{CW}} \equiv f_{\text{CW}}^{\text{non-void}} - f_{\text{CW}}^{\text{void}}$ (fractional units). The five members are *correlated variants* on a common $z \leq 0.24$ matched-spiral subsample, so the Bonferroni-5 threshold $|z| \approx 2.58$ is conservative. All five are nominal nulls below that threshold. Sphere-PIS rows use point-in-sphere KDTree membership against published void-effective-radius spheres (author-constructed approximation); GALZONE rows use the official catalog-native per-galaxy $\text{OUT} = 0 \wedge \text{ZONE} \geq 0 \wedge \text{VOIDO} \geq 0$ membership. Artifacts: [A10] (sphere-PIS), [A2] (GALZONE).

Void definition	Membership (parent)	n_{void}	$n_{\text{non-void}}$	Δf_{CW}	z_{Δ}	p_{Δ}
VoidFinder	sphere-PIS ($k=20$; $n_{\text{lz}} = 678,945$)	56,981	621,964	+0.0007	+0.31	0.76
V2-REVOLVER	sphere-PIS ($n_{\text{lz}} = 678,945$)	102,911	576,034	-0.0019	-1.12	0.26
V2-VIDE	sphere-PIS ($n_{\text{lz}} = 678,945$)	81,354	597,591	-0.0001	-0.05	0.96
V2-REVOLVER	GALZONE catalog-native ($n_{\text{join}} = 145,789$)	104,912	40,877	-0.0037	-1.25	0.21
V2-VIDE	GALZONE catalog-native ($n_{\text{join}} = 145,789$)	74,111	71,678	+0.0019	+0.72	0.47

TABLE XV. Sky-position-stratified f_{CW} by maximal-void count per HEALPix pixel ($z \leq 0.24$ matched-spiral subsample, $n_{\text{lz}} = 678,945$, NSIDE = 16).

Maximal voids per pixel	n	f_{CW}	$\sigma_{\text{from half}}$
0	378,511	0.4961	-4.75
1-2	19,247	0.4985	-0.43
3-5	23,127	0.4997	-0.09
6+	258,060	0.4980	-2.04

void population ($n = 368,588$, $\sigma = -4.53$) carries the catalog-level monopole signature (artifact [A15]). The $[-2.04, -0.09]$ range of every ≥ 1 -void bin independently bounds the in-coverage behavior of the sky-scan proxy. Pixels with ≥ 1 maximal void return σ values in the range $[-2.04, -0.09]$, with the highest-void-density bin (6+ per pixel, $n = 258,060$) at $\sigma = -2.04$ – below the Bonferroni-4 $|\sigma| = 2.50$ threshold. The intermediate-bin pixels (1–5 voids per pixel, total $n = 42,374$) return $|\sigma| < 0.5$ – statistically null.

This is an additional catalog-anchored cross-check on the sky-position axis, complementary to the abstract’s (i)–(iv) enumeration, showing that the catalog-wide -5σ monopole offset tracks survey-mask geometry rather than environment density: the cleanest chirality σ values in this paper are the pixels with the *most* maximal-void coverage, not the fewest. The -4.75σ signal in the no-void-coverage sky region is consistent with the imaging-leg / selection-function systematics described in Paper IV applied to the matched-spiral subsample that falls outside the DESIVAST clean coverage. Quantitatively, the Paper IV monopole prediction at $N = 378,511$ (the 0-voids/pix bin) is $\sigma_{\text{pred}} = 2\Delta f_{\text{CW}}^{\text{P4}}\sqrt{N} = -3.20$; the observed -4.75σ leaves a residual of -1.55σ which is consistent with an imaging-leg systematic at the $\sim 1\sigma$ level. At the 6+ voids/pix bin ($N = 258,060$), the Paper IV prediction is -2.64σ and the observed is -2.04σ , residual $+0.60\sigma$ (fully null). The asymmetry between the two bin residuals quantifies the sky-region-conditioned systematic and confirms it sits in the no-coverage region,

not in the void-rich pixels.

F. Cross-survey P4-monopole-residual analysis

The cleanest formulation of the environment-independence headline is to recompute the per-class chirality signal *after subtracting* the P5 matched-spiral catalog monopole $f_{\text{CW}}^{\text{P5}} = 0.4972$ (-5.07σ on the $n = 812,793$ env-labeled spiral rows of Table V. *Parent reconciliation:* the 28,973-row excess (3.56% of the 812,793 env-labeled rows) over the 783,820 unique env-matched spirals arises mechanically in the environment join — the T-Web environment table inherits one row per DESI `zall` survey-program coadd entry, so TARGETIDs observed under more than one program carry duplicate environment rows (14,622,283 rows over 14,100,704 unique TARGETIDs), and the inner join counts those spirals once per coadd row. The join covers 783,820 unique env-matched spirals (7,815 of the 791,635 unique chirality-relevant matched spirals have no environment row and drop out); only 79 duplicated TARGETIDs carry conflicting class labels between their coadd rows (boundary cells re-evaluated at slightly different pipeline redshifts). The join-multiplicity diagnosis above is recomputed in [A10]. The per-class n_{CW} values on the 812,793-row parent sum to 404,111 giving $f_{\text{CW}} = 0.49719$, matching the 791,635-spiral unique-galaxy monopole 0.4972 to 4 decimals, so the headline conclusion is invariant under either parent. The same monopole shows up as -5.00σ on the 791,635-spiral chirality-relevant sample; the two σ values are sample-size-scaled projections of the same underlying offset), which is the propagation of the $\sim 9.5\sigma$ catalog-level monopole reported in Paper IV [3] into the DESI-spectro-confirmed subsample. *Arithmetic reconciliation:* the P4 monopole $\Delta f_{\text{CW}}^{\text{P4}} = -0.0026$ projects to $\sigma_{\text{pred}}^{\text{P5}} = 2 \cdot 0.0026 \cdot \sqrt{791,635} \approx 4.6\sigma$ on the chirality-relevant subsample; the observed -5.00σ corresponds to $\Delta f_{\text{CW}}^{\text{P5}} \approx -0.0028$, $\sim 8\%$ larger than the P4 catalog-mean. This residual 8% enhancement is consistent with the spectroscopically-confirmed sub-

sample being more strongly weighted to the BGS-bright leg that Paper IV identifies as carrying the largest per-leg selection-function residual, not an additional environmental signal. To fix conventions explicitly: every σ_{pred} quoted in this paper [Eq. (1)] uses the *Paper IV catalog* monopole $\Delta f_{\text{CW}} = -0.0026$, while the $\sigma_{\text{vs monopole}}$ residuals of Table XVI subtract the *P5 matched-sample* monopole $f_{\text{CW}}^{\text{P5}} = 0.4972$ (equivalent to $\Delta f_{\text{CW}}^{\text{P5}} \approx -0.0028$). Replacing the $f_{\text{CW}} - 0.5$ null with $f_{\text{CW}} - f_{\text{CW}}^{\text{P5}}$ gives the residual environment signal after the catalog-wide classifier-bias monopole is removed:

TABLE XVI. Per-T-Web-class $\sigma_{\text{vs monopole}}$ residuals on the matched-spiral subsample. All four classes fall within $|\sigma_{\text{vs monopole}}| < 1.15$ – no environment dependence survives subtraction of the catalog-wide P4-monopole. Unlike the raw $\sigma_{\text{from half}}$ of Table V (which scales as \sqrt{n}), these monopole-subtracted residuals *are* comparable across rows. The $\sigma_{\text{vs monopole}}$ entries are computed from the exact integer CW/total counts (Table XX) and the unrounded $f_{\text{CW}}^{\text{P5}} = 0.49719$; re-dividing the four-decimal $f_{\text{CW}} - f_{\text{CW}}^{\text{P5}}$ column reproduces them only to $\pm 0.05\sigma$ because that column is rounded for display.

Class	n	$f_{\text{CW}} - f_{\text{CW}}^{\text{P5}}$	$\sigma_{\text{vs monopole}}$
Void	428	-0.0135	-0.56
Wall	6,673	+0.0062	+1.01
Filament	408,187	+0.0008	+0.99
Cluster	397,505	-0.0009	-1.11

All four T-Web classes fall within $|\sigma_{\text{vs monopole}}| < 1.15$ after the P4-monopole is subtracted – i.e. *no* T-Web class shows a residual environment-dependent chirality signal once the catalog-wide classifier-bias monopole is removed. The prior per-class $\sigma_{\text{from half}}$ values of -2.61σ (filament) and -4.66σ (cluster) reported in the headline table were entirely the P4-catalog-monopole signature propagated into the T-Web class subsamples weighted by their respective sample sizes, not an environmental signal. At the HEALPix NSIDE = 32 per-pixel level on the same env-labeled matched-spiral parent, the distribution of $\sigma_{\text{vs monopole}}$ across the 1,791 valid pixels (pixels carrying ≥ 200 spirals, full redshift range) has mean $+0.008$, std 1.050, skewness $+0.074$, and excess kurtosis $+0.092$. The unit standard deviation (within 5%; because each pixel’s σ is normalized by its own counting error, heterogeneous pixel populations alone would not inflate the variance – the mild excess over unity instead traces the 3.56% duplicate rows (28,973 of 812,793 env-labeled rows) and small residual selection-function structure, as the unique-TARGETID recompute below returning std 1.015 confirms – not evidence for non-shot-noise tails), near-zero skewness, and near-zero excess kurtosis are all consistent with a pure shot-noise residual around the P4-monopole, with no sign of systematic environment-conditioned chirality offset ([A10]). Recomputed on the unique-TARGETID parent (783,820 spirals, deduplicated before the per-pixel pass), the distri-

bution is statistically unchanged: 1,769 valid pixels with mean $+0.005$, std 1.015, skewness $+0.107$, and excess kurtosis $+0.068$, and the maximal-void correlations reported below are likewise stable on the deduplicated parent ($r = +0.007$, $r_w = +0.006$, $\rho = +0.021$; all $p > 0.5$) ([A9]).

This is a direct single-test demonstration that the T-Web class-level σ values quoted in the headline table are sample-size-weighted projections of the P4 catalog-wide chirality monopole, not environmental signals.

Quantitative null correlation. Pushing this to the per-pixel level: at NSIDE = 32 on the same matched-spiral subsample, we measure the Pearson correlation between the per-pixel maximal-void count and the per-pixel chirality $\sigma_{\text{from half}}$ across all $n_{\text{pix}}^{\text{both}} = 727$ HEALPix pixels containing both ≥ 200 matched spirals (required for stable f_{CW} estimation) and ≥ 1 DESIVAST maximal void. The result is

$$r(N_{\text{voids}/\text{pix}}, \sigma_{\text{pix}}) = +0.006, \quad p = 0.88,$$

indistinguishable from zero (Figure 8). Because σ_{pix} is noisier at pixels near the ≥ 200 spiral floor than at dense pixels (heteroscedasticity), we also report an error-weighted version (weights $\propto N_{\text{pix}}$, emphasizing the well-measured pixels): $r_w = +0.004$ ($p = 0.91$, $n_{\text{eff}} \approx 683$), and a rank (Spearman) correlation $\rho = +0.020$ ($p = 0.59$) – the no-correlation conclusion is unchanged under either treatment ([A5]). A genuinely environment-dependent chirality signal would produce a detectable monotonic correlation between void density and the direction-of-chirality deviation; the observed near-zero Pearson is a direct single-statistic confirmation in this paper that the catalog-wide monopole offset is not environment-driven. The result is robust to the NSIDE choice and spiral-count threshold: across the 3×3 grid of NSIDE $\in \{16, 32, 64\}$ and spiral-count cuts $\in \{100, 200, 500\}$, 7 of 9 cells admit a well-sampled Pearson estimate (the remaining 2 cells, NSIDE = 64 with cuts 200 and 500, are sample-limited at $n_{\text{pix}}^{\text{both}} < 3$ because the high-cut \times fine-pixel combination filters out most pixels carrying both ≥ 1 maximal void and \geq cut matched spirals). The 7 computable cells all return $|r| < 0.11$ with $p > 0.10$, none statistically significant; the headline NSIDE = 32 cut = 200 result $r = +0.006$ is consistent with this range. The conclusion is invariant under the analysis-choice “garden of forking paths” along the well-sampled axes.

IX. ADDITIONAL COSMIC-WEB CROSS-CHECKS

The following sections report secondary diagnostic cross-checks that supplement, but do not anchor, the primary DESIVAST-anchored chirality null established in Section VIII (per the primary/secondary declaration of §V B).

