Skepthical review: *A Low-Significance Measurement of the kSZ $-\tau-M$ Scaling Relation from Wiener-Filtered Simulated CMB Maps*

Summary

This paper presents an end-to-end feasibility study for measuring the cluster kSZ optical-depth–mass (τ–M) scaling relation with a pipeline intended to resemble current ground-based CMB temperature surveys. Using simulated single-frequency maps over $100\ \mathbf{deg}^2$ ($1.4'$ beam, $20\ \mu$K-arcmin white noise) and a catalog of 5,000 halos (10^{13}–$10^{15}\ M_\odot$; Sec. 2.1), the analysis Wiener-filters the map to estimate/subtract the primary CMB (Sec. 2.2), then applies a mass-weighted pairwise kSZ estimator using halo line-of-sight velocities with jackknife covariances and shuffled-velocity nulls (Sec. 2.3–2.4). The Wiener filtering step substantially reduces map variance and improves the pairwise SNR (Sec. 3.1–3.3), but velocity reconstruction from the sparse halo catalog fails ($r \approx -0.026$; Sec. 3.2), so the τ $-M$ fit proceeds only under an optimistic "true velocities" assumption. Even in this upper-bound scenario the detection is marginal (SNR ≈ 1.56) and the τ–M slope is essentially unconstrained ($\alpha = 0.38 \pm 7.23$; Sec. 3.4). The negative result is potentially valuable, but several key implementation details (filter conventions/transfer function, temperature extraction at halo positions, estimator normalization/units, regression under low SNR, and velocity-reconstruction specifics) need to be made explicit and stress-tested to establish whether the limitations are intrinsic to the stated survey/catalog combination or partly methodological/optimizable.

Strengths

- Clear motivation connecting kSZ measurements to intracluster gas physics and the τ–M relation (Introduction/Sec. 1).
- Survey-like simulation choices are stated up front ($100\ \mathbf{deg}^2$, $1.4'$ beam, $20\ \mu$K-arcmin noise, 5,000 halos; Sec. 2.1), making the exercise easy to interpret as a "current-generation" realism check.
- Modular pipeline (Wiener filtering → pairwise estimator → covariance/nulls → τ–M fit) is conceptually straightforward and, with added details, could be reproducible (Sec. 2.2–2.4).
- Useful quantitative diagnostics are already included (map RMS reduction; SNR before/after filtering; velocity-reconstruction correlation coefficient; Sec. 3.1–3.3).
- Appropriately cautious framing: the paper does not oversell the marginal detection and is transparent that "true velocities" represent an upper bound (Sec. 3.2–3.4, Sec. 4).
- Figures convey the dynamic-range challenge and show multiple cross-checks (e.g., shuffled-velocity null) that are standard and appropriate for kSZ analyses (Fig. 1–2, Sec. 2.4).

Major issues

1. **Velocity reconstruction is central to the paper's feasibility conclusion, but the reconstruction method is only sketched and its failure is not yet diagnostic (Sec. 2.3, Sec. 3.2, Sec. 4).** Key missing specifics include the exact $\delta \to v$ mapping (e.g., continuity-equation/Fourier-space form), grid/pixel and k-space resolution, smoothing/regulariza-

tion, halo bias model and redshift treatment, boundary conditions on a $100\ \mathrm{deg}^2$ patch, and how shot noise is handled. Without this, it is unclear whether $r \approx -0.026$ reflects an unavoidable sparsity/geometry limitation or a particular (possibly suboptimal) implementation choice.

Recommendation: Expand Sec. 2.3 and Sec. 3.2 with a fully specified reconstruction recipe: the equation used to infer v from δ (including fH factors), how δ is built from halos (mass weighting? bias correction?), smoothing scale(s), grid size, and treatment of edges (apodization/zero-padding/periodic). Provide scale-dependent diagnostics (e.g., $r(k)$ or r after low-pass filtering) in addition to a single Pearson r. Add at least two controlled tests: (i) an "oracle" reconstruction using the true matter density field (to isolate tracer sparsity vs. algorithmic issues), and (ii) a tracer-density scaling study (subsample/augment halos or include lower-mass tracers if available) showing how r changes with number density for this geometry.

