

Signatures of Suppressed Matter Clustering revealed by Fast Radio Bursts

Corresponding Author: Ms Kritti Sharma

This manuscript has been previously reviewed at another journal. This document only contains information relating to versions considered at Nature Astronomy.

This file contains all reviewer reports in order by version, followed by all author rebuttals in order by version.

Version 0:

Reviewer comments:

Reviewer #1

(Comments for the Author)

This is a review of Sharma et al. "Signatures of Suppressed...Bursts."

This is a very long, dense and highly technical paper on a very timely and important topic, namely the development of FRBs as a new tool for cosmology. The paper's importance does not lie so much in the results that it reports. Indeed, my impressions from the opening paragraph were that an improvement in an unspecified prior and ruling out extreme models at the 2σ level are not major leaps forward. Rather, the value of the paper lies in showing how such an analysis can be done, and how FRBs will be very powerful in the future when there are large samples (although getting redshifts for 10^4 FRBs may be problematic — an issue not admitted or discussed by the authors, see below). The bulk of the important and novel analysis is done and explained in the Methods of the paper, which is presently very dense and very long.

I think of Nature Astronomy as a journal meant to showcase results that are of a somewhat broader impact and comprehensible to a broader audience of astronomers not just in the field reported on in the paper. I leave it therefore to the Editor's discretion regarding whether the paper is suitable for Nature Astronomy, given that the results are not (yet) very constraining, but given that the techniques being used are novel, interesting, and are pushing the field forward in a very important and impressive manner.

Text:

Paragraph 1 refers to scales via k , whereas paragraph 2 in Mpc. It would help people like me to not jump back and forth; actually my preference would be for each mention, provide both.

lines 70, 81

"X-ray gas mass fractions" is a bit of a word salad...

how about "cluster gas mass fractions measured through X-ray observations"

line 73

enables —> enable

lines 96-97 “Hydrodynamical simulations...”

The authors have just explained why these are presently not well constrained at the relevant scales so how can we believe this statement?

line 98: “does not impact current measurements”

Can the authors summarize here in a few words why? This is a very important statement so it is a bit frustrating to have to go to Methods to find out why

lines 100-101: the statement “Likewise, selection effects” is an important one and relies entirely on the single study by the same authors, which unless I’m mistaken, assumes FRBs all detected with the same selection biases ie with a single instrument. Here, the authors cobble together a sample from many instruments. I gather the argument is that if selection biases don’t matter for any one instrument then there should be no problem combining. I think if that is the reasoning being used, it needs to be stated explicitly. In the limit where one FRB survey is very insensitive and sees only nearby FRBs and another is vastly more sensitive and sees FRBs to high redshift, even if selection biases don’t matter for either, it’s not obvious to me combining them doesn’t pose a problem.

line 108: M_{200} is not defined. What does the 200 mean?

Fig 1:

For plots b and c, in two cases top one-third and two-thirds contours are plotted, and for two others, 86th percentile contours are plotted, and it is not stated what is plotted for the DM case. First, the latter needs rectifying, but more generally, it is hard to robustly compare all of these when such different contours are plotted.

Also, the y-axis of panel c is not properly explained. Please clarify what this quantity is. I also found the dual method for specifying M_0 (legend with 3 line styles and color bar) confusing. I guess it is meant to show how things vary with M_0 , so the color bar should suffice, but the authors wanted to also specify clearly for the round numbers. To make it a bit less confusing, I suggest that the lines used for the round numbers be colored according to where they are in the color bar.

Lines 124-125

CHIME/FRB is also undergoing an upgrade that could increase their detection rate by a factor of 3-5, and together with its Outriggers, could also make it a player in this game.

Lines 128,129: CHIME should be CHIME/FRB to distinguish from the cosmology experiment. Also please fix all occurrences in Extended Table 1.

Fig 2: The title of this figure confused me. The plot shows DM vs z so I’d have expected that should be the title. I don’t see how the plot is telling me clearly how baryon clustering is probed. In the caption it says “quantified by the degree of clustering of baryons” is also confusing. Perhaps the authors mean that the variance quantifies the degree of clustering rather than than is being quantified by it?
And again CHIME --> CHIME/FRB

line 131: provide a reference for BCEMU7 halo model

line 134: specify what NFW is for non-specialists

line 139: please provide equation numbers rather than just saying “key equations in the Methods”

Paragraph starting at line 140:

The text here suddenly becomes impenetrable to me.
I feel the authors need to do a better job at explaining what they are constraining for a journal like Nature Astronomy that is meant to be comprehensible to people outside this narrow field.
e.g. why is “with a 224% reduction in the variance of the 1 dominant eigenmode relative to prior” important?
and what is a “Karhunen-Loeve mode” ?

lines 152-153: the number of significant figures quoted for DM_{host} and sigma are too high. Looks like at most you should be quoting 3 sig figs.

Also and more importantly, the value for DM_{host} is somewhat higher than that inferred by Shin et al. 2023 based on a very large, uniform sample of FRBs albeit without redshifts. The authors should comment on the difference.

Fig 3:
Left — I guess the hatched region are the FRBs? The legend in the right plot explicitly says this; why isn't that also done in the left plot?

Fig 4:
Re the sentence “The entropy...” I don't see where and in which plot this is shown and suspect it is not, so it's unclear to me why the statement is made in the plot caption here.

line 172
Commenting on a 0.6-0.8sigma preference is a stretch. The final sentence in this paragraph “Together, these results...” is extremely speculative and not based on very strong evidence. This should be better emphasized so as not to give the impression the authors have actually demonstrated any of this. I realize they say “suggest” but I find that to be a stretch given the actual results.

line 195
Similar to above, the word “discrepant” seems an overstatement. There is mild tension here, not discrepancy.

line 218-219
There is some concern that these surveys will be limited by ability to measure redshifts for a large population especially for faint hosts where very large telescopes are needed, even at moderate redshifts. That needs to be acknowledged and could ultimately be the limiting factor in achieving the goals of this study.

Methods

The authors need to provide some preamble at the start of the “Sensitivity Analysis” section to help the reader understand to what this pertains in the main text, and what the goal of the section is.
Indeed, every Methods subsection should have this.

line 565: As asked above, Shin et al. 2023 find a much lower mean DM_{host}.
Also the reference provided here [81] doesn't seem appropriate for this statement.

(Remarks on code availability)

Reviewer #2

(Comments for the Author)

The astrophysical feedback produced by the star-formation processes and/or AGN activity are key ingredients in the hydrodynamical simulations of the galaxy formation and evolution. Moreover, they are considered to be responsible for the redistribution of cosmic baryons between the diffuse intergalactic medium (IGM), multiphase circumgalactic medium (CGM) and the interstellar medium (ISM). Subsequently, due to its effect on the matter clustering, feedback directly affects any

cosmological measurements. Though a plethora of observational techniques have been employed by the community in order to study and quantify the feedback in data at $z < 2$, all of these methods require specific assumptions about the density, metallicity and temperature of the probed gas that bias the interpretation of the results. In that respect, the FRB dispersion measure constitutes a unique probe of all ionized baryons irrespective of such assumptions.

This work is a continuation of authors' previous theoretical efforts to quantify the sensitivity of the FRB DMs to the matter power spectrum suppression, feedback and halo models. However, this is the first attempt to apply the developed model and inference algorithm to the observed sample of FRB DMs. Moreover, authors assess the sensitivity of the FRB-based constraints and compare it to those by other astrophysical probes. The authors utilize a robust and well-defined statistical algorithm with careful treatment of uncertainties. Additionally, this work lays a foundation for the future usage of FRBs for cosmological applications. While this manuscript definitely has merit and deserves a publication in Nature Astronomy, first, before proceeding, the authors must address some of the comments/concerns outlined in what follows.

[Comments/concerns/suggested improvements]
(numbers (if any) indicate specific lines in the text)

[Main text]

1. 104, 108, 119: variance should be σ^2 , not just σ
2. 131: you need to reference the BCEmu7 model - at this point the reader has no idea what it is.
3. 153: this is not a variance, but standard deviation.

[Methods/Sensitivity Analysis]

1. 430-432: Please, make it explicit that results for other probes presented in Fig. 1 are computed via eq.(1). This section generally lacks references to Fig.1 of the main text.
2. 435: What is the effect of fixed source redshift distribution on the analysis? The authors should expand their statement and comment on any potential effects on the sensitivity function.
3. 448: The authors have not indicated the number of FRBs that was used for the sensitivity prediction, neither in this section nor in the main text (this includes the Fig. 1 caption).
4. 454: Did the authors mean standard deviation? Variance (assuming of a log-normal distribution) in DM_host of 100 pc/cc seems to be rather tiny implying std.dev ~ 10 pc/cc. Please clarify and give reference for this number, and be explicit of the quantities and why the specific values have been chosen. I think this value is quoted correctly in line 819.

[Sensitivity synopsis of baryon probes]

1. The text in this section will benefit from cross-referencing with Fig. 1. Please include description of the color/format of the contours in Fig.1 that correspond to the described measurement. In addition, authors should consider including a short description of the FRB DM sensitivity as well for consistency.

[Halo mass sensitivity of matter power spectrum suppression]

1. Authors did not explain how the corresponding power spectra were calculated. Presumably, this is related to the discussion in section "Halo Model Approach", but no references are made. Please make it clear how the power spectra in panel C of Fig. 1 is computed.

[Inference Framework]

1. 495: "Equations 9 and 10 establish that both, the mean DM_exgal-z relation and the scatter in DM_exgal at a given z carry cosmological and astrophysically valuable information.". Technically these equations describe DM_cosmic, not DM_exgal. Please, fix. Overall, this statement is too generic and unnecessary as it is clear that any detected astrophysical quantity carries astrophysically valuable information, unless, of course there is some information that the authors assume to be astrophysically useless.
2. Overall, the PDF in equation 8 establishes the connection between the DM_cosmic and gas/electron power spectrum of the intervening halos via equations 9 and 10. In turn, the power spectrum is related to various feedback processes. However, the second term under the integral in eq.8, describing DM_host PDF, does not depend on the feedback by construction. FRB host galaxies, in principle, are subjected to the same feedback processes and, therefore the corresponding PDF can (should) similarly be written as a function of the electron power spectrum. Authors should explain why they chose to assume no feedback dependence for the host galaxies or model it through their numerical scheme by estimating the electron power spectra for the hosts.

[Halo model]

1. 511-512: «...two halo model prescription...» is a bit ambiguous: do you mean a two halo model for the power spectrum analysis or two models to describe a halo? Obviously, the latter, but it helps to rephrase for better clarity.
2. 519,522: check the formatting of the references.
3. 529: "Due to its inherent limitations, we use this model primarily as a cross-check against our baseline results obtained with BCEMU model". It is not clear how the authors perform this cross-check between seemingly completely different inferred parameters: temperature of the AGN-influenced gas on one hand, and the halo mass and geometric terms on the other. Only the parameters governing the DM_host PDF are the same between two models. Please, explain.
4. 551: for consistency, this should be variance, i.e., σ^2

[Plausible DM_host redshift evolution]

1. 559: "However, the apparent shift in posteriors between the fiducial and extended FRB sample may be indicative of a

potential $\{\rm DM_{host}\}$ redshift evolution.” This statement is misleading as the apparent shift can also indicate the presence of some bias once extra FRB sightlines are included/excluded. I suggest the authors do a simple test by splitting the fiducial sample into two redshift bins and repeating the analysis. Moreover, to control possible bias one can do the jackknife analysis repeatedly with $N-1$ sightlines in each sample, excluding a random sightline from each run.

2. 561: define ‘extreme $\rm DM_{host}$ evolution’;

3. 564: define ‘local FRB samples’;

4. 566: missing citations: e.g., Mannings+2021, Khrykin+2026

5. 569-570: missing citations: e.g., James+2022a

6. 572: one can argue that instead of ‘robust’ the word ‘insensitive’ can be used.

[Constrained Baryonic Feedback Parameters]

1. Not surprisingly, parameters of the full 7-parameter halo model are not constrained with the current dataset. Moreover, the halo models with reduced number of parameters are also unconstrained as the respective posterior PDFs span all parameter ranges with comparable amplitudes. I would argue that only the BCEmu1 model has (at this point) the constraining power, which is indeed backed by your DOF analysis. Have authors considered less sophisticated analytical models? It would be beneficial if authors discuss why other analytical models are worse (if they are) than the BCEmu set and shouldn’t be considered.

2. While the FRBs clearly have very limited sensitivity to the specific parameters of a particular BCEmu halo model, I suggest the authors also show how other probes can/cannot recover BCEmu model parameters. Such comparison will help to understand what FRBs bring to the table.

[Stellar-to-Halo Mass Relation Prior]

1. 667: it would be useful if authors provide a reference for this calculation;

2. 692: add reference to equation 15 to remind the reader about the functional form of the $f_{\star}$ dependence on SHM and its normalization. In addition, the discussion in this and next sections revolves around equation 20, which is, however, defined only in the end of the next section. The authors might consider re-arranging the text such that the baryon budget equation appears earlier. This can improve the flow of the discussion for the reader and clearly show how SHM changes $f_{\rm diffuse}$ as well.

[Our feedback measurement in the landscape of existing studies]

1. The authors should explain how different baryon fraction are calculated. Presumably, this is done by taking the marginalized posterior PDFs from MCMC fits to BCEmu4 or BCEmu7 models and plugging these numbers back into the corresponding halo models. However, the discussion lacks this important information.

[Constraining Power across Scales]

1. 748-750: discussion explains two different mocks in terms of presence/absence of the variance in $\rm DM_{cosmic}$, however the labels in sup.fig.2 describe the mocks in terms of extreme/no feedback. For clarity, I suggest the authors expand the labels of the panels in the top row of sup.fig 2 by adding « $\sigma_{\rm cosmic}^2 = 0$ » and «high $\sigma_{\rm cosmic}^2$ », respectively.

2. 749: please be mathematically precise and define ‘randomly high variance’.

3. KL divergence test: the authors do not specify which «data» they are discussing (one of the mock samples, fiducial observed sample or extended observed sample? Instead it is simply called ‘observed data’). Please describe it. Does the result of the KL divergence test change if fiducial vs extended samples are used?

4. In principle, the KL divergence test is not related to the mocks described in the first paragraph of this section. The authors, therefore, might consider splitting the discussion.