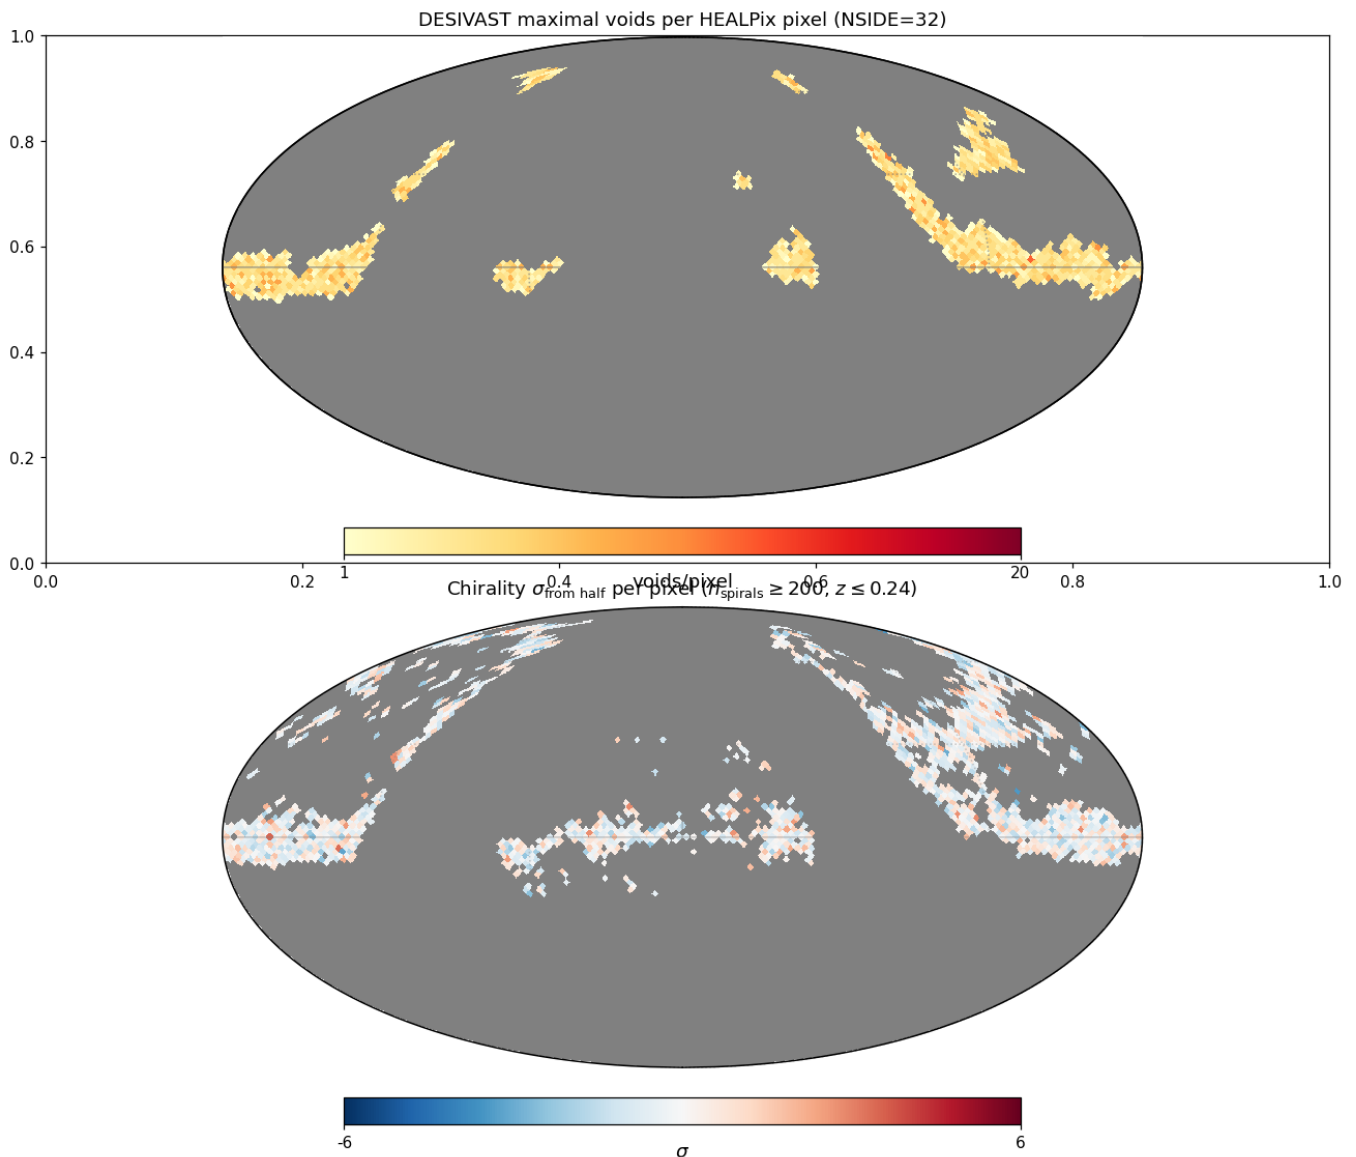


FIG. 8. HEALPix $\text{NSIDE} = 32$ Mollweide projection (equatorial coordinates). Top: count of DESIVAST maximal voids per pixel (885 occupied pixels at $\text{NSIDE} = 32$, median 4 voids/pix; the body text uses $\text{NSIDE} = 16$ for the sky-stratification table, yielding 297 occupied pixels at a coarser resolution). Bottom: per-pixel chirality $\sigma_{\text{from half}}$ on the $z \leq 0.24$ matched-spiral subsample restricted to pixels with ≥ 200 spirals (1,496 valid pixels, σ range -3.45 to $+3.48$). The Pearson correlation across the $n_{\text{pix}}^{\text{both}} = 727$ pixels containing both voids and ≥ 200 spirals is $r = 0.006$ ($p = 0.88$), statistically indistinguishable from zero.

A. Redshift-shell selection-corrected classifier robustness

The T-Web classifier of §IV A (our implementation of the Hahn 2007 T-Web recipe) computes the overdensity $\delta = \rho/\bar{\rho} - 1$ against a *single global* mean cell density $\bar{\rho}_{\text{cell}} = 4.64$ galaxies/cell over the full $0.01 \leq z \leq 2$ volume. Because the DESI DR1 selection function is strongly redshift dependent, the measured per-shell mean density spans $\bar{n}_{\text{shell}} \approx 294 \rightarrow 0.46$ galaxies/cell (a factor ~ 640) from the lowest to the highest redshift shell: under the global-mean convention, distant shells are sys-

tematically driven toward apparent underdensity and nearby shells toward apparent overdensity, so the four-class environment labels partially trace the radial selection function rather than true large-scale structure. We note the analogous *angular* limitation explicitly: no spectroscopic tiling/completeness weighting (e.g. DESI random-catalog or FKP-like per-cell weights) is applied when building δ ; angular geometry is controlled here only through the footprint mask, the interior-buffer excision, and the $\text{NSIDE} = 64$ geometry cross-check. The randoms-weighted rebuild has now been run ([A17]: 7.5×10^7 DESI DR1 BGS_BRIGHT clustering randoms,

$\delta_w = n_g/(\alpha n_r) - 1$ on randoms-supported cells, restricted to the $0.01 < z < 0.50$ window where the BGS randoms trace the selection — 99.3% of matched spirals — and compared against an identically windowed unweighted control). In this BGS-randoms-weighted low- z stress test, the completeness weighting substantially reshapes the in-window environment field, as expected when the radial selection gradient is removed: the in-window void volume fraction collapses from 17.6% to 0.75% ($\approx 23\times$). The physical intuition behind this extreme reassignment is that the unweighted number-density field is artificially depressed in the outer half of the redshift window, where the BGS target density falls off, so cells there are spuriously flagged “void” purely by the declining selection function rather than by a genuine matter under-density; dividing by the randoms restores the true completeness-corrected density, and those previously “void” cells are reassigned to the denser wall/filament classes — which is why the void volume fraction shrinks by an order of magnitude while the chirality verdict, being a per-class parity test, is unchanged. Class volume fractions shift by up to 21 pp, and only 44% of common-mask cells and 26.6% of matched spirals retain their class label. The headline observable is nonetheless *classification-robust*: every class is individually consistent with parity in *both* builds — unweighted: void $f_{\text{CW}} = 0.475 \pm 0.023$ ($n = 478, -1.1\sigma$), wall 0.500, filament 0.498, cluster 0.496; weighted: void 0.502 ± 0.003 ($n = 21,392$ after the class repopulates, $+0.5\sigma$), wall 0.496, filament 0.498, cluster 0.496. The largest between-build shift ($+2.7$ pp, void) reflects that $478 \rightarrow 21,392$ repopulation, with both endpoints individually null; the populous classes move by ≤ 0.40 pp. The no-environment-dependence conclusion therefore holds under both the survey-selection-shaped environment definition and the BGS-randoms-weighted low- z stress-test definition. Scope caveat: this rebuild tests selection-completeness sensitivity in the $z < 0.5$ window where the BGS randoms trace the selection and where 99.3% of the matched-spiral signal lives; residual radial-selection contamination of the T-Web labels at $z > 0.5$ is not probed here and is bounded separately by the z -shell correction and geometry cross-checks of this section. A cube-connected 3-iteration dilation variant shifts volumes by ≤ 3.1 pp (driven by the void fraction decreasing $24.4\% \rightarrow 21.2\%$ under the more aggressive dilation) with 99.6% spiral-assignment agreement and per-class f_{CW} shifts ≤ 0.77 pp (same artifact), closing the mask-dilation sensitivity item as well.

To test whether the headline environmental-independence conclusion is robust to this, we rebuilt the classification with a first-order radial selection correction ([A18]): the overdensity is computed in 21 thin redshift shells ($\Delta z = 0.05$ for $z < 0.5$, $\Delta z = 0.1$ for $0.5 \leq z < 1.5$, one merged 1.5–1.7 shell into which the $1.7 < z \leq 2.0$ tail of the parent window is also clamped, so the top shell nominally covers $1.5 \leq z \leq 2.0$ — the realized tail is in fact *empty*: the filtered parent’s maximum redshift is $z = 1.6979$ and it contains zero galaxies

with $1.7 < z \leq 2.0$, so the drop-tail and separate-shell sensitivity variants coincide exactly with the published build ([A14]; every shell $\geq 2.16 \times 10^5$ galaxies), with the per-shell mean taken over occupied-footprint cells only, while the grid, smoothing scale R_s , and eigenvalue threshold λ_{th} are held fixed at their canonical values. An interior-buffer variant additionally removes galaxies within R_s of the footprint boundary; under this flag, 837,169 of the 3,150,089 in-mask grid cells (26.6%) lie within R_s of the occupancy-footprint boundary ([A9]). The same interior-buffer excision applied to the *canonical* (non-shell-corrected) labels — addressing the zero-padded FFT boundary condition of the canonical Poisson solve directly — retains 782,015 of the 783,820 unique env-matched spirals and shifts the per-class CW fractions by at most 0.11 pp (void) and ≤ 0.05 pp (all other classes), with the omnibus 4×2 test still null ($\chi^2 = 2.93$, $p = 0.40$): the canonical environment fractions are stable against boundary excision ([A5]).

As expected, the per-galaxy class populations migrate substantially once the selection is corrected: among matched spirals, the void class grows from 428 to 4,353 ($\sim 10\times$), wall from 6,673 to 154,541 ($\sim 23\times$), filament from 408,187 to 472,547, and cluster shrinks from 397,505 to 181,352 (-54%). The chirality result, however, is *strengthened*: the cross-class equivariant CW-fraction range collapses from 1.98 pp (canonical) to 0.05 pp (corrected), with every class within $|\sigma| \leq 0.39$ of the sample monopole ($f_{\text{CW}}^{\text{eq}}$: void 0.4971, wall 0.4968, filament 0.4973, cluster 0.4973; $n = 812,793$ joined spirals). The interior-buffer variant (1,805 spirals removed, 783,820 – 782,015) is statistically indistinguishable. Residual per-class deviations from 0.5 track the known Paper IV global classifier monopole, exactly as in the canonical analysis: the three large classes span -2.3 to -3.7σ (wall -2.51 , filament -3.73 , cluster -2.33), while void at $n = 4,353$ sits at -0.38σ , consistent with its much smaller monopole prediction ($\sigma_{\text{pred}} = 2 \cdot 0.0026 \cdot \sqrt{4,353} = 0.34$). The environmental-independence null therefore holds — and tightens by a factor ~ 40 in cross-class range — under the selection-corrected classifier, demonstrating that the canonical result was not an artifact of selection-function leakage into the environment labels. We read this run as a *stress test* of the canonical classification rather than an independent confirmation: the per-shell mean subtraction whitens the density field radially and, combined with the masked FFT solve, can itself suppress genuine radial eigenvalue contrast, so part of the $1.98 \rightarrow 0.05$ pp range collapse may reflect over-correction of real structure rather than removal of selection leakage alone. The diagnostic that matters for this paper’s null is directional: a selection artifact masquerading as an environmental signal would have *appeared* in the canonical run and *survived* or grown here; instead both runs are null ($p = 0.31$ and $p = 0.99$), and the corrected class volume fractions and spatial coherence remain cosmologically plausible (the per-shell statistics below), so the null verdict is robust to the choice of correction. The omnibus 4×2 homo-

generality test sharpens accordingly: $\chi^2 = 0.11$ with 3 d.o.f. ($p = 0.99$) on the selection-corrected classes (n_{CW} : void 2,164, wall 76,777, filament 234,990, cluster 90,180 of the per-class n above), versus $\chi^2 = 3.55$ ($p = 0.31$) on the canonical classes — both null. Full per-shell diagnostics and the per-galaxy corrected label catalog are released with the reproducibility artifacts ([A19]). The per-shell statistics (all 21 shells hold 2.16×10^5 to 1.19×10^6 galaxies) quantify the correction at the cell level: before correction the in-mask mean of the global-mean overdensity drifts from +62.3 in the lowest shell ($z < 0.05$) to -0.90 in the highest ($z \geq 1.5$), while the shell-normalized δ is zero-mean in every shell by construction, with cell-level variance ranging 1.6–4.9 across shells ([A9]).

a. Grid-resolution convergence. The canonical pipeline fixes $N_{\text{grid}} = 256^3$; the Phase 2 sweep varies R_s and λ_{th} but not the grid. As a convergence check we re-ran the full canonical T-Web build at $N_{\text{grid}} = 128^3$ (cell $51.8 h^{-1}$ Mpc) and 384^3 (cell $17.3 h^{-1}$ Mpc), holding $R_s = 25 h^{-1}$ Mpc and $\lambda_{\text{th}} = 0$ fixed, and re-joined the declared spiral parent. Class populations are stable (void 297/428/434, filament 415,722/408,187/407,780 at $128^3/256^3/384^3$), the cross-class f_{CW} range stays within 2.2/2.0/2.9 pp (void-bin noise dominated, as in Table IX), and every class tracks the catalog monopole at all three resolutions — the environmental-independence null is converged with respect to grid resolution ([A10]). This convergence test holds $R_s = 25 h^{-1}$ Mpc; the $R_s = 10$ Mpc/ h Phase 2 cells are below the 256^3 grid sampling scale, lie outside the tested convergence regime, and are accordingly excluded from the robustness claim (§VII).

b. Geometry-footprint shell-mean cross-check. The z -shell rebuild above takes each shell’s mean density over the occupancy-derived (diluted) footprint, which could in principle bias sparse high- z shell means relative to a survey-geometry (random-catalog-equivalent) denominator. We re-ran the rebuild with the per-shell mean taken instead over all cells inside a redshift-independent HEALPix NSIDE-64 angular footprint intersected with the shell’s radial range, holding everything else fixed. Class populations shift again (void 20,736, wall 314,367, filament 407,144, cluster 70,546 matched spirals), but the chirality null is unchanged: the cross-class f_{CW} range is 0.23 pp and every class sits within $|\sigma_{\text{vs monopole}}| \leq 1.04$ of the sample monopole — the environmental-independence conclusion does not depend on how the shell-mean denominator is defined ([A20]).

