Skepthical review: *Mapping the Optimal Sensitivity of the 21 cm Forest to Dark Matter-Baryon Scattering*

Summary

This paper explores how the high-redshift 21 cm forest (IGM + minihalo absorption at $z \approx 7$–15) can constrain elastic DM–baryon scattering that suppresses small-scale structure. Using the HAYASHI semi-analytic framework (Sec. 2.1), the authors implement DM–baryon effects phenomenologically through (i) a halo-mass cutoff $M_{\rm cut}$ applied to the halo mass function (Sec. 2.2) and (ii) an IGM temperature rescaling $f_{\rm cool} = T_k/T_k^{\rm ad}$ to represent non-standard cooling (Sec. 2.2). They quantify observable imprints via the optical-depth distribution $dN/d\tau$ (and cumulative $N(>\tau)$), and use the KL divergence as a shape-based "fingerprint" statistic (Sec. 2.3). A Fisher forecast over ${M_{\rm cut}, f_{\rm cool}, \bar{x}{\rm HI}} *isconstructedfrombinnedabsorbercounts, withPoissonuncertaintiesdeterminedbyanassumedredshift − dependentnumberofbackgroundsightlinesNintoprojectedconstraintsonvelocity − independent*(n=0)*andCoulomb − like*(n=-4)*crosssections, claimingpotentialimprovementsoverCMBlimitsby*4–5 *ordersofmagnitude(Sec.*3.4*, Conclusions).Theoverallnarrativeistimelyandtheemphasisonshapeinformationin* dN/d\tau*iscompelling, butseveralcentralmodelingandforecastassumptions(sharpcutoff, mappingto* \sigma_0/m_\chi, thermal/spin-temperature treatment, and observational/noise realism) require substantially more documentation and robustness testing for the quantitative headline claims to be convincing and reproducible.}(z) (Sec. 2.4). The main conclusions are that suppressing low-mass minihalos dominates the detectable impact (removing preferentially low-τ absorbers and changing the shape of $dN/d\tau$), that varying $f_{\rm cool}$ has little effect on absorber statistics, and that the optimal redshift window is $z \approx 8$–10 once $N_{\rm los}(z)$ is included (Sec. 3.1–3.3). The work then maps sensitivity to $M_{\rm cut

Strengths

- Clear end-to-end narrative from DM microphysics → small-scale structure suppression → 21 cm forest observables, with a logically organized presentation (Introduction, Sec. 2–4).
- Using the *shape* of $dN/d\tau$ (and information-theoretic diagnostics like KL divergence) as a discriminator is a strong and potentially powerful idea (Sec. 2.3, Sec. 3.2).
- Helpful separation of two conceptually distinct DM–baryon channels—structure suppression ($M_{\rm cut}$) vs thermal effects ($f_{\rm cool}$)—and a quantitative comparison of their impacts (Sec. 2.2, Sec. 3.1).
- Forecast framework includes nuisance parameters ($\bar{x}*{\rm HI}andf*$) and explores redshift dependence and an "optimal" observing window once source statistics are modeled (Sec. 2.4, Sec. 3.3).
- Figures generally communicate the main trends clearly (especially the differential/cumulative τ-statistics and the redshift dependence of forecast sensitivity).
- The paper makes an explicit attempt to connect phenomenological small-scale suppression ($M_{\rm cut}$) to particle-physics parameters ($\sigma(v) = \sigma_0 v^n$), which is valuable if fully specified (Sec. 2.2, Sec. 3.4).

Major issues

1. **Structure suppression is modeled as an *infinitely sharp* halo-mass cutoff $M_{\rm cut}$ (Sec. 2.2), i.e. all halos with $M < M_{\rm cut}$ are removed. For DM–baryon scattering (drag/heat exchange), realistic suppression is typically smooth and redshift/scale dependent (via a transfer function), and can also alter collapse thresholds, concentrations/formation times, and gas content. A step-function truncation risks (i) exaggerating or distorting the low-τ shape change in $dN/d\tau$ (the central "fingerprint"), and (ii) producing over-optimistic or mis-centered Fisher constraints on $M_{\rm cut}$ (Sec. 3.1–3.3), which then propagate directly to σ_0/m_χ (Sec. 3.4, Conclusions).**

Recommendation: In Sec. 2.2, replace the step cutoff with a physically motivated smooth suppression, or at minimum bracket the step model with a smoothed cutoff (e.g., logistic/erf in $\log M$ with tunable width) calibrated to published transfer-function/HMF results for DM–baryon scattering. In Sec. 3.1–3.2, add a robustness test showing how $dN/d\tau$, $N(>\tau)$, and $D_{\rm KL}$ change when the cutoff is smoothed and/or when a standard "half-mode mass" style mapping is used. Quantify how the inferred $\sigma(M_{\rm cut})$ (Sec. 3.3) shifts under these alternatives and temper claims that depend sensitively on the sharpness of the cutoff.

