Local bandpower estimation of hemispherical power asymmetry in CMB temperature and E-mode polarisation

Robert A. Lynch
*Independent Researcher, Miami, FL**

The hemispherical power asymmetry (HPA) in temperature is one of the most persistent statistical anomalies in the cosmic microwave background. Extending the measurement to E-mode polarisation requires correcting for E-to-B leakage introduced by incomplete sky coverage; existing approaches address this through inpainting of masked regions. We present a local bandpower estimator (LBE) that corrects for leakage deterministically through per-patch coupling matrix deconvolution, without inpainting. Applied to overlapping patches on the masked sky, the estimator yields simultaneous, deconvolved EE, BB, and EB bandpower estimates at a specified multipole range within each patch, a spectral decomposition not available from existing local estimators. We introduce half-mission directional consistency as a test statistic for spatial anisotropy, exploiting the independent noise realisations in Planck half-mission data splits. The LBE is independent of the test statistic: the deconvolved bandpower maps can be analysed with any statistical test. We demonstrate the framework on Planck PR3 data at $\ell = 15$–40. The temperature HPA direction is recovered with sub-degree half-mission agreement. In E-mode polarisation, directions are consistent with Gimeno-Amo et al. (2023), with p-values of 3.0% for SEVEM and $< 0.33\%$ for Commander. Without deconvolution, the recovered directions shift to the galactic plane and the half-mission agreement becomes consistent with isotropy ($p = 16.7\%$), confirming the necessity of leakage correction. The spectral decomposition yields EE coherence with null BB and EB, consistent with scalar primordial modulation. Signal injection tests validate the pipeline but establish that the directional consistency test is noise-limited at Planck sensitivity, motivating application to LiteBIRD.

I. INTRODUCTION

The assumption of statistical isotropy — that the statistical properties of the cosmic microwave background (CMB) are invariant under rotations — is a foundational prediction of the standard ΛCDM cosmological model, supported by the simplest models of inflation [1]. Observations by WMAP [2] and Planck [3] have confirmed this assumption to remarkable precision across most angular scales. However, a number of unexpected features on large angular scales have been identified in the temperature fluctuation field, collectively termed "CMB anomalies" [4, 5]. These include the hemispherical power asymmetry [6, 7], the alignment of the quadrupole and octopole [8] and the CMB Cold Spot [9].

Among these anomalies, the hemispherical power asymmetry (HPA) is one of the most persistent and well-characterised. First identified in the first-year WMAP data [6, 7], the HPA manifests as a dipolar modulation of the local CMB power spectrum: one hemisphere of the sky exhibits systematically more power than the other. The anomaly has been confirmed in WMAP data releases [10] and in Planck data [5, 11, 12], with a statistical significance exceeding 3σ in temperature. The preferred direction is consistently found near $(l, b) \approx (209°, -15°)$ in galactic coordinates. The amplitude of the modulation is scale-dependent, decreasing with multipole and apparently vanishing above $\ell \sim 600$ [5, 10, 13], though anomalous clustering of local bandpower directions has been detected at multipoles up to $\ell \sim 2000$ [14].

Temperature measurements of the CMB have reached the cosmic variance limit at the scales relevant to HPA. The E-mode polarisation of the CMB provides a largely independent probe of the same primordial fluctuations. If the HPA originates from a modulation of the primordial power spectrum — as proposed in several inflationary scenarios [15] — the same modulation should imprint a corresponding asymmetry in the E-mode polarisation, although the temperature and polarisation transfer functions differ, potentially mapping the signal to different angular scales [16]. For a scalar modulation, the expected signal is concentrated in EE; BB is zero at leading order [17] and affected only at second order with amplitudes equivalent to $r \simeq 0.005$ [18]; EB vanishes by parity conservation [19]. A confirmed detection in polarisation with the expected spectral pattern would constitute strong evidence for a primordial origin. Future experiments with improved polarisation sensitivity, notably LiteBIRD [20], are expected to provide definitive measurements.

Several analyses have searched for the HPA in E-mode polarisation. Gimeno-Amo et al. (2023) [21] applied the local variance estimator (LVE) to Planck PR3 and PR4 E-mode maps after inpainting the masked regions with constrained Gaussian realisations. They reported p-values of 0.22% for SEVEM on PR3 data, with a direction of $(l, b) = (232°, -9°)$, and 2.8% on PR4 data. The significance varies across inpainting realisations, reflecting the stochastic nature of the approach. Aluri & Shafieloo (2017) [22] similarly applied the LVE to Planck 2015 Commander E-mode maps and found 2.6–3.9% significance, though with a direction pointing toward the kinematic dipole rather than the established HPA direction, suggesting possible contamination from residual

* r.a.lynch@gmail.com

systematic effects. In a subsequent analysis, Gimeno-Amo et al. (2025) [14] applied the MASTER pseudo-C_ℓ estimator in 12 independent sky regions to test for angular clustering of bandpower dipole directions in polarisation across multipole bins using a Rayleigh statistic. No significant clustering was found in E-mode polarisation. This tests a different question — coherence of directions across multipole bins — from the half-mission directional consistency we probe. The LBE extends the local pseudo-C_ℓ approach with denser spatial sampling (600 overlapping patches), per-patch spectral decomposition into EE, BB, and EB, and a different test statistic. Ghosh & Jain (2020) [16] analysed the squared polarisation amplitude $|P|^2 = Q^2 + U^2$ in pixel space, but found that correlated noise and foreground residuals dominated the Planck polarisation signal, precluding a detection. They did, however, identify the importance of excluding multipoles below $\ell \sim 10$–15 for reliable direction recovery, a finding we confirm independently.

A fundamental challenge for polarisation analyses on incomplete skies is the E–B leakage problem: on a masked sky, the decomposition of Q and U into E-mode and B-mode components is ambiguous, and power leaks between the two [23]. This leakage is correlated with the mask geometry and can introduce spurious directional signals in analyses that first construct an E-mode map and then measure its local properties. The pseudo-C_ℓ framework [24, 25] offers a principled solution: by computing power spectra directly from the masked data and deconvolving the mask-induced mixing through an analytically computed coupling matrix, the true E-mode power can be recovered without ever constructing an intermediate E-mode map.

In this work, we introduce a local bandpower estimator (LBE) that addresses the E–B leakage problem through a different strategy: rather than inpainting masked regions before computing local power, we estimate pseudo-C_ℓ spectra within overlapping patches and correct for mask-induced mode coupling through coupling matrix deconvolution. This provides a deterministic estimator — the same input always produces the same output — that yields separate EE, BB, and EB bandpowers at a specified multipole range within each patch. We also introduce half-mission directional consistency as a test statistic for spatial anisotropy, exploiting the independent noise realisations in Planck half-mission data splits to test whether two independent measurements of the same sky agree on the preferred direction.

We demonstrate the method on Planck PR3 temperature and E-mode polarisation data at $\ell = 15$–40. In temperature, the pipeline recovers the published HPA direction with sub-degree half-mission agreement. In E-mode polarisation, the recovered directions are consistent with the independent analysis of Gimeno-Amo et al. (2023) [21], providing consistency with a complementary methodology. An ablation test in which the deconvolution step is removed confirms that E–B leakage correction is essential: without it, the signal is not recovered. The spectral decomposition — directional coherence in EE with null results in BB and EB — is consistent with a modulation of the scalar primordial power spectrum. Signal injection tests validate the pipeline and characterise its sensitivity at Planck noise levels.