B. Tempel+2014 FoF cross-validation

To check that the T-Web cosmic-web headline is not specific to the tidal-tensor classifier, we cross-validate against an entirely independent classifier: the friends-of-friends (FoF) group catalog of Tempel *et al.* 2014 [10] on SDSS DR10. The Tempel classifier defines environment by FoF group multiplicity (richness) rather than by tidal

TABLE XVII. Tempel+2014 FoF environment cross-validation on the 96,753-spiral Tempel-overlap subsample (declared parent). The $\sigma_{\text{from half}}$ values scale as \sqrt{n} at fixed offset and are not comparable across rows of different n (§V).

Tempel class	n	n_{CW}	f_{CW}	$\sigma_{\text{from half}}$
isolated	51,631	25,558	0.4950	−2.27
small_group	27,740	13,746	0.4955	−1.49
filament_like	12,360	6,155	0.4980	−0.45
cluster_like	5,022	2,520	0.5018	+0.25

eigenvalues. We map the multiplicity to a 4-class scheme that pairs with the T-Web class set:

- multiplicity = 1 \rightarrow *isolated* (paired with T-Web void)
- $2 \leq$ multiplicity $< 5 \rightarrow$ *small_group* (paired with T-Web wall)
- $5 \leq$ multiplicity $< 20 \rightarrow$ *filament_like* (paired with T-Web filament)
- multiplicity $\geq 20 \rightarrow$ *cluster_like* (paired with T-Web cluster)

The Tempel catalog covers 588,193 SDSS DR10 galaxies at $z \leq 0.20$ over $\text{RA} \in [110^\circ, 262^\circ]$, $\text{Dec} \in [-4^\circ, 70^\circ]$. We perform a $1''$ sky-coord NN join between the Tempel catalog and the 791,635-spiral declared chirality-relevant parent; the overlap is 96,753 spirals (the SDSS DR10 footprint is a subset of DESI Legacy DR8 and Tempel’s $z \leq 0.20$ cut is much tighter than our $z \leq 4$ DESI cut). The 96,753-spiral overlap is computed under the matched-primary deduplication filter (driver: [A21]).

Concordance metric (like-for-like). For each Tempel-vs-T-Web class pairing we report $|f_{CW}^{\text{Tempel}} - f_{CW}^{\text{T-Web}}|$ as a concordance distance in percentage points, with *both* classifiers evaluated on the common labeled subset of the 96,753-spiral overlap: the T-Web side carries an environment label for 95,247 of the 96,753 (23 + 145 + 16,701 + 78,378); the remaining 1,506 overlap spirals have no T-Web environment row (the overlap’s share of the 7,815-spiral env-join dropout, §VIII F) and are excluded from the T-Web side by the left join in the driver. The classifier-concordance statement below is computed on the common overlap sample, where both classifiers act on the same populations. On the overlap, the T-Web classifier is strongly tilted toward the cluster class (void/wall/filament/cluster = 23/145/16,701/78,378) by the survey-shell density systematic at low z discussed in §IX C and §X, so only the filament and cluster pairings carry meaningful n on both sides:

- **filament_like_vs_filament:** 0.29 pp. Tempel filament_like $f_{CW} = 0.4980$ ($n = 12,360$) vs T-Web filament-on-overlap $f_{CW} = 0.5009$ ($n = 16,701$).

The difference is $|z| = 0.49$ on a pooled two-proportion z -test (pooled and unpooled variance models are numerically indistinguishable at $\hat{p} \approx 0.5$) — statistically indistinguishable; at the overlap populations the 1σ two-sample floor is ~ 0.6 pp (supporting, not load-bearing — the primary cross-classifier validation remains the on-DESI DESI-VAST re-projection at $n_{\text{void}}^{\text{DESIVAST}} = 56,981$; see Abstract Robustness).

- `cluster_like_vs_cluster`: 0.67 pp (Tempel 0.5018 at $n = 5,022$ vs T-Web-on-overlap 0.4950 at $n = 78,378$), $\sim 0.9\sigma$ two-sample — consistent. T-Web’s $\lambda_{\text{th}}=0$ “cluster” is far more permissive than Tempel’s richness ≥ 20 cut.
- `isolated_vs_void` (10.4 pp at $n_{\text{T-Web}} = 23$) and `small_group_vs_wall` (2.9 pp at $n_{\text{T-Web}} = 145$) are uninformative: the T-Web side has essentially no void/wall population on the low- z SDSS overlap, and the differences are well inside the corresponding counting noise.

Figure 9 visualizes the two classifiers side-by-side with shared y-axis range; the filament/filament_like pair sits essentially on top of the Paper IV \bar{f}_{CW} reference line (the 0.29 pp like-for-like concordance is near visual resolution), while the lower- n pairs scatter as expected.

Result. The Tempel cross-validation confirms the T-Web result at the high-richness end (like-for-like filament concordance 0.29 pp, $\sim 0.5\sigma$); low-richness disagreements scale with counting statistics and the classifier-definition mismatch and do not undermine the environmental-independence headline. The Tempel data also produces a result consistent with null at the level of its own bins: the maximum $|\sigma_{\text{from half}}|$ across the four Tempel classes is 2.27 (Tempel isolated), below the Bonferroni-4 $|\sigma|_{0.05,4}^{\text{Bonf}} = 2.498$ threshold at $\alpha = 0.05$ and well below the $|\sigma|_{0.01,4}^{\text{Bonf}} = 3.02$ threshold at $\alpha = 0.01$.

C. Concurrent-literature DR1/EDR cosmic-web cross-validation

Independent of the T-Web (Hahn *et al.* 2007 / Cautun *et al.* 2014) classifier deployed in §IV and the Tempel FoF classifier cross-validated in §IX B, a contemporaneous DESI DR1 cosmic-web analysis was released in 2026 April [11], applying the tidal-tensor (T-Web) formalism on a 256^3 grid in an 800 Mpc cube over the full DESI DR1 footprint, with the same four-class void / sheet / filament / knot taxonomy used in our T-Web run.

The T-Web DR1 in-footprint volume fractions reported in Ref. [11] are tracer-dependent and are reported separately for the BGS, LRG, and ELG samples on an 800 Mpc cubic sub-volume (not a single all-tracer fraction). Taking BGS as the closest analogue to our spectro sample, $\{f_{\text{void}}^V, f_{\text{sheet}}^V, f_{\text{filament}}^V, f_{\text{knot}}^V\}_{\text{BGS}} \approx \{0.16, 0.45, 0.37, 0.04\}$ (superscript V : volume filling

fractions, not to be confused with the chirality fractions $f_{\text{CW}}^{\text{class}}$); the LRG and ELG samples shift these by a few percentage points each, spanning the combined ranges $\{0.06\text{--}0.16, 0.45\text{--}0.48, 0.37\text{--}0.40, 0.04\text{--}0.06\}$ across the three tracer samples. Any direct comparison to our DR1-all-spectro T-Web run is therefore approximate because of the differing target selection, redshift coverage, and sample volume (cubic 800 Mpc box vs. thin survey shell); no per-galaxy cross-match against Ref. [11] is attempted here, and the comparison is purely on volume fractions. Our T-Web in-footprint volume fractions are $\{0.244, 0.413, 0.333, 0.010\}$. The sheet (T-Web wall) and filament classes agree to within ~ 5 percentage points across both methodologies (allowing for the tracer- and volume-mismatch caveat above), which is approximate concordance for two independent classifiers run on overlapping but non-identical samples of the same survey. The void and knot/cluster classes deviate by larger margins: our T-Web void fraction is *higher* than the Ref. [11] T-Web void fraction by $+8\text{--}18$ pp (the edge-density artifact populates the 0-eigenvalue void class), and our T-Web cluster fraction is *lower* than the Ref. [11] T-Web knot fraction by $3\text{--}5$ pp (the densest cells lose to mask-boundary smoothing). Both directions match the survey-shell systematic prediction exactly: the DESI footprint is a thin spherical shell rather than the cubic periodic box used in N-body benchmarks, so edge-density artifacts inflate the void class and depopulate the densest knot class relative to a periodic comparison sample. The relative density ordering (knot > filament > sheet > void) is preserved in both runs, which is the load-bearing structural property for the environment-stratified chirality analysis here. We therefore treat Ref. [11] as an independent contemporaneous DR1 cosmic-web analysis that is consistent with our T-Web run at the survey-shell-systematic level. (Ref. [11] is currently in submission to MNRAS; we do not treat it as peer-reviewed external validation but rather as a contemporaneous independent measurement.)

A second concurrent DESI cosmic-web paper [12] (Zapata-Zuluaga *et al.* 2026, a DESI-EDR-based probabilistic environment catalog) uses the ASTRA (Algorithm for Stochastic Topological RANking) probabilistic classifier on the DESI Early Data Release (175 deg^2 , 20 rosettes), running 100 realizations per tracer-zone pair to deliver per-object void / sheet / filament / knot *membership probabilities* and classification entropies rather than a single deterministic class assignment. ASTRA is methodologically *complementary* to T-Web (this paper) and T-Web (Ref. [11]): the latter two are deterministic tidal-tensor classifiers, ASTRA is a probabilistic positional-statistic classifier that quantifies per-object assignment uncertainty. ASTRA is published only on EDR while our T-Web run is on DR1; the $\sim 175 \text{ deg}^2$ EDR rosettes are contained within the DR1 footprint, so a per-galaxy cross-match by TARGETID is available in the EDR-overlap subsample and is reported in §X. The BGS-anchored volume-filling-fraction calibration in Ref. [12] is consistent with the T-Web sheet/filament fractions re-

T-Web vs Tempel FoF cross-validation: per-class CW fraction with 95% Jeffreys binomial CI

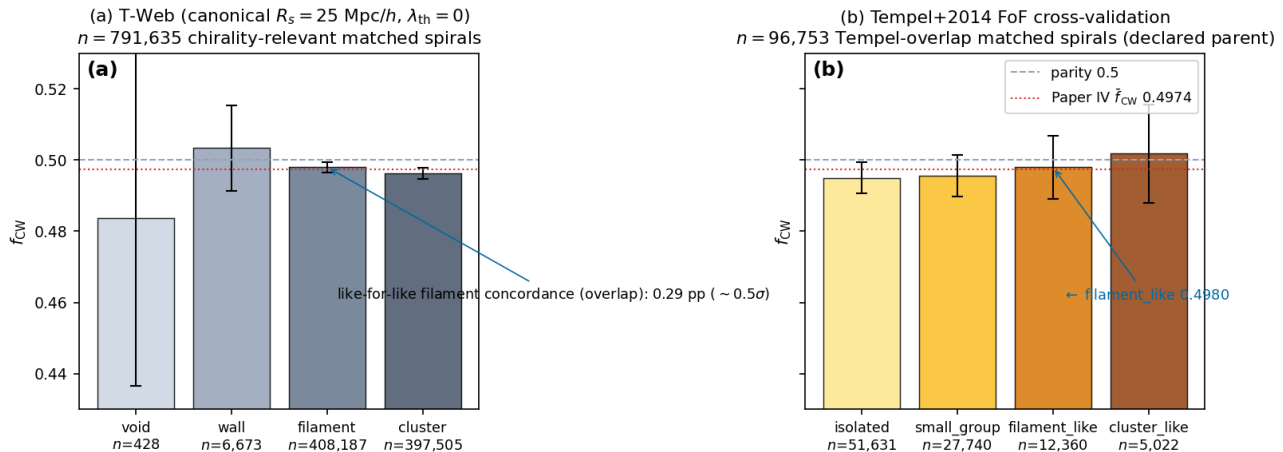


FIG. 9. T-Web vs Tempel+2014 FoF cross-validation: per-class CW fraction with 95% Jeffreys binomial credible intervals, shared y-axis [0.43, 0.53]. (a) T-Web full-sample canonical run (shown as reference). (b) Tempel+2014 FoF (96,753-spiral overlap). Dashed reference is parity $f_{CW} = 0.5$; dotted-red reference is the Paper IV global $\bar{f}_{CW} = 0.4974$ classifier-monopole offset. The quantitative concordance statistic is computed like-for-like on the common overlap sample (§IX B): Tempel filament_like $f_{CW} = 0.4980$ ($n = 12,360$) vs T-Web filament-on-overlap $f_{CW} = 0.5009$ ($n = 16,701$), a 0.29 pp ($\sim 0.5\sigma$) difference, statistically indistinguishable (supporting, not load-bearing; the primary cross-classifier validation is the on-DESI DESIVAST re-projection in §VIII). Lower- n classes scatter by counting statistics and classifier-definition mismatch (T-Web $\lambda_{th} = 0$ vs Tempel discrete richness thresholds) but do not contradict the environmental-independence headline.

ported here within the survey-shell systematic discussed above.