2. **The headline constraints assume true halo velocities, but the paper does not quantify how kSZ SNR and τ–M uncertainties degrade under realistic velocity errors (Sec. 3.2–3.4, Sec. 4).** As written, the reader learns that reconstruction fails, but not what reconstruction quality would be required for τ–M science, or how "close" the setup is to feasibility if r were modest (e.g., 0.3–0.7).

   Recommendation: Add a simple propagation model linking velocity-reconstruction fidelity to kSZ amplitude/SNR (e.g., signal suppression $\propto r$ between reconstructed and true pairwise velocities, or an equivalent multiplicative calibration factor). Using the measured $r \approx -0.026$ and a few representative literature values (e.g., $r = 0.3, 0.5, 0.7$), forecast the resulting pairwise SNR and the expected errors on A and α (Sec. 3.3–3.4). Present this as a compact table/figure and summarize explicitly in Sec. 4 what reconstruction performance would be necessary for meaningful τ–M constraints under the assumed map noise/area.

3. **The Wiener-filter/CMB-subtraction step is not fully specified in terms of conventions and its kSZ transfer function, yet it directly affects the recovered τ normalization and potentially the τ–M slope (Sec. 2.2, Sec. 3.1–3.4; Fig. 1).** Eq. (1) defines $W(\ell)$ as a CMB Wiener filter, but the science signal is measured on the residual map, which effectively applies $(1 - W)$ to all components; this can attenuate kSZ in a scale-dependent way. The manuscript states attenuation is "accounted for," but does not show how (e.g., multiplicative calibration, profile-dependent correction, or simulation-based debiasing). Beam/noise conventions in the filter (whether C_ℓ^{noise} is pre/post-beam; whether spectra are multiplied by $B(\ell)^2$ consistently) also remain ambiguous.

   Recommendation: In Sec. 2.2, write the cleaned-map relation explicitly in Fourier space (e.g., $T_{\mathrm{clean}}(\ell) = [1 - W(\ell)]T_{\mathrm{obs}}(\ell)$) and state the conventions for beam convolution and the noise power spectrum (pre/post-beam) unambiguously. Quantify the effective transfer function for a cluster kSZ signal by injecting either (i) a known kSZ-only map or (ii) a parametric cluster profile convolved with the beam, passing it through the same filtering/temperature-extraction step, and measuring the recovered amplitude. Use this to either (a) debias τ (reporting de-filtered A and α) or (b) clearly define that you are fitting a "filtered τ" and provide the mapping to physical τ. Update Fig. 1/caption to match the actual operation (CMB estimation via W and residual via $1 - W$) and state how C_ℓ^{kSZ} in Eq. (1) is obtained (measured from the simulated kSZ-only map, modeled, or effectively negligible).

4. **The kSZ temperature statistic at halo positions is underspecified and may be far from optimal, which matters because the paper's main conclusion is that instrumental noise dominates (Sec. 2.3, Sec. 3.3–3.4).** It is unclear whether T_i is a single pixel value, an interpolated sample, aperture photometry, or a matched-filter output, and how this choice relates to the $1.4'$ beam and expected halo angular sizes.

   Recommendation: Clarify in Sec. 2.3 exactly how T_i is computed from the cleaned map (pixelization/projection; interpolation; aperture radius; any additional spatial filtering). Then add a targeted optimization/robustness check: compare the baseline choice to at least one more standard kSZ photometry option (e.g., aperture photometry with a few radii tied to beam FWHM and/or θ_{500}, and/or an "oracle" matched filter using the injected profile if available). Report how the pairwise SNR in Sec. 3.3 changes; this will help determine whether the reported SNR ≈ 1.56 is close to the best achievable for the stated map specs.

5. **Estimator definition/units and its connection to τ are currently ambiguous and may be dimensionally inconsistent (Sec. 2.3; Eq. (2)).** If the map is in μK (Sec. 2.1), Eq. (2) yields τ only if T_i denotes $\Delta T/T_{\mathrm{CMB}}$ (dimensionless); otherwise a missing $1/T_{\mathrm{CMB}}$ factor (and possibly other normalization details) makes the reported τ–M normalization hard to interpret and compare to literature. The pair-sum indexing ($i \neq j$ vs unique pairs) and pair selection (separation cuts, geometry factors) are also not fully specified, limiting interpretability and reproducibility.