The paper is organised as follows. Section II describes the Planck data. Section III presents the LBE, the normalisation procedure, the directional consistency test, and the FFP10 simulation ensemble. Section IV presents temperature validation, the E-mode polarisation results including the deconvolution ablation test and spectral decomposition, noise characterisation, and signal injection validation. Section V provides robustness tests across pipeline parameters and component separation methods. Section VI discusses the results in the context of previous analyses and characterises the sensitivity of the test. Section VII summarises our conclusions.

II. DATA

We use the Planck 2018 (PR3) maps produced by the SEVEM component separation pipeline [26]. We analyse the half-mission data splits (HM1 and HM2), which share the same sky signal but have independent noise realisations. For polarisation we use the Stokes Q and U maps; for temperature the intensity map, with monopole and dipole removed. We also construct a cross-map as the average of the two half-mission maps: $Q_{\rm cross} = (Q_{\rm HM1} + Q_{\rm HM2})/2$ and likewise for U and T. All maps are analysed at HEALPix $N_{\rm side} = 64$ [27], smoothed with a Gaussian beam of FWHM $= 160'$, matching the resolution used in previous large-scale polarisation HPA analyses [21]. We apply the Planck 2018 confidence masks for temperature and polarisation [26]; the polarisation mask retains approximately 96% of the sky.

For calibration and significance assessment we use 300 simulations from the Planck Full Focal Plane 10 (FFP10) suite [28] processed through the SEVEM pipeline. Matched FFP10 simulation suites are used for Commander (300 simulations), SMICA (300), and NILC (300) to enable cross-checks across component separation methods. Each simulation consists of a CMB realisation drawn from the Planck best-fit ΛCDM power spectrum plus independent half-mission noise realisations matching the noise properties and scanning strategy of the real data:

$$d_i^{\rm HM1} = s_i^{\rm CMB} + n_i^{\rm HM1}, \quad d_i^{\rm HM2} = s_i^{\rm CMB} + n_i^{\rm HM2} \qquad (1)$$

where $s_i^{\rm CMB}$ is the i-th CMB realisation (common to both half-missions) and $n_i^{\rm HM1}$, $n_i^{\rm HM2}$ are independent noise realisations. Under the null hypothesis of statistical isotropy, any directional consistency between simulated half-missions arises solely from the shared CMB cosmic variance and not from an anisotropic signal.

III. METHOD

A. Local bandpower estimator

One standard approach to measuring hemispherical power asymmetry divides the sky into overlapping circular patches and assigns a scalar measure of CMB power to each patch. In the LVE used by previous analyses [11], this scalar is the pixel variance of an E-mode map within the patch. By Parseval's theorem, the pixel variance is equivalent to a weighted sum of the angular power spectrum over all multipoles present in the map:

$$\sigma_i^2 = \sum_\ell \frac{2\ell+1}{4\pi} C_\ell^{\rm patch} \tag{2}$$

This is a valid measure of total power, but it has two limitations when applied to E-mode polarisation on a masked sky.

First, the sum is broadband, weighted by the number of modes per multipole. For a map at resolution $N_{\rm side} = 64$, it includes multipoles up to $\ell \sim 191$, well above the range where the HPA has been detected in temperature ($\ell \lesssim 64$) [5, 13]. A broadband sum therefore dilutes the directional signal with isotropic power from scales where no asymmetry is present.

Second, the LVE requires a global E-mode map as input, from which patches are then extracted. On a masked sky, the decomposition of Q and U into E and B modes is ambiguous: what would be pure E-mode power on the full sky leaks into the B-mode estimate and vice versa. This leakage is determined by the global mask geometry and is baked into the E-mode map before any local measurement takes place. The pixel variance within each patch therefore includes leaked B-mode power that the local measurement cannot separate from the true E-mode signal. Existing analyses address this through inpainting — filling masked pixels with constrained Gaussian realisations before decomposing into E and B [21] — but this introduces a dependence on the assumed covariance and on the particular realisation used to fill the mask.

We take a different approach. Rather than constructing a global E-mode map and measuring its properties locally, we perform the full harmonic analysis within each patch independently, working directly with the Stokes parameters Q and U. No E-mode map is ever constructed. For each patch, we compute the pseudo-power spectrum $\tilde{C}_\ell$ of the masked Q/U field treated as a spin-2 quantity. The pseudo-spectrum is related to the true underlying spectra C_ℓ^{EE}, C_ℓ^{BB}, and C_ℓ^{EB} through a coupling matrix $M_{\ell\ell'}$ that is computed analytically from that patch's specific mask geometry. The coupling matrix encodes both the overall power suppression from incomplete sky coverage and the mixing between E and B modes induced by the mask. Inverting it yields the deconvolved true power spectra for each patch independently:

$$C_\ell^{EE,\rm patch} = \left[M^{-1}\tilde{C}\right]_\ell^{EE} \tag{3}$$

We implement this using NaMaster [25], which computes both the pseudo-spectrum and the coupling matrix for arbitrary spin-2 fields and masks. One matrix inversion yields the full set of deconvolved EE, BB, and EB power spectra within each patch — three spectra, deterministically, from a single operation.

We compress the EE spectrum to a single scalar by averaging over a chosen multipole range with uniform weight:

$$\hat{C}_i^{EE} = \frac{1}{\ell_{\rm max} - \ell_{\rm min} + 1} \sum_{\ell=\ell_{\rm min}}^{\ell_{\rm max}} \left[M^{-1}\tilde{C}\right]_\ell^{EE} \tag{4}$$

The fiducial range $\ell = 15$–40 is motivated by two considerations: the need to exclude the lowest multipoles, which introduce large cosmic-variance bias in pixel-space estimators [16], and consistency with the angular scales at which the temperature hemispherical power asymmetry is most statistically significant [13]. The temperature and polarization transfer functions differ, so the ℓ-space mapping of the same primordial signal could be different for E-mode polarization [16]. Optimisation of the ℓ-range for E-mode sensitivity is deferred to future work. The bandpower averaging also regularises the coupling matrix inversion: the deconvolution operates on a coarsely binned coupling matrix rather than inverting at individual multipoles, which would be poorly conditioned for the small sky fractions typical of local patches.

The LBE can be understood as a refinement of the LVE in which the broadband, $(2\ell+1)$-weighted sum is replaced by a flat-weighted, deconvolved sum over a chosen range. However, the two estimators differ more fundamentally: the LVE extracts patches from a global E-mode map in which E–B separation has already been performed, while the LBE performs the E–B separation independently within each patch using that patch's own mask geometry. The deconvolution corrects E–B leakage and mask-induced power suppression; the band restriction isolates the angular scales of interest and stabilises the matrix inversion. Both ingredients are essential: deconvolution without band restriction amplifies noise from poorly conditioned individual multipoles, while band restriction without deconvolution leaves the E–B leakage uncorrected.

B. Patch layout and selection

We define circular patches of radius $R = 15^\circ$ centred on the pixel centres of a HEALPix grid at $N_{\rm side} = 16$, giving 3072 candidate patch locations. We retain patches in which at least 60% of pixels pass the Planck PR3 confidence mask, yielding 2281 candidates for temperature and 2288 for polarisation. From these, we select $N_p = 600$ by greedy distance maximisation to ensure approximately uniform sky coverage. The patches overlap, which introduces correlations between neighbouring estimates. These correlations do not bias the dipole fit

but affect its variance; we account for this through calibration against simulations rather than analytic error propagation.

The patch radius $R = 15°$ is chosen so that modes at $\ell_{\rm max} = 40$ (corresponding to angular scales of $\sim 4.5°$) fit comfortably within the patch diameter. Alternative radii are tested in section V.

