

Skeptical review: Cross-Frequency Temperature Coherence of ACT DR6 Maps: Pair-Specific Diagnostics and Scale-Cut Recommendations for Multi-Frequency Analyses

Summary

This manuscript presents a model-light, pairwise cross-frequency coherence study of the six ACT DR6 AA-night temperature channels (90/150/220 GHz). Using noise-bias-free split-cross pseudo- C_ℓ estimates on a common mask ($f_{\text{sky}} \approx 0.46$), the authors compute the multipole-dependent cross-correlation coefficient ρ_ℓ for all channel pairs and use it as a practical diagnostic of inter-channel consistency. The paper emphasizes that coherence is inherently pair-specific (beams, filtering, and foreground SED mixtures differ by channel), and it delivers two practically useful products for downstream analyses: (i) recommended pair-dependent stability windows $[\ell_{\min}, \ell_{\max}]$ summarized in Table I, and (ii) a conservative per-channel beam-shape systematic envelope $\Delta b_\ell/B_\ell$ (Fig. 6, Table II) together with simple propagation formulas.

The main qualitative results are sensible and potentially impactful: same-band pairs (notably 150×150 and 90×90) are highly coherent across a broad mid- ℓ range, 90×150 retains high coherence to lower ℓ_{\max} before decohering at high ℓ (plausibly from foreground-mixture differences), and 220-related pairs decorrelate earlier and are positioned primarily as foreground-correlation diagnostics to $\ell \lesssim 1000$. The paper is clearly motivated and generally well written, but several central methodological points are currently specified only qualitatively (what spectra enter ρ_ℓ , how uncertainties are computed, how windows are chosen, how the beam envelope is constructed and should be interpreted). Clarifying and quantifying these elements (with a small set of robustness checks) would substantially strengthen reproducibility and the credibility/portability of the proposed scale-cut recommendations.

Strengths

- Clear, focused goal: provide a standalone, model-light inter-channel diagnostic for ACT DR6 temperature maps with direct practical value for power-spectrum analyses, lensing pipelines, and component separation (Sec. I, Sec. IV, Sec. VII).
- Appropriate use of split-cross spectra to avoid noise bias, and use of a common mask to ensure like-for-like sky comparison across all channels (Sec. II.1–II.2, Sec. III.1–III.2).
- Correct emphasis that coherence is pair-specific (different beams/foreground mixtures), motivating pair-dependent stability windows rather than a single universal ℓ -cut (Sec. III.4, Sec. IV.2–IV.7; Table I).
- Beam-systematics deliverable is potentially very useful to the community: a conservative beam-shape envelope plus simple propagation formulas for autos/crosses and a discussion of how beam uncertainty impacts high- ℓ usage (Sec. II.3, Sec. V.A, Sec.

VI.2; Fig. 6; Table II).

- Good positioning within ACT/Planck/SPT multi-frequency context, and a sensible qualitative interpretation of where decoherence is likely foreground-driven vs instrument-driven (Sec. I, Sec. VI.1–VI.3).
- Figures are generally well organized for cross-comparison across channel pairs, with stability-window shading that makes the intended use clear.

Major issues

1. **Ambiguity about what spectra enter ρ_ℓ (Eq. (2)) and how beam/transfer functions are handled (Sec. III.B–III.C; Eqs. (2)–(3); Figs. 1–4).** It is not explicit whether $C_\ell^{aa}, C_\ell^{bb}, C_\ell^{ab}$ used in ρ_ℓ are (i) raw pseudo-spectra, (ii) mode-coupling-corrected but beam-convolved spectra, or (iii) beam-deconvolved spectra \tilde{C}_ℓ . This choice directly affects the interpretation of high- ℓ behavior, the role of beam uncertainties, and the logic connecting ℓ_{\max} to the beam envelope.

Recommendation: In Sec. III.B/III.C, state unambiguously the full data model and processing sequence for the spectra used in Eq. (2): whether mode coupling is corrected (and how), whether pixel/window/filter transfer functions are applied, and whether nominal beams are deconvolved before forming ρ_ℓ . Use a compact notation (e.g., \hat{C}_ℓ pseudo, \bar{C}_ℓ mode-decoupled, \tilde{C}_ℓ beam-deconvolved) and reference which one is used in each figure. If ρ_ℓ uses deconvolved spectra, add a short note on numerical stability when B_ℓ becomes small and any regularization/binning choices used to control noise amplification.

