
Signatures of suppressed matter clustering revealed by fast radio bursts

In the format provided by the
authors and unedited

Supplementary Information

DM_{host}: Comparison with Literature and Plausibility of Redshift Evolution

We discuss the inferred host galaxy DM contribution, and its potential sources of systematic uncertainty. We first compare our inferred $\langle \text{DM}_{\text{host}} \rangle$ with values reported in the literature, and discuss how differences in Milky Way DM modeling can account for much of the apparent mild tensions between studies. We then assess the insensitivity of our feedback constraints to a plausible redshift evolution of DM_{host} , as illustrated in Extended Data Fig. 1. While this redshift evolution does not impact current measurements, we discuss why it will need to be carefully modeled as FRB samples grow with extended redshift baselines.

Our inferred mean of the rest-frame host galaxy DM contribution, $\langle \text{DM}_{\text{host}} \rangle = 149_{-19}^{+23} \text{ pc cm}^{-3}$, is consistent with existing observational constraints^{31,32,73}; the differences are all at less than 1σ significance, and shifts of order tens of DM units in the inferred $\langle \text{DM}_{\text{host}} \rangle$ across analyses can arise from variations in the modeling of the Milky Way DM contribution, as explicitly demonstrated by reference⁷³. A cleaner measurement of DM_{host} can in principle 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. Reference³⁶ conducted one such study using a sample of FRBs at redshifts $z \leq 0.2$, finding $\langle \text{DM}_{\text{host}} \rangle \sim 150 \text{ pc cm}^{-3}$. 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 local FRB samples required for such measurements.

Alongside the question of the absolute value of DM_{host} , its potential redshift evolution is an important systematic to consider. If FRBs preferentially trace star-forming galaxies²⁴, one might expect DM_{host} to evolve with redshift in a manner tracking the cosmic star-formation rate history. An extreme scenario in which the observed increase of $\sigma^2[\text{DM}_{\text{exgal}}](z_s)$ with redshift is driven entirely by redshift evolution in DM_{host} is disfavored by direct observational evidence. Excess $\text{DM}_{\text{cosmic}}$ around intervening halos has been robustly detected through sightline resolved studies of FRBs viewed through galaxy clusters⁷⁴, stacking analyses of FRBs on foreground galaxies⁷⁵, and through statistical cross-correlations with galaxies^{48,76} at $\sim 5\sigma$ significance – detections that would not be expected if the variance in DM_{exgal} as a function of redshift were dominated by redshift evolution of DM_{host} alone.

To assess any evidence for redshift evolution of DM_{host} , we split the full sample into two equal subsets: a lower-redshift subset of 57 FRBs at $z \leq 0.25$ with $\bar{z} = 0.08$ and a higher-redshift subset of 57 FRBs at $z \geq 0.25$ with $\bar{z} = 0.37$. We show the posterior distributions from HMCODE halo model inference in Extended Data Fig. 1. The inferred $\langle \text{DM}_{\text{host}} \rangle$ is consistent across the two subsets. When comparing with the posteriors from the full sample analysis, consistent with the findings of reference³², 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 DM_{host} inference (also observed in reference²⁵ Fig. 3). We also examine a model in which the median DM_{host} follows the cosmic star-formation rate history^{55,77,24}. Crucially, $\log T_{\text{AGN}}$ remains stable at the 1σ level across both redshift subsets and under the star-formation-rate-motivated redshift evolution model, indicating that our current constraints on baryonic feedback are insensitive to potential redshift evolution of DM_{host} . However, these results highlight the need to develop the redshift lever arm required to directly constrain DM_{host} evolution, which may become an important systematic for next-generation FRB experiments.

