capstract:2611.00011v1 [astro-ph.CO] 5 Nov 2026

Matched-circle degeneracy from a compact foreground in locally Euclidean skies

Anja Miretskova,[1,*] Idriss Benkouri,[2] Mateo Junquera Salvat,[3] and Reiko Satomura[1,4]

[1] *Velt Centre for Theoretical Cosmology, Utrecht, Netherlands*
[2] *Laboratoire de Géométrie Observationnelle, Institut Tassiri, Oujda, Morocco*
[3] *Departamento de Física del Universo, Instituto de la Ría Alta, Vigo, Spain*
[4] *Division of Space Geometry, Seihoku Institute of Fundamental Science, Sendai, Japan*

(Dated: November 5, 2026)

A match between two circles on a sky map is evidence for spatial identification if the shared structure samples the emitting surface and foreground repeatability is adequately controlled. We construct a flat, simply connected mock sky containing two images of one spatially compact foreground. The images are produced using an independently surveyed optical glide and are injected upstream of smoothing and pixelization. At a threshold calibrated to a 5% false-alarm rate on Gaussian skies, 873 of 1,024 foreground realizations are classified as matches. The recovered separation and relative circle orientation agree with those expected from a fixed-point-free, orientation-reversing Euclidean holonomy. A mask defined from the foreground acquisition reduces the number of matches to 49, while retaining 902 of 1,024 diffuse topology injections. We give the geometric conditions under which a compact foreground reproduces the inferred holonomy, and discuss which follow-up observables can distinguish a shared patch from a common emitting surface.

I. INTRODUCTION

The distance to last scattering limits a direct search for noncontractible loops. In a flat model the presence of an isometry relating the observer to an image at distance $d < 2R$ produces an intersection between the sphere of last scattering and its image, where R is the comoving radius to emission. Following [1], the intersection can be sought by correlating the anisotropy measured on two circles. The center separation, angular radius and relative phase carry different information about the isometry. In particular, the antipodal restriction used in several searches trades coverage of spatial topologies for a considerable saving in computation [2–4].

Foreground masks in this analysis are usually derived from the intensity or spectrum of emission per sky pixel. A mask serves a different purpose from a requirement on the number of independent patches supporting the match. Consider a sharp feature with two images. Its contribution after azimuthal filtering can dominate both ring variances, even when most pixels around the rings sample uncorrelated sky. In a search with enough geometric freedom, points on either image are inexpensive to bring into correspondence. Once the images have appreciable angular extent, the relative orientation and the curvature of the image centrelines matter. The extent therefore offers an opportunity to reject the match, and also gives a repeated foreground the means to select a well-defined isometry.

We test this possibility with a deliberately narrow question. Can one compact object, acquired independently of the sky generator, return the circle geometry of a specified glide, including the sign of the azimuthal mapping? We use a fiducial distance in the optical problem and report its cosmological equivalent exclusively as d/R. There is a direct dimensional ambiguity between a small covering space and a small foreground when the analysis is carried out on angular data. Predictions for the common sky statistic and for the geometry require separate calculations. Earlier numerical work on compact defects largely used shapelets or points with an assigned parity [5, 6]. That removes uncertainty in the input covariance along an arc but also removes most of the structure from which the search might infer the phase.

Our acquisition and geometry are described in Sec. II. Sec. III gives the map ensembles and fixes the searched family. All rates in Sec. IV refer to maxima over that family. The search is deliberately smaller than an all-sky blind survey. This is a test of a mechanism under known observing conditions. A probability derived for the whole sphere in a topology survey would require independent calibration of the all-sky maximum [7, 8].

