

Skeptical review: Deprojection-Response Diagnostics for ACT DR6 –times NILC Cross-Spectra: Beam-Amplification Systematics and Scale-Cut Recommendations

Summary

The manuscript presents an end-to-end validation of ACT DR6 single-frequency temperature (TT) cross-spectra with ACT+Planck NILC CMB temperature maps, comparing a standard NILC solution to a tSZ-deprojected NILC variant. The central diagnostic is the deprojection-response ratio $R_b(\ell)$ between cross-spectra formed with NILC(deproj) and NILC(std), supported by split-cross pseudo- C_ℓ estimation, Monte Carlo-calibrated transfer functions, explicit beam deconvolution, a beam-amplification proxy $A_b(\ell)$, split-difference null tests, and propagation of beam-envelope uncertainties (Secs. 2–4). Over $200 \leq \ell \leq 1500$, five of six ACT channels yield inverse-variance-weighted mean ratios consistent with unity at the sub-percent level with acceptable χ^2/dof , while pa4_f220 shows a mild $\approx 1\sigma$ mean excess and a pronounced rise at high ℓ that the paper attributes primarily to beam/deconvolution amplification (Secs. 4.1–4.5). The work is timely and practically useful for ACT DR6 cross-spectrum analyses, but several key elements need clarification or correction for internal consistency and reproducibility—especially what cancels (and what does not) in $R_b(\ell)$, how transfer functions and beam uncertainties are treated in the ratio, and stronger quantitative support for attributing the pa4_f220 behavior to beam-related effects rather than NILC-configuration/foreground differences (Secs. 3–5).

Strengths

- Clear, well-motivated validation question with direct relevance to ACT DR6 cross-spectrum pipelines and the practical choice of NILC configuration (Sec. 1, Sec. 5.1).
- Cohesive methodological pipeline combining split-cross pseudo- C_ℓ estimation, MC-calibrated transfer functions, explicit beam treatment, and multiple complementary diagnostics (Secs. 3.1–3.6).
- Use of both scale-dependent ($R_b(\ell)$) and summary (\bar{R}_b , χ^2/PTE) statistics, together with null tests, gives a multi-angle view of potential systematics (Secs. 4.1–4.4).
- The identification of pa4_f220 as the only problematic channel at high ℓ , and the resulting conservative scale-cut guidance, is operationally valuable for downstream analyses (Secs. 4.2, 5.2, 5.4).
- Generally clear presentation with helpful figures and a discussion section that attempts to situate results relative to Planck/NILC and known limitations (Secs. 5–6).

Major issues

1. **Internal algebra/logic of what cancels in the ratio $R_b(\ell)$ is not made explicit, leading to inconsistent interpretation of beam-driven effects (Secs. 3.3, 4.2, 5.2).** As defined (Eqs. (1) and (4)), any multiplicative factor applied identically to both numerator and denominator (e.g., the ACT beam deconvolution $1/b_\ell^{(b)}$, or a common transfer function) should largely cancel in $R_b(\ell)$. However, the text attributes the pa4_f220 high- ℓ excursion to ACT beam-deconvolution amplification using $A_b(\ell)$, without clearly specifying which non-cancelling terms differ between NILC(std) and NILC(deproj) (e.g., different NILC effective beams, different transfer functions, different beam handling in the estimator).

Recommendation: Add an explicit expression for $R_b(\ell)$ obtained by substituting Eq. (1) into Eq. (4), and show term-by-term which factors cancel and which remain. Then align the interpretation in Secs. 4.2 and 5.2 with that expression: if the driver is a difference in $b_\ell^{(\text{NILC,std})}$ vs $b_\ell^{(\text{NILC,deproj})}$, or configuration-dependent F_ℓ , state this clearly and quantify its expected size. If instead some steps apply different effective deconvolution/transfer corrections between the two configurations, document exactly where that asymmetry enters the pipeline.

2. **Beam-uncertainty propagation for $R_b(\ell)$ appears incorrect/incomplete as written (Sec. 3.6, Eq. (7)) and is in tension with the cancellation structure of $R_b(\ell)$.** Eq. (7) propagates ACT channel beam uncertainty into the ratio with a $\sqrt{2}$ factor, but the ACT beam is common to both spectra entering $R_b(\ell)$ and should be highly correlated (and potentially cancel), while uncertainty in the NILC effective beam—more likely to differ between std/deproj—does not appear explicitly.

