



Publication Year	2022
Acceptance in OA	2022-09-02T08:57:51Z
Title	HERA Phase I Limits on the Cosmic 21 cm Signal: Constraints on Astrophysics and Cosmology during the Epoch of Reionization
Authors	Abdurashidova, Zara, Aguirre, James E., Alexander, Paul, Ali, Zaki S., Balfour, Yanga, Barkana, Rennan, Beardsley, Adam P., BERNARDI, GIANNI, Billings, Tashalee S., Bowman, Judd D., Bradley, Richard F., BULL, PHILIP, Burba, Jacob, Carey, Steve, Carilli, Chris L., Cheng, Carina, DeBoer, David R., Dexter, Matt, de Lera Acedo, Eloy, Dillon, Joshua S., Ely, John, Ewall-Wice, Aaron, Fagnoni, Nicolas, Fialkov, Anastasia, Fritz, Randall, Furlanetto, Steven R., Gale-Sides, Kingsley, Glendenning, Brian, Gorthi, Deepthi, Greig, Bradley, Grobbelaar, Jasper, Haldy, Ziyaad, Hazelton, Bryna J., Heimersheim, Stefan, Hewitt, Jacqueline N., Hickish, Jack, Jacobs, Daniel C., Julius, Austin, Kern, Nicholas S., Kerrigan, Joshua, Kittiwisit, Piyanat, Kohn, Saul A., Kolopanis, Matthew, Lanman, Adam, La Plante, Paul, Lekalake, Telalo, Lewis, David, Liu, Adrian, Ma, Yin-Zhe, MacMahon, David, Malan, Lourence, Malgas, Cresshim, Maree, Matthys, Martinot, Zachary E., Matsetela, Eunice, Mesinger, Andrei, Mirocha, Jordan, Molewa, Mathakane, Morales, Miguel F., Mosiane, Tshagofalang, Muñoz, Julian B., Murray, Steven G., Neben, Abraham R., Nikolic, Bojan, Nunhokee, Chuneeta D., Parsons, Aaron R., Patra, Nipanjana, Pieterse, Samantha, Pober, Jonathan C., Qin, Yuxiang, Razavi-Ghods, Nima, Reis, Itamar, Ringuette, Jon, Robnett, James, Rosie, Kathryn, Santos, Mario G., Sikder, Sudipta, Sims, Peter, Smith, Craig, Syce, Angelo, Thyagarajan, Nithyanandan, Williams, Peter K. G., Zheng, Haoxuan
Publisher's version (DOI)	10.3847/1538-4357/ac2ffc
Handle	http://hdl.handle.net/20.500.12386/32537
Journal	THE ASTROPHYSICAL JOURNAL
Volume	924



HERA Phase I Limits on the Cosmic 21 cm Signal: Constraints on Astrophysics and Cosmology during the Epoch of Reionization

Zara Abdurashidova¹, James E. Aguirre² , Paul Alexander³, Zaki S. Ali¹, Yanga Balfour⁴, Rennan Barkana⁵ , Adam P. Beardsley^{6,7} , Gianni Bernardi^{4,8,9} , Tashalee S. Billings², Judd D. Bowman⁷ , Richard F. Bradley¹⁰ , Philip Bull¹¹ , Jacob Burba¹², Steve Carey³, Chris L. Carilli¹³ , Carina Cheng¹, David R. DeBoer¹ , Matt Dexter¹, Eloy de Lera Acedo³, Joshua S. Dillon^{1,31} , John Ely³, Aaron Ewall-Wice¹ , Nicolas Fagnoni³, Anastasia Fialkov^{14,15} , Randall Fritz⁴, Steven R. Furlanetto¹⁶ , Kingsley Gale-Sides³, Brian Glendenning¹³, Deepthi Gorthi¹ , Bradley Greig^{17,18} , Jasper Grobbelaar⁴, Ziyaad Haldar⁴, Bryna J. Hazelton^{19,20} , Stefan Heimersheim¹⁴ , Jacqueline N. Hewitt²¹ , Jack Hickish¹, Daniel C. Jacobs⁷ , Austin Julius⁴, Nicholas S. Kern²¹ , Joshua Kerrigan¹² , Piyanat Kittiwisit²² , Saul A. Kohn² , Matthew Kolopanis⁷ , Adam Lanman¹², Paul La Plante² , Telalo Lekalake⁴, David Lewis⁷, Adrian Liu^{1,23} , Yin-Zhe Ma^{24,25} , David MacMahon¹, Lourence Malan⁴, Cresshim Malgas⁴, Matthys Maree⁴, Zachary E. Martinot², Eunice Matsetela⁴, Andrei Mesinger²⁶ , Jordan Mirocha²³ , Mathakane Molewa⁴, Miguel F. Morales¹⁹ , Tshagofalang Mosiane⁴, Julian B. Muñoz²⁷ , Steven G. Murray⁷, Abraham R. Neben²¹ , Bojan Nikolic³, Chuneeta D. Nunhokee¹ , Aaron R. Parsons¹, Nipanjana Patra¹ , Samantha Pieterse⁴, Jonathan C. Pober¹² , Yuxiang Qin^{17,18,26} , Nima Razavi-Ghods³, Itamar Reis⁵ , Jon Ringuette¹⁹, James Robnett¹³, Kathryn Rosie⁴, Mario G. Santos^{4,28} , Sudipta Sikder⁵, Peter Sims¹² , Craig Smith⁴, Angelo Syce⁴, Nithyanandan Thyagarajan^{7,13,29} , Peter K. G. Williams^{27,30}, and Haoxuan Zheng²¹

¹ Department of Astronomy, University of California, Berkeley, CA, USA

² Department of Physics and Astronomy, University of Pennsylvania, Philadelphia, PA, USA

³ Cavendish Astrophysics, University of Cambridge, Cambridge, USA

⁴ South African Radio Observatory (SARAO), 2 Fir Street, Observatory, Cape Town, 7925, South Africa

⁵ School of Physics and Astronomy, Tel-Aviv University, Tel-Aviv, 69978, Israel

⁶ Department of Physics, Winona State University, Winona, MN, USA

⁷ School of Earth and Space Exploration, Arizona State University, Tempe, AZ, USA

⁸ Department of Physics and Electronics, Rhodes University, PO Box 94, Grahamstown, 6140, South Africa

⁹ INFN-Istituto di Radioastronomia, via Gobetti 101, I-40129 Bologna, Italy

¹⁰ National Radio Astronomy Observatory, Charlottesville, VA, USA

¹¹ School of Physics & Astronomy, Queen Mary University of London, London, UK

¹² Department of Physics, Brown University, Providence, RI, USA

¹³ National Radio Astronomy Observatory, Socorro, NM, USA

¹⁴ Institute of Astronomy, University of Cambridge, Madingley Road, Cambridge CB3 0HA, UK; heimersheim@ast.cam.ac.uk

¹⁵ Kavli Institute for Cosmology, Madingley Road, Cambridge CB3 0HA, UK

¹⁶ Department of Physics and Astronomy, University of California, Los Angeles, CA, USA

¹⁷ School of Physics, University of Melbourne, Parkville, VIC 3010, Australia

¹⁸ ARC Centre of Excellence for All-Sky Astrophysics in 3 Dimensions (ASTRO 3D), Australia

¹⁹ Department of Physics, University of Washington, Seattle, WA, USA

²⁰ eScience Institute, University of Washington, Seattle, WA, USA

²¹ Department of Physics, Massachusetts Institute of Technology, Cambridge, MA, USA

²² School of Chemistry and Physics, University of KwaZulu-Natal, Westville Campus, Durban, South Africa

²³ Department of Physics and McGill Space Institute, McGill University, 3600 University Street, Montreal, QC H3A 2T8, Canada; jordan.mirocha@mcgill.ca

²⁴ School of Chemistry and Physics, University of KwaZulu-Natal, Westville Campus, Private Bag X54001, Durban 4000, South Africa

²⁵ NAO-CUKZN Computational Astrophysics Centre (NUCAC), University of KwaZulu-Natal, Durban 4000, South Africa

²⁶ Scuola Normale Superiore, I-56126 Pisa, PI, Italy; yuxiang.l.qin@gmail.com

²⁷ Center for Astrophysics | Harvard & Smithsonian, Cambridge, MA, USA; julianmunoz@cfa.harvard.edu

²⁸ Department of Physics and Astronomy, University of Western Cape, Cape Town, 7535, South Africa

²⁹ CSIRO, Space & Astronomy, P.O. Box 1130, Bentley, WA 6102, Australia

³⁰ American Astronomical Society, Washington, DC, USA

Received 2021 August 13; revised 2021 October 5; accepted 2021 October 13; published 2022 January 11

Abstract

Recently, the Hydrogen Epoch of Reionization Array (HERA) has produced the experiment's first upper limits on the power spectrum of 21 cm fluctuations at $z \sim 8$ and 10. Here, we use several independent theoretical models to infer constraints on the intergalactic medium (IGM) and galaxies during the epoch of reionization from these limits. We find that the IGM must have been heated above the adiabatic-cooling threshold by $z \sim 8$, independent of uncertainties about IGM ionization and the radio background. Combining HERA limits with complementary observations constrains the spin temperature of the $z \sim 8$ neutral IGM to $27 \text{ K } \langle \bar{\tau}_s \rangle$ $630 \text{ K } (2.3 \text{ K } \langle \bar{\tau}_s \rangle$ $640 \text{ K})$ at 68%

³¹ NSF Astronomy and Astrophysics Postdoctoral Fellow.



(95%) confidence. They therefore also place a lower bound on X-ray heating, a previously unconstrained aspects of early galaxies. For example, if the cosmic microwave background dominates the $z \sim 8$ radio background, the new HERA limits imply that the first galaxies produced X-rays more efficiently than local ones. The $z \sim 10$ limits require even earlier heating if dark-matter interactions cool the hydrogen gas. If an extra radio background is produced by galaxies, we rule out (at 95% confidence) the combination of high radio and low X-ray luminosities of $L_{r,\nu}/\text{SFR} > 4 \times 10^{24} \text{ W Hz}^{-1} M_{\odot}^{-1} \text{ yr}$ and $L_X/\text{SFR} < 7.6 \times 10^{39} \text{ erg s}^{-1} M_{\odot}^{-1} \text{ yr}$. The new HERA upper limits neither support nor disfavor a cosmological interpretation of the recent Experiment to Detect the Global EOR Signature (EDGES) measurement. The framework described here provides a foundation for the interpretation of future HERA results.

Unified Astronomy Thesaurus concepts: [Reionization \(1383\)](#); [Intergalactic medium \(813\)](#); [Galaxy formation \(595\)](#)

1. Introduction

One of the final frontiers of observational cosmology is the cosmic dawn, during which the first luminous sources formed and grew into galaxies. This era ended with the reionization of the intergalactic medium (IGM), when ultraviolet photons from these sources ionized virtually all of the neutral hydrogen—and hence when stars and black holes affected every baryon in the universe. This constitutes the last baryonic phase transition in the universe’s history and has important implications for later generations of galaxies.

Observations are now beginning to probe this era. Measurements of the large-scale polarization of the cosmic microwave background (CMB) imply that reionization reached its midpoint at $z \sim 7\text{--}8$ (Planck Collaboration 2020; de Belsunce et al. 2021; Heinrich & Hu 2021). Models of Ly α emission lines of galaxies (Stark et al. 2010; Schenker et al. 2012; Jensen et al. 2013a; Caruana et al. 2014; Pentericci et al. 2014; Mesinger et al. 2015; Mason et al. 2018, 2019) and quasars (Mesinger & Haiman 2004; Bolton et al. 2011; Greig et al. 2017; Davies et al. 2018; Greig et al. 2019; Wang et al. 2020; Yang et al. 2020) also suggest a relatively large neutral fraction at $z \sim 7$. While the conventional wisdom has long held that the reionization process ends at $z \sim 6$ (e.g., McGreer et al. 2015; though see Lidz et al. 2006; Mesinger 2010), recent measurements of the Ly α forest suggest that it may continue to somewhat later times (Becker et al. 2015; Bosman et al. 2018; Kulkarni et al. 2019; Keating et al. 2020; Nasir & D’Aloisio 2020; Qin et al. 2021a; Choudhury et al. 2021b).

However, our understanding of this era is still incomplete: models and empirical extrapolations suggest that even the deepest Hubble Space Telescope (and upcoming James Webb Space Telescope) observations probe only a fraction of the total star formation in the early universe (Behroozi & Silk 2015; Mason et al. 2015; Robertson et al. 2015; Furlanetto et al. 2017; Gillet et al. 2020). This could mean that the galaxies providing most of the reionizing photons will remain *unseen*. Moreover, while reionization is the most dramatic effect of the first galaxies, their X-ray and ultraviolet radiation fields can affect the IGM even while it remains neutral—a phase that cannot be observed directly by many cosmological probes.

A complete understanding of the cosmic dawn therefore requires complementary measurements of the IGM gas. The most powerful potential probe is the 21 cm spin-flip line of neutral hydrogen (Field 1959; Madau et al. 1997). The 21 cm line is particularly sensitive to (Furlanetto 2006; Morales et al. 2012; Pritchard & Loeb 2012): (1) structure formation in the universe, which can be observed through density fluctuations; (2) the reionization process, which eliminates the 21 cm signal inside the large ionized bubbles that grow throughout that era; (3) the X-ray background (or other exotic heating or cooling

mechanisms), which likely sets the IGM temperature before reionization and hence determines whether the 21 cm line is seen in absorption or emission; (4) the nonionizing ultraviolet background, as photons that redshift into the hydrogen Ly α transition mix the hyperfine level populations; and (5) the radio background at high redshifts, including the CMB but also potential contributions from astrophysical sources or exotic processes.

Because the spin-flip cosmological signal is very weak compared to other astrophysical radio backgrounds, mapping these IGM fluctuations is extremely challenging, and early efforts to observe it have focused on two complementary directions. One is the “global” all-sky signal, measuring the sky averaged spectral signature of the line, covering Gpc³-sized comoving volumes (Shaver et al. 1999; Muñoz & Cyr-Racine 2021). Several such experiments are underway (Voytek et al. 2014; Price et al. 2018; Singh et al. 2018; Philip et al. 2019; DiLullo et al. 2020). Of these, only the Experiment to Detect the Global EOR Signature (EDGES) collaboration has made a tentative detection (Bowman et al. 2018); although, the cosmological interpretation of the measurement is subject to significant instrumental and systematic uncertainties (e.g., Hills et al. 2018; Bradley et al. 2019; Singh & Subrahmanyan 2019; Sims & Pober 2019; Tauscher et al. 2020). Interestingly, the claimed signal is much stronger than expected, requiring either that the IGM temperature is smaller than allowed by adiabatic cooling (from, e.g., energy exchange with dark matter; Barkana 2018; Berlin et al. 2018; Kovetz et al. 2018; Muñoz & Loeb 2018; Slatyer & Wu 2018), or that an additional radio background (beyond the CMB) is present in the early universe (e.g., Ewall-Wice et al. 2018; Feng & Holder 2018; Pospelov et al. 2018; Fialkov & Barkana 2019; Mebane et al. 2020).

A number of other experiments hope to use interferometers to measure statistical fluctuations in the 21 cm background, most often quantified through the power spectrum, which measures the variance in the field as a function of smoothing scale. Several experiments have now published upper limits from $z \sim 6\text{--}10$, though these limits so far probe only a small fraction of the parameter space spanned by “standard” models of early galaxies (Ghara et al. 2020; Mondal et al. 2020; Ghara et al. 2021; Greig et al. 2021a, 2021b).

Recently in HERA Collaboration (2021, hereafter H21), we presented the first upper limits on the 21 cm power spectrum from a new experiment, the Hydrogen Epoch of Reionization Array (HERA). HERA is now under construction in the Karoo Desert of South Africa (DeBoer et al. 2017). Its phased construction allowed for an initial observing campaign in 2017–2018, the results of which are considered here. Note that we base these results on data from just 39 antennas; HERA is

now expanding to ~ 350 antennas, so the interpretation here provides a framework for improved analyses in the future.

This paper is organized as follows. We introduce the physics of the 21 cm signal in Section 2. Then, in Section 3, we describe HERA’s limits and our inference tools. In Section 4, we use a very simple model to motivate the most important implications of HERA’s upper limit. In the following four sections, we present several complementary interpretations to elucidate these results: we use the 21cmMC code to infer constraints on early galaxy populations and the IGM (Section 5) and a phenomenological model that directly parameterizes IGM properties to better understand the IGM constraints (Section 6). Then, we examine the implications of the HERA limits for exotic dark-matter models (Section 7), and finally we consider constraints derived from models with an enhanced radio background (Section 8). In Section 9, we summarize these results and their implications for the epoch of reionization.

Throughout this work, we assume a standard flat Λ cold dark matter (CDM) cosmology, consistent with the latest CMB measurements (Planck Collaboration 2020). The separate analyses use slightly different cosmological parameters, but these have little effect on our constraints. We denote co-moving megaparsecs with “cMpc.”

2. The 21 cm signal

HERA and other low-frequency instruments aim to observe emission or absorption of the neutral-hydrogen hyperfine transition at an observed wavelength of $\lambda_{\text{obs}} = 21(1+z)$ cm. The intensity of this line is conventionally expressed as the differential brightness temperature, δT_{21} , relative to the low-frequency radio background, which we assume has a brightness temperature T_{rad} at the relevant frequency. Then the brightness of a patch of the IGM can be expressed approximately as (Madau et al. 1997; Furlanetto 2006)

$$\delta T_{21}(\nu) = \frac{T_S - T_{\text{rad}}}{1+z} (1 - e^{-\tau_{\nu}}) \approx T_0 x_{\text{HI}} (1 + \delta) \left(\frac{H}{dv_r/dr + H} \right) \left(1 - \frac{T_{\text{rad}}}{T_S} \right), \quad (1)$$

where $T_0 = 27[(1+z)/10]^{1/2}$ mK is the overall normalization, H is the Hubble parameter at the appropriate redshift, and we have assumed $\Omega_m h^2 = 0.15$ and $\Omega_b h^2 = 0.023$ (with $H_0 = 100h$ km s $^{-1}$ Mpc $^{-1}$). Here, x_{HI} is the neutral fraction of the patch, $\delta = (\rho - \bar{\rho})/\bar{\rho}$ is its fractional overdensity, T_S is the spin temperature (or the excitation temperature of the 21 cm transition), and dv_r/dr is the gradient of the proper velocity along the line of sight.

The spin temperature is determined by (e.g., Madau et al. 1997; Furlanetto 2006; Pritchard & Loeb 2012; Venumadhav et al. 2018)

$$T_S^{-1} = \frac{x_{\text{rad}} T_{\text{rad}}^{-1} + x_c T_K^{-1} + x_\alpha T_K^{-1}}{x_{\text{rad}} + x_c + x_\alpha}, \quad (2)$$

where x_{rad} , x_α , and x_c are coupling constants describing the strength of the relevant interactions. This equation reflects the competition between several processes: (1) interactions with radio photons tend to drive T_S to T_{rad} , with a coupling constant x_{rad} ; (2) collisions drive T_S toward the kinetic temperature of

the gas, T_K with a coupling constant x_c ; and (3) absorption and re-emission of Ly α photons mixes the hyperfine states and also drives T_S toward T_K with a coupling x_α , through a process known as the Wouthuysen–Field effect (Wouthuysen 1952; Field 1958, 1959; Hirata 2006). Meanwhile, the kinetic temperature is affected by the expansion cooling of the IGM and interactions with several radiation backgrounds—most importantly, prior to reionization, any X-ray background generated by early sources. A proper accounting of the temperature requires tracking both the IGM properties and the radiation backgrounds generated by galaxy formation or exotic processes in the early universe.

It is important to note that, in the standard picture, reionization by UV photons is an inhomogeneous process—(nearly) completely ionized regions around the first galaxies expand into (nearly) completely neutral IGM patches as the source population grows. The values of $\bar{x}_{\text{HI}}$ we quote below can therefore be considered as approximately corresponding to the volume-filling factor of the remaining neutral IGM patches during the epoch of reionization (EoR). However, X-rays have much longer mean free paths than UV photons and can deposit their energy in the neutral IGM, partially ionizing and heating that phase, so the relation between the true neutral fraction and the filling factor of the ionized bubbles is not exact.

Given the sensitivity of current experiments, the focus of interferometric observations to date has been on measuring the spatial power spectrum of the 21 cm signal,

$$\langle \tilde{\delta T_{21}}(\mathbf{k}_1) \tilde{\delta T_{21}}(\mathbf{k}_2) \rangle = (2\pi)^3 \delta^D(\mathbf{k}_1 + \mathbf{k}_2) P_{21}(\mathbf{k}_1), \quad (3)$$

where tildes denote Fourier transforms, angular brackets denote ensemble averages, and δ^D is the Dirac delta function. We will typically plot $\Delta_{21}^2(\mathbf{k}) \equiv k^3 P_{21}(\mathbf{k}) / (2\pi^2)$, with units of mK 2 .

The velocity term in Equation (1) accounts for the mapping between redshift and real space, which is complicated by redshift-space distortions (RSDs; Kaiser 1987; Bharadwaj & Ali 2004; Barkana & Loeb 2005a). Crudely, overdense regions expand more slowly than the average universe, so they appear compressed along the radial direction, while under-dense regions appear larger in that direction. Because these distortions occur only along the line of sight, they make the power spectrum anisotropic. The modes used in the HERA analysis are mostly aligned along the line of sight, and care must be taken when comparing to the theoretical models, as we discuss further below.

3. HERA Phase I Power Spectrum Limits

Next we establish the formalism that will be used to interpret our observables. Section 3.1 describes HERA’s data products that are used in this paper, and Section 3.2 defines the likelihood that links these observable quantities to theoretical models. The goal is therefore to provide the necessary machinery to interpret our measurements in a model-agnostic way before introducing our theoretical models in subsequent sections.

3.1. Observational Campaign

The power-spectrum upper limits analyzed in this paper have been published in H21. Here, we describe some of the essential features of the data for convenience, but we refer the reader to H21 for more details.