   Recommendation: In Sec. 2.3, define T_i explicitly as either ΔT or $\Delta T/T_{\mathrm{CMB}}$ and revise Eq. (2) accordingly (include a $1/T_{\mathrm{CMB}}$ factor if using temperature units). State whether the sum is over ordered pairs ($i \neq j$) or unique unordered pairs ($i < j$) and ensure the text/equation agree. Describe pair selection choices: whether all separations are included, whether there is a maximum separation, and whether any geometric/projection factor enters the estimator (or justify why not). Add a brief statement connecting the estimator's expectation value to an average τ under your filtering/photometry choices so that A in Sec. 3.4 has a clear operational meaning.

6. **The τ–M regression methodology is not well-defined for a low-SNR measurement and may be unstable or biased (Sec. 2.4, Sec. 3.4; Fig. 2a).** The text mentions a log–log linear regression, but binned τ estimates can be noise-dominated and may cross zero, making log transforms ill-defined and potentially biasing slope/normalization. In addition, the stability of the jackknife covariance (and bin-to-bin correlations) is not demonstrated, yet α has an enormous uncertainty (± 7.23), raising concerns about ill-conditioned fits rather than purely lack of information.

   Recommendation: Specify the exact fitting procedure in Sec. 2.4/Sec. 3.4: whether the fit is performed in linear space or log space; how bins with $\hat{\tau} \leq 0$ are treated; whether the full jackknife covariance is used (generalized least squares) and whether any regularization is applied. Provide at least one covariance diagnostic (correlation matrix heatmap or key summary statistics such as condition number/eigenvalue spectrum). Add robustness checks: repeat the fit with diagonal-only covariance, with fewer/more mass bins, and with different numbers of jackknife regions (e.g., 50/100/200) to show that A and α are not artifacts of covariance noise or binning choices. If possible, include an end-to-end recovery test in a higher-SNR toy setup (e.g., lower noise or larger area) with a known injected τ–M law to verify the pipeline can recover an unbiased slope when information exists.

7. **The paper's "fundamentally limited" conclusion is based on a single survey/catalog configuration, with limited scaling/forecast context and limited comparison to existing kSZ detections/forecasts (Abstract, Sec. 4).** This makes it hard to generalize: is the limitation primarily (i) area ($100\ \mathrm{deg}^2$), (ii) noise ($20\ \mu$K-arcmin), (iii) single-frequency temperature-only (no multifrequency tSZ cleaning), (iv) halo number density/mass threshold, or (v) the chosen photometry/estimator?

   Recommendation: Either temper the language to explicitly restrict conclusions to the studied configuration or add a simple scaling/forecast section in Sec. 4: vary (even analytically) map noise, area, and halo density and report expected SNR and σ_α trends (e.g., SNR $\propto \sqrt{\mathrm{area}}$ and $\propto 1/\mathrm{noise}$ as a baseline, plus an empirical dependence on N_{halo} from resampling). Provide 1–2 comparison points to the literature (ACT/SPT/Planck pairwise kSZ detections/forecasts and typical map depths/areas/catalog densities) to position why this configuration underperforms and what improvements (larger area, lower noise, denser tracers, better velocity recon) would be required for meaningful τ–M science.

8. **Figures 1–2 contain several ambiguities that directly affect interpretation of the main methodological claims (Fig. 1 caption vs definition of the Wiener filter; beam/noise treatment; and missing uncertainty/context in Fig. 2 SNR and τ–M panels).**

   Recommendation: Revise Fig. 1 to (i) clearly state whether the plotted filter is $W(\ell)$ (CMB estimator) or $1 - W(\ell)$ (residual map response), (ii) annotate the analysis multipole range used, and (iii) clarify beam and noise conventions (and optionally show $B(\ell)$). For Fig. 2, add/clarify: uncertainty visualization for τ–M fit (credible band or fit covariance), explicit handling/meaning of negative $\hat{\tau}$ points (sign conventions and whether they enter the fit), and uncertainty or definition for the SNR bars (including what exactly differs between bars: filtering on/off; true vs reconstructed velocities). Reduce overplotting in the velocity-scatter panel (e.g., density/hexbin) and add necessary plot metadata (axes units, colorbars where relevant).

Minor issues

1. Reproducibility-critical simulation details are incomplete (Sec. 2.1–2.2): halo redshift range/distribution, cosmological parameters, map pixelization/projection and pixel size, whether the map is a light cone or snapshot, and how the halo catalog is generated/selected. These are particularly important for interpreting velocity reconstruction on a $100\ \mathrm{deg}^2$ patch.