5. Likewise, it is not clear what BCEmu model is referenced when authors are quoting $D_{\rm KL} \sim 15 \sim \{\rm bits\}$ at $k \sim 1 \sim \{\rm Mpc^{-1}\}$. Does this information gain correspond to all halo models?

6. KL divergence: the information gain is typically quoted in bits if the base 2 logarithm is used. Eq. 21 seems to imply the natural logarithm instead, which gives the units of ‘nats’. If that is indeed the case, the quoted $D_{\rm KL}$ value must be converted from nats to bits. Please double check the number and correct if needed. Also, in Figure 4 the quoted value is 8 bits - why?

7. The bottom row of Sup.Fig.2, showing the gas mass fraction ratios, is not really discussed in the text.

[Survey specifications]

1.812-816: please add reference for the quoted parameter values. The projected DSA performance values should also be referenced.

2.817: Likewise, the line ‘representative of next-generation FRB surveys’ should be backed up by a reference or calculation.

[Sensitivity forecast]

1. 827: the authors should explain the chosen $\rm DM-M$ approximation and the choice of the scaling relation between N and M .

[Fisher forecast]

1. What is the number of FRBs that were used in 6×20 point analysis? This might give an idea if when such sensitivity can be achieved.

[Figures]

Figure 1: please add labels «a», «b», and «c» to the panels if you reference them.

Figure 2: while the requirement for the PATH probability is indicated ($P_{\rm host} > 0.9$), the PATH paper is not referenced. Please add it either here or in the corresponding discussion.

Figure 3: please add labels «a», «b», and «c» to the panels if you reference them.

Extended Data Figure 1: σ_{cosmic}^2 is variance, please correct the mathematical notation.

Extended Data Figure 2: same as previously — be careful with mathematical notation for variance/std.deviation. The last statement is somewhat questionable for parameter $\log M_c$.

Extended Data Figure 3: the colour chosen for the prior is very hard to distinguish from the white background. I suggest you choose a different colour scheme.

Extended Data Figure 4: please add labels «a», «b», and «c» to the panels if you reference them. Similar to previous figure, the colour for prior can be hard to see on some computers.

Supplementary Fig.2: the caption references «panel a», but it is not indicated on the actual plot. Same is true about panels b and c. What's the meaning of different hatches used in the panels in the middle row? Consider making them homogeneous across the panels.

Supplementary Fig. 4: f_{baryons} have not been introduced anywhere in the text. Please, define it.

Supplementary Fig. 5: please add labels «a», «b», and «c» to the panels if you reference them.

(Remarks on code availability)

The code archive is supplied with an extensive README file, explaining the main functions and how-to compute different quantities. While I have not reproduced the results of the manuscript myself, the code covers all calculations described in the text. Specific functions somewhat lack description.

Reviewer #3

(Comments for the Author)

This manuscript is the latest, and easily most detailed, work that uses fast radio bursts (FRBs) to study the distribution of ionised, baryonic matter in the Universe. This topic addresses many related questions regarding galaxy formation and cosmology - from the σ_8 tension, to feedback processes at different mass scales, to name a few. FRBs offer a unique probe of this physics, since their dispersion measure (which traces the column density of ionised gas along their line of sight) is equally sensitive to low and high values, whereas other probes typically have sensitivity scaling with density or density squared. This work uses a sample of 109 FRBs from a variety of instruments to perform inferences on this distribution, finding halo gas fractions broadly consistent with other measures, such as X-ray observations, and studies based on the Sunyaev-Zeldovich effect. Their results are consistent with other measures, and do not offer significantly improved constraints. However, their methodology is extremely interesting, and links FRB observables to more traditional cosmological methods in a far more direct way than some other works. Ultimately, this is a very technically interesting manuscript, and it absolutely deserves to be published, but I am unsure whether the format of Nature Astronomy is best suited to this kind of work.

Before proceeding to my specific comments, I'll expand on the above.

This manuscript builds a model for the cosmological distribution of ionised gas, based explicitly on models for the gas distribution in Galactic halos. It then fits this to FRB dispersion measure (DM) data as a function of model parameters - in particular, the variance in DM as a function of redshift is the key observable. As the authors explain (very well, I might say), this gives the model one primary degree of freedom for describing the shape of the cosmological DM contribution, being a linear combination of model parameters governing the mass-dependence of the shape of galaxy halo profiles. From a broad perspective, the work is then performing the same kind of fitting to data as Refs 15, 16, and 74. The former two works effectively model variance through distributing gas between halos and the intergalactic medium, while the latter uses a purely parametric description of the DM variance without linking this to physical details.

This current work improves on the above studies in that it gives highly detailed and well-motivated bottom-up descriptions of how gas physics - from star-formation to feedback - leads to a predicted power spectrum of ionised gas, and hence, how that leads to variance in variations in DM for a given z . It therefore adds significantly to the current literature by describing how parameters more familiar to those who study galaxy evolution and cosmology influence FRB observables, and hence acts as a guide to the FRB community on how the latter can influence the former. However, at a fundamental level, it is still turning a single dial (that of variance in FRB dispersion measure) and matching that to observables. While the data used to constrain the models is a larger set than that of previous studies, this is not a great leap forward, and the primary interest comes from the details found in the methods.

To be clear, the results of this paper are still useful - not all the gas probed by FRBs can be probed by other measures - but fundamentally do not add to existing constraints. What the paper does do extremely well is to analyse in detail both the parameter space that FRBs do constrain, and the sensitivity of these constraints to variation in what might be called nuisance parameters. I found the details explored in the Methods section to be extremely interesting, and of high relevance to those in FRB DM studies - the authors are absolutely to be congratulated on their excellent analysis.

In this sense, I found the paper ill-suited to the format of a Nature journal, which focusses on the result much more so than the method. The methodology is of primary interest here, and the methodology is better explained in traditional article format (and in any case, needs to be better explained here - see below). However, the authors try to make a bit too much of too little - the 0.6-0.8 σ significance of larger gas densities within R_{200} compared to x-ray models takes lines 160-179, when even that 0.6-0.8 σ conclusion is tied to other aspects of their model (as a reminder, the authors effectively vary one parameter). Then again, the methodology itself is very good. I leave it to the Editor to make a final call on whether or not this is suitable for Nature Astronomy.

Below, I make specific comments on the manuscript. Most are relatively minor details - the main comment is to give a better overview of the method in the main body.

MAIN MANUSCRIPT

In the main body of the manuscript, the authors jump rather quickly from the very broad overview of the topic - that of halo studies and how FRBs help - to the very specific, e.g. the exact parameters used to describe their halo model. However, there is no broad overview of the principle of their analysis, which is (ordering could be changed, depending on phrasing)

- develop a halo model
- use this to predict the electron power spectrum, P_{ee}
- use this to model the variance of the cosmological component of the DM budget of FRBs
- Include Galactic and host DM contributions
- fit to the $P(\text{DM}|z)$ distribution of data from 109 FRBs

I found I had to read between the lines to determine this, and even these basics of their approach wasn't clear until delving into the Methods. This will likely be seen as stating the obvious by the authors, but I believe it is very important to avoid confusing readers. I think this requires re-working some aspects of the paragraph 126-139 to really ensure this broad overview of the method is made clearer.

The terminology "Stage IV experiments" is specialised, and should be expanded upon in the main body

Line 118-120. FRBs probe "FRBs offer... and probes electron distributions on scales $0.1\text{-}3\text{ h Mpc}^{-1}$ ". This needs a reference.

l178 Evidence for 'with feedback processes likely becoming progressively more efficient at higher redshifts' seems to be based on the insignificant (0.6-0.8 sigma) "preference for systematically larger gas densities within R200 relative to comparable halo\ model-based analyses of stacked kSZ measurements from ACT DR5 and Compton y-cosmic shear cross-correlations". This is not worthwhile commenting on further - or if the authors do, then even the word "suggest" is too strong. They should rather note that the results do not constrain the redshift dependence of feedback processes, or at best "hint at a possible redshift dependence". I want to be clear - it's good for the authors to note this as a possible cause, but these are <1 sigma effects, so the authors should take care not to over-sell this.

The only definition of the "fiducial" and "extended" samples of FRBs I can find is in the caption to Figure 2, but the motivation for these being separated is not clear to me - if they have confident host associations, then surely those FRBs can be included in the main sample? Sample selection should certainly be discussed somewhere in methods, which would allow for further discussion of biasing effects, and what actually contributes to it. For instance, the biases addressed in Ref 32 refer to selection effects against high-DM sources due to e.g. scattering, DM smearing etc and their interaction with the FRB luminosity function; but it does not consider selection effects due to a telescope's follow-up procedure or localisation, or reporting biases. I have not gone through the references for all 109 FRBs used in this work, but presumably they are only taken from "catalog" papers that report a complete sample of FRBs detected by a given instrument in a given time period, rather than cherry-picked "interesting" FRBs published because of their unusual nature. A short note on sample selection should be added to the Methods.

204-216: this comparison could also compare the authors' derived $\sigma^2(z)$ values to those of Ref 74, and the implied baryon spread metric from Ref 75. That would perhaps be an appropriate plot in the Methods.

Figure 3 a: Currently, ACT kSZ: Bigwood+2024 is repeated in both legends, while FRBs are only given in panel b, which is confusing, so I could not immediately tell what the constraints from this work were.

The last paragraph, "Looking ahead," provides an excellent summary of what FRBs can do. However, I do not think it belongs in this style of paper. It essentially reviews existing work. I would prefer the current method to be expanded upon in proper detail in the main body.

I will note that all the discussion from line 180 onwards is essentially identical to what is already known about feedback scenarios and the matter power spectrum. It is useful in the sense that it tells a reader what the current constraining power of FRBs is. But it also gives a slightly false impression in the sense that the authors are using one observable ($\sigma^2(z)$) and their model is used to tie this to a large range of physics properties of the Universe, so one might imagine scenarios where different models produce different constraints.

METHODS

Overall comment: Reading the methods, I found myself scrolling upwards to check that I hadn't accidentally skipped over a page. The ordering of the methods is quite unclear, and they are not at all described in the main body. Looking back, I

suspect this is because there is not a good overview of the methods given in the main body, making it difficult to place the relevance of specific subsections in the Methods. But I do suggest that the authors revisit the ordering of the methods sections for a clearer progression.

Line 519: /citemethods2017MNRAS.465.2936M.

518-531: good discussion of caveats

It is not clear in Eq 10 for σ^2 how the $(1+z)^{-1}$ suppression of DM with z is included in this expression. Is that part of the calculation of $P_{ee}(k,z)$? However, normally, $P_{ee}(k,z)$ is not calculated with DM and the $(1+z)^{-1}$ in mind.

Line 546: "The stellar content locked in central galaxy indirectly affect" (sounds strange - please check the meaning)

Line 548 "The scenario of extreme DMhost evolution is unlikely". What exactly is "extreme" here? A quantifiable example that can be excluded would be good to see. There is not standard "extreme" example - unless there is a paper proposing this? I am also not sure that "If DMhost exhibited strong redshift evolution, corresponding trends in [RM and scattering] observables would be anticipated, but are not seen." is correct. There is so much intrinsic variation in these parameters that even if they were strongly correlated with DMhost, one might not detect it; and furthermore, predictions for correlations do not seem to hold out in experimental data. I don't think this topic is particularly key to the paper however, but the authors should tone down their certainty. After all, the test the authors perform gives a reasonable account of the sensitivity of their results to SFR.

Line 551: "exhibit < 4% accuracy" - maybe better phrased as "exhibit < 4% error"?

Line 557 "the extended FRB sample and the fiducial subset". What are these samples and subsets? Discussion of DMhost redshift evolution is difficult to unpack here, because there is no discussion of how it relates to the samples used. Do the samples come from different redshifts? Are they simulated samples with different truths?

Analysis of the number of effective degrees of freedom in "Constrained Baryonic Feedback Parameters" is excellent. Can the authors interpret a little more their result that primarily it is M_c which is constrained? Following through the effects of M_c on the equations right through to $p(\text{DM}|z)$ is a little tricky - perhaps the authors can illustrate how $p(\text{DM}|z)$ changes, for 1-2 example z -values, when shifting either M_c alone, or shifting along the principle component vector? (Eq 19) The authors do this nicely with Extended Data Figure 1.

Line 648: The sentence "However, the inability of FRBs (and indeed, any probe of baryons) to constrain η due to a lack of constraining power necessitates a better-motivated prior on the SHM relation" needs a bit more explanation. Firstly, Secondly, being unable to "constrain" something due to a lack of "constraining power" is quite a tautology! Secondly, claiming that FRBs or other baryonic probes are simply unable to constrain η - which governs the slope of the stellar-to-halo mass relation, i.e., the fraction of baryons in stars vs gas as a function of total halo mass - seems unlikely. After all, one can imagine correlating FRB lines of sight against specific halos and statistically untangling the relative contributions as a function of halo mass. The next paragraph begins with "Since we find that the choice of fiducial η impacts the final measurement ..." also seems to contradict the notion that η can't be constrained.

Line 652 - should refer to *extended data* figure 4

The parameter-by-parameter discussion of "Stellar-to-Halo Mass Relation Prior" is very useful.

Caption to supp Figure 5 should remind a reader what η is.

Line 660 - missing bracket, or the sentence is cut off

(Remarks on code availability)

When unzipping the file, I got the following error:

```
dmz_spk_sharma2026-main/data/hydrosims_spk/fgas-8???.hdf5: write error (disk full?). Continue? (y/n/C) y fchmod (file
attributes) error: Bad file descriptor
(warning) cannot set modif./access times
No such file or directory bad CRC 17501865 (should be 350496f6)
```

Needless to say, this was *not* because my local disk was full. I was using a mac os Tahoe 26.2. The authors should check if there was some error in the file upload to google drive. Most of the rest of the files checked out.

I have not performed a full install of the codebase, nor run through all individual notebooks to check reproduction of the results - I was not sure how critical that hdf5 file error was above, and did not want to fight to install a run a code with potentially missing info. However, the README file appears clear in how to install and run the code, so I believe this is sufficient for a researcher to be able to check the authors' results - once the unzip process operates correctly.

Version 1:

Reviewer comments:

Reviewer #1

(Comments for the Author)

The authors have made substantive improvements to the paper and address practically all my comments.