Recommendation: Re-derive $\sigma_{\text{beam}}[R_b(\ell)]$ from the explicit expanded form of $R_b(\ell)$, including correlated-error assumptions. If ACT beam factors are common, treat their uncertainty as (nearly) fully correlated and show the residual sensitivity (if any). If NILC effective beams differ between std and deproj, propagate $\delta b_\ell^{(\text{NILC,std})}$ and $\delta b_\ell^{(\text{NILC,deproj})}$ (and their covariance) explicitly. Justify or remove the $\sqrt{2}$ factor by stating the assumed correlation structure. Update Sec. 4.5's error-budget discussion accordingly.

3. **Transfer-function definition and application are not sufficiently specified to verify that corrections are applied exactly once and consistently across configurations (Sec. 3.2; Eqs. (1)–(2)).** Eq. (2) states the MC output spectrum includes “beam convolution,” while Eq. (1) separately divides by beam transfer functions; it is therefore unclear whether beams (and/or pixel windows) are absorbed into F_ℓ or deconvolved separately, and whether F_ℓ is calibrated per channel and per NILC configuration (std vs deproj) or shared—an important point because shared corrections would largely cancel in $R_b(\ell)$.

Recommendation: Rewrite Sec. 3.2 to define precisely: (i) whether F_ℓ is computed from beam-convolved maps, beam-deconvolved maps, or maps smoothed to a common target beam; (ii) whether pixel window functions are included in F_ℓ or in the beam terms; (iii) whether $F_\ell \equiv F_\ell^{(b \times X)}$ depends on ACT channel b and NILC configuration $X \in \text{std, deproj}$. Ensure the estimator in Eq. (1) and the cali-

bration in Eq. (2) are consistent so that each correction (mask/filters/beams/pixels) is applied once. If a common F_ℓ is used for both NILC configurations, state this explicitly and explain why configuration dependence is negligible.

4. Simulation suite description is insufficient for reproducibility and to support sub-percent claims (Sec. 3.2, Sec. 5.5).

The paper states $N_{\text{MC}} = 480$ signal-only simulations but does not fully specify the input cosmology/ C_ℓ , whether any noise/foreground components are included, how channel/NILC-specific beams and filtering are implemented, and whether the same simulation realizations are used for both NILC configurations (which affects covariance and cancellation in the ratio). In addition, the reliance on Gaussian, CMB-only simulations leaves unquantified the potential impact of non-Gaussian foregrounds on pseudo- C_ℓ mode coupling and transfer functions at the $\leq 1\%$ level being tested.

Recommendation: Expand Sec. 3.2 (and/or add an appendix) to document: fiducial cosmology and CMB spectrum; map-level processing (beams, filtering, reprojection, masks) applied to each simulated product; whether F_ℓ is calibrated separately for std/deproj; and whether the same realizations are shared between the two pipelines. Add at least one quantitative bound on non-Gaussian foreground impact: e.g., a small ensemble with foregrounds (dust/CIB/point sources; optionally tSZ) passed through the same pipeline and the resulting shift in F_ℓ and/or $R_b(\ell)$, or a literature-supported validation argument tied specifically to the ℓ -range 200–1500 used for conclusions.

5. The pa4_f220 high- ℓ excursion is plausibly beam/deconvolution related but is not yet demonstrated to be uniquely (or “unambiguously”) beam-driven (Secs. 4.2, 4.5, 5.2). Other channel- and configuration-dependent effects could contribute (e.g., differences in NILC effective beam or weights between std and deproj, or frequency-dependent foreground leakage that changes under tSZ deprojection). The current evidence is mainly qualitative correlation with $A_b(\ell)$ and the prominence of beam-envelope errors.

Recommendation: Either soften language in Sec. 5.2 (e.g., “strongly consistent with a beam/deconvolution origin”) or add targeted quantitative checks. Examples: (i) perturb the pa4_f220 beam within its envelope/eigenmodes in simulations and show the induced spread in $R_b(\ell)$ matches the observed excursion; (ii) repeat the measurement after smoothing all maps to a common resolution (or computing spectra without deconvolution but with forward-modeling) and show the feature is reduced/removed; (iii) vary masking (more aggressive point-source/cluster masking; different sky regions) to bound a foreground-driven contribution. Report the resulting changes in the highest- ℓ bins and in \bar{R}_b .

6. Physical interpretation of “ $R_b(\ell) \approx 1$ ” under tSZ deprojection is not fully articulated in the bigger-picture sense (Secs. 1, 5.1, 5.5). Since NILC(deproj) explicitly alters the map by nulling tSZ, one might expect a predictable differential effect on cross-spectra—especially at higher frequencies—depending on residual foreground correlations, masking, and the extent to which the cross-spectrum is CMB-dominated over $200 \leq \ell \leq 1500$. Without a short expectation argument, it is harder to interpret unity as a meaningful validation rather than an accidental cancellation.