2. **Uncertainty/error-bar construction for C_ℓ and especially ρ_ℓ is under-specified, limiting interpretability and reproducibility (Sec. III.1–III.2; Figs. 1–3).** The paper refers to “empirical split-cross scatter”, but does not define precisely how scatter is computed (over which cross combinations; with what weighting; per-bin or after smoothing), nor does it address that ρ_ℓ is a ratio of correlated estimators (numerator/denominator share sky, mask, and often noise/foreground fluctuations).

Recommendation: Expand Sec. III.1–III.2 to give a precise, reproducible recipe for uncertainties: (i) define the multipole binning (edges/widths and any within-bin weighting); (ii) specify how the N split-cross spectra are combined into a final binned C_b for each pair; (iii) define how the covariance of C_b is estimated (scatter across cross combinations, jackknife/bootstrap over splits, or analytic Gaussian approximation), and then how this is propagated to $\text{Var}(\rho_b)$ (either via error propagation with correlations stated, or via resampling that recomputes ρ_b per realization). Clarify whether the plotted error bars include only split-to-split noise variability or also masked-sky sample variance, and briefly justify why sample variance largely cancels (or does not) in ρ .

3. **The algorithm/criteria used to define the stability windows $[\ell_{\min}, \ell_{\max}]$ are currently heuristic and not quantitatively reproducible (Sec. III.4; Sec. IV.2–IV.7; Table I).** Since Table I is a key deliverable, readers need to know what objective thresholds (S/N, beam-systematic tolerance, coherence threshold, slope change, etc.) were applied and how sensitive the windows are to reasonable variations of those thresholds.

Recommendation: In Sec. III.4, formalize the stability-window selection with explicit quantitative criteria. For example, define ℓ_{\min} via a minimum S/N on the relevant spectra (or on ρ), and define ℓ_{\max} using one or more thresholds such as: (a) ρ dropping below a stated value (e.g., 0.99 for same-band, 0.98 for cross-band) or deviating from unity by $> X\sigma$; (b) beam-envelope-propagated fractional uncertainty exceeding Y%; and/or (c) a foreground-mixture indicator exceeding a tolerance. Then, in Sec. IV (or an appendix), include a brief sensitivity test showing how ℓ_{\max} shifts when thresholds change (even ± 100 – 200 in ℓ is informative). Annotate Figs. 1–4 with vertical lines and numeric labels for ℓ_{\min}, ℓ_{\max} per pair to make Table I directly traceable.

4. **Beam-shape envelope construction and its statistical meaning are not sufficiently specified, despite being central to the paper’s recommended ℓ -cuts and systematics guidance (Sec. II.3; Sec. V.A; Fig. 6; Table II).** The envelope is described as a maximum variation across beam splits, but the exact inputs (which splits), normalization, smoothing/regularization, and whether it is intended as an upper bound vs a 1σ -like uncertainty are unclear.

Recommendation: In Sec. II.3 and Sec. V.A, explicitly document: (i) which beam-split products are used (elevation/PWV/time/detector subsets; ideally with DR6 product names/paths); (ii) how each split beam is normalized before comparison (e.g., $B_{\ell=0} = 1$, fixed solid angle, or other); (iii) whether the per- ℓ maxima are smoothed or percentile-based to avoid noise spikes at high ℓ ; and (iv) whether the envelope includes only shape differences or also amplitude/normalization (calibration-like) differences. In Sec. VI.2 add a short “how to use” paragraph: what downstream users should assume about correlation of beam-shape errors across arrays/bands, and whether the envelope is best treated as a conservative bound or an effective uncertainty.

5. **Beam-to-coherence propagation (Eq. (6) and surrounding discussion) is currently not well justified and may be misleading without an explicit beam-error model (Sec. VI.2; Eq. (6)).** For purely multiplicative, map-level beam transfer-function errors applied consistently to auto- and cross-spectra, ρ_{ℓ} cancels beam factors to first order; a residual impact generally requires split-dependent effective beams, non-factorizable beam systematics, or mismatched deconvolution/transfer functions across maps.

Recommendation: State explicitly the beam systematics model under which Eq. (6) is intended (e.g., different effective beams per split/map; beam errors that do not factorize as a single B_{ℓ} ; or anisotropic/elliptical effects that couple to masking/filtering).