My original comments regarding whether this paper -- which is primarily a methods paper -- is suitable for Nature Astronomy in the absence of new and highly constraining astrophysical scientific results stand. The authors provide useful arguments of why they feel it is suitable, but not anything that wasn't clear at the outset. I feel is ultimately an Editorial decision depending on what the journal aims to be.

That said, I definitely think the paper is important and will be well cited for its methodology and laying out a strategy for future high impact in the field.

(Remarks on code availability)

Reviewer #2

(Comments for the Author)

This is the second referee report for the article «Signatures of Suppressed Matter Clustering revealed by Fast Radio Bursts». The authors addressed all of my comments and concerns raised in the original report. As it stands right now, the article has been significantly improved: both the clarity of the discussion and presentation of the key results. Therefore, I am happy to recommend it for publication in Nature Astronomy.

(Remarks on code availability)

Reviewer #3

(Comments for the Author)

I thank the authors for their detailed response to the queries of myself, and the other two referees. I'd first like to re-iterate again that while the most substantive comments made by myself, and in particular referee #1 also, questioned the general suitability of these results, and their significance, for Nature Astronomy, this was meant in no way as a fundamental critique of the work itself.

However - in response to this critique, the authors first argue that their constraints are remarkable given the few FRBs used - and I agree with that. But this is an argument emphasising the promise of the technique, not the result. Second, the authors argue that FRBs target a critical and complementary redshift/halo mass range, noting x-ray measurements. However, they currently do not provide any significant constraints, only 0.6-0.8 sigma hints, which are not at all significant - i.e. it is about the future promise of the technique. Third, the authors argue that they establish an observational link between FRB observables and matter power spectrum suppression - aside from the fact that this claim probably should go to Ref 40 (Medlock et al), establishing such a link is a methods paper / theoretical study, not an observational result. Fourthly, the authors state that they introduce a new analysis methodology - and I fully agree that this methodology is indeed the value of the manuscript.

Therefore, I disagree with the arguments put forward by the authors on pages 1-4 of their reply as to why it is the results of the manuscript that are of broad significance, and not the method.

However - given that the Editor has decided to pass the original set of comments back to the authors, I am going to drop all concerns to this effect. My opinion was, and still remains, that this work is ill-suited to Nature Astronomy. But the authors have gone to significant effort to update their manuscript, and respond to the particular detailed queries of myself and other referees. Given this, I believe the responsible thing is to continue towards publication in Nature Astronomy, with the aim of making the manuscript as well-suited to the format as possible. And I also want to re-emphasise that I believe the authors to have produced a fantastic work with very interesting results, and it is only the framing of those results, and their suitability for NatAstron, I am critiquing.

Indeed, the authors have done a good job of addressing the specific queries of all referees. I have two further minor points, but also a significant concern regarding the methods used to curate the FRB sample used, which is now described in the new Methods section "The Localized FRB Sample".

- The single sentence from lines 78-28 should be broken up.

- the new statements on lines 195-199 about how remarkable the FRB measurement is compares ~100 FRBs to an

unspecified operational lifetime. For this to make sense to a reader, a quantity of FRBs should be compared to e.g. a quantity of clusters probed by x-ray observations, or an experimental lifetime of x-ray obs to the duration over which these FRBs were detected. However, in general, I think that a better use of such comparisons would be to note that the number of localised FRBs is increasing rapidly, and comparing this to e.g. the number of x-ray clusters expected to be studied, and to make such a comparison in the future prospects section. But this is mostly already done (line 263), so I would just remove these lines.

- I like the new caption to figure 3, which takes a reader progressively through the plots

- The new section explaining the data sample is generally good. However, the statement "we exclude 18 FRBs due to either insecure host galaxy association, evidence for excess DM_{host}, negative DM_{exgal}, or DM_{exgal} less than..." needs to be expanded. Firstly, the authors mention the "secure host association" criteria of $P_{\text{host}} > 0.9$ - presumably, "insecure host galaxy association" just means $P_{\text{host}} < 0.9$? If so, why repeat it? If not, please specify. And the criteria for evidence of excess/deficit DM contributions should be expanded upon. Indeed, to my knowledge, there is no good evidence existing in the current literature. The methodology used by the authors here is critical to avoiding biases in the data! In fact, cherry-picking data points here may invalidate the entire analysis, so the authors need to ensure their choices are well-justified.

(Remarks on code availability)

I have successfully built the code according to the installation instructions included in the README from the github site at which the code is based. It appears to be in working order and well-documented.

Open Access This Peer Review File is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

In cases where reviewers are anonymous, credit should be given to 'Anonymous Referee' and the source.

The images or other third party material in this Peer Review File are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

To view a copy of this license, visit <https://creativecommons.org/licenses/by/4.0/>

REPLY TO REFEREE’S COMMENTS

Title: Signatures of Suppressed Matter Clustering revealed by Fast Radio Bursts

Manuscript #: NATASTRON-26030375-T

We begin this document with a section intended to address common concerns of all referees, and then proceed with responses to specific comments from each referee.

We thank the referees for their insightful comments and valuable suggestions, which have significantly contributed to improving our manuscript. In this document, we address each of the referees’ concerns and highlight the corresponding changes made to the paper. The line numbers referenced below correspond to the revised manuscript, where changes are marked in red for clarity.

The referees have raised some concerns about the constraining power of the FRB dataset. We address these concerns in four parts. First, we emphasize that it is remarkable that our analysis framework, applied to just 109 FRBs, already delivers constraints competitive with conventional baryon probes, whose measurements were obtained using instruments that have operated for several years with enormous investments. Second, FRBs are not merely competitive with existing probes, but they also fill a critical and otherwise inaccessible redshift and halo mass gap that no other probe can currently reach. For example, this enables our analysis to arbitrate unresolved tensions in systematics-limited X-ray measurements, making them indispensable rather than just complementary. Third, this work establishes a direct and quantitative observational link between FRB observables and the matter power spectrum suppression – a result of broad relevance to the entire observational cosmology community, not just the FRB field, and one that already places FRBs at the threshold of enabling joint analyses with weak lensing surveys to address the longstanding S_8 tension. Fourth, beyond the immediate scientific results, this work introduces a new analysis methodology for the FRB community that directly connects DM variance to feedback-dependent matter clustering, and which we hope – alongside the baryon density field reconstruction approach of the FLIMFLAM collaboration – will become a new analysis standard as FRB samples grow by orders of magnitude with next-generation facilities. We have revised the summary and relevant paragraphs of the manuscript to make these points clearer to a broad audience of astrophysicists. Below, we discuss each of these points in detail.

- **Competitive constraints from just 109 FRBs: a remarkable milestone.** We respectfully push back on the characterization that the results are not very constraining. The key point – which was insufficiently emphasized in the original manuscript – is not the absolute level of constraint, but rather what our analysis framework makes possible to achieve with such a dataset. With just 109 FRBs, our framework delivers constraints on the matter power spectrum suppression that are **competitive with what eROSITA and Atacama Cosmology Telescope (ACT) have achieved over their entire operational lifetimes** using traditional probes

such as X-ray observations, kinetic (kSZ) and thermal (tSZ) Sunyaev-Zel'dovich effects. The fact that our methodology, applied to a sample of just over a hundred FRBs, already delivers competitive constraints – and that next-generation facilities such as CHIME/Outriggers, CHORD, DSA, and SKA will exceed this dataset by orders of magnitude – is a remarkable, groundbreaking result. Our analysis framework unlocks the constraining power of the current FRB sample, and will scale accordingly as larger samples become available. We have revised the manuscript (L25-26, L195-199) to make this comparison explicit. Furthermore, there is no planned next-generation X-ray instrument with substantially improved sensitivity over eROSITA to conduct these measurements. While the kSZ effect measurements from the Simons Observatory will excel at sub-Mpc scales, they will be limited by inevitable contamination from primary CMB anisotropies. Furthermore, tSZ measurements from Simons Observatory are limited by cosmic infrared background contamination. This makes the science case for FRBs not merely complementary, but arguably irreplaceable for constraining baryonic feedback at Mpc scales in the coming decade.

- **FRBs independently fill a critical redshift and halo mass regime inaccessible to current SZ measurements – necessary for weak lensing analyses and to arbitrate tensions in X-ray measurements.** Even though current FRB, tSZ and kSZ constraints appear competitive in terms of the matter power spectrum suppression amplitude, our sensitivity analysis reveals that they are probing **fundamentally different regimes of halo mass and redshift**. As illustrated in Fig. 1, DM variance measurements from current FRB sample are primarily sensitive to low-redshift ($\lesssim 0.3$) group and cluster-scale halos, whereas kSZ measurements preferentially probe higher-redshift groups and tSZ measurements probe higher-redshift clusters. No existing work has demonstrated that current SZ measurements can deliver meaningful constraints in the low-redshift group and cluster regime at the required precision.

This distinction has two important consequences. First, to support and interpret current weak lensing analyses from surveys such as DECADE, DES, KiDS and HSC, what is needed are measurements of power spectrum suppression from redshifts $z \sim 0 - 0.5$ for group and cluster-scale halos – a regime that SZ observations alone do not currently cover. Second, the only probe other than FRBs sensitive to low-redshift groups and clusters are X-ray observations, which are currently systematics-limited: the X-ray measurements of gas mass fractions (shown in Fig. 4) – often used to calibrate subgrid feedback physics in hydrodynamical simulations – reveal an $\sim 8\sigma$ tension between pre-eROSITA versus eROSITA data. Without an independent probe sensitive to the same halo mass and redshift regime, there is no definitive way to adjudicate between conflicting measurements. Our framework therefore positions FRBs as currently the only probe that can serve the role of arbitrating the tension in systematics-limited X-ray observations of low-redshift groups and clusters (L21-22, L173-

183), offering sensitivity to the total ionized gas content – including the cool, diffuse phase invisible to X-rays. This makes the present analysis of broad relevance, not only to the FRB field, but to the wider observational cosmology and galaxy formation theorists communities.

- **A direct observational link between FRBs and the matter power spectrum.** This work represents one of the first measurements that quantitatively links FRB observables to the quantity of central interest to observational cosmologists: the suppression of the matter power spectrum due to baryonic feedback. Previous FRB-based analysis have focused on estimates of the mean diffuse baryon density or the partitioning of baryons between halos and the intergalactic medium. Here, we go a step further by **inferring the spatial fluctuations in the baryon density field** through the gas power spectrum and connecting them to the matter power spectrum suppression across a range of scales. The systematic comparison of halo mass, redshift, and scale sensitivity across multiple probes (Fig. 1) – and the corresponding comparisons of inferred power spectrum suppression (Fig. 4) – is of direct relevance to all researchers interested in the impact of feedback on matter clustering, far beyond the FRB community. Crucially, the current FRB-based matter power spectrum suppression constraints are already competitive enough to begin working towards the next milestone: joint analyses with weak lensing measurements that can directly address the question of whether baryonic feedback is a viable solution to the longstanding S_8 tension. As a point of reference, Bigwood et al. (2024) conducted a joint kSZ and weak lensing analysis using ACT DR5 stacks – which we demonstrate in Fig. 4 have comparable constraining power to our current FRB DM– z relation-based measurements – to place constraints on feedback and cosmology simultaneously. This demonstrates that the FRB constraints presented here are not just a proof of concept, but are already at the threshold of enabling the kind of multi-probe joint analyses that will be central to precision cosmology in the Stage-IV era.
- **Novel inference framework: a new standard methodology for FRB analyses.** Beyond the immediate scientific results, this work introduces a novel methodology for modeling the FRB DM– z relation that is new to the FRB community as a whole. By directly connecting the sightline-to-sightline variance of FRB dispersion measures to the feedback-dependent matter power spectrum suppression through a state-of-the-art halo model framework, we develop an analysis pipeline that goes significantly beyond previous FRB-based baryon studies. We have substantially revised the manuscript to provide more intuitive and accessible details of the working principles of this framework (Fig. 3, L113-149), so that it can be adopted by the broader FRB community. We hope that over the coming years, this DM– z modeling methodology – alongside the baryon density field reconstruction technique developed Khrykin et al. (2024) – will become a new analysis standard within the FRB community, enabling increasingly precise and robust constraints on baryonic feedback as FRB samples grow by orders of

magnitude with next-generation facilities.

The framing above – competitive constraints from just 109 FRBs, filling a critical and otherwise inaccessible redshift-halo mass regime, providing direct link between FRB observables and the matter power spectrum, and the development of a novel methodology for the FRB $DM-z$ analysis – together make a compelling case for broad impact.

Referee 1

This is a review of Sharma et al. “Signatures of Suppressed. . . Bursts.”

This is a very long, dense and highly technical paper on a very timely and important topic, namely the development of FRBs as a new tool for cosmology. The paper’s importance does not lie so much in the results that it reports. Indeed, my impressions from the opening paragraph were that an improvement in an unspecified prior and ruling out extreme models at the 2 sigma level are not major leaps forward. Rather, the value of the paper lies in showing how such an analysis can be done, and how FRBs will be very powerful in the future when there are large samples (although getting redshifts for 10^4 FRBs may be problematic - an issue not admitted or discussed by the authors, see below). The bulk of the important and novel analysis is done and explained in the Methods of the paper, which is presently very dense and very long.

I think of Nature Astronomy as a journal meant to showcase results that are of a somewhat broader impact and comprehensible to a broader audience of astronomers not just in the field reported on in the paper. I leave it therefore to the Editor’s discretion regarding whether the paper is suitable for Nature Astronomy, given that the results are not (yet) very constraining, but given that the techniques being used are novel, interesting, and are pushing the field forward in a very important and impressive manner.

We thank the referee for their careful reading of the manuscript, and for their thoughtful and encouraging assessment of the novelty and importance of the techniques presented. We also appreciate the referee’s candid comments regarding the framing of the results. We believe there are several important points that were not sufficiently emphasized in the original manuscript, and we discuss them in the common preamble addressed to all referees in the beginning of this document. In summary, we articulated in detail four distinct contributions of this work that we believe establish its broad relevance to observational cosmology and astrophysics community: 1) the remarkable constraining power of just 109 FRBs relative to years of investment in traditional baryon probes (L25-26, L195-199); 2) the unique and irreplaceable role of FRBs in filling a critical redshift and halo mass gap that no other current probe can access to help arbitrate the tensions in systematics-limited X-ray measurements which are often used to calibrate subgrid feedback physics in hydrodynamical simulations (L21-22, L173-183); 3) the direct observational link between FRB observables and the matter power spectrum suppression, a quantity of central interest to observational cosmologists, with current constraints already capable to addressing the question of S_8 tension in joint weak lensing analyses; and 4) a novel analysis methodology for the FRB community that enhances the impact of FRB-based measurements by explicitly addressing feedback-driven uncertainties in standard cosmological analysis framework (Fig.3, L113-149). We have also revised the opening paragraph of the manuscript to make these points clearer to a broader audience.