   Recommendation: Expand Sec. 2.1 with: cosmological parameters; halo redshift range and distribution; how masses are defined (e.g., M_{500c}) and measured; map-making details (flat-sky vs HEALPix/CAR; pixel size; projection); whether periodic boundaries/apodization are used for FFTs; and the provenance of the halo catalog (simulation volume/resolution or halo-model prescription) including any selection cuts.

2. The construction of the Wiener filter includes C_ℓ^{kSZ} in the denominator but its source is not specified, and its practical relevance is unclear (Sec. 2.2).

   Recommendation: State whether C_ℓ^{kSZ} is taken from the input kSZ-only simulation, a theoretical template, or neglected/approximated. If it is negligible over the ℓ-range used, explicitly show or state that omitting it changes results negligibly (e.g., $\Delta\mathrm{SNR}$).

3. Jackknife covariance construction lacks operational details (geometry/halo counts per region), and the null test is summarized only by a p-value without describing the null distribution (Sec. 2.4, Sec. 3.3).

   Recommendation: In Sec. 2.4, describe how the 200 jackknife regions are tiled over $100\ \mathbf{deg}^2$, their typical area, and the typical number of halos removed per resample. In Sec. 3.3, add basic null-distribution diagnostics from shuffled velocities (mean, width, approximate Gaussianity) and state where the measured signal lies within that distribution.

4. Edge effects/flat-sky approximations and FFT boundary conditions can bias harmonic-space filtering on a finite $100\ \mathbf{deg}^2$ patch, but treatment is not discussed (Sec. 2.1–2.2).

   Recommendation: Briefly document whether the analysis assumes flat-sky Fourier transforms, whether maps are apodized/padded, and whether periodic boundaries are assumed. Provide a short justification (or a simple check) that these choices do not materially affect the filtered RMS or the pairwise estimator for the separations used.

5. Additional real-world contaminants (tSZ, CIB/radio sources, Galactic emission, atmosphere/scan systematics) are not included and only lightly discussed, yet the paper aims to speak to “current-generation survey” feasibility (Sec. 1, Sec. 4).

   Recommendation: Add a concise paragraph (Sec. 4 is fine) explaining that the simulation includes only primary CMB + kSZ + white noise, and discuss qualitatively how adding tSZ/foregrounds (especially in single-frequency temperature-only analyses) would likely further reduce kSZ detectability or complicate filtering, reinforcing that the presented results are optimistic.

6. Key performance numbers are scattered and would be easier to absorb if summarized (Sec. 3.1–3.3).

   Recommendation: Add a small table summarizing: pre/post-filter map RMS (by component if available), pairwise SNR (unfiltered vs filtered; true vs reconstructed velocities), and velocity-reconstruction r. This would make the pipeline “budget” immediately clear.

Very minor issues

1. Minor presentation/LaTeX formatting inconsistencies: heading levels (stray hash markers), inconsistent unit typography ($\mathbf{deg}^2$ vs $\mathbf{deg}^2$, arcmin symbols), and equation referencing/numbering (Sec. 1–3).

   Recommendation: Standardize headings and unit formatting throughout, and format the main equations (Wiener filter, estimator, τ–M law) as displayed, consistently numbered equations that are referenced uniformly in the text.

2. Equation/pair notation clarity: the text says “unique pairs” but Eq. (2) uses $i \neq j$; and some symbols (e.g., M_{500}, v_r) are not always defined at first appearance (Sec. 2.3).

   Recommendation: Align the summation convention ($i < j$ or $i \neq j$ with an explicit note about ordered pairs) and ensure all symbols are defined at first use with units where applicable.

3. Figure styling/readability (color distinguishability, panel label size, legend clutter) occasionally makes it harder to parse key curves/points (Fig. 1–2).

   Recommendation: Adopt colorblind-safe palettes and/or distinct line styles, increase panel-label font sizes, and reduce legend clutter (e.g., inset/ratio panels or fewer overlapping curves).