Provide a short derivation or a conservative bound with clearly stated assumptions. If Eq. (6) is meant as a heuristic upper limit rather than a first-principles result, label it as such and explain when beam effects cancel in ρ_ℓ to leading order.

6. **Pseudo- C_ℓ pipeline details and validation of the flat-sky approximation on a large sky fraction are insufficiently documented/quantified (Sec. III.1; Sec. VI.C).** With $f_{\text{sky}} \approx 0.46$ and $\ell_{\text{min}} \sim 400\text{--}500$, curvature/mode-coupling/pixelization choices can matter for low- ℓ behavior and thus for window boundaries and coherence trends.

Recommendation: In Sec. III.1, specify the pseudo- C_ℓ implementation details needed to reproduce results (binning, apodization beyond “10’ cosine taper”, treatment of $M_{\ell\ell}$ for cross-spectra, pixel window, filtering/transfer functions, and map projection/patching if any). In Sec. VI.C, add at least one quantitative validation: e.g., compare ρ_ℓ for a representative pair (such as pa4_f150×pa5_f150) between the current flat-sky pipeline and a curved-sky pseudo- C_ℓ estimator over $400 \leq \ell \leq 1500$, or use simulations/literature bounds to show the induced bias is negligible relative to quoted deviations.

7. **Foreground-vs-instrument attribution of decoherence (especially for 90×150 high- ℓ and all 220-related pairs) remains largely qualitative; calibration/bandpass mismatch is also not explicitly bounded (Sec. IV.2–IV.6; Sec. VI.3).** The paper’s “bigger picture” conclusions (why ℓ_{max} differs by pair, and why 220 should be used mainly as a foreground diagnostic) would be more convincing with minimal quantitative cross-checks.

Recommendation: Add lightweight quantitative diagnostics to separate contributions to $|1 - \rho_\ell|$: (i) a simple two-/three-component toy model (CMB + dust/CIB + tSZ/PS) showing expected $\rho_\ell(\ell)$ trends for 90×150 and 150×220 given plausible component fractions (anchored to Planck/SPT literature already cited); (ii) a mask-dependence test (e.g., slightly more/less aggressive Galactic/dust cut) demonstrating that 220-pair decoherence responds as expected if dust-driven; (iii) a null test using split-difference maps to bound noise/analysis-induced decoherence; and (iv) a short statement (with numbers, if available from DR6) of relative calibration uncertainty and how a small calibration/bandpass mismatch would manifest in ρ_ℓ and in the deconvolved spectral ratios. These do not need a full parametric fit, but should quantitatively support the dominant-effect claims used to motivate Table I.

Minor issues

1. Mask construction and data-product usage are not described with enough implementation detail to reproduce the common mask and to assess sensitivity of ρ_ℓ to footprint choices (Sec. II.2; Sec. III.1). The common mask is constrained by pa4_f220; it is unclear how much this choice, rather than physics, affects coherence trends and window recommendations for 90/150-only analyses.

Recommendation: In Sec. II.2/Sec. III.1, document: map resolution/projection, exact inverse-variance/cross-linking thresholds, point-source and Galactic masking, and handling of small holes/isolated regions; provide a pointer (URL/DOI/repo) to the exact mask file or a script to regenerate it. Add one robustness check (appendix is fine): repeat a subset of ρ_ℓ curves using a 90/150-only common mask (larger f_{sky}) to isolate footprint-driven effects, and comment on whether Table I windows change.

2. Limited quantitative summary statistics for within-band consistency claims (e.g., “150 GHz ratios within $\sim 10\%$ ”) (Sec. IV.3). Without a metric, it is hard to compare to other experiments or to judge whether observed ratio scatter is dominated by noise amplification, transfer-function mismatch, or residual systematics.

Recommendation: In Sec. IV.3, add a compact quantitative summary over the stated window (e.g., mean \pm std of $\tilde{C}_\ell^X/\tilde{C}_\ell^{\text{ref}} - 1$, max deviation, and/or reduced χ^2 to a constant ratio). If the scatter is indeed $\sim 10\%$ while the beam envelope is $\sim 0.2\%$ at $\ell = 2000$ (Table II), briefly explain what dominates the ratio scatter (noise deconvolution, filtering/transfer functions, residual mode-coupling, etc.).

3. Figures do not consistently make stability-window boundaries and uncertainties visually legible; some panels claim completeness (e.g., “full 6×6 matrix”) but omit pairs, and error bars are often invisible (Figs. 1–6).