The patch mask $w_i(\hat{n})$ for patch i is the product of the galactic mask and a circular boundary:

$$w_i(\hat{n}) = M_{\rm gal}(\hat{n}) \cdot \Theta(\hat{n} \cdot \hat{n}_i - \cos\alpha_R) \tag{5}$$

where $M_{\rm gal}$ is the galactic mask, Θ is the Heaviside step function, α_R is the angular radius of the patch, and $\hat{n}_i$ is the patch centre. We adopt the unapodised configuration as fiducial for simplicity, noting that apodisation at the scales tested neither improves nor degrades the result.

C. Normalisation

The raw bandpower $\hat{C}_i^{EE}$ in each patch depends on the local noise level, the effective beam, and the patch mask geometry. To isolate the cosmological signal, we normalise by subtracting the mean bandpower obtained from FFP10 simulations processed through the identical pipeline:

$$\hat{C}_i^{\rm norm} = \hat{C}_i^{\rm data} - \langle \hat{C}_i^{\rm sims} \rangle \tag{6}$$

where the average is over 300 FFP10 simulations. The mean is computed separately for each half-mission: HM1 data is normalised by the mean of HM1 simulations, and likewise for HM2. This normalisation removes the spatially varying baseline due to mask geometry and noise, while preserving any dipolar modulation of the cosmological signal. A patch with excess E-mode power relative to the isotropic expectation will have $\hat{C}_i^{\rm norm} > 0$. A dipolar modulation of the primordial power spectrum would imprint a coherent dipole pattern on the bandpower map.

D. Dipole fitting

We fit a dipole to the normalised bandpower map. The N_p patch centres $\hat{n}_i$ and their normalised bandpowers $\hat{C}_i^{\rm norm}$ are fit to:

$$\hat{C}_i^{\rm norm} = c_0 + \mathbf{c} \cdot \hat{n}_i \tag{7}$$

where c_0 is the monopole and $\mathbf{c} = (c_x, c_y, c_z)$ is the dipole vector. The dipole direction in galactic coordinates is $l = \arctan(c_y/c_x)$ and $b = \arcsin(c_z/|\mathbf{c}|)$, and the dipole amplitude is $|\mathbf{c}|$. We fit the dipole using equal weights across all selected patches, with significance calibrated against simulations processed through the identical pipeline.

E. Ablation design

The LVE computes the variance of a reconstructed E-mode map within each patch, without spectral deconvolution or coupling matrix correction. The LBE instead estimates pseudo-C_ℓ spectra from the Q/U maps within each patch and inverts the coupling matrix to obtain deconvolved bandpowers. The two estimators share the same patch layout, normalisation, and dipole fitting procedure; they differ in how the per-patch power is estimated and in their treatment of E–B leakage. Comparing them demonstrates that the LBE pipeline recovers a signal that the LVE does not. We emphasise that this comparison tests the LBE pipeline as a whole within our framework. The LVE used by Gimeno-Amo et al. (2023) [21] incorporates inpainting to mitigate E–B leakage, and its demonstrated sensitivity to the E-mode HPA confirms that leakage correction — whether through inpainting or deconvolution — is the essential ingredient.

F. Half-mission directional consistency test

At $\ell = 15$–40, the small number of independent modes means that an isotropic CMB realisation produces substantial patch-to-patch power fluctuations, and a dipole fit to these fluctuations yields a nonzero amplitude even under isotropy. For individual half-mission maps, the amplitude of the fitted dipole under isotropy is comparable to the expected HPA modulation at $A = 0.07$, making the half-mission dipole amplitude alone a poor discriminator in this framework. However, the *direction* of the dipole carries additional information. We apply the LBE pipeline independently to HM1 and HM2, obtaining dipole directions $\hat{d}_1$ and $\hat{d}_2$, and quantify their agreement through the angular distance:

$$\theta = \arccos(\hat{d}_1 \cdot \hat{d}_2) \tag{8}$$

The behaviour of the test statistic can be understood intuitively. Under isotropy, each half-mission bandpower map is the sum of a shared CMB realisation and an independent noise realisation. The dipole direction recovered from each map reflects the relative contributions of the CMB signal fluctuations (common to both half-missions) and the instrumental noise (independent). In temperature, where the per-patch CMB signal-to-noise ratio is high, both half-missions recover directions driven by the same underlying realisation; the resulting angular separations are small even under isotropy, leaving little room for a modulation to further tighten the agreement. In E-mode polarisation at Planck sensitivity, instrumental noise dominates the per-patch bandpower, scattering the two half-mission directions approximately independently across the sky; the typical angular separation under isotropy is $\sim 76°$. A coherent modulation of the underlying power spectrum adds a fixed dipole component to both bandpower maps, pulling both directions toward

the true preferred axis and reducing the angular separation below the isotropic baseline. The test is therefore most sensitive in the regime where noise is large enough to produce broad scatter under isotropy but a real signal can still produce anomalous agreement — a regime that Planck E-mode data occupy marginally. We calibrate the expected distribution of angular separations empirically using 300 FFP10 simulations rather than analytically, since the overlapping patch geometry and non-uniform noise make this impractical.

The p-value is the fraction of FFP10 simulations with angular distance smaller than or equal to that observed in the data:

$$p = \frac{1}{N_{\rm sim}} \sum_{j=1}^{N_{\rm sim}} \Theta(\theta_{\rm data} - \theta_j) \tag{9}$$

IV. RESULTS

A. Temperature

We first validate the pipeline on the Planck PR3 SEVEM temperature maps, where the HPA is well established with significance exceeding 3σ and a dipole direction of $(l, b) \approx (209°, -15°)$ [5, 11]. The temperature analysis uses identical pipeline settings to the polarisation analysis — the same patch selection, the same multipole range (ℓ = 15–40), and the same normalisation procedure — differing only in the use of spin-0 fields (intensity) rather than spin-2 fields (Q, U).

The LBE recovers a dipole direction of $(l, b) = (222.9°, -31.7°)$ from HM1 and $(222.3°, -31.8°)$ from HM2, an angular separation of 0.48° (table I). This level of directional agreement is exceeded by 26.7% of isotropic FFP10 simulations (80/300). As discussed in section III F, the high per-patch signal-to-noise ratio in temperature means that both half-missions are dominated by the same CMB realisation rather than by independent noise; the resulting null distribution is concentrated at small angular separations (simulation mean 1.02°), leaving little room for a modulation to further tighten the agreement. The recovered direction lies within $\sim 21°$ of the established temperature HPA direction, within the $\pm 31°$ directional uncertainty reported by the Planck Collaboration for the ℓ = 2–64 measurement [5].

The dipole amplitude from both half-missions is individually anomalous: the HM1 amplitude is exceeded by only 7% of simulations in both HM1 and HM2. This reflects the high signal-to-noise ratio of the temperature HPA at these scales, even within the restricted ℓ = 15–40 range. In contrast, the E-mode amplitude will be shown to be consistent with simulations (section IV B), motivating the use of the directional consistency test as the primary probe of anisotropy in polarisation.

The LVE applied to the intensity map yields half-mission directions of $(221.6°, -37.5°)$ and $(221.6°, -37.6°)$ with an angular distance of 0.08° and p = 1.7%. In temperature, the LVE outperforms the LBE on the directional consistency test, as the pixel-space variance captures the full broadband signal while the LBE restricts to ℓ = 15–40. The distinction between the estimators emerges in polarisation, where E–B leakage correction is required.