Regarding the concern that obtaining redshifts for 10^4 FRBs may be problematic, we address this

in detail in our response to the referee’s specific question on this topic below.

Main

Paragraph 1 refers to scales via k , whereas paragraph 2 in Mpc. It would help people like me to not jump back and forth; actually my preference would be for each mention, provide both.

In the revised manuscript, we have adopted a consistent convention of quoting scales in Fourier space (k , in units of $h \text{ Mpc}^{-1}$). The corresponding configuration-space scales are simply the inverse, i.e., $r \sim 1/k$. This ensures a uniform description of scales across the manuscript.

lines 70, 81: “X-ray gas mass fractions” is a bit of a word salad... how about “cluster gas mass fractions measured through X-ray observations”

We thank the referee for this suggestion and we have updated the text for the same (L50-51, L61).

line 73: enables $\rightarrow$ enable

Fixed (L53).

lines 96-97: “Hydrodynamical simulations...” The authors have just explained why these are presently not well constrained at the relevant scales so how can we believe this statement?

We thank the referee for this careful observation. The log-normal DM_{host} distribution is an approximation that has been shown to work well in a handful of hydrodynamical simulations, rather than a robustly established result across the full range of plausible feedback scenarios. We have revised the text to make it explicit that this is an approximation (L115). While one may consider formally connecting the DM_{host} PDF to the electron power spectrum because the FRB host galaxies are, in principle, subject to the same feedback processes as intervening halos, we expect the connection to be fairly non-trivial due to the sensitivity of DM_{host} to source-specific physics, such as the circum-burst environment and the FRB’s location within the host – scales (sub-parsec) and physics that are not captured by halo model frameworks. A self-consistent treatment within the same formalism is therefore beyond the scope of this work. We instead follow the standard approach of approximating DM_{host} as a nuisance component described by a log-normal distribution and marginalize over it during inference (L592-599).

line 98: “does not impact current measurements” Can the authors summarize here in a few words why? This is a very important statement so it is a bit frustrating to have to go to Methods to find out why

We thank the referee for this suggestion. We have revised this sentence to note that the insensitivity to DM_{host} redshift evolution arises because the redshift sensitivity of current FRB sample spans a

relatively narrow and low range, over which any plausible redshift evolution of DM_{host} – for example, one tracking the cosmic star-formation rate history – produces shifts in the inferred feedback parameters that remain within 1σ , as demonstrated in Extended Data Fig. 1 (L79-81).

lines 100-101: the statement “Likewise, selection effects” is an important one and relies entirely on the single study by the same authors, which unless I’m mistaken, assumes FRBs all detected with the same selection biases i.e. with a single instrument. Here, the authors cobble together a sample from many instruments. I gather the argument is that if selection biases don’t matter for any one instrument then there should be no problem combining. I think if that is the reasoning being used, it needs to be stated explicitly. In the limit where one FRB survey is very insensitive and sees only nearby FRBs and another is vastly more sensitive and sees FRBs to high redshift, even if selection biases don’t matter for either, it’s not obvious to me combining them doesn’t pose a problem.

We thank the referee for this careful and valid point. The referee is correct that our previous work was conducted assuming a single instrument’s selection function. We agree that the assumption that selection biases are negligible for any individual instrument does not necessarily guarantee that combining instruments with vastly different properties is unproblematic. We have revised the relevant text in the paper to make our assumption explicit (L84-86). We however note that we have conducted mock tests to validate the ability of the current observed FRB sample to discriminate between extreme feedback scenarios (Supplementary Fig. 2). Crucially, these mock tests use the actual observed redshift distribution of our 109 FRBs, which is a product of combining multiple instruments with different selection properties. The unbiased recovery of both the no-feedback and strong-feedback scenarios from this multi-instrument redshift distribution provides confidence that, at current sample sizes, selection biases from combining instruments are unlikely to play a dominant role in our inference. We nonetheless fully agree with the referee that as FRB samples grow by orders of magnitude with next-generation multiple facilities, a more careful treatment of multi-instrument selection effects will be essential, and we flag this explicitly as an important systematic to address in future analyses (L84-86).

line 108: M_{200} is not defined. What does the 200 mean?

We have now added a brief definition of M_{200} and R_{200} at its first appearance in the main text, clarifying that it denotes the total mass enclosed within a sphere whose mean interior density is 200 times the critical density of the Universe at the halo’s redshift (L92-93).

Fig 1: For plots b and c, in two cases top one-third and two-thirds contours are plotted, and for two others, 86th percentile contours are plotted, and it is not stated what is plotted for the DM case. First, the latter needs rectifying, but more generally, it is hard to robustly compare all of these when such different contours are plotted. Also, the y-axis of panel c is not properly explained. Please

clarify what this quantity is. I also found the dual method for specifying M_0 (legend with 3 line styles and color bar) confusing. I guess it is meant to show how things vary with M_0 , so the color bar should suffice, but the authors wanted to also specify clearly for the round numbers. To make it a bit less confusing, I suggest that the lines used for the round numbers be colored according to where they are in the color bar.

We agree that plotting different percentile contours for different probes makes a robust visual comparison difficult. We have revised Fig. 1 to show uniform contours across all probes – including the FRB DM variance – corresponding to the top one-third and two-thirds of sensitivity. This ensures a consistent and fair comparison across all baryon tracers shown in the figure. We have also added a description of the y-axis of panel c to the caption, where we explain that the modified power spectrum ($P_{\text{hydro}}^{M \geq M_0} + P_{\text{gravity}}^{M < M_0}$) is constructed such that baryonic feedback affects only halos with masses $M \geq M_0$, for which the matter power spectrum is computed using feedback-modified (hydrodynamic) profiles, contributing $P_{\text{hydro}}^{M \geq M_0}$, while halos with $M < M_0$ are assumed to be unaffected by baryonic feedback and are instead described using gravity-only profiles in the power spectrum calculation, contributing $P_{\text{gravity}}^{M < M_0}$. We have now also removed the additional M_0 legend for clarity because we agree with the referee that the color bar should be sufficient for interpreting this plot.

Lines 124-125: CHIME/FRB is also undergoing an upgrade that could increase their detection rate by a factor of 3-5, and together with its Outriggers, could also make it a player in this game.

We have added CHIME/FRB Outriggers to the list of next-generation FRB experiments in the manuscript, alongside DSA, SKA, and CHORD, as a facility that will contribute substantially to the growth of well-localized FRB samples in the coming years (L249-250, Fig. 6 caption).

Lines 128,129: CHIME should be CHIME/FRB to distinguish from the cosmology experiment. Also please fix all occurrences in Extended Table 1.

We have now updated all occurrences of “CHIME” to “CHIME/FRB” consistently throughout the manuscript (L111, L112, Fig. 2 caption, Extended Data Table 1).

Fig 2: The title of this figure confused me. The plot shows DM vs z so I'd have expected that should be the title. I don't see how the plot is telling me clearly how baryon clustering is probed. In the caption it says “quantified by the degree of clustering of baryons” is also confusing. Perhaps the authors mean that the variance quantifies the degree of clustering rather than that is being quantified by it? And again CHIME $\rightarrow$ CHIME/FRB

We have revised the figure title to “The observed FRB dispersion measure–redshift relation.” We have also reworded the sentence in caption to read: “The variance (i.e. spread in DM for fixed z), arising from intervening halos, quantifies the degree of baryon clustering, and is sensitive to the strength of feedback.”

line 131: provide a reference for BCEMU7 halo model

We have now added the Giri & Schneider (2021) reference (L127).

line 134: specify what NFW is for non-specialists

We have removed the usage of NFW in the Main text and added a brief clarification at the first mention of the NFW profile in Methods, noting that NFW refers to the Navarro-Frenk-White profile, the standard analytical description of the dark matter density distribution within gravitationally collapsed halos (L131, L581-583).

line 139: please provide equation numbers rather than just saying “key equations in the Methods”

We have revised the text to explicitly reference the relevant equation numbers with corresponding parameter descriptions in the Main text (L129-134).

Paragraph starting at line 140: The text here suddenly becomes impenetrable to me. I feel the authors need to do a better job at explaining what they are constraining for a journal like Nature Astronomy that is meant to be comprehensible to people outside this narrow field. E.g. why is “with a 224% reduction in the variance of the 1 dominant eigenmode relative to prior” important? and what is a “Karhunen-Loeve mode”?

We have revised the beginning of the paragraph to convey the same essential information in plain language, removing all specialized terminology, and leaving the technical details in the Methods section. The revised beginning of the paragraph now clearly communicates two key messages: (1) despite the complexity of the 7-parameter halo model, the FRB data effectively constrain one combination of parameters; (2) this constrained combination is dominated by $\log M_c$, the characteristic halo mass scale governing gas expulsion efficiency, while the remaining parameters are largely prior-dominated (L150-152).

lines 152-153: the number of significant figures quoted for DM_{host} and sigma are too high. Looks like at most you should be quoting 3 sig figs.

Done (L162-163).

Also and more importantly, the value for DM_{host} is somewhat higher than that inferred by Shin et al. 2023 based on a very large, uniform sample of FRBs albeit without redshifts. The authors should comment on the difference.

We thank the referee for raising this point. Our inferred value of $\langle DM_{\text{host}} \rangle = 129^{+18}_{-18} \text{ pc cm}^{-3}$ is broadly consistent with other FRB-based measurements in the literature: Shin et al. (2023) report $84^{+69}_{-49} \text{ pc cm}^{-3}$, Khrykin et al. (2024) report $90^{+29}_{-19} \text{ pc cm}^{-3}$, and Connor et al. (2025) report $130^{+25}_{-23} \text{ pc cm}^{-3}$. The differences between these measurements are all at less than 1σ significance,

which can arise from differences in the modeling of the Milky Way DM contribution, as demonstrated explicitly by Shin et al. (2023) in their Figure 5 and Section 5.3.1. These methodological differences in Milky Way subtraction alone are sufficient to account for much of the apparent mild tensions in DM_{host} values across studies.

We note, however, that a cleaner measurement of DM_{host} can be obtained from local FRB samples, where the observed DM is known to be dominated by the host galaxy contribution rather than the cosmic component. Leung et al. (2025) conducted one such study, finding an average DM_{host} of $\sim 150 \text{ pc cm}^{-3}$. The ultimate gold standard for measuring DM_{host} would be spatially resolved spectroscopy of host galaxies, which directly quantifies the host contribution without any modeling assumptions. This approach has so far been limited by the scarcity of FRBs with accurate localizations in the local universe, but is expected to improve substantially in the coming years as (sub)arcsecond-scale localizations become routine – particularly with CHIME/FRB Outriggers, which, owing to CHIME’s wide field of view and sensitivity, is uniquely positioned to build the statistical local FRB samples required for such measurements.

We have added a discussion of these points to the Methods section (L863-875).

Fig 3: Left - I guess the hatched region are the FRBs? The legend in the right plot explicitly says this; why isn’t that also done in the left plot?

Done (Fig. 4 caption).

Fig 4: Re the sentence “The entropy...” I don’t see where and in which plot this is shown and suspect it is not, so it’s unclear to me why the statement is made in the plot caption here.

We agree with the referee and have removed this statement (Fig. 5 caption).

line 172: Commenting on a 0.6-0.8 sigma preference is a stretch. The final sentence in this paragraph “Together, these results...” is extremely speculative and not based on very strong evidence. This should be better emphasized so as not to give the impression the authors have actually demonstrated any of this. I realize they say “suggest” but I find that to be a stretch given the actual results.

We thank the referee for raising this concern. We agree that $0.6 - 0.8\sigma$ differences between probes hints for a preference and are not yet a strong evidence. We have now revised the phrasing to reflect the same (L191, L193, L203-L211).

line 195: Similar to above, the word “discrepant” seems an overstatement. There is mild tension here, not discrepancy.

We agree with the referee and have now revised the phrasing to reflect the same (L222-224).

line 218-219: There is some concern that these surveys will be limited by ability to measure redshifts for a large population especially for faint hosts where very large telescopes are needed, even at moderate redshifts. That needs to be acknowledged and could ultimately be the limiting factor in achieving the goals of this study.

We thank the referee for raising this important practical concern. The referee is absolutely correct that obtaining spectroscopic or photometric redshifts for large FRB host galaxy samples – particularly for faint hosts at moderate to high redshifts with luminosity/magnitudes beyond the imaging depths of mentioned galaxy surveys – will require substantial telescope time on very large facilities, and could ultimately become a limiting factor in realizing the full potential of FRB-based cosmological analyses. We acknowledge this in the revised manuscript. However, we note that there is one path forward that can help address this limitation, at least partially.

Cross-correlation redshift distributions offer a powerful alternative to spectroscopic or photometric redshifts for individual FRBs. With large FRB samples, one of the ultimate goals is to conduct joint analyses with weak lensing and galaxy clustering in an $n \times 2$ -point framework, as we detailed in the Methods. These analyses only require the sample redshift distributions to model the signal during inference. While having per-source redshift information is valuable for accurately removing the mean DM contribution, even an ensemble average mean subtraction can serve as a reasonable approximation; the recent detections of CHIME FRB DM cross-correlations with a range of large-scale structure tracers by Wang et al. (2025) and Sharma et al. (2026) corroborates the feasibility of this approach. By cross-correlating FRB positions with photometric galaxy catalogs in various tomographic bins and/or cosmic infrared background light at different frequencies – which probes different redshift distributions of its source star-forming galaxies (see Fig. 2 of Chiang et al. (2025)) – one can in principle statistically infer the redshift distribution for an FRB population without requiring individual host identifications. Currently, the biggest practical limitation of this approach is the size of localization uncertainties, which are expected to improve with next-generation FRB experiments significantly as arcsecond-scale localizations become routine.

We have added this discussion to the revised manuscript to acknowledge this concern and outline this path forward (L252-254, L259-262, L1010-1023).

Methods

The authors need to provide some preamble at the start of the “Sensitivity Analysis” section to help the reader understand to what this pertains in the main text, and what the goal of the section is. Indeed, every Methods subsection should have this.