Recommendation: Add a brief discussion (Sec. 5.1 or 5.5) of the expected sign/magnitude of the change in $b \times \text{NILC TT}$ cross-spectra when tSZ is deprojected: why the effect should be small over 200–1500 (primary CMB dominance; tSZ subdominant in cross with a CMB-cleaned map; masking), and what residual foreground terms (CIB/radio/dust/kSZ) could in principle change under the deprojection constraint. If possible, include an order-of-magnitude estimate or a simple simulation/analytic forecast for $\Delta C_\ell/C_\ell$ to contextualize the measured constraints.

7. Scale-cut recommendations are useful but the cosmological/analysis impact is not quantified (Secs. 5.1, 5.4, 6). The paper recommends $\ell_{\text{max}} = 1500$ (and possible extensions) and notes a $\approx 4\%$ effect for pa4_f220 at high ℓ , but does not translate these into expected parameter shifts, information loss, or robustness impact for representative downstream use cases (e.g., TT-only parameter constraints or TT-based lensing).

Recommendation: Augment Secs. 5.1/5.4/6 with a quantitative impact estimate. Even a simple Fisher-style calculation or a pipeline-based reweighting that answers: (i) how much pa4_f220 contributes to total TT information in current DR6 analyses; (ii) the parameter impact of a constant 1–4% multiplicative distortion over a given ℓ -range; and (iii) the net effect of excluding pa4_f220 or cutting at $\ell = 1500$ vs 2000. Approximate numbers (with clear assumptions) are sufficient and would make the recommendations substantially more actionable.

8. Null-test reporting is not fully comprehensive across channels/configurations (Secs. 3.5, 4.3; Fig. 2). Figure 2 shows only a representative channel while the text states similar results hold for all channels; however, readers cannot easily assess whether any borderline PTEs exist, nor whether null consistency is equally strong for both NILC(std) and NILC(deproj) cross-spectra.

Recommendation: Add a compact table (Sec. 4.3 or an appendix) listing χ^2/dof and PTEs for the split-difference null spectra for all six channels (and clarify whether they are computed for both NILC configurations or for the ratio pipeline specifically). This can replace the need for many extra plots while making the null-test claim verifiable.

Minor issues

1. Presentation of pa4_f220 uncertainties is not fully uniform across abstract/text/table/figure (Abstract; Sec. 4.4; Table 1; Fig. 3). In places \bar{R}_b is quoted with a single total uncertainty, while Table 1 emphasizes statistical uncertainties and Fig. 3 shows both “stat” and “stat+beam,” which can confuse readers about what uncertainty is being compared across channels.

Recommendation: Standardize reporting: in Table 1 either (i) include both σ_{stat} and σ_{beam} as separate columns plus σ_{total} , or (ii) clearly state in the caption which column is tabulated and where the complementary term is provided. Use the same convention in the abstract and Sec. 4.4.

2. Definition and interpretation of the beam-amplification diagnostic $A_b(\ell)$ is not well matched to the stated cross-spectrum deconvolution (Sec. 3.3; Eq. (3)). The text motivates amplification by $1/(b_\ell^{(b)}b_\ell^{\text{NILC}})$, but Eq. (3) uses only $1/[b_\ell^{(b)}]^2$ normalized at ℓ_{\min} , which resembles an auto-spectrum proxy and omits NILC beam effects (which may differ between std and deproj).

Recommendation: Either redefine $A_b(\ell)$ to track the actual deconvolution relevant for the cross-spectrum (e.g., proportional to $1/[b_\ell^{(b)}b_\ell^{\text{NILC},X}]$ with $X = \text{std/deproj}$, or provide two curves), or explicitly justify why the ACT-only proxy is sufficient for the comparisons being made (including discussion of whether NILC effective beams are identical between configurations).

3. Several implementation details needed for reproducibility are too terse (Secs. 2.1–2.3, 3.1). In particular: how ACT CAR maps are combined with HEALPix NILC maps for harmonic analysis (reprojection/interpolation choices); mask resolution and apodization; and explicit listing of the split-cross combinations and how variances/weights w_ℓ in Eq. (5) are constructed.

Recommendation: Add concise, explicit descriptions in Secs. 2.1–2.3 and 3.1: reprojection method and resolution; mask construction/apodization and treatment of point sources/clusters; enumerate the split-cross pairings used (all $i \neq j$ or a subset) and state whether variances are analytic, from simulations, or from data-based scatter, and how they enter w_ℓ .

4. Goodness-of-fit testing is referenced (χ^2/dof and PTEs) without defining the χ^2 used or whether bin-to-bin covariance is neglected (Sec. 4.4; also relevant to summary statistics in Eq. (5)).