2. **The mapping between the phenomenological cutoff $M_{\rm cut}$ and the microscopic scattering parameter σ_0/m_χ (for $n=0$ and $n=-4$) is asserted but not transparently derived or reproducible (Sec. 2.2, Sec. 3.4). The manuscript does not specify: the perturbation calculation/transfer function used, how the suppression scale is defined (e.g., 50% suppression in $T(k)$, filtering mass, half-mode), how it is translated to a halo-mass scale at relevant redshifts, dependence on m_χ and cosmology, or what baryonic target (protons/electrons/neutral H) is assumed. Since the headline result is the σ_0/m_χ reach and the "4–5 orders of magnitude over CMB" comparison (Sec. 3.4, Conclusions), this opacity undermines credibility and prevents verification.**

   Recommendation: Add a dedicated subsection (end of Sec. 2.2 or new Sec. 2.5 / Appendix) detailing the full pipeline $\sigma_0/m_\chi \to$ transfer-function suppression $\to$ halo-mass-function modification $\to$ an effective $M_{\rm cut}$ used in HAYASHI. Explicitly state: assumed DM mass m_χ, scattering targets, cosmological parameters, the definition of the cutoff scale (e.g., $k_{1/2}$ where $T(k)=1/2$), and the conversion to a characteristic mass (e.g., half-mode mass). Provide a table/figure of $M_{\rm cut}$ versus σ_0/m_χ for representative redshifts (e.g., $z=8,10$) and both $n=0$ and $n=-4$. In Sec. 3.4, overlay or tabulate the corresponding $M_{\rm cut}$ values for the quoted σ_0/m_χ sensitivities and include an estimate of systematic uncertainty from the mapping.

3. **The thermal/spin-temperature and minihalo gas modeling is too compressed for the paper's strong conclusion that even extreme $f_{\rm cool}$ has negligible impact and that the $M_{\rm cut}$-induced *shape* change robustly breaks astrophysical degeneracies (Sec. 2.1–2.2, Sec. 3.1–3.2). It is unclear how HAYASHI computes T_s (collisional vs. Wouthuysen–Field coupling, radiation backgrounds), whether the regime $T_s \approx T_k$ is always realized over $z=7$–15, whether absorption is in the linear $\tau \ll 1$ regime, and whether $f_{\rm cool}$ is applied to IGM only or also to minihalo gas (where line profiles/central densities set τ). The current parameter set $M_{\rm cut}, f_{\rm cool}, \bar{x}_{\rm HI}$ may not capture astrophysical effects that also change the *shape* of $dN/d\tau$ (e.g., photoheating feedback reducing gas in low-mass halos, mass-dependent gas fractions, X-ray/Lyα fluctuations).**

   Recommendation: Expand Sec. 2.1 with an explicit (but concise) summary of the ingredients that determine τ in the IGM and minihalos: gas density/temperature profiles, how T_s is computed (collisions/Ly α coupling), what backgrounds are assumed, and the halo mass/redshift range driving the statistics. In Sec. 2.2 and Sec. 3.1, clarify exactly where $f_{\rm cool}$ is applied (IGM only vs also halos), whether it is redshift dependent, and report representative T_k and T_s ranges to support the "saturation" argument. To support degeneracy-breaking claims (Sec. 3.2), add at least one additional nuisance parameter that can preferentially suppress low-τ absorbers (e.g., a feedback/photoheating suppression mass, a low-mass gas-fraction suppression parameter, or a temperature floor) and show how it compares to changing $M_{\rm cut}$ in $dN/d\tau$ and in the Fisher constraints.