We thank the referee for this helpful suggestion. We have added a preamble to each Methods subsection that orients the reader and points to the relevant figures or context (L461-464, L527-529, L547-554, L615-617, L630-632, L638-642, L707-711, L739-745, L857-862, L921-925).

line 565: As asked above, Shin et al. 2023 find a much lower mean DM_{host} . Also the reference provided here [81] doesn't seem appropriate for this statement.

We have addressed the comparison with Shin et al. (2023) and other literature values of DM_{host} as part of our response to the earlier comment on this topic. We thank the referee for flagging reference [81], which we have now removed from this location in the text.

Referee 2

The astrophysical feedback produced by the star-formation processes and/or AGN activity are key ingredients in the hydrodynamical simulations of the galaxy formation and evolution. Moreover, they are considered to be responsible for the redistribution of cosmic baryons between the diffuse intergalactic medium (IGM), multiphase circumgalactic medium (CGM) and the interstellar medium (ISM). Subsequently, due to its effect on the matter clustering, feedback directly affects any cosmological measurements. Though a plethora of observational techniques have been employed by the community in order to study and quantify the feedback in data at $z < 2$, all of these methods require specific assumptions about the density, metallicity and temperature of the probed gas that bias the interpretation of the results. In that respect, the FRB dispersion measure constitutes a unique probe of all ionized baryons irrespective of such assumptions.

This work is a continuation of authors' previous theoretical efforts to quantify the sensitivity of the FRB DMs to the matter power spectrum suppression, feedback and halo models. However, this is the first attempt to apply the developed model and inference algorithm to the observed sample of FRB DMs. Moreover, authors assess the sensitivity of the FRB-based constraints and compare it to those by other astrophysical probes. The authors utilize a robust and well-defined statistical algorithm with careful treatment of uncertainties. Additionally, this work lays a foundation for the future usage of FRBs for cosmological applications. While this manuscript definitely has merit and deserves a publication in Nature Astronomy, first, before proceeding, the authors must address some of the comments/concerns outlined in what follows.

We thank the referee for their careful and thorough reading of the manuscript, and for their encouraging assessment. We have addressed all of the concerns and comments raised, as detailed in our point-by-point responses below.

Main

1. L104, 108, 119: variance should be σ^2 , not just σ .

We thank the referee for flagging this. We have corrected all occurrences of variance throughout the manuscript (L88, L94, L104, Fig. 1 caption, L492).

2. L131: you need to reference the BCEmu7 model - at this point the reader has no idea what it is.

We have now added the Giri & Schneider (2021) reference (L127).

3. L153: this is not a variance, but standard deviation.

Fixed (L163).

Methods/Sensitivity Analysis

1. 430-432: *Please, make it explicit that results for other probes presented in Fig. 1 are computed via eq.(1). This section generally lacks references to Fig.1 of the main text.*

We have now added explicit references to Fig. 1 in this section (L461-464, L498-499) and clarify that the sensitivity of other probes shown in Fig. 1 is computed using Eqn. 1 (L465-466).

2. 435: *What is the effect of fixed source redshift distribution on the analysis? The authors should expand their statement and comment on any potential effects on the sensitivity function.*

We want to clarify that the goal of the sensitivity calculations is to characterize the halo mass, redshift, and scale sensitivity of the imprint of feedback on existing measurements from various baryon probes (shown in Fig. 4 and 5). These calculations specifically correspond to current measurements of cosmic shear from DES-Y3, shear- γ cross-correlations from DES-Y3+ACT-DR5, and the ACT DR5 tSZ Y - M relation – all of which have been used successfully to measure matter power spectrum suppression presented in Fig. 4. In this context, fixing the source/cluster redshift distribution is the correct approach: we are interested in understanding which halo masses and redshifts contribute most to the feedback signal in these already-made measurements, given the actual source/cluster redshift distributions used in those analyses. The sensitivity contours would of course shift significantly as instruments become more sensitive and redshift baselines extend, as illustrated for next-generation surveys in Fig. 6. We have expanded the relevant sentence in the Methods to make this distinction explicit (L478-481).

3. 448: *The authors have not indicated the number of FRBs that was used for the sensitivity prediction, neither in this section nor in the main text (this includes the Fig. 1 caption).*

We thank the referee for this comment. We clarify that this is not a sensitivity prediction for a future FRB sample, but rather the sensitivity of existing sample of 109 FRBs used in this work. We now clarify this in the manuscript to avoid confusion (L104, L485-486, Fig. 1 caption).

4. 454: *Did the authors mean standard deviation? Variance (assuming of a log-normal distribution) in DM_{host} of 100 pc/cc seems to be rather tiny implying $\text{std.dev} \sim 10$ pc/cc. Please clarify and give reference for this number, and be explicit of the quantities and why the specific values have been chosen. I think this value is quoted correctly in line 819.*

The quoted value of 100 pc cm^{-3} is indeed the standard deviation of the host galaxy DM distribution, not the variance. We have corrected the terminology throughout the text to explicitly state standard deviation wherever this quantity appears (L162-163, L495, Extended Data Fig. 1 caption). The value of 100 pc cm^{-3} is an order-of-magnitude estimate motivated by the measurements of Khrykin et al. (2024) and Connor et al. (2025), for which we have now added references (L496).

Sensitivity synopsis of baryon probes

1. The text in this section will benefit from cross-referencing with Fig. 1. Please include description of the color/format of the contours in Fig.1 that correspond to the described measurement. In addition, authors should consider including a short description of the FRB DM sensitivity as well for consistency.

Done (L498-511).

Halo mass sensitivity of matter power spectrum suppression

1. Authors did not explain how the corresponding power spectra were calculated. Presumably, this is related to the discussion in section “Halo Model Approach”, but no references are made. Please make it clear how the power spectra in panel C of Fig. 1 is computed.

We have now added explanation of the modified power spectrum calculation with reference to the following section (L516-522).

Inference Framework

1. 495: “Equations 9 and 10 establish that both, the mean $DM_{\text{exgal}}-z$ relation and the scatter in DM_{exgal} at a given z carry cosmological and astrophysically valuable information.”. Technically these equations describe DM_{cosmic} , not DM_{exgal} . Please, fix. Overall, this statement is too generic and unnecessary as it is clear that any detected astrophysical quantity carries astrophysically valuable information, unless, of course there is some information that the authors assume to be astrophysically useless.

We agree that this statement is unnecessary, and have removed it from the revised manuscript.

2. Overall, the PDF in equation 8 establishes the connection between the DM_{cosmic} and gas/electron power spectrum of the intervening halos via equations 9 and 10. In turn, the power spectrum is related to various feedback processes. However, the second term under the integral in eq.8, describing DM_{host} PDF, does not depend on the feedback by construction. FRB host galaxies, in principle, are subjected to the same feedback processes and, therefore the corresponding PDF can (should) similarly be written as a function of the electron power spectrum. Authors should explain why they chose to assume no feedback dependence for the host galaxies or model it through their numerical scheme by estimating the electron power spectra for the hosts.

We thank the referee for this insightful comment. We agree that FRB host galaxies are, in principle, subject to the same feedback processes as intervening halos, and that the host DM PDF could formally be connected to the electron power spectrum. In practice, however, the host DM contribution depends sensitively on unconstrained source-specific physics, such as the sub-parsec scale circum-burst environment and the FRB source’s progenitor formation channel-dependent location within

the host – scales and physics that are not captured by halo model frameworks. A self-consistent treatment within our formalism is therefore non-trivial and beyond the scope of this work. We instead follow the standard approach of modeling DM_{host} as a nuisance component described by a log-normal distribution and marginalize over it during inference. We clarify this in the Methods section of revised manuscript (L592-599).

Halo model

1. 511-512: *..two halo model prescription.. is a bit ambiguous: do you mean a two halo model for the power spectrum analysis or two models to describe a halo? Obviously, the latter, but it helps to rephrase for better clarity.*

We have rephrased all occurrences in the text for clarity (L549-550, Extended Data Table 2 caption).

2. 519,522: *check the formatting of the references.*

Fixed.

3. 529: *“Due to its inherent limitations, we use this model primarily as a cross-check against our baseline results obtained with BCEMU model”. It is not clear how the authors perform this cross-check between seemingly completely different inferred parameters: temperature of the AGN-influenced gas on one hand, and the halo mass and geometric terms on the other. Only the parameters governing the DM_{host} PDF are the same between two models. Please, explain.*

We thank the referee for flagging this ambiguity. We clarify that the cross-check is not performed at the level of individual model parameters – which are indeed physically distinct between HMCODE and BCEMU – but rather at the level of the derived matter power spectrum suppression posterior. The purpose is simply to verify that the inferred suppression is stable across two different gas profile prescriptions, providing confidence that our results are not sensitive to the specific parameterization of baryonic feedback adopted. We have clarified this in the revised text (L568-573).

4. 551: *for consistency, this should be variance, i.e., σ^2*

Fixed (L602).

Plausible DM_{host} redshift evolution

1. 559: *“However, the apparent shift in posteriors between the fiducial and extended FRB sample may be indicative of a potential DM_{host} redshift evolution.” This statement is misleading as the apparent shift can also indicate the presence of some bias once extra FRB sightlines are included/excluded. I suggest the authors do a simple test by splitting the fiducial sample into two redshift bins and repeating the analysis. Moreover, to control possible bias once can do the jack-*

knife analysis repeatedly with $N-1$ sightlines in each sample, excluding a random sightline from each run.

We thank the referee for this suggestion. We agree that the apparent shift in posteriors between the fiducial and extended samples could, in principle, reflect biases introduced by including/excluding certain sightlines, rather than genuine DM_{host} redshift evolution. To address this, we have performed two tests:

(1) Redshift-split analysis (L885-898): We split the FRB sample into two subsets of 54 and 55 FRBs, divided at redshift $z = 0.25$, with median redshifts of $\bar{z} = 0.08$ and $\bar{z} = 0.37$ for the low- and high-redshift subsets, respectively. As shown in the updated Extended Data Fig. 1, the inferred mean host DM contribution $\langle \text{DM}_{\text{host}} \rangle$ is consistent across the two redshift bins. When comparing with the posteriors from the full sample analysis, consistent with the findings of Connor+2025, we find that $\sigma[\text{DM}_{\text{host}}]$ is primarily constrained by the lower-redshift subset. This is physically expected: at low redshifts, $\langle \text{DM}_{\text{cosmic}}(z) \rangle$ is small, so the total variance in DM_{exgal} is dominated by host galaxy contributions, providing a direct handle on $\sigma[\text{DM}_{\text{host}}]$. At higher redshifts, the growing standard deviation $\sigma[\text{DM}_{\text{cosmic}}(z_s)]$ increasingly dominates the total scatter, creating a degeneracy between host and cosmic contributions. To break this degeneracy, the lower-redshift FRB sample is needed to anchor the DM_{exgal} -redshift relation and enable unbiased inference of DM_{host} contribution (also observed in Sharma+2026, Fig. 3). Crucially, $\log T_{\text{AGN}}$ remains stable at the 1σ level across both redshift subsets, indicating that our current constraints on baryonic feedback are insensitive to potential redshift evolution of DM_{host} .

(2) Jackknife resampling (L630-636): As described in Supplementary Fig. 1, we performed a leave-one-out jackknife resampling analysis, demonstrating that the posterior constraints are stable across all realizations, ruling out bias from individual sightlines.

2. 561: define ‘extreme DM_{host} evolution’;

We have clarified the term “extreme” in the revised text (L878-880). By “extreme DM_{host} redshift evolution”, we specifically mean the scenario in which there is no variance in $\text{DM}_{\text{cosmic}}$ at all – i.e., all of the observed scatter in the $\text{DM}-z$ relation arises from the DM_{host} redshift evolution alone. This limiting case comes up frequently in discussions with non-FRB experts and represents a natural worst-case scenario to consider. The most direct evidence against this extreme scenario is the robust detection of excess $\text{DM}_{\text{cosmic}}$ around intervening halos – through sightline resolved studies of FRBs viewed through galaxy clusters, stacking analyses of FRB sightlines intersecting foreground galaxies, and through cross-correlations of FRB DMs with large-scale structure – which together demonstrate that cosmic structure contributes meaningfully to the observed DM variance. This directly supports our use of DM variance as a probe of feedback-regulated structure formation.

3. 564: define ‘local FRB samples’;

Done (L869).

4. 566: missing citations: e.g., Mannings+2021, Khrykin+2026

We thank the referee for flagging the missing citations. We note that the statements about rotation measures and scattering delays have been removed from the revised manuscript in response to a comment from another referee.

5. 569-570: missing citations: e.g., James+2022a

We have added the missing citation in the revised manuscript (L896).

6. 572: one can argue that instead of ‘robust’ the word ‘insensitive’ can be used.

Yes, we agree that ‘insensitive’ is the more appropriate word in this context, and we have updated the text accordingly (L898).

Constrained Baryonic Feedback Parameters

1. Not surprisingly, parameters of the full 7-parameter halo model are not constrained with the current dataset. Moreover, the halo models with reduced number of parameters are also unconstrained as the respective posterior PDFs span all parameter ranges with comparable amplitudes. I would argue that only the BCEmu1 model has (at this point) the constraining power, which is indeed backed by your DOF analysis. Have authors considered less sophisticated analytical models? It would be beneficial if authors discuss why other analytical models are worse (if they are) than the BCEmu set and shouldn't be considered.

We thank the referee for this comment. We agree that the current dataset has limited constraining power in the full parameter space, as quantified by our principal component analysis. However, we emphasize that marginalizing over the full parameter space of the reduced-complexity models (BCEMU4, BCEMU5) is preferable over fixing parameters to fiducial values, for the following reasons (which we now also describe at L696-705 in the manuscript).

The CPC analysis demonstrates that θ_{ej} , δ , and μ_β are mildly degenerate with $\log M_c$ along the first constrained mode (which is mildly also visible in Extended Data Fig. 2). Fixing these parameters to fiducial values – as in BCEMU1 – effectively slices through these degeneracy directions, which can artificially shift and narrow the $\log M_c$ posterior. The somewhat unusual behavior of the BCEMU1 posteriors visible in Extended Data Fig. 2 is consistent with exactly this effect. It is therefore safer to marginalize over a wide prior on these parameters rather than risk biasing the inference with all parameters fixed. Furthermore, the one-parameter model varying only $\log M_c$ is known to be an inaccurate representation of the full scale-dependent complexity of feedback-induced matter power spectrum suppression (Giri & Schneider (2021) and Bigwood et al. (2024) discuss this in detail).