Recommendation: Provide the explicit χ^2 expression and clarify the assumed covariance model (diagonal-only vs including correlations). If diagonal-only is used, briefly justify why correlations are negligible for these binnings/masks or provide a sensitivity check.

5. Scope statements occasionally risk overgeneralizing beyond TT (Secs. 5.1, 6). The analysis is TT-only, but some phrasing could be read as validating ACT×NILC spectra more broadly, despite different systematics/weights in polarization.

Recommendation: Add explicit language in Secs. 5.1 and 6 that conclusions apply to TT cross-spectra only, and briefly outline what would need to change to extend the framework to TE/EE (e.g., different noise/splits/beam handling).

6. The comparison to Planck NILC validation is currently mostly qualitative (Sec. 5.3) and does not map Planck-era tests onto the specific diagnostics used here, nor highlight differences in regime (beam sizes, noise, ℓ -range).

Recommendation: Strengthen Sec. 5.3 with a short, concrete comparison: summarize Planck NILC validation tests and their typical consistency levels and contrast them with the present ℓ -range/beam regime and diagnostics ($R_b(\ell)$, $A_b(\ell)$, null tests). A small table or tightly written paragraph would suffice.

7. Figure clarity and accessibility can be improved (Fig. 2–3). Figure 2 lacks an explicit legend mapping series to tests and becomes cluttered; Figure 3 relies heavily on color to distinguish uncertainty components and the narrow x-range can visually exaggerate deviations.

Recommendation: For Fig. 2 add a clear legend and reduce overplotting (marker shapes, slight ℓ offsets, split into panels, or transparency), and state what error bars include. For Fig. 3 use style (line/marker) in addition to color for stat vs total, and consider plotting $(\bar{R}_b - 1) \times 10^3$ or adding an inset/wider axis to reduce perceptual exaggeration.

Very minor issues

1. Reference list and formatting contain errors/inconsistencies (References), including the MASTER citation year and malformed repeated bullet prefixes, which may confuse readers and citation tooling.

Recommendation: Clean up the References: correct “Hivon et al. (2025)” to Hivon et al. (2002), remove malformed repeated prefixes, and ensure consistent formatting of Planck Collaboration entries and other key NILC/ACT references.

2. Minor LaTeX/typography issues affect readability (Sec. 2; Sec. 3.2; Sec. 5): raw LaTeX for arcminutes, inconsistent $\backslash\mathrm{mathrm}$ usage, occasional broken equation formatting, and an out-of-place mid-text banner/heading around Sec. 5.

Recommendation: Proofread the LaTeX source to standardize unit typography (e.g., arcminutes), unify roman-font conventions, fix equation line breaks/parentheses, and ensure headings/banners are consistent with the rest of the manuscript structure.

3. Small definitional omissions: ℓ_{\min} in Eq. (3) is implied but not explicitly stated; binning conventions (bin centers; whether R_b is computed from binned or unbinned spectra) are not fully specified.

Recommendation: Define ℓ_{\min} explicitly near Eq. (3) (likely $\ell_{\min} = 200$) and add one sentence describing the binning convention (top-hat bins, bin centers, and whether ratios are formed before/after binning).

Key statements and references

- ✓ The ACT DR6 beam transfer functions $b_\ell^{(b)}$ for each channel are measured from dedicated planet observations and released in Naess et al. (2025), while the effective NILC beam b_ℓ^{NILC} is computed from the scale-dependent NILC weights and the individual input-map beams as described in Coulton et al. (2024), enabling a deconvolution factor $1/(b_\ell^{(b)}b_\ell^{\text{NILC}})$ that grows rapidly at high ℓ for narrow-beam channels and motivates the beam-amplification diagnostic $A_b(\ell)$ closely related to the beam-error propagation formalism developed for ACT DR4 and extended in DR6 (Choi et al. 2020; Louis et al. 2025).
- *Reference(s):* Naess et al., 2025, Coulton et al., 2024, Choi et al., 2020