Recommendation: For Figs. 1–6: (i) ensure the set of pairs shown matches the caption (or explicitly state that only a subset is plotted); (ii) add vertical lines with numeric ℓ_{min} and ℓ_{max} labels matching Table I; (iii) increase error-bar visibility or use shaded uncertainty bands; (iv) ensure legends/linestyles unambiguously map to channels; and (v) use colorblind-safe palettes and/or distinct linestyles.

4. Domain issues for ρ_ℓ : the definition in Eq. (2) requires positive C_ℓ^{aa}, C_ℓ^{bb} , but split-cross estimates (especially at low ℓ or noise-dominated bins) can be negative; the manuscript does not state how such bins are handled (Sec. III.2).

Recommendation: Add a short clarification in Sec. III.2: e.g., exclude bins where either auto estimate is ≤ 0 , or form ρ from averaged/regularized autos, or impose an S/N cut before computing ρ . State how many bins are affected and whether this influences ℓ_{min} .

5. The “model-light” claim could be misread as model-independent inference, while the interpretation of decoherence and the chosen windows rely on standard foreground expectations and DR6 beam products (Sec. I; Sec. III.2; Sec. VI.1).

Recommendation: Tighten wording in Sec. I/Sec. III.2: explicitly state that no parametric cosmology/foreground fit is performed, but that physical interpretation of ρ_ℓ trends and the window recommendations are informed by standard foreground SED behavior and beam/transfer-function knowledge.

6. Calibration and bandpass mismatch are not explicitly discussed, though they can affect cross-frequency comparisons at the percent level (Sec. II.1; Sec. IV; Sec. VI.3).

Recommendation: Add a brief note stating whether DR6 maps are used with native DR6 calibration in thermodynamic temperature units and whether any additional relative calibration between arrays/bands was performed for this analysis. Provide (or cite) an estimated relative calibration uncertainty and a sentence on how such mismatches would propagate into the plotted ratios and ρ .

7. Table II reports beam-envelope values at $\ell = 2000$, while the recommended ℓ_{\max} values are typically 1000–1500 (Table I), making it harder for readers to connect the numbers to the adopted cuts.

Recommendation: Augment Table II (or add a short companion table) with $\Delta b_\ell/B_\ell$ at $\ell = 1000$ and $\ell = 1500$ (or at each pair’s ℓ_{\max}), and/or provide a machine-readable file of $\Delta b_\ell/B_\ell(\ell)$ referenced in the text.

8. Opportunity for broader context is underused: the manuscript could more concretely compare its recommended cuts/coherence ranges to ACT DR6 likelihood choices and to comparable Planck/SPT cross-frequency coherence behavior (Sec. VI.B).

Recommendation: In Sec. VI.B, add a short quantitative comparison: note whether Table I ℓ -cuts align with (or differ from) official DR6 power-spectrum/likelihood scale cuts, and cite one or two concrete Planck/SPT cross-frequency coherence or spectral-ratio benchmarks to place ACT DR6 results in context.

Very minor issues

1. Notation for the beam transfer function and its uncertainty is inconsistent (b_ℓ vs B_ℓ ; $\Delta b_\ell/B_\ell$ vs $\Delta B_\ell/B_\ell$) (Sec. II.3; Sec. V.A; Eqs. (4)–(6); Table II; Fig. 6).

Recommendation: Unify notation throughout (or explicitly define b_ℓ vs B_ℓ if both are intended) and ensure Eqs. (4)–(6), Fig. 6, and Table II use the same symbols.

2. Hat/tilde notation around Eq. (3) is only implicitly defined (Sec. III.C). Readers must infer which spectra are pseudo-/measured vs deconvolved.

Recommendation: Add one sentence near Eq. (3) defining the hat and tilde quantities (and whether mode-coupling correction is included before/after deconvolution).

3. A few captions/axis labels omit explicit definitions/units (e.g., ρ_ℓ as unitless) and some text relies on color-only descriptions that may not translate to grayscale or be accessible (Figs. 1–6).

Recommendation: Ensure each axis/caption states the observable (auto/cross, binned vs unbinned, deconvolved vs convolved), include units where applicable, and avoid color-only identification by adding linestyles/markers or direct labels.

4. Minor style consistency: section/caption formatting and citation punctuation vary (Sec. II–VI).

Recommendation: Normalize headings, figure/table caption style, and citation punctuation to the target journal style, and fix any lingering typographic artifacts in equations.