The upper limits relevant to the paper are reproduced in Figure 1. These were based on 18 nights of data (Julian Dates




Figure 1. Reported limits on the 21 cm power spectrum from H21 used to place constraints on the various models explored in this work (top panel). The bottom panel shows the derived window functions of the limits, showing a peaked sensitivity with compact support around each k mode. Note that in the present analysis we only include every other k bin from the limits quoted in H21 in order to mitigate the effect of nonzero covariance between neighboring k modes. We start this decimation at $k = 0.179 \text{ cMpc}^{-1}$ and $k = 0.134 \text{ cMpc}^{-1}$ for Bands 1 and 2, respectively. Filled points represent positive measurements, and error bars without points represent negative measurements. The error bars show $\pm 1\sigma$ uncertainty.

2,458,098 to 2,458,116) taken as part of an observing campaign from 2017 October to 2018 April when HERA was in its Phase I observing configuration. In Phase I, HERA observed with “hybrid” antenna elements, which consisted of HERA’s 14-m parabolic antennae with modified cross-dipole feeds and a front-end from the Precision Array to Probe the Epoch of Reionization (PAPER) experiment (Parsons et al. 2010; DeBoer et al. 2017). HERA Phase I also inherited PAPER’s back-end system, which processed 100 MHz of bandwidth from 100–200 MHz. For these observations, HERA consisted of 52 operating antennas, 39 of which were deemed science-ready after passing our data quality metrics (H21). Note that these 52 antennas make up a small fraction of the experiment at full capacity of ~ 350 antennas, which will observe from 50–225 MHz (Dillon & Parsons 2016; DeBoer et al. 2017).

The analysis and reduction of these data are discussed in H21 and in several supporting papers in more detail (Kern et al. 2020b; Dillon et al. 2020; Kern et al. 2020a; Aguirre et al. 2021; Tan et al. 2021). For the purposes of this work, the important takeaway is that, while nearly the full band is processed in the data reduction pipeline, only two portions of the band are largely free of radio frequency interference, which sets the redshift ranges studied in this work (Band 2, centered at $z = 7.9$, and Band 1, centered at $z = 10.4$). Additionally, the power spectra (PS) studied in this work come from only one of the fields reported in H21. Because HERA observes in a drift scan mode, it surveys a $\sim 10^\circ$ wide stripe centered on decl. -30.7° . However, to avoid the brightest portions of the sky (including foregrounds from our Galaxy as well as bright sources such as Fornax A), H21 made further cuts to the data in local sidereal

time (LST). This yields three fields’ (with LST ranges from 1.25 to 2.7 hr, 4.5 to 6.5 hr, and 8.5 to 10.75 hr) worth of data that were propagated through to the power-spectrum pipeline. The parameter inference discussed in this work comes solely from the limits presented from the first cut (Field 1; see Figure 1 in H21), as these showed the least amount of foreground contamination and therefore produced the most stringent limits.

For the $z \sim 8$ band, the data presented in H21 provide the most sensitive upper limits on the 21 cm power spectrum to date, improving upon previous limits at that redshift by roughly one order of magnitude. Another important feature of the H21 analysis is that they report measurements consistent with the thermal noise floor at intermediate and high Fourier k wavevectors. The dynamic range between that noise floor and the peak measured foreground signal is $\sim 10^9$ in power, in spite of the fact that they perform no explicit foreground subtraction in their analysis. Upper limits on the 21 cm power spectrum are also currently best constrained by the Murchison Widefield Array (MWA) at lower redshifts (Trott et al. 2020) and by Low Frequency Array (LOFAR) at higher redshifts (Gehlot et al. 2019; Mertens et al. 2020).

3.2. Data Likelihood

To relate our power-spectrum measurements to theoretical models, we first group our data at all k bins and redshifts into a column vector, i.e.,

$$\mathbf{d} = \begin{pmatrix} \Delta_{21}^2(k_1, z_1) \\ \Delta_{21}^2(k_1, z_2) \\ \Delta_{21}^2(k_2, z_1) \\ \Delta_{21}^2(k_2, z_2) \\ \vdots \end{pmatrix}, \quad (4)$$

which has a length of $N_d = N_k \times N_z$. In this work, we use the power-spectrum data tabulated in H21 Tables 3 and 4 for Field 1 only, spanning a k range of $0.13\text{--}0.64 \text{ cMpc}^{-1}$ and the two redshift bins $z = 10.4$ and $z = 7.9$. Furthermore, we also make use of the associated window function and covariance matrices, which are included with the data and will be publicly accessible. In this work, we assume the thermal noise on the data to be Gaussian distributed and thus adopt a Gaussian likelihood. This is a fair approximation as the large amounts of averaging performed in the analysis Gaussianize the data due to the central limit theorem. Having adopted a model $\mathcal{M}$ for the cosmic 21 cm signal (e.g., one of the simulations described in later sections), $\mathbf{m}$, and a model for any extant systematics, $\mathbf{u}$, we can write the probability distribution for the data given the parameters (i.e., the likelihood function) as

$$\mathcal{L}(\mathbf{d}|\theta, \mathcal{M}, \mathbf{u}) \propto \exp\left(-\frac{1}{2}\mathbf{r}(\theta, \mathbf{u})^T \mathbf{\Gamma} \mathbf{r}(\theta, \mathbf{u})\right), \quad (5)$$

where $\mathbf{r}(\theta, \mathbf{u}) = \mathbf{d} - \mathbf{u} - \mathbf{W}\mathbf{m}(\theta)$, θ are the parameters of $\mathcal{M}$, $\mathbf{m}$ is the simulation’s deterministic prediction of the data vector mean given θ , $\mathbf{W}$ is the $N_d \times N_d$ window function matrix of the data,³² and $\mathbf{\Gamma} = \mathbf{\Sigma}^{-1}$ is the $N_d \times N_d$ precision matrix, which is

³² In general, the window function matrix can be of shape $N_d \times N_m$, where N_m is the number of k bins predicted by the model. In our case, we estimate the window function along with the data power spectrum and discretize into the same space.

the inverse of the covariance matrix of the data. The window functions account for the corrections to the predicted mean vector due to the telescope measurement and data reduction process (see Tegmark 1997; Liu & Tegmark 2011; Dillon et al. 2014; Liu & Shaw 2020; Kern & Liu 2021); in other words, it is the point-spread function of the power-spectrum measurement in Fourier k space. The covariance matrix accounts for the variance of the measured power spectrum and the correlation of that uncertainty between band powers, irrespective of non-thermal systematics. This covariance is assumed to be diagonal given the analysis methods in H21 (see Section 3.2.1 for details). The on-diagonal elements (i.e., the variances) are estimated using antenna auto-correlation data to model the instrument noise. Because the power spectrum is a quadratic statistic, the sky signal enters in various signal-noise cross-terms even if our variance model is due entirely to instrumental noise. For this contribution, it is the total sky signal (including foregrounds) that matters, and we model this using the empirically measured power spectrum, as detailed in Tan et al. (2021). Note that while we write Σ as model independent, there are some terms that can be model dependent, and thus it can take on an explicit dependence on θ (cosmic variance, for example, is dependent on the amplitude of the predicted mean signal). For the current limits, we do not expect cosmic variance to be important (H21).

Ultimately, one is interested in the probability distribution of the parameters θ given the data $\mathbf{d}$, i.e., the posterior probability distribution $p(\theta, \mathbf{u}|\mathbf{d}, \mathcal{M})$. This is related to the likelihood via Bayes' theorem:

$$p(\theta, \mathbf{u}|\mathbf{d}, \mathcal{M}) \propto \mathcal{L}(\mathbf{d}|\theta, \mathcal{M}, \mathbf{u})p(\theta|\mathcal{M})p(\mathbf{u}), \quad (6)$$

where $p(\theta|\mathcal{M})$ is our prior distribution on the parameters, and $p(\mathbf{u})$ is our prior on the systematics (assumed to be independent of the physical parameter prior).

3.2.1. Marginalizing Over Systematics

The likelihood as expressed in Equation (5) (and thus the posterior in Equation (6)) has a dependence on the systematics, $\mathbf{u}$. In this paper, we have no explicit way of modeling $\mathbf{u}$, so we desire a likelihood that is dependent only on the astrophysical parameters. This suggests *marginalizing* over the prior range of the unknown systematics. In principle, we would express $\mathbf{u} = \mathbf{u}(\phi)$, i.e., we would have some physically motivated set of parameters ϕ that produce a set of systematics in $\Delta_{21}^2(k, z)$, and we would marginalize the posterior over these parameters. In the absence of such a physically motivated model, we marginalize directly over the binned values $\mathbf{u}$:

In particular, taking a multivariate uniform prior on $\mathbf{u}$ gives

$$p(\theta|\mathbf{d}, \mathcal{M}) \propto p(\theta|\mathcal{M}) \int_{\mathbf{u}_{\min}}^{\mathbf{u}_{\max}} \exp\left(-\frac{1}{2}\mathbf{r}(\theta, \mathbf{u})^T \mathbf{\Gamma} \mathbf{r}(\theta, \mathbf{u})\right) d\mathbf{u}. \quad (7)$$

Note the assumptions that have been made in obtaining this expression. Here, our multivariate uniform prior on $\mathbf{u}$ allows each k and z bin to vary independently, thus allowing random fluctuations of arbitrary form. Although this is the form we employ for this paper, future analyses would be considerably

improved with detailed physical models for systematics that might be present, for example by imposing smoothness priors (in k and/or z) when appropriate.

If we also assume that $\mathbf{\Gamma}$ is diagonal, and writing $\mathbf{t} = \mathbf{d} - \mathbf{W}\mathbf{m}(\theta)$, then this equation reduces to

$$\begin{aligned} p(\theta|\mathbf{d}, \mathcal{M}) &\propto p(\theta|\mathcal{M}) \int_{\mathbf{u}_{\min}}^{\mathbf{u}_{\max}} \exp\left(-\frac{1}{2}\sum_i^{N_d} (t_i - u_i)^2 \Gamma_{ii}\right) d\mathbf{u} \\ &= p(\theta|\mathcal{M}) \prod_i^{N_d} \int_{u_{i,\min}}^{u_{i,\max}} \exp\left(-\frac{[t_i - u_i]^2}{2\sigma_i^2}\right) du_i, \end{aligned} \quad (8)$$

where the second line follows due to the separability of the factors in u_i , and $\sigma_i = (\mathbf{\Gamma}^{-1/2})_{ii}$ is the standard deviation of d_i . In this paper, we utilize band powers that are widely separated in wavenumber (see below) so that a diagonal covariance matrix is a good approximation.

In order to provide a systematics-marginalized likelihood, we must choose prior ranges for the systematics on each (k, z) bin. Allowing for unbounded (possibly negative) systematics would not allow us to constrain the cosmological signal, as the systematic would be completely degenerate with the model. Thus, we should look to our understanding of the data analysis process to set this prior. Calibration errors causing residual phase differences or chromatic effects can lead to biases that are positive or negative. Though negative biases or systematics in the power spectrum have been observed in previous experiments (see, e.g., Kolopanis et al. 2019), for the present H21 data set and analysis pipeline, the most likely causes of such issues have been mitigated by the application of absolute calibration and other improvements; large negative detections are not observed in null tests or validation simulations. The most likely remaining source of systematic bias is un-modeled signal chain chromaticity common to all elements, and this would couple *positive* foreground power beyond the wedge. Given this expectation of positive-only systematics, we set the prior constraint that $\mathbf{u} \geq 0$, yielding

$$p(\theta|\mathbf{d}, \mathcal{M}) \propto p(\theta|\mathcal{M}) \prod_i^{N_d} \frac{1}{2} \left(1 + \operatorname{erf}\left[\frac{t_i}{\sqrt{2}\sigma_i}\right]\right), \quad (9)$$

where erf is the error function. It is worth making clear that this form of the posterior is relatively flat once $t_i \gtrsim \sigma_i$. Since t_i represents the data minus the theory model, in effect this means that our posterior produces close to equal probability for any scenario in which the model is less than or equal to measured values (within error bars). Our treatment of systematics therefore leads to a well-defined posterior that naturally treats data points as “upper limits.” This result is the same form as that derived in Appendix B of Ghara et al. (2020) in the interpretation of LOFAR data. A similar derivation of the marginal upper-limit likelihood can also be found in Appendix A of Li et al. (2019).

If the off-diagonal components of $\mathbf{\Gamma}$ are not zero, the integral of Equation (7) is not tractable in closed form. For the HERA data used in this work, we specifically use band powers that are widely separated in k such that their error correlations are

negligibly small; concretely, we use only every second k bin (“decimation”). Quantitatively, Figure 20 of HERA Collaboration (2021) shows an example of the normalized covariance between k bins, demonstrating that after decimation, the remaining modes have negligible covariance, on the order of 1%–2%. In decimating, we could in principle choose either the even or odd k bins from each band, and as each of the four choices would be a slight underestimation of the constraints, we choose the combination providing the strongest limits. This includes $k = 0.17 \text{ cMpc}^{-1}$ in Band 1 and $k = 0.13 \text{ cMpc}^{-1}$ in Band 2; as a matter of convention, we refer to the former as even and the latter as odd k bins. As we will show later (see Figures 3 and 19), the constraints on realistic models are primarily driven by the two most stringent limits, so we can expect the decimation to have a negligible effect.

3.2.2. “Inverse” Likelihood

In practice, given that the upper limits presented in H21 are still roughly two orders of magnitude above fiducial 21 cm models, the majority of the parameter space for standard models is left unconstrained. One way to illustrate how the new limits help is by combining them with existing constraints. Alternatively, to provide a clearer picture of the model parameter choices that *exceed* the HERA limits, we also consider an “inverse likelihood” defined as

$$\mathcal{L}^{\text{inv}}(\mathbf{d}|\theta) \equiv 1 - \frac{\mathcal{L}_m(\mathbf{d}|\theta)}{L_0} \quad (10)$$

where L_0 is the maximum of $\mathcal{L}_m$.³³ With the inverse likelihood, the resulting marginalized distributions identify the parameter combinations that can be ruled out by the HERA limits alone (see Figure 4 for an example). However, these distributions must be treated with caution: models that lie inside of the projections of the full distribution are not necessarily excluded. The inverse likelihood should only be used to gain intuition about the utility of the HERA limits and parameters that are necessary (but not sufficient) to drive a power spectrum beyond the HERA limits.

4. Building Physical Intuition: A Density-driven Bias Approach

Before studying galaxy-driven models, we begin with a simple bias analysis, which will allow us to build intuition about the implications of the HERA measurements. As these limits are well above predictions of “vanilla” models of the reionization era, the most important parameter that can be constrained is the IGM temperature, as the temperature ratio term in Equation (1) can become arbitrarily large for gas that is very cold. In the spirit of simplicity, throughout this section we will assume efficient Wouthuysen–Field coupling ($T_K = T_S$), sourced by a nonionizing ultraviolet background from early star formation. We then infer constraints from the most stringent HERA k bin at each redshift, which we will interpret in terms of changes to the gas kinetic temperature (i.e., we will set $T_{\text{rad}} = T_{\text{CMB}}$).

Let us begin by supposing that the 21 cm power spectrum traces the matter power spectrum,

$$\Delta_{21}^2(k, z) = b_m^2(z) \Delta_m^2(k, z), \quad (11)$$

which is appropriate when the fluctuations are sourced by the matter fluctuations. The key assumption here is that the bias parameter is scale-independent—which is exact for ionization and temperature that vary linearly with density, and can be extended beyond this approximation (McQuinn et al. 2005). We then use the HERA measurements to constrain the bias parameter b_m . We compute the linear matter power spectrum from CAMB³⁴ (Lewis & Bridle 2002) at each z , and we find the 95% confidence level (CL) limits of $b_m < \{156, 529\}$ mK for $z = \{7.9, 10.4\}$ (an analogous analysis for the relative-velocity power spectrum can be found in Appendix A). These can be translated into lower limits on the ratio of the spin-to-radio temperatures through the relation

$$b_m = T_0 x_{\text{HI}} \left\{ \left(1 + \mu \right) - \frac{T_{\text{rad}}}{\bar{T}_S} [(1 + \mu) - C_T] \right\} \quad (12)$$

derived from Equation (1) (see, e.g., Pritchard & Loeb 2008), where C_T is the adiabatic index, which accounts for the preferential cooling of under-dense regions; and μ is the line-of-sight cosine of the wavenumbers observed, which accounts for RSDs. Here, and throughout this text, an overbar represents average over volume. We obtain the adiabatic index $C_T = \delta T_K / (T_K \delta) \approx 0.6$ as a function of z following Muñoz et al. (2015a), which we correct for kinetic temperatures above the adiabatic threshold ($T_K > T_K^{\text{ad}}$) by writing $C_T(T_K) \rightarrow C_T \times \min(1, T_K^{\text{ad}}/T_K)$, assuming homogeneous heating. We further assume negligible ionizations, setting $x_{\text{HI}} = 1$, and for RSDs, we take spherically averaged modes ($\mu = 0.6$) through this section to match the common procedure done in simulations. We will show how the constraints shift when altering these two assumptions later in Section 7.

Under our assumptions, the b_m upper limits translate into lower limits for the spin temperature of

$$\bar{T}_S \geq \{7.8, 1.9\} \text{ K} \quad \text{for } z = \{7.9, 10.4\} \quad (13)$$

at 95% confidence, where we reemphasize we have assumed $x_{\text{HI}} = 1$. We show these limits in Figure 2, along with the adiabatic-cooling prediction in the standard CDM model. These $\bar{T}_S$ values have interesting implications for the thermal state of the IGM at high redshifts. As is clear from Figure 2, the HERA Band 2 ($z = 7.9$) 95% confidence limit is above the adiabatic-cooling prediction, which demands that some heating must have occurred before $z = 7.9$. Moreover, the HERA limits for Band 1 ($z = 10.4$), while below the adiabatic limit at that z , can be used to clarify the state of the IGM in comparison with the claimed EDGES detection (also shown in Figure 2), which we will explore in Section 7.

We emphasize that these limits rest on three strong assumptions, which we will highlight here and will, in the following sections, explore with more physics-rich models. First, the limits assume full Wouthuysen–Field coupling ($T_S = T_K$), which is all but guaranteed by the redshifts we

³³ This typically is the likelihood of a model power spectrum equal to zero. Including L_0 here makes sure $\mathcal{L}^{\text{inv}}$ is independent of the normalization of $\mathcal{L}_m$, e.g., of the number of data points used.

³⁴ <https://camb.info/>

consider ($z \leq 10.4$; see Section 5). Second, they assume a value of $x_{\text{HI}} = 1$, which can be varied at each z in Equation (12), though only homogeneously. Lastly, in this analysis we have performed a spherical average of RSDs, whereas HERA data mostly contains modes along the line of sight ($\mu \approx 1$, see Section 2). Properly accounting for RSDs can result in stronger limits, as we will show in Section 7.

In summary, the bias approach here outlined is useful for building intuition, although reionization models and observations suggest that the spatial fluctuations in the ionization field (rather than the matter field) should drive the 21 cm signal at $z \sim 8$. We will explore such models in detail in the following sections, but for now, we show the limit from 21cmMC in Figure 2. We describe below how this limit was obtained, but we see already that there is general agreement ($\sim$ factor of few) with the density-driven bias limit at $z \approx 8$ (indeed our density-driven bias limit is very close to the analogous density-driven 21cmMC limit, denoted by the red contours in Figure 5). The reason these two approaches yield similar results—despite their vastly different assumptions about the EoR—is that the density and ionization power spectra (PS) are of the same magnitude at $z \sim 8$ and $k \sim 0.1 \text{ Mpc}^{-1}$ (e.g., Furlanetto et al. 2004). Astrophysical models can only modify the peak power during the EoR by a factor of a few (e.g., Greig & Mesinger 2015). The only way to reach the power-spectrum amplitudes probed by the HERA limits is by having a large $\sim(1 - T_{\text{rad}}/T_S)^2$ pre-factor, i.e., requiring $\bar{T}_S \ll T_{\text{rad}}$. In this regime, model differences can be easily compensated by relatively small changes in $\bar{T}_S$. Therefore, in the regime of current HERA limits, constraints on $\bar{T}_S$ are of the same magnitude whether the 21 cm power spectrum tracks density or ionization fluctuations.

5. Galaxy and IGM Properties Inferred from HERA Observations

We next consider the HERA limits in light of “standard” galaxy formation models using data-constrained 21cmFAST semi-numerical simulations.

5.1. Galaxy-driven Models of the Cosmic 21 cm Signal

Here we briefly summarize how the 21 cm signal is computed using the galaxy-driven models of 21cmFAST³⁵ (Mesinger & Furlanetto 2007; Mesinger et al. 2011; Murray et al. 2020). The main ansatz of these 21 cm models is that *cosmic radiation fields are sourced by galaxies, hosted by dark-matter halos* (whose relation to the large-scale matter field is comparably well understood). We generate Eulerian density and velocity fields with second-order Lagrangian perturbation theory (2LPT; e.g., Scoccimarro 1998). Galaxy properties are then assigned to dark-matter halos via scaling relations with halo mass. In Section 6, we explore toy models in which radiation fields are not directly associated with galaxies in order to study the robustness of our inferences and the “value-added” by explicit models of structure formation.

Specifically, we use the empirical galaxy relations of Park et al. (2019), capable of reproducing the observed UV luminosity functions of galaxies during the EoR ($z = 6\text{--}10$), as well as the spatial distribution of IGM opacities seen in Ly α forest spectra at $z = 5\text{--}6$ (Qin et al. 2021a). Consistent with semi-analytic models and hydrodynamic simulations of high- z




Figure 2. Lower limits on $\bar{T}_S$ from HERA data (95% CL, green and purple arrows) compared to the upper limit from EDGES (blue arrow). The green HERA limits have been obtained by assuming that the IGM is fully neutral and at a constant temperature (aside from small fluctuations due to adiabatic expansion). The purple HERA limit is from full galaxy models at $z = 7.9$ with 21cmMC (see Section 5); note that by construction, the 21cmMC models cannot cool below the adiabatic prediction, so we do not show a 21cmMC limit at $z = 10.4$. Both assume spherically symmetric RSDs. The black line shows the standard-model prediction assuming full Wouthuysen-Field coupling ($T_S = T_K$) without any heating. The HERA Band 2 data can rule out adiabatic cooling in both approaches, requiring some heating to take place before $z = 7.9$.

galaxies (e.g., Moster et al. 2013; Mutch et al. 2016; Sun & Furlanetto 2016; Xu et al. 2016; Tacchella et al. 2018; Behroozi et al. 2019; Yung et al. 2019; Ma et al. 2020), we describe the mean stellar-to-halo-mass relation, M_*/M_h , with a power law:

$$\frac{M_*}{M_h} = f_{*,10} \left(\frac{M_h}{10^{10} M_\odot} \right)^{\alpha_*} \left(\frac{\Omega_b}{\Omega_m} \right), \quad (14)$$

where (Ω_b/Ω_m) is the mean baryon fraction, and the stellar fraction, $f_* = f_{*,10} (M_h/10^{10} M_\odot)^{\alpha_*}$, is restricted to be between 0 and 1. The corresponding star formation rate assumes a characteristic star formation timescale that scales with the Hubble time, H^{-1} (which, during matter domination, is equivalent to scaling with the halo freefall time): $\dot{M}_* = M_*/(t_* H^{-1})$. Furthermore, we assume only a fraction $f_{\text{duty}} = \exp[-M_{\text{turn}}/M_h]$ of halos host galaxies; the free parameter M_{turn} encodes the mass scale below which inefficient cooling and/or feedback suppresses efficient star formation (e.g., Hui & Gnedin 1997; Springel & Hernquist 2003; Okamoto et al. 2008; Sobacchi & Mesinger 2013; Xu et al. 2016; Ocvirk et al. 2018; Ma et al. 2020).