- *Justification*: Verification failed with gpt-5: Error code: 400 - {'error': {'message': 'Your input exceeds the context window of this model. Please adjust your input and try again.', 'type': 'invalid_request_error', 'param': 'input', 'code': 'context_length_exceeded'}}
- ✓ **The Monte Carlo-calibrated transfer function F_ℓ used in the pseudo- C_ℓ cross-spectrum estimator is obtained from $N_{\text{MC}} = 480$ signal-only simulations processed through the full ACT DR6 map-making, masking, and beam-convolution pipeline following Naess et al. (2025) and Louis et al. (2025), yielding a transfer-function uncertainty $\lesssim 0.2\%$ at $\ell < 2000$ that is subdominant in the deprojection-response ratio error budget.**
- *Reference(s)*: Naess et al., 2025, Louis et al., 2025, Alonso et al., 2019
- *Justification*: Verification failed with gpt-5: Error code: 400 - {'error': {'message': 'Your input exceeds the context window of this model. Please adjust your input and try again.', 'type': 'invalid_request_error', 'param': 'input', 'code': 'context_length_exceeded'}}
- **△ The ACT+Planck NILC CMB temperature maps used here, including the standard and tSZ-deprojected variants, are constructed in Coulton et al. (2024) by combining ACT DR6 and Planck PR4 data (Planck Collaboration 2020a) with scale-dependent needlet weights in the HEALPix pixelization (Górski et al. 2005), imposing a unit-response constraint to the CMB blackbody spectrum and, for the deprojected map, an additional constraint that nulls the tSZ spectral signature (Sunyaev & Zeldovich 1972; Remazeilles et al. 2011) at the cost of $\approx 10\text{--}30\%$ extra noise power depending on angular scale and resulting in an effective beam FWHM $\approx 5'$ at the scales of interest.**
- *Reference(s)*: Coulton et al., 2024, Planck Collaboration, 2020a, Górski et al., 2005
- *Justification*: Supported: Coulton et al. (2024) combine ACT DR4/DR6 with Planck PR4 (NPIPE) using a needlet ILC with scale-dependent weights, enforce unit CMB response, and provide a tSZ-deprojected CMB map via constrained ILC (Coulton et al., 2024). Not supported: the maps are processed and the needlet analysis is carried out using CAR pixelization after projecting Planck from HEALPix to CAR, not 'in the HEALPix pixelization' (Coulton et al., 2005 only defines HEALPix). Also, Coulton et al. use a common $1.6'$ FWHM beam, not $\approx 5'$, and they do not quantify a $10\text{--}30\%$ noise-power increase for the tSZ-deprojected CMB map (they only note a noise increase qualitatively). Thus the statement is only partially supported.

Mathematical consistency audit

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

Maths relevance: substantial

The paper's analytic content centers on a pseudo- C_ℓ cross-spectrum estimator with beam and transfer-function corrections (Eq. (1)), a Monte Carlo transfer-function definition (Eq. (2)), diagnostics for beam deconvolution amplification (Eq. (3)) and the deprojection-response ratio (Eq. (4)), weighted averaging (Eq. (5)), split-difference null tests (Eq. (6)), and a beam-envelope uncertainty propagation for the ratio (Eq. (7)). The main internal consistency problems arise from how beam factors should cancel in the ratio and how beam uncertainties are propagated.

Checked items

1. **△ Pseudo- C_ℓ cross-spectrum estimator structure** (Eq. (1), Sec. 3.1, p.2-3)
 - **Claim:** Defines the estimated cross-spectrum as $\hat{C}_\ell = (1/F_\ell) (1/b_\ell) b_\ell$ (where b_ℓ is the beam transfer function of the two maps entering the cross-spectrum, \tilde{C}_ℓ is computed from masked maps (pseudo-spectrum)).
 - **Checks:** symbolic cancellation/sanity, dimension/units consistency, definition consistency with later ratio
 - **Verdict:** UNCERTAIN; confidence: medium; impact: critical
 - **Assumptions/inputs:** F_ℓ is a multiplicative transfer function correcting filtering/masking/pixelization., $b_\ell^{(b)}$ and $b_\ell^{(\text{NILC})}$ are the beam transfer functions of the two maps entering the cross-spectrum., \tilde{C}_ℓ is computed from masked maps (pseudo-spectrum).
 - **Notes:** Dimensionally plausible (F_ℓ and beams dimensionless). However, consistency with Eq. (2) is unclear because Eq. (2) states $\tilde{C}_\ell^{\text{out}} = \tilde{C}_\ell^{\text{in}} / F_\ell$ (where $\tilde{C}_\ell^{\text{in}}$ includes beamconvolution while Eq.(1) also divides by beam factors. Without specifying whether Eq.(2) includes or excludes beamconvolution, it is not possible to verify that corrections are not double-counted).
2. **△ Monte Carlo transfer function definition** (Eq. (2), Sec. 3.2, p.3)
 - **Claim:** Defines F_ℓ as the Monte Carlo mean of $\tilde{C}_\ell^{\text{out}} / \tilde{C}_\ell^{\text{in}}$
 - **Checks:** definition consistency with Eq. (1), dimensional consistency
 - **Verdict:** UNCERTAIN; confidence: medium; impact: critical
 - **Assumptions/inputs:** Signal-only simulations are used ($N_{\text{MC}} = 480$)., $\tilde{C}_\ell^{\text{in}}$ represents the simulation input spectrum (not otherwise specified as full-sky C_ℓ or pseudo- C_ℓ)., $\tilde{C}_\ell^{\text{out}}$ is computed after map-making, masking, and beam convolution.
 - **Notes:** As a definition, Eq. (2) is dimensionless. But because $\tilde{C}_\ell^{\text{out}}$ is said to include beamconvolution, and Eq.(1) separately divides by beams, the pipeline's analytic correction factor cannot be validated, or whether F_ℓ is later used with/without explicit beam deconvolution).
3. **✓ Beam-amplification diagnostic definition** (Eq. (3), Sec. 3.3, p.3)
 - **Claim:** Defines $A_b(\ell) = (1/b_\ell^{(b)})^2$ normalized by its value at ℓ_{min} , as a measure of deconvolution amplification growth.
 - **Checks:** algebraic simplification, consistency with stated purpose