Key statements and references

- • **Ground-based experiments including the Atacama Cosmology Telescope (ACT) and the South Pole Telescope (SPT) have demonstrated over the past three decades that robust cosmological inference from CMB data requires quantitative characterization of inter-channel consistency as a function of angular scale, building on earlier space-based measurements from WMAP and Planck that established the importance of multi-frequency observations for separating primordial CMB from astrophysical foregrounds and controlling instrumental systematics.**
 - *Reference(s):* [4], [5], [10]
- • **At 220 GHz, thermal dust emission dominates the foreground budget over a broad multipole range, and the steeply rising dust spectral energy distribution produces a much larger foreground-signal ratio than at lower frequencies, leading to earlier foreground-driven decoherence in cross-frequency coherence measurements with 90 and 150 GHz channels, consistent with Planck’s diffuse component-separation and high-frequency instrument analyses.**
 - *Reference(s):* [3], [5], [6]
- • **The pa4_f220 beam-shape systematic envelope is approximately three times larger than the 90 and 150 GHz envelopes, reflecting greater difficulty in beam characterization at higher frequencies where the beam solid angle is smaller and atmospheric contamination more pronounced, as established in prior ACT instrument and DR6 map-characterization work.**
 - *Reference(s):* [7], [9], [12]

Mathematical consistency audit

This section audits **symbolic/analytic** mathematical consistency (algebra, derivations, dimensional/unit checks, definition consistency).

Maths relevance: light

The paper contains a small number of central analytic definitions and propagation approximations: the coherence coefficient ρ_ℓ , split-cross spectrum averaging, harmonic-space beam deconvolution, and first-order beam-error propagation formulas. Most other content is interpretive

and procedural rather than derivation-heavy.

Checked items

1. ✓ Split-cross combinatorics (auto and cross) (Sec. III.A, p.2)

- **Claim:** With four splits, averaging the six $i < j$ within-channel cross-spectra yields a noise-bias-free auto-spectrum estimate; for a cross-channel pair (a,b), averaging twelve cross-split spectra with $i \neq j$ avoids correlated noise bias.
- **Checks:** algebra/combinatorics, definition consistency
- **Verdict:** PASS; confidence: high; impact: moderate
- **Assumptions/inputs:** Each channel has exactly 4 independent time-domain splits., Within-channel: distinct split pairs $i < j$ are used., Cross-channel: all 4×4 split pairs exist and the $i = j$ subset is excluded.
- **Notes:** Counts are consistent: $C(4, 2) = 6$ within-channel cross pairs. Cross-channel pairs: $4 \times 4 = 16$ total minus 4 with $i = j$ gives 12. The claim about avoiding correlated noise bias depends on how splits are constructed, but the counting/math is correct.

2. ✓ Sky fraction definition (Eq. (1), Sec. II.B, p.2)

- **Claim:** The common-mask sky fraction is $f_{\text{sky}} \approx 0.4604$.
- **Checks:** dimensional/units, notation consistency
- **Verdict:** PASS; confidence: high; impact: minor
- **Assumptions/inputs:** f_{sky} denotes the masked sky fraction, dimensionless.
- **Notes:** No algebraic structure to verify beyond dimensional consistency; f_{sky} is properly treated as dimensionless throughout.

3. ✓ Cross-correlation coefficient definition (Eq. (2), Sec. III.B, p.2)

- **Claim:** ρ_ℓ is defined as $C_\ell^{ab} / \sqrt{C_\ell^{aa} C_\ell^{bb}}$.
- **Checks:** algebra, dimensional/units, symbol/definition consistency
- **Verdict:** PASS; confidence: high; impact: critical
- **Assumptions/inputs:** C_ℓ^{aa} and C_ℓ^{bb} are (split-cross) auto-spectra for channels a and b ., C_ℓ^{ab} is the (split-cross) cross-spectrum between channels a and b ., Auto-spectra are positive in the range where ρ_ℓ is evaluated.
- **Notes:** The expression is dimensionless and symmetric under $a \leftrightarrow b$. The definition implicitly requires $C_\ell^{aa}, C_\ell^{bb} > 0$; the handling of possible negative estimates is not specified (flagged separately as a minor issue).