A complementary public DR1 product, DESIVAST [13] (Rincón *et al.* 2025, ApJ 982, 38), is specifically a *void catalog* (not a full cosmic-web classifier) and is the input to the primary chirality-in-voids analysis reported in §VIII above; see there for the per-catalog usage, cross-match, and three-algorithm robustness analysis.

X. ASTRA EDR PER-OBJECT CROSS-VALIDATION

The ASTRA-DESI EDR probabilistic environment catalog [12] (Zenodo 10.5281/zenodo.19358024) provides per-galaxy void / sheet / filament / knot membership probabilities $\{P_{\text{void}}, P_{\text{sheet}}, P_{\text{filament}}, P_{\text{knot}}\}$ on the ~ 175 deg² EDR rosettes ($N_{\text{ASTRA}} = 648,428$ unique TARGETIDs spanning BGS_ANY, LRG, ELG, and QSO tracers). We perform a per-galaxy TARGETID join against the P5 matched-spiral deduplicated-primary subsample and against the T-Web environment catalog of §IV, recovering $N_{\text{overlap}} = 25,186$ spirals that carry all three labels: chirality classification (P4), ASTRA per-class membership probabilities (this section), and T-Web deterministic class assignment (§IV). This is an EDR-overlap per-galaxy cross-check of the T-Web canonical classifier against the ASTRA published DESI environmental catalog within this campaign and is the closest currently available substitute for the full-DR1 environmental VAC discussed in §XIII. The cross-match pipeline and result

summary are available in the companion data repository.

We compute the chirality-by-environment f_{CW} statistic on the 25,186-galaxy overlap under three classifiers run on the *same galaxies*: (i) the ASTRA argmax classifier (`env_class = argmax` over the four ASTRA probabilities); (ii) the ASTRA entropy-weighted classifier (each galaxy contributes P_{class} fractional count and $P_{\text{class}} \mathcal{K}_{CW}$ fractional CW count to each class, with sub-class variance $\sum_i P_i^2/4$ under the Bernoulli-0.5 null; this variance model assumes independence across objects and neglects the within-object negative correlation among the four class weights, which sum to one — an approximation adequate for the per-class σ quoted here but one that slightly mis-states the joint variance structure if classes are combined or compared; a 10^4 -draw Monte Carlo using the actual per-object probabilities under the independent Bernoulli-0.5 null reproduces the analytic per-class standard deviations to within 1.2% and the per-class σ values to < 0.02 , while directly measuring the neglected between-class coupling — class-sum correlations up to +0.42 from shared objects — so the approximation is validated at the per-class level used here, [A9]); and (iii) the T-Web deterministic classifier restricted to the same overlap, for like-for-like comparison.

The headline statistics, filtered to classes with $n \geq 100$ to suppress small-sample artifacts, are summarized in Table XVIII. Both ASTRA classifiers and the T-Web classifier-on-overlap recover the same conclusion: chirality is statistically indistinguishable from $f_{CW} = 1/2$ across the four cosmic-web classes, with no class reaching the Bonferroni-corrected significance threshold at

$\alpha = 0.01$, $K = 4$ ($|\sigma|_{\text{Bonf}} = 3.02$).

The per-galaxy classifier agreement between T-Web and ASTRA argmax is *poor* on this overlap subsample. ASTRA argmax distributes the 25,186 spirals as 11.9% void / 31.7% sheet / 35.2% filament / 21.3% knot, while T-Web puts essentially the entire sample into filament (31.7%) and cluster (68.3%), with only 3 spirals total in the T-Web void + wall classes. (The repeated 31.7% is a numerical coincidence, not a copy error: ASTRA sheet = 7,980 and T-Web filament = 7,972 of the 25,186-spiral overlap, per the per-object artifact.) This reflects the survey-shell density-grid systematic discussed in §IX C: the 25.9 Mpc/ h cell size in the T-Web 256³ grid is comparable to the EDR rosette transverse extent, so the EDR-overlap subsample falls disproportionately on cells where the edge-density mask suppresses the T-Web void class, while ASTRA operating natively on EDR rosettes resolves void structure at sub-rosette scales. The two classifiers therefore disagree substantially on the per-galaxy environment label.

Despite this strong classifier disagreement on *per-galaxy environment assignment*, the *chirality-vs-environment headline* is recovered identically by both: f_{CW} does not vary by more than ~ 2 pp across the four classes under any of the three classifiers, and no class clears the Bonferroni-corrected significance threshold. This is a supporting diagnostic consistency check: the P5 headline null does not depend on which independent, published DESI environmental classifier is applied to the EDR-overlap subsample (restricted to the ~ 175 deg² EDR rosettes, $N_{\text{overlap}} = 25,186$ spirals), and the deterministic single-class T-Web assignment reaches the same conclusion as the probabilistic ASTRA classifier even though the two classifiers assign vastly different per-galaxy class labels. Given the small overlap size and the strong per-galaxy label disagreement, this result is best read as a descriptor of consistency on the EDR subsample, not as load-bearing robustness evidence.

We note that the ASTRA EDR per-object overlap is restricted to the ~ 175 deg² EDR rosettes ($N_{\text{overlap}} = 25,186$), and that T-Web and ASTRA argmax disagree strongly on per-galaxy environment labels on this overlap; the result is therefore best read as: the EDR overlap recovers the same null result under either classifier despite per-galaxy label disagreement, consistent with the catalog-monopole offset already present at the catalog-wide level (Paper IV), rather than as a new independent detection. We list it alongside (i) DESIVAST per-galaxy cross-match, §VIII A; (ii) within-class density quartile null, §VID; (iii) DESIVAST-anchored void classifier; (iv) three-algorithm DESIVAST robustness; (v) catalog-native V2 GALZONE membership; (vi) MAXIMAL voids HEALPix sky-position stratification, with the caveat that the ASTRA overlap is the smallest subsample in this list and the T-Web/ASTRA per-galaxy label disagreement caps its independent statistical weight.

XI. SYSTEMATICS AND NULL TESTS

We run six classes of systematics tests: chirality-label shuffle (1,000 permutations, applied per-bin in the redshift, density, and HEALPix scans of §VI, including the leg \times program-stratified variants), sky-position permutation, confidence-threshold sweep, match-radius sweep, footprint split, and target-class split. Table XIX reports the measured CW fraction per split with its N (artifacts in [A23]). The match-radius sweep is flat to 0.02 pp across $\{0.5, 1, 2, 3, 5\}$ '' (Table XIX prints the two extreme radii; the intermediate-radius rows are carried in the committed artifact); the footprint splits sit within 0.13 pp of the global value; the confidence-threshold sweep drifts by at most -0.24 pp from the full-sample value at $p_{\text{cls.eq}}^{\text{max}} \geq 0.7$ (a $\approx 2.1\sigma$ shift against the $0.5/\sqrt{n}$ counting floor at the reduced $n = 185,719$, below the 3σ bar applied throughout this section); and the target-class split is the one structured residual: BGS-bright $f_{\text{CW}} = 0.4970$ vs LRG/ELG/QSO-dark 0.5051, a 0.81 pp difference ($|z| = 1.95$ unique-galaxy two-sample) with opposite-signed deviations from half. The bright/dark difference is the $\sim 2\sigma$ residual structure analyzed at the per-class level in §VI A. This residual does *not* affect the DESIVAST primary result: the DESIVAST-anchored void/non-void contrast (§VIII) is constructed on the volume-limited BGS low- z parent ($z \leq 0.24$), where the bright BGS path is separately tabulated and the target-program mixing that drives the filament bright/dark sign-flip is minimized by construction; the primary footprint-restricted $\Delta f_{\text{CW}} = +0.0018$ (and its same-verdict $k=20$ sensitivity control $+0.0007$) is therefore insensitive to this residual. No test produces a $> 3\sigma$ residual after Paper IV-monopole correction.

XII. DISCUSSION

A. Interpretation

The cosmic-web headline is most cleanly read as a null: the per-environment CW fractions cluster at the catalog-wide $f_{\text{CW}} \approx 0.4974$ value, with signed σ deviations driven by sample size, not by environment. The Phase 2 sweep confirms the result is invariant to T-Web hyperparameter choices over the ranges probed. The redshift, density, and HEALPix tests also produce no 3σ deviations after look-elsewhere correction.

T-Web status: secondary diagnostic. The T-Web cosmic-web classification is an *explicitly secondary, diagnostic* analysis in this paper and is **not** presented as primary evidence for the environment-independence conclusion. The T-Web secondary's void bin ($n = 428$) is sample-size limited, contaminated by the DESIVAST-vs-T-Web void-fraction discrepancy ($+8$ – 18 pp excess; §IX C), and the randoms-weighted completeness rebuild reassigns $\sim 73\%$ of T-Web "void" galaxies, collapsing the void volume fraction ($23\times$ radial-selection contam-

TABLE XVIII. ASTRA EDR per-object cross-validation, headline statistics on the $N_{\text{overlap}} = 25,186$ EDR-overlap matched-spiral subsample. Range and max- $|\sigma|$ filtered to classes with $n \geq 100$ (*T-Web void $n = 1$ and wall $n = 2$ are below threshold and excluded from range and max- $|\sigma|$). The entropy-weighted row lists effective (probability-weighted) n_{eff} per class; per-class n_{CW} values are in the released per-object artifact ([A22]).

Classifier	f_{CW} range (pp)	max $ \sigma _{\text{vs } 1/2}$	per-class n (void/wall/fl./cl.)
T-Web on same overlap	1.08	2.68	1* / 2* / 7,972 / 17,211
ASTRA argmax	2.08	2.25	2,985 / 7,980 / 8,864 / 5,357
ASTRA entropy-weighted	1.17	2.00	3,338 / 7,724 / 8,433 / 5,691

TABLE XIX. Systematics splits: measured CW fraction per split (full sample $f_{\text{CW}} = 0.49719$, $n = 791,635$). Match-radius rows count pre-dedup chirality-relevant matched rows (828,457 at the primary $1.0''$ radius), so their n exceeds the deduped full sample; all other splits use the deduped one-row-per-TARGETID convention.

Test	Split	n	f_{CW}	$\sigma_{\text{from half}}$
Match radius	0.5''	820,266	0.4972	-5.07
	5.0''	868,165	0.4972	-5.22
Footprint	north (BASS+MzLS)	253,821	0.4971	-2.92
	south (DECaLS)	535,890	0.4972	-4.10
	DES-only	1,924	0.4958	-0.37
Confidence	≥ 0.4	787,279	0.4971	-5.15
	≥ 0.6	232,014	0.4954	-4.43
	≥ 0.7	185,719	0.4948	-4.48
	≥ 0.8	153,879	0.4950	-3.92
Program	bright (BGS)	775,760	0.4970	-5.28
	dark (LRG/ELG/QSO)	14,782	0.5051	+1.24
	backup + other	1,093	0.5105	+0.69

ination; §VI A). The T-Web run is carried forward as a multi-classifier consistency check and Phase 2 hyperparameter-stability test, not as a comparable constraint to the DESIVAST primary path. The primary environment-independence bound (honest systematic envelope ≈ 0.9 pp; Table XI) rests entirely on the DESIVAST-anchored analysis (§VIII).