- **Verdict:** PASS; confidence: high; impact: minor
 - **Assumptions/inputs:** ℓ_{\min} corresponds to the lowest multipole in the analysis range (implied), $b_\ell^{(b)}$ is the ACT channel beam transfer function.
 - **Notes:** Algebraically $A_b(\ell) = (b_{\min}^{(b)}/b_\ell^{(b)})^2$, so $A_b(\ell_{\min}) = 1$ and it grows as the beam decays. However, it is only a proxy for the cross-spectrum deconvolution factor stated elsewhere ($1/(b_\ell^{(b)} b_\ell^{\text{NILC}})$).
4. ✓ **Deprojection-response ratio definition** (Eq. (4), Sec. 3.4, p.3)
- **Claim:** Defines $R_b(\ell)$ as the ratio of deprojected-to-standard estimated cross-spectra for channel b .
 - **Checks:** definition consistency, dimension/units consistency
 - **Verdict:** PASS; confidence: high; impact: moderate
 - **Assumptions/inputs:** Both numerator and denominator are computed with the same estimator form (Eq. (1)) but different NILC configurations.
 - **Notes:** $R_b(\ell)$ is dimensionless and the definition is unambiguous.
5. ✓ **Cancellation of common beam factors in $R_b(\ell)$** (Implied by Eqs. (1) and (4), Secs. 3.1 & 3.4, p.2–3)
- **Claim:** Given Eq. (1), the ACT channel beam factor $b_\ell^{(b)}$ should cancel in $R_b(\ell)$ if the same $b_\ell^{(b)}$ is used in both numerator and denominator.
 - **Checks:** symbolic cancellation
 - **Verdict:** PASS; confidence: high; impact: critical
 - **Assumptions/inputs:** Eq. (1) applies identically to both $X = \text{std}$ and $X = \text{deproj}$ for a fixed channel b . The ACT beam transfer function $b_\ell^{(b)}$ used in deconvolution is identical in both computations.
 - **Notes:** Substituting Eq. (1) into Eq. (4) gives $R_b(\ell) = [\tilde{C}_{\text{deproj}}/\tilde{C}_{\text{std}}] \times [F_{\text{std}}/F_{\text{deproj}}] \times [b_{\text{std}}/b_{\text{deproj}}]$. *cancel exactly. Therefore any analytic attribution of $R_b(\ell)$ bias to ACT-beam-only deconvolution requires additional non-cancelling effects to be explicitly introduced.* The factor $b_{\text{std}}/b_{\text{deproj}}$ is $b_{\text{std}}/b_{\text{deproj}}$.
6. ✓ **Inverse-variance weighted mean ratio** (Eq. (5), Sec. 3.4, p.3)
- **Claim:** Defines $\bar{R}_b = 1/\sigma^2[R_b(\ell)]$ as the inverse-variance weighted mean of $R_b(\ell)$ with weights $w_\ell = 1/\sigma^2[R_b(\ell)]$.
 - **Checks:** algebraic correctness, definition consistency
 - **Verdict:** PASS; confidence: high; impact: minor
 - **Assumptions/inputs:** $\sigma^2[R_b(\ell)]$ is the variance per multipole/bin used for weighting.
 - **Notes:** Standard weighted-mean form; no internal algebra issues.
7. ✓ **Split-difference half-difference maps** (Sec. 3.5, p.3)
- **Claim:** Defines $d_{01} = (\text{set0} - \text{set1})/2$ and $d_{23} = (\text{set2} - \text{set3})/2$; these contain no sky signal to first order.
 - **Checks:** algebraic cancellation
 - **Verdict:** PASS; confidence: high; impact: minor
 - **Assumptions/inputs:** set0 and set1 observe the same sky signal with independent noise/systematics contributions (similarly for set2 and set3).
 - **Notes:** If $\text{set } i = s + n_i$, then $d_{01} = (n_0 - n_1)/2$ and the sky term cancels exactly. Same for d_{23} .
8. ✓ **Null cross-spectra should be zero** (Eq. (6), Sec. 3.5, p.3)
- **Claim:** $\hat{C}_{d01} \times \hat{C}_{d23}$ should be consistent with zero, and $\hat{C}_{d01} \times \hat{C}_{d23}$.
 - **Checks:** probabilistic expectation sanity check
 - **Verdict:** PASS; confidence: medium; impact: minor
 - **Assumptions/inputs:** Split-difference maps contain only noise/systematics, uncorrelated with NILC and between d_{01} and d_{23} (under the null).
 - **Notes:** Under the stated independence assumptions, the expected cross-power is zero. The paper does not provide covariance expressions, but the qualitative null expectation is consistent.
9. ✗ **Beam-envelope propagation into the ratio** (Eq. (7), Sec. 3.6, p.3–4)
- **Claim:** $\sigma_{\text{beam}}[R_b(\ell)] = R_b(\ell) \times \sqrt{2} \times \delta b_\ell^{(b)}/b_\ell^{(b)}$, with $\sqrt{2}$ because the beam appears in numerator and denominator.
 - **Checks:** symbolic consistency with Eq. (4), error propagation sanity check
 - **Verdict:** FAIL; confidence: high; impact: critical
 - **Assumptions/inputs:** $\delta b_\ell^{(b)}$ is the 1σ uncertainty on the ACT channel beam transfer function., Beam uncertainties in numerator and denominator are treated as contributing additively in variance (implied by $\sqrt{2}$).
 - **Notes:** From Eqs. (1) and (4), the ACT beam factor $b_\ell^{(b)}$ cancels exactly in $R_b(\ell)$ if the same ACT beam deconvolution is used in both numerator and denominator. In that case, the derivative $\partial R_b/\partial b_\ell^{(b)}$ is zero (to first order), so an ACT-beam-only $\delta b_\ell^{(b)}/b_\ell^{(b)}$ term should not appear. Conversely, non-cancelling beam terms would come from NILC effective beam differences/uncertainties, which are not included in Eq. (7). The $\sqrt{2}$ justification conflicts with the cancellation implied by the ratio definition.