We then compute the galactic emissivities (soft UV, ionizing UV, and X-ray), assuming they scale with the star formation rates. We identify ionized regions with an excursion set approach (Furlanetto et al. 2004), comparing the cumulative (local) numbers of emitted photons and recombinations. We slightly adjust the number of emitted photons to correct for the nonconservation of ionizing photons in excursion set algorithms (e.g., Zahn et al. 2007; Paranjape & Choudhury 2014; for details see J. Park et al. 2021, in preparation). Sub-grid IGM recombinations are tracked according to Sobacchi & Mesinger (2014). We assume Population II stellar spectral energy distributions (SEDs) for the ionizing and soft UV emission, corresponding to ~ 5000 ionizing photons produced per stellar

³⁵ <https://github.com/21cmfast/21cmFAST>

baryon (e.g., Leitherer et al. 1999; Barkana & Loeb 2005b).³⁶ A fraction $1 - f_{\text{esc}}$ of these photons is absorbed within the galaxy itself, and does not reach the IGM. We allow the ionizing escape fraction to also scale with the halo mass:

$$f_{\text{esc}}(M_h) = f_{\text{esc},10} \left(\frac{M_h}{10^{10} M_\odot} \right)^{\alpha_{\text{esc}}}, \quad (15)$$

where $f_{\text{esc},10}$ is the normalization, and α_{esc} is a power-law index. The ionizing escape fraction is also restricted to values between 0 and 1. Although there is currently no consensus on the ionizing escape fraction or its dependence on galaxy properties, simulations suggest that such a generic power law is an acceptable characterization of the population-averaged values (e.g., Paardekooper et al. 2015; Kimm et al. 2017; Lewis et al. 2020).

In contrast to ionizing UV photons, the soft UV and X-ray photons responsible for coupling the gas and spin temperatures and heating the gas can have long mean free paths through even the neutral IGM. We follow the corresponding ionization and heating rates for each simulation cell by integrating the specific emissivities back along the light cone, attenuated by the corresponding opacities. Our simulations track the spatial fluctuations in the X-ray and Lyman series backgrounds, with the IGM opacity computed assuming a standard “picket-fence” absorption for Lyman series photons and absorption from partially ionized hydrogen and helium in a two-phased IGM for X-ray photons (e.g., Mesinger et al. 2011, 2013; Qin et al. 2020a). The X-ray SED emerging from galaxies is approximated as a power law whose luminosity scales with the star formation rate (SFR). This is consistent with theoretical models and observations of local star-forming galaxies, for which X-ray emission is dominated by high-mass X-ray binaries (HMXBs) and/or the hot interstellar medium (ISM; e.g., Mineo et al. 2012; Fragos et al. 2013; Brorby et al. 2014; Pacucci et al. 2014; Lehmer et al. 2016). Specifically, we parameterize the typical emerging X-ray SED of high- z galaxies via their integrated soft-band (<2 keV) luminosity per SFR (in units of $\text{erg s}^{-1} M_\odot^{-1} \text{yr}$),

$$L_{X<2 \text{ keV}}/\text{SFR} = \int_{E_0}^{2 \text{ keV}} dE_e L_X/\text{SFR}, \quad (16)$$

where L_X/SFR is the specific X-ray luminosity per unit star formation escaping the host galaxies in units of $\text{erg s}^{-1} \text{keV}^{-1} M_\odot^{-1} \text{yr}$, taken here to be a power law with energy index α_X , and E_0 is the minimum energy for X-rays to be able to emerge from the galaxy and not be absorbed locally in the ISM. For reference, the typical value of $E_0 \sim 0.5 \text{ keV}$ found in the simulations of Das et al. (2017) corresponds to an HI column density of $\sim 10^{21.4} \text{ cm}^{-2}$, assuming zero metallicity.

In summary, our 21cmFAST galaxy models have nine free parameters:

1. $f_{*,10}$, the normalization of the stellar mass–halo mass relation, evaluated at $M_h = 10^{10} M_\odot$;

2. α_* , the power-law index of the stellar mass–halo mass relation;
3. $f_{\text{esc},10}$, the normalization of the ionizing escape fraction–halo mass relation, evaluated at $M_h = 10^{10} M_\odot$;
4. α_{esc} , the power-law index of the ionizing escape fraction–halo mass relation;
5. M_{turn} , the characteristic halo-mass scale below which the abundance of active galaxies is exponentially suppressed;
6. t_* , the characteristic star formation timescale, expressed in units of the Hubble time;
7. $L_{X<2 \text{ keV}}/\text{SFR}$, the soft-band X-ray luminosity per unit SFR;
8. E_0 , the minimum X-ray energy of photons capable of escaping their host galaxies; and
9. α_X , the energy power-law index of the X-ray SED.

We emphasize that this flexible galaxy parameterization used in 21cmFAST enables us to set physically meaningful priors over the free parameters and use high- z galaxy observations in our inference. For instance, the common simplification of a constant stellar-to-halo-mass relation is inconsistent with galaxy SFR and luminosity function (LF) observations, and can thus bias parameter inference (see Mirocha et al. 2017; Figure 1 in Park et al. 2019). Our galaxy model therefore allows us to use existing high- z observations, in addition to HERA, when computing the model likelihood (see Section 5.4). This quantifies the “added value” of HERA, given that existing observations already exclude a significant prior volume (e.g., Park et al. 2019). Without them, our posterior would strongly depend on our priors.

5.2. Inference

To perform Bayesian inference, we use 21cmMC³⁷ (Greig & Mesinger 2015, 2017a, 2018) with the recently implemented Multinest-based (Feroz et al. 2009) sampler (Qin et al. 2021b; see also Binnie & Pritchard 2019). For a given sample of astrophysical parameters, we compute 4D realizations of the 21 cm signal in a cubic volume with a periodic boundary condition and a length of 250 cMpc. The initial conditions and 2LPT are calculated on a 512^3 grid, while the final radiation fields are computed on a 128^3 grid. Choosing a line-of-sight axis, we account for nonlinear RSDs via the real-to-redshift space sub-grid transformation described in Greig & Mesinger (2018), and first introduced in Mao et al. (2012) and Jensen et al. (2013b).³⁸

When evaluating the likelihood according to Equation (5), we add in quadrature a conservative 20% modeling error (see Zahn et al. 2011) as well as the sample variance from our simulation. In contrast to other simulation-based inference codes, 21cmMC forward models 4D realizations of the 21 cm signal. We compute the PS on-the-fly from these 4D realizations *without* emulators, over our 9D parameter space; we therefore do not include emulator error/bias in our likelihood.

³⁷ <https://github.com/21cmfast/21CMMC>

³⁸ We note, however, that we spherically average the model PS before comparing with the data, which does not match the line-of-sight selection performed by the HERA analysis. In regimes dominated by density fluctuations, the HERA mode selection can substantially enhance the power (La Plante et al. 2014; Pober et al. 2015; Jensen et al. 2016). However, as we will see in Section 5.4, these regimes are excluded by current observations requiring reionization to be underway at $z \sim 8$. We therefore do not expect our main conclusions in Section 5.4 to be impacted by these selection effects; nevertheless, in future analysis, we will compare the forward-modeled PS to the data in a like-to-like fashion, using the same mode sampling.

³⁶ Specifically, the Population II SEDs were generated with the Starburst99 code (Leitherer et al. 1999), assuming a Scalo (1998) IMF and 0.05 solar metallicity. The spectra in the Lyman bands were interpolated using broken power laws between each Lyman transition according to Barkana & Loeb (2005b).




Figure 3. Two examples of galaxy-driven models that are ruled out by HERA 2021 limits. The rows correspond to Band 1 ($z = 10.4$; top) and Band 2 ($z = 7.9$; bottom). The columns correspond to (from left to right): (i) the IGM density; (ii) the brightness temperature of a “reionization-driven” model ($\bar{x}_{\text{H I}} = 0.73$, $\bar{T}_S = 1.42$ at $z = 7.9$; black patches are cosmic H II regions); (iii) brightness temperature of a “density-driven” model ($\bar{x}_{\text{H I}} = 0.98$, $\bar{T}_S = 1.99$ at $z = 7.9$); and (iv) the corresponding power spectra together with the H21 limits. Note that the power spectra of the models are much flatter (with k) than the observational limits; thus the constraining power is entirely provided by the two filled squares at low k . In the bottom-right panel, we also show power spectra ignoring redshift-space distortions (RSDs). RSDs are important for the density-driven models but much less so for the reionization-driven models. The slices are 1 cGpc on a side and 2 cMpc thick and were generated with 21cmFAST v3.

When we include other observational constraints in our inference procedure (see Section 5.4), we calculate the total likelihood with $\mathcal{L}_{\text{total}} = \mathcal{L}_m \times \mathcal{L}_{\text{LFs}} \times \mathcal{L}_{\text{DF}} \times \mathcal{L}_{\tau_e}$, where the last three terms reflect the comparison between the modeled results against (i) the observed faint galaxy ($M_{\text{UV}} < -20$) UV luminosity functions at $z = 6-10$ from Bouwens et al. (2015, 2016) and Oesch et al. (2017); (ii) the upper limit on the neutral-hydrogen fraction at $z \sim 5.9$ measured by the dark fraction on high-redshift quasar spectra (McGreer et al. 2015), $x_{\text{H I}} < 0.06 + 0.05(1\sigma)$ where we consider a one-sided Gaussian likelihood function;³⁹ and (iii) the Thomson scattering optical depth of CMB photons, using Planck Collaboration (2020) data analyzed by Qin et al. (2020a), $\tau_e = 0.0569^{+0.0081}_{-0.0066}$.

5.3. Models That Exceed the HERA Limits

In this section we highlight the astrophysical models *disfavored* by the current HERA limits. To do this, we use the inverse likelihood from Equation (10). Because the inverse likelihood is only illustrative, we also confine the analysis to the two most stringent limits at $z = 7.9$ and $z = 10.3$. In any case, these two data points provide all of the constraining power because the observed limits rise much more steeply with k than the model predictions. This allows us to compare to similar analysis of recent LOFAR

and MWA data, which also used an inverse likelihood and the same galaxy models (Greig et al. 2021a, 2021b).

Before showing the full distribution of models, in Figure 3 we show examples of two classes of models capable of exceeding the HERA upper limits. The top row corresponds to Band 1 ($z = 10.4$) and the bottom to Band 2 ($z = 7.9$). Slices through the density field and 21 cm brightness temperature fields are shown on the left, with the 21 cm PS shown together with the data in the rightmost panels. For visualization purposes, the maps are generated from larger boxes than used in the inference, corresponding to 1 cGpc on a side, but with the same 2 cMpc resolution.⁴⁰

The 21 cm PSs in the two classes of models exceeding these upper limits are driven by spatial fluctuations in either: (i) the IGM ionized fraction, which we will refer to as “reionization driven” (also referred to as “cold reionization” in the literature; e.g., Mesinger 2014); or (ii) the gas density, which we will refer to as “density driven” (see Section 4, and Greig et al. 2021a for the same qualitative result using recent MWA limits).⁴¹ Both scenarios *require* a cold IGM, which sets a lower limit on the heating rate (and hence on the X-ray emissivity within these models).

³⁹ A revision of these dark fraction limits from a larger QSO sample (S. Campo et al. 2022, in preparation) as well as inference from the large-scale Lyman forest opacity fluctuations (Qin et al. 2021a) seem to favor a slightly later end to reionization (a delay of $\Delta z \sim 0.5$). Since reionization-driven PS amplitudes are maximized around the midpoint of the EoR, which for current observations occurs right around HERA’s Band 2 at $z \sim 8$, we expect that shifting the EoR toward later times could slightly weaken the HERA constraints we derive below.

⁴⁰ We confirm that the PSs in the 1 cGpc and 250 cMpc runs are converged to the percent level or better for the relevant wavenumbers, $k > 0.1 \text{ cMpc}^{-1}$. This level of convergence is consistent with the results of Kaur et al. (2020), who quantified the bias and scatter in the 21 cm signal resulting from missing large-scale modes (see also Iliev et al. 2014), and is orders of magnitude smaller than the observational uncertainties.

⁴¹ Ghara et al. (2020) also consider a model in which highly biased active galactic nuclei with luminous, soft X-ray SEDs but negligible UV emission dominate the radiation background. Such extreme scenarios might also produce very strong temperature fluctuations, capable of exceeding the HERA limits; however, such models are not inside our prior volume.

In the right panels, we confirm that the 21 cm PSs of both scenarios are flatter than the observational limits. Thus when the observational limits are consistent with thermal noise, the constraining power comes entirely from the deepest limits (see Mertens et al. 2020; Trott et al. 2020; Ghara et al. 2021; in our case, primarily the deepest data point at $z = 7.9$).

In the bottom-right panel, we also show how the PSs depend on RSDs. For the reionization-driven model, RSDs are not important since the first H II regions are highly biased, zeroing out the signal from the densest regions with the strongest RSDs (Mesinger et al. 2011; Jensen et al. 2013a; Ghara et al. 2015; Ross et al. 2021). However, the density-driven models have a negligible contribution from ionization and heating, with the 21 cm power spectrum driven entirely by the nonlinear matter field. By comparing the solid and dashed red curves, we see that nonlinear RSDs can boost the spherically averaged power by factors of ~ 2 – 3 , in excess of the linear prediction of 1.87 (e.g., Bharadwaj & Ali 2004; Barkana & Loeb 2005a). Indeed without RSDs, this density-driven model is consistent with the data at $\sim 1\sigma$. We explore the effect of different RSD assumptions for density-driven models in Section 7.1.

In Figure 4 we show a corner plot corresponding to the inverted likelihood from Equation (10). We caution that our parameter ranges in this figure/subsection do not correspond to a “prior” belief of the distribution of disfavored models, and marginalizing over an inverse likelihood is different from an inversion of the 2D marginalized Bayesian posteriors. Therefore Figure 4 should not be interpreted as a Bayesian posterior of disfavored models, and it is difficult to formally relate it to the normal likelihood results in the next section. However, the figure illustrates where the models that exceed HERA reside in our parameter space. In the top right, we draw from these distributions the redshift evolution of the mean neutral fraction, the mean 21 cm signal, and 21 cm power spectrum at $k = 0.13 \text{ cMpc}^{-1}$.

Here we highlight the two modes discussed above: red and blue curves denote the density-driven and reionization-driven models, respectively, classified on the basis of whether the universe is mostly neutral or mostly ionized at $z = 6$.⁴² The shaded regions enclose 68% of the distributions. Astrophysically, the two modes are most easily distinguished by the ionizing escape fraction parameter, $f_{\text{esc},10}$, and to a lesser degree by their star formation efficiencies, here parameterized by the ratio $f_{*,10}/t_*$. All of the models require that the IGM was not heated significantly, as seen by the upper limits on the X-ray luminosity per SFR, $L_{X,<2 \text{ keV}}/\text{SFR}$.

The upper-right panels show that the density-driven models are already ruled out by other observations, since they fail to reionize the universe early enough. In particular, we show the observed upper limit on the neutral fraction from the dark pixels in the Lyman forests (McGreer et al. 2015), as well as the Compton scattering optical depth from Planck 2018 (Qin et al. 2020a). Note that these observations were not used in computing the inverse likelihood. However, some of the

reionization-driven models are consistent with current observations. We return to this in the next subsection.

In Figure 5, we show where these HERA-disfavored models sit in the marginalized 2D space of $\bar{x}_{\text{H I}}$ versus $\bar{T}_S$.⁴³ The left and right panels correspond to $z = 7.9$ and 10.4. The two modes discussed above are clearly seen to emerge by $z = 8$; at present, the lower-redshift data provide most of the constraining power.

At $z \approx 8$, we see that HERA-disfavored models have low spin temperatures: $\bar{T}_S \lesssim 3 \text{ K}$ (or more generically for any uniform radio background $\bar{T}_S/\bar{T}_{\text{radio}} \lesssim 0.1$ for $0.1 \lesssim \bar{x}_{\text{H I}} \lesssim 0.9$). These constraints are somewhat tighter than analogous ones based on recent LOFAR (Mertens et al. 2020) and MWA (Trott et al. 2020) upper limits: $\bar{T}_S \lesssim 2$ – 2.5 K , over narrower ranges in $\bar{x}_{\text{H I}}$ (see Figure 4 in Greig et al. 2021b and Figure 6 in Greig et al. 2021a). Thus, as expected from the stronger PS upper limits, the H21 limits rule out more models than previous power-spectrum limits. Furthermore, the density-driven modes were not ruled out by the previous LOFAR limits, which had a larger amplitude and were performed at a higher redshift ($z = 9.1$; at which the adiabatic-cooling temperature is larger by a factor of $(1 + 9.1)^2/(1 + 7.9)^2$).

At $\bar{x}_{\text{H I}} \gtrsim 0.9$, the range of temperatures for the disfavored models broadens. This is due to the negative contribution of the ionization-density cross power term, which dominates the large-scale 21 cm power in this regime (Lidz et al. 2008; Zahn et al. 2011). The first galaxies drive H II regions that are very biased in the early stages of the EoR. These quickly cover up the largest matter over-densities, which had earlier dominated the 21 cm power spectrum. Thus for models with negligible temperature fluctuations, the large-scale power drops in the early stages of the EoR before rising again as it transitions from being sourced by the matter fluctuations to ionization fluctuations.

5.4. How Do the HERA Limits Improve Upon Previous Complementary Data?

As already mentioned, many of the models that are disfavored by the current HERA limits are already inconsistent with existing observations of the $z > 6$ universe. Here we put the HERA constraints in context with these other observations by computing the Bayesian posterior over our parameter space with and without the new HERA limits. In particular, we run two inferences (see also Section 5.2):

1. *without HERA*: This run corresponds roughly to our current state of knowledge, without including 21 cm observations.⁴⁴ As detailed in Section 5.2, the likelihood

⁴³ Note that the values of $\bar{T}_S$ we quote throughout Section 5 are averaged only over the neutral IGM component (T_S is undefined for ionized gas). Because in the standard picture reionization is approximately “inside-out” on large scales, averaging over the neutral IGM means that the $\bar{T}_S$ limits are slightly biased toward under-dense volumes.

⁴⁴ Here we restrict ourselves to arguably the most model-independent EoR constraints. In the future, as the 21 cm data improves, we will fold in additional constraints from Ly α emitting galaxies, QSO damping wing analysis, opacity fluctuations in the Lyman forests, and the patchy kinetic Sunyaev–Zeldovich effect (e.g., Stark et al. 2010; Schenker et al. 2012; Pentericci et al. 2014; Becker et al. 2015; Mason et al. 2018; Bañados et al. 2018; Bosman et al. 2018; Reichardt et al. 2021; Wang et al. 2021). These require more subtle modeling of associated systematics, but could have a non-negligible impact on the recovered EoR history (e.g., Greig & Mesinger 2017b; Dai et al. 2019; Choudhury et al. 2021a; Qin et al. 2021a). We also do not include previous 21 cm upper limits from MWA and LOFAR since these weaker PS limits would not change our *with HERA* posterior (see the PS evolution inset in Figure 6). Thus by comparing *without HERA* and *with HERA*, we highlight the impact of 21 cm measurements.

⁴² There is a clear bimodality in the $z \sim 6$ neutral fraction of disfavored models (see top-right panel of Figure 4), allowing us to easily distinguish the density-driven and reionization-driven modes. Models with intermediate values $\bar{x}_{\text{H I}}(z = 6) \sim 0.5$ would generally have $\bar{x}_{\text{H I}}(z = 8) \sim 0.1$ – 0.2 . In this early EoR regime, the negative contribution of the ionization-density cross-correlation can result in a decrease of large-scale 21 cm power (e.g., Lidz et al. 2008; Zahn et al. 2011), making it difficult for those models to exceed the HERA limits. Thus the highest power is achieved when only one variable is dominating the fluctuations, and the cross-terms can be ignored.




Figure 4. Distribution of models disfavored by **H21**, calculated using the inverse likelihood (Equation (10)) and using only the two most constraining data points in Bands 1 and 2 (see filled squares in Figure 3). The 1D and 2D marginalized distributions were generated by assuming flat priors over the ranges shown by the figure; we caution that these marginalized inverse likelihood results should not be interpreted strictly as a “posterior,” but instead they serve to illustrate where the models disfavored by **H21** sit in astrophysical parameter space. In the bottom left panels, we show the 2D and 1D distributions, while in the top-right panels, we show the EoR history, global signal evolution, and power-spectrum evolution at $k = 0.13 \text{ cMpc}^{-1}$. Red/blue curves denote density-driven/reionization-driven models, classified according to the value of the neutral fraction at $z = 6$. Shaded regions enclose 68% of the disfavored models for each mode. In the power-spectrum evolution plot, we also show the two **H21** data points used to compute these distributions (note that the Band 1 data point is at a slightly higher wavenumber of $k = 0.17 \text{ cMpc}^{-1}$). This highlights that the Band 2 ($z = 8$) data point has all of the constraining power. In the EoR history panel, we also include the QSO dark fraction upper limits from McGreer et al. (2015; empty square). In the bottom-right panels, we also include the corresponding probability density functions (PDFs) of the CMB optical depth, τ_e , from both modes; the gray region spans 68% CL of the observed value, implied by the galaxy-model recovery of Planck Collaboration (2020) EE power spectra described in Qin et al. (2020a). These two EoR observations were not used in the inverted likelihood; unlike the density-driven modes, the reionization-driven modes are largely consistent with these limits.

incorporates observations of (i) the galaxy UV luminosity functions at $z = 6\text{--}10$; (ii) the upper limit on $\bar{x}_{\text{H I}}$ at $z \sim 5.9$, inferred from the Lyman forest dark fraction; and (iii) the CMB optical depth τ_e .