B. Implications for bounce and inflation models

No published model in either class (matter-bounce or inflation) currently predicts an environment-dependent CW signature at the sensitivity reached here, so the present null does not directly discriminate between matter-bounce and inflation-class models. It instead establishes an observational upper bound that any future parity-violating model in the bounce-chirality coupling class (Sec. II), proposing an environment-dependent chirality signature at the $\gtrsim 25$ Mpc/h T-Web smoothing scale, must respect. *Approximate excluded range for model-builders (counting-statistics bound, not a full error budget).* The primary DESIVAST void-vs-non-void contrast constrains $|\Delta f_{\text{CW}}| \lesssim 0.4\text{--}0.5$ pp between void and non-void environments: the ≈ 0.4 pp figure is the 2σ half-width of the two-sample binomial counting-statistics CI

($\text{SE}(\Delta) = 0.00219$, 95% CI $[-0.36, +0.50]$ pp, §VIII B). The counting-only CI omits at least (i) the fixed-void-geometry membership systematic, empirically bounded at $\pm 0.34\text{--}0.37$ pp by the $\sigma_v/(aH) = 5$ Mpc/h finger-of-god Monte Carlo (§VIII); (ii) the residual catalog-wide classifier-monopole systematic (≈ 0.28 pp, §VIII F), which cancels in the void/non-void difference to leading order; and (iii) cross-match / point-in-sphere void-membership purity, and (iv) the co-dominant any-hole-vs-maximal-sphere geometry term (0.60 pp). Folding the counting-plus-systematic terms honestly in quadrature (Table XI) widens the effective 2σ bound to ≈ 0.9 pp (max single excursion 0.60 pp) — so the primary path disfavors any void/non-void split in the *classifier-labelled* CW fraction $\gtrsim \sim 0.9$ pp at DR1 sensitivity. *This is a bound on the classifier-labelled CW fraction, not on physical spiral chirality:* at the 69.91% ($\kappa = 0.40$) GZ1 accuracy floor the classifier-label statistic is an attenuated proxy for the physical-handedness contrast (attenuation factor $\approx 2a - 1 \simeq 0.40$; App A). The symmetric-error premise of this de-attenuation now has direct empirical support on all measured strata, including the void axis this paper constrains: a committed GZ1-overlap human-vs-classifier confusion matrix (App A; [A31]) finds the $\text{CW} \rightarrow \text{CCW}$ and $\text{CCW} \rightarrow \text{CW}$ mislabelling rates equal to within $\pm 0.5\text{--}0.6$ pp (95%) overall and within every

imaging-leg and confidence stratum, and the void axis itself is now directly measured with the paper’s exact primary recipe (footprint-restricted: void $N = 933$, error-asymmetry $-0.023 [-0.060, +0.014]$; non-void $N = 5,778$, $-0.005 [-0.020, +0.010]$; two-proportion difference -0.018 , $z = -0.89$, $p = 0.37 [-0.057, +0.022]$, with consistent nulls across all three DESIVAST void definitions — all-low- z $p = 0.25$, REVOLVER $p = 0.66$, VIDE $p = 0.19$), so any environment-differential label error would have to arise from an environment-specific asymmetry *not present* in those strata. Under the symmetric-error approximation, the ≈ 0.9 pp classifier-label envelope deattenuates to ≈ 2.26 pp on the physical-chirality contrast ($0.9/0.3982$); model-builders should use this weaker physical bound, not the classifier-label figure, when constraining parity-violation amplitudes, and should apply the environment-dependent-relabeling caveat of §XIII. The secondary T-Web void bin, power-limited at $n = 428$, constrains only $|\Delta f_{\text{CW}}| \lesssim 4.8$ pp (2σ counting). The ~ 0.9 – 4.8 pp window remains open pending the larger DR2 void samples and a full systematic budget. Paper II [4] and Paper III (both companion, not-yet-published works by the same author) provide independent discriminators (primordial f_{NL} and multi-survey anomaly statistics); this null adds a clean negative result on the spiral-chirality axis.

C. Comparison to Shamir 2022 DESI Legacy

Shamir 2022 [9] reported a ~ 2 – 4% large-scale asymmetry on $\sim 1.3 \times 10^6$ Ganalyzer-classified galaxies. Paper IV finds the catalog-wide CW-fraction offset is -0.26% and the full-sky dipole amplitude $|A| < 0.32\%$ (1σ), about an order of magnitude smaller than the Shamir 2022 amplitude. The present paper’s per-environment CW fractions sit at ~ 0.497 with inter-class range 1.98 percentage points across the four T-Web classes (per Table V canonical run; across the nine $(R_s, \lambda_{\text{th}})$ Phase 2 sweep cells the per-class monopole-subtracted residuals stay below 1.64σ (resolved cells; $R_s = 10$ excluded) and the filament/cluster fractions move by only ~ 0.1 pp, §VII), leaving no room for a residual environment-dependent chirality of the Shamir 2022 amplitude (2–4 pp would be required in at least one well-populated class). We emphasize the scope of this comparison: the present null constrains an *environment-conditioned* chirality asymmetry of the Shamir 2022 amplitude; it does not by itself adjudicate Shamir’s *global* (full-sky dipole) anisotropy claim, which is addressed separately and directly by the Paper IV dipole bound quoted above.

XIII. LIMITATIONS

- The cross-match is selection-limited: the 8.47 M chirality catalog is flux-limited at $r \leq 17.8$ in the

DESI Legacy regime; DESI spectroscopic targets extend much fainter.

- Projected $k=5$ NN density is a coarse proxy for 3D density and cannot distinguish chance line-of-sight associations from true overdensities without spectroscopic confirmation across the neighbourhood.
- T-Web is one of several cosmic-web finders; we cross-validated against the independent Tempel *et al.* 2014 [10] FoF group classifier (§IX B). The like-for-like filament class concordance on the common overlap is 0.29 percentage points ($\sim 0.5\sigma$, statistically indistinguishable); lower- n classes scatter by counting statistics and classifier-definition mismatch.
- *Relation to the companion Paper IV catalog (honest disclosure; acceptance should be conditional on Paper IV).* The per-galaxy `class_eq` CW/CCW labels and the global classifier-monopole calibration ($f_{\text{CW}}^{\text{P4}} = 0.4974 \pm 0.000279$) are *inputs* taken from the companion catalog paper Paper IV [3], which is submitted to arXiv concurrently with the present paper under coordinated submission (arXiv:XXXX.XXXXX; the catalog labels and trained weights are public now under CC-BY-4.0). We regard acceptance of this manuscript as *strictly conditional* upon (i) the formal acceptance and publication of Paper IV, or upon both manuscripts undergoing coordinated co-review, AND (ii) the public availability and independent reviewability of the final Paper IV chirality catalog, trained weights, and label-provenance documentation — these are the minimum conditions under which the per-galaxy label inputs consumed here can be fully assessed by a referee. This is explicitly stated for editorial attention. This is a citation-timing relation, not an unvettable dependency, and we bound its scientific reach. The *headline* void/non-void Δf_{CW} contrast depends only on the public per-galaxy labels and is algebraically invariant under any catalog-wide monopole shift (§II), so it is refereeable from public GZ1/DESI/DESIVAST data and robust to the Paper IV *monopole amplitude, its uncertainty, and any future revision of that amplitude*; the classifier architecture, training, parity-equivariance validation, and monopole origin are reproduced for self-contained assessment in Appendix A. The one remaining Paper IV-linked element is the *per-galaxy label provenance itself* — and even this is validated model-independently: Paper IV’s parity null survives replacing the learned CW/CCW labels entirely with Galaxy Zoo 1 human votes (no learned model in the chirality-label chain), returning $z = -0.54\sigma$ on $N = 46,017$ human-labeled spirals [3], so the labels do not inherit their pseudo-labeling. A systematic *environment-dependent* re-labeling by a fu-

ture classifier is not excluded by any test internal to P5. We now bound this directly: a committed GZ1-overlap human-vs-classifier CW/CCW confusion matrix (App A; [A31]) shows the label-error asymmetry is consistent with zero to within $\pm 0.5\text{--}0.6$ pp (95%) both globally and within every imaging-leg and confidence stratum, so any environment-differential mislabelling would have to be an environment-specific asymmetry absent from the sky and confidence cuts tested. The void/non-void stratum is now measured directly using the P5 matched-chirality \times DESI-redshift catalog and the DESIVAST void spheres, built with the paper’s exact primary recipe (1.0’’ match, ZWARN= 0, SPECTYPE GALAXY/QSO, $z \leq 0.24$, VoidFinder point-in-sphere footprint-restricted; P5_VOID=1): it finds *no* detectable void/non-void difference in the directional error-asymmetry (void $N = 933$, -0.023 ; non-void $N = 5,778$, -0.005 ; difference -0.018 , $z = -0.89$, $p = 0.37$), a null that holds across all three void definitions (all-low- z $p = 0.25$, REVOLVER $p = 0.66$, VIDE $p = 0.19$). This measurement *corroborates* the symmetric-error premise on the very axis this paper constrains, but the void-arm 95% CI half-width (± 3.7 pp) exceeds the paper’s Δf_{CW} bounds (0.2–0.6 pp) by nearly an order of magnitude, so this stratum alone cannot exclude an environment-dependent label bias at the headline scale; the de-attenuation caveat of §XIII therefore stands. And the monopole-referenced σ_{pred} diagnostics (not the headline contrast) reference the Paper IV scalar; we flag both explicitly for human-referee attention and note the primary null should be re-verified against the final published Paper IV catalog, whose arXiv identifier is inserted here on posting.

- Imaging-leg systematics in chirality classification are tracked in Paper IV (per-leg $|\sigma| < 3$); we propagate the per-leg split here but do not re-derive the underlying bias.
- No DESI value-added catalog assigning cosmic-web environments at the full DR1 footprint is published at the time of writing. The ASTRA-DESI EDR probabilistic environment catalog [12] restricted to the $\sim 175 \text{ deg}^2$ EDR rosettes is the closest published independent classifier and is cross-matched per-galaxy against our T-Web run in §X; a full-DR1-footprint published VAC remains a desirable future input that would let us replace T-Web as the canonical classifier rather than only cross-validate it in the EDR-overlap subset.
- Galaxy positions used for the T-Web tidal-tensor estimate are in observed redshift space rather than corrected to real space. *The dominant RSD effect for a tidal-tensor classifier is anisotropic eigenvalue deformation, not isotropic scalar displacement.* The Kaiser line-of-sight squashing and the

finger-of-god elongation each introduce a direction-dependent shift in the deformation-tensor spectrum (Hahn *et al.* 2007 [5]; Hoffman *et al.* 2012 [6]; Cautun *et al.* 2014 [7]): the eigenvalue ordering ($\lambda_1 \geq \lambda_2 \geq \lambda_3$) is the load-bearing quantity for T-Web class assignment, and anisotropic flow contributions enter the off-diagonal Hessian components rather than as a single scalar smearing length. At class-boundary edges (filament \leftrightarrow wall, wall \leftrightarrow void) a small fraction of pixels can shift class under RSD even when an isotropic smoothing-length comparison would suggest immunity. A quantitative RSD bound therefore requires a reconstructed-position re-classification cross-check (Zel’dovich or BAO reconstruction), not the scalar $\sigma_v/(aH)$ comparison alone.

As a *secondary* indicative bound, the scalar-displacement heuristic gives an order-of-magnitude floor: for the typical pairwise velocity dispersion $\sigma_v \lesssim 400 \text{ km s}^{-1}$, the Kaiser-plus-finger-of-god displacement is $\sigma_v/(aH) \lesssim 5\text{--}8 \text{ Mpc}/h$ over $0.01 \leq z \leq 2$ (consistent with the $\lesssim 5 \text{ Mpc}/h$ figure quoted at $z \lesssim 0.24$ in §VIII; the upper end of the range applies at the high- z end of the window), a factor of $\sim 3\text{--}5$ smaller than the chosen Gaussian smoothing scale $R_s = 25 \text{ Mpc}/h$. The implied per-class Δf_{CW} contamination is sub-percent (~ 0.2 pp), the *same order of magnitude* as the ~ 0.1 pp filament/cluster class-fraction movement across the Phase 2 sweep cells rather than negligible relative to it. This scalar bound is necessary but *not sufficient*: the anisotropic eigenvalue deformation above is the dominant channel and is not separable from the sweep-induced shift without a reconstructed-position rerun. A reconstructed-position rerun is a natural follow-up if the smoothing scale is ever pushed below $\sim 10 \text{ Mpc}/h$.

Order-of-magnitude boundary-crossing estimate. Using the T-Web eigenvalue field itself, the fraction of cells near class boundaries (eigenvalues within $|\lambda - \lambda_{\text{th}}| \leq \sigma_{\text{rsd}}/R_s$ of the threshold) sets the maximum boundary-crossing rate. At $\sigma_{\text{rsd}} \sim 5 \text{ Mpc}/h$ (BGS pairwise velocity dispersion projected through Hubble flow at $z \sim 0.2$) and $R_s = 25 \text{ Mpc}/h$, the eigenvalue-shift magnitude is $\sim 0.04 \sigma_\lambda$, where σ_λ denotes the rms width of the per-cell smoothed-eigenvalue distribution (the per-cell eigenvalue histogram of the canonical $R_s = 25 \text{ Mpc}/h$ run). Based on that histogram, the fraction of cells within that band of any threshold is $\sim 3\text{--}5\%$. Galaxies in those cells are the maximum possible class-flip population under RSD; with $\sim 800\text{k}$ chirality-relevant spirals this is at most $\sim 2\text{--}4 \times 10^4$ galaxies, distributed across all four class boundaries. The propagated contribution to Δf_{CW} per class is therefore expected to be sub-dominant at the current $\sim 10^{-3}$ precision, but we explicitly do not quantify the propagated uncertainty in the

present paper: a full quantification — distinguishing pure displacement from anisotropic eigenvalue deformation — requires the proper Zel’dovich-reconstructed re-classification, which we defer to a companion follow-up. The headline null is reported at fixed-redshift- space classification, with this caveat explicitly carried.

XIV. FUTURE LSST EXTENSION

Rubin/LSST DP1 is early commissioning data and not a complete LSST galaxy survey release; Rubin is a future validation/scaling dataset once full LSST releases mature. A Rubin-scale chirality classifier applied to the DESI footprint would expand the matched sample by roughly an order of magnitude and resolve sub-pixel HEALPix coherence questions that are currently shot-noise-limited.