10. Δ **Interpretation linking $A_b(\ell)$ growth to bias in $R_b(\ell)$** (Sec. 3.3 (text), Secs. 4.2 & 5.2, p.3–5)

- **Claim:** High- ℓ deviation in $R_b(\ell)$ for pa4 f220 is 'driven by beam-deconvolution amplification' quantified by $A_b(\ell)$.
- **Checks:** consistency with algebraic cancellation of common factors
- **Verdict:** UNCERTAIN; confidence: medium; impact: critical
- **Assumptions/inputs:** $A_b(\ell)$ reflects the relevant amplification affecting $R_b(\ell)$. The deconvolution amplification is dominated by the ACT channel beam.
- **Notes:** Given the exact cancellation of $b_\ell^{(b)}$ in $R_b(\ell)$ implied by Eqs. (1) and (4), an ACT-beam-only amplification proxy $A_b(\ell)$ cannot by itself produce a systematic offset in the ratio. Such an effect would require additional non-cancelling terms (e.g., different $b_\ell^{(\text{NILC}, X)}$ or F_X) to be explicitly incorporated. The paper does not provide the missing analytic link, so the stated causal attribution is not verifiable as written.

Limitations

- Audit is restricted to the provided PDF text; several needed clarifications are absent (exact definition of $\{\tilde{C}\}_{\text{in/out}}$ in Eq.(2), whether NILC beams differ between standard and deproj, whether F_{ell} is configuration-dependent, and covariance assumptions for beam uncertainties).
- No numerical validation was performed; the audit only checks symbolic/algebraic consistency and whether stated uncertainty/diagnostic formulas follow from the definitions given.

Numerical results audit

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

No candidate numerical checks were executed; therefore, no PASS/FAIL determinations can be reported for C1–C10 based on computed results.