TABLE I. Temperature (TT) primary results from the LBE at ℓ = 15–40.

	Direction (l, b)	Amplitude	Ang. dist.	p-value
HM1	$(222.9°, -31.7°)$	3.23	—	—
HM2	$(222.3°, -31.8°)$	3.21	—	—
HM1–HM2	—	—	0.48°	26.7%

B. E-mode polarisation

The primary empirical test is whether E–B leakage correction changes the recovered signal. We compare the full LBE pipeline — which takes Q/U maps as input, estimates pseudo-C_ℓ spectra within each patch, and inverts the coupling matrix to obtain deconvolved EE bandpowers — to an LVE that computes the variance of a reconstructed E-mode map within each patch without spectral deconvolution (section III E). The two estimators share the same patch layout, normalisation, and dipole fitting procedure; they differ in how the per-patch power is estimated and in their treatment of E–B leakage.

Applying the LVE without deconvolution to the E-mode half-mission maps yields dipole directions of $(l, b) = (2.8°, -23.7°)$ from HM1 and $(329.3°, -41.6°)$ from HM2. These directions point near the galactic plane, more than 90° from the published HPA direction. The angular distance of 33.1° between them has a p-value of 16.7% — consistent with isotropic simulations. Without leakage correction, the recovered directions are dominated by mask geometry and galactic contamination residuals. This is the expected outcome: the ablation deliberately omits leakage correction, and the failure demonstrates that some form of E–B correction — whether coupling matrix deconvolution or inpainting, as employed by Gimeno-Amo et al. (2023) [21] — is required for polarisation HPA analysis on incomplete skies (table II).

With coupling matrix deconvolution, the LBE recovers a dipole direction of $(l, b) = (234.5°, -19.3°)$ from HM1 and $(219.2°, -13.7°)$ from HM2. The angular distance of 15.8° between them is exceeded by only 3.0% of isotropic FFP10 simulations (9 out of 300); the binomial 95% confidence interval on 9/300 is $[1.4\%, 5.6\%]$. The effect of deconvolution on the recovered signal is substantial: the angular distance drops from 33.1° (p = 16.7%) to 15.8° (p = 3.0%; table III), and the recovered directions shift from near the galactic plane to within $\sim 5°$ of the direc-

TABLE II. Effect of E–B leakage correction: LBE (with deconvolution) versus LVE (without deconvolution) for temperature and E-mode polarisation. The LVE ablation tests the role of deconvolution within the LBE framework.

Band	Estimator	HM1 direction	HM2 direction	Ang. dist.	p-value
TT	LBE	$(222.9°, -31.7°)$	$(222.3°, -31.8°)$	.48°	26.7%
TT	LVE	$(221.6°, -37.5°)$	$(221.6°, -37.6°)$	0.08°	1.7%
EE	LBE	$(234.5°, -19.3°)$	$(219.2°, -13.7°)$	15.8°	3.0%
EE	LVE	$(2.8°, -23.7°)$	$(329.3°, -41.6°)$	33.1°	16.7%

TABLE III. E-mode polarisation (EE) primary results from the LBE at $\ell = 15$–40 on SEVEM data.

	Direction (l, b)	Amplitude	Ang. dist.	p-value
HM1	$(234.5°, -19.3°)$	8.5×10^{-5}	—	—
HM2	$(219.2°, -13.7°)$	2.4×10^{-4}	—	—
HM1–HM2	—	—	15.8°	3.0%

TABLE IV. Spectral diagnostics: half-mission directional consistency for EE, BB, and EB bandpowers from the same coupling matrix deconvolution. EE is the primary result; BB and EB serve as null checks.

Spectrum	HM1 (l, b)	HM2 (l, b)	Ang. dist.	p-value
EE	$(234.5°, -19.3°)$	$(219.2°, -13.7°)$	15.8°	3.0%
BB	$(297.7°, 28.4°)$	$(237.9°, 57.1°)$	50.2°	31.3%
EB	$(143.9°, 60.3°)$	$(16.3°, 7.3°)$	100.9°	58.7%

tion reported by Gimeno-Amo et al. (2023) [21].

The same coupling matrix deconvolution that yields the EE bandpower simultaneously produces BB and EB estimates for each patch. The BB half-mission dipole directions are $(297.7°, 28.4°)$ and $(237.9°, 57.1°)$, with an angular distance of 50.2° ($p = 31.3\%$). The EB half-mission directions are $(143.9°, 60.3°)$ and $(16.3°, 7.3°)$, with an angular distance of 100.9° ($p = 58.7\%$). Neither spectrum shows directional coherence. For a dipolar modulation of the scalar primordial power spectrum, the expected signal is concentrated in EE [17]; BB is zero at leading order for scalar modes and affected only at second order with amplitudes equivalent to $r \simeq 0.005$ [18]; EB vanishes by parity conservation [19]. The observed pattern is consistent with this expectation (table IV). BB and EB bandpowers have not previously been examined in the HPA context with Planck data; the most recent such measurement used WMAP data [29, 30]. The spectral decomposition also provides a qualitative check on the deconvolution: residual E-to-B leakage from incomplete coupling matrix inversion would elevate BB bandpowers preferentially in patches with complex mask geometry, which could introduce spurious structure in the BB bandpower map. The absence of directional coherence in BB is consistent with adequate leakage correction, though we note that this check is not a formal test of deconvolution completeness.

The HM1 direction of $(234.5°, -19.3°)$ is within 0.5° in longitude and 5° in latitude of the $(234°, -14°)$ reported by Gimeno-Amo et al. (2023) [21], despite the two analyses employing independent leakage correction strategies: deconvolution on overlapping patches (this work) versus inpainting on the masked sky. Gimeno-Amo et al. report $p = 0.22\%$ for SEVEM on PR3 data using an amplitude-based test statistic, with a range of $[<0.22\%, 0.44\%]$ across inpainting realisations; on PR4 data the corresponding result is $p = 2.8\%$ with range $[1.8\%, 3.8\%]$. The present analysis yields a single deterministic value of $p = 3.0\%$ on PR3 data using a directional consistency test statistic. The difference in p-values reflects different test statistics applied to different quantities: Gimeno-Amo et al. test whether the amplitude of the local variance dipole in the full E-mode map exceeds isotropic expectations, while the present analysis tests whether two independent half-mission measurements agree on the dipole direction. The directional agreement between the two analyses suggests that the recovered preferred axis is not an artefact of either leakage correction strategy (figure 1).

For Commander, the half-mission directions of $(214.4°, -31.8°)$ and $(208.3°, -24.6°)$ are separated by only 8.74°, with none of the 300 simulations producing an angular distance this small ($p < 0.33\%$; 95% upper bound $p < 1.0\%$). The four methods form a gradient from Commander ($p < 0.33\%$) through SEVEM (3.0%) and SMICA (5.7%) to NILC (25.3%) (table V). The dispersion across methods is expected: component separation pipelines differ in their treatment of large-scale polarised foregrounds and instrumental systematics, producing different residual contamination at low multipoles. The gradient observed here — with Commander and SEVEM showing the strongest signal — is consistent with that reported by the Planck Collaboration (2020) [26] and Gimeno-Amo et al. (2023) [21] for the same methods.

The dipole amplitude from individual half-missions is consistent with the isotropic simulations: HM1 has an amplitude p-value of 86.0% and HM2 of 15.3%. The directional consistency test is sensitive to the coherent signal even when the amplitude test is not (section VI C).