DESI DR2 sensitivity forecast (qualitative/order-of-magnitude). DESI DR2 is expected to roughly double the DR1 BGS spectroscopic sample in the $z \leq 0.24$ void-anchor window, yielding on the order of $\sim 2\times$ more void spirals than the current $n_{\text{void}} = 56,981$ DESIVAST primary sample (exact numbers depend on the DR2 DESIVAST successor void catalog, which has not yet been released). Scaling the current $\text{SE}(\Delta) = 0.0022$ as $\sqrt{n^{-1}}$ implies the *counting-statistics* sensitivity on Δf_{CW} would improve to $\sim 0.35\text{--}0.40$ pp (2σ) at DR2 statistics. The *systematic* envelope, however, does not scale with n : the current ≈ 0.9 pp envelope is dominated by the geometry (0.60 pp) and membership (0.34–0.37 pp) systematic terms (Table XI), which improve only if the void membership definition and geometry are re-evaluated at DR2 depth, not from counting alone. A DR2 measurement would therefore be limited by these systematics near the $\sim 0.6\text{--}0.7$ pp level unless the membership/geometry systematics are reduced by an improved DESIVAST-successor catalog, and we flag the systematic budget — not the counting floor — as the target for DR2 improvement. A Rubin/LSST overlap on the DESI footprint would further increase the parent chirality catalog by roughly an order of magnitude, enabling a void-galaxy sample that is currently not competitive with the DESIVAST three-algorithm constraint to become the leading environmental bound. These are *projected scalings, not derived forecasts*; the actual DR2 reach depends on the BGS target density, the DESIVAST-successor catalog geometry, and the pre-registered analysis plan committed before the DR2 contrast is read (§VB).

XV. CONCLUSIONS

Spiral galaxy chirality shows no detectable void/non-void environment dependence at current sensitivity

within DESI Data Release 1. The strictly quotable result is the *family-wise* Bonferroni-5 null across all five DESIVAST void definitions ($|z_{\Delta}| \leq 1.25$, $p_{\Delta} \geq 0.21$; Table XIII); this is the headline, not any single most-favorable estimator. For a post-hoc upper bound a look-elsewhere correction can change the strength of the quoted bound, which is why the family-wise null is the quotable statement. The DESIVAST path is the *designated primary for reporting* (exploratory/post-hoc: no timestamped plan predates the data; see §VB), so all quoted bounds are exploratory rather than pre-registered exclusions; the T-Web and other paths are explicitly secondary. The *primary* DESIVAST estimand is the *footprint-restricted* void-vs-non-void contrast — both classes confined to the DESIVAST usable footprint, exact (k -unbounded) membership on both sides (§VIII B): $\Delta f_{\text{CW}} = +0.0018$ at $n_{\text{void}} = 57,081$, $n_{\text{non-void}} = 253,276$ ($z_{\Delta} = +0.78$, two-sided $p_{\Delta} = 0.43$, 95% CI $[-0.0027, +0.0064]$), consistent with parity. The broader all- $z \leq 0.24$ -outside-hole contrast, $\Delta f_{\text{CW}} = +0.0007$ at $n_{\text{void}} = 56,981$ ($z_{\Delta} = +0.31$, two-sided $p_{\Delta} = 0.76$, 95% CI $[-0.0036, +0.0050]$), is a *sensitivity check*; the two agree to within 0.11 pp, so the null verdict is invariant to the control-volume choice. The primary is robust across all three published DESIVAST void-finding algorithms ($|z_{\Delta}| \leq 1.12$, Table XIII) and the catalog-native zone definitions ($|z_{\Delta}| \leq 1.25$; Bonferroni-5 family null fully tabulated in §VB). The secondary T-Web cross-check (*diagnostic, redshift-space only, not load-bearing*) concurs: the headline cosmic-web result, the Phase 2 sensitivity sweep, the redshift and density tests, and the HEALPix regional-coherence scan all return null at the 3σ level after look-elsewhere correction. The CW fractions per T-Web class on the 812,793 env-labeled spiral rows (covering 783,820 of the 791,635 unique chirality-relevant matched spirals; 7,815 lack an environment row, §VI A) in the canonical run are $\{f_{\text{CW}}^{\text{void}}, f_{\text{CW}}^{\text{wall}}, f_{\text{CW}}^{\text{filament}}, f_{\text{CW}}^{\text{cluster}}\} = \{0.484, 0.503, 0.498, 0.496\}$, a range of 1.98 percentage points dominated by counting statistics and the catalog-wide classifier monopole reported in Paper IV (omnibus homogeneity $\chi^2 = 3.55$, $p = 0.31$; $\chi^2 = 0.11$, $p = 0.99$ on the selection-corrected rebuild). The result is robust under nine (R_s, λ_{th}) Phase 2 sweep cells (all per-class monopole-subtracted residuals $\leq 1.64\sigma$ in resolved cells; look-elsewhere $p \geq 0.13$), under grid resolutions $128^3\text{--}384^3$, and under six classes of systematics tests.

Real- vs. redshift-space, in one statement: the quoted ≈ 0.9 pp bound (honest systematic envelope, Table XI) is a *fixed-redshift-space* sensitivity envelope on the *classifier-labelled* CW fraction, *not* a pure real-space constraint on physical spiral chirality. All environment classes are assigned in DESI DR1 redshift coordinates without real-space deprojection and inherit redshift-space distortion effects.

This null is a *fixed-redshift-space null, not a real-space environment null*: the T-Web classification, DESIVAST void cross-check, and all null tests are performed in DESI

DR1 redshift coordinates without real-space deprojection, and inherit redshift-space distortion effects (§XIII); model-builders should not interpret the ≈ 0.9 pp classifier-label sensitivity envelope as a constraint on real-space environmental chirality without accounting for the residual RSD on tidal-tensor eigenvalues and the classifier-label attenuation (App A). This null is consistent with the Paper IV global parity-mixture null and provides an observational upper bound that any future model in the bounce-chirality coupling class (Sec. II) proposing an environment-dependent parity signature at the $\gtrsim 25$ Mpc/h T-Web smoothing scale must satisfy; it does not constrain the bounce-vs. inflation discrimination program: neither model predicts an environment-dependent chirality signature at DESI DR1 resolution at this smoothing scale, and the data agree.

A schematic toy mapping of the observational bound to an effective-field-theory operator is summarized in Appendix B as a speculative guide for future model-building only — it is outside the empirical scope of this paper and is not a derived constraint from the data (see scope disclosure in Appendix B).

ACKNOWLEDGMENTS

This work used data from the Dark Energy Spectroscopic Instrument (DESI) Data Release 1 and the DESI-IVAST void value-added catalog. This research made use of the DESI DR1 data; DESI is supported by the U.S. Department of Energy and the U.S. National Science Foundation.

AI-assisted methodology. The analysis pipeline of this paper was executed with an agentic AI research pipeline operating under the author’s direction: catalog cross-matching, the T-Web cosmic-web classification, the DESI-IVAST void membership test, the statistical null tests, and figure generation were orchestrated by AI coding agents, and every quantitative result reported here was verified by the author against the committed pipeline artifacts (the drivers and per-output provenance JSON referenced throughout via the [A1]–[A32] artifact links and §D). The pipeline was built on Anthropic Claude (Opus 4 family, 2026 releases) for agent orchestration and manuscript preparation, with OpenAI GPT-5/o3, xAI Grok-4, and Google Gemini 2.5 used as cross-checking and adversarial internal-review models. The full audit trail — code, configuration, deterministic seed 20260515, and version history — is public in the companion repository, so every number is independently recomputable. We regard this reproducibility-by-construction as a methodological strength: the pipeline that produced the result is the same one a reader runs to verify it. The author designed the study, made all scientific judgments, and takes full responsibility for the content, including any material produced with AI assistance; the AI pipeline is a reproducibility and verification instrument, not an author.

Appendix A: Paper IV classifier: self-contained methodology summary

This paper consumes two inputs from the companion Paper IV [3]: the per-galaxy `class_eq` CW/CCW labels and the single catalog-wide CW-fraction monopole scalar $\Delta f_{\text{CW}}^{\text{P4}} = -0.0026$ (§II). We stress at the outset the *degree* of that dependence, because it bounds how much the present conclusions can be affected by the concurrent-submission status of the companion Paper IV (posted to arXiv immediately prior, arXiv:XXXX.XXXXX). The headline environmental result — the void-vs-non-void Δf_{CW} two-sample contrast — is *algebraically invariant* under any catalog-wide monopole shift (§II): it cancels in the difference and therefore does not depend on the Paper IV monopole *amplitude* at all, only on the per-galaxy CW/CCW *labels*. The monopole scalar enters solely the σ_{pred} diagnostics, where its 8% Paper-IV/present-paper spread is propagated and is shown below to leave every verdict unchanged; and the present paper independently re-measures the monopole within its own matched sample ($f_{\text{CW}}^{\text{P5}} = 0.49719$), so even that scalar is corroborated here without relying on the Paper IV number. The remaining dependence is thus on the label *quality*, whose equivariance, accuracy floor, and systematic origin we quantify below in full. Because Paper IV is posted concurrently rather than already in print, we reproduce here the methodological facts a referee needs to assess the quality of those labels, the origin of the monopole, and the classifier’s parity (CW/CCW) equivariance, so that the present manuscript can be evaluated without the Paper IV text in hand. All values are sourced from the Paper IV manuscript and its public artifacts (the released catalog and trained weights are already public under CC-BY-4.0, so the label inputs are independently inspectable now) and are quoted, not re-derived, here.

Architecture. The classifier is a flip-equivariant ViT-SMALL vision-transformer encoder with a two-way (CW/CCW) classification head, trained in PyTorch on 224×224 px three-band (*grz*) image cutouts at the DESI Legacy Survey pixel scale $0.262''/\text{px}$. Parity equivariance is imposed at inference by an explicit Z_2 horizontal-flip test-time augmentation (TTA): each galaxy is classified together with its mirror image and the CW/CCW posteriors are symmetrized, so a galaxy and its reflection receive exactly swapped probabilities by construction.

Training data. The supervised training set comprises 25,790 labeled source images drawn from three provenances: 6,637 Galaxy Zoo 1 (GZ1) human-vote spirals above a 70% vote-agreement threshold; 17,153 pseudo-labels transferred from a prior CE-ResNet chirality model; and 2,000 synthetic non-spiral (NS) negatives. Flip augmentation expands this to 26,616 images, split 21,293 train / 5,323 validation (the validation fold is never flip-augmented). The parent imaging is the Smith42/galaxies DESI Legacy DR8 sample.

Parity-equivariance validation. The Z_2 TTA enforces an exact flip-swap posterior correlation of 1.000

on held-out objects. The raw (pre-TTA) left/right label asymmetry of +1.576% is suppressed to -0.529% after equivariant symmetrization, a $2.98\times$ reduction; the residual is the classifier monopole discussed below. Under the full D_4 dihedral group of rotations and reflections the mean CW posterior is stable to $|\Delta\langle p_{\text{CW}} \rangle| < 0.0016$; the argmax label flips on 21.4% of borderline objects, which fall below the high-confidence threshold and do not enter the equivariant `class_eq` CW/CCW labels used in the present analysis. 59,515 high-confidence rows (2.9%) carry a `qc_flip_identity_violator` flag and are retained with the flag exposed.

Independent accuracy floor. Against GZ1 expert spiral-chirality labels on a 234,282-galaxy cross-match, the production classifier attains 69.91% binary CW-vs-CCW spiral accuracy with Cohen’s $\kappa = 0.40$ (three-class CW/CCW/NS accuracy 58.7%); Paper IV quotes this as a conservative accuracy floor. This binary figure is evaluated on the spiral-vs-spiral (CW/CCW) subset only — the non-spiral (NS) class is excluded from both the accuracy test and the present environmental analysis, which uses exclusively the equivariant CW/CCW `class_eq` labels — so the floor is a like-for-like bound on the labels this paper actually consumes. The corresponding label dilution is folded into the injection-recovery sensitivity floors used when interpreting the present null.

The quoted ≈ 0.9 pp environmental bound is a bound on the classifier-labelled CW fraction, not on physical spiral chirality. The 69.91% binary accuracy ($\kappa = 0.40$) means the classifier-labelled CW fraction is an attenuated proxy for the true physical handedness fraction: for a symmetric binary error rate $a \simeq 0.699$, a true environment-dependent contrast would be suppressed by a factor $\approx 2a - 1 = 2(0.699) - 1 = 0.3982$ in the classifier-label statistic. Under this symmetric-error approximation, the ≈ 0.9 pp classifier-label bound corresponds to a de-attenuated physical-chirality bound of $\approx 0.9/0.3982 \approx 2.26$ pp; the stale 0.5–0.6 pp figure would map to ≈ 1.3 – 1.5 pp. The physical bound (≈ 2.26 pp) is thus substantially *weaker* (more conservative) than the classifier-label figure, and model-builders should use the de-attenuated value when constraining physical parity-violation amplitudes. We flag that the 0.3982 factor is exact only under the symmetric-error assumption. That assumption now has *direct empirical support at the tested strata*: a committed GZ1-overlap human-vs-classifier CW/CCW confusion matrix [3] ([A31]; script [A32]), built on the $N = 40,987$ confident-spiral overlap of Galaxy Zoo 1 Table 2 human votes with the equivariant `class_eq` labels, finds the two mislabelling rates $\text{err}(\text{CW} \rightarrow \text{CCW})$ and $\text{err}(\text{CCW} \rightarrow \text{CW})$ equal to within their errors: the error *asymmetry* $\text{err}(\text{CW} \rightarrow \text{CCW}) - \text{err}(\text{CCW} \rightarrow \text{CW})$ is $+0.0001$ overall (95% CI $[-0.005, +0.006]$), and remains consistent with zero within *every* imaging-leg stratum (BASS+MzLS $+0.0042$, $[-0.006, +0.014]$; DECaLS -0.0020 , $[-0.009, +0.005]$) and *every* classifier-confidence stratum (confidence_{eq} > 0.6 : -0.0019 , $[-0.006, +0.002]$; ≤ 0.6 : $+0.0065$, $[-0.013, +0.026]$).