Our feedback measurement in the landscape of existing studies

Given the current constraining power, there are two robust and unbiased probes of feedback strength in the Universe: (1) halo gas mass fractions within R_{200} , and (2) the matter power spectrum suppression at intermediate scales. As illustrated in Supplementary Fig. 5, both of these measurements are essentially insensitive to stellar physics uncertainties. With this in mind, we compare our measurements of gas mass fractions within R_{200} as a fraction of the baryon budget (Ω_b/Ω_m) with previous works in literature. From the BCEMU5 MCMC posterior distributions of the halo gas profile parameters, the gas mass fractions in galaxy groups and clusters are derived by integrating the analytical gas density profile (equation 12) within R_{200} for halos in the relevant mass bins. The posterior samples are propagated through this integral to obtain the 68% confidence intervals on the gas mass fractions reported in Fig. 4 and in the baryon partition results. For clusters ($M_{200} \geq 10^{14} M_\odot$), we infer an intracluster medium baryon fraction of $f_{\text{ICM}} = 0.038^{+0.006}_{-0.008}$, in good agreement with $f_{\text{ICM}} = 0.0375^{+0.005}_{-0.005}$ ³², derived from Chandra and eROSITA cluster mass functions combined with halo baryon fraction relations. At group scales ($5 \times 10^{12} \leq M < 10^{14} M_\odot$), we find $f_{\text{IGrM}} = 0.055^{+0.014}_{-0.033}$, consistent with the Chandra-based estimate $f_{\text{IGrM}} = 0.054^{+0.010}_{-0.010}$ ³². For lower-mass halos ($10^9 \leq M < 5 \times 10^{12} M_\odot$), we measure a CGM fraction of $f_{\text{CGM}} = 0.017^{+0.031}_{-0.014}$, in agreement with reference³². The combined fraction of diffuse gas in halos, $f_{\text{CGM}} + f_{\text{IGrM}} = 0.07^{+0.03}_{-0.04}$, is statistically consistent with reference³¹, who reconstruct the foreground density field of FRBs. Collectively, these results demonstrate coherence between diverse FRB-based measurement techniques.


Adopting our conservative approach to f_{diffuse} , where we impose no external priors on the SHM relation, we obtained $f_{\text{diffuse}} = 0.95^{+0.03}_{-0.04}$. Using baryon budgeting,

$$f_{\text{diffuse}} = f_{\text{IGM}} + f_{\text{ICM}} + f_{\text{IGrM}} + f_{\text{CGM}}, \quad (38)$$


we infer the fraction of baryons beyond R_{200} of halos, $f_{\text{IGM}} = 0.85^{+0.05}_{-0.05}$. This is in excellent agreement with the hydrodynamical simulation-calibrated inference of reference³², despite employing an entirely independent methodology. The modest apparent tension ($\sim 2.3\sigma$) with reference³¹ arises primarily from differences in the treatment of f_{diffuse} . In our conservative analysis, we explicitly model $f_{\text{diffuse}}(z)$ without invoking external priors on the SHM relation. In contrast, reference³¹ adopt a flat prior $f_{\text{diffuse}} \in [0.75, 0.95]$, meaning the lower bound effectively truncates the posterior and can bias the inferred f_{IGM} downward. This under-

scores that when f_{diffuse} is not fitted directly, the choice of prior, particularly its edge, needs to be physically motivated. At the level of halo gas fraction measurements across FRB studies, all results remain consistent; discrepancies in f_{IGM} stem from f_{diffuse} prior assumptions (when f_{diffuse} is not explicitly modeled or fit for during the analyses) rather than the underlying data. These results are encouraging and highlight that future FRB samples with improved statistics will allow us to measure f_{diffuse} directly and precisely, further advancing our understanding of baryon partition in the Universe.


We further contextualize our measurement within the landscape of existing FRB DM variance measurements. Translating our inferred $\sigma^2[\text{DM}_{\text{cosmic}}(z_s)]$ to the F -parameter formulation of reference²¹, we measure $F = 0.21^{+0.03}_{-0.03}$, with a 95% confidence lower limit of $F > 0.17$. This is consistent with the lower limit of $F > 0.13$ reported by reference⁴¹. Converting our F -parameter measurement into the physically intuitive baryon spread metric of reference⁴⁰ – which directly quantifies the displacement of baryons from their initial positions in hydrodynamical simulations due to feedback – is currently limited by the small 25 Mpc box sizes of the CAMELS simulations used to calibrate this metric, which systematically underestimate the DM variance by missing the contribution from massive halos (there are hardly any $M_{200} \gtrsim 10^{13} M_{\odot}$ halos in CAMELS boxes) that dominate the observed DM variance (Fig. 1). This comparison is, however, an exciting avenue for future work, once larger simulation volumes provide a more reliable calibration of the baryon spread metric, incorporating DM variance from massive halos, since it directly links the FRB measurement to the impact of feedback on particles in hydrodynamical simulations.




Supplementary Fig.1: Jackknife resampling validation of matter power spectrum suppression constraints. We perform jackknife resampling, sequentially excluding individual FRBs from the analysis to assess potential bias. This is essential to test whether the current constraining power is driven by a single FRB sightline intersecting a massive galaxy cluster. The resulting posterior constraints (68% confidence interval) remain stable across all jackknifed realizations, demonstrating that our results are not driven by any particular sightline.