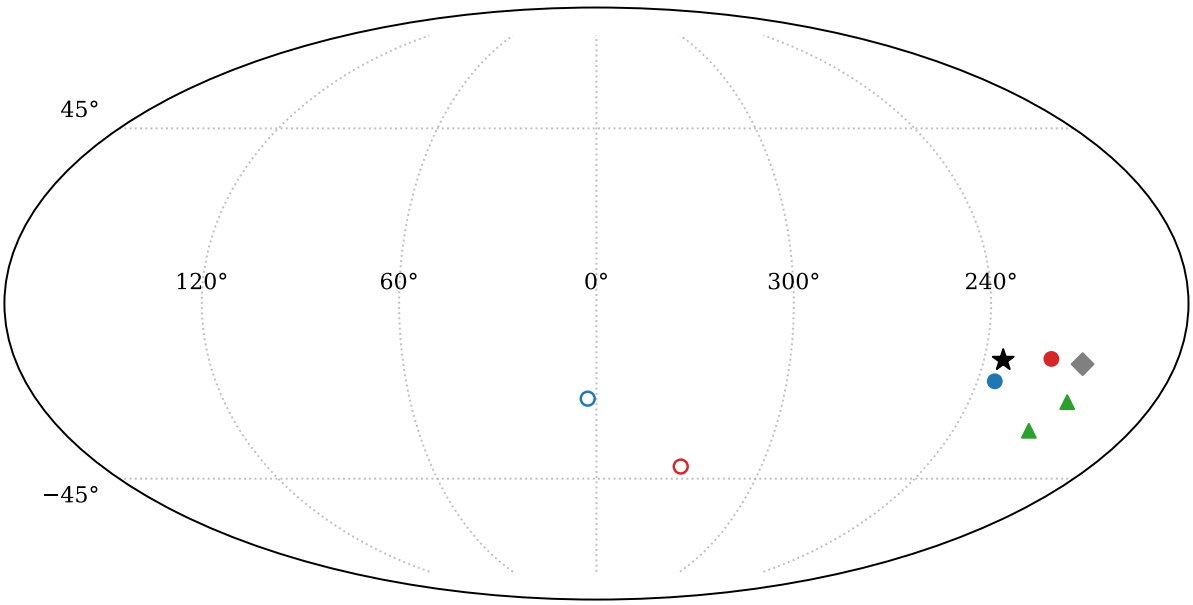

FIG. 1. Recovered dipole directions in galactic coordinates (Mollweide projection). Blue filled circle: LBE SEVEM HM1. Red filled circle: LBE SEVEM HM2. Blue open circle: LVE SEVEM HM1. Red open circle: LVE SEVEM HM2. Green filled triangles: Commander HM1 and HM2. Black star: Gimeno-Amo et al. (2023) E-mode HPA direction $(234°, -14°)$. Grey diamond: published temperature HPA direction $(209°, -15°)$. The LBE directions cluster near the published HPA direction, while the LVE directions (without deconvolution) scatter toward the galactic plane.

TABLE V. E-mode directional consistency test across component separation methods. Each method is calibrated against its own matched FFP10 simulation suite (300 for SEVEM, SMICA, and NILC; 300 for Commander).

Method	HM1 (l, b)	HM2 (l, b)	Ang. dist.	$N_{\rm sim} \leq$ data	p-value
Commander	$(214.4°, -31.8°)$	$(208.3°, -24.6°)$	8.74°	0/300	$< 0.33\%$
SEVEM	$(234.5°, -19.3°)$	$(219.2°, -13.7°)$	15.8°	9/300	3.0%
SMICA	$(168.3°, -42.8°)$	$(191.4°, -22.9°)$	27.6°	17/300	5.7%
NILC	$(271.9°, -60.2°)$	$(187.2°, -40.7°)$	53.1°	76/300	25.3%

C. Noise characterisation

The simulation mean bandpower maps subtracted during normalisation have residual dipole amplitudes of 3.2% of the data dipole in temperature (both half-missions) and 62.8% (HM1) and 153% (HM2) in E-mode polarisation. In temperature, the normalisation baseline is small relative to the signal. In polarisation, the baseline is comparable to or larger than the data dipole, reflecting the dominance of mask geometry and noise patterns over the cosmological signal in the raw per-patch bandpower. The subtraction removes this half-mission-specific baseline, isolating the residual cosmological signal.

To characterise the systematic floor, we apply the pipeline to FFP10 noise-only half-mission maps containing no CMB component and subtract the same simulation mean. In temperature, the noise-only dipole amplitude is 0.04% of the data value, with direction $(l, b) = (37.9°, -37.8°)$, confirming that noise contributes negligibly to the temperature result. In polarisation, the noise-only dipole amplitude is 148.5% of the data value — noise alone produces dipoles larger than the data signal in individual realisations. This is expected in the noise-dominated E-mode regime and is precisely why significance is assessed through comparison to simulations that include the same noise properties. The noise-only dipole direction $(176.8°, 34.1°)$ is 76.6° from the data direction and 72.6° from the published HPA direction, confirming that the directional coherence observed in the data is not produced by a systematic noise pattern. Across the 300 FFP10 noise-only realisations, the recovered dipole directions are uniformly distributed on the sky, with no clustering near the published HPA direction.

D. Signal injection validation

We validate the pipeline through signal injection tests, in which a known dipolar modulation is applied to

TABLE VI. Temperature signal injection results. A scalar dipolar modulation of amplitude A is applied at $(l, b) = (45°, 60°)$ to each of 300 FFP10 simulations. HM→tgt: mean angular distance from the recovered half-mission directions to the injection target. HM1↔HM2: mean half-mission angular distance. $P \leq$ data: fraction of realisations with half-mission angular distance $\leq 0.48°$ (the observed value). All statistics are per-simulation means across 300 realisations.

A	HM→tgt	HM1↔HM2	$P \leq$ data
0	88.3°	1.02°	26.7%
0.07	45.7°	0.79°	36.0%
0.09	37.7°	0.68°	45.0%
0.20	18.4°	0.33°	81.7%
0.50	8.3°	0.13°	100%
1.0	5.6°	0.07°	100%

FFP10 simulations and the modulated maps are processed through the identical analysis pipeline. For each of the 300 simulations, we inject a modulation at the direction $(l, b) = (45°, 60°)$ — deliberately offset from the reported HPA direction — with amplitudes ranging from $A = 0$ to $A = 2$. For each amplitude, we process all 300 modulated simulations and compare the resulting half-mission angular distances to the isotropic null distribution.

1. Temperature

For temperature, we apply a scalar dipolar modulation $T(\hat{n}) \to T(\hat{n})[1 + A\,\hat{d} \cdot \hat{n}]$. Table VI and figure 2 summarise the results. At $A = 0.07$, the mean distance from the recovered direction to the injection target is 45.67°: individual realisations do not reliably recover the injected direction, as the isotropic CMB fluctuations at $\ell = 15$–40 produce a fitted dipole of comparable amplitude. Per-realisation direction recovery becomes effective at $A \gtrsim 0.5$, where the mean distance to the target drops to 8.33°.

The mean half-mission angular distance decreases from 1.02° at $A = 0$ to 0.07° at $A = 1.0$, and $P \leq$ data rises from 26.7% to 100% over the same range (table VI). This response validates the pipeline's sensitivity to dipolar modulations. The modest shift at the physical amplitude ($P \leq$ data rises from 26.7% to 36.0% at $A = 0.07$) reflects the high signal-to-noise ratio of the temperature data: under isotropy, the mean half-mission angular distance is already only 1.02°, leaving little room for a modulation to further tighten the agreement.