It should therefore always be interpreted with caution. The multi-parameter models (BCEMU4, BCEMU5, BCEMU7) provide a more faithful and flexible description of the feedback-dependent gas profiles, and it is across these that our constraints are reported and compared with other probes and hydrodynamical simulations.

For completeness, we also show results from all reduced complexity models in Extended Data Fig. 2 and Supplementary Fig. 2, and note that the constraints on $\log M_c$ remain stable across BCEMU4, BCEMU5, and BCEMU7.

2. While the FRBs clearly have very limited sensitivity to the specific parameters of a particular BCEmu halo model, I suggest the authors also show how other probes can/cannot recover BCEmu model parameters. Such comparison will help to understand what FRBs bring to the table.

First, we want to emphasize that the contribution of FRBs as a probe is not only manifested in the exact BCEMU halo model parameters being constrained (which is nominally the same across all probes), but instead in 1) the fundamentally different regime of halo mass, redshift, and physical scales probed by them (which are required to interpret weak lensing measurements), compared to traditional baryon probes, such as kSZ and tSZ effect (demonstrated in Fig. 1), and 2) the ability to arbitrate the tensions in conflicting X-ray measurements (often used to calibrate subgrid feedback physics in hydrodynamical simulations), which are currently systematics limited (illustrated in Fig. 4). While we agree that a direct comparison of the specific BCEMU halo model parameters constrained by FRBs with those constrained by other baryon probes would indeed be illuminating, a direct quantitative comparison is limited by the unavailability of the posterior chains from other baryon probe analyses. To our knowledge, the only other analysis that explicitly reports the first constrained eigenmode (computed using the same level of rigor as our analysis) is of Dalal et al. (2025). Unsurprisingly, their ability to constrain the gas profile parameters is diluted compared to FRBs because of their increased sensitivity to thermal pressure profile parameters.

Stellar-to-Halo Mass Relation Prior

1. 667: it would be useful if authors provide a reference for this calculation;

We thank the referee for this suggestion. We have added additional panels to Supplementary Fig. 5 to directly demonstrate this calculation (L768).

2. 692: add reference to equation 15 to remind the reader about the functional form of the f_{star} dependence on SHM and its normalization. In addition, the discussion in this and next sections revolves around equation 20, which is, however, defined only in the end of the next section. The authors might consider re-arranging the text such that the baryon budget equation appears earlier. This can improve the flow of the discussion for the reader and clearly show how SHM changes

f_{diffuse} as well.

We thank the referee for these suggestions. We have added a reference to equation 14 (L786). We have also moved the baryon budget equation $f_{\text{diffuse}} = 1 - f_{\text{star}} - f_{\text{ISM}}$ earlier in the text (equation 20), so that the reader has this decomposition in hand before the discussion of how SHM variations affect f_{diffuse} . The detailed partitioning of baryons among different halo mass scales, $f_{\text{diffuse}} = f_{\text{IGM}} + f_{\text{ICM}} + f_{\text{IGrM}} + f_{\text{CGM}}$, appears later in the text as before (equation 21).

Our feedback measurement in the landscape of existing studies

1. The authors should explain how different baryon fraction are calculated. Presumably, this is done by taking the marginalized posterior PDFs from MCMC fits to BCEmu4 or BCEmu7 models and plugging these numbers back into the corresponding halo models. However, the discussion lacks this important information.

We thank the referee for flagging this. We have added a description of how the gas fractions are calculated to the Methods section (L817-821).

Constraining Power across Scales

1. 748-750: discussion explains two different mocks in terms of presence/absence of the variance in DM_{cosmic} , however the labels in sup.fig.2 describe the mocks in terms of extreme/no feedback. For clarity, I suggest the authors expand the labels of the panels in the top row of sup.fig 2 by adding “ $\sigma_{\text{cosmic}}^2 = 0$ ” and “high σ_{cosmic}^2 ”, respectively.

Done (Supplementary Fig. 2).

2. 749: please be mathematically precise and define ‘randomly high variance’.

Done (L619).

3. KL divergence test: the authors do not specify which data they are discussing (one of the mock samples, fiducial observed sample or extended observed sample? Instead it is simply called ‘observed data’). Please describe it. Does the result of the KL divergence test change if fiducial vs extended samples are used?

We thank the referee for this comment. The KL divergence is computed using the full observed FRB sample. The results are not susceptible to the choice of sample subsets – we have verified that the information gain is stable across different sample configurations. We note that, as advised by the referee in an earlier comment, we have removed the fiducial/extended subset distinction from the revised manuscript in favor of a clearer narrative, so this ambiguity no longer arises in the revised text. We have updated the relevant sentence in the Methods to explicitly state that the full

observed FRB sample is used for the KL divergence calculation (L902).

4. In principle, the KL divergence test is not related to the mocks described in the first paragraph of this section. The authors, therefore, might consider splitting the discussion.

Done (L901).

5. Likewise, it is not clear what BCEmu model is referenced when authors are quoting $D_{\text{KL}} \sim 15$ bits at $k \sim 1 \text{ Mpc}^{-1}$. Does this information gain correspond to all halo models?

The quoted information gain of $D_{\text{KL}} \sim 15$ bits at $k \sim 1 h \text{ Mpc}^{-1}$ corresponds to our baseline BCEMU5 model, which is used consistently throughout the Main text. The information gain is comparable across all halo model variants, as can be visually confirmed from the posterior widths shown in Supplementary Fig. 2. We have updated the Methods to explicitly state this (L915-918).

6. KL divergence: the information gain is typically quoted in bits if the base 2 logarithm is used. Eq. 21 seems to imply the natural logarithm instead, which gives the units of ‘nats’. If that is indeed the case, the quoted D_{KL} value must be converted from nats to bits. Please double check the number and correct if needed. Also, in Figure 4 the quoted value is 8 bits - why?

Our calculations use the base-2 logarithm throughout, and we have corrected equation 22 in the revised manuscript to explicitly reflect this. We apologize for the typo in the original submission where we quoted ~ 8 bits in Fig. 5 caption: we note that this statement has been removed from the revised manuscript in response to a comment from another referee, as it was deemed not directly relevant to the main conclusions presented in the figure. We confirm that the KL divergence value (in bits) reported in the Methods section is correct.

7. The bottom row of Sup.Fig.2, showing the gas mass fraction ratios, is not really discussed in the text.

We have added a brief discussion on the same (L622-627).

Survey specifications

1.812-816: please add reference for the quoted parameter values. The projected DSA performance values should also be referenced.

We added references to LSST DESC requirements document and DSA Project White Paper (L964).

2.817: Likewise, the line ‘representative of next-generation FRB surveys’ should be backed up by a reference or calculation.

Done (L968).

Sensitivity forecast

1. 827: the authors should explain the chosen DM-M approximation and the choice of the scaling relation between N and M .

We have added an explanation of both approximations to the revised manuscript (L977-982). We have also corrected a typo (with guidance of Prof. Matthew McQuinn) in the manuscript, where we replaced $M^{2/3}$ by $M^{1/3}$ in the DM scaling relation.

Fisher forecast

1. What is the number of FRBs that were used in 6x2-point analysis? This might give an idea if when such sensitivity can be achieved.

We assume 10^4 localized FRBs overlapping with the Rubin footprint, corresponding to a sky coverage of $\sim 40\%$, which is achievable with DSA (L968-969).

Figures

Figure 1: please add labels $\langle a \rangle$, $\langle b \rangle$, and $\langle c \rangle$ to the panels if you reference them.

Done (Fig. 1).

Figure 2: while the requirement for the PATH probability is indicated ($P_{\text{host}} > 0.9$), the PATH paper is not referenced. Please add it either here or in the corresponding discussion.

Done (L528).

Figure 3: please add labels $\langle a \rangle$, $\langle b \rangle$, and $\langle c \rangle$ to the panels if you reference them.

Done (Fig. 4).

Extended Data Figure 1: σ_{host}^2 is variance, please correct the mathematical notation.

Done (Extended Data Fig. 1 caption).

Extended Data Figure 2: same as previously - be careful with mathematical notation for variance/std.deviation. The last statement is somewhat questionable for parameter $\log M_c$.

We have corrected our usage of standard deviation versus variance (Extended Data Fig. 2 caption). Regarding the last statement, we agree that this requires qualification – we have revised the caption to note that this statement holds for the multi-parameter models, but that the one-parameter model, which varies only $\log M_c$, 1) is known to be an inaccurate representation of the full complexity of the baryonic feedback-induced scale-dependent matter power spectrum suppression, and 2) in light of known parameter degeneracies (equation 19), the posteriors from BCEMU1 can be significantly

biased and should therefore be interpreted with caution (L696-705).

Extended Data Figure 3: the color chosen for the prior is very hard to distinguish from the white background. I suggest you choose a different color scheme.

Done (Extended Data Fig. 3).

Extended Data Figure 5: please add labels «a», «b», and «c» to the panels if you reference them. Similar to previous figure, the color for prior can be hard to see on some computers.

Done (Extended Data Fig. 5).

Supplementary Fig.2: the caption references «panel a», but it is not indicated on the actual plot. Same is true about panels b and c. What's the meaning of different hatches used in the panels in the middle row? Consider making them homogeneous across the panels.

We have now added panel labels (Supplementary Fig. 2). The different hatch styles are for different models, which we have kept consistent across figures – Fig. 5, Extended Data Fig. 3, Extended Data Fig. 4, and Supplementary Fig. 2.

Supplementary Fig. 3: f_{baryons} have not been introduced anywhere in the text. Please, define it.

Done (Supplementary Fig. 3 caption).

Supplementary Fig. 4: please add labels «a», «b», and «c» to the panels if you reference them.

Done (Supplementary Fig. 4).

Code availability

The code archive is supplied with an extensive README file, explaining the main functions and how-to compute different quantities. While I have not reproduced the results of the manuscript myself, the code covers all calculations described in the text. Specific functions somewhat lack description.

We have now added more detailed descriptions of the functions throughout the code.

Referee 3

This manuscript is the latest, and easily most detailed, work that uses fast radio bursts (FRBs) to study the distribution of ionised, baryonic matter in the Universe. This topic addresses many related questions regarding galaxy formation and cosmology - from the σ_8 tension, to feedback processes at different mass scales, to name a few. FRBs offer a unique probe of this physics, since their dispersion measure (which traces the column density of ionised gas along their line of sight) is equally sensitive to low and high values, whereas other probes typically have sensitivity scaling with density or density squared. This work uses a sample of 109 FRBs from a variety of instruments to perform inferences on this distribution, finding halo gas fractions broadly consistent with other measures, such as X-ray observations, and studies based on the Sunyaev-Zeldovich effect. Their results are consistent with other measures, and do not offer significantly improved constraints. However, their methodology is extremely interesting, and links FRB observables to more traditional cosmological methods in a far more direct way than some other works. Ultimately, this is a very technically interesting manuscript, and it absolutely deserves to be published, but I am unsure whether the format of Nature Astronomy is best suited to this kind of work.

Before proceeding to my specific comments, I'll expand on the above.

This manuscript builds a model for the cosmological distribution of ionised gas, based explicitly on models for the gas distribution in Galactic halos. It then fits this to FRB dispersion measure (DM) data as a function of model parameters - in particular, the variance in DM as a function of redshift is the key observable. As the authors explain (very well, I might say), this gives the model one primary degree of freedom for describing the shape of the cosmological DM contribution, being a linear combination of model parameters governing the mass-dependence of the shape of galaxy halo profiles. From a broad perspective, the work is then performing the same kind of fitting to data as Refs 15, 16, and 74. The former two works effectively model variance through distributing gas between halos and the intergalactic medium, while the latter uses a purely parametric description of the DM variance without linking this to physical details.

This current work improves on the above studies in that it gives highly detailed and well-motivated bottom-up descriptions of how gas physics - from star-formation to feedback - leads to a predicted power spectrum of ionised gas, and hence, how that leads to variance in variations in DM for a given z . It therefore adds significantly to the current literature by describing how parameters more familiar to those who study galaxy evolution and cosmology influence FRB observables, and hence acts as a guide to the FRB community on how the latter can influence the former. However, at a fundamental level, it is still turning a single dial (that of variance in FRB dispersion measure) and matching that to observables. While the data used to constrain the models is a larger set than that of previous studies, this is not a great leap forward, and the primary interest comes from the details found in the methods.

To be clear, the results of this paper are still useful - not all the gas probed by FRBs can be probed by other measures - but fundamentally do not add to existing constraints. What the paper does do extremely well is to analyse in detail both the parameter space that FRBs do constrain, and the sensitivity of these constraints to variation in what might be called nuisance parameters. I found the details explored in the Methods section to be extremely interesting, and of high relevance to those in FRB DM studies - the authors are absolutely to be congratulated on their excellent analysis.

In this sense, I found the paper ill-suited to the format of a Nature journal, which focuses on the result much more so than the method. The methodology is of primary interest here, and the methodology is better explained in traditional article format (and in any case, needs to be better explained here - see below). However, the authors try to make a bit too much of too little - the 0.6-0.8 sigma significance of larger gas densities within R200 compared to x-ray models takes lines 160-179, when even that 0.6-0.8 sigma conclusion is tied to other aspects of their model (as a reminder, the authors effectively vary one parameter). Then again, the methodology itself is very good. I leave it to the Editor to make a final call on whether or not this is suitable for Nature Astronomy.

Below, I make specific comments on the manuscript. Most are relatively minor details - the main comment is to give a better overview of the method in the main body.

We thank the referee for their careful reading of the manuscript, and for their thoughtful and encouraging assessment of the novelty and importance of the techniques presented. Several of the broad concerns raised in the opening assessment are shared with another referee, and we respond to them in detail in the common preamble addressed to all referees at the beginning of this document. Below, we address three specific concerns stated in the referee’s opening statement in turn.