2. *with HERA*: Here the likelihood is computed using both the complementary observations in *without HERA* above, as well as the HERA limits from Bands 1 and 2. Specifically, we use the (regular) HERA likelihood as defined in Equation (5).




Figure 5. Distributions of $\bar{x}_{\text{HI}}$ and $\bar{T}_S$ corresponding to the two classes of HERA-disfavored models from Figure 4, obtained using the inverse likelihood. The left (right) panel corresponds to $z = 7.9$ (10.3). The adiabatic-cooling limit is shown with the dashed horizontal line, and the two example models from Figure 1 are denoted with “x” and “+.” The left axis corresponds to the usual assumption of the radio background being dominated by the CMB, while the right axis translates these values into the more generic ratio, $\bar{T}_S/T_{\text{rad}}$, valid for any homogeneous radio background. For $\bar{T}_S$, we perform averaging only over the neutral IGM (specifically those cells with $x_{\text{HI}} > 0.95$).

5.4.1. Galaxy Properties: Disfavoring X-Ray Faint Galaxies with HERA

The corner plot of these two posteriors is shown in Figure 6, with tan (purple) denoting *without HERA* (*with HERA*). As discussed in detail in Park et al. (2019), we see that current observations (*without HERA*) already rule out a significant fraction of our prior volume, which highlights the power of our 21cmMC approach’s inclusion of complementary galaxy observations. Observations of high- z UV luminosity functions shown in the top-middle subpanels of Figure 6 constrain the stellar-to-halo-mass relation and its scaling with halo mass ($f_{*,10}$ and α_*), as well as place an upper limit on the characteristic turnover scale (M_{turn}). On the other hand, observations of the EoR timing through the CMB optical depth (see bottom-right subpanel) and the Lyman forest dark fraction (see upper limit in the EoR history subpanel) constrain the ionizing escape fraction normalization ($f_{\text{esc},10}$) to within 1 dex and place very weak constraints on its evolution with halo mass (α_{esc}). Using such complementary observations in the likelihood is especially important when sampling from a high-dimensional parameter space with flat priors, for which most of the prior volume is sourced by “extreme” corners of parameter space that are already ruled out by existing observations (as is immediately evident from Figure 6).

Comparing the *without HERA* and *with HERA* posteriors, we see that the H21 limits do not have a notable impact over most of the astrophysical parameter space. The new models that HERA rules out, discussed in the previous section, occupy a modest prior volume.⁴⁵

However, note that the three X-ray parameters ($L_{X<2\text{ keV}}/\text{SFR}$, E_0 , and α_X) are largely unconstrained by the complementary observations over our prior ranges, because none of the *without HERA* observations are sensitive to the IGM temperature, the observable most strongly affected by the X-ray emissivity. In this part of parameter space, HERA does

have a notable impact by ruling out models with weak X-ray heating, which in our parameterization, is predominately determined by the integrated soft-band X-ray luminosity to SFR, $L_{X<2\text{ keV}}/\text{SFR}$. The exclusion of these models is also evident in the 21 cm panels at the upper right, where the recovered signal ranges decrease significantly when including HERA data.

We show a zoom-in of the marginalized 1D probability density functions (PDFs) of $L_{X<2\text{ keV}}/\text{SFR}$ in Figure 7. The marginalized *without HERA* posterior is consistent with the flat prior over the range shown. Current observations do not constrain this quantity aside from disfavoring extreme values of $L_{X<2\text{ keV}}/\text{SFR} \gtrsim 10^{42} \text{ erg s}^{-1} M_{\odot}^{-1} \text{ yr}$, which is so large that X-rays can significantly contribute to reionization (e.g., Mesinger et al. 2013), making it too early in many models. However, the *with HERA* posterior is able to rule out the lower end of this range, resulting in a 68% highest posterior density (HPD) confidence interval of $L_{X<2\text{ keV}}/\text{SFR} = \{10^{40.2}, 10^{41.9}\} \text{ erg s}^{-1} M_{\odot}^{-1} \text{ yr}$. H21 is the first observation to place constraints over this range; the analogous analysis of MWA and LOFAR observations (see Figure 1 in Greig et al. 2021b and Figure 2 in Greig et al. 2021a) disfavored models with lower luminosities.⁴⁶

In Figure 7 we also compare the *with HERA* limits with estimates based on HMXBs, thought to be the dominant X-ray sources in high- z galaxies (e.g., Fragos et al. 2013). The left vertical line denotes the average value observed from HMXBs in local, metal-enriched, star-forming galaxies (Mineo et al. 2012; see also, e.g., Lehmer et al. 2010). Because the HMXB luminosity increases with decreasing metallicity (e.g., Basu-Zych et al. 2013; Douna et al. 2015; Brorby et al. 2016), we do not expect the first, metal-poor galaxies to sit on the left side of this line. And indeed, this local scaling relation is outside of the *with HERA* 68% confidence interval; thus, HERA data already suggests that the first galaxies were more X-ray luminous than their local counterparts. In contrast, the right vertical line in Figure 7 corresponds to the theoretical result from Fragos et al. (2013) for a metal-free HMXB population, expected to be more

⁴⁵ We use a narrower prior range on $L_{X<2\text{ keV}}/\text{SFR}$ and M_{turn} in Figure 6 compared to Figure 4. This is because Figure 6 is a true posterior requiring physically reasonable prior ranges, which we discuss further below when presenting galaxy inference. In contrast, Figure 4 is only meant to illustrate where HERA-disfavored models are expected to reside in our parameter space.

⁴⁶ This comparison is only approximate, because the earlier analyses were based on the inverse likelihood rather than the proper marginalized posterior shown here.




Figure 6. Posteriors with and without the H21 limits. The 1D and 2D marginalized posteriors are shown in the bottom left panels, while the corresponding UV LFs, EoR histories, global 21 cm signal, evolution of the power spectrum at $k = 0.13 \text{ cMpc}^{-1}$, and the CMB optical depth are shown in the top-right panels (clockwise from the top middle). The *without HERA* posterior (tan) is computed using previous observations: (i) galaxy UV LFs from $z = 6 - 10$ (filled squares in LF panels); (ii) upper limit on $\bar{x}_{\text{H I}}$ from the QSO dark fraction (filled square at $z = 5.9$ in the EoR history panel); and (iii) CMB optical depth from Planck (shaded region in the τ_e panel). The *with HERA* posterior (purple) uses the HERA limits from Figure 1 in addition to (i)–(iii). Although we use all data points in the HERA likelihood, we show the two deepest limits from Band 1 ($k = 0.17 \text{ cMpc}^{-1}$) and Band 2 ($k = 0.13 \text{ cMpc}^{-1}$) in the PS evolution inset panel. Here, for comparison, we also show the recent 1σ limits at $k = 0.1 \text{ cMpc}^{-1}$ from MWA (pentagons; Trott et al. 2020) and $k \approx 0.1$ (0.05) cMpc^{-1} from LOFAR (upper/lower circle; Mertens et al. 2020); the MWA and LOFAR limits are not included in the likelihood. We assume flat priors over the astrophysical parameter ranges shown in the subplot axes. This figure illustrates two important points: (i) current observations already exclude a large majority of our prior volume; (ii) HERA limits constrain the X-ray luminosities of the first galaxies.

representative of the first galaxies. Our recovered 1D posterior of $L_{X,<2\text{keV}}/\text{SFR}$ supports theoretical predictions (e.g., Fragos et al. 2013) and the observed evolution with metallicity and redshift (Basu-Zych et al. 2013; Douna et al. 2015; Brorby et al. 2016; Lehmer et al. 2016) that this quantity increases toward high redshifts.

Finally, we caution that our limits on $L_{X,<2\text{keV}}/\text{SFR}$ could weaken if alternate heating mechanisms play a significant role. Although we include adiabatic, ionization, X-ray, and Compton heating/cooling, in some extreme models, alternate heating sources could dominate. These could include shock heating (e.g., Furlanetto 2006; McQuinn & O’Leary 2012), dark-matter




Figure 7. The marginalized 1D PDFs of the soft-band X-ray luminosity to SFR, $L_X < 2 \text{ keV} / \text{SFR}$, from the *with HERA* and *without HERA* posteriors. The highest posterior density (HPD) 68% (95%) confidence intervals are denoted under the *with HERA* posterior with dark (light) shading. The left vertical line denotes the average value of this quantity observed from HMXBs in the sample of local, star-forming galaxies from Mineo et al. (2012). The right vertical line corresponds to the theoretical result from Fragos et al. (2013) for a metal-free HMXB population, expected to be more representative of the first galaxies. HERA is the first observation to constrain the X-ray luminosities of cosmic dawn galaxies over this range, disfavoring the values seen in local, metal-enriched galaxies at $> 1\sigma$.

annihilation heating (e.g., Evoli et al. 2014; though see Lopez-Honorez et al. 2016), CMB heating (e.g., Venumadhav et al. 2018; though see Meiksin 2021), and $\text{Ly}\alpha$ heating (e.g., Chuzhoy & Shapiro 2007). However, the amount of heating required by the HERA limits at $z \approx 8$ is generally beyond what most of these alternate sources can achieve without violating constraints from other high- z observations in our model likelihood. For example, $\text{Ly}\alpha$ heating only dominates for a relatively large, slowly evolving star formation density coupled with a low X-ray efficiency. This region of parameter space is ruled out by the combination of complementary observations and HERA limits (e.g., compare the narrower range of our *with HERA* posterior in the top-right panels of Figure 6 to the range of blue curves in Figure 10 of Reis et al. 2021). Thus it is unimportant for the data-constrained *with HERA* posterior in this section, though it can be important in ruling out extreme models when not considering complementary observational data (see Reis et al. 2021 and Section 8).

5.4.2. IGM Properties: Disfavoring a Cold IGM with HERA

In Figure 8 we show the marginalized *without HERA* and *with HERA* posteriors in the space of $(\bar{x}_{\text{H I}}, \bar{T}_S)$; tan and purple regions, respectively). In gray we also show the prior distribution over this space. Comparing the tan to the gray regions, we see that previous observations disfavor a notable prior volume also in the space of IGM properties.⁴⁷ Most notably, current observations shift the posterior so that the midpoint of the EoR is occurring around $z \sim 8$ to match EoR constraints from Planck and QSO spectra.

Now introducing the H21 limits with the purple curves, we see that the HERA disfavors this region of low temperatures for $0.4 \lesssim \bar{x}_{\text{H I}} \lesssim 0.8$ at $z = 8$. These are the previously mentioned

reionization-driven models: having large fluctuations in the ionization field combined with a cold IGM. The impact of HERA is most strongly seen in the marginalized temperature PDFs in the right side panel: *with HERA* and *without HERA* exhibit qualitatively different distributions, with the HERA limits strongly disfavoring the low $\bar{T}_S$ peak seen in the posterior without HERA. This demonstrates that the HERA limits are ruling out otherwise viable models.

In Figure 9, we further investigate the physical origins of the temperature PDFs, plotting the spin temperature distributions in the bottom panel and the corresponding kinetic temperature distributions in the top panel. Both are averaged only over the neutral IGM, specifically those cells with $x_{\text{H I}} > 0.95$. We see that the kinetic temperature of the neutral IGM smoothly extends to $T_K \sim 10^4 \text{ K}$, without the bimodality seen in the spin temperature distributions for *without HERA*. This is because the spin temperature is inversely weighted between the kinetic (T_K) and radio background (T_{rad}) temperatures (see Equation (2)). As $T_K \rightarrow \infty$, the spin temperature asymptotes to $T_S \rightarrow (1 + x_\alpha) T_{\text{rad}} \sim (1 + x_\alpha) \times 24 \text{ K}$ for the standard assumption of a CMB-dominated radio background at $z = 8$. Although x_α scales with the $\text{Ly}\alpha$ background, it cannot exceed values of $x_\alpha \sim 300$ in our data-constrained models of f_{esc} without the gas in the simulation cell becoming ionized. This results in the sharp upper limit of $\bar{T}_S \lesssim 600\text{--}10^3 \text{ K}$ for the neutral IGM seen in Figure 9.⁴⁸

Comparing the purple and the tan curves in Figure 9, we reach the main conclusions of this subsection. H21 observations substantially improve our understanding of the $z = 8$ neutral IGM temperatures,⁴⁹ allowing us to place 68% (95%) high posterior density confidence intervals on the spin temperature of $27 \text{ K} < \bar{T}_S < 630 \text{ K}$ ($2.3 \text{ K} < \bar{T}_S < 640 \text{ K}$) and the kinetic temperature of $8.9 \text{ K} < \bar{T}_K < 1.3 \times 10^3 \text{ K}$ ($1.5 \text{ K} < \bar{T}_K < 3.3 \times 10^3 \text{ K}$). Other observations of the early universe and high- z galaxies are unable to constrain these temperatures on the low end.

Indeed because these temperature constraints of the neutral IGM come almost exclusively from the 21 cm signal (where they depend only on the ratio T_S/T_{rad} ; see Equation (1)), we can generalize our temperature limits for any *homogeneous* radio background even if the standard assumption of $T_{\text{rad}} = T_{\text{CMB}}$ is incorrect. In the regime of $T_{\text{rad}} > T_{\text{CMB}}$, our *with HERA* limits can thus be generalized as $1.1 \text{ (0.095)} < \bar{T}_S/T_{\text{rad}} < 26 \text{ (26)}$ and $0.37 \text{ (0.062)} < \bar{T}_K/T_{\text{rad}} < 54 \text{ (140)}$ at 68% (95%) confidence.

6. Constraints on IGM Properties Using a Reionization-driven Phenomenological Model

Here we introduce simple, phenomenological models for reionization-driven 21 cm PS and compare the resulting constraints on IGM properties to those obtained with

⁴⁷ We note that our priors over galaxy parameters do not translate into flat priors over $(\bar{x}_{\text{H I}}, \bar{T}_S)$. It is easier to theoretically and empirically motivate priors on (fundamental) galaxy properties than on (derived) IGM properties. Thus choosing flat priors directly over mean IGM properties could result in biased posteriors when using weakly constraining data (e.g., Ghara et al. 2020, 2021).

⁴⁸ Indeed the marginalized prior on $\bar{T}_S$ (shown with the gray curve in the bottom panel of Figure 9) extends out to $\bar{T}_S \sim 10^4 \text{ K}$ as the prior volume includes low values of f_{esc} that do not reionize the universe. Observations exclude these models from the posterior.

⁴⁹ We want to reemphasize that our temperatures are averaged over the *neutral* IGM component, for which the spin temperature is a defined quantity. The ionized IGM component, likely comprising tens of percent of the IGM volume at $z \sim 8$ (see the EoR history panel in the top right of Figure 6) would have $T_K \sim 10^4 \text{ K}$ (e.g., D’Aloisio et al. 2019). Thus the kinetic temperature averaged over all volume would be roughly $\langle T_K \rangle_V \approx (1 - Q_{\text{H II}}) \bar{T}_K + Q_{\text{H II}} \times 10^4 \text{ K}$, where $Q_{\text{H II}} \approx (1 - \bar{x}_{\text{H I}})$ is the volume-filling factor of the ionized IGM component, and $\bar{T}_K$ is the average IGM temperature of the neutral IGM component (plotted in the top panel of Figure 9).




Figure 8. Marginalized IGM properties corresponding to the posteriors from Figure 6. As throughout, $\bar{T}_S$ is computed by averaging only over the neutral IGM. Note that unlike in Figure 5, these are true Bayesian posteriors, as they were generated using a regular likelihood and are marginalizing over physical priors. The left (right) panel corresponds to $z = 7.9$ (10.4). **H21** limits increase the preference for hotter temperatures of the neutral IGM component. The gray curves shown in the 1D marginalized panels show our prior distribution. Our galaxy priors do not result in flat priors over $\bar{x}_{\text{HI}}$ and $\bar{T}_S$.




Figure 9. Marginalized 1D PDFs of the spin temperature (bottom) and the kinetic temperature (top) of the neutral IGM at $z = 7.9$. The colors are the same as in the previous figure. The HPD 68% (95%) confidence intervals are denoted under the *with HERA* posterior with dark (light) shading. Since by definition, the averaging is performed only over the neutral cells (with $x_{\text{HI}} > 0.95$), there are no kinetic temperatures above $\bar{T}_K \sim 10^4$ K. For models without any neutral cells at $z = 7.9$ (13% of the *with HERA* posterior), we take the mean temperatures from the last snapshot that contains neutral cells. The gray region on the left denotes values below the adiabatic-cooling limit for the IGM at mean density.

21cmFAST in the previous section. Although very simple, these phenomenological models help build physical intuition for the most important effects to consider when interpreting upper limits on the 21 cm power spectrum. We summarize the functionality of this model briefly below, and we defer a more complete description to J. Mirocha et al. (2022, in preparation).

Our principal goal is to examine a model built directly from IGM structures rather than galaxy models, so that we do not

make any *explicit* assumptions about the heating and ionizing sources during reionization. To that end, we parameterize the process not with galaxy properties but with the IGM temperature and with the ionized bubble size distribution (BSD) directly. Note that this approach does make *implicit* assumptions about the sources of reionization, e.g., through the assumed BSD parameterization; it is just nontrivial to determine what these assumptions are. However, they are certainly different from physical models like 21cmFAST, and as a result, help to determine how robust IGM constraints are to modeling assumptions.

For an idealized two-phase IGM in which the BSD is known, the two-point statistics of the ionization field can be worked out analytically following Furlanetto et al. (2004). In 21cmFAST and similar models, the excursion set approach is used to forward model the BSD, but we parameterize it more flexibly here with a log-normal distribution and allow the characteristic bubble size, R_b , and the width of the distribution, σ_b , to vary as free parameters. Note that BSDs derived from excursion set or semi-numerical models generally have broader tails to low R_b than even a log-normal (Furlanetto et al. 2004; Paranjape & Choudhury 2014; Ghara et al. 2020). But for fits to a single k and a wide prior on σ_b , as we perform here, we do not expect the detailed shape of the BSD to be important. We further assume that the “bulk IGM” outside of bubbles is fully neutral and is of uniform spin temperature, $\bar{T}_S$. The fourth and final free parameter is the volume-filling fraction of ionized gas, $Q_{\text{H II}} \equiv 1 - \bar{x}_{\text{HI}}$, which normalizes the BSD.

To model the 21 cm power spectrum within this simplified framework, one must model the ionization field *and* its correlation (or anticorrelation) with the density field. Because we abstract away assumptions about astrophysical sources completely, and instead work in terms of the BSD and mean IGM properties $Q_{\text{H II}}$ and $\bar{T}_S$, it is not immediately obvious how to do this. While it is possible to estimate the behavior of cross-terms using the halo model (Furlanetto et al. 2004) or perturbation theory (Lidz et al. 2007), here, we take a simpler approach that avoids explicit assumptions about astrophysical sources. If we assume for simplicity that the structure of the density field mirrors that of the ionization field, i.e., it is a binary field, cross-terms involving ionization and density can

be rewritten in terms of the ionization power spectrum given the typical density of ionized regions, $\langle \delta \rangle_i$. To estimate $\langle \delta \rangle_i$, we assume that reionization is “inside-out,” or in other words that the ionized volume fraction $Q_{\text{H II}}$ is made up of the densest fraction $Q_{\text{H II}}$ of the volume. Then, to complete this “volume matching” procedure, we assume the density PDF is log-normal (Coles & Jones 1991) with a variance given by the density field smoothed on the scale at which the BSD peaks. This naturally leads to a model in which the typical bubble density declines with time, so the importance of cross-terms is greatest in the early stages of reionization. Finally, as in Section 4, we assume $\mu = 0.6$ to match the spherical averaging done in 21cmMC simulations.

We perform two MCMC fits using EMCEE (Foreman-Mackey et al. 2013)—one using the inverse likelihood (Equation (10)) and one with the regular likelihood (Equation (9))—to the $k = 0.134 \text{ cMpc}^{-1}$ limit from Band 2 at $z = 7.9$ using 192 walkers for a total of $\sim 500,000$ steps. We adopt flat priors on each model parameter: $0 \leq \bar{T}_s / \text{K} \leq 10^3$, $0 \leq Q_{\text{H II}} \leq 1$, $0 \leq R_b / \text{cMpc} \leq 30$, and $0.5 \leq \sigma \leq 2$. Note that while the 21 cm signal is insensitive to $\bar{T}_s$ once $\bar{T}_s \gg T_r$, our lower limits on $\bar{T}_s$ are sensitive to the prior range. Our choice of $\bar{T}_s \leq 10^3 \text{ K}$ is motivated by the maximum allowed spin temperature in standard scenarios (see Section 5.4 and Figure 9), though we broaden the lower bound from the expected adiabatic-cooling limit of $\bar{T}_s \simeq 1.7 \text{ K}$ to zero so that more exotic scenarios may be considered.

In the top panel of Figure 10, we show constraints on the mean spin temperature and ionized fraction of the IGM obtained from this model after marginalizing over the parameters of the BSD (R_b and σ_b). We obtain 95% (68%) lower limits on the spin temperature of the $z = 7.9$ IGM of $\bar{T}_s \geq 5.3 \text{ K}$ ($\bar{T}_s \geq 25 \text{ K}$). Qualitatively, these results are in good agreement with those derived using 21cmMC (the *with HERA* posterior is shown with purple contours; see also Figure 8). As discussed in the previous section, the data-constrained 21cmMC posterior is dominated by reionization-driven fluctuations, since the density-driven models have a neutral fraction at $z \sim 8$ that is too high and are disfavored by EoR observations. It is therefore encouraging that our reionization-driven phenomenological model is broadly consistent with the *with HERA* posterior from 21cmMC. This implies that our claims of HERA’s upper limits disfavoring models in which the IGM has not been heated at $z \sim 8$ are not sensitive to the nature of the EoR fluctuations.