Checked items

1. Δ **C1** (Abstract (page 1))
 - **Claim:** “Over the multipole range $200 \leq \ell \leq 1500$, five of six channels yield inverse-variance-weighted mean ratios consistent with unity at the sub-percent level ($|\bar{R}_b - 1| < 0.005$). The remaining channel, pa4 f220, exhibits a mild $\sim 1\sigma$ excess ($\bar{R}_b = 1.042 \pm 0.040$).”
 - **Checks:** cross-check repeated summary vs Table 1; compute $|\bar{R}_b - 1|$ for the five channels; check pa4 f220 combined uncertainty matches ± 0.040
 - **Verdict:** UNCERTAIN
 - **Notes:** Check not executed due to execution error; no computed comparison available.
2. Δ **C2** (Results §4.4 (page 4) vs Table 1 (page 6) vs Conclusions bullet (page 7))
 - **Claim:** “pa4 f220 channel gives $\bar{R}_b = 1.042 \pm 0.018$ (stat) ± 0.035 (beam)” and Conclusions also report “ $\bar{R}_b = 1.042 \pm 0.040$ ” for pa4 f220.
 - **Checks:** quadrature combination and rounding consistency
 - **Verdict:** UNCERTAIN
 - **Notes:** Check not executed due to execution error; no computed comparison available.
3. Δ **C3** (Results §4.5 (page 4))
 - **Claim:** “for pa4 f220 the beam envelope is the dominant uncertainty source ($\sigma_{\text{beam}} = 0.035$ versus $\sigma_{\text{stat}} = 0.018$)”
 - **Checks:** inequality/dominance check
 - **Verdict:** UNCERTAIN
 - **Notes:** Check not executed due to execution error; no computed comparison available.
4. Δ **C4** (Methods §3.2 (page 3) and Limitations §5.5 (page 6))
 - **Claim:** “ $N_{\text{MC}} = 480$ ” simulations; “transfer-function uncertainty ... $\lesssim 0.2\%$ at $\ell < 2000$ ” and again “ $\lesssim 0.2\%$ at all scales below $\ell = 2000$ ”
 - **Checks:** repeated-constant consistency check
 - **Verdict:** UNCERTAIN
 - **Notes:** Check not executed due to execution error; no computed comparison available.
5. Δ **C5** (Methods §3.1 (page 3))
 - **Claim:** “We use split-cross spectra—averaging over the six independent pairs from the four time splits...”
 - **Checks:** combinatorics check
 - **Verdict:** UNCERTAIN
 - **Notes:** Check not executed due to execution error; no computed comparison available.
6. Δ **C6** (Methods §3.5 (page 3))
 - **Claim:** “Define the half-differences $d_{01} = (\text{set0} - \text{set1})/2$ and $d_{23} = (\text{set2} - \text{set3})/2$.”

- **Checks:** algebraic scaling check (half-difference definition)
 - **Verdict:** UNCERTAIN
 - **Notes:** Check not executed due to execution error; no computed comparison available.
7. \triangle **C7** (Data §2.2 (page 2))
- **Claim:** “tSZ-deprojected variant adds a second constraint ... at the cost of $\sim 10\text{--}30\%$ additional noise power depending on angular scale.”
 - **Checks:** range sanity check
 - **Verdict:** UNCERTAIN
 - **Notes:** Check not executed due to execution error; no computed comparison available.
8. \triangle **C8** (Data §2.1 (page 2))
- **Claim:** “The beam full-width at half-maximum (FWHM) ranges from $\approx 2.1'$ at 90 GHz to $\approx 1.0'$ at 220 GHz.”
 - **Checks:** monotonicity check of stated range endpoints
 - **Verdict:** UNCERTAIN
 - **Notes:** Check not executed due to execution error; no computed comparison available.
9. \triangle **C9** (Data §2.3 (page 2))
- **Claim:** “The resulting sky fraction is $f_{\text{sky}} \approx 0.40$.”
 - **Checks:** probability range check
 - **Verdict:** UNCERTAIN
 - **Notes:** Check not executed due to execution error; no computed comparison available.
10. \triangle **C10** (Methods §3.1 (page 3))
- **Claim:** “The binned estimator averages \hat{C}_ℓ into bandpowers of width $\Delta\ell = 40$.”
 - **Checks:** positive-integer check
 - **Verdict:** UNCERTAIN
 - **Notes:** Check not executed due to execution error; no computed comparison available.

Limitations

- Only the provided parsed text of the PDF was used; no external datasets, code, or beam/spectrum arrays are available for recomputation-heavy checks.
- Figure-based numerical claims (e.g., R_b at specific ℓ , A_b thresholds) cannot be verified without underlying data; plot-pixel extraction is explicitly avoided.
- Several statements reference per- ℓ uncertainties, transfer functions, and beam envelopes that are not numerically specified in the text, limiting verification to algebraic/consistency and simple recombinations.
- Execution errors prevented running any checks: Sandbox policy violation: from-import of 'typing' is not allowed