4. **Forecast/statistical methodology is under-specified and likely optimistic: the Fisher analysis (Sec. 2.4) appears to assume independent Poisson counts per τ-bin with $\mathrm{Var}(N_k)=N_k$, but the paper does not fully define the data vector (counts vs normalized PDF), τ-binning and τ-range, redshift binning/path length, or how $N_{\rm los}(z)$ enters N_k. Instrumental effects**

and analysis systematics that are central for low-τ features—finite spectral resolution, thermal noise, continuum fitting, line blending/confusion, completeness/detection thresholds in τ—are neglected, yet much of the claimed information is in the low-τ regime (Sec. 3.1–3.3). This makes $\sigma(M_{\rm cut})$ and the optimal $z \approx 8$–10 window potentially fragile (Sec. 3.3), and the resulting σ_0/m_χ reach likely too optimistic (Sec. 3.4, Conclusions).

Recommendation: In Sec. 2.3–2.4, explicitly define the observable: whether N_k are total absorber counts across all sightlines in (z, τ) bins (preferred for Poisson likelihood), or a normalized shape-only statistic; provide the τ-bin edges, τ-range, and any redshift binning/path length per sightline. Write $N_k(z) = N_{\rm los}(z) \times \lambda_k(z, \theta)$ (or equivalent) so the $N_{\rm los}$ scaling is explicit. Then incorporate a minimal observational realism layer: e.g., (i) a τ detection threshold/completeness curve, (ii) an effective reduction in usable $N_{\rm los}$, and/or (iii) an added fractional variance term per bin to represent non-Poisson systematics. Recompute (or at least bracket) $\sigma(M_{\rm cut})$ and the preferred redshift window under optimistic vs pessimistic thresholds, and soften conclusions where they depend on Poisson-limited low-τ performance.

5. **KL divergence is used as a central "intrinsic distinguishability" metric (Sec. 2.3, Fig. 3, Sec. 3.2) but the manuscript does not define the normalized distributions entering $D_{\rm KL}$. KL divergence requires proper probability distributions $p(\tau)$, $q(\tau)$ (or discrete probabilities p_k, q_k). If computed from $dN/d\tau$ or raw N_k without normalization, the interpretation changes (shape-only vs amplitude+shape), and the connection to detectability with finite counts is unclear.**

   Recommendation: In Sec. 2.3, add the explicit definition of $D_{\rm KL}$ and specify precisely how p and q are constructed from $dN/d\tau$ or binned N_k (including normalization over the chosen τ-range). If the intent is "shape-only," state that explicitly and show how amplitude information is treated separately. In Sec. 3.2, add a brief interpretive bridge: for representative $N_{\rm los}$ and total counts, what $D_{\rm KL}$ values correspond to a meaningful likelihood-ratio/$\Delta\chi^2$ separation? (A simple approximation or illustrative conversion is sufficient.)

6. **The assumed redshift evolution of background radio-loud sources $N_{\rm los}(z) = 10[(1+z)/8]^{-2.5}$ (Sec. 2.4) is a key driver of the forecast and the claimed optimal $z \approx 8$–10 window (Sec. 3.3), but is currently schematic and not clearly justified by a cited luminosity-function model, flux limit, or survey strategy. Likewise, the "4–5 orders of magnitude better than CMB" statement (Sec. 3.4, Conclusions/Abstract) is presented without a like-for-like comparison (model assumptions, m_χ, n, ionization history) and without plotting the referenced CMB limits alongside the forecast.**

   Recommendation: In Sec. 2.4, cite the provenance of the $N_{\rm los}(z)$ scaling (data or simulations) and add at least one optimistic and one pessimistic alternative (varying normalization and exponent, or tying $N_{\rm los}$ to a flux limit and a luminosity function). In Sec. 3.3, show how $\sigma(M_{\rm cut})$ and the optimal redshift shift under these alternatives. In Sec. 3.4, state explicitly which CMB limits are used (citations, m_χ, targets, n, and assumptions) and overlay them on the same σ_0/m_χ plot; then restate the improvement factor conditional on those assumptions, and temper the language in the Abstract/Conclusions if it only holds in optimistic source/noise scenarios.

Minor issues

1. Several plots/figures lack enough normalization/unit/definition detail to be quantitatively interpretable and reproducible (notably Fig. 1, Fig. 3, Fig. 4, Fig. 5 as referenced in the structured report). For example: units/normalization of $N(>\tau)$ and τ, how derivatives were taken for $dN/d\tau$, how $D_{\rm KL}$ was computed, and what exactly the Fisher curves assume are not always explicit in captions.