II. GEOMETRY AND FOREGROUND ACQUISITION

A. Comparison of the two ring phases

Let g be an isometry of Euclidean covering space and write $g(X) = AX + t$, with A orthogonal. Let O denote the observer. The intersection point X obeys $|X - O|^2 = |X - gO|^2 = R^2$. Subtracting the two sphere equations and repeating the construction at $g^{-1}X$ gives the pair of ring centers and a common radius,

* a.miretskova@velt-cosmology.example

$$c_+ = \frac{gO - O}{d}, \qquad c_- = \frac{g^{-1}O - O}{d}, \qquad \alpha = \arccos\!\left(\frac{d}{2R}\right), \quad d = |gO - O|.$$

We measure azimuth in bases oriented by the two outward centre normals. Since $c_- = -A^{\mathsf{T}} c_+$, the basis in the second circle related by A^{T} alone has the opposite orientation if $\det A = +1$. The relation between the two recorded azimuths is therefore

$$\phi_- = -(\det A)\phi_+ + \phi_0 \pmod{2\pi}. \tag{2}$$

The additive phase ϕ_0 depends on the convention used to place the polar axis of each ring. Throughout this paper a *preserved* phase means a positive coefficient of ϕ_+ in Eq. (2). The designation is made after fixing the two local bases. In particular, reversing both scan tapes upstream of map making leaves this coefficient invariant.

We use a glide in the xy plane, $g(x, y, z) = (-x, y + L, z)$. At $O = (a, 0, 0)$, the two displacements are $(-2a, L, 0)$ and $(-2a, -L, 0)$. Their separation on the observer's sky and the displacement norm are

$$\beta = 2 \arctan \frac{L}{2a}, \qquad d = \sqrt{4a^2 + L^2}. \tag{3}$$

An individual plane reflection has order two, and its two centre directions are equal. We require two distinct centers and use the glide, for which g^2 translates by $2L$. For our choice of scale $2L > 2R$. The two rings predicted by g and g^{-1} constitute one unordered pair. In a compact quotient with two further translations of lengths 8.40 and 7.20 m along x and z, respectively, all additional group elements give displacements larger than $2R$ at the selected observer. Such a quotient can generate one detectable pair without violating closure. The comparison simulations in Sec. III use that quotient.

B. Optical input

The calibration fixture belonged to a fitting-room demonstrator. It consisted of a side mirror carried on a surveyed rail and an upholstered stool of the type supplied with the enclosure. We removed the curtain and illuminated the stool diffusely. The rail had two locating holes separated by L parallel to the mirror. A direct exposure at one hole and a folded exposure at the other fixed the translation and reflection in Eq. (3). The ray tracer used the measured plane, rest-frame range and image handedness. The back of the seat was smooth and the front had a woven check. Moving the source to bring the check into both views was allowed only before acquisition. All reported injections use the front of the same seat.

In the optical problem $R = 3.10$ m is the radius of a virtual emission shell in the ray tracer. Light from the stool originates over a finite range and the seat is not physically at this radius. Its radiance image is carried to the shell along the associated rays. We use the measured angular brightness field of the direct view as a two-dimensional template on a tangent plane, transport it to the shell at one ring crossing, and use the folded projection for the other crossing. This operation preserves the local surface-brightness pattern to first order in the source size. It does not require brightness conservation in the presence of unequal mirror throughput. Each acquisition was flat-fielded against a diffuse reference, and the relative gain in the map injections was measured on independent exposures of the seat. Direct and folded images were normalized to the same filtered ring variance when making the comparison in Fig. 2.

The source center was set 0.3 degrees inside the nominal first ring to avoid exact placement of the fabric seam on the ring. The ring passed through 17 degrees of the azimuthal span of the visible upholstered edge. Most of the common structure in that section came from the difference between the elongated seam and its two adjacent rows of checks. The support mask, fixed before generating any skies, covers the full measured template support plus 1.5 degrees of radial margin around it. The stool base and mirror edges were outside this mask because they were outside the input illumination aperture. Tests injecting the base in isolation were conducted using a separate plate. The median raw score in that test was 0.48, compared with 0.72 for the front edge, and 116 of 1,024 realizations passed. The source surface entered the model through the brightness measurement alone. No stool dimension was fitted against the sky statistic.