Supplementary Fig.2: Validation of the ability of current FRB dataset to distinguish between extreme feedback scenarios. The $\text{DM}_{\text{exgal}} - z$ relation of FRBs in universes with extreme feedback (zero variance in $\text{DM}_{\text{cosmic}}$), no feedback (large variance in $\text{DM}_{\text{cosmic}}$) and feedback strength indicated by current dataset is illustrated in panel **a**. In the two mock universes (extreme and no feedback), for redshifts of observed FRB sample, we generate synthetic DM_{exgal} and fit for the $\text{DM}_{\text{exgal}} - z$ relation to measure the suppression in matter power spectrum ($P_{\text{hydro}}/P_{\text{gravity}}$) and gas mass fractions ($M_{\text{gas},500}/M_{500}$). The posteriors for the two mock scenarios for models, including the hydrodynamical simulations-calibrated halo model (HMcode) and variants of flexible analytical model (BCEmu1, BCEmu4, BCEmu5, BCEmu7), are shown in panels **b** and **c** (68% confidence interval). The mock FRBs from extreme feedback scenario predict a strong suppression/low gas fractions, whereas the mock FRBs from no feedback scenario predict a close to zero suppression/high gas fractions – both consistent with our expectations. This validates the constraining power of current FRB sample.



Supplementary Fig.3: Sensitivity of halo gas, stellar, and baryon mass fractions and matter power spectrum suppression to stellar and gas physics. The baryon mass fraction is defined as $f_{\text{baryons}} = f_{\text{gas}} + f_{\text{star}}$. The parameters governing the stellar-to-halo mass relation (A and η) primarily affect the gas fractions integrated out to infinity, but within R_{200} , gas fractions are largely insensitive to stellar physics. A similar trend is seen for matter power spectrum suppression: the influence of stellar physics at intermediate-scales is below the current sensitivity of baryon probes. In contrast, the parameters controlling the halo gas profile ($\log M_c$, θ_{ej} , μ_β , and δ) have a strong impact on both, the matter power spectrum suppression and the gas fractions within R_{200} . This demonstrates that both, the matter power spectrum suppression and gas fractions within R_{200} are robust metrics of gas physics.



Supplementary Fig.4: An evaluation of the impact of stellar initial mass function on the matter power spectrum suppression measurement when using observed stellar-to-halo mass relation data. The stellar-to-halo mass (SHM) relation measurements from DES, WISE/Spitzer and SPT SZ scaling relation of galaxy clusters (with uncertainties denoting the 68% confidence interval)⁷¹, assuming Chabrier and Salpeter initial mass functions (IMFs), is shown in panel **a**. The corresponding constraints (68% confidence interval) on the matter power spectrum suppression, as illustrated in panel **b**, are robust and unbiased by a change in the normalization (consistent with the two IMFs) of SHM relation. However, we show in panel **c** that the choice of stellar IMF introduces bias in the inferred global stellar (f_{star}) and diffuse (f_{diffuse}) baryon fractions. This motivates the need to directly fit for diffuse gas fraction, without imposing external priors, to avoid biased inference, as larger FRB samples become accessible in near future. Such an inference will also provide direct constraints on the stellar IMF.



Supplementary Fig.5: Sensitivity of gas and stellar mass fractions enclosed within R_{200} as a function of halo mass to stellar physics. The variations of slope of the stellar-to-halo mass (SHM) relation at high mass end, η (upper panels) impacts the gas mass fractions in cluster-scale ($\gtrsim 10^{14} M_{\odot}$) halos, and consequently the small scale ($\gtrsim 1 h \text{ Mpc}^{-1}$) matter power spectrum suppression at $\sim 5\text{-}10\%$ level. Varying the SHM relation normalization parameter, A (middle panels) is equivalent to changing stellar initial mass function (IMF). The impact of stellar IMF on gas mass fractions within R_{200} across a broad halo mass range and matter power spectrum suppression at small scales is at $\sim 1\text{-}5\%$ level, which is below the current constraining power of baryon probes. The variations of SHM slope at low mass end, τ (bottom panels) do not impact the gas mass fractions in group and cluster-scale halos, and consequently the impact on small scale matter power spectrum suppression is at $\lesssim 1\%$ level.