Recommendation: Strengthen figure captions to include (i) axis units and normalization (per sightline? per Δz? total survey counts?), (ii) binning and numerical differentiation method (for $dN/d\tau$), (iii) the precise definition/normalization used for KL, and (iv) Fisher assumptions (τ-range, binning, included parameters). Where possible, add uncertainty bands (Poisson-only and/or systematic bracketing) rather than only showing noiseless curves.

2. Instrumental/detection accessibility is not indicated on key distributions, especially at low τ where much of the information resides (Sec. 3.1–3.3; figures showing $dN/d\tau$ or $N(>\tau)$).

   Recommendation: Overlay illustrative detectability thresholds: e.g., a $\tau_{\rm min}$ corresponding to a representative spectral resolution/channel width and S/N, or show completeness curves. Alternatively, shade regions of τ that are likely inaccessible under plausible noise/continuum-fitting performance.

3. The Fisher setup uses ${M_{\rm cut}, f_{\rm cool}, \bar{x}{\rm HI}}(Sec.2.4), but\bar{x}$ as a single uniform parameter may be inadequate near the end of reionization: patchiness/topology and fluctuating ionizing/X-ray backgrounds can change absorber statistics in a nontrivial, potentially shape-changing way.

   Recommendation: Add a short discussion (Sec. 3.3 or Conclusions) of how patchy reionization and spatially varying backgrounds could impact $dN/d\tau$ and degeneracies with $M_{\rm cut}$, with citations to relevant 21 cm forest/reionization modeling literature. If feasible, include a simple bracketing scenario (e.g., two reionization histories or an added "patchiness variance" nuisance) to test stability of the main conclusions.

4. The text's statement about Poisson scaling is somewhat inconsistent/implicit: absolute uncertainties scale as $\sqrt{N_k}$, while fractional uncertainties scale as $1/\sqrt{N_k}$, and N_k should scale with $N_{\rm los}$ (Sec. 2.4).

   Recommendation: In Sec. 2.4, explicitly distinguish absolute vs fractional uncertainties, and write $N_k \propto N_{\rm los}$ so the scaling of Fisher information with $N_{\rm los}$ is transparent.

5. The focus on particular redshifts (e.g., $z = 10$ examples in Sec. 3.1 and Sec. 3.4) is not fully aligned with the paper's own statement that the optimal window depends on $N_{\rm los}(z)$ and may prefer $z \approx 8$–9 in some cases (Sec. 3.3).

   Recommendation: Briefly justify choosing $z = 10$ as a benchmark (e.g., literature convention or balancing intrinsic signal and source counts), and/or add one additional benchmark redshift within the preferred window ($z = 8$ or 9) to show that key qualitative statements about $dN/d\tau$ shape, $D_{\rm KL}$ trends, and $\sigma(M_{\rm cut})$ persist.

6. Cosmological parameters used for the HAYASHI runs and for the $M_{\rm cut}$–σ_0/m_χ mapping are not explicitly listed (Sec. 2.1–2.2), limiting reproducibility and making it hard to judge sensitivity to, e.g., σ_8 and n_s (important for low-mass halo abundance).

   Recommendation: Add a short paragraph in Sec. 2.1 stating the fiducial cosmological parameters used. Optionally comment qualitatively (or with a quick test) on how varying σ_8/n_s within current uncertainties would shift $dN/d\tau$ and inferred $M_{\rm cut}$ sensitivity.

7. A numerical consistency check flags ambiguous "percent reduction vs remaining fraction" statements (reported in the structured report as $50.2\%/47.0\%$ reductions not matching implied remaining fractions).

   Recommendation: Re-check the relevant text/annotations and ensure reductions vs remaining fractions are defined and computed consistently (and that figure annotations match the intended definition).

8. Some claims in Sec. 3.2 and the Conclusions that astrophysical uncertainties "typically rescale the amplitude" (and thus are easily distinguished from $M_{\rm cut}$ shape changes) are stronger than what is demonstrated with the current nuisance parameter set.

Recommendation: Soften/qualify this language unless additional nuisance parameters/tests are added. Explicitly state the scope of the robustness claim: e.g., “for the limited set of astrophysical degrees of freedom modeled here.”