The values needed to predict the sky geometry are listed in Table I. Uncertainties on a and L are standard uncertainties from independent surveys. R is an exact input to the numerical projection. The mean azimuth difference $\phi_0 = 0$ in the transported bases was fixed by the measured mirror normal. An offset of 0.35° between the projected image axes serves as the rms angular registration uncertainty. We use direction cosines from the surveyed plane for placement and sample this uncertainty for each map. Multiple optical placements in one realization would introduce extra foreground correlations. We retain the original pair.

TABLE I. Optical configuration and resulting sky geometry. The angular uncertainties propagate the survey errors in the first two rows.

Quantity	Value	Unit
Distance to plane, a	0.920 ± 0.005	m
Translation, L	3.820 ± 0.010	m
Shell radius, R	3.100	m
Displacement, d	4.240 ± 0.010	m
Circle radius, α	46.85 ± 0.13	degree
Center separation, β	128.56 ± 0.27	degree

The third digit in d is retained to define the simulation. The fractional displacement is 1.368. In a cosmology with $R = 14$ Gpc the same angular solution would be assigned a displacement of 19.15 Gpc. Without the range measurement, the circle geometry gives no preference between the latter value and 4.240 m. This equality holds for the inferred displacement. It cannot establish a correspondence between the surface in the optical fixture and the perturbations at last scattering.

III. MAP ENSEMBLES AND SEARCH

A. Acquisition model and covariance

We use maps on the sphere with $N_{\mathrm{side}} = 128$ and truncate the spherical harmonics at $\ell = 256$ [9]. The null anisotropy is a mean-zero Gaussian sky with

$$C_\ell \propto \frac{\exp[-\ell(\ell+1)/180^2]}{\ell(\ell+1)}, \qquad 2 \le \ell \le 256. \quad (4)$$

The normalization is set to $40\,\mu\mathrm{K}$ pointwise rms before convolution, and we add independent $4\,\mu\mathrm{K}$ pixel noise after convolving with a circular beam of FWHM 0.75°. The choice in Eq. (4) is for a locally smooth primary field with finite small-scale power. It keeps the ring covariance well conditioned at the pixel size. It omits the acoustic peaks, gravitational lensing, Doppler projection and the integrated Sachs–Wolfe effect. These enter a quantitative search for topology with CMB data, including matched polarization maps [10, 11]. Our μK conversion of the optical signal fixes an amplitude relative to a mock sky. The conversion does not set a microwave emissivity.

We draw 2,048 calibration skies, then 1,024 independent base skies for evaluation. Calibration and evaluation use independent streams, initialized by the integer seeds 105610 and 105612. The Gaussian case and the foreground case in Table II use the same evaluation base skies, so their difference uses paired observations. Diffuse topology maps use independent quotient-compatible fields with the same covariance law projected to the available modes, followed by the same noise and beam. We match unsmoothed rms to the null ensemble before map projection. For the quotient simulation we symmetrize modes on the double cover, periodic with side lengths $(8.40, 7.64, 7.20)$ m, by averaging a periodic realization with its image under g. The mode amplitudes follow a three-dimensional radial spectrum proportional to $k^{-3}\exp[-(kR/180)^2]$ for $k > 0$, sampled through $kR = 256$. Its projection is similar in slope to Eq. (4). Each set of power comparisons is made within its own ensemble. We do not use agreement in the mean angular power to establish identical distributions.

In the foreground simulations the two image sections are added in linear units prior to the beam. A gain multiplies both templates such that the mean weighted foreground power at the nominal rings is 2.57 times the mean weighted power of the corresponding Gaussian rings. That factor applies after beam, interpolation and azimuthal filtering. The amplitude was determined on a separate set of 128 trials with a uniform bath and held fixed for the primary analysis. It corresponds to local contrast much larger than the pointwise sky rms because the image occupies less than a quarter of the circumference. A full-sky estimate of the brightness variance represents a different weighting. Using it to choose injection amplitude after selecting a candidate could artificially sharpen the consistency with the fixture.