In the bottom panel, we show the results obtained via the inverse likelihood. Note again that we require only that the mean temperature of the IGM be positive, which is why the disfavored region in this panel extends all the way to $\bar{T}_s = 0$. This is one of the advantages of the phenomenological approach: it can constrain more exotic scenarios without invoking a particular physical model (see Section 7, where we introduce some such physically motivated models). Here again, our results are broadly consistent with the analogous ones from 21cmMC (the reionization-driven models are shown with the blue curves; note the red contours are density-driven models that are not considered by our phenomenological BSD model).

To further explore this agreement, in the bottom panel of Figure 10 we show isopower contours for several different bubble sizes, holding fixed the width of the BSD at $\sigma_b = 0.8$. The rationale here is simple: isolikelihood contours should trace isopower contours for inference based on a single k mode.




Figure 10. Constraints on the mean properties of the $z \sim 8$ IGM using phenomenological models (see Section 6) compared to the 21cmFAST results. Top: filled cyan contours are 68% (dark) and 95% (light) confidence levels obtained with the phenomenological model, while purple contours are those from 21cmMC (as in Figure 8). Bottom: as in Figure 6, the blue and red open contours in each panel correspond to reionization- and density-driven scenarios, respectively, while the filled contours show the disfavored region determined by the phenomenological model. Additional black contours in this panel trace the phenomenological model’s predictions at fixed $k = 0.134$ power, $\Delta_{21}^2 = 24.07^2 \text{ mK}^2$, corresponding to the HERA measurement $+1\sigma$, for three different bubble sizes assuming a fixed $\sigma_b = 0.6$. The cross-hatching along the bottom of each panel indicates regions with temperatures below that of a homogeneous and unheated high- z IGM.

From this plot, we see that if bubbles are generally small, $R_b \lesssim 2 \text{ cMpc}$, the phenomenological model predicts that warmer temperatures are needed to preserve the large-scale power as $x_{\text{HI}} \rightarrow 1$. However, if bubbles are generally larger, with $R_b \gtrsim 2 \text{ cMpc}$, this trend is reversed. These results suggest that physical models like 21cmFAST effectively have a low prior probability assigned to models with large bubbles at early times. Indeed, excursion set calculations suggest that typical bubbles sizes $R_b \gtrsim 2 \text{ cMpc}$ generally do not emerge until reionization is underway at the $\sim 20\%$ – 30% level (Furlanetto et al. 2004). However, because the phenomenological model can have arbitrarily large bubbles at any time, the density-driven mode is washed out when marginalizing over R_b and σ_b . Though the density-driven models are ultimately disfavored given that they do not complete reionization by $z \sim 6$ (see Figure 4), they serve as interesting test cases nonetheless (see Section 4).

7. Constraints on Dark Matter and Adiabatic Cooling Using Density-driven Models

In previous sections we have obtained limits on the IGM spin temperature $\bar{T}_s$ using different approaches. Here we study how these limits compare to predictions in the standard CDM cosmological model, as well as models of millicharged DM (mQDM). We will also briefly explore how our assumptions

about RSDs affect the limits imposed on the IGM. Throughout this section we will use our density-driven phenomenological model, Equation (12) (in all cases assuming $T_S = T_K$). While this approach has limited validity, it provides a useful test bed of our assumptions, as it allows us to obtain analytic limits under different RSD assumptions, as well as extend the temperature range studied below the adiabatic-cooling threshold to probe mQDM models, neither of which are currently included in the usual 21 cmFAST simulation-based approach of Section 5.

7.1. The Impact of RSDs on the $\bar{T}_S$ Limits

First we study how our analytic limits change under different RSD assumptions. Within our bias approach, this can be readily implemented by varying μ in Equation (12). For simulations, on the other hand, it is challenging to study the $\mu \rightarrow 1$ limit, given the geometry of the Fourier modes populating a square box. The analytic limits obtained in Section 4 ($\bar{T}_S \geq \{7.8, 1.9\}$ K at $z = \{7.9, 10.4\}$), assumed $\mu = 0.6$ to match the spherical averaging done in 21 cmMC simulations. Under the assumption that modes lie predominantly along the line of sight ($\mu \approx 1$), as actually observed by HERA (see Section 3.1), these limits strengthen to $\bar{T}_S \geq \{11, 2.6\}$ K (for $z = \{7.9, 10.4\}$) at 95% confidence, which are $\sim 50\%$ stronger, as shown in Figure 11. If we had ignored RSDs ($\mu = 0$), but kept the same assumptions otherwise, the 95% CL limits would weaken to $\bar{T}_S > \{3.1, 0.74\}$ K at $z = \{7.9, 10.4\}$, a factor of ~ 3 smaller. The difference between these three assumptions highlights the importance of properly modeling RSDs in 21 cm power-spectrum analyses. We note, however, that these results assume the density field drives the 21 cm fluctuations, in which case RSDs always increase the 21 cm power spectrum. This trend can be reversed if radiation fields are the main source of anisotropy (e.g., in reionization-driven scenarios as in Figure 3), though it is not expected to change our conclusions (see Section 5.2).

We also show the impact of varying the neutral-hydrogen fraction $\bar{x}_{\text{HI}}$ on our analytic results. Unlike the galaxy-driven models of previous sections—in which patchy reionization *enhances* the 21 cm power spectrum because of the bubble structure—here we assume uniform reionization (which could result from exotic processes; e.g., Evoli et al. 2014; Lopez-Honorez et al. 2016), in which case $\bar{x}_{\text{HI}} < 1$ *suppresses* the 21 cm power spectrum, as is clear from Equation (12). Had we assumed $\bar{x}_{\text{HI}} = 0.5$ (instead of fixing $\bar{x}_{\text{HI}} = 1$), we would arrive at the 95% confidence limits $\bar{T}_S > \{4.1, 1.2\}$ K at $z = \{7.9, 10.4\}$ (both with $\mu \approx 1$). While it is unlikely that $\bar{x}_{\text{HI}}$ deviates significantly from unity at $z = 10.4$, a global value of $\bar{x}_{\text{HI}} = 0.5$ is in line with our expectations for $z = 7.9$.

These limits have interesting implications for the thermal state of the IGM at high redshifts, as well as for the first EDGES claimed detection (Bowman et al. 2018). We compare all of the $\bar{T}_S$ limits (divided by T_{CMB}) in Figure 11 against the $\bar{T}_S/T_{\text{rad}}$ prediction for the standard CDM model, both in the absence of heating and with a fiducial X-ray heating model, akin to the ones implemented within 21 cmFAST in previous sections. The HERA Band 2 95% confidence limit is above the adiabatic-cooling prediction at $z = 7.9$, both for $\bar{x}_{\text{HI}} = 0.5$ and 1 (and in fact for any $\bar{x}_{\text{HI}} \geq 0.3$ in this bias approach). Thus, HERA requires some heating by $z = 7.9$ given our assumptions. Moreover, the HERA limits for Band 1 ($z = 10.4$), while




Figure 11. Analytically derived lower limits on $\bar{T}_S/T_{\text{rad}}$ from HERA data (95% confidence, green arrows) compared to the upper limit from EDGES (blue arrow, which also implies the two lower limits at $z = 15$ and 21 given their profile shape). The H21 limits have been obtained by assuming density-driven fluctuations (see Section 4 for details) and two different constant values of the neutral-hydrogen fraction $\bar{x}_{\text{HI}} = 1$ (dark green) and 0.5 (light green). The black line shows the standard CDM prediction assuming full coupling ($T_S = T_K$), with the dashed line corresponding to zero heating and the solid line to a toy model of X-ray heating. The HERA Band 2 data can rule out adiabatic cooling under our assumptions, requiring some X-ray heating to take place before $z = 7.9$. The red line includes a fraction $f_{\text{dm}} = 0.5\%$ of millicharged DM (mQDM), so as to explain the EDGES depth. Without any heating (dashed), this curve is ruled out by HERA Band 1, showing that there must be heating between $z = 17$ and $z = 10.4$ if EDGES is explained by mQDM. The conclusion is similar for a model with excess radio emission (with a radio fraction $f_r = 9000$; see Section 8), shown as the tan line. The empty symbols represent spherically averaged RSDs (purple, which were shown in Figure 2).

below the adiabatic limit at that z , can be used to set constraints on dark-matter induced cooling of the gas, which we now explore.

7.2. Dark Matter–Baryon Interactions

The first claimed 21 cm detection from the cosmic dawn in Bowman et al. (2018) shows a surprisingly deep absorption feature at $z \sim 17$. The depth of this absorption, if interpreted to be cosmological (see, however, e.g., Hills et al. 2018; Sims & Pober 2019; Tauscher et al. 2020), can be translated into a requirement that $\bar{T}_S/T_{\text{rad}} \leq 0.08$ at $z = 17$, a factor of two smaller than allowed by the standard cosmological model. Reducing the Wouthuysen–Field coupling in this case only exacerbates the tension, as it would bring the spin temperature closer to that of the CMB, producing shallower absorption.

A possible explanation for this anomalous depth consists of lowering the temperature of baryons in the IGM by allowing them to interact with the cosmologically abundant—and kinetically cold—DM. Elastic scattering between these two fluids would bring them closer to thermal equilibrium, cooling down the baryons and heating up the DM. These interactions could take the form of a new fundamental force (Tashiro et al. 2014; Muñoz et al. 2015b; Barkana 2018; Fialkov et al. 2018), which however would be in conflict with fifth-force constraints and stellar-cooling bounds. Alternatively, part of the DM can be electrically charged, for instance through a dark-photon portal (Holdom 1986), a scenario dubbed millicharged DM (mQDM). In this case there are no new charges for baryons, and therefore fifth-force and stellar-cooling bounds are naturally evaded (Barkana et al. 2018; Berlin et al. 2018; Kovetz et al. 2018; Muñoz & Loeb 2018; Muñoz et al. 2018;

Slatyer & Wu 2018; Liu et al. 2019). Here we briefly study how well DM–baryon interactions, in the form of mQDM,⁵⁰ can be constrained by the H21 limits.

To illustrate the effect of mQDM, we show in Figure 11 the prediction for an example model using the software developed in Muñoz & Loeb (2018). We solve the coupled differential equations for the mQDM and hydrogen-gas temperatures starting at recombination. The interactions due to millicharges produce thermalization of the (initially cold) DM and the hydrogen gas, therefore cooling the latter. For this figure we have chosen mQDM with a charge $Q_\chi = 1.3 \times 10^{-5} e$, where e is the electron charge, and mass $m_\chi = 10$ MeV, composing a fraction 0.5% of the total DM. These parameters are chosen to (barely) explain the EDGES depth and, as is clear from the figure, the cooling induced at later times lowers $\bar{T}_S$ below the HERA limit both at $z = 10.4$ and $z = 7.9$, ruling out this model in the absence of heating.

We generalize this result by performing a 2D scan of mQDM charges Q_χ and masses m_χ , assuming that mQDM particles compose a 0.5% fraction of the DM, which is at the edge of the 95% confidence interval region allowed by CMB constraints (Boddy et al. 2018), with the remainder being neutral and noninteracting CDM. We show the results in Figure 12, where we also show the region that produces enough cooling to explain the EDGES depth (defined to be $\bar{T}_S < 4$ K as in Muñoz & Loeb 2018). This region is entirely contained by the HERA Band 1 constraint ($\bar{T}_S > 2.6$ K at 95% confidence), which shows that all of the mQDM models that explain the EDGES depth also require heating before $z = 10.4$ in order to avoid conflict with HERA. Our conclusions hold for all other mQDM fractions in the relevant range $f_{\text{dm}} = 0.05\%–5\%$.

8. Astrophysical Constraints in Models with an Extra Radio Background

8.1. The 21 cm Signal in the Presence of Radio Sources

In this section we use HERA data to constrain models in which either astrophysical or exotic high-redshift radio sources contribute to the total radio background, in addition to the CMB. Such an excess radio background above the CMB level has been observed at $z = 0$, with the data consistent with a synchrotron radio background of a spectral index -2.58 and a brightness temperature ~ 603 K at the rest-frame 21 cm frequency (Fixsen et al. 2011; Seiffert et al. 2011; Dowell & Taylor 2018). The nature of this excess is still undetermined (e.g., Subrahmanyan & Cowsik 2013), and it could partially be accounted for by a population of unresolved high-redshift sources of either astrophysical or exotic origin (Ewall-Wice et al. 2018; Fraser et al. 2018; Pospelov et al. 2018; Brandenberger et al. 2019; Jana et al. 2019; Thériault et al. 2021).

An excess high- z radio background would have important implications for the 21 cm signal, because the stronger background amplifies the absorption (via the temperature term in Equation (1) including an effect on coupling coefficients in Equation (2); see complete discussion in Fialkov & Barkana 2019; Reis et al. 2020). Such models have been presented as potential explanations of the anomalously strong EDGES low band detection (Bowman et al. 2018); for example, Fialkov &




Figure 12. HERA constraints on the millicharged-DM (mQDM) parameter space. Charge Q_χ of the particles (divided by the electron charge e) vs. their mass m_χ in gigaelectronvolts. Hatched regions are ruled out by different experiments; in brown we show limits from the SLAC experiment (Prinz et al. 1998), in orange we show the most conservative Big Bang Nucleosynthesis (BBN) constraints (Jaekel & Ringwald 2010), in red we show the constraint from cooling of SN1987A from Chang et al. (2018; see, however, Bar et al. 2020 for criticisms), and in cyan we show the region disfavored if there is a dark photon A' mediating the millicharge (Muñoz & Loeb 2018; see also Vogel & Redondo 2014). Blue shows the EDGES-preferred region ($z \approx 17$; following Muñoz & Loeb 2018), and the green region is ruled out by HERA band 1 ($z = 10.4$) at 95% confidence, assuming density-driven fluctuations. An EDGES detection of millicharged DM is only compatible with HERA if heating takes place between $z \approx 17$ and $z = 10.4$. We have taken a fraction $f_{\text{dm}} = 0.5\%$ of DM to be millicharged, but our conclusions extend to all relevant fractions. The red star is the point that gives rise to the red line in Figure 11. We remind the reader that the HERA Band 2 data ($z = 7.9$) already rules out adiabatic cooling at 95% confidence, so by construction it also rules out any DM model that produces additional cooling.

Barkana (2019) found that the EDGES signal can be explained if the cosmological (high-redshift) contribution of such a background is between 0.1% and 22% of the CMB at 1.42 GHz (see also Ewall-Wice et al. 2018; Jana et al. 2019; Mirocha & Furlanetto 2019; Ewall-Wice et al. 2020; Mebane et al. 2020; Reis et al. 2020; Thériault et al. 2021). These explanations are challenging, however, they require either unconfirmed exotic sources or astrophysical sources that are far stronger than expected based on local observations (Ewall-Wice et al. 2018; Mirocha & Furlanetto 2019; Ewall-Wice et al. 2020; Mebane et al. 2020) and necessitating rapid X-ray heating to match the steep recovery in the EDGES signal (Reis et al. 2020).

More interestingly for our purposes, the presence of a radio background can also enhance fluctuations in the 21 cm signal (Fialkov & Barkana 2019; Reis et al. 2020), so that H21 can place limits on such a background (whether generated by discrete sources or more exotic processes) at $z \sim 8$ and 10. In this section we will consider both such scenarios, including the resulting limits in the context of other observations of the low-frequency radio background (Fixsen et al. 2011; Dowell & Taylor 2018) and X-ray background (Lehmer et al. 2012).

8.2. Modeling

We generate spherically averaged 21 cm PS as a function of astrophysical parameters using the semi-numerical simulation method described in Visbal et al. (2012), Fialkov et al. (2014), Fialkov & Barkana (2019), Cohen et al. (2020), Reis et al. (2020), and Reis et al. (2021). Our simulations are 384 cMpc on a side and have a resolution of 3 cMpc. Initial large-scale perturbations in density and relative velocity between dark

⁵⁰ We note that current HERA data does not allow us to place limits on DM annihilation or decay (Evoli et al. 2014; Lopez-Honorez et al. 2016; Liu & Slatyer 2018), as we only have lower bounds on the gas kinetic temperature. A 21 cm detection is required for those analyses.

matter and gas (Tseliakevich & Hirata 2010) are linearly evolved from the Dark Ages ($z \sim 60$) to $z = 5$. Using the modified Press-Schechter mass function, which takes into account the effect of large-scale overdensity and velocity fields (Barkana & Loeb 2004; Fialkov et al. 2012), we calculate the halo abundance in each voxel of the simulation. Each halo is then populated by stars, and emissivities in different bands are calculated (see, e.g., Cohen et al. 2020, for details). RSDs are computed by multiplying the real space isotropic 21 cm signal by $(dv_r/dr)^{-1}$, which is the radial component of the velocity gradient created by structure formation (Fialkov et al. 2020). Using coeval simulation cubes, we calculate the spherically averaged power spectrum at every redshift.

The key radiation backgrounds are all driven by the cosmic star formation rate, which, in the simulations, depends on two parameters. First, we choose a minimum circular velocity of star-forming halos V_c , which determines the halo population that can form stars. We then choose a star formation efficiency f_* , which measures the amount of collapsed gas that turns into stars for halos above the atomic cooling limit, imposing an extra suppression in smaller halos (e.g., Fialkov et al. 2013; Cohen et al. 2020). The code includes the suppression of star formation by Lyman–Werner feedback (Fialkov et al. 2013), relative velocities between dark matter and gas (Fialkov et al. 2012), and photoheating feedback (Cohen et al. 2016). To calculate the Ly α background, we assume Population II star formation following Barkana & Loeb (2005b). For completeness we note that here we include multiple scattering of Ly α photons and Poisson fluctuations in the number of first star-forming halos (however, these effects are not significant at the redshift range observed by HERA; Reis et al. 2021).

X-ray heating of the IGM is powered by a population of X-ray binaries with the ratio of bolometric luminosity to SFR of

$$\frac{L_{X, 0.2-95 \text{ keV}}}{\text{SFR}} = 3 \times 10^{40} f_X \frac{\text{erg s}^{-1}}{M_\odot \text{yr}^{-1}} \quad (17)$$

calculated between 0.2 and 95 keV assuming a hard X-ray SED of X-ray binaries (Fragos et al. 2013). The free parameter f_X normalizes the X-ray efficiency relative to a population of present-day binaries (but including an order-of-magnitude increase in this ratio at the low metallicity expected for high-redshift galaxies; Fragos et al. 2013). The assumed SED is relatively hard, peaking at ~ 1 keV (Fragos et al. 2013; Fialkov et al. 2014). Note that assuming a different SED could affect our final results (Fialkov et al. 2014; Pacucci et al. 2014; Das et al. 2017; Reis et al. 2020). The unresolved X-ray background observed by Chandra imposes an upper limit on this contribution (Fialkov et al. 2017), as we will explore later, but we allow a broad range of f_X between 10^{-4} and 10^3 in the estimation framework.

This model also includes heating by scattering of Ly α photons (Chen & Miralda-Escudé 2004; Chuzhoy & Shapiro 2007; Ciardi et al. 2010; Mittal & Kulkarni 2021; Reis et al. 2021) and the CMB (Hirata & Sigurdson 2007; Venumadhav et al. 2018; Fialkov & Barkana 2019; Reis et al. 2021, though see Meiksin 2021). With the onset of the first stellar population, the extra heating processes raise the IGM temperature above the adiabatic limit even in the absence of X-ray heating, reducing the 21 cm background at the relevant redshifts in

some scenarios by a factor of a few (see Reis et al. 2021, for more details).

The process of reionization is implemented using the excursion set formalism (Furlanetto et al. 2004) and is described by two parameters: the ionizing efficiency of sources ζ , which is normalized via the total CMB optical depth τ , and the horizon of ionizing photons, R_{mfp} . Although the latter parameter does affect the intensity of the 21 cm fluctuations at the end of the EoR, we fix it at 40 Mpc here as it plays a secondary role in our constraints.

Finally, we explore two types of radio backgrounds (beyond the CMB):

1. A fluctuating, time-variable radio background generated by galaxies, parameterized by f_r . We assume that the galaxy radio luminosity per unit frequency in units of W Hz^{-1} is proportional to the SFR (following Mirocha & Furlanetto 2019)

$$\frac{L_{r,\nu}}{\text{SFR}} = 10^{22} \left(\frac{\nu}{150 \text{ MHz}} \right)^{-\alpha_r} f_r \frac{\text{W Hz}^{-1}}{M_\odot \text{yr}^{-1}} \quad (18)$$

where α_r is the spectral index in the radio band, which we set to the typical value of 0.7 (Mirocha & Furlanetto 2019), which is compatible with observations (Hardcastle et al. 2016; Gürkan et al. 2018). We calculate T_{rad} at redshift z by summing up over the past light-cone contribution of all of the radio galaxies (see Reis et al. 2020, for more details).

2. A smooth synchrotron background that decays with time, for which we replace T_{CMB} by

$$T_{\text{rad}} = T_{\text{CMB}}(1+z) \left[1 + A_r \left(\frac{\nu_{\text{obs}}}{78 \text{ MHz}} \right)^\beta \right] \quad (19)$$

where ν_{obs} is the observed frequency, A_r is defined relative to the CMB temperature, and $\beta = -2.6$ is the spectral index in agreement with the ARCADE2 (Fixsen et al. 2011) and LWA1 (Dowell & Taylor 2018) observations. Here we treat this background as phenomenological, but it could have been produced by exotic radio sources, e.g., radiative decay of relic neutrinos into sterile neutrinos (Chianese et al. 2019), light dark matter decays (Fraser et al. 2018; Pospelov et al. 2018), and superconducting cosmic strings (Brandenberger et al. 2019).

In this work, we allow a broad range of f_r and A_r parameters. However, as we discuss later (and as was shown by Fialkov & Barkana 2019; Reis et al. 2020), models with strong radio backgrounds, e.g., $f_r \times f_* > 10^3 - 10^4$ for the radio from galaxies, are constrained by ARCADE2/LWA1 data.

To summarize, the models considered here include an extra radio background in addition to the CMB either produced by radio galaxies or emitted by exotic sources. Our models build on the following parameters: V_c varied between 4.2 and 100 km s^{-1} , f_* between 0.001 and 0.5, f_X in the range between 10^{-4} and 10^3 , τ between 0.035 and 0.088, f_r from 1 to 10^5 , and A_r between 10^{-2} and 10^5 . Owing to the large dynamic ranges, we assume uniform priors on the parameters $\log_{10} f_*$, $\log_{10} V_c$, $\log_{10} f_X$, τ , and $\log_{10} f_r$ or $\log_{10} A_r$. For completeness we also include our constraints on the standard models (i.e., with no extra radio background above the CMB).