4. ✓ Claim that $\rho_\ell = 1$ for identical sky signals (Immediately after Eq. (2), Sec. III.B, p.2)

- **Claim:** $\rho_\ell = 1$ when both channels observe identical sky signals convolved with known beams and free of noise bias.
- **Checks:** sanity/limiting case, algebra

- **Verdict:** PASS; confidence: high; impact: moderate
- **Assumptions/inputs:** Observed spectra factorize with beams: $C^{\{ab, \text{rm obs}\}} \ell = B_a(\ell) B_b(\ell) C^{\{\text{rm sky}\}} \ell, C^{\{aa, \text{rm obs}\}} \ell = B_a^2(\ell) C^{\{\text{rm sky}\}} \ell$, etc., No additive contaminants/noise bias in the cross estimates.
- **Notes:** Under multiplicative, isotropic beam transfer functions and identical underlying sky, the beam factors cancel exactly in the ratio, giving $\rho_\ell = 1$.

5. ✓ **Beam deconvolution formula** (Eq. (3), Sec. III.C, p.2)

- **Claim:** Beam deconvolution is performed via $\tilde{C}^{\{ab\}} \ell = \hat{C}^{\{ab\}} \ell / (B_a \ell B_b \ell)$.
- **Checks:** algebra, dimensional/units, definition consistency
- **Verdict:** PASS; confidence: high; impact: moderate
- **Assumptions/inputs:** Measured (beam-convolved) cross-spectrum satisfies $\hat{C}^{\{ab\}} \ell \approx B_a \ell B_b \ell C_{\ell}^{\{\text{true}\}}$ (*ignoring mode – coupling and binning details*), $B_a \ell, B_b \ell$ are the beam transfer functions for channels a and b .
- **Notes:** The deconvolution factor $B^a B^b$ is correct for cross-spectra; for auto-spectra it reduces to division by $(B^a)^2$.

6. ✓ **Auto-spectrum beam-error propagation (factor of two)** (Eq. (4), Sec. III.C, p.2)

- **Claim:** For an auto-spectrum, fractional beam-induced uncertainty satisfies $\Delta C_\ell^{aa} / C_\ell^{aa} \approx \pm 2 \Delta B_\ell^a / B_\ell^a$.
- **Checks:** algebra (differential/propagation), approximation check
- **Verdict:** PASS; confidence: high; impact: moderate
- **Assumptions/inputs:** C^{aa} refers to a beam-deconvolved auto-spectrum estimate, scaling as $C^{aa, \text{deconv}} \propto (B_\ell^a)^{-2}$. $\Delta B^a / B^a$ is small (linear approximation).
- **Notes:** If $C \propto B^{-2}$, then $d \ln C = -2 d \ln B$, giving $|\Delta C / C| \approx 2 |\Delta B / B|$. The \pm sign is consistent with an envelope/unknown sign systematic.

7. ✓ **Cross-spectrum beam-error propagation (sum of fractional terms)** (Eq. (5), Sec. V.B, p.5)

- **Claim:** For a cross-spectrum, $\Delta C_\ell^{ab} / C_\ell^{ab} \approx \pm (\Delta B_\ell^a / B_\ell^a + \Delta B_\ell^b / B_\ell^b)$.
- **Checks:** algebra (differential/propagation), approximation check
- **Verdict:** PASS; confidence: medium; impact: moderate
- **Assumptions/inputs:** C^{ab} refers to a beam-deconvolved cross-spectrum estimate scaling as $C^{ab, \text{deconv}} \propto (B_\ell^a B_\ell^b)^{-1}$. $\Delta B^a / B^a$ and $\Delta B^b / B^b$ are small (linear approximation). The use of a linear sum (not quadrature) is intended as a conservative bound/envelope, not a statistical error propagation.