Very minor issues

1. Notation inconsistency: $f_{\rm cool}$ is defined as $T_k/T_k^{\rm ad}$ (Sec. 2.2) but appears as $T_k/T_k^{\rm ud}$ in Sec. 3.1; also $f_{\rm cool}$ is called a “cooling fraction” though it is a temperature ratio and could exceed **1** if heating is allowed.

   Recommendation: Standardize to $T_k^{\rm ad}$ everywhere and either (i) state the intended domain (e.g., \$01\$ are allowed.

2. The cross-section parameterization $\sigma = \sigma_0 v^n$ does not specify whether v is dimensional or normalized by c, which determines the units of σ_0 (Sec. 2.2).

   Recommendation: State the convention for v (e.g., v/c) and the resulting units of σ_0; ensure this is consistent with the σ_0/m_χ values quoted in Sec. 3.4.

3. Minor presentation/typography inconsistencies (quotation marks around “fingerprint”, hyphenation such as “Coulomb-like”/“velocity-independent”, “21 cm” formatting, occasional formatting artifacts in headings, acronym re-definitions).

   Recommendation: Proofread for consistent typography and acronym definitions (CMB, CDM, KL) and ensure figure labels/legends are readable and consistent (including colorblind-safe styling where possible).

Key statements and references

- **• For a velocity-independent ($n = 0$) dark matter–baryon scattering cross-section, achieving a 10% suppression in the total number of 21 cm forest absorbers at $z = 10$ corresponds to a sensitivity of $\sigma_0/m_\chi \approx 2.13 \times 10^{-26}$ cm^2/GeV, while for a Coulomb-like ($n = -4$) interaction the same 10% suppression corresponds to $\sigma_0/m_\chi \approx 2.13 \times 10^{-43}$ cm^2/GeV; these projected sensitivities are four to five orders of magnitude stronger than current constraints from the Cosmic Microwave Background.**
- *Reference(s):* Cosmic Microwave Background
- **• A modest cutoff in the halo mass function of $M_{\rm cut} = 10^4\, M_\odot$ induced by dark matter–baryon scattering reduces the total number of 21 cm forest absorbers by $\approx 50.2\%$ at $z = 7$ and 47.0% at $z = 10$, while stronger interactions with $M_{\rm cut} \geq 10^6\, M_\odot$ eliminate over 92% of absorbers across the studied redshift range, demonstrating that kinematic suppression of minihalo formation overwhelmingly dominates over thermal cooling effects in shaping the signal.**
- *Reference(s):* (none)
- **• Using a Fisher matrix forecast for the parameter set ${M_{\rm cut}, f_{\rm cool}, \bar{x}{\rm HI}}thatincorporatesaredshift−dependentevolutionofthenumberofbackgroundradiosourcesN) \approx 1127~M_\odotatz = 7 and\approx 5018~M_\odotatz = 15, withanoptimalobservationalwindowidentifiedatz \approx 8–10 where the trade-off between intrinsic signal distinctness and statistical power is most favorable.}(z) = 10[(1+z)/8]^{-2.5}, the marginalized 1-σ uncertainty on the cutoff mass is found to be $\sigma(M_{\rm cut**
- *Reference(s):* (none)

Mathematical consistency audit

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

Maths relevance: light

The paper's analytical content is primarily definitional and forecasting-based: it introduces a phenomenological halo-mass cutoff M_{cut}, a cooling parameter f_{cool}, uses a Poisson-count Fisher matrix (Eq. 1), an empirical redshift evolution for the number of lines of sight $N_{\mathrm{los}}(z)$ (Eq. 2), and invokes KL divergence as a shape distinguishability metric. Several central mappings/definitions needed to verify key conclusions (KL definition/normalization; $M_{\mathrm{cut}} \leftrightarrow$cross-section relation) are not provided in the text shown.

Checked items

1. ✔ **Cross-section velocity power law** (Sec. 2.2, p.3)
    - **Claim:** The DM–baryon elastic scattering cross-section can be modeled as $\sigma = \sigma_0 v^n$, focusing on $n = 0$ and $n = -4$.
    - **Checks:** notation consistency, dimensional/units sanity
    - **Verdict:** PASS; confidence: medium; impact: minor
    - **Assumptions/inputs:** v denotes the DM–baryon relative velocity under some convention (not specified)., σ_0 is a normalization constant whose units depend on the choice of v units.
    - **Notes:** Form is internally consistent, but the units of σ_0 cannot be checked without specifying whether v is dimensionless or has physical units.