2. E-mode polarisation

We consider two injection approaches: pixel-space modulation of the Q/U maps, which applies $Q(\hat{n}) \to Q(\hat{n})[1 + A\,\hat{d} \cdot \hat{n}]$ and equivalently for U; and harmonic-space modulation of the E-mode spherical harmonic coefficients through the ℓ-to-$\ell \pm 1$ coupling induced by a dipole field.

The two methods produce nearly identical results across all amplitudes (table VII), confirming that the choice of injection pathway does not affect the results and that the limiting factor is the per-patch instrumental noise. At $A = 0.07$, the mean distance from the recovered direction to the injection target is 82°, essentially indistinguishable from the unmodulated value of 89.7°. $P \leq$ data rises from 3.0% to only 3.7%. Per-realisation direction recovery becomes effective at $A \gtrsim 2$, where the mean distance to the target drops below 15°.

3. Interpretation

The injection tests characterise the sensitivity of the directional consistency test across signal-to-noise regimes. In temperature, the high per-patch signal-to-noise ratio compresses the null distribution to a mean angular distance of 1.02°; a 7% modulation tightens this only modestly to 0.79°, and $P \leq$ data shifts from 26.7% to 36.0%. In E-mode polarisation, the low per-patch signal-to-noise ratio produces a broad null distribution mean 76°; a 7% modulation shifts this negligibly to 75.6°, and $P \leq$ data rises from 3.0% to 3.7% (likelihood ratio 1.23). In both cases, the directional consistency test at Planck sensitivity cannot distinguish a 7% modulation from isotropy per realisation — in temperature because the null distribution is already compressed, in polarisation because noise dominates. The test is most informative in an intermediate regime where the signal measurably tightens a broad null distribution; LiteBIRD's improved per-patch polarisation sensitivity will move E-mode measurements toward this regime [20].

The $p = 3.0\%$ is calibrated against isotropic simulations and quantifies the probability of the observed directional coherence under the null hypothesis. This calibration is independent of the injection model. The observed result is therefore consistent with two interpretations: a noise fluctuation that occurs 3% of the time under isotropy, or a real signal that the directional consistency test lacks the per-realisation power to confirm at Planck sensitivity. Distinguishing between these requires either a more powerful test statistic or reduced instrumental noise. We note that Gimeno-Amo et al. (2023) [21] found that pixel-space Q/U injection at $A = 0.09$ yielded a detection rate of 62% with the LVE and an amplitude-based test statistic, compared to 4.0% for the directional consistency test at the same amplitude.

V. ROBUSTNESS TESTS

Table VIII summarises the results for variations of the fiducial pipeline configuration ($R = 15°$, $\ell = 15$–40,

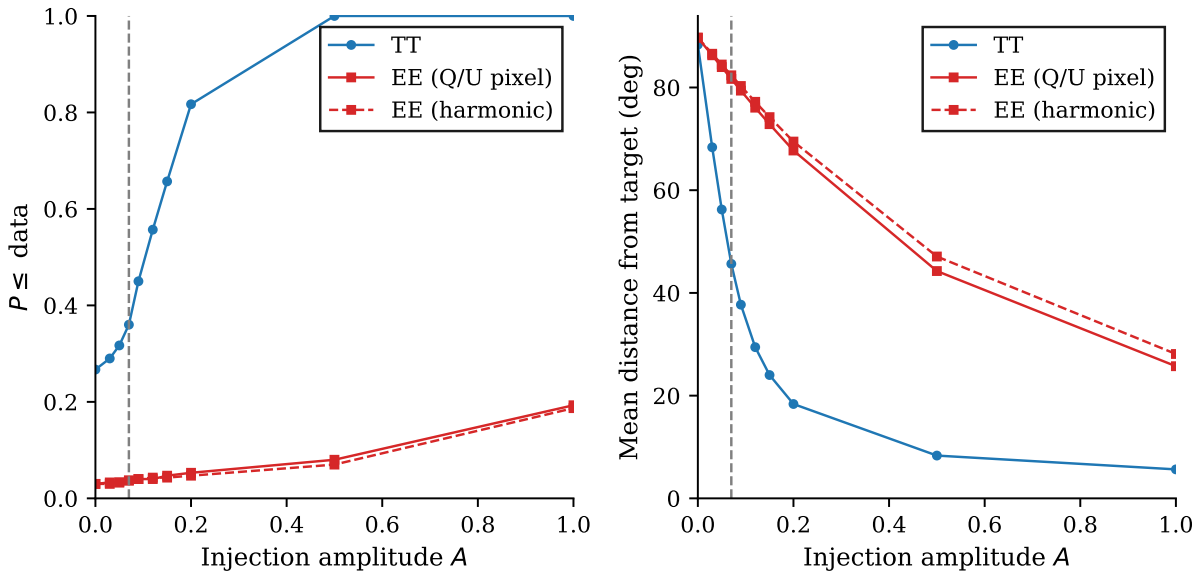

FIG. 2. Signal injection sensitivity for amplitudes $A = 0$–1. Left panel: fraction of 300 modulated realisations with half-mission angular distance at least as small as the observed data value ($P \leq$ data) as a function of injection amplitude A. Temperature (blue solid) rises steeply from 26.7% at $A = 0$ to 100% at $A = 0.50$. E-mode polarisation Q/U pixel injection (red solid) and harmonic injection (red dashed) rise slowly, overlapping at all amplitudes. The vertical dashed line marks $A = 0.07$. Right panel: mean angular distance from the recovered half-mission directions to the injection target.

TABLE VII. E-mode polarisation signal injection results. Two methods are compared: pixel-space Q/U modulation and harmonic-space E-mode modulation, both applied at $(l, b) = (45°, 60°)$ to 300 FFP10 simulations. All columns report per-simulation means. $P \leq$ data: fraction of realisations with half-mission angular distance $\leq 15.8°$ (the observed value). The two methods produce nearly identical results.

	Q/U pixel			E-mode harmonic		
A	HM→tgt	HM1↔2	$P \leq$data	HM→tgt	HM1↔2	$P \leq$data
0	89.7°	76.3°	3.0%	89.7°	76.3°	3.0%
0.07	81.7°	75.6°	3.7%	82.4°	75.6°	3.7%
0.09	79.4°	74.9°	4.0%	80.3°	75.1°	4.0%
0.50	44.3°	51.8°	8.0%	47.1°	54.2°	7.0%
1.0	25.7°	31.6°	19.3%	28.1°	33.9°	18.7%
2.0	14.1°	17.1°	53.0%	15.6°	18.6°	44.0%

no apodisation, SEVEM). For each configuration, the p-value is calibrated against FFP10 simulations processed through the identical pipeline.

Applying C2 apodisation at 1.5° and 2.0° produces negligible changes in both directions and significance (table VIII). The fiducial unapodised configuration was adopted after confirming this insensitivity.

Extending to $\ell = 15$–60 and 15–80 yields consistent directions. The significance improves slightly at $\ell = 15$–80 ($p = 2.0\%$), suggesting the signal extends beyond $\ell = 40$, though the fiducial range is chosen to match the established HPA scale.

Varying patch radius from 12° to 20° produces consistent directions with p-values between 5.0% and 7.3%. The fiducial $R = 15°$ balances angular resolution against per-patch signal-to-noise.

The component separation dependence is discussed in section IV B and table V.

The significance is calibrated against simulations processed through the identical pipeline for each configuration, ensuring that the p-value is valid independent of the specific parameter choices. Optimisation of the pipeline configuration for maximum sensitivity is deferred to future work.

