Monthly Notices of the ROYAL ASTRONOMICAL SOCIETY

MNRAS **514,** 2010–2030 (2022) https://doi.org/10.1093/mnras/stac1479
Advance Access publication 2022 May 31

A galaxy-free phenomenological model for the 21-cm power spectrum during reionization

Jordan Mirocha [1]★ Julian B. Muñoz,[2] Steven R. Furlanetto [3] Adrian Liu [1] and Andrei Mesinger [4]

[1]*Department of Physics & McGill Space Institute, McGill University, 3600 Rue University, Montréal, QC H3A 2T8, Canada*
[2]*Center for Astrophysics | Harvard & Smithsonian, 60 Garden St, Cambridge, MA 02138, USA*
[3]*Department of Physics and Astronomy, University of California, Los Angeles, CA 90095, USA*
[4]*Scuola Normale Superiore, Piazza dei Cavalieri 7, I-56126 Pisa, Italy*

Accepted 2022 May 25. Received 2022 May 25; in original form 2022 January 17

ABSTRACT

Upper limits from the current generation of interferometers targeting the 21-cm signal from high redshifts have recently begun to rule out physically realistic, though still extreme, models of the Epoch of Reionization (EoR). While inferring the detailed properties of the first galaxies is one of the most important motivations for measuring the high-z 21-cm signal, they can also provide useful constraints on the properties of the intergalactic medium (IGM). Motivated by this, we build a simple, phenomenological model for the 21-cm power spectrum that works directly in terms of IGM properties, which bypasses the computationally expensive 3D semi-numerical modeling generally employed in inference pipelines and avoids explicit assumptions about galaxy properties. The key simplifying assumptions are that (i) the ionization field is binary, and composed of spherical bubbles with an abundance described well by a parametric bubble size distribution, and (ii) that the spin temperature of the 'bulk' IGM outside bubbles is uniform. Despite the simplicity of the model, the mean ionized fraction and spin temperature of the IGM recovered from mock 21-cm power spectra generated with 21cmFAST are generally in good agreement with the true input values. This suggests that it is possible to obtain comparable constraints on the IGM using models with very different assumptions, parameters, and priors. Our approach will thus be complementary to semi-numerical models as upper limits continue to improve in the coming years.

Key words: galaxies: high-redshift – intergalactic medium – dark ages, reionization, first stars – diffuse radiation.

1 INTRODUCTION

The 21-cm background (Madau, Meiksin & Rees 1997) has long been recognized as a powerful probe of the intergalactic medium (IGM) before cosmic reionization is complete (see reviews by, e.g. Furlanetto, Oh & Briggs 2006; Morales & Wyithe 2010; Pritchard & Loeb 2012; Liu & Shaw 2020). Because ionized regions are transparent at redshifted 21-cm wavelengths, maps of the 21-cm background during reionization will consist of 'holes' in regions where there are many galaxies, and a mostly neutral 'bulk' IGM beyond with a 21-cm signal that traces the gas density, temperature, and Lyman α background intensity. While the ultimate goal is to map this patchy structure in detail with, e.g. the Square Kilometer Array (SKA), the current generation of interferometers are seeking a statistical detection of the 21-cm power spectrum (LOFAR, MWA, HERA, GMRT, LWA; van Haarlem et al. 2013; Tingay et al. 2013; DeBoer et al. 2017; Paciga et al. 2013; Eastwood et al. 2019), while a complementary suite of current and planned experiments (Bowman & Rogers 2010; Burns et al. 2017; Singh et al. 2017; de Lera Acedo 2019; Philip et al. 2019) are targeting the sky-averaged 'global' 21-cm signal (Shaver et al. 1999), which traces the average properties of the IGM as a function of redshift.

In the last few years, several experiments have reported upper limits on the power spectrum of 21-cm fluctuations during reionization (Parsons et al. 2014; Patil et al. 2017; Barry et al. 2019; Mertens et al. 2020; The HERA Collaboration 2021b) and the earlier cosmic-dawn era (Eastwood et al. 2019; Gehlot et al. 2019, 2020; Garsden et al. 2021; Yoshiura et al. 2021). Scenarios in which the bulk IGM is still colder than the cosmic microwave background (CMB) during reionization give rise to the strongest fluctuations and so will be the first models to be tested as upper limits continue to improve (e.g. Parsons et al. 2014; Pober et al. 2015; Greig, Mesinger & Pober 2016). Similarly, stronger-than-expected 21-cm signals can arise if the cosmic radio background has contributions other than the CMB (Feng & Holder 2018), e.g. synchrotron emission from accreting black holes (Ewall-Wice et al. 2018), star-forming galaxies (Mirocha & Furlanetto 2019), or from decaying particles (Fraser et al. 2018; Pospelov et al. 2018). Indeed, constraints from MWA, HERA, and LoFAR disfavour models with