Although similar in spirit to the 21cmFAST simulation suite described in Section 5, there are differences between the two

sets of simulations. We refer the reader to the relevant papers for details on the physics and implementation differences. Broadly, the simulations described in this section include some additional heating processes, such as Ly α heating (which can affect the IGM temperature when it is very cold in some models) and (of course) radio emission, but they have a more prescriptive source model with fewer free parameters (and they are not constrained by ancillary observations such as the galaxy luminosity functions). A detailed code comparison is beyond the scope of this paper; instead, we focus on how these distinct codes can address the issues to which they are each best suited.

8.3. Parameter Estimation

We explore the parameter space of models compatible with the HERA upper limits based on the likelihood $\mathcal{L}_m$ defined in Equation (9). We also decimate the PS as described in Section 3.2.1, using the even wavenumbers ($k = 0.086, 0.17, \dots$ cMpc $^{-1}$) of Band 1 and odd wavenumbers ($k = 0.13, 0.21, \dots$ cMpc $^{-1}$) of Band 2. Note that we only compute the power spectrum up to $k = 1.1$ cMpc $^{-1}$, limited by the simulation resolution. Larger wavenumbers (smaller scales), however, are irrelevant as the HERA limits rise much more steeply at small scales than realistic models, so those scales do not contribute toward the constraints.

Because individual simulations take a few hours to complete, we instead use an emulator to interpolate the PS from $\sim 10^4$ existing simulation runs (for each of the two types of radio background that we investigate here). We implement the emulator using neural networks: taking the astrophysical parameters as an input, a network predicts the logarithm of the PS for the HERA redshifts ($z = 7.93$ and 10.37) and wavenumbers (from $k = 0.086$ to 1.1 cMpc $^{-1}$). The architecture of the emulator includes a multilayer perceptron with four hidden layers of 100 nodes each, implemented using scikit-learn (Pedregosa et al. 2011). The PS are predicted with a relative error of 20%; we take this uncertainty into account by adding it in quadrature to the observational error σ_i in Equation (9). Although this is an approximation, the associated error is negligible in the context of current analysis. A detailed discussion of the emulator and its accuracy can be found in Appendix B.

We explore the parameter space using the MCMC Ensemble sampler *emcee* (Foreman-Mackey et al. 2013), and we visualize and analyze the results using *anesthetic* (Handley 2019) and *GetDist* (Lewis 2019).

8.4. Results: A Radio Background Generated by Galaxies

As was alluded to above, models with an additional radio background can easily, unlike most standard scenarios, exceed the HERA upper limits. To illustrate this point, we show a random subset of the simulated PSs for the case of a radio background from galaxies in Figure 13. The PSs are shown at $z = 7.93$ and colored with respect to their compatibility with HERA constraints, indicating the difference in log-likelihood $\Delta \log \mathcal{L}_m$ compared to the best fit (which is $\Delta_{21}^2 \approx 0$ mK 2). For comparison, we also plot the current HERA limits marked by data points with error bars. Clearly, a substantial fraction of the models (shades of orange) are excluded by the current HERA limits with high significance. In comparison, corresponding standard models (no additional radio background) have much lower amplitudes. We show the envelope of these models (i.e., the maximal possible amplitude of the ensemble of standard




Figure 13. Power spectra of 1000 randomly selected models with an extra radio background created by galaxies, at $z = 7.93$ (HERA Band 2). A number of these models can be ruled out by the HERA data as shown by the color of the lines indicating the likelihood of each model. We use the decimated data points as described in Section 3.2.1, taking only the “even” points of Band 1 and “odd” points of Band 2. We show the Band 2 points with black circles for the points we take into account and white circles for the unused data. The error bars show 1σ errors, and the crosses indicate the 2σ upper limits. For comparison, the thick dashed line shows the maximal possible amplitude of the ensemble of standard models at each wavenumber k (i.e., the envelope).

models at each k) with the thick dashed line in Figure 13 and discuss astrophysical implications of HERA for these cases in Section 8.6.

Using the HERA likelihood alone ($\mathcal{L}_m$), we show the marginalized constraints on the parameters in Figure 14. The diagonal panels show the 1D marginalized posterior PDFs while the others show the 2D marginalized PDFs with the dotted and dashed lines indicating the 68% and 95% confidence contours, respectively (containing 68% and 95% of the 2D posterior probability, respectively). The 2D marginalized posteriors involving f_* , V_c , and τ are relatively flat with the ratio between minimum and maximum posterior probability⁵¹ being between 0.3 and 0.6. This results in confidence contours that could be easily affected by fluctuations due to the random sampling and are strongly dependent on the prior. On the contrary, we find the 2D posterior in the f_X – f_r plane to show a strong contrast between minimum and maximum regions (with minimum/maximum posterior ratio of 0.02, i.e., dropping by more than three e-folds). There is a vanishing probability for models with both a strong radio background (large f_r) and weak X-ray heating (low f_X). The large contrast in the probability across this subspace indicates that the constraints on the combination of f_X and f_r are expected to be robust, i.e., even for different priors (which would cause a small shift in the contour lines) most of the high- f_r low- f_X region will still be excluded.

Marginalizing over the rest of the model parameters, we constrain f_X to be greater than 0.25 and f_r less than 397 at the 68% confidence level individually, which maps to the excluded (at 68%) $L_{X,0.2-95 \text{ keV}}/\text{SFR} < 7.6 \times 10^{39} \text{ erg s}^{-1} M_\odot^{-1} \text{ yr}$ and $L_{r,\nu}/\text{SFR} > 4 \times 10^{24} \text{ W Hz}^{-1} M_\odot^{-1} \text{ yr}$ (calculated at reference frequency $\nu = 150 \text{ MHz}$). The 2D region where both $f_X < 0.25$

⁵¹ These are measured from the bins in Figure 14 using the minimum/maximum bin sample count ($\propto$ posterior value). A value close to one shows that the PDF is largely flat and therefore does not provide a strong constraint.




Figure 14. HERA constraints on models with extra radio background from galaxies. The triangle plot shows the marginalized 1D and 2D posteriors for all model parameters, with 68% and 95% confidence contours indicated by black dotted and dashed lines, respectively. The backgrounds show the 2D histograms of posterior values to illustrate how strongly a region is excluded. The white background indicates a posterior close to zero while dark blue corresponds to the maximum posterior value of a panel. The gray areas in the diagonal panels, and the solid orange lines show the individual 68% confidence limits on f_X and f_r (numbers quoted in text). Combined, these approximately correspond to our main result, the 95% confidence excluded region in the f_X - f_r space, as shown in the center bottom panel. The inset in the top-right corner shows the corresponding constraints on the derived parameters, $\bar{T}_{\text{rad}}$ and $\bar{T}_K$, at $z = 8$ for both prior and posterior. The prior is uniform in the model parameters ($\log_{10} f_*$, $\log_{10} V_c$, $\log_{10} f_X$, τ , $\log_{10} f_r$) but therefore nonuniform in the derived parameters, $\bar{T}_{\text{rad}}$ and $\bar{T}_K$. The blue and orange dashed lines show upper bounds (at 95% confidence) on $\log_{10}(\bar{T}_{\text{rad}}/\bar{T}_K)$ from prior and posterior, respectively.

and $f_r > 397$ (marked by orange solid lines on the corresponding 2D PDF in Figure 14) approximately corresponds to the region excluded at the 95% confidence level (exact contours shown as dashed black line).

To relate the constraints on model parameters to the physical state of the IGM at the observed redshifts, we calculate limits on the radiation temperature $\bar{T}_{\text{rad}}$ and gas temperature $\bar{T}_K$ at redshift $z = 8$ (note that these are derived, i.e., not sampled, parameters in the context of this model). The temperatures are

averaged over all of the regions that are not fully ionized (i.e., over the volume outside of H II regions).⁵² These temperatures

⁵² For completeness, here the model suite includes scenarios that have extremely strong heating corresponding to $f_* \times f_X \gg 1$, which are ruled out by Chandra observations. In such scenarios, gas outside of the H II regions could be partially ionized up to $\sim 50\%$ by X-ray photons (e.g., Fialkov et al. 2017), which leads to a suppressed 21 cm signal. The maximal possible temperature of such partially ionized gas is set to be the ionizing temperature of hydrogen, 1.57×10^5 K, equivalent to the energy of 13.6 eV.




Figure 15. HERA constraints compared to the radio background constraints from LWA1 and the X-ray background constraints from Chandra. The plot color-codes the models by their HERA likelihood $\mathcal{L}_m$ (orange corresponds to the models excluded at $\Delta \log \mathcal{L}_m \geq 20$ relative to the best fit, and purple shows the allowed models), and models excluded by LWA and Chandra are shown as empty circles and crosses, respectively. The latter constraints are computed from the level of X-ray and radio background produced by the population of sources at $z > 8$ and approximately correspond to $f_X \cdot f_* \lesssim 1$ and $f_r \cdot f_* \lesssim 10^3$, respectively. We see that HERA can exclude a substantial fraction of the otherwise unconstrained parameter space.

are stored for every simulation and emulated similarly to the PS (see Appendix B for details). We obtain 2D posterior PDFs for these temperatures, marginalized over all other parameters, and we show the results in the top-right panel of Figure 14 (blue contours). Because the sampling is uniform in terms of the simulation parameters (or their log for f_* , V_c , f_X , τ , f_r), the priors on the derived parameters (temperatures) are not uniform and are shown as orange contours in the background.

The H21 limits significantly affect the temperature distribution, excluding high values of $\bar{T}_{\text{rad}}$ when $\bar{T}_K \lesssim 1000$ K. This is expected, as, typically, a combination of high $\bar{T}_{\text{rad}}$ and low $\bar{T}_K$ leads to a strong 21 cm signal exceeding the observed HERA limits. Exceptions are scenarios in which gas is either significantly ionized already by $z = 8$ or in which the radio background has a non-negligible effect on the spin temperature $\bar{T}_S$, driving it toward the background radiation temperature $\bar{T}_{\text{rad}}$, and, as a result, suppressing the total power spectrum. In such models, the 21 cm power spectrum is below the HERA limits even when gas is predominantly neutral and is colder than the radio background. These models populate the lower-left corner of the temperature plot and are not ruled out by HERA even when $\bar{T}_K \ll \bar{T}_{\text{rad}}$. Overall we derive a constraint on the ratio $\log_{10}(\bar{T}_{\text{rad}}/\bar{T}_K) < 1.1$ at $z = 8$ (at 95% confidence), compared to the prior constraint of $\log_{10}(\bar{T}_{\text{rad}}/\bar{T}_K) < 2.6$ (95% confidence), noting that these limits can be exceeded at low $\bar{T}_K$ due to the effect described above. The bounds are shown as dashed lines in the temperature plot in Figure 14.

Reis et al. (2020) showed that models with a strong radio background, corresponding to large values of the product $f_* \times f_r$, violate the ARCADE2/LWA1 measurement with only a weak dependence on the values of other astrophysical parameters. Additionally, models with large values of $f_* \times f_X$ are ruled out by the Chandra X-ray background (Fialkov et al. 2017); though in detail, the constraints are sensitive to the assumed SED at high energies. To compare HERA constraints

to the ARCADE2/LWA1 and Chandra limits, we calculate the contributions to the observed present-day radio and X-ray backgrounds created by a population of sources at $z > 8$ and flag as rejected models in which these high-redshift contributions exceed the observations. In Figure 15 we show that the H21 limits are complementary to both the ARCADE2/LWA1 and Chandra X-ray background limits in that they rule out a different section of the parameter space, which is characterized by a high radio background and low X-ray heating (orange filled circles).

8.5. Synchrotron Radio Background

Next, we calculate limits on the model parameters for the case of a smooth phenomenological synchrotron radio background (given by Equation (19) with $\nu_{\text{obs}} = 1420/(1+z)$ MHz) that is stronger at higher redshifts. Here we build on the formalism developed by Fialkov & Barkana (2019) of the excess radio background, taking into account its effect on the background intensity, spin temperature and coupling coefficients, as well as heating by radio/CMB photons. We expand over the model used in Fialkov & Barkana (2019) by adding Ly α heating, RSDs, multiple scatterings of Ly α photons, and Poisson fluctuations as described in Reis et al. (2021). We use the HERA likelihood alone and obtain qualitatively similar results (Figure 16) to the ones found in models where the radio background is generated by galaxies. Differences arise because the synchrotron background decays with time, while the radio background from galaxies grows. In addition, the former background is smooth, while the latter traces galaxies and, thus, includes spatial fluctuations.

We find that models with a high excess radio background of $A_r > 31$ (calculated relative to the CMB at the reference frequency of 78 MHz, corresponding to 1.6% of the CMB at 1.42 GHz) are excluded at 68% confidence; and so are models with a low X-ray efficiency of $f_X < 0.20$ (corresponding to the X-ray luminosity per SFR of $L_{X, 0.2-95 \text{ keV}}/\text{SFR} < 5.9 \times 10^{39} \text{ erg s}^{-1} M_{\odot}^{-1} \text{ yr}$). Models fulfilling both criteria are clearly excluded by the data as seen in the f_X - A_r panel in Figure 16; they mostly lie beyond the 95% 2D confidence contour (dashed line).

We also constrain the ratio $\log_{10}(\bar{T}_{\text{rad}}/\bar{T}_K) < 1.7$ at $z = 8$ at 95% confidence, compared to the prior constraint of $\log_{10}(\bar{T}_{\text{rad}}/\bar{T}_K) < 3.5$, again noting that these limits are less strict due to the shape of the posterior distribution as shown in the temperature plot in Figure 16.

A similar analysis with these models was done by Mondal et al. (2020) using the upper limit at $z = 9.1$ derived from 141 hr of LOFAR HB data (Mertens et al. 2020). These data rule out (at 68% confidence) a strong contribution above 0.83% of the CMB at 1.42 GHz (corresponding to $A_r > 15.9$) of the high-redshift universe to the ARCADE2 and LWA1 measurements as well as scenarios with weak X-ray heating $f_X < 0.01$. Because the assumed synchrotron radio background decays with time, tighter constraints at higher redshifts are expected. However, the theoretical framework used here has been upgraded since the analysis of the LOFAR data (specifically, RSDs and extra heating by Ly α and radio were added). This difference impedes a quantitative comparison between the two works. Moreover, compared to Mondal et al. (2020), here we used a broader prior range on A_r , which could partially explain our higher limit.




Figure 16. HERA constraints on models with a smooth synchrotron extra radio background, with analogous notation to Figure 14. Dotted and dashed lines indicate 68% and 95% confidence contours, and the background shows the histogram of the 2D posterior. The gray areas and orange solid lines show the 1D 68% percent limits on f_X and A_r individually; the combination of the latter approximately corresponds to the 95% combined exclusion region, our main result. The upper-right inset shows the prior (orange) and posterior (blue) distributions of the derived model parameters $\bar{T}_K$ and $\bar{T}_{\text{rad}}$ at $z = 8$, with the dashed lines indicating the respective 95% confidence limits on $\log_{10}(\bar{T}_{\text{rad}}/\bar{T}_K)$.

8.6. Implications for Standard Models

For completeness we note that only a small fraction of the prior volume used for this simulation suite is ruled out if the only source of the background radio photons is the CMB. We find that every excluded model (at the level of $\Delta \log \mathcal{L}_m > 1$) belongs to the region of parameter space with $f_X \lesssim 1$ (corresponding to X-ray luminosity per unit SFR of $3 \times 10^{40} \text{ erg s}^{-1} M_\odot^{-1} \text{ yr}$), $V_c \gtrsim 50 \text{ km s}^{-1}$, $\tau \lesssim 0.064$, and $f_* \lesssim 0.07$. Note that varying R_{mfp} , which is fixed here to 40 Mpc, could slightly modify the standard model's constraints (see Mondal et al. 2020). The excluded models are characterized by a low

X-ray heating efficiency and star formation rate, resulting in a partially reionized IGM (close to 50%) at $z = 8$ and creating the strongest reionization-driven power spectra at the frequencies observed by HERA. We do not rule out density-driven models with HERA observations, likely due to Ly α heating, which lowers the power-spectrum amplitude in those models.

8.7. EDGES-motivated Scenarios

As mentioned in Section 8.1, models with a strong extra radio background were proposed as a potential explanation of the anomalously deep absorption trough detected by EDGES.




Figure 17. Models compatible with a cosmological interpretation of the EDGES detection in the context of HERA limits. Top: global signals of randomly selected models (gray) and those that are compatible with EDGES (purple), selected by their absorption feature close to 78 MHz and the narrowness of the profile. The models here include an extra radio background from radio galaxies. Bottom: corresponding power spectra at $k = 0.13 \text{ cMpc}^{-1}$. HERA measurements at $k = 0.13$ and 0.17 cMpc^{-1} with 1σ error bars are shown for comparison. These wavenumbers are the most constrained by HERA but still cannot constrain the EDGES-compatible models, as their power spectra peak at a much higher redshift.

In such models, an accompanying enhancement of the PS (compared to the corresponding cases with the CMB as the background radiation) is expected (Fialkov & Barkana 2019; Reis et al. 2020; see also Figure 17), and, thus, the EDGES detection could be verified by interferometers. However, in the case of the EDGES-compatible scenarios, the expected enhancement of the power spectrum mostly falls around the central frequency of the detected global signal (i.e., 78 MHz, corresponding to $z \sim 17$), and, thus, far outside the HERA band. This is because, in order to create a deep and narrow signal at $z \sim 17$ to explain the EDGES detection, a combination of a strong radio background and strong X-ray heating is required (Fialkov & Barkana 2019; Reis et al. 2020), which is a different corner of the astrophysical parameter space than that excluded by HERA. Models with strong X-ray heating tend to be suppressed at low redshifts, and, therefore, the EDGES-compatible models are also compatible with the limits set by HERA.

To illustrate this behavior, we check which of our simulated global signals are broadly compatible with EDGES⁵³ and mark them as purple lines in Figure 17. The rest of the models are

painted in gray. In the bottom panel, we plot corresponding PSs. As we can see from the figure, indeed, the PSs of the EDGES-compatible models (purple) peak at much higher redshifts than those probed by HERA and have low power in the HERA band. None of the EDGES-compatible models (neither with the radio background from galaxies as shown here, nor with the phenomenological synchrotron radio background) are excluded by the HERA data.

9. Discussion

The HERA Collaboration (2021) has recently presented our first power-spectrum limits at $z = 7.9$ and $z = 10.4$ in H21. In this paper, we have used a suite of theoretical models to quantify the implications of these upper limits for the landscape of IGM and galaxy formation models of the cosmic dawn. The most fundamental of our conclusions, about the IGM temperature at $z \sim 8$, are illustrated in Figure 18. We have shown that:

1. The HERA limits at $z = 7.9$ strongly disfavor otherwise viable models with weak or no IGM heating by astrophysical sources. In the fiducial scenario where the CMB dominates the high- z radio background, the HERA limits combined with galaxy and EoR observations constrain the spin temperature of the neutral IGM to $27 \text{ (2.3) K} < \bar{T}_S < 630 \text{ (640) K}$ at the 68% (95%) HPD confidence interval. This corresponds to kinetic temperatures of the neutral gas of $8.9 \text{ K} < \bar{T}_K < 1.3 \times 10^3 \text{ K}$ at 68% confidence. We have demonstrated this with flexible galaxy evolution models (using 21cmFAST; Figure 9), a phenomenological model that directly parameterizes the IGM properties (Figure 10), and a simple bias model (Figure 11).
2. Using 21cmFAST and 21cmMC in which astrophysical heating pre-reionization is sourced by galactic X-ray emission, we found that the combination of other high- z observations and the recent limits from HERA allows us to constrain the X-ray luminosity per unit star formation rate of cosmic dawn galaxies, $L_{X, < 2 \text{ keV}}/\text{SFR}$. Our resulting 68% confidence interval of $L_{X, < 2 \text{ keV}}/\text{SFR} = \{10^{40.2}, 10^{41.9}\} \text{ erg s}^{-1} M_{\odot}^{-1} \text{ yr}$ supports theoretical predictions (e.g., Fragos et al. 2013) and the observed evolution with metallicity and redshift (Basu-Zych et al. 2013; Douna et al. 2015; Brorby et al. 2016; Lehmer et al. 2016) that this quantity increases toward high redshifts.
3. Using a different set of simulations (Section 8), we showed that HERA's limit (on its own) places a constraint on an early radio background for models in which galaxies have substantial radio emission but weak X-ray heating. We find that the models with $L_{r,\nu}/\text{SFR} > 4 \times 10^{24} \text{ W Hz}^{-1} M_{\odot}^{-1} \text{ yr}$ (calculated at reference frequency $\nu = 150 \text{ MHz}$) and $L_{X,0.2-95 \text{ keV}}/\text{SFR} < 7.6 \times 10^{39} \text{ erg s}^{-1} M_{\odot}^{-1} \text{ yr}$ are excluded at 95% confidence. The ratio of radio to gas temperatures must satisfy $\log_{10}(\bar{T}_{\text{rad}}/\bar{T}_K) < 1.1$ at $z = 8$ to be consistent with HERA at 95% confidence.

Considering a phenomenological model where the radio background (in addition to the CMB) has a synchrotron spectrum, we find that an extra radio background of 1.6% of the CMB at 1.42 GHz or stronger and weak heating by X-ray binaries with luminosity $L_{X,0.2-95 \text{ keV}}/\text{SFR} < 5.9 \times 10^{39} \text{ erg s}^{-1} M_{\odot}^{-1} \text{ yr}$, are robustly excluded by the data (at 95 % confidence). For

⁵³ Following Fialkov et al. (2018), we adopt the criterion that EDGES-compatible global signals should have a drop of 300–1000 mK at frequencies between 68 and 88 MHz, compared to the adjacent regions (60–68 and 88–96 MHz, indicated by the dotted lines).

this model, we find that the ratio $\log_{10}(\bar{T}_{\text{rad}}/\bar{T}_k) < 1.7$ at $z = 8$ is required for consistency with HERA at 95% confidence.

The HERA constraints on X-ray and radio efficiencies are complementary to the constraints from the unresolved X-ray background (measured by Chandra) and the low-frequency radio background (detected by ARCADE2/LWA), ruling out otherwise unconstrained parameter space.