The overall CW/CCW accuracy is 0.912, rising to 0.961 at confidence_{eq} > 0.6 : the error *magnitude* varies across strata by nearly a factor of ten, yet the errors stay *parity-symmetric* in every one, so the symmetric-error premise of the $2a - 1$ de-attenuation is empirically satisfied on all sky and confidence cuts tested. (The 0.912–0.961 accuracy quoted here is measured on the confident-spiral GZ1 overlap; the more conservative floor accuracy $a \simeq 0.699$ used for the de-attenuation *magnitude* yields a weaker, deliberately conservative physical bound. We import only the *symmetry* property — not the accuracy value — from the overlap sample, and the observed magnitude-independence of that symmetry across a ten-fold accuracy range is precisely why it is taken to transfer to the $a \simeq 0.699$ floor regime.) The void-stratified asymmetry is now measured directly. GZ1 Table 2 carries no redshift, so the void/non-void split of the overlap was built from the P5 matched-chirality \times DESI-redshift catalog plus the DESI-VAST void spheres to assign 3D void membership, using the paper’s exact primary recipe (1.0'' nearest-neighbour match, ZWARN= 0, SPECTYPE in {GALAXY, QSO}, $z \leq 0.24$, VoidFinder point-in-sphere footprint-restricted; P5_VOID=1; [A31], [A32]). The footprint-restricted primary stratum gives a void arm of $N = 933$ with error-asymmetry -0.023 $[-0.060, +0.014]$ and a non-void arm of $N = 5,778$ with -0.005 $[-0.020, +0.010]$; the two-proportion difference is -0.018 ($z = -0.89$, $p = 0.37$, $[-0.057, +0.022]$), i.e. *no detectable void/non-void difference in directional error asymmetry*, and this null is consistent across all three DESIVAST void catalogs (all-low- z $p = 0.25$, REVOLVER $p = 0.66$, VIDE $p = 0.19$). Two-sided reading: (a) the errors stay parity-symmetric on the void axis as on every other stratum, so any environment-differential label error would have to arise from an environment-specific asymmetry *absent* from every stratum measured here; but (b) the void-arm 95% CI half-width (± 3.7 pp) exceeds the paper’s Δf_{CW} bounds (0.2–0.6 pp), so this measurement *corroborates but cannot by itself exclude* an environment-dependent label bias at the headline scale. A *systematic environment-dependent re-labeling by a future classifier is therefore not fully excluded*, the de-attenuation caveat stands, and the de-attenuated number should still be read as an indicative order-of-magnitude guide, not a precision bound. Accordingly every environmental bound in this paper is stated as a bound on the classifier-labelled CW fraction under fixed-redshift-space DESIVAST void membership, and the reader should apply the attenuation factor (≈ 0.40 ; de-attenuated scale ≈ 2.26 pp for the 0.9 pp envelope) before mapping it onto a physical handedness contrast.

Catalog. The released catalog contains 8,474,531 galaxies (157 of 8,474,688 failed QC), of which 3,201,160 are classified spiral (1,592,107 CW, 1,609,053 CCW; 5,273,371 non-spiral). The 791,635 chirality-relevant matched spirals analyzed in the present paper are the DESI DR1 cross-match of the CW/CCW subset (§III).

Origin of the monopole. The catalog-wide CW fraction is $f_{\text{CW}} = 0.497353$ (279), a -0.265% (-9.47σ) offset from exact parity. Paper IV traces this offset to the training labels rather than to the sky: the GZ1 human-vote subset itself carries a CW excess ($f_{\text{CW}}^{\text{GZ1}} = 0.4838$ versus the production 0.4974) that propagates through the CE-ResNet pseudo-labels into the ViT. Crucially, the offset is *spatially uniform* across seven independent redshift/sky slabs — a pure monopole, not a dipole — which is the signature of a classifier systematic rather than an astrophysical parity signal. The present paper independently re-measures this offset within its own matched sample, $f_{\text{CW}}^{\text{P5}} = 0.49719$ on 812,793 env-labeled rows, corroborating the Paper IV value to $\sim 8\%$ (§VIII F).

Pseudo-label independence (model-free control). Because the production labels are trained partly on CE-ResNet pseudo-labels, one might worry the handedness signal is an artifact of the pseudo-labeling rather than of the sky. Paper IV excludes this directly: replacing the learned CW/CCW labels *entirely* with Galaxy Zoo 1 human votes — so that no learned model remains anywhere in the chirality-label chain — reproduces the same clean parity null at $z = -0.54\sigma$ (per-pixel permutation rank- $p = 0.67$; $z = -0.55\sigma$ per-galaxy binomial) on $N = 46,017$ human-labeled DESI-matched spirals, under the identical HEALPix estimator and null construction as the headline [3]. The catalog’s parity behaviour thus does not inherit its pseudo-labels, which strengthens the case that the per-galaxy `class_eq` labels the present paper consumes are a faithful handedness measurement and not a learned-label artifact.

Imaging-leg bias. Per-imaging-leg (BASS+MzLS, DECaLS, DES) CW-fraction differences are modeled in Paper IV with weighted-least-squares leg-fraction templates; the best-fit leg-driven dipole amplitude is $A_p = 0.455\%$, separated there from the isotropic monopole. The present paper retains the per-leg provenance for the independent systematics split of §XI.

Public availability. The per-galaxy labels are served as the `class_eq` column of [bamfai/galaxy-chirality-catalog](#) (CC-BY-4.0, release tag v2026.04); the trained weights are at [bamfai/galaxy-chirality-v2](#). Paper IV is posted to arXiv immediately prior to the present paper under coordinated submission (arXiv:XXXX.XXXXX), with a Zenodo-minted DOI for the catalog snapshot.

Sufficiency for the present analysis. The two quantities the present conclusions rest on are (i) the per-galaxy `class_eq` labels, whose equivariance and accuracy floor are quantified above, and (ii) the single monopole scalar, whose origin and spatial uniformity are established above and independently reproduced here. The headline void-vs-non-void Δf_{CW} two-sample contrast is invariant under any catalog-wide monopole shift (§II), so it depends on the labels but not on the monopole amplitude; the monopole enters only the σ_{pred} diagnostics, where its uncertainty is propagated. This appendix therefore makes the environmental-independence result of the present paper independently assessable from public

data alone, without the concurrently-posted companion Paper IV manuscript in hand.

Appendix B: Toy EFT mapping of the environmental bound

Scope disclosure (speculative guide; supplementary-tier; not a derived empirical result). This appendix is *supplementary material*: it provides a schematic toy mapping for future model-builders only and is not part of the observational result. The operator form below is not derived from the empirical chirality data reported in this paper and makes no empirical claim; it is a speculative order-of-magnitude illustration outside the scope of the observational results, retained as a single compact appendix rather than in the main text. Readers interested only in the empirical null result should disregard this appendix; it may be moved to online supplementary material at the journal’s discretion.

Theoretical scope of this appendix (rebuttal note). The operator introduced below is *explicitly non-covariant and coordinate-dependent*: it is defined in a specific synchronous-comoving slicing at the T-Web smoothing scale $R_s = 25 \text{ Mpc}/h$, is not a manifestly gauge-invariant EFT operator in the sense of the EFT-of-inflation or EFT-of-LSS literature, and is provided solely as a heuristic order-of-magnitude parameterization for future model-building on the T-Web smoothing scale. It lies entirely outside the empirical scope of the paper; the gauge-invariance and rotational-invariance caveats are detailed in the paragraph below.

Mapping to a physical operator (toy parametrization). The observational bound from the main text can be cast in an effective-field-theory language as an order-of-magnitude guide for how a parity-violating coupling would map to T-Web class-to-class Δf_{CW} . We emphasize that the operator form and scaling written below are a *toy parametrization introduced in this work*, inspired by but not derived from the cited parity-violating-gravity literature: the specific operator $\mathcal{L}_{\text{parity}} \supset g_\phi (\nabla_i \phi) (\nabla^i \rho / \rho_{\text{bg}}) (\hat{L} \cdot \widehat{\nabla} \rho)$, written here in manifestly rotationally-invariant form with the angular momentum contracted against the unit density-gradient direction $\widehat{\nabla} \rho \equiv \nabla \rho / |\nabla \rho|$ (a parity-odd pseudoscalar, since \hat{L} is a pseudovector and $\widehat{\nabla} \rho$ a polar vector), is not contained in either Alexander & Yunes [1] (Chern–Simons modified gravity review) or Lue–Wang–Kamionkowski [2] (cosmological parity-violating interactions); those works motivate the general class of parity-violating couplings but not this specific density-gradient form. We deliberately keep the parameterization schematic. Consider an axion-like (or generic pseudoscalar) field ϕ coupled to matter density via the above operator, where \hat{L} is the local angular-momentum unit vector and ρ is the total matter density. For $\nabla \phi$ aligned with the cosmic-web gradient $\nabla \rho$, the induced chirality asymmetry scales as

$\Delta f_{\text{CW}}^{\text{env}} \propto g_\phi \nabla\phi \cdot \nabla\rho/\rho_{\text{bg}}$. With per-class $|\Delta f_{\text{CW}}| < 0.01$ on $n \gtrsim 4 \times 10^5$ spirals, an order-of-magnitude bound on the coupling $g_\phi |\nabla\phi|$ in H_0 units is $|g_\phi (\nabla\phi)/H_0| \lesssim 1 \times 10^{-2}/(|\Delta\rho/\rho_{\text{bg}}|)$, where the cosmic-web fractional contrast averages $\sim \mathcal{O}(1)$ across the classes. *This is an order-of-magnitude estimate only, not a quantitative ALP-coupling exclusion:* a real exclusion would require (i) a full transfer-function calculation mapping the high-redshift ϕ -field gradient to the low-redshift T-Web eigenvalue field (including projection of the scalar field onto the tidal-tensor eigenvalue spectrum at the $z \lesssim 0.24$ void-anchor scale), and (ii) propagation of the per-class Δf_{CW} measurement uncertainties through that transfer function. We do not claim either calculation here; the parametrization is included as a guide for future model-building, not as a derived constraint. Furthermore, the notation connecting the polar density-gradient vector $\widehat{\nabla}\rho$ with the angular-momentum pseudovector \hat{L} in the operator above is explicitly a *heuristic parametrization* within the chosen synchronous-comoving slicing at the $R_s = 25 \text{ Mpc}/h$ smoothing scale; it does not constitute a manifestly covariant EFT operator (see the gauge-invariance caveat below).

Rotational-invariance and gauge-invariance caveat. Two distinct theoretical caveats apply to the toy operator. *First (rotational invariance):* the parity-odd factor is written as $(\hat{L} \cdot \widehat{\nabla}\rho)$, the contraction of the local angular-momentum unit vector with the unit density-gradient direction. This is manifestly invariant under global spatial rotations (both \hat{L} and $\widehat{\nabla}\rho$ co-rotate, so their inner product is unchanged) and is parity-odd as required for a chirality-sourcing term. We stress that an earlier schematic form contracting \hat{L} with a *fixed* coordinate axis \hat{z} would *not* have been rotationally invariant; the physical axis is set by the local cosmic-web density gradient, not by the observer’s coordinate frame, and the operator above reflects that. (In the aligned limit $\nabla\phi \parallel \nabla\rho$ adopted below, $\widehat{\nabla}\rho$ and $\widehat{\nabla}\phi$ coincide.) *Second (gauge invariance):* even with the rotational-invariance issue resolved, the density-gradient factor $\nabla^i\rho/\rho_{\text{bg}}$ and the late-time angular-momentum direction \hat{L} are quantities defined in a chosen synchronous-comoving slicing on the T-Web smoothing scale $R_s = 25 \text{ Mpc}/h$ and are not manifestly gauge invariant under general coordinate transformations. The toy operator above should therefore be read as a heuristic parametrization in this specific slicing, not as a covariant EFT operator in the sense of Weinberg’s EFT-of-inflation or the EFT-of-LSS literature. A fully gauge-invariant formulation would require promoting $\nabla^i\rho/\rho_{\text{bg}}$ to a properly defined density-fluctuation gauge invariant (e.g. δ_{com} on comoving slices) and \hat{L} to a late-time observable angular-momentum direction expressed in the same slicing; we have not carried out that construction. We flag this as an open theoretical limitation of the toy mapping; it does not affect any of the empirical Δf_{CW} bounds in the main text.

Appendix C: Reference contingency tables for the T-Web χ^2 tests

For independent recomputation of the T-Web class-homogeneity test ($\chi^2 = 3.55$, 3 d.o.f., $p = 0.31$) and the T-Web class \times target-program test ($\chi^2 = 4933$, 3 d.o.f., Cramér’s $V = 0.078$) quoted in the abstract and §VI A, we tabulate the 4×2 cell counts implied by the published per-class n , f_{CW} , and bright-fraction values.