2. ✔ **Cooling fraction definition** (Sec. 2.2, p.3)
    - **Claim:** Thermal impact is parameterized by $f_{\mathrm{cool}} = T_k/T_k^{\mathrm{ad}}$.
    - **Checks:** definition consistency, dimensional/units sanity
    - **Verdict:** PASS; confidence: high; impact: minor
    - **Assumptions/inputs:** T_k and T_k^{ad} are both kinetic temperatures defined at the same redshift/epoch.
    - **Notes:** Dimensionless ratio is consistent. Domain/interpretation (cooling-only vs possible heating) is not specified.

3. ✔ **Fisher matrix for Poisson counts** (Eq. (1), Sec. 2.4, p.4)
    - **Claim:** $F_{ij} = \sum_k \frac{1}{N_k}\left(\frac{\partial N_k}{\partial \theta_i}\right)\left(\frac{\partial N_k}{\partial \theta_j}\right)$ for binned absorber counts N_k.
    - **Checks:** algebra/derivation sanity, assumption check, symbol consistency
    - **Verdict:** PASS; confidence: medium; impact: moderate
    - **Assumptions/inputs:** Bins k are statistically independent., Counts in each bin are Poisson with $\mathrm{Var}(N_k) = N_k$., N_k represents expected counts at fiducial parameters (typical Fisher convention).
    - **Notes:** Equation is standard under Poisson assumptions, but the paper does not explicitly state whether N_k is expected or observed, nor how N_{los} enters N_k. These omissions reduce verifiability but do not by themselves contradict Eq. (1).

4. ✔ **Lines-of-sight evolution model** (Eq. (2), Sec. 2.4, p.4)
    - **Claim:** Number of lines of sight evolves as $N_{\mathrm{los}}(z) = 10((1+z)/8)^{-2.5}$.
    - **Checks:** dimensional/units sanity, symbol consistency, limiting/sanity check
    - **Verdict:** PASS; confidence: high; impact: minor
    - **Assumptions/inputs:** N_{los} is dimensionless and intended as an effective number of usable background sources/LOS.

- **Notes:** Dimensionless and monotone-decreasing with z as intended. No internal inconsistencies found.

5. ✖ **Poisson scaling statement vs Poisson statistics** (Sec. 2.4, p.4)
   - **Claim:** Statistical uncertainty on N_k is Poissonian, scaling with $1/\sqrt{N_{\rm los}}$.
   - **Checks:** consistency with definitions, algebraic/statistical scaling sanity
   - **Verdict:** FAIL; confidence: high; impact: minor
   - **Assumptions/inputs:** N_k is a count aggregated over $N_{\rm los}$ independent lines of sight., Per-LOS expected counts are approximately constant so $N_k \propto N_{\rm los}$.
   - **Notes:** For Poisson counts, absolute uncertainty scales as $\sqrt{N_k}$ (thus typically $\propto \sqrt{N_{\rm los}}$ if $N_k \propto N_{\rm los}$), while fractional uncertainty scales as $1/\sqrt{N_k}$ (thus $\propto 1/\sqrt{N_{\rm los}}$). The text appears to conflate absolute and fractional uncertainty.

6. ⚠ **KL divergence applicability to $dN/d\tau$** (Sec. 2.3 and Sec. 3.2, pp.3–5; Fig. 3 caption p.7)
   - **Claim:** $D_{\rm KL}({\rm CDM} \parallel M_{\rm cut})$ quantifies distinguishability between the differential optical depth distributions.
   - **Checks:** definition consistency, normalization/constraints, missing-steps audit
   - **Verdict:** UNCERTAIN; confidence: high; impact: critical
   - **Assumptions/inputs:** A KL divergence is computed between two probability distributions over τ (or over bins)., The 'differential optical depth distribution' is converted into a normalized distribution (not shown).
   - **Notes:** KL divergence requires normalized probabilities; $dN/d\tau$ is not inherently normalized (it is a differential count/rate). The paper does not provide the formula for $D_{\rm KL}$ or the normalization procedure, so the claimed monotonicity/interpretation cannot be verified analytically.