Key statements and references

- **• Simple, self-similar models of the intracluster medium predict a scaling of Thomson optical depth with halo mass of the form $\tau \propto M^{2/3}$, which serves as the theoretical expectation for the τ–M relation used for comparison in this work.**
- *Reference(s):* (none)
- **• The theoretical ΛCDM angular power spectrum is used for C_ℓ^{CMB} in constructing the Wiener filter, with all theoretical power spectra multiplied by the instrumental beam transfer function $B(\ell)^2$, where $B(\ell) = \exp(-\ell^2\sigma_b^2/2)$ and $\sigma_b = \mathrm{FWHM}/\sqrt{8\ln 2}$.**
- *Reference(s):* (none)
- **• The recovered slope of the τ–M scaling relation, $\alpha = 0.38 \pm 7.23$, is statistically consistent with the theoretical expectation of $\alpha = 2/3$ from self-similar models, implying that the present analysis cannot meaningfully constrain deviations from the self-similar prediction for intracluster baryonic physics.**
- *Reference(s):* (none)

Mathematical consistency audit

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

Maths relevance: substantial

The paper’s analytic core comprises (i) the kSZ temperature–optical-depth–velocity relation, (ii) a Wiener-filter construction used to estimate/subtract the primary CMB component in harmonic space, and (iii) a mass-weighted pairwise/velocity-difference estimator intended to infer a mean optical depth per mass bin and fit a power-law τ–M scaling. Most definitions are coherent, but the τ estimator has a key units/normalization ambiguity (missing T_{CMB} if map temperatures are in μK), and the stated log-log regression is not fully specified in the presence of potentially non-positive τ estimates.

Checked items

1. ✔ **kSZ temperature shift relation** (Sec. 1–2 (description), p.2)

  - **Claim:** The kSZ temperature fluctuation satisfies $\Delta T/T_{\rm CMB} = -\tau v_r/c$.
  - **Checks:** symbol/definition consistency, dimensional/units consistency, sign convention sanity check
  - **Verdict:** PASS; confidence: high; impact: moderate
  - **Assumptions/inputs:** v_r is the line-of-sight peculiar velocity, τ is the Thomson optical depth (dimensionless), ΔT denotes the kSZ-induced temperature fluctuation
  - **Notes:** Given τ dimensionless and v_r/c dimensionless, $\Delta T/T_{\rm CMB}$ is dimensionless as required. The negative sign is a convention consistent with Doppler shift direction; the paper stays self-consistent with this sign later via the leading minus sign in Eq. (2).

2. ✔ **Optical depth definition** (Sec. 1–2 (description), p.2)

  - **Claim:** $\tau = \int n_e \sigma_T dl$ along the line of sight through the halo.
  - **Checks:** dimensional/units consistency, notation consistency
  - **Verdict:** PASS; confidence: high; impact: minor
  - **Assumptions/inputs:** n_e has units of $1/\mathrm{length}^3$, σ_T has units of length^2, dl has units of length
  - **Notes:** Integral produces a dimensionless τ. Notation is consistent with later use of τ as a scalar per halo/bin.

3. ✔ **Wiener filter form for CMB estimation** (Eq. (1), Sec. 2.2, p.3)

  - **Claim:** $W(\ell) = C_\ell^{\rm CMB}/(C_\ell^{\rm CMB} + C_\ell^{\rm kSZ} + C_\ell^{\rm noise})$ and $T_{\rm CMB,est}(\ell) = W(\ell) T_{\rm obs}(\ell)$.
  - **Checks:** algebraic form sanity check, definition consistency
  - **Verdict:** PASS; confidence: high; impact: moderate
  - **Assumptions/inputs:** Observed map decomposes as $T_{\rm obs} = T_{\rm CMB} + T_{\rm kSZ} + T_{\rm noise}$ in harmonic space, Fields are treated as stationary with power spectra C_ℓ, kSZ and instrumental noise are treated as 'noise' when estimating the CMB
  - **Notes:** As written, this is the standard Wiener-weight for the linear minimum-variance estimator of the CMB component given total contamination power $C_\ell^{\rm kSZ} + C_\ell^{\rm noise}$.

4. ✔ **Effective cleaned-map filter implied by subtraction** (Sec. 2.2, p.3)

  - **Claim:** Subtracting $T_{\rm CMB,est}$ from $T_{\rm obs}$ yields a cleaned map for kSZ extraction.
  - **Checks:** algebra between steps, operator consistency
  - **Verdict:** PASS; confidence: high; impact: moderate
  - **Assumptions/inputs:** Subtraction is performed consistently between harmonic and real space, $T_{\rm CMB,est}(\ell) = W(\ell) T_{\rm obs}(\ell)$
  - **Notes:** In harmonic space, $T_{\rm cleaned}(\ell) = T_{\rm obs}(\ell) - W(\ell) T_{\rm obs}(\ell) = (1 - W(\ell)) T_{\rm obs}(\ell)$. This is consistent with the stated procedure, though the paper does not write this explicit relation.