VI. DISCUSSION

A. Comparison with previous analyses

The LBE and the inpainting-based approach of Gimeno-Amo et al. (2023) [21] address the same fundamental challenge — E–B leakage in local polarisation

TABLE VIII. Robustness of the E-mode directional consistency test to pipeline parameters. Each row varies one parameter from the fiducial configuration ($R = 15°$, $\ell = 15$–40, no apodisation, SEVEM). All p-values are calibrated against 300 FFP10 simulations processed through the identical pipeline configuration.

Parameter	Value	HM1 (l, b)	HM2 (l, b)	Ang. dist.	p-value
Fiducial	—	$(234.5°, -19.3°)$	$(219.2°, -13.7°)$	15.8°	3.0%
Apodisation	1.5° C2	$(235.3°, -20.7°)$	$(221.4°, -13.0°)$	15.3°	3.3%
Apodisation	2.0° C2	$(235.8°, -22.7°)$	$(224.0°, -12.2°)$	15.4°	3.3%
ℓ-range	15–60	$(201.5°, -13.6°)$	$(219.3°, -15.7°)$	17.4°	5.7%
ℓ-range	15–80	$(210.8°, -10.5°)$	$(220.1°, -14.1°)$	9.8°	2.0%
Radius	12°	$(216.5°, -23.3°)$	$(226.4°, -10.4°)$	16.0°	5.0%
Radius	20°	$(235.0°, -13.0°)$	$(210.5°, -15.2°)$	23.8°	7.3%

estimation — through independent strategies. The LBE corrects leakage through coupling matrix deconvolution: for a given patch configuration and input maps, the output is uniquely determined with no stochastic component. The inpainting approach reconstructs the masked field using constrained realisations drawn from an estimated noise covariance, yielding results that vary across realisations. The LBE provides scale-isolated bandpower at a defined multipole range with separate EE, BB, and EB spectra; the LVE provides a broadband measure of E-mode power. These are complementary rather than competing approaches, with different strengths: the amplitude test applied to the inpainted LVE has greater statistical power at Planck sensitivity, while the LBE provides deterministic reproducibility and spectral decomposition. Their directional agreement on the same data suggests that the recovered preferred direction is not an artefact of either leakage correction strategy.

Despite these methodological differences, the two approaches recover consistent directions: our HM1 direction of $(234.5°, -19.3°)$ is within 0.5° in longitude and 5° in latitude of the $(234°, -14°)$ reported by Gimeno-Amo et al.

The amplitude-based test statistic employed by Gimeno-Amo et al. has greater statistical power than the directional consistency test at Planck sensitivity, yielding $p = 0.22\%$ on PR3 data compared to our $p = 3.0\%$. This difference is expected: the injection tests in section IV D demonstrate that the consistency test is underpowered for per-realisation detection at Planck E-mode noise levels. The complementary strengths of the two approaches — statistical power for amplitude estimation, deterministic reproducibility and spectral decomposition for the LBE — motivate their joint application to future data sets with improved polarisation sensitivity.

In a subsequent analysis, Gimeno-Amo et al. (2025) [14] applied the MASTER pseudo-C_ℓ estimator in 12 independent sky regions to test for angular clustering of bandpower dipole directions in polarisation across multipole bins using a Rayleigh statistic. No significant clustering was found in E-mode polarisation. This tests a different question — coherence of directions across multipole bins — from the half-mission directional consistency we probe. The LBE extends the local pseudo-C_ℓ approach with denser spatial sampling (600 overlapping patches), per-patch spectral decomposition into EE, BB, and EB, and a different test statistic.

Ghosh & Jain (2020) [16] identified the importance of excluding multipoles below $\ell \sim 10$ for reliable direction recovery in pixel-space polarisation analyses, a consideration consistent with our choice of $\ell_{\rm min} = 15$. Their $|P|^2$ estimator mixes E and B power inseparably; the LBE separates E from B through the coupling matrix.

The gradient across component separation methods could in principle reflect residual foreground power that is preserved by some pipelines and suppressed by others, rather than a cosmological signal. However, the directions recovered by Commander and SEVEM — which employ fundamentally different foreground cleaning strategies — agree to within 20°, and both point toward the published HPA direction rather than toward known polarised foreground structures such as the Galactic dust emission concentrated along the plane. The NILC null result is consistent with this method's known excess large-scale polarisation noise rather than with foreground removal of a spurious signal. A definitive separation of cosmological signal from residual foregrounds at these angular scales will require the multi-frequency polarisation coverage of LiteBIRD [20].

B. Multiple comparisons

We note that the fiducial configuration (SEVEM, $\ell = 15$–40, $R = 15°$, no apodisation) was specified before the robustness and component separation cross-checks were performed. The robustness table (table VIII) demonstrates stability across configurations but was not used to select the fiducial. The component separation cross-checks were similarly performed after the SEVEM analysis was complete. No correction for multiple comparisons is applied, as the additional configurations serve as robustness tests of a single pre-specified analysis rather than independent searches. We also note that the choice of SEVEM as fiducial is conservative: the Commander result ($p < 0.33\%$) is substantially more significant, but

was obtained as a cross-check rather than the primary analysis.

C. Why directional consistency works where amplitude does not

To our knowledge, this is the first application of half-mission directional consistency as a probe of statistical anisotropy in the CMB. While half-mission splits are routinely used in CMB power spectrum estimation to avoid noise bias through cross-spectra, their use as a statistical test for isotropy violation — measuring whether the angular separation between independently recovered dipole directions is anomalously small — has not previously been explored.

At $\ell = 15$–40, the small number of independent modes means that an isotropic CMB realisation produces a dipolar pattern in the bandpower map with amplitude comparable to the expected HPA modulation. Comparing half-mission dipole amplitudes therefore provides little discriminating power at these scales in this framework, motivating the use of directional agreement as the test statistic. The directional consistency test circumvents this limitation. Both half-missions observe the same CMB realisation and the same HPA signal (if present); only the instrumental noise differs. If a preferred direction exists, the combined signal provides a directional anchor in both half-mission maps, resulting in tighter agreement than noise alone would produce. The test is thus sensitive to the directional coherence of the signal rather than its amplitude, accessing information that single-map amplitude tests cannot exploit.

An analogy clarifies the logic. Consider detecting a faint astronomical source in two independent exposures. The flux in each exposure is dominated by background noise, rendering a flux-excess test insensitive. However, if a source is present, both exposures will show excess flux at the same position. Positional coincidence between exposures detects the source even when its flux is individually unremarkable. The half-mission directional consistency test exploits the same principle: a shared signal anchors both half-mission dipole directions to the same sky location, producing tighter angular agreement than would arise from isotropic fluctuations alone.

D. Future applications

The LBE generalises to any spectrum, multipole range, and data split. The framework is directly applicable to LiteBIRD [20], which will provide E-mode maps with substantially lower noise than Planck at the scales relevant to HPA. Detailed forecasts for LiteBIRD sensitivity to the directional consistency test, including optimisation of the pipeline configuration, are left for future work. At LiteBIRD sensitivity, the transfer matrix — quantifying how much injected EE power appears in BB after deconvolution — becomes measurable, enabling a direct test of the leakage correction. The per-patch spectral decomposition also has potential applications beyond anisotropy: local dust spectral index mapping through bandpower ratios across frequencies, and cross-spectrum analyses between temperature and polarisation within each patch.