B. Search statistic and decision rule

For a trial pair, we interpolate 512 samples per circle and retain Fourier modes with $8 \le |m| \le 128$. Subtracting modes $|m| < 8$ reduces the effect of coherent diffuse structure. Given the sign ϵ in Eq. (2) and the tested phase δ, we use

$$S(\epsilon, \delta) = \frac{2\sum_m |m|\,\mathrm{Re}[u_m v_{\epsilon m}^{*} e^{-im\delta}]}{\sum_m |m|(|u_m|^2 + |v_m|^2)}. \quad (5)$$

Here u_m and v_m are coefficients in the first and second oriented bases. This is the high-azimuthal-frequency weighted statistic used in related searches [2, 4]. Its symmetry under exchange of the two rings is useful when the pointing record does not retain an ordering. The weights also give the image edge more influence than an equally bright feature with vanishing gradient. For masked evaluations, we apply the union of the two exclusion arcs after phase alignment in real space, taper the new edges by 2°, and recompute the coefficients for each tested phase. Repeating the full procedure on the calibration maps accounts for taper-induced mode coupling.

Trials extend over the ring radii 42° through 52° at 0.5° spacing and over center separations 115° through 140° at 1° spacing. At each pair of angles, orientations

vary about the surveyed bisector and plane normal, with a tangent-plane grid of radius 3° and spacing 1° for the bisector and rotations of -2° through 2° in 1° steps about it. Phase is optimized within 12° of the transported reference on a grid with spacing 1°. Both signs are fitted. We retain the discrete maximum for the detection decision. For parameter recovery at fixed sign, we evaluate a finer grid around it and use the highest value for reporting α and β. Finer sampling does not revise the trial count used to set the threshold. Figures that show the antipodal comparison optimize phase on the same angular window after transporting the second basis to its new center.

The calibration sets the false-alarm threshold at the 95th empirical percentile of the maximum, giving 0.56 unmasked and 0.60 masked after rounding. Decisions use unrounded thresholds. Mask geometry is retained on the feature-free calibration trials. For each sky we rotate the pair of templates and its reference orientation together by an independent uniform rotation to average over the relation to the numerical pixel grid. The foreground placement and trial grid rotate with the same operation. This conditioning is the reason the rates in Table II cannot be compared directly to thresholds from a free search over circle centers. Sky scanning introduces strong dependence between candidates even with the calibration window held fixed [12, 13].

IV. RECOVERY TESTS

A. Frequencies and recovered geometry

The detection rates, quoted with Wilson 95% binomial intervals [14], are given in Table II. The Gaussian evaluation set checks calibration and is statistically compatible with 5% both before and after application of the placebo mask. The foreground produces a much higher raw pass fraction. The excess nearly disappears when the two copies of the seat are excluded independently of the statistic. In the quotient simulation, the same mask reduces the pass fraction by three percentage points. Both phenomena can be attributed to the relative distribution of common structure around the two circles.

TABLE II. Fraction of maxima above threshold in independent evaluation skies. Foreground-bearing and Gaussian cases are paired before masking. Intervals in parentheses are marginal Wilson intervals in percent.

Map ensemble	Raw count	Raw rate (%)	Excised count	Excised rate (%)
Gaussian sky	51 / 1024	5.0 (3.8–6.5)	48 / 1024	4.7 (3.6–6.2)
Repeated foreground	873 / 1024	85.3 (83.0–87.3)	49 / 1024	4.8 (3.6–6.3)
Diffuse topology	934 / 1024	91.2 (89.3–92.8)	902 / 1024	88.1 (86.0–89.9)

In the repeated-foreground case, 864 of the 873 detections preferred preserved phase. Among the detections, the median fitted α was 46.8° and the middle 68% of fitted values spanned 45.9° to 47.9°. The central separation estimate was 128.7° with middle 68% range 126.9° to 130.6°. This scatter describes recovery conditional on a detection. It is larger than the uncertainty from metrology, chiefly because displacing a ring slightly does not remove the seat pattern if the ring remains within the edge's beam footprint. The two effects are distinguishable when the survey tolerance and simulation width are quoted separately. Coverage of the reported central 68% ranges for individual fits is not asserted.