5. ✔ **Gaussian beam transfer function and σ_b–FWHM relation** (Sec. 2.2, p.3)

  - **Claim:** $B(\ell) = \exp(-\ell^2 \sigma_b^2/2)$ with $\sigma_b = \mathrm{FWHM}/\sqrt{8 \ln 2}$, and power spectra are multiplied by $B(\ell)^2$.

  - **Checks:** dimensional/units consistency, definition consistency
  - **Verdict:** PASS; confidence: medium; impact: minor
  - **Assumptions/inputs:** Flat-sky/2D Fourier convention for ℓ, FWHM and σ_b are in consistent angular units (radians implied)
  - **Notes:** The σ_b–FWHM relation is correct for a Gaussian. The $B(\ell)$ functional form is consistent under a flat-sky convention; the paper does not specify flat-sky vs full-sky, but the $10^\circ \times 10^\circ$ patch context supports flat-sky usage.

6. ⚠ **Beam treatment of the noise power in Eq. (1)** (Sec. 2.1–2.2, pp.3)
   - **Claim:** C_ℓ^{noise} is a constant 'white' level and 'all theoretical power spectra are multiplied by $B(\ell)^2$'.
   - **Checks:** definition consistency, modeling assumption consistency
   - **Verdict:** UNCERTAIN; confidence: medium; impact: moderate
   - **Assumptions/inputs:** Map-making order: whether beam convolution is applied before/after adding noise is not specified
   - **Notes:** If noise is added after beam smoothing (a common simulation order), C_ℓ^{noise} would be flat without $B(\ell)^2$. If noise is convolved by the beam (less common for instrumental white noise), then it picks up $B(\ell)^2$. The paper's phrasing leaves this ambiguous, so the internal consistency of Eq. (1)'s components cannot be verified from the document alone.

7. ✖ **Pairwise optical-depth estimator normalization/units** (Eq. (2), Sec. 2.3, p.4 (with map/unit descriptions in Sec. 2.1, p.3))
   - **Claim:** $\hat{\tau} = -c\, \frac{\sum w_{ij}(T_i - T_j)(v_{r,i} - v_{r,j})}{\sum w_{ij}(v_{r,i} - v_{r,j})^2}$ estimates the mean optical depth.
   - **Checks:** dimensional/units consistency, consistency with stated kSZ relation, algebraic sanity check
   - **Verdict:** FAIL; confidence: high; impact: critical
   - **Assumptions/inputs:** Underlying per-halo kSZ model is $\Delta T/T_{\text{CMB}} = -\tau v_r/c$, T_i is described as a temperature value taken from the (cleaned) map, v_r is a physical velocity
   - **Notes:** If T_i denotes a temperature (e.g., μK as described in Sec. 2.1), then $(T_i - T_j)(v_{r,i} - v_{r,j})$ has units of temperature×velocity and dividing by velocity^2 leaves temperature/velocity; multiplying by c yields temperature, not a dimensionless τ. Eq. (2) estimates τ correctly only if T_i is actually the dimensionless quantity ($\Delta T/T_{\text{CMB}}$). As written, the estimator is missing an explicit $1/T_{\text{CMB}}$ factor (or an explicit definition that T_i is normalized by T_{CMB}).

8. ✔ **Pair summation over 'unique pairs' vs $i \neq j$** (Eq. (2) and text, Sec. 2.3, p.4)
   - **Claim:** The sum is over all unique pairs (i, j) but is written with $i \neq j$.
   - **Checks:** notation consistency, algebraic invariance to double-counting
   - **Verdict:** PASS; confidence: medium; impact: minor
   - **Assumptions/inputs:** Weights satisfy $w_{ij} = w_{ji}$, Numerator and denominator use identical pair indexing

    - **Notes:** If interpreted literally, $i \neq j$ double-counts unordered pairs, but since both numerator and denominator would be multiplied by the same factor (2) under symmetric weights, $\hat{\tau}$ is unchanged. Nonetheless, the notation conflicts with the prose and should be clarified.