VII. CONCLUSIONS

We have introduced a local bandpower estimator (LBE) for measuring hemispherical power asymmetry in CMB polarisation, based on per-patch pseudo-C_ℓ estimation with coupling matrix deconvolution using NaMaster [25]. The estimator operates directly on the masked Stokes Q and U maps without inpainting, corrects for E–B leakage deterministically, and yields simultaneous deconvolved EE, BB, and EB bandpowers at a specified multipole range within each patch. The LBE is independent of the test statistic used to assess anisotropy: the deconvolved bandpower maps it produces can be analysed with any statistical test.

We have introduced half-mission directional consistency as a test statistic for spatial anisotropy, exploiting the independent noise realisations in Planck half-mission data splits to measure whether two independent measurements of the same sky agree on the preferred direction.

Applied to Planck PR3 temperature data at $\ell = 15$–40, the pipeline recovers the established HPA dipole direction with sub-degree half-mission agreement (0.48°), validating the framework on a well-established signal.

In E-mode polarisation, removing the deconvolution step causes the recovered directions to shift to the galactic plane and the half-mission agreement to become consistent with isotropy ($p = 16.7\%$), confirming that E–B leakage correction is essential at these angular scales. With deconvolution, the recovered directions are consistent with the independent analysis of Gimeno-Amo et al. (2023) [21], yielding $p = 3.0\%$ (9/300) for SEVEM and $p < 0.33\%$ (0/300) for Commander. The spectral decomposition yields EE directional coherence with null BB ($p = 31.3\%$) and EB ($p = 58.7\%$), consistent with the expectation for a scalar primordial modulation [17] and providing a qualitative check on the deconvolution.

Signal injection tests validate the pipeline as unbiased but establish that the directional consistency test is noise-limited at Planck E-mode sensitivity: a 7% modulation is indistinguishable from the null in individual realisations. These p-values quantify the rarity of the observed directional coherence under isotropy, not a per-realisation detection.

The LBE framework is directly applicable to LiteBIRD [20], which will provide the per-patch sensitivity needed to move E-mode measurements into the signal-to-noise regime where the directional consistency test has greatest discriminating power.

ACKNOWLEDGMENTS

We acknowledge the use of the Planck Legacy Archive, the HEALPix/`healpy` package, and the `NaMaster` code. This work is based on publicly available data from the Planck Legacy Archive (`https://pla.esac.esa.int`). The analysis code is available at `https://github.com/ralynch4916/LBEHPA`. Claude (Anthropic) and Cursor were used to assist with code development and manuscript preparation.

[1] Planck Collaboration, Astronomy & Astrophysics **641**, A10 (2020), arXiv:1807.06211 [astro-ph.CO].

[2] C. L. Bennett *et al.*, The Astrophysical Journal Supplement Series **208**, 20 (2013), arXiv:1212.5225 [astro-ph.CO].

[3] Planck Collaboration, Astronomy & Astrophysics **641**, A1 (2020), arXiv:1807.06205 [astro-ph.CO].

[4] D. J. Schwarz *et al.*, Classical and Quantum Gravity **33**, 184001 (2016), arXiv:1510.07929 [astro-ph.CO].

[5] Planck Collaboration, Astronomy & Astrophysics **641**, A7 (2020), arXiv:1906.02552 [astro-ph.CO].

[6] H. K. Eriksen *et al.*, The Astrophysical Journal **605**, 14 (2004), arXiv:astro-ph/0307507 [astro-ph].

[7] F. K. Hansen *et al.*, Monthly Notices of the Royal Astronomical Society **354**, 641 (2004), arXiv:astro-ph/0404206 [astro-ph].

[8] A. de Oliveira-Costa *et al.*, Physical Review D **69**, 063516 (2004), arXiv:astro-ph/0307282 [astro-ph].

[9] M. Cruz *et al.*, The Astrophysical Journal **655**, 11 (2007), arXiv:astro-ph/0603859 [astro-ph].

[10] J. Hoftuft *et al.*, The Astrophysical Journal **699**, 985 (2009), arXiv:0903.1229 [astro-ph.CO].

[11] Y. Akrami *et al.*, The Astrophysical Journal Letters **784**, L42 (2014), arXiv:1402.0870 [astro-ph.CO].

[12] Planck Collaboration, Astronomy & Astrophysics **594**, A16 (2016), arXiv:1506.07135 [astro-ph.CO].

[13] D. Hanson and A. Lewis, Physical Review D **80**, 063004 (2009), arXiv:0908.0963 [astro-ph.CO].

[14] C. Gimeno-Amo *et al.*, Journal of Cosmology and Astroparticle Physics **2025** (12), 057, arXiv:2504.05597 [astro-ph.CO].

[15] A. L. Erickcek, M. Kamionkowski, and S. M. Carroll, Physical Review D **78**, 123520 (2008), arXiv:0806.0377 [astro-ph].

[16] S. Ghosh and P. Jain, Monthly Notices of the Royal Astronomical Society **492**, 3994 (2020), arXiv:1807.02359 [astro-ph.CO].

[17] M. H. Namjoo, A. A. Abolhasani, H. Assadullahi, S. Baghram, H. Firouzjahi, and D. Wands, JCAP **2015** (05), 015, arXiv:1411.5312 [astro-ph.CO].

[18] J. Khodagholizadeh, R. Mohammadi, and S. M. S. Movahed, Eur. Phys. J. C **83**, 756 (2023), arXiv:2307.16194 [gr-qc].

[19] J. Grain, M. Tristram, and R. Stompor, Phys. Rev. D **86**, 076005 (2012), arXiv:1207.5344 [astro-ph.CO].

[20] LiteBIRD Collaboration, E. Allys, *et al.*, Progress of Theoretical and Experimental Physics **2023**, 042F01 (2023), arXiv:2202.02773 [astro-ph.IM].

[21] C. Gimeno-Amo, R. B. Barreiro, E. Martínez-González, and A. Marcos-Caballero, Journal of Cosmology and Astroparticle Physics **2023** (12), 023, arXiv:2306.14880 [astro-ph.CO].

[22] P. K. Aluri and A. Shafieloo, J. Cosmol. Astropart. Phys. arXiv:1710.00580 [astro-ph.CO].

[23] A. Lewis, A. Challinor, and N. Turok, Physical Review D **65**, 023505 (2002), arXiv:astro-ph/0106536 [astro-ph].

[24] E. Hivon *et al.*, The Astrophysical Journal **567**, 2 (2002), arXiv:astro-ph/0105302 [astro-ph].

[25] D. Alonso *et al.*, Monthly Notices of the Royal Astronomical Society **484**, 4127 (2019), arXiv:1809.09603 [astro-ph.CO].

[26] Planck Collaboration, Astronomy & Astrophysics **641**, A4 (2020), arXiv:1807.06208 [astro-ph.CO].

[27] K. M. Górski *et al.*, The Astrophysical Journal **622**, 759 (2005), arXiv:astro-ph/0409513 [astro-ph].

[28] Planck Collaboration, Astronomy & Astrophysics **594**, A12 (2016), arXiv:1509.06348 [astro-ph.CO].

[29] F. Paci, A. Gruppuso, F. Finelli, P. Cabella, A. De Rosa, N. Mandolesi, and P. Natoli, MNRAS **407**, 399 (2010), arXiv:1002.4745 [astro-ph.CO].

[30] F. Paci, A. Gruppuso, F. Finelli, A. De Rosa, N. Mandolesi, and P. Natoli, MNRAS **434**, 3071 (2013), arXiv:1301.5195 [astro-ph.CO].