- **Notes:** For $C^{ab} \propto (B^a B^b)^{-1}$, $d \ln C^{ab} = -(d \ln B^a + d \ln B^b)$, giving magnitude $|\Delta C^{ab}/C^{ab}| \approx |\Delta B^a/B^a + \Delta B^b/B^b|$. The paper presents it as $\approx \pm$ (sum), consistent with an envelope interpretation.
8. \triangle **Beam-error impact on coherence and bound on $|\Delta \rho_\ell|$** (Eq. (6) and preceding paragraph, Sec. V.B, p.5)
- **Claim:** If beam errors are perfectly correlated between arrays, beam-induced uncertainty on ρ_ℓ cancels; with partial correlation, the residual is bounded by $|\Delta \rho_\ell| \lesssim \Delta B_\ell^a/B_\ell^a + \Delta B_\ell^b/B_\ell^b$.
 - **Checks:** algebraic propagation, assumption sufficiency check
 - **Verdict:** UNCERTAIN; confidence: medium; impact: critical
 - **Assumptions/inputs:** ρ_ℓ is computed from spectra (either convolved or deconvolved) in a manner sensitive to beam systematics., Beam systematics can be represented by per-channel fractional errors $\Delta B/B$ and may be partially correlated across arrays.
 - **Notes:** Given $\rho = C_{ab}/\sqrt{C_{aa}C_{bb}}$, purely multiplicative beam transfer-function errors that enter C_{ab} as $B_a B_b$ and C_{aa}, C_{bb} as B_a^2, B_b^2 cancel in ρ to first order even if δ_a and δ_b are uncorrelated, leaving only higher-order residuals. To justify an $O(\Delta B/B)$ bound like Eq. (6), the paper needs an explicit model where beam systematics affect auto- and cross-spectra differently (e.g., split-dependent effective beams, non-factorizable beam errors, or anisotropic beam effects). That model/derivation is not provided.
9. \triangle **Consistency of beam-uncertainty notation (Δb vs ΔB)** (Sec. II.C ($\Delta b_\ell/B_\ell$), Sec. V.A, Eqs. (4)–(6), Table II, pp.2–5)
- **Claim:** The beam-shape systematic envelope is $\Delta b_\ell/B_\ell$ and is used in propagation expressions written with $\Delta B_\ell/B_\ell$.
 - **Checks:** notation/definition consistency
 - **Verdict:** UNCERTAIN; confidence: high; impact: minor
 - **Assumptions/inputs:** b_ℓ and B_ℓ refer to the same beam transfer function quantity.
 - **Notes:** The paper does not define b_ℓ separately from B_ℓ , so the mixed notation is most likely a typographical inconsistency. If b_ℓ is intended to mean a perturbation (e.g., beam variation mode) distinct from B_ℓ , then Eqs. (4)–(6) should be rewritten accordingly.

Limitations

- Audit is restricted to the provided PDF text/images; no external beam/foreground models or pipeline details were consulted.
- The paper omits derivations connecting the beam envelope to coherence deviations; items depending on that connection are marked UNCERTAIN rather than inferred.

- No validation of numerical values, plots, or empirical scale-cut choices was performed (symbolic/analytic audit only).

Numerical results audit

This section audits **numerical/empirical** consistency: reported metrics, experimental design, baseline comparisons, statistical evidence, leakage risks, and reproducibility.

No candidate numeric relationships were computationally verified in this run because the execution produced an error and returned no check results.

Checked items

1. \triangle **C1_fsky_to_area_deg2** (Page 3, Sec. IV.A (Common-mask sky fraction))
 - **Claim:** The six-channel common mask yields $f_{\text{sky}} = 0.4604$, corresponding to approximately 19,000 deg².
 - **Checks:** unit-consistent recomputation (fraction of full sky area)
 - **Verdict:** UNCERTAIN
 - **Notes:** Intended check: $\text{area_deg2} \approx f_{\text{sky}} \times \text{full_sky_area_deg2}$; not computed in this run.
2. \triangle **C2_fsky_consistency_across_pages** (Page 2, Eq. (1) and Page 3, Sec. IV.A)
 - **Claim:** Sky fraction is reported as $f_{\text{sky}} \approx 0.4604$ (Eq. 1) and later as $f_{\text{sky}} = 0.4604$.
 - **Checks:** repeated constant equality
 - **Verdict:** UNCERTAIN
 - **Notes:** Intended check: parse both as floats and assert equality; not computed in this run.
3. \triangle **C3_channel_count_by_band** (Page 2, Sec. II.A (Channel selection))
 - **Claim:** Six channels total: two at 90 GHz, three at 150 GHz, and one at 220 GHz.
 - **Checks:** parts-to-total count consistency
 - **Verdict:** UNCERTAIN
 - **Notes:** Intended check: $2 + 3 + 1 = 6$; not computed in this run.
4. \triangle **C4_four_splits_cross_spectra_counts** (Page 2, Sec. III.A (Split-cross power spectra))
 - **Claim:** With four time-domain splits: (i) six independent within-channel cross-spectra C_ℓ^{ij} ($i < j$), and (ii) twelve cross-channel cross-split spectra $C_\ell^{ai \times bj}$ ($i \neq j$).
 - **Checks:** combinatorial count verification
 - **Verdict:** UNCERTAIN