9. ✔ **Mass-weighting specification in the pairwise estimator** (Sec. 2.3, p.4)
    - **Claim:** Use $w_{ij} \propto M_i M_j$ to up-weight massive halo pairs.
    - **Checks:** algebraic invariance to weight scaling, definition consistency
    - **Verdict:** PASS; confidence: high; impact: minor
    - **Assumptions/inputs:** w_{ij} enters both numerator and denominator of Eq. (2)
    - **Notes:** A proportionality constant in w_{ij} cancels from $\hat{\tau}$. The document does not define an exact normalization, but none is needed for $\hat{\tau}$ as written.

10. ⚠ **Power-law τ–M model and log-log regression applicability** (Sec. 2.4 (model) and Sec. 3.4 (fit results), pp.4,7)
    - **Claim:** Fit $\tau = A(M/10^{14}M_\odot)^\alpha$ via linear regression in log-log space to obtain A and α.
    - **Checks:** definition consistency, mathematical well-definedness of transformation
    - **Verdict:** UNCERTAIN; confidence: medium; impact: critical
    - **Assumptions/inputs:** τ values used in the regression are strictly positive if taking logarithms, The fitted τ corresponds to the same τ estimated by Eq. (2)
    - **Notes:** The paper does not state how it handles potentially non-positive $\hat{\tau}k(expectedwhenmeasurementsarenoise − dominated).Withoutthis, theclaimed'log − log'regressionprocedureisnotanalyticallyverifiable.Additionally, ifEq.(2)ismissing T(asflagged), thenthefittedquantitywouldnotbe\tauunlessT_i$ is dimensionless.

Limitations

- Only the content visible in the provided PDF text/images was audited; no external references or undocumented implementation details were used.
- Several key analytic steps are described qualitatively (e.g., 'attenuation accounted for' after filtering; details of the regression and jackknife covariance formula), preventing full symbolic verification beyond consistency/units checks.
- The audit does not assess numerical correctness, simulation fidelity, statistical optimality, or whether the chosen estimators are standard/appropriate—only internal mathematical consistency.

Numerical results audit

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

Twelve candidate numerical statements were identified for verification, but none were computationally checked because the execution produced an error. Four additional items are explicitly unverified because they require underlying data products not present in the provided inputs.

Checked items

1. ⚠ **C1_area_consistency** (Page 3, Sec. 2.1 (Dataset and simulations))
   - **Claim:** The temperature map covers a $10^\circ \times 10^\circ$ ($100\ \mathrm{deg}^2$) patch of the sky.
   - **Checks:** unit_consistent_recomputation
   - **Verdict:** UNCERTAIN
   - **Notes:** Check not executed due to execution error; intended recomputation: 10×10 vs stated $100\ \mathrm{deg}^2$.
2. ⚠ **C2_rms_quadrature_total_map** (Page 3, Sec. 2.1 (end of paragraph) and Page 5, Sec. 3.1 (repeats))
   - **Claim:** Total observed map RMS is $60.41\ \mu$K, dominated by primary CMB ($56.98\ \mu$K RMS) and noise ($19.99\ \mu$K RMS); kSZ RMS is $0.14\ \mu$K.
   - **Checks:** quadrature_sum_check
   - **Verdict:** UNCERTAIN
   - **Notes:** Check not executed due to execution error; intended recomputation: $\sqrt{56.98^2 + 19.99^2 + 0.14^2}$ vs 60.41.
3. ⚠ **C3_cleaned_map_rms_vs_noise** (Page 5, Sec. 3.1)
   - **Claim:** After Wiener filtering and subtraction, the cleaned map has RMS $19.55\ \mu$K , very close to the $20\ \mu$K RMS of the instrumental noise component.
   - **Checks:** difference_check
   - **Verdict:** UNCERTAIN
   - **Notes:** Check not executed due to execution error; intended computation: absolute/relative difference between 19.55 and 20.
4. ⚠ **C4_jackknife_subregions_area_per_region** (Page 4, Sec. 2.4)
   - **Claim:** The $100\ \mathrm{deg}^2$ map is divided into 200 contiguous spatial subregions.
   - **Checks:** unit_consistent_recomputation
   - **Verdict:** UNCERTAIN
   - **Notes:** Check not executed due to execution error; intended recomputation: implied area per subregion $= 100/200$.
5. ⚠ **C5_mass_bins_equal_count** (Page 4, Sec. 2.4)
   - **Claim:** Divide the 5,000 halos into 10 bins of approximately equal halo count.
   - **Checks:** integer_division_consistency
   - **Verdict:** UNCERTAIN
   - **Notes:** Check not executed due to execution error; intended recomputation: implied nominal halos per bin $= 5000/10$.
6. ⚠ **C6_fwhm_to_sigma_conversion** (Page 3, Sec. 2.2 (beam definition))
   - **Claim:** Beam transfer uses $\sigma_b = \mathrm{FWHM}/\sqrt{8 \ln 2}$ with FWHM $= 1.4$ arcmin.
   - **Checks:** formula_recomputation
   - **Verdict:** UNCERTAIN