4. There has been considerable recent interest in exotic cosmological models that might explain the EDGES measurement at $z \sim 18$ (Bowman et al. 2018), either through cooling the IGM below the adiabatic limit or imposing an additional radio background. The H21 limits cannot confirm or disprove the EDGES signal, simply because it observes at a lower redshifts. However, the limits do independently rule out a very cold IGM at $z = 10.4$ (as well as an IGM at or below the adiabatic limit at $z = 7.9$), so that heating must occur before $z \sim 10$ if millicharged dark matter is behind the anomalous EDGES depth. The EDGES signal itself, if taken at face value, already requires such heating at $z \gtrsim 15$, so cosmological explanations of that signal are compatible with the H21 limits. Future measurements from HERA at lower frequencies will be important for further testing such scenarios.

These conclusions, validated by multiple independent approaches, demonstrate that the H21 upper limits provide important astrophysical constraints on early galaxies. This is the first time IGM heating has been required by observations, independent of any assumptions about galaxy formation during this era. Nevertheless, the upper limits are sufficient to show how state-of-the-art galaxy models can offer new insights, as demonstrated by the difference in temperature constraints in density-driven models.

Moreover, our results emphasize how 21 cm measurements are highly complementary to other probes of the reionization era. No other existing probe can constrain the X-ray or radio emissivity of early galaxy populations as effectively. The requirement of at least minimal heating (or substantially more if radio emission is strong or if dark-matter interactions are at play) translates to a lower limit on radiative heating by star-forming galaxies (conventionally interpreted as X-ray heating, though at very low temperatures, Ly α heating also plays a role). While the current upper limits do not strongly constrain the expected parameter space of early galaxies, they show the promise of 21 cm interferometers. We also note that the combination of HERA data with complementary data is most powerful when they probe similar redshifts, so that assumptions about redshift evolution can be avoided.

Of course, there are several caveats to this interpretation and ways in which our methods can be improved: (i) they are driven largely by a very small number of measurements (single band powers at $z = 7.9$ and $z = 10.4$) so are more subject to systematic concerns; (ii) our constraints do not fully capture the correlations between band powers; and (iii) our simulation-based analyses also do not treat the angular dependence of the signal properly, as current observations are sensitive only to modes nearly along the line of sight, where the power spectrum is most strongly affected by RSDs (see Figure 11). Our assumptions are strictly conservative for density-driven scenarios, but the effect of RSDs is more complicated during




Figure 18. Summary of lower limits on the IGM spin temperature at $z = 7.9$ enabled by HERA (H21) limits. From top to bottom, we show the 95% lower limits derived via density-driven models (Section 4), phenomenological reionization-driven models (Section 6), 21cmMC (Section 5), and radio background models (Section 8). The red region shows temperatures below the adiabatic-cooling limit. Black points show our final constraints, while gray points indicate priors, or constraints derived from simplifying assumptions or without HERA; e.g., the top panel adopts two contrasting scenarios: a fully neutral case and a case with *uniform* 50% ionization. The bottom three panels all indicate the prior range, as well as the limits obtained via independent measurements alone in the 21cmMC case. The bottom panel includes both homogeneous (Section 8.5) and inhomogeneous (Section 8.4) radio backgrounds. In the latter case, the correspondence between top and bottom x-axes assumes a uniform $T_{\text{rad}} = T_{\gamma}(z = 7.9)$, but we caution that T_s/T_{rad} constraints cannot simply be converted to $\bar{T}_s$ constraints.

reionization. We expect to improve the match between simulations and the data for future observing campaigns.

The H21 constraints are based on a single observing season with a small number of antennae and the first iteration of HERA’s infrastructure. HERA has continued to expand and improve, and later campaigns will offer much higher sensitivity. This analysis provides a foundation for the interpretation of these forthcoming data sets.

This material is based upon work supported by the National Science Foundation under grant Nos. 1636646 and 1836019 and institutional support from the HERA collaboration partners. This research is funded by the Gordon and Betty Moore Foundation through grant GBMF5215 to the

Massachusetts Institute of Technology. HERA is hosted by the South African Radio Astronomy Observatory, which is a facility of the National Research Foundation, an agency of the Department of Science and Innovation. A.M. acknowledges funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No. 638809–AIDA). The results presented here reflect the authors’ views; the ERC is not responsible for their use.

J.S.D. gratefully acknowledges the support of the NSF AAMP award #1701536. A.L. acknowledges support from the New Frontiers in Research Fund Exploration grant program, the Canadian Institute for Advanced Research (CIFAR) Azrieli Global Scholars program, a Natural Sciences and Engineering Research Council of Canada (NSERC) Discovery Grant and a Discovery Launch Supplement, the Sloan Research Fellowship, and the William Dawson Scholarship at McGill. N.K. acknowledges support from the MIT Pappalardo fellowship. G.B. acknowledges funding from the INAF PRIN-SKA 2017 project 1.05.01.88.04 (FORECAST), support from the Ministero degli Affari Esteri della Cooperazione Internazionale–Direzioe Generale per la Promozione del Sistema Paese Progetto di Grande Rilevanza ZA18GR02 and the National Research Foundation of South Africa (grant No. 113121) as part of the ISARP RADIOSKY2020 Joint Research Scheme, from the Royal Society and the Newton Fund under grant NA150184 and from the National Research Foundation of South Africa (grant No. 103424). A.F. is supported by the Royal Society University Research Fellowship. J.B.M. is supported by a Clay Fellowship at the Smithsonian Astrophysical Observatory. Y.Q. would like to thank Cathryn M. Trott for providing MWA data and acknowledge the support from the High Performance Computing centers of the Scuola Normale Superiore (Italy), the Council for Scientific and Industrial Research (CHPC South Africa), and the OzSTAR national facility (Australia) for computational resources. Parts of this research were supported by the Australian Research Council Centre of Excellence for All Sky Astrophysics in 3 Dimensions (ASTRO 3D), through project No. CE170100013. M.G.S. acknowledges support from the South African Radio Astronomy Observatory (SARAO) and the South Africa National Research Foundation (grant No. 84156).

Contributions of the Authors

S.R.F. leads the HERA Theory team and organized this effort. Y.Q. led the 21cmFAST and 21cmMC simulation and analyses, with particular support from B.G. and A.M. (Section 5). J.M. led the analysis using phenomenological models (Section 6), with particular support from J.B.M. and S.R.F. J.B.M. led the bias analysis (Section 4) and the dark-matter interpretation (Section 7). S.H. led the analysis described in Section 8 using models created by A.F., and developed the emulators based on the architecture by S.S. Additionally, S.G.M. and N.K. developed the derivation of the likelihood (Section 3.2). R.B. is in the Builder’s list of the 21 cm code used in Section 8. I.R. led the development of the radio background from the galaxies model used in Section 8. The HERA Builder’s list is included as authors because of the dependence of this work on the combined efforts of collaborators on the various software repositories as well as the necessity of using HERA data to constrain the models used in this paper.

Appendix A Fluctuations Tracing the Velocity Field

In Section 4, we assumed that the 21 cm signal traced the matter power spectrum for simplicity. However, in some regimes, the 21 cm power spectrum may more closely trace the velocity-induced acoustic oscillations (VAOs) from the dark matter–baryon relative velocities v_{cb} . This regime can be triggered by a strong dependence of galaxy formation on these streaming velocities at small scales (Dalal et al. 2010; Tseliakhovich & Hirata 2010; Naoz et al. 2012), which, however, are not expected to play a key role at low redshifts (Fialkov et al. 2012; McQuinn & O’Leary 2012; Muñoz 2019a; see however Cain et al. 2020; Park et al. 2021 for the effect during the epoch of reionization). Exotic interactions between dark matter and baryons can modulate the gas temperature with the same velocity feature (Muñoz et al. 2015b; Fialkov et al. 2018; Muñoz et al. 2018). Regardless of whether its origin is astrophysical or exotic, a detection of VAOs would provide us with a standard ruler (Muñoz 2019b).

To model such a scenario, we replace the matter power spectrum in Equation (11) with $\Delta_{v_{cb}}^2(k)$ and constrain the corresponding b_v . The quantity $\Delta_{v_{cb}}^2(k)$ is the power spectrum of $\delta_{v_{cb}} = \sqrt{3/2} (v_{cb}^2/v_{rms}^2 - 1)$, which is dimensionless, z -independent, and has unit variance (Dalal et al. 2010; Ali-Ha ‘moud et al. 2014). Using the H21 data, we can set the 95% confidence limits $b_{v_{cb}} < \{49, 175\}$ mK at $z = \{7.9, 10.4\}$, respectively. For comparison, a typical VAO contribution to the 21 cm power spectrum found in the simulations of Muñoz (2019a) and Fialkov et al. (2013) is roughly $\Delta_{21}^2 \approx 10$ mK² (which peaks at $z \approx 15$ during their X-ray heating epoch). This would convert into a bias $b_{v_{cb}} \approx 5$ mK, far below the limits above.

Appendix B Emulator for Extra Radio Background Simulations

Emulators are widely used to rapidly evaluate 21 cm PS across the broad astrophysical parameter space (Kern et al. 2017; Schmit & Pritchard 2018; Jennings et al. 2019; Ghara et al. 2020; Mondal et al. 2020). This technique allows us to interpolate existing simulations instead of running a new simulation at every point of the parameter space. Here we choose to emulate the power spectrum using a neural network (multilayer perceptron) regression, implemented in `scikit-learn` (Pedregosa et al. 2011).

We train emulators for the two cases of radio backgrounds separately: one data set corresponding to the models with radio background created by galaxies (10,700 samples), the other corresponding to the smooth synchrotron radio background (10,300 samples). These sets were created by randomly sampling the parameters V_c , f_* , f_X , τ , and f_r or A_r within the bounds described in Section 8.2. To calculate the likelihood of each parameter set, we go through a two-step process. First, we emulate the power spectrum assuming no RSDs (emulators based on the corresponding $\sim 10,000$ models mentioned above), then we use another emulator to account for the effect of RSDs (using another ~ 2000 simulations calculated separately for each type of radio background), boosting or suppressing the PS by a factor between 0.7 and 3 depending on the values of the astrophysical parameters. The final PSs are used in the likelihood calculation. We also emulate the values of spin, gas, and radiation temperatures (training data $\sim 10,000$




Figure 19. Top: accuracy of the neural network predictions. The triangle plots show the relative accuracy of predicting the power spectrum at $z = 7.93$, $k = 0.13 \text{ cMpc}^{-1}$, for all parameters in the test set. The relative emulator error is shown by the color of the points, from yellow (no error) to red (50% error). For comparison we add the black contour lines from Figure 14 corresponding to the 68% and 95% parameter confidence intervals, respectively. The histogram insert shows the overall distribution of the error $\sigma_{\text{rel}} = (\Delta_{\text{predict}}^2 - \Delta_{\text{test}}^2) / (\Delta_{\text{predict}}^2 + 1 \text{ mK}^2)$, for $z = 7.93$, $k = 0.13 \text{ cMpc}^{-1}$ (orange histogram) and for all points combined (blue). The orange histogram can be approximated by a normal distribution with scale $\sigma = 0.2$ (pink curve). Bottom: the left panel shows the covariance matrix of the relative emulator error (σ_{rel} as used above, variance of 0.04 corresponds to 20% error) over both bands and the first few wavenumbers. The emulator error correlation between the different bands is negligible, but the errors within a band can be strongly correlated with a high correlation coefficient $r \approx 0.9$. The middle panel shows the overall impact of taking the emulator uncertainty into account by looking at the marginalized f_{X} , f_{X} , and τ constraints. We see that assuming no emulator error (orange), 20% (black, as used in results), or 40% (green) does not noticeably affect the parameter constraints, as the observational uncertainty dominates the error budget in the relevant region. The right panel shows parameter constraints using a restricted likelihood, based on just a single data point per band (red contours, using $k = 0.17$ and 0.13 cMpc^{-1} for Bands 1 and 2, respectively) compared to the full likelihood (black contours, as used for our scientific results). The constraints with the single data points are just slightly weaker than using the full likelihood, which shows that the parameter constraints are mainly driven by the lowest point in each band.

cases for each type of radio background) to infer constraints on the physical properties of the IGM (e.g., Figures 11 and 14, insert). For completeness we also create a data set (~ 2000 models) with no radio background to test the emulator in this limit and to calculate the envelope of these models in Figure 13 (black dashed line).

The input parameters for all of our emulators are taken to be natural logarithms of model parameters ($\ln f_{\text{star}}$, $\ln V_c$, $\ln f_{\text{X}}$, $\ln \tau$, and $\ln f_r$ or $\ln A_r$) to cover the dynamic range (using $\ln \tau$ for simplicity; we use linear τ in the analysis). As part of the emulator, the inputs are also shifted and scaled such that the training input values have zero mean and unit variance. The

network itself consists of four hidden layers for PSs and RSDs and two hidden layers for temperatures. Each layer has 100 nodes, employs an ReLU activation function (Nair & Hinton 2010), and is trained using the adam optimizer (Kingma & Ba 2014). The PS emulators return 78 outputs corresponding to the two log-power spectra (or RSD boost factors) at $z=7.93$ and 10.37 and at the 39 wavenumbers measured by HERA; the temperature emulators have a single output, which is the logarithm of temperature. The simulation PSs are only computed at the wavenumbers $\leq 1.1 \text{ cMpc}^{-1}$ due to the limited resolution of the simulation. However, values at smaller scales are not relevant as the HERA limits there exceed the expected PSs. To take into account the large dynamic range of possible PSs, we use the logarithm $\log(\Delta_{21}^2 + 1 \text{ mK}^2)$ as the target for the emulator. Adding the baseline 1 mK^2 term helps to improve the performance of the emulator for the values of PSs that can be constrained by HERA.

The accuracy in predicting the PS is evaluated using a test data set (2000 test samples for power spectrum and temperatures, 1000 for RSD boost factor and final PS). Since we train the emulator on logarithms of the PS, the emulator errors on these logarithmic values are approximately the same over the parameter space. This implies that the error on the power spectrum Δ_{21}^2 is proportional to the power spectrum value itself. We find that the deviations are well within the 20% relative error in the theoretical modeling that we assume here, as seen from the histogram insert in Figure 19, which shows the distribution of the relative error (in $\Delta_{21}^2 + 1 \text{ mK}^2$) for all of the predicted PS points (blue) at the most relevant $k=0.13 \text{ cMpc}^{-1}$, $z=7.93$ band (orange), along with the assumed Gaussian uncertainty (pink curve, $\sigma=0.2$).

The triangle plot in Figure 19 shows the relative emulator error (including RSD, and again as $\Delta_{21}^2 + 1 \text{ mK}^2$) for every parameter point in the test set, together with the contour lines from Figure 14 in the background. This shows that no parameter region has a particularly high emulator error, and that the test (and also training) data are evenly distributed over the whole parameter space.

We assume this relative emulator uncertainty to be approximately Gaussian and independent, and add it in quadrature to the observational error. We note that the emulator error is strongly correlated within each band as shown in the lower-left panel of Figure 19. However, we neglect this correlation to derive the marginal likelihood in Section 3.2.1. This is justified as (i) the emulator uncertainty is subdominant compared to the large observational errors and barely affects the parameter constraints (lower-central panel of Figure 19), and (ii) the constraints are mainly driven by the single lowest data points in each band (the lower-right panel of Figure 19 shows that using just these two points only slightly weakens the constraints to $f_X > 0.15$ and $f_r < 442$ compared to $f_X > 0.25$ and $f_r < 397$ when the complete data set is used). This implies that the correlations have only a small, secondary effect.

ORCID iDs

James E. Aguirre  <https://orcid.org/0000-0002-4810-666X>
 Rennan Barkana  <https://orcid.org/0000-0002-1557-693X>
 Adam P. Beardsley  <https://orcid.org/0000-0001-9428-8233>
 Gianni Bernardi  <https://orcid.org/0000-0002-0916-7443>
 Judd D. Bowman  <https://orcid.org/0000-0002-8475-2036>
 Richard F. Bradley  <https://orcid.org/0000-0003-1172-8331>
 Philip Bull  <https://orcid.org/0000-0001-5668-3101>

Chris L. Carilli  <https://orcid.org/0000-0001-6647-3861>
 David R. DeBoer  <https://orcid.org/0000-0003-3197-2294>
 Joshua S. Dillon  <https://orcid.org/0000-0003-3336-9958>
 Aaron Ewall-Wice  <https://orcid.org/0000-0002-0086-7363>
 Anastasia Fialkov  <https://orcid.org/0000-0002-1369-633X>
 Steven R. Furlanetto  <https://orcid.org/0000-0002-0658-1243>
 Deepthi Gorthi  <https://orcid.org/0000-0002-0829-167X>
 Bradley Greig  <https://orcid.org/0000-0002-4085-2094>
 Bryna J. Hazelton  <https://orcid.org/0000-0001-7532-645X>
 Stefan Heimersheim  <https://orcid.org/0000-0001-9631-4212>
 Jacqueline N. Hewitt  <https://orcid.org/0000-0002-4117-570X>
 Daniel C. Jacobs  <https://orcid.org/0000-0002-0917-2269>
 Nicholas S. Kern  <https://orcid.org/0000-0002-8211-1892>
 Joshua Kerrigan  <https://orcid.org/0000-0002-1876-272X>
 Piyanat Kittiwisit  <https://orcid.org/0000-0003-0953-313X>
 Saul A. Kohn  <https://orcid.org/0000-0001-6744-5328>
 Matthew Kolopanis  <https://orcid.org/0000-0002-2950-2974>
 Paul La Plante  <https://orcid.org/0000-0002-4693-0102>
 Adrian Liu  <https://orcid.org/0000-0001-6876-0928>
 Yin-Zhe Ma  <https://orcid.org/0000-0001-8108-0986>
 Andrei Mesinger  <https://orcid.org/0000-0003-3374-1772>
 Jordan Mirocha  <https://orcid.org/0000-0002-8802-5581>
 Miguel F. Morales  <https://orcid.org/0000-0001-7694-4030>
 Julian B. Muñoz  <https://orcid.org/0000-0002-8984-0465>
 Abraham R. Neben  <https://orcid.org/0000-0001-7776-7240>
 Chuneeta D. Nunhokee  <https://orcid.org/0000-0002-5445-6586>
 Nipanjana Patra  <https://orcid.org/0000-0002-9457-1941>
 Jonathan C. Pober  <https://orcid.org/0000-0002-3492-0433>
 Yuxiang Qin  <https://orcid.org/0000-0002-4314-1810>
 Itamar Reis  <https://orcid.org/0000-0002-6203-7496>
 Mario G. Santos  <https://orcid.org/0000-0003-3892-3073>
 Peter Sims  <https://orcid.org/0000-0002-2871-0413>
 Nithyanandan Thyagarajan  <https://orcid.org/0000-0003-1602-7868>

References

- Aguirre, J. E., Murray, S. G., Pascua, R., et al. 2021, arXiv:2104.09547
 Ali-Ha "imoud, Y., Meerburg, P. D., & Yuan, S. 2014, *PhRvD*, 89, 083506
 Bañados, E., Venemans, B. P., Mazzucchelli, C., et al. 2018, *Natur*, 553, 473
 Bar, N., Blum, K., & D'Amico, G. 2020, *PhRvD*, 101, 123025
 Barkana, R. 2018, *Natur*, 555, 71
 Barkana, R., & Loeb, A. 2004, *ApJ*, 609, 474
 Barkana, R., & Loeb, A. 2005a, *ApJL*, 624, L65
 Barkana, R., & Loeb, A. 2005b, *ApJ*, 626, 1
 Barkana, R., Outmezguine, N. J., Redigolo, D., & Volansky, T. 2018, *PhRvD*, 98, 103005
 Basu-Zych, A. R., Lehmer, B. D., Hornschemeier, A. E., et al. 2013, *ApJ*, 762, 45
 Becker, G. D., Bolton, J. S., Madau, P., et al. 2015, *MNRAS*, 447, 3402
 Behroozi, P., Wechsler, R. H., Hearin, A. P., & Conroy, C. 2019, *MNRAS*, 488, 3143
 Behroozi, P. S., & Silk, J. 2015, *ApJ*, 799, 32
 Berlin, A., Hooper, D., Krnjaic, G., & McDermott, S. D. 2018, *PhRvL*, 121, 011102
 Bharadwaj, S., & Ali, S. S. 2004, *MNRAS*, 352, 142
 Binnie, T., & Pritchard, J. R. 2019, *MNRAS*, 487, 1160
 Boddy, K. K., Gluscevic, V., Poulin, V., et al. 2018, *PhRvD*, 98, 123506
 Bolton, J. S., Haehnelt, M. G., Warren, S. J., et al. 2011, *MNRAS*, 416, L70
 Bosman, S. E. I., Fan, X., Jiang, L., et al. 2018, *MNRAS*, 479, 1055
 Bouwens, R. J., Illingworth, G. D., Oesch, P. A., et al. 2015, *ApJ*, 803, 34
 Bouwens, R. J., Oesch, P. A., Labbé, I., et al. 2016, *ApJ*, 830, 67
 Bowman, J. D., Rogers, A. E. E., Monsalve, R. A., Mozdzen, T. J., & Mahesh, N. 2018, *Natur*, 555, 67