Fig. 1 compares the sky scores of the foreground and diffuse classes as a function of the ring separation. Here we profile at $\alpha = 47^\circ$ to separate the effects of ring placement and preferred sign. A wide nonzero floor arises because the sky score at zero signal is an optimized quantity. At the antipodal comparison we force the rings to a new relative separation and repeat the local orientation and phase search. In that test the foreground and Gaussian profiles are nearly indistinguishable. The large-separation profiles show no detectable excess.

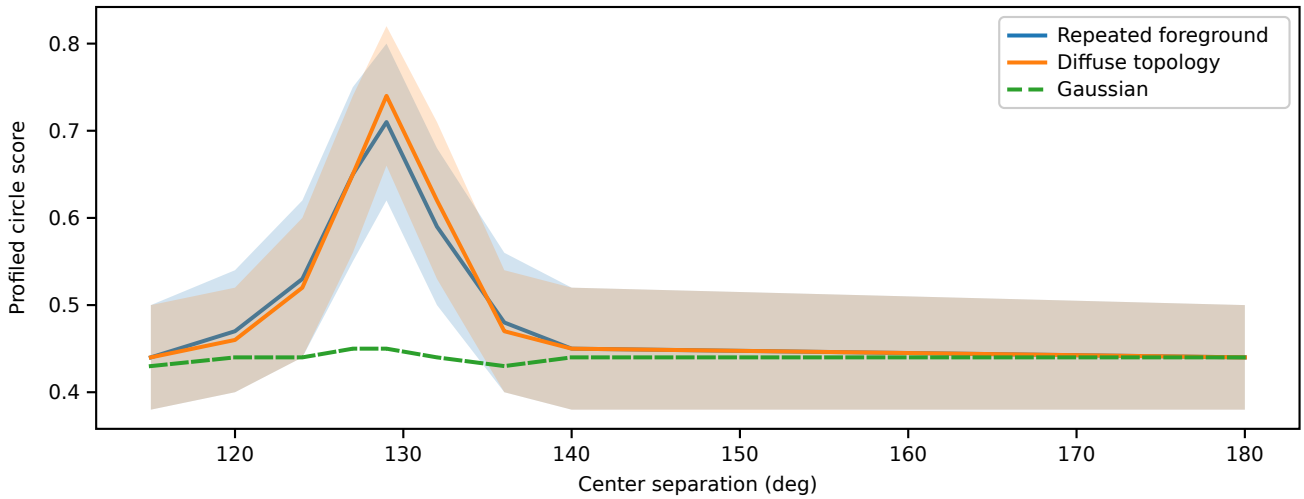


FIG. 1. Median profiled score at fixed angular radius. Bands extend one half of the central 68% interval width on either side of the median. The 180-degree comparison uses its own recentered grid.

B. Amplitude and participation along a ring

If the two rings contain the same weighted template F at scale λ, and weighted backgrounds with equal ensemble power B and negligible mean cross power, the ratio of their expected numerator and expected denominator in Eq. (5) is

$$\bar{S}_{\mathrm{pow}}(\lambda) = \frac{\lambda^2 P_F}{B + \lambda^2 P_F}. \tag{6}$$

This is a population ratio. For finite sampled Gaussian rings, the expectation of the realized ratio differs from the ratio of the expectations. The calculation also neglects the minor misalignment introduced by the mirror normal perturbation. In Fig. 2 we measure the fixed-geometry score without selecting on the maximum, and compare it to Eq. (6). The relation explains why the paired image can dominate the score well before it is a large contribution to the sky mean. It also predicts a saturation that depends on the coherent power fraction. Scaling the whole sky to maintain fixed pointwise rms leaves Eq. (5) invariant. It does nothing to cure the foreground.