7. ⚠ **Sharp cutoff implementation in halo mass function** (Sec. 2.2, p.3)
   - **Claim:** Structure suppression is modeled by introducing a sharp cutoff in the halo mass function below $M_{\rm cut}$ (no halos for $M < M_{\rm cut}$).
   - **Checks:** definition consistency, missing-steps audit
   - **Verdict:** UNCERTAIN; confidence: medium; impact: moderate
   - **Assumptions/inputs:** Baseline halo mass function exists in the HAYASHI framework., Cutoff is implemented as a Heaviside-like truncation.
   - **Notes:** Mathematically plausible, but the exact functional form (e.g., dn/dM multiplied by $\Theta(M - M_{\rm cut})$) and any smoothing are not specified, so the analytic consequences for $dN/d\tau$ cannot be checked from the provided text.

8. ⚠ **Mapping $M_{\rm cut}$ to σ_0/m_χ constraints** (Sec. 2.2 (statement) and Sec. 3.4 (results), pp.3 and 8)
   - **Claim:** $M_{\rm cut}$ is directly related to σ, enabling constraints on σ_0/m_χ from sensitivity to $M_{\rm cut}$.
   - **Checks:** missing-steps audit, symbol/definition consistency
   - **Verdict:** UNCERTAIN; confidence: high; impact: critical
   - **Assumptions/inputs:** There exists a deterministic relation $M_{\rm cut}(\sigma_0/m_\chi, n, z, ...)$.
   - **Notes:** No equation or derivation is provided for $M_{\rm cut}(\sigma_0/m_\chi)$. Without this mapping, the translation in Sec. 3.4 is not verifiable on symbolic/analytic grounds.

Limitations

- Only the equations and definitions explicitly present in the provided PDF text (Eq. (1) and Eq. (2) plus a few inline definitions) can be audited; key mathematical definitions (e.g., explicit KL divergence formula, normalization) are absent.

- Figures are referenced for derivatives and trends, but no underlying analytic expressions for $N(>\tau)$, $dN/d\tau$, or their dependence on parameters are provided, preventing step-by-step derivation checks.
- The paper's central physical-to-phenomenological mapping (DM scattering → power-spectrum suppression → M_{cut}) is asserted but not written as mathematics in the provided text, so internal consistency of that chain cannot be assessed.

Numerical results audit

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

10 candidate numerical checks were executed: 8 PASS and 2 FAIL. Passed items include a monotonicity evaluation of $N_{\mathrm{los}}(z)$ from Eq. (2), exact percent-to-fraction conversion for 0.002%, parameter-range sanity for $f_{\mathrm{cool}} = 0.02$, internal consistency/ordering of quoted Fisher uncertainties, scientific-notation parsing for two cross-section sensitivities, and a 17-orders-of-magnitude comparison derived from those sensitivities. The only failures occurred in percent-reduction-to-remaining-fraction conversions for two stated reductions ($z = 7$ and $z = 10$).

Checked items

1. ✔ **C1_Nlos_powerlaw_eq2** (Page 4, Eq. (2))
   - **Claim:** The number of lines of sight is modeled as $N_{\mathrm{los}}(z) = 10 \times ((1+z)/8)^{-2.5}$.
   - **Checks:** formula_evaluation
   - **Verdict:** PASS
   - **Notes:** Computed $N_{\mathrm{los}}(7) = 10.0$, $N_{\mathrm{los}}(10) = 4.510692841903824$, $N_{\mathrm{los}}(15) = 1.7677669529663689$; monotonic decreasing and positive.
2. ✖ **C2_absorber_reduction_Mcut1e4_z7** (Page 5, Section 3.1 (text))
   - **Claim:** A modest cutoff of $M_{\mathrm{cut}} = 10^4\,M_\odot$ reduces the absorber count by 50.2% at $z = 7$.
   - **Checks:** percent_to_fraction
   - **Verdict:** FAIL
   - **Notes:** Computed remaining fraction $r = 1 - 50.2/100 = 0.498$; the check's reported expected_remaining_fraction was 1.0, producing a mismatch.
3. ✖ **C3_absorber_reduction_Mcut1e4_z10** (Page 5, Section 3.1 (text))
   - **Claim:** A modest cutoff of $M_{\mathrm{cut}} = 10^4\,M_\odot$ reduces the absorber count by 47.0% at $z = 10$.
   - **Checks:** percent_to_fraction
   - **Verdict:** FAIL
   - **Notes:** Computed remaining fraction $r = 1 - 47.0/100 = 0.53$; the check's reported expected_remaining_fraction was 1.0, producing a mismatch.
4. ✔ **C4_cooling_fractional_change_bound** (Page 5, Section 3.1 (text))
   - **Claim:** Across all redshifts studied, the maximum fractional change in the total number of absorbers due to cooling remains below 0.002%.
   - **Checks:** threshold_interpretation
   - **Verdict:** PASS
   - **Notes:** 0.002% converts exactly to a fraction of 2×10^{-5}.
5. ✔ **C5_extreme_cooling_value** (Page 5, Section 3.1 (text))
   - **Claim:** Even extreme cooling ($f_{\mathrm{cool}} \to 0.02$) produces no discernible change in $N(>\tau)$.
   - **Checks:** parameter_range_sanity
   - **Verdict:** PASS