- Bradley, R. F., Tauscher, K., Rapetti, D., & Burns, J. O. 2019, *ApJ*, **874**, 153
- Brandenberger, R., Cyr, B., & Shi, R. 2019, *JCAP*, **09**, 009
- Brorby, M., Kaaret, P., & Prestwich, A. 2014, *MNRAS*, **441**, 2346
- Brorby, M., Kaaret, P., Prestwich, A., & Mirabel, I. F. 2016, *MNRAS*, **457**, 4081
- Cain, C., D'Aloisio, A., Iršič, V., McQuinn, M., & Trac, H. 2020, *ApJ*, **898**, 168
- Caruana, J., Bunker, A. J., Wilkins, S. M., et al. 2014, *MNRAS*, **443**, 2831
- Chang, J. H., Essig, R., & McDermott, S. D. 2018, *JHEP*, **09**, 051
- Chen, X., & Miralda-Escudé, J. 2004, *ApJ*, **602**, 1
- Chianese, M., Di Bari, P., Farrag, K., & Samanta, R. 2019, *PhLB*, **790**, 64
- Choudhury, T. R., Mukherjee, S., & Paul, S. 2021a, *MNRAS*, **501**, L7
- Choudhury, T. R., Paranjape, A., & Bosman, S. E. I. 2021b, *MNRAS*, **501**, 5782
- Chuzhoy, L., & Shapiro, P. R. 2007, *ApJ*, **655**, 843
- Ciardi, B., Salvaterra, R., & Di Matteo, T. 2010, *MNRAS*, **401**, 2635
- Cohen, A., Fialkov, A., & Barkana, R. 2016, *MNRAS*, **459**, L90
- Cohen, A., Fialkov, A., Barkana, R., & Monsalve, R. A. 2020, *MNRAS*, **495**, 4845
- Coles, P., & Jones, B. 1991, *MNRAS*, **248**, 1
- Dai, W.-M., Ma, Y.-Z., Guo, Z.-K., & Cai, R.-G. 2019, *PhRvD*, **99**, 043524
- Dalal, N., Pen, U.-L., & Seljak, U. 2010, *JCAP*, **1011**, 007
- D'Aloisio, A., McQuinn, M., Maupin, O., et al. 2019, *ApJ*, **874**, 154
- Das, A., Mesinger, A., Pallottini, A., Ferrara, A., & Wise, J. H. 2017, *MNRAS*, **469**, 1166
- Davies, F. B., Hennawi, J. F., Bañados, E., et al. 2018, *ApJ*, **864**, 142
- de Belsunce, R., Gratton, S., Coulton, W., & Efstathiou, G. 2021, *MNRAS*, **507**, 1072
- DeBoer, D. R., Parsons, A. R., Aguirre, J. E., et al. 2017, *PASP*, **129**, 045001
- Dillon, J. S., Lee, M., Ali, Z. S., et al. 2020, *MNRAS*, **499**, 5840
- Dillon, J. S., & Parsons, A. R. 2016, *ApJ*, **826**, 181
- Dillon, J. S., Liu, A., Williams, C. L., et al. 2014, *PhRvD*, **89**, 023002
- DiLullo, C., Taylor, G. B., & Dowell, J. 2020, *JAI*, **9**, 2050008
- Douna, V. M., Pellizza, L. J., Mirabel, I. F., & Pedrosa, S. E. 2015, *A&A*, **579**, A44
- Dowell, J., & Taylor, G. B. 2018, *ApJL*, **858**, L9
- Evoli, C., Mesinger, A., & Ferrara, A. 2014, *JCAP*, **11**, 024
- Ewall-Wice, A., Chang, T. C., Lazio, J., et al. 2018, *ApJ*, **868**, 63
- Ewall-Wice, A., Chang, T.-C., & Lazio, T. J. W. 2020, *MNRAS*, **492**, 6086
- Feng, C., & Holder, G. 2018, *ApJL*, **858**, L17
- Feroz, F., Hobson, M. P., & Bridges, M. 2009, *MNRAS*, **398**, 1601
- Fialkov, A., & Barkana, R. 2019, *MNRAS*, **486**, 1763
- Fialkov, A., Barkana, R., & Cohen, A. 2018, *PhRvL*, **121**, 011101
- Fialkov, A., Barkana, R., & Jarvis, M. 2020, *MNRAS*, **491**, 3108
- Fialkov, A., Barkana, R., Tselikhovich, D., & Hirata, C. M. 2012, *MNRAS*, **424**, 1335
- Fialkov, A., Barkana, R., & Visbal, E. 2014, *Natur*, **506**, 197
- Fialkov, A., Barkana, R., Visbal, E., Tselikhovich, D., & Hirata, C. M. 2013, *MNRAS*, **432**, 2909
- Fialkov, A., Cohen, A., Barkana, R., & Silk, J. 2017, *MNRAS*, **464**, 3498
- Field, G. B. 1958, *PIRE*, **46**, 240
- Field, G. B. 1959, *ApJ*, **129**, 536
- Fixsen, D. J., Kogut, A., Levin, S., et al. 2011, *ApJ*, **734**, 5
- Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, *PASP*, **125**, 306
- Fragos, T., Lehmer, B., Tremmel, M., et al. 2013, *ApJ*, **764**, 41
- Fragos, T., Lehmer, B. D., Naoz, S., Zezas, A., & Basu-Zych, A. 2013, *ApJL*, **776**, L31
- Fraser, S., Hektor, A., Hütsi, G., et al. 2018, *PhLB*, **785**, 159
- Furlanetto, S. R. 2006, *MNRAS*, **371**, 867
- Furlanetto, S. R., Mirocha, J., Mebane, R. H., & Sun, G. 2017, *MNRAS*, **472**, 1576
- Furlanetto, S. R., Zaldarriaga, M., & Hernquist, L. 2004, *ApJ*, **613**, 1
- Gürkan, G., Hardcastle, M. J., Smith, D. J. B., et al. 2018, *MNRAS*, **475**, 3010
- Gehlot, B. K., Mertens, F. G., Koopmans, L. V. E., et al. 2019, *MNRAS*, **488**, 4271
- Ghara, R., Choudhury, T. R., & Datta, K. K. 2015, *MNRAS*, **447**, 1806
- Ghara, R., Giri, S. K., Ciardi, B., Mellema, G., & Zaroubi, S. 2021, *MNRAS*, **503**, 4551
- Ghara, R., Giri, S. K., Mellema, G., et al. 2020, *MNRAS*, **493**, 4728
- Gillet, N. J. F., Mesinger, A., & Park, J. 2020, *MNRAS*, **491**, 1980
- Greig, B., & Mesinger, A. 2015, *MNRAS*, **449**, 4246
- Greig, B., & Mesinger, A. 2017a, *MNRAS*, **472**, 2651
- Greig, B., & Mesinger, A. 2017b, *MNRAS*, **465**, 4838
- Greig, B., & Mesinger, A. 2018, *MNRAS*, **477**, 3217
- Greig, B., Mesinger, A., & Bañados, E. 2019, *MNRAS*, **484**, 5094
- Greig, B., Mesinger, A., Haiman, Z., & Simcoe, R. A. 2017, *MNRAS*, **466**, 4239
- Greig, B., Trott, C. M., Barry, N., et al. 2021a, *MNRAS*, **500**, 5322
- Greig, B., Mesinger, A., Koopmans, L. V. E., et al. 2021b, *MNRAS*, **501**, 1
- Handley, W. 2019, *JOSS*, **4**, 1414
- Hardcastle, M. J., Gürkan, G., van Weeren, R. J., et al. 2016, *MNRAS*, **462**, 1910
- Heinrich, C., & Hu, W. 2021, *PhRvD*, **104**, 063505
- HERA Collaboration, Abdurashidova, Z., Aguirre, J. E., et al. 2021, *ApJ*, in press (arXiv:2108.02263)
- Hills, R., Kulkarni, G., Meerburg, P. D., & Puchwein, E. 2018, *Natur*, **564**, E32
- Hirata, C. M. 2006, *MNRAS*, **367**, 259
- Hirata, C. M., & Sigurdson, K. 2007, *MNRAS*, **375**, 1241
- Holdom, B. 1986, *PhLB*, **166**, 196
- Hui, L., & Gnedin, N. Y. 1997, *MNRAS*, **292**, 27
- Iliev, I. T., Mellema, G., Ahn, K., et al. 2014, *MNRAS*, **439**, 725
- Jaackel, J., & Ringwald, A. 2010, *ARNPS*, **60**, 405
- Jana, R., Nath, B. B., & Biermann, P. L. 2019, *MNRAS*, **483**, 5329
- Jennings, W. D., Watkinson, C. A., Abdalla, F. B., & McEwen, J. D. 2019, *MNRAS*, **483**, 2907
- Jensen, H., Datta, K. K., Mellema, G., et al. 2013b, *MNRAS*, **435**, 460
- Jensen, H., Laursen, P., Mellema, G., et al. 2013a, *MNRAS*, **428**, 1366
- Jensen, H., Majumdar, S., Mellema, G., et al. 2016, *MNRAS*, **456**, 66
- Kaiser, N. 1987, *MNRAS*, **227**, 1
- Kaur, H. D., Gillet, N., & Mesinger, A. 2020, *MNRAS*, **495**, 2354
- Keating, L. C., Weinberger, L. H., Kulkarni, G., et al. 2020, *MNRAS*, **491**, 1736
- Kern, N. S., Dillon, J. S., Parsons, A. R., et al. 2020b, *ApJ*, **890**, 122
- Kern, N. S., & Liu, A. 2021, *MNRAS*, **501**, 1463
- Kern, N. S., Liu, A., Parsons, A. R., Mesinger, A., & Greig, B. 2017, *ApJ*, **848**, 23
- Kern, N. S., Parsons, A. R., Dillon, J. S., et al. 2020a, *ApJ*, **888**, 70
- Kimm, T., Katz, H., Haehnelt, M., et al. 2017, *MNRAS*, **466**, 4826
- Kingma, D. P., & Ba, J. 2014, arXiv:1412.6980
- Kolopanis, M., Jacobs, D. C., Cheng, C., et al. 2019, *ApJ*, **883**, 133
- Kovetz, E. D., Poulin, V., Gluscevic, V., et al. 2018, *PhRvD*, **98**, 103529
- Kulkarni, G., Keating, L. C., Haehnelt, M. G., et al. 2019, *MNRAS*, **485**, L24
- La Plante, P., Battaglia, N., Natarajan, A., et al. 2014, *ApJ*, **789**, 31
- Lehmer, B. D., Alexander, D. M., Bauer, F. E., et al. 2010, *ApJ*, **724**, 559
- Lehmer, B. D., Basu-Zych, A. R., Mineo, S., et al. 2016, *ApJ*, **825**, 7
- Lehmer, B. D., Xue, Y. Q., Brandt, W. N., et al. 2012, *ApJ*, **752**, 46
- Leitherer, C., Schaerer, D., Goldader, J. D., et al. 1999, *ApJS*, **123**, 3
- Lewis, A. 2019, arXiv:1910.13970
- Lewis, A., & Bridle, S. 2002, *PhRvD*, **66**, 103511
- Lewis, J. S. W., Ocvirk, P., Aubert, D., et al. 2020, *MNRAS*, **496**, 4342
- Li, W., Pober, J. C., Barry, N., et al. 2019, *ApJ*, **887**, 141
- Lidz, A., Oh, S. P., & Furlanetto, S. R. 2006, *ApJL*, **639**, L47
- Lidz, A., Zahn, O., McQuinn, M., et al. 2007, *ApJ*, **659**, 865
- Lidz, A., Zahn, O., McQuinn, M., Zaldarriaga, M., & Hernquist, L. 2008, *ApJ*, **680**, 962
- Liu, A., & Shaw, J. R. 2020, *PASP*, **132**, 062001
- Liu, A., & Tegmark, M. 2011, *PhRvD*, **83**, 103006
- Liu, H., Outmezguine, N. J., Redigolo, D., & Volansky, T. 2019, *PhRvD*, **100**, 123011
- Liu, H., & Slatyer, T. R. 2018, *PhRvD*, **98**, 023501
- Lopez-Honorez, L., Mena, O., Moliné, A., Palomares-Ruiz, S., & Vincent, A. C. 2016, *JCAP*, **8**, 004
- Ma, X., Quataert, E., Wetzel, A., et al. 2020, *MNRAS*, **498**, 2001
- Madau, P., Meiksin, A., & Rees, M. J. 1997, *ApJ*, **475**, 429
- Mao, Y., Shapiro, P. R., Mellema, G., et al. 2012, *MNRAS*, **422**, 926
- Mason, C. A., Trenti, M., & Treu, T. 2015, *ApJ*, **813**, 21
- Mason, C. A., Treu, T., Dijkstra, M., et al. 2018, *ApJ*, **856**, 2
- Mason, C. A., Fontana, A., Treu, T., et al. 2019, *MNRAS*, **485**, 3947
- McGreer, I. D., Mesinger, A., & D'Odorico, V. 2015, *MNRAS*, **447**, 499
- McQuinn, M., Furlanetto, S. R., Hernquist, L., Zahn, O., & Zaldarriaga, M. 2005, *ApJ*, **630**, 643
- McQuinn, M., & O'Leary, R. M. 2012, *ApJ*, **760**, 3
- Mebane, R. H., Mirocha, J., & Furlanetto, S. R. 2020, *MNRAS*, **493**, 1217
- Meiksin, A. 2021, *RNAAS*, **5**, 126
- Mertens, F. G., Mevius, M., Koopmans, L. V. E., et al. 2020, *MNRAS*, **493**, 1662
- Mesinger, A. 2010, *MNRAS*, **407**, 1328
- Mesinger, A., Aykutalp, A., Vanzella, E., et al. 2015, *MNRAS*, **446**, 566
- Mesinger, A., Ewall-Wice, A., & Hewitt, J. 2014, *MNRAS*, **439**, 3262
- Mesinger, A., Ferrara, A., & Spiegel, D. S. 2013, *MNRAS*, **431**, 621
- Mesinger, A., & Furlanetto, S. 2007, *ApJ*, **669**, 663

- Mesinger, A., Furlanetto, S., & Cen, R. 2011, *MNRAS*, **411**, 955
- Mesinger, A., & Haiman, Z. 2004, *ApJL*, **611**, L69
- Mineo, S., Gilfanov, M., & Sunyaev, R. 2012, *MNRAS*, **419**, 2095
- Mirocha, J., & Furlanetto, S. R. 2019, *MNRAS*, **483**, 1980
- Mirocha, J., Furlanetto, S. R., & Sun, G. 2017, *MNRAS*, **464**, 1365
- Mittal, S., & Kulkarni, G. 2021, *MNRAS*, **503**, 4264
- Mondal, R., Fialkov, A., Fling, C., et al. 2020, *MNRAS*, **498**, 4178
- Morales, M. F., Hazelton, B., Sullivan, I., & Beardsley, A. 2012, *ApJ*, **752**, 137
- Moster, B. P., Naab, T., & White, S. D. M. 2013, *MNRAS*, **428**, 3121
- Muñoz, J. B. 2019a, *PhRvD*, **100**, 063538
- Muñoz, J. B. 2019b, *PhRvL*, **123**, 131301
- Muñoz, J. B., Ali-Ha “imoud, Y., & Kamionkowski, M. 2015a, *PhRvD*, **92**, 083508
- Muñoz, J. B., & Cyr-Racine, F.-Y. 2021, *PhRvD*, **103**, 023512
- Muñoz, J. B., Dvorkin, C., & Loeb, A. 2018, *PhRvL*, **121**, 121301
- Muñoz, J. B., Kovetz, E. D., & Ali-Ha “imoud, Y. 2015b, *PhRvD*, **92**, 083528
- Muñoz, J. B., & Loeb, A. 2018, *Natur*, **557**, 684
- Murray, S., Greig, B., Mesinger, A., et al. 2020, *JOSS*, **5**, 2582
- Mutch, S. J., Geil, P. M., Poole, G. B., et al. 2016, *MNRAS*, **462**, 250
- Nair, V., & Hinton, G. E. 2010, in ICM 10, Proc. 27th Int. Conf. Int. Conf. Mach. Learn. (Madison, WI: Omnipress), <https://dl.acm.org/doi/proceedings/10.5555/3104322>
- Naoz, S., Yoshida, N., & Gnedin, N. Y. 2012, *ApJ*, **747**, 128
- Nasir, F., & D’Aloisio, A. 2020, *MNRAS*, **494**, 3080
- Ocvirk, P., Aubert, D., Sorce, J. G., et al. 2018, *MNRAS*, **496**, 4087
- Oesch, P. A., Bouwens, R. J., Illingworth, G. D., Labbe, I., & Stefanon, M. 2017, *ApJ*, **855**, 105
- Okamoto, T., Gao, L., & Theuns, T. 2008, *MNRAS*, **390**, 920
- Paardekooper, J.-P., Khochfar, S., & Dalla Vecchia, C. 2015, *MNRAS*, **451**, 2544
- Pacucci, F., Mesinger, A., Mineo, S., & Ferrara, A. 2014, *MNRAS*, **443**, 678
- Paranjape, A., & Choudhury, T. R. 2014, *MNRAS*, **442**, 1470
- Park, H., Shapiro, P. R., Ahn, K., Yoshida, N., & Hirano, S. 2021, *ApJ*, **908**, 96
- Park, J., Mesinger, A., Greig, B., & Gillet, N. 2019, *MNRAS*, **484**, 933
- Parsons, A. R., Backer, D. C., Foster, G. S., et al. 2010, *AJ*, **139**, 1468
- Pedregosa, F., Varoquaux, G., Gramfort, A., et al. 2011, *J. Mach. Learn. Res.*, **12**, 2825, <http://jmlr.org/papers/v12/pedregosa11a.html>
- Pentericci, L., Vanzella, E., Fontana, A., et al. 2014, *ApJ*, **793**, 113
- Philip, L., Abdurashidova, Z., Chiang, H. C., et al. 2019, *JAI*, **8**, 1950004
- Planck Collaboration, Aghanim, N., Akrami, Y., et al. 2020, *A&A*, **641**, A6
- Pober, J. C., Ali, Z. S., Parsons, A. R., et al. 2015, *ApJ*, **809**, 62
- Pospelov, M., Pradler, J., Ruderman, J. T., & Urbano, A. 2018, *PhRvL*, **121**, 031103
- Price, D. C., Greenhill, L. J., Fialkov, A., et al. 2018, *MNRAS*, **478**, 4193
- Prinz, A. A., Baggs, R., Ballam, J., et al. 1998, *PhRvL*, **81**, 1175
- Pritchard, J. R., & Loeb, A. 2008, *PhRvD*, **78**, 103511
- Pritchard, J. R., & Loeb, A. 2012, *RPPh*, **75**, 086901
- Qin, Y., Mesinger, A., Bosman, S. E. I., & Viel, M. 2021a, *MNRAS*, **506**, 2390
- Qin, Y., Mesinger, A., Greig, B., & Park, J. 2021b, *MNRAS*, **501**, 4748
- Qin, Y., Mesinger, A., Park, J., Greig, B., & Muñoz, J. B. 2020a, *MNRAS*, **495**, 123
- Reichardt, C. L., Patil, S., Ade, P. A. R., et al. 2021, *ApJ*, **908**, 199
- Reis, I., Fialkov, A., & Barkana, R. 2020, *MNRAS*, **499**, 5993
- Reis, I., Fialkov, A., & Barkana, R. 2021, *MNRAS*, **506**, 5479
- Robertson, B. E., Ellis, R. S., Furlanetto, S. R., & Dunlop, J. S. 2015, *ApJL*, **802**, L19
- Ross, H. E., Giri, S. K., Mellema, K. L., et al. 2021, *MNRAS*, **506**, 3717
- Scalo, J. 1998, in ASP Conf. Ser. 142, The Stellar Initial Mass Function, ed. G. Gilmore & D. Howell (San Francisco, CA: ASP), 201
- Schenker, M. A., Stark, D. P., Ellis, R. S., et al. 2012, *ApJ*, **744**, 179
- Schmit, C. J., & Pritchard, J. R. 2018, *MNRAS*, **475**, 1213
- Scoccimarro, R. 1998, *MNRAS*, **299**, 1097
- Seiffert, M., Fixsen, D. J., Kogut, A., et al. 2011, *ApJ*, **734**, 6
- Shaver, P. A., Windhorst, R. A., Madau, P., & de Bruyn, A. G. 1999, *A&A*, **345**, 380
- Sims, P. H., & Pober, J. C. 2019, *MNRAS*, **488**, 2904
- Singh, S., & Subrahmanyan, R. 2019, *ApJ*, **880**, 26
- Singh, S., Subrahmanyan, R., Shankar, N. U., et al. 2018, *ExA*, **45**, 269
- Slatyer, T. R., & Wu, C.-L. 2018, *PhRvD*, **98**, 023013
- Sobacchi, E., & Mesinger, A. 2013, *MNRAS*, **432**, L51
- Sobacchi, E., & Mesinger, A. 2014, *MNRAS*, **440**, 1662
- Springel, V., & Hernquist, L. 2003, *MNRAS*, **339**, 312
- Stark, D. P., Ellis, R. S., Chiu, K., Ouchi, M., & Bunker, A. 2010, *MNRAS*, **408**, 1628
- Subrahmanyan, R., & Cowsik, R. 2013, *ApJ*, **776**, 42
- Sun, G., & Furlanetto, S. R. 2016, *MNRAS*, **460**, 417
- Tacchella, S., Bose, S., Conroy, C., Eisenstein, D. J., & Johnson, B. D. 2018, *ApJ*, **868**, 92
- Tan, J., Liu, A., Kern, N. S., et al. 2021, *ApJS*, **255**, 26
- Tashiro, H., Kadota, K., & Silk, J. 2014, *PhRvD*, **90**, 083522
- Tauscher, K., Rapetti, D., & Burns, J. O. 2020, *ApJ*, **897**, 132
- Tegmark, M. 1997, *PhRvD*, **55**, 5895
- Thériault, R., Mirocha, J. T., & Brandenberger, R. 2021, *JCAP*, **10**, 046
- Trott, C. M., Jordan, C. H., Midgley, S., et al. 2020, *MNRAS*, **493**, 4711
- Tselikhovich, D., & Hirata, C. 2010, *PhRvD*, **82**, 083520
- Venumadhav, T., Dai, L., Kaurov, A., & Zaldarriaga, M. 2018, *PhRvD*, **98**, 103513
- Visbal, E., Barkana, R., Fialkov, A., Tselikhovich, D., & Hirata, C. M. 2012, *Natur*, **487**, 70
- Vogel, H., & Redondo, J. 2014, *JCAP*, **02**, 029
- Voytek, T. C., Natarajan, A., Jáuregui García, J. M., Peterson, J. B., & López-Cruz, O. 2014, *ApJL*, **782**, L9
- Wang, F., Davies, F. B., Yang, J., et al. 2020, *ApJ*, **896**, 23
- Wang, F., Yang, J., Fan, X., et al. 2021, *ApJL*, **907**, L1
- Wouthuysen, S. A. 1952, *AJ*, **57**, 31
- Xu, H., Wise, J. H., Norman, M. L., Ahn, K., & O’Shea, B. W. 2016, *ApJ*, **833**, 84
- Yang, J., Wang, F., Fan, X., et al. 2020, *ApJL*, **897**, L14
- Yung, L. Y. A., Somerville, R. S., Popping, G., et al. 2019, *MNRAS*, **490**, 2855
- Zahn, O., Lidz, A., McQuinn, M., et al. 2007, *ApJ*, **654**, 12
- Zahn, O., Mesinger, A., McQuinn, M., et al. 2011, *MNRAS*, **414**, 727