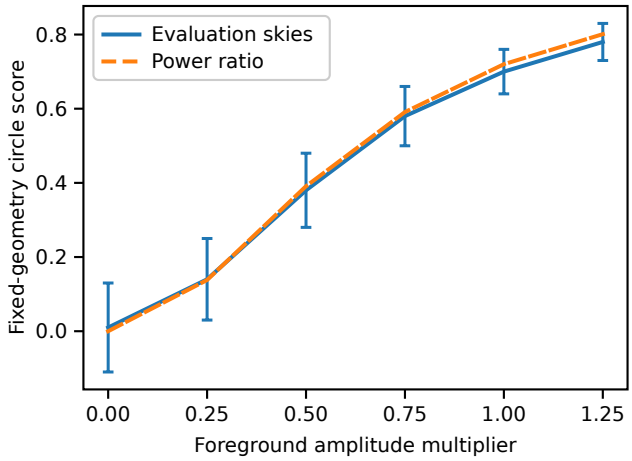


FIG. 2. Fixed-geometry mean score in the foreground ensemble as a function of the amplitude multiplier. Bars give ensemble standard deviations. The dashed line gives the ratio of ensemble powers.

A matched pair dominated by the same object on both rings has a concentrated contribution to the correlation numerator. To measure this, we divide the phase-aligned pair into 16 azimuth sectors and compute low-mode-subtracted cross power in each, using a localized version of the ring filter. Let p_j be the positive part of that cross power. Our participation statistic is

$$D = \frac{(\sum_j p_j)^2}{\sum_j p_j^2}. \tag{7}$$

In Eq. (7) we set $D = 0$ if all p_j vanish. D is a summary of where the correlation is, not a sample size. Overlap of the local strips and the common fitted phase would invalidate a test based on treating the 16 numbers as independent [15, 16]. The unmasked detections had median $D = 2.3$ in the foreground case and $D =$

11.8 in the quotient case. The ratio differs by more than the change induced by replacing the azimuth weighting $|m|$ with uniform weights over the same band. Uniform weighting reduced the median difference in D from 9.5 to 8.1. The latter filter requires new maximum-score calibration if used for decisions.

Fig. 3 gives the mean sector contributions for the detected realizations, scaled by the sum of positive contributions for each fit. For the foreground plots we align sector 0 to the independently registered seat edge. For the diffuse plots the same operation defines sector 0 for the null placement of the hypothetical source. A median of 74% of the positive correlation power was carried by the two sectors nearest the registered upholstery edge in the foreground sample. The equivalent median for the diffuse case was 13%. We report the selected distributions because the practical question is what to do once a match is found.

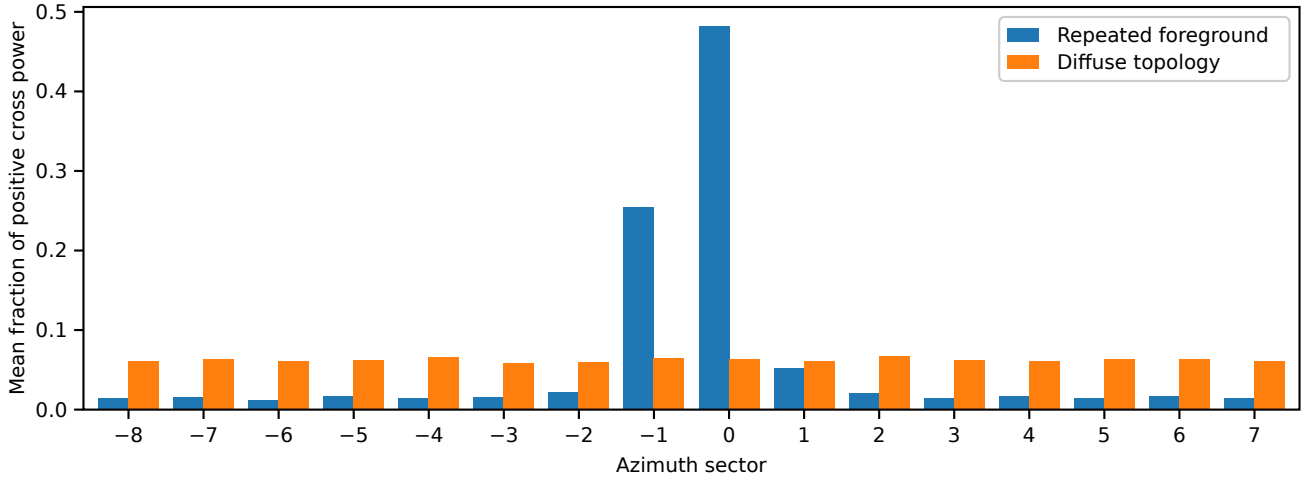


FIG. 3. Sector participation at the maximizing geometry, conditional on detection. Sector zero contains the registered seam. Negative local contributions enter with zero weight.