- **Notes:** Sanity check only: $f_{\text{cool}} = 0.02$ satisfies $0 < f_{\text{cool}} \leq 1$.

6. ✔ **C6_Fisher_uncertainty_growth_ratio** (Page 7, Section 3.3 (text) and Page 8, Figure 4 caption)
   - **Claim:** $\sigma(M_{\text{cut}}) \approx 1127\ M_\odot$ at $z = 7$ and $\approx 5018\ M_\odot$ at $z = 15$ (uncertainty increases monotonically with redshift).
   - **Checks:** ratio_and_monotonicity_from_reported_points
   - **Verdict:** PASS
   - **Notes:** Positivity and ordering hold ($5018 > 1127$). Computed ratio $5018/1127 = 4.452528837622006$.
7. ✔ **C7_Fisher_uncertainty_z10_vs_z7** (Page 7, Section 3.3 (text))
   - **Claim:** $\sigma(M_{\text{cut}}) \approx 1907\ M_\odot$ at $z = 10$ while $\sigma(M_{\text{cut}}) \approx 1127\ M_\odot$ at $z = 7$.
   - **Checks:** ordering_check
   - **Verdict:** PASS
   - **Notes:** Ordering is consistent with monotonic increase in redshift: $1907 > 1127$ for $z = 10 > 7$.
8. ✔ **C8_sigma0_over_mchi_n0_value** (Page 8, Section 3.4 (text))
   - **Claim:** For velocity-independent ($n = 0$), a 10% suppression measurement at $z = 10$ corresponds to $\sigma_0/m_\chi \approx 2.13 \times 10^{-26}\ \text{cm}^2/\text{GeV}$.
   - **Checks:** scientific_notation_parse
   - **Verdict:** PASS
   - **Notes:** Parsed numeric value matches 2.13×10^{-26}.
9. ✔ **C9_sigma0_over_mchi_nneg4_value** (Page 8, Section 3.4 (text))
   - **Claim:** For Coulomb-like ($n = -4$), the sensitivity reaches $\sigma_0/m_\chi \approx 2.13 \times 10^{-43}\ \text{cm}^2/\text{GeV}$ (at $z = 10$ in the same sentence context).
   - **Checks:** scientific_notation_parse
   - **Verdict:** PASS
   - **Notes:** Parsed numeric value matches 2.13×10^{-43}.
10. ✔ **C10_orders_of_magnitude_claim_from_two_numbers** (Page 8, Section 3.4 (text))
    - **Claim:** Compare the two quoted sensitivities: 2.13×10^{-26} vs $2.13 \times 10^{-43}\ \text{cm}^2/\text{GeV}$; they differ by 17 orders of magnitude.
    - **Checks:** orders_of_magnitude_difference
    - **Verdict:** PASS
    - **Notes:** Computed $\log_{10}$ ratio equals 17.0 exactly given equal mantissas.

Limitations

- Only parsed text was available; figures are images without machine-readable numeric tables, so plot-based quantitative checks were not feasible.
- Many central results (absorber counts, KL divergence values, Fisher derivatives) rely on model outputs not enumerated numerically in the text, preventing recomputation.
- Feasible fast checks were limited to arithmetic conversions, inequality/monotonicity assertions from stated point estimates, and scientific-notation parsing.