- **Notes:** Intended check: $C(4, 2) = 6$ and $4 \times (4 - 1) = 12$; not computed in this run.
5. \triangle **C5_150GHz_effective_frequency_span** (Page 2, Sec. II.A and Table II (page 5))
- **Claim:** The three 150 GHz channels span 145.38–147.04 GHz (Table II).
 - **Checks:** min/max range recomputation from tabulated values
 - **Verdict:** UNCERTAIN
 - **Notes:** Intended check: min/max of [145.53, 147.04, 145.38] equals reported endpoints; not computed in this run.
6. \triangle **C6_overall_effective_frequency_range** (Page 2, Sec. II.A and Table II (page 5))
- **Claim:** Effective CMB frequencies range from 93.42 GHz (pa6 f090) to 219.6 GHz (pa4 f220).
 - **Checks:** min/max range recomputation from table
 - **Verdict:** UNCERTAIN
 - **Notes:** Intended check: min/max across the Table II ν_{eff} list equals 93.42 and 219.6; not computed in this run.
7. \triangle **C7_90GHz_effective_frequency_difference** (Page 3, Sec. IV.D and Table II (page 5))
- **Claim:** The effective-frequency difference between the two 90 GHz arrays (95.00 vs. 93.42 GHz) is small.
 - **Checks:** difference recomputation
 - **Verdict:** UNCERTAIN
 - **Notes:** Intended check: $\Delta\nu = 95.00 - 93.42 = 1.58$ GHz; not computed in this run.
8. \triangle **C8_fwhm_range_claim_vs_table** (Page 2, Sec. II.C and Table II (page 5))
- **Claim:** FWHM range from $\approx 2.1'$ at 90 GHz to $\approx 1.0'$ at 220 GHz (Table II).
 - **Checks:** approximate range consistency with tabulated endpoints
 - **Verdict:** UNCERTAIN
 - **Notes:** Intended check: compare claimed values to Table II endpoints (90 GHz near $2.10'$, 220 GHz $0.98'$); not computed in this run.
9. \triangle **C9_pa4f220_envelope_three_times_larger** (Page 5, Sec. V.C and Table II (page 5))
- **Claim:** The pa4 f220 beam-shape systematic envelope is approximately three times larger than the 90 and 150 GHz envelopes (Table II).
 - **Checks:** ratio comparison from tabulated envelope values at $\ell = 2000$

- **Verdict:** UNCERTAIN
 - **Notes:** Intended check: ratio = $0.0060/0.0020 \approx 3$; not computed in this run.
10. \triangle **C10_tableII_identical_envelopes_for_5_channels** (Table II (page 5))
- **Claim:** Table II lists $\Delta b_\ell/B_\ell|_{\ell=2000} = 0.0020$ for five channels (two 90 GHz + three 150 GHz) and 0.0060 for the 220 GHz channel.
 - **Checks:** table internal consistency / repeated constants
 - **Verdict:** UNCERTAIN
 - **Notes:** Intended check: envelope values set $0.0020, 0.0060$ with counts $5, 1$; not computed in this run.
11. \triangle **C11_eq4_factor_of_two_logic** (Page 2, Sec. III.C, Eq. (4))
- **Claim:** For an auto-spectrum, the fractional beam-induced uncertainty is $\Delta C_\ell^{aa}/C_\ell^{aa} \approx \pm 2 \Delta B_\ell^a/B_\ell^a$, where the factor of two reflects $C_\ell \propto (B_\ell)^{-2}$.
 - **Checks:** symbolic-numeric sanity check with small perturbation
 - **Verdict:** UNCERTAIN
 - **Notes:** Intended check: log-derivative/differential consistency for $y = B^{-2}$; not computed in this run.

Limitations

- Only the provided PDF text was used; no external ACT DR6 data products or numerical arrays were available for verification.
- Figures contain qualitative/graphical claims (ρ_ℓ curves, ratio points, envelopes vs ℓ) but do not provide machine-readable numbers; plot-pixel extraction is out of scope.
- Most quantitative results are inequalities over multipole ranges (e.g., ρ_ℓ thresholds) that cannot be recomputed without the underlying spectra.
- Execution error prevented running automated checks: Sandbox policy violation: import of 'ast' is not allowed.