V. DISCUSSION

An orientation-reversing matched circle is often easier to recognize if the common structure resolves several distinct elements. A single sharp spot gives little evidence about which way around the circle the second spot is traversed. The source in this experiment had an ordered seam and resolved brightness structure on both sides. Exposure geometry and relative smoothing determine whether the angular resolution is sufficient to read the ordering. Rotating the second image by 180° in its tangent plane, while leaving its center fixed, reduced the proportion of preserved-phase detections in the 1,024 trials from 864 to 45. Most new maxima were at reversed phase. This confirms that relative handedness in our result comes from the source image. A uniform disc with the measured integrated flux gave little discrimination between the signs. This behavior is relevant when an estimated holonomy is supported primarily by one bright region. The ranking then depends on the local image orientation.

A prediction of repeated polarization should be considered in any application to CMB data. Stokes components are basis-dependent; the spin transformation and propagation through the optical system need to be included when forming an independent match [17, 18]. Our intensity templates alone give no prediction for the magnitude or orientation of a repeated polarized signal. Even a high polarization statistic on the same pixels would require a check of whether it arose from the shared primary source or from the optical response on those pixels [19]. In the diffuse calculation a common distant field supplies independent spatial structure in additional sectors. The sector balance is therefore more informative in this fixture [20].

There are limitations to comparing power fractions across these examples. The diffuse simulations have spatially extended quotient fields after symmetry projection, whereas the image case has a deterministic contribution with known position. Foreground structure in a survey is neither exactly deterministic nor independent of its mask, and the beam may depend on scan direction. Strong residuals need not be optical images. An independently traced pair from a catalog of repeated intensity structures could serve as the calibration for a more realistic null [21]. The forward model would then need physical signal levels and a frequency dependence. The present range and brightness calibrations suffice only for the local geometry test. Moving the observer or the seat changes the relationship between the two arc masks and the ring pair, and we do not average across fitting-room layout.

APPENDIX A: CIRCLE AND EXTRACTION CONVENTIONS

For the two ring centers in Eq. (1), let $e_{+,1}$ be the direction of the projected vertical on the first tangent plane. We make the second basis vector $e_{+,2} = c_+ \times e_{+,1}$. The first basis vector on the second ring is the image of $e_{+,1}$ under A^{T} and the second vector again follows from a cross product with the outward centre normal. Near the best-fit location the projected vertical has norm close to unity, so the chart of phase is continuous across the whole search region. With these conventions the sampled direction on a circle is

$$n(\phi) = c\cos\alpha + (e_1\cos\phi + e_2\sin\phi)\sin\alpha. \tag{8}$$

To obtain Eq. (2), use $A^{\mathsf{T}}(u \times v) = (\det A)(A^{\mathsf{T}}u) \times (A^{\mathsf{T}}v)$ and $A^{\mathsf{T}}c_+ = -c_-$. Thus the mapping of $e_{+,2}$ gives the sign $-\det A$. Assigning phase from two north-pole meridians independently generally introduces a nonzero ϕ_0, even when no registration perturbation is present. The basis from Eq. (8) is transported with the fixture rotation. It does not use information from the image itself.

APPENDIX B: SCALE AND NUMERICAL CHECKS

Differentiating Eq. (3) gives the two error-propagation weights for β, $-4L/(4a^2 + L^2)$ and $4a/(4a^2 + L^2)$ with respect to a and L. For α, the derivative with respect to d is $-[(2R)^2 - d^2]^{-1/2}$. These derivatives and the independent survey uncertainties give the angle errors quoted in Table I. The d/R scaling has no additional uncertainty from range because the numerical shell radius is fixed. The table entries are rounded separately.