  - **Notes:** Check not executed due to execution error; intended recomputation: $1.4/\sqrt{8\ln 2}$.

7. ⚠ **C7_snr_improvement_ratio** (Page 7, Sec. 3.3 and Figure 2 caption (Page 9))
   - **Claim:** SNR improves from 0.51 (no filter) to 1.56 after applying the Wiener filter.
   - **Checks:** ratio_and_difference_check
   - **Verdict:** UNCERTAIN
   - **Notes:** Check not executed due to execution error; intended computations: $\Delta = 1.56 - 0.51$ and ratio $= 1.56/0.51$.
8. ⚠ **C8_pvalue_vs_sigma_consistency** (Page 7, Sec. 3.3)
   - **Claim:** Measured signal has 1.56σ significance and corresponds to a p-value of 0.146 (14.6%).
   - **Checks:** distribution_tail_probability_check
   - **Verdict:** UNCERTAIN
   - **Notes:** Check not executed due to execution error; intended comparison: 0.146 vs one- or two-sided Gaussian tail for $z = 1.56$.
9. ⚠ **C9_shuffle_count_to_pvalue_resolution** (Page 4, Sec. 2.4 and Page 7, Sec. 3.3)
   - **Claim:** Null test shuffles velocities 500 times; measured p-value is 0.146.
   - **Checks:** discrete_resolution_check
   - **Verdict:** UNCERTAIN
   - **Notes:** Check not executed due to execution error; intended check: closeness of 0.146 to multiples of $1/500$ (or $(k+1)/(N+1)$ with $N = 500$).
10. ⚠ **C10_alpha_uncertainty_matches_abstract** (Page 1 Abstract vs Page 7 Eq. (4) and Page 8 Conclusions)
    - **Claim:** Power-law slope reported as 0.38 ± 7.23 in Abstract and in results (Eq. 4) and conclusions.
    - **Checks:** cross_section_numeric_match
    - **Verdict:** UNCERTAIN
    - **Notes:** Check not executed due to execution error; intended check: exact match of central value and uncertainty across sections.
11. ⚠ **C11_A_parameter_parsing_and_sanity** (Page 7, Eq. (3))
    - **Claim:** Normalization parameter $A = (1.32 \pm 14.7) \times 10^{-4}$.
    - **Checks:** numeric_parsing_and_relative_uncertainty
    - **Verdict:** UNCERTAIN
    - **Notes:** Check not executed due to execution error; intended computation: relative uncertainty $(14.7/1.32)$ after consistent $\times 10^{-4}$ parsing.
12. ⚠ **C12_theory_slope_value** (Page 2 (self-similar $\tau \propto M^{2/3}$); Page 7 compares to $\alpha = 2/3$)
    - **Claim:** Theoretical expectation quoted as $\alpha = 2/3$.
    - **Checks:** fraction_to_decimal_conversion

  - **Verdict:** UNCERTAIN
  - **Notes:** Check not executed due to execution error; intended recomputation: **2/3** as decimal.

Limitations

- Sandbox policy violation: from-import of 'typing' is not allowed
- Only the provided PDF text was used; no underlying maps, catalogs, power spectra arrays, or intermediate numerical outputs are available for recomputation.
- No values were extracted from plot pixels or image-only figure elements; checks are limited to explicit numeric statements in the text.
- > **Some relationships (e.g., RMS quadrature) assume statistical independence; discrepancies may reflect correlations or pipeline details not specified in the PDF.**