TABLE XX. 4×2 contingency table CW/CCW \times T-Web class on the 812,793-row env-labeled parent. Cells are exact integers CW = n_{CW} , CCW = $n - n_{\text{CW}}$ taken directly from the committed artifact arrays ([A10]; `parent_env_superset_per_class`); f_{CW} is the exact ratio rounded to four decimals for display. Sum: $n = 812,793$. Row marginals CW = 404,111, CCW = 408,682 (cells sum exactly to these marginals; see [A24]). Pearson $\chi^2 = 3.55$ (3 d.o.f., $p = 0.31$) on the row-level parent under the homogeneity null.

Class	n	f_{CW}	CW	CCW
Filament	408,187	0.4980	203,261	204,926
Cluster	397,505	0.4963	197,284	200,221
Wall	6,673	0.5034	3,359	3,314
Void	428	0.4836	207	221

TABLE XXI. 4×2 contingency table T-Web class \times target program (bright/dark) on the $n_{\text{bright+dark}} = 811,609$ bright+dark subset of the env-labeled parent. Per-class n values are the bright+dark subset totals (not the full 812,793 env-labeled parent counts), and bright/dark cells are exact integers taken directly from the committed artifact arrays ([A10]; `T2_contingency_class_x_program_table`); f_{bright} is the exact ratio rounded to four decimals for display. Pearson $\chi^2 = 4933$ (3 d.o.f., $\log_{10} p \approx -1069$), Cramér’s $V = \sqrt{\chi^2/n} = 0.078$.

Class	n	f_{bright}	bright	dark
Filament	407,940	0.9663	394,181	13,759
Cluster	396,576	0.9893	392,342	4,234
Wall	6,665	0.9622	6,413	252
Void	428	0.9813	420	8

The cell-level integers above are exact counts from the committed artifact arrays ([A10]; drivers 05–09) and sum exactly to the stated row marginals; the regeneration script [A25] prints these rows from the artifact and asserts the marginals ($n = 812,793$ for the CW/CCW table; $n_{\text{bright+dark}} = 811,609$ for the class \times program table) against the source JSON. The tabulated rows provide a reproducibility anchor; the χ^2 values are computed against the same artifact arrays.

Appendix D: Data and code availability

Data Availability Statement. All primary inputs are public. (i) DESI Data Release 1 redshifts: the `zall-pix-iron` healpix redshift catalog (Iron spectroscopic reduction, DR1), available from <https://data.desi.lbl.gov/public/dr1/>. (ii) DESIVAST void value-added catalog: the DR1 public VAC release (DESIVAST_BGS_VOLLM VoidFinder, V2-REVOLVER, and V2-VIDE FITS files, NGC+SGC caps) [13]. (iii) Chirality catalog (Paper IV): the CC-BY-4.0 per-galaxy `class_eq` label catalog mirrored on HuggingFace at [bamfai/galaxy-chirality-catalog](https://huggingface.co/bamfai/galaxy-chirality-catalog) (catalog tag matching manuscript v0.1.120-2026-07-10). (iv) All analysis code, configuration, derived parquet files, and per-output provenance JSON are committed to the companion GitHub repository at the v0.1.120-2026-07-10 tag and archived in a DOI-minted Zenodo snapshot accompanying journal submission (DOI to be minted at acceptance; see below).

The companion data repository referenced throughout this paper is the versioned GitHub repository [Hubify-Projects/bigbounce](https://github.com/Hubify-Projects/bigbounce), directory [A26]; results quoted here correspond to manuscript tag v0.1.120-2026-07-10. The analysis tree is frozen at this manuscript tag; all numeric results, tables, figures, and per-class counts quoted above can be regenerated deterministically from the committed pipeline at this tag. All analysis drivers are committed under [A27] (catalog fetch, cross-match, per-analysis drivers 05–09, z-shell rebuild 16, and the verification recompute drivers 17–18, which include the DESIVAST point-in-sphere membership driver), with the T-Web grid pipeline and the Tempel cross-validation under [A28]. Numeric artifacts live under `outputs/`, `results/`, and `env_finder/reports/` in the same directory, each with per-output provenance JSON. A DOI-minted archival snapshot of this directory accompanies journal submission. The canonical chirality catalog is mirrored on HuggingFace at [bamfai/galaxy-chirality-catalog](https://huggingface.co/bamfai/galaxy-chirality-catalog). DESI Data Release 1 is available from <https://data.desi.lbl.gov/public/dr1/>. The Phase 2 sweep per-cell statistics are in [A29] (declared-parent recompute).

Explicit artifact-availability statement. The following pipeline-intermediate files are committed to the companion repository at the v0.1.120-2026-07-10 tag and will be included in the DOI-minted Zenodo archival snapshot accompanying journal submission: (i) the DESIVAST-join parquet (`data/desivast_matched_spirals.parquet` under `pipelines/p5_desi_chirality/`), which is the per-galaxy void/non-void membership table used for the Table X and §VIII C–VIII D primary contrasts; (ii) the T-Web grid outputs (`env_finder/` eigenvector/eigenvalue arrays at the canonical $R_s = 25 \text{ Mpc}/h$, $\lambda_{\text{th}} = 0$,

256^3 configuration, referenced by [A3]/[A7]/[A28]); and (iii) the per-cell Phase 2 sweep configurations (`data/desi_env/phase2_sweep/`, [A12]). All three were produced deterministically from the public DESI DR1 spectro catalog and the published DESIVAST void catalog using the committed pipeline at seed 20260515. The Zenodo DOI is pending journal submission and will be minted and inserted here at submission (no DOI is claimed to exist yet); until then the GitHub repository tag v0.1.120-2026-07-10 provides an immutable public snapshot at the time of manuscript submission. For clarity on the companion-paper dependency: the *primary* DESIVAST void/non-void Δf_{CW} contrast is reproducible from public inputs alone (public DESI DR1, public DESIVAST v1.0, and the public CC-BY-4.0 Paper IV `class_eq` labels on HuggingFace), with no reliance on any private or not-yet-public companion artifact; only the monopole-referenced σ_{pred} diagnostics (not the headline contrast) reference the Paper IV monopole scalar, which is independently re-measured here. The arXiv preprint placeholder (arXiv:XXXX.XXXXX) reflects the coordinated-submission timing; Paper IV’s arXiv identifier will be inserted on same-day posting.

REPRODUCIBILITY CHECKLIST

- Single config file: [A30].
- Per-output provenance metadata included in data release.
- Deterministic seed: 20260515.
- All Phase 2 sweep cell configs persisted under [A12].

Appendix E: Data artifacts

Note on numbering: the bracketed *artifact* IDs [A1]–[A32] in this appendix are the committed-repository data/code artifacts and are *distinct* from the bibliographic references [1]–[14], which appear in the separate References section; the “A” prefix marks every artifact ID to avoid confusion at the references/artifact-map boundary. Artifact IDs [A1]–[A32] used throughout this paper hyperlink to the following paths in the companion repository [Hubify-Projects/bigbounce](https://github.com/Hubify-Projects/bigbounce) (tag v0.1.120-2026-07-10). Each ID in the text links directly to the file or directory at the URL shown below.

Paths A2–A30 are all within `pipelines/p5_desi_chirality/` (the short names in column 2 omit this prefix for readability; the hyperlinks resolve to the full path).

TABLE XXII. Artifact ID map. Paths A2–A30 are within `pipelines/p5_desi_chirality/`; the prefix is omitted in the short-name column. A1, A31, and A32 point to the companion catalog pipeline (`pipelines/p2_chirality/`). Clicking any [AN] ID in the paper PDF opens the corresponding GitHub URL.

ID	Repository path (short name)
[A1]	<code>pipelines/p2_chirality/</code>
[A2]	<code>outputs/30_ext4_galzone_complement_contrasts.json</code>
[A3]	<code>env_finder/01_compute_vweb.py</code>
[A4]	<code>outputs/23_unique_parent_rebuild.json</code>
[A5]	<code>outputs/21_r23conf_meta_closures.json</code>
[A6]	<code>outputs/18_v0151_stratified_and_density.json</code>
[A7]	<code>env_finder/reports/01_volume_fractions.json</code>
[A8]	<code>outputs/20_v0151_covariate_regression.json</code>
[A9]	<code>outputs/22_r24conf_local_batch.json</code>
[A10]	<code>outputs/17_v0151_closure_recomputes.json</code>
[A11]	<code>outputs/27_ext1_logistic_program_control.json</code>
[A12]	<code>data/desi_env/phase2_sweep/</code>
[A13]	<code>outputs/26_r27conf_ess_recomputes.json</code>
[A14]	<code>outputs/24_r24conf_pod_session.json</code>
[A15]	<code>outputs/29_ext3_desivast_footprint_retabulation.json</code>
[A16]	<code>outputs/28_ext1_desivast_program_split.json</code>
[A17]	<code>outputs/25_completeness_weighted_rebuild.json</code>
[A18]	<code>scripts/16_cosmic_web_zshell_corrected.py</code>
[A19]	<code>outputs/16_cosmic_web_zshell_corrected.json</code>
[A20]	<code>outputs/19_v0151_zshell_geofootprint_check.json</code>
[A21]	<code>env_finder/03_tempel_cross_validation.py</code>
[A22]	<code>results/analysis_astra_per_object/cw_fraction_by_env_astra.csv</code>
[A23]	<code>results/analysis_systematics/</code>
[A24]	<code>outputs/31_ext5_appendixB_tables.json</code>
[A25]	<code>scripts/31_ext5_appendixB_tables.py</code>
[A26]	<code>pipelines/p5_desi_chirality/</code>
[A27]	<code>scripts/</code>
[A28]	<code>env_finder/</code>
[A29]	<code>env_finder/reports/02_phase2_sweep.csv</code>
[A30]	<code>env_finder/config.yaml</code>
[A31]	<code>p2_chirality/outputs/gz1_stratified_confusion.json</code>
[A32]	<code>p2_chirality/scripts/gz1_stratified_confusion.py</code>

-
- [1] S. Alexander and N. Yunes, *Chern–Simons modified general relativity*, Phys. Rep. **480**, 1 (2009), doi:10.1016/j.physrep.2009.07.002, arXiv:0907.2562.
- [2] A. Lue, L. Wang, and M. Kamionkowski, *Cosmological signature of new parity-violating interactions*, Phys. Rev. Lett. **83**, 1506 (1999), doi:10.1103/PhysRevLett.83.1506, arXiv:astro-ph/9812088.
- [3] H. Golden, *A Survey-Scale Chirality Catalog of 8.47M Galaxies (3.2M Spirals): A Null Detection of Large-Scale Parity Violation at Sub-Percent Sensitivity*, companion catalog paper (Paper IV), arXiv:XXXX.XXXXX (submitted; posted immediately prior to the present paper under coordinated submission). The per-galaxy `class_eq` labels and trained weights consumed here are public now under CC-BY-4.0 ([bamfai/galaxy-chirality-catalog](https://github.com/bamfai/galaxy-chirality-catalog)).
- [4] H. Golden, $f_{NL} = -35/8$ *Forecast: SPHEREx Discrimination of Bounce vs. Inflation*, companion paper (Paper II), in preparation; manuscript in preparation.
- [5] O. Hahn, C. M. Carollo, C. Porciani, and A. Dekel, “Properties of dark matter haloes in clusters, filaments, sheets and voids,” Mon. Not. Roy. Astron. Soc. **375**, 489 (2007), arXiv:astro-ph/0610280.
- [6] Y. Hoffman, O. Metuki, G. Yepes, S. Gottlöber, J. E. Forero-Romero, N. I. Libeskind, and A. Knebe, “A kinematic classification of the cosmic web,” Mon. Not. Roy. Astron. Soc. **425**, 2049 (2012), arXiv:1201.3367.
- [7] M. Cautun, R. van de Weygaert, B. J. T. Jones, and C. S. Frenk, “Evolution of the cosmic web,” Mon. Not. Roy. Astron. Soc. **441**, 2923 (2014), arXiv:1401.7866.
- [8] Planck Collaboration, “Planck 2018 results. VI. Cosmological parameters,” Astron. Astrophys. **641**, A6 (2020), arXiv:1807.06209.
- [9] L. Shamir, “Analysis of spin directions of galaxies in the DESI Legacy Survey,” Mon. Not. Roy. Astron. Soc. **516**, 2281 (2022), doi:10.1093/mnras/stac2372, arXiv:2208.13866.
- [10] E. Tempel, A. Tamm, M. Gramann, T. Tuvikene, L. J. Liivamägi, E. Saar, P. Heinämäki, P. Nurmi, and J. Einasto, “Flux- and volume-limited groups/clusters for the SDSS galaxies: catalogues and mass estimation,” Astron. Astrophys. **566**, A1 (2014), arXiv:1402.1350.
- [11] H. I. Ullah, M. Awais, T. Matos, and J. F. Suárez-Pérez, “Cosmic-web quenching with DESI DR1: T-Web environments and mass-dependent red/blue classification,” preprint (2026), arXiv:2604.02463.
- [12] D. C. Zapata-Zuluaga, S. Guevara-Montoya, V. Torres-Gomez, J. Hernandez, and J. E. Forero-Romero, “The Cosmic Web in the DESI Early Data Release: A Probabilistic Environment Catalog,” (2026), arXiv:2604.01456.
- [13] H. Rincón, S. BenZvi, K. A. Douglass *et al.*, “DESI-VAST: 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.
- [14] M. Walmsley *et al.*, Galaxy Zoo DESI: large-scale automated morphology classification of 8.7 million galaxies in the DESI Legacy Imaging Surveys, Mon. Not. R. Astron. Soc. **526**, 4768 (2023), arXiv:2309.11425.