[1] On the novelty of the measurement. We respectfully disagree with the characterization that this work represents only an incremental step. We show that 109 FRBs already deliver constraints competitive with instruments that have operated for years with enormous investment in traditional baryon probes (L25-26, L195-199), while accessing a redshift and halo mass regime no current probe can reach to help arbitrate the tensions in systematics-limited X-ray measurements which are often used to calibrate subgrid feedback physics in hydrodynamical simulations (L21-22, L173-183). Beyond this, the work establishes a direct observational link between FRB observables and the matter power spectrum suppression – a quantity of central interest to observational cosmologists – with constraints already capable of informing the S_8 tension in joint weak lensing analyses. We agree, however, that this framing was insufficiently prominent in the original text, and have restructured the Main text accordingly.

[2] On the 0.6–0.8 σ measurements. We agree entirely. We have revised the manuscript to flag this as a tentative hint rather than a definitive conclusion (L191-194, L203-211). We emphasize, however, that this mild preference was not the main result of our work – it was an incidental ob-

servation made in passing. The primary result is that FRB-based constraints on the matter power spectrum suppression and halo gas mass fractions are already competitive with those from traditional baryon probes, while probing a complementary and otherwise inaccessible redshift and halo mass regime (L173-183, L195-199). We have revised the manuscript to ensure this hierarchy of results is clearly communicated, so that the reader is not left with the impression that these $0.6\text{--}0.8\sigma$ differences between probes are being presented as a primary finding.

[3] On the methodology and format. We have substantially revised the Main text (Fig. 3, L112-149) to provide a more complete and intuitive description of the working principles of our analysis framework, so that the methodological contribution is legible to the broader observational cosmology and astrophysics readership. We believe these revisions address the referee’s concern that the method was underexplained relative to its importance.

MAIN MANUSCRIPT

In the main body of the manuscript, the authors jump rather quickly from the very broad overview of the topic - that of halo studies and how FRBs help - to the very specific, e.g. the exact parameters used to describe their halo model. However, there is no broad overview of the principle of their analysis, which is (ordering could be changed, depending on phrasing) - develop a halo model - use this to predict the electron power spectrum, P_{ee} - use this to model the variance of the cosmological component of the DM budget of FRBs - Include Galactic and host DM contributions - fit to the $P(\text{DM}|\text{z})$ distribution of data from 109 FRBs. I found I had to read between the lines to determine this, and even these basics of their approach wasn’t clear until delving into the Methods. This will likely be seen as stating the obvious by the authors, but I believe it is very important to avoid confusing readers. I think this requires re-working some aspects of the paragraph 126-139 to really ensure this broad overview of the method is made clearer.

We thank the referee for this important advice. We have substantially revised the Main text (Fig. 3, L73-74, L112–149) to provide a more complete description of the working principles of our analysis framework in the Main text.

The terminology “Stage IV experiments” is specialized, and should be expanded upon in the main body

We have replaced “Stage IV experiments” with “next-generation experiments” throughout the text.

Line 118-120. FRBs probe “FRBs offer... and probes electron distributions on scales $0.1\text{--}3\text{ h Mpc}^{-1}$ ”. This needs a reference.

The quoted sensitivity ranges are derived from the sensitivity calculations performed in this work to construct Fig. 1. We have added a reference to the Methods and Fig. 1 to make these numbers

for all probes as contributions of this work clearer (L94-95).

L178 Evidence for “with feedback processes likely becoming progressively more efficient at higher redshifts” seems to be based on the insignificant (0.6-0.8 sigma) “preference for systematically larger gas densities within R200 relative to comparable halo model-based analyses of stacked kSZ measurements from ACT DR5 and Compton y-cosmic shear cross-correlations”. This is not worthwhile commenting on further - or if the authors do, then even the word “suggest” is too strong. They should rather note that the results do not constrain the redshift dependence of feedback processes, or at best “hint at a possible redshift dependence”. I want to be clear - it’s good for the authors to note this as a possible cause, but these are < 1 sigma effects, so the authors should take care not to over-sell this.

We thank the referee for raising this concern. We completely agree that $0.6 - 0.8\sigma$ differences between probes constitute at best a hint of a preference for redshift-dependent feedback. We have revised the corresponding discussion in the manuscript to reflect the same (L191-194, L203-211).

The only definition of the “fiducial” and “extended” samples of FRBs I can find is in the caption to Figure 2, but the motivation for these being separated is not clear to me - if they have confident host associations, then surely those FRBs can be included in the main sample? Sample selection should certainly be discussed somewhere in methods, which would allow for further discussion of biasing effects, and what actually contributes to it. For instance, the biases addressed in Ref 32 refer to selection effects against high-DM sources due to e.g. scattering, DM smearing etc and their interaction with the FRB luminosity function; but it does not consider selection effects due to a telescope’s follow-up procedure or localization, or reporting biases. I have not gone through the references for all 109 FRBs used in this work, but presumably they are only taken from “catalog” papers that report a complete sample of FRBs detected by a given instrument in a given time period, rather than cherry-picked “interesting” FRBs published because of their unusual nature. A short note on sample selection should be added to the Methods.

We thank the referee for raising these important points regarding sample selection and potential biasing effects. We have now added “The Localized FRB Sample” section to the Methods (L526-545). Regarding the distinction between the “fiducial” and “extended” samples: we apologize for the confusing terminology in the original manuscript. In the revised manuscript, we define a single “full sample” (109 FRBs, including 22 (sub)arcminute-scale and 87 FRBs (sub)arcsecond-scale localizations with confident host associations) that is used to derive the baseline results (L526-535). The bulk of this sample originates from systematic survey programs with well-defined follow-up strategies, including ASKAP, CHIME/FRB, DSA-110, and MeerKAT, with a smaller number drawn from individual source studies in the literature (L536-540).

Regarding the concern about reporting biases from optical follow-up completeness: our inference is

based on the $p(\text{DM}|z)$ distribution – the probability of observing a given DM at a known redshift – rather than the $p(\text{DM}, z)$ joint distribution. Optical follow-up completeness affect the latter, where the observed FRB rate is a critical component of the analysis framework, but not the former: once an FRB is localized and a spectroscopic redshift is secured, it enters our analysis on equal footing regardless of the survey strategy that produced it. We have added a clear statement of this important distinction to the Methods (L540-545).

204-216: this comparison could also compare the authors’ derived $\sigma^2(z)$ values to those of Ref 74, and the implied baryon spread metric from Ref 75. That would perhaps be an appropriate plot in the Methods.

We thank the referee for this constructive suggestion. Translating our inferred $\sigma^2[\text{DM}_{\text{cosmic}}(z)]$ to the F -parameter formulation of Macquart et al. (2020), we measure $F = 0.20^{+0.02}_{-0.02}$, with a 95% confidence lower limit of $F > 0.16$. This is consistent with the lower limit of $F > 0.13$ reported by Baptista et al. (2024). Converting this F -parameter measurement into the baryon spread metric of Medlock et al. (2025) is currently not possible because the F -parameter calibration from small box sizes (~ 25 Mpc) of CAMELS simulations systematically underestimate the DM variance by missing the contribution from groups and cluster-scale halos (there are hardly any halos with $M_{200} \gtrsim 10^{13} M_{\odot}$ in CAMELS) – precisely the halo mass regime that dominates the DM variance (see Fig. 1) – making their calibration not directly comparable to observations. We have now added a comparison of our inferred DM variance with the measurements of Baptista et al. (2024), highlighting the prospects of the baryon spread metric formulation of Medlock et al. (2025) as an exciting avenue for future work once larger simulation boxes provide a more accurate calibration of the DM variance from group and cluster-scale halos, as it directly link the FRB observables to displacements of particles in hydrodynamical simulations due to feedback (L242-246, L844-855).

Figure 3 a: Currently, ACT kSZ: Bigwood+2024 is repeated in both legends, while FRBs are only given in panel b, which is confusing, so I could not immediately tell what the constraints from this work were.

We thank the referee for flagging this. We have now added legend for FRBs to both panels of Fig. 4 for clarity.

The last paragraph, “Looking ahead,”, provides an excellent summary of what FRBs can do. However, I do not think it belongs in this style of paper. It essentially reviews existing work. I would prefer the current method to be expanded upon in proper detail in the main body.

We thank the referee for this comment. We note that the “Looking ahead” paragraph was added specifically in response to the Editor’s comments on our pre-submission inquiry, where the Editor advised the inclusion of a prognosis for the sort of improvements that can be achieved with upcom-

ing FRB surveys. We therefore conducted sensitivity and fisher forecasts, and opted to retain this paragraph in the interest of addressing the Editor’s advice, but we are happy to defer to the Editor’s discretion on whether it is appropriate for the final published version. We have also taken on board the referee’s broader concern about the level of methodological detail in the main body, and have now included a concise but complete description of the underlying principles of our analysis in the main text, as described in our response to the earlier comment on this subject.

I will note that all the discussion from line 180 onwards is essentially identical to what is already known about feedback scenarios and the matter power spectrum. It is useful in the sense that it tells a reader what the current constraining power of FRBs is. But it also gives a slightly false impression in the sense that the authors are using one observable ($\sigma^2(z)$) and their model is used to tie this to a large range of physics properties of the Universe, so one might imagine scenarios where different models produce different constraints.

We agree that the original discussion could give the impression that the FRB constraints are simply reconfirming what is already known about feedback scenarios from other probes. However, the subtle but crucial point that we have attempted to convey early on in the revised manuscript is not merely that the FRB constraints are consistent with other probes, but that they are derived from a fundamentally independent baryon probe that traces a completely different regime of halo mass, redshift, and physical scale compared to traditional probes such as the kSZ and tSZ effects (Fig.1), and that the constraints from a modest FRB sample are already competitive with other probes that have accumulated their constraining power after decades of observations. This independence is not incidental – it is essential. Every baryon probe carries its own systematic uncertainties, and the pre-eROSITA versus eROSITA tension in X-ray gas mass fractions is a vivid reminder that without independent probes, there is no reliable way to adjudicate between conflicting measurements. Having a completely independent and competitive probe that accesses a complementary halo mass and redshift regime is therefore not merely useful, but indispensable for the multi-tracer analyses required for precision cosmology in the Stage-IV era. We have revised the manuscript to make this contribution and perspective explicit, as detailed in responses to previous questions.

Regarding the referee’s concern that different halo models could produce different constraints from the same observable: at the level of current constraining power, we find that our inferred matter power spectrum suppression is stable across halo model choices – the constraints from HMCODE and the various BCEMU variants of different complexity are all consistent with each other (see Supplementary Fig. 2). This stability reflects the fact that, with current sample sizes, the data effectively constrain only a single well-defined combination of parameters that is robustly captured by all model variants. As constraining power improves with larger FRB samples, however, the impact of model choice will inevitably become more significant, and model-agnostic approaches –

such as the nulling technique proposed by Leung et al. (2025) – will play an increasingly important role in enabling model-independent cosmological inference directly from FRB observables.

METHODS

Overall comment: Reading the methods, I found myself scrolling upwards to check that I hadn't accidentally skipped over a page. The ordering of the methods is quite unclear, and they are not at all described in the main body. Looking back, I suspect this is because there is not a good overview of the methods given in the main body, making it difficult to place the relevance of specific subsections in the Methods. But I do suggest that the authors revisit the ordering of the methods sections for a clearer progression.

We thank the referee for this important feedback. We agree that the original ordering of the Methods section lacked a clear progression, and that the absence of a broad methodological overview in the main body made it difficult for the reader to follow the relevance of individual subsections. We have addressed this in two ways. First, we have completely revised the ordering of the Methods to follow a clearer narrative:

- (i) Sensitivity of Probes to Baryonic Feedback (L460)
- (ii) The Localized FRB Sample (L526)
- (iii) A Halo Model Approach for DM– z Analysis (L546)
- (iv) Validation of Inference Framework (L614)
- (v) Evaluating Bias via Jackknife Resampling (L629)
- (vi) Baryonic Feedback Parameters Constrained by FRB Data (L637)
- (vii) Susceptibility to Fixed Baryonic Feedback and Cosmology Parameters (L706)
- (viii) Anchoring the Stellar-to-Halo Mass Relation with Cluster Observations (L738)
- (ix) DM_{host}: Comparison with Literature and Plausible Redshift Evolution (L856)
- (x) Constraining Power on the Power Spectrum Suppression (L901)
- (xi) Sensitivity and Fisher Forecasts (L920).

Second, we have added a preamble to each Methods subsection that orients the reader and points to the relevant figures or context (L461-464, L527-529, L547-554, L615-617, L630-632, L638-642, L707-711, L739-745, L857-862, L921-925). We have also added references to the Methods at appropriate points throughout the Main text (L82, L94, L110, L128, L138, L160, L164, L169, L170, L172, L265, L273).

Line 519: /citemethods2017MNRAS.465.2936M.

Fixed.

518-531: good discussion of caveats

We are glad that the referee found it useful.

It is not clear in Eq 10 for σ^2 how the $(1+z)^{-1}$ suppression of DM with z is included in this expression. Is that part of the calculation of $P_{ee}(k, z)$? However, normally, $P_{ee}(k, z)$ is not calculated with DM and the $(1+z)^{-1}$ in mind.

The $(1+z)^{-1}$ suppression of DM with redshift is indeed accounted for in equation 3 and is folded into the radial kernel $W_{DM}(\chi)$. Specifically, $W_{DM}(\chi)$ is the integrand of the equivalent comoving-distance form of equation 2:

$$\begin{aligned} \langle DM_{\text{cosmic}}(z_s) \rangle &= \int_0^{\chi(z_s)} W_{DM}(\chi) d\chi, \\ W_{DM}(\chi) &= \frac{3c\chi_e\Omega_b H_0}{8\pi G m_p} \frac{f_{\text{diffuse}}(z)(1+z)}{\sqrt{\Omega_m(1+z)^3 + \Omega_\Lambda}} \frac{dz}{d\chi}. \end{aligned} \quad (1)$$

Here, we have – (i) $(1+z)^{-1}$ suppression of DM with redshift, (ii) $(1+z)^{-1}$ factor from conversion of the differential of proper distance to comoving distance, and (iii) $(1+z)^3$ from the redshift evolution of baryon mass density – resulting in $(1+z)$ factor in the numerator (similar to Macquart et al. (2020) equation 2!). The electron power spectrum $P_{ee}(k, z)$ is therefore the standard power spectrum of the electron density field and does not itself carry any DM-specific $(1+z)$ factors; all such factors are contained within $W_{DM}(\chi)$. This expression has been derived in Reischke et al. (2023) and Sharma et al. (2025) in greater detail, accounting for all redshift factors step-by-step.