Table III summarizes independent checks on the raw maximum. The median shifts use the same base skies. In the beam test we alternate larger and smaller beams and report the shift for the larger beam. Changing the beam modifies both background and foreground variance. The template gain remains fixed in this test, so the comparison includes the actual loss of seam contrast. A reflected image was created from a fresh frame with the stool cover reversed only in the indicated acquisition check. Removing the total intensity monopole from either input before ray projection changed the median maximum by less than 0.002. The mode cut accounts for most of that insensitivity.

TABLE III. Sensitivity of the median raw foreground maximum to individual changes. Values are paired shifts relative to the primary reduction.

Change	Median shift in S
Double azimuthal sample count at fixed mode cutoff	$+0.001$
Increase beam FWHM from 0.75° to 0.90°	-0.024
Suppress the measured registration scatter	$+0.009$
Use a fresh acquisition with cover reversed and identical masks	-0.081

DATA AND REPRODUCIBILITY

The direct and folded intensity plates, fixture registration, input masks and seeds are maintained with the Velt Centre's release VT-48. Pixels are indexed in ring order, angular coordinates are in radians in the release, and the camera registration uses right-handed room coordinates. Monte Carlo tables count complete maps. No pixel, circle or phase trial is resampled as an independent observation.

ACKNOWLEDGMENTS

We thank the Elmen Court Technical Collection for access to the enclosure and its original seat cover, and V. Alzeta for confirming the right-handed rail convention.

[1] N. J. Cornish, D. N. Spergel, and G. D. Starkman, Classical and Quantum Gravity **15**, 2657 (1998).
[2] N. J. Cornish, D. N. Spergel, G. D. Starkman, and E. Komatsu, Physical Review Letters **92**, 201302 (2004).
[3] A. Riazuelo, J.-P. Uzan, R. Lehoucq, and J. Weeks, Physical Review D **69**, 103518 (2004).
[4] P. Bielewicz and A. J. Banday, Monthly Notices of the Royal Astronomical Society **412**, 2104 (2011).
[5] I. Sardel and S. Ostadpour, Circulars of Observational Geometry **12**, 33 (2018).
[6] M. Wakayori and C. Fenaud, Journal of Spatial Cosmology **7**, 89 (2021).
[7] Planck Collaboration, Astronomy and Astrophysics **594**, A18 (2016).
[8] R. Aurich and S. Lustig, Monthly Notices of the Royal Astronomical Society **433**, 2517 (2013).
[9] K. M. Górski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelmann, The Astrophysical Journal **622**, 759 (2005).
[10] R. K. Sachs and A. M. Wolfe, The Astrophysical Journal **147**, 73 (1967).
[11] U. Seljak and M. Zaldarriaga, Physical Review Letters **78**, 2054 (1997).
[12] E. Netrel and A. Bangoura, Annals of Survey Inference **3**, 227 (2020).

[13] A. Banet and L. Pecar Vrdoljak, Journal of Spatial Cosmology **5**, 180 (2019).
[14] E. B. Wilson, Journal of the American Statistical Association **22**, 209 (1927).
[15] R. A. Tornel and M. Skorvan, Annals of Survey Inference **5**, 73 (2022).
[16] S. Merak and A. Ostervall, Circulars of Observational Geometry **18**, 127 (2024).
[17] M. Zaldarriaga and U. Seljak, Physical Review D **55**, 1830 (1997).
[18] P. Bielewicz, A. J. Banday, and K. M. Górski, Monthly Notices of the Royal Astronomical Society **421**, 1064 (2012).
[19] N. Izumiya and C. Prévot, Transactions on Precision Radiometry **17**, 91 (2023).
[20] E. Sarela and M. Angot, Annals of Survey Inference **1**, 41 (2017).
[21] D. Siret and B. Cardeña, Circulars of Observational Geometry **19**, 5 (2025).