Line 546: “The stellar content locked in central galaxy indirectly affect” (sounds strange - please check the meaning)

We thank the referee for flagging this. We have replaced the original phrasing (L590).

Line 548 “The scenario of extreme DM_{host} evolution is unlikely”. What exactly is “extreme” here? A quantifiable example that can be excluded would be good to see. There is not standard “extreme” example - unless there is a paper proposing this? I am also not sure that “If DM_{host} exhibited strong redshift evolution, corresponding trends in [RM and scattering] observables would be anticipated, but are not seen.” is correct. There is so much intrinsic variation in these parameters that even if they were strongly correlated with DM_{host} , one might not detect it; and furthermore, predictions for correlations do not seem to hold out in experimental data. I don’t think this topic is particularly key to the paper however, but the authors should tone down their certainty. After all, the test the authors perform gives a reasonable account of the sensitivity of their results to SFR.

By “extreme DM_{host} redshift evolution”, we specifically mean the scenario in which there is no variance in DM_{cosmic} at all – i.e., all of the observed scatter in the $DM-z$ relation arises from

the host galaxy contribution and its redshift evolution alone. This limiting case comes up frequently in discussions with non-FRB experts and represents a natural worst-case scenario to consider. The most direct evidence against this extreme scenario is the robust detection of excess DM_{cosmic} around intervening halos – through sightline resolved studies of FRBs viewed through galaxy clusters, stacking analyses of FRB sightlines intersecting foreground galaxies, and through cross-correlations of FRB DMs with large-scale structure – which together demonstrate that cosmic structure contributes meaningfully to the observed DM variance. This directly supports our use of DM variance as a probe of feedback-regulated structure formation (L878-880).

We agree with the referee that our statements regarding rotation measures and scattering delays were overstated. As the referee notes, the intrinsic scatter in these observables is large, and the absence of detected trends does not strongly constrain DM_{host} evolution. Since these arguments are not of direct relevance to the main conclusions of this work, we have removed them from the revised manuscript.

Line 551: “exhibit < 4% accuracy” - maybe better phrased as “exhibit < 4% error”?

Done (L602).

Line 557: “the extended FRB sample and the fiducial subset”. What are these samples and subsets? Discussion of DM_{host} redshift evolution is difficult to unpack here, because there is no discussion of how it relates to the samples used. Do the samples come from different redshifts? Are they simulated samples with different truths?

We thank the referee for flagging this confusion. In the revised manuscript, we have removed the fiducial/extended distinction, and replaced it with a cleaner redshift-split test in which the full sample is divided into two equal subsets: 54 FRBs at $z \leq 0.25$ with $\bar{z} = 0.08$, and 55 FRBs at $z \geq 0.25$ with $\bar{z} = 0.37$. The purpose of each subsample and its relation to the DM_{host} redshift evolution test is now explicitly stated (L885-898, Extended Data Fig. 1).

Analysis of the number of effective degrees of freedom in “Constrained Baryonic Feedback Parameters” is excellent. Can the authors interpret a little more their result that primarily it is M_c which is constrained? Following through the effects of M_c on the equations right through to $p(DM|z)$ is a little tricky - perhaps the authors can illustrate how $p(DM|z)$ changes, for 1-2 example z -values, when shifting either M_c alone, or shifting along the principle component vector? (Eq 19) The authors do this nicely with Extended Data Figure 1.

We thank the referee for this excellent advice. Following through the effects of the constrained parameter combination on $p(DM|z_s)$ is indeed non-trivial. Variations of the FRB observable with $\log M_c$ alone are discussed and illustrated in Sharma et al. (2026); here we show variations along the full first CPC mode direction, which additionally captures subdominant contributions from θ_{ej} ,

μ_β and δ . We have added Fig. 3 (L137-149) that illustrates how variations along this constrained direction propagates through the full analysis chain: from the gas profile $\rho_{\text{gas}}(r|M)$ (Fig. 3a), to the gas power spectrum $P_{\text{gas}}(k)$ (Fig. 3b), to the standard deviation $\sigma[\text{DM}_{\text{cosmic}}(z_s)]$ (Fig. 3c), to the conditional probability distribution $p(\text{DM}_{\text{cosmic}}|z_s)$ at two example redshifts (Fig. 3d), to the two physical quantities of interest: the gas mass fractions (Fig. 3e) and the matter power spectrum suppression (Fig. 3f). This figure demonstrates clearly and visually how the sightline-to-sightline variance in FRB DMs is highly sensitive to this single constrained parameter combination, and why it emerges as the dominant constrained direction in our analysis.

Line 648: The sentence “However, the inability of FRBs (and indeed, any probe of baryons) to constrain η due to a lack of constraining power necessitates a better-motivated prior on the SHM relation” needs a bit more explanation. Firstly, being unable to “constrain” something due to a lack of “constraining power” is quite the tautology! Secondly, claiming that FRBs or other baryonic probes are simply unable to constrain η - which governs the slope of the stellar-to-halo mass relation, i.e., the fraction of baryons in stars vs gas as a function of total halo mass - seems unlikely. After all, one can imagine correlating FRB lines of sight against specific halos and statistically untangling the relative contributions as a function of halo mass. The next paragraph begins with “Since we find that the choice of fiducial η impacts the final measurement ...” also seems to contradict the notion that η can’t be constrained.

We thank the referee for this careful reading. On the first point, we apologize for the tautological phrasing. Our intention was simply to convey that current FRB samples are too small to constrain η directly – a limitation of present sample sizes rather than a fundamental barrier (L728-730). Indeed, as we discuss later in the Methods, the ability to precisely constrain f_{diffuse} with larger FRB samples in future will directly translate into constraints on the parameters governing the stellar-to-halo mass relation, including η (L802-811).

On the second point, the key subtlety is that the matter power spectrum suppression depends not only on the gas distribution, but also on stellar physics through the boundary condition $f_{\text{gas}}(M, R \rightarrow \infty) + f_{\text{star}}(M, R \rightarrow \infty) = \Omega_b/\Omega_m$. Even when current FRB data cannot directly constrain η , the posterior on power spectrum suppression retains dependence on our assumed stellar physics priors. While the shifts induced by varying η are currently within 1σ and may appear tolerable, we caution that this will not remain the case as FRB samples grow. With samples an order of magnitude larger, the posterior on power spectrum suppression will tighten substantially, and the same systematic uncertainty from η will become a dominant source of bias – analogous to how baryonic uncertainties have become the limiting systematic for Stage-IV weak lensing analyses that were previously statistics-limited. In this regard, our goal in the current work is to set up a path forward: by joining hands with independent cluster stellar mass fraction measurements, we can currently

mitigate this limitation while FRB samples remain small, paving the way for a fully self-consistent analysis as samples grow (L731-737).

*Line 652 - should refer to *extended data* figure 4*

Fixed (L747).

The parameter-by-parameter discussion of “Stellar-to-Halo Mass Relation Prior” is very useful. We are glad that the referee found it useful.

Caption to supp Figure 5 should remind a reader what eta is.

Done (Supplementary Fig. 5).

Line 660 - missing bracket, or the sentence is cut off

We thank the referee for pointing out the error with compilation, we have now fixed it (L756-760).

Code availability:

*When unzipping the file, I got the following error. Needless to say, this was *not* because my local disk was full. I was using a mac os Tahoe 26.2. The authors should check if there was some error in the file upload to google drive. Most of the rest of the files checked out.*

I have not performed a full install of the codebase, nor run through all individual notebooks to check reproduction of the results - I was not sure how critical that hdf5 file error was above, and did not want to fight to install a run a code with potentially missing info. However, the README file appears clear in how to install and run the code, so I believe this is sufficient for a researcher to be able to check the authors' results - once the unzip process operates correctly.

We are able to reproduce this issue when using the UNZIP command-line tool, which appears to stem from the troublesome file naming. In contrast, extracting the files via a simple double-click works without issue. As a workaround, we have now renamed that hdf5 file. Our code is publicly available at: https://github.com/krittisharma/dmz_spk_sharma2026.

REPLY TO REFEREE’S COMMENTS

Title: Signatures of Suppressed Matter Clustering revealed by Fast Radio Bursts

Manuscript #: NATASTRON-26030375-T

Referee 1

The authors have made substantive improvements to the paper and address practically all my comments.

My original comments regarding whether this paper – which is primarily a methods paper – is suitable for Nature Astronomy in the absence of new and highly constraining astrophysical scientific results stand. The authors provide useful arguments of why they feel it is suitable, but not anything that wasn’t clear at the outset. I feel is ultimately an Editorial decision depending on what the journal aims to be.

That said, I definitely think the paper is important and will be well cited for its methodology and laying out a strategy for future high impact in the field.

Thank you very much for your thoughtful re-review of our manuscript and for your generous assessment of its value to the field. We are grateful for your recognition of the promise of analysis developed here, and thank you again for the time and care you have put into reviewing our work across both rounds.

Referee 2

This is the second referee report for the article “Signatures of Suppressed Matter Clustering revealed by Fast Radio Bursts”. The authors addressed all of my comments and concerns raised in the original report. As it stands right now, the article has been significantly improved: both the clarity of the discussion and presentation of the key results. Therefore, I am happy to recommend it for publication in Nature Astronomy.

We sincerely thank the referee for their positive assessment of our manuscript. We are grateful that the discussion and presentation of our results now meet the referee’s expectations. We thank the referee again for their insightful and constructive review across both rounds, which has meaningfully strengthened this work.

Referee 3

I thank the authors for their detailed response to the queries of myself, and the other two referees. I’d first like to re-iterate again that while the most substantive comments made by myself, and in particular referee 1 also, questioned the general suitability of these results, and their significance, for Nature Astronomy, this was meant in no way as a fundamental critique of the work itself.

However - in response to this critique, the authors first argue that their constraints are remarkable given the few FRBs used - and I agree with that. But this is an argument emphasising the promise of the technique, not the result. Second, the authors argue that FRBs target a critical and complementary redshift/halo mass range, noting x-ray measurements. However, they currently do not provide any significant constraints, only 0.6-0.8 sigma hints, which are not at all significant - i.e. it is about the future promise of the technique. Third, the authors argue that they establish an observational link between FRB observables and matter power spectrum suppression - aside from the fact that this claim probably should go to Ref 40 (Medlock et al), establishing such a link is a methods paper / theoretical study, not an observational result. Fourthly, the authors state that they introduce a new analysis methodology - and I fully agree that this methodology is indeed the value of the manuscript.

Therefore, I disagree with the arguments put forward by the authors on pages 1-4 of their reply as to why it is the results of the manuscript that are of broad significance, and not the method.

However - given that the Editor has decided to pass the original set of comments back to the authors, I am going to drop all concerns to this effect. My opinion was, and still remains, that this work is ill-suited to Nature Astronomy. But the authors have gone to significant effort to update their manuscript, and respond to the particular detailed queries of myself and other referees. Given this, I believe the responsible thing is to continue towards publication in Nature Astronomy, with the aim of making the manuscript as well-suited to the format as possible. And I also want to re-emphasise that I believe the authors to have produced a fantastic work with very interesting results, and it is only the framing of those results, and their suitability for NatAstron, I am critiquing.

Indeed, the authors have done a good job of addressing the specific queries of all referees. I have two further minor points, but also a significant concern regarding the methods used to curate the FRB sample used, which is now described in the new Methods section "The Localized FRB Sample".

We thank the referee for their continued careful engagement with our manuscript across multiple rounds of review, and for taking the time to detail their perspective so thoroughly. We are grateful for their acknowledgment that we have made a genuine effort to address the specific queries raised by all referees, and we deeply appreciate their generous characterization of this work as containing interesting results. We also thank the referee for their constructive approach, and for their continued investment in helping us refine the manuscript. We now turn to the specific points raised below.

- The single sentence from lines 78-28 should be broken up.

We thank the referee for this suggestion. We have broken up the identified sentence into three shorter sentences for improved readability (L78-83).

- The new statements on lines 195-199 about how remarkable the FRB measurement is compares

100 FRBs to an unspecified operational lifetime. For this to make sense to a reader, a quantity of FRBs should be compared to e.g. a quantity of clusters probed by x-ray observations, or an experimental lifetime of x-ray obs to the duration over which these FRBs were detected. However, in general, I think that a better use of such comparisons would be to note that the number of localised FRBs is increasing rapidly, and comparing this to e.g. the number of x-ray clusters expected to be studied, and to make such a comparison in the future prospects section. But this is mostly already done (line 263), so I would just remove these lines.

We thank the referee for this comment. Rather than removing the comparison, we have revised it to make it quantitatively meaningful, as advised by the referee. L199-204 now explicitly compare the ~ 100 FRBs (collected over $\sim$ five years of operations of radio interferometers) to the $\sim 12,000$ optically-confirmed eRASS1 clusters and the millions of DESI Y1 galaxy groups underlying the ACT kSZ constraints.

- I like the new caption to figure 3, which takes a reader progressively through the plots.

We are glad that the referee found it useful.

- The new section explaining the data sample is generally good. However, the statement "we exclude 18 FRBs due to either insecure host galaxy association, evidence for excess DM_{host} , negative DM_{exgal} , or DM_{exgal} less than..." needs to be expanded. Firstly, the authors mention the "secure host association" criteria of $P_{\text{host}} > 0.9$ - presumably, "insecure host galaxy association" just means $P_{\text{host}} < 0.9$? If so, why repeat it? If not, please specify. And the criteria for evidence of excess/deficit DM contributions should be expanded upon. Indeed, to my knowledge, there is no good evidence existing in the current literature. The methodology used by the authors here is critical to avoiding biases in the data! In fact, cherry-picking data points here may invalidate the entire analysis, so the authors need to ensure their choices are well-justified.

We thank the referee for this important comment. We have removed the redundant statement regarding insecure host galaxy association, since this simply corresponds to $P_{\text{host}} < 0.9$, already defined via our secure host association criterion. We also agree with the referee that the indication for excess DM_{host} in almost all of these FRBs is, on reflection, ambiguous in the current literature, with no concrete evidence supporting their exclusion on these grounds. To address this concern and avoid the risk of inadvertently biasing our sample through an under-justified cut, we have included these 5 FRBs in our revised manuscript with updated analysis.

Remarks on code availability: I have successfully built the code according to the installation instructions included in the README from the github site at which the code is based. it appears to be in working order and well-documented.

We are glad to know this, and we thank the referee for taking the time to test our reproduction package and confirm that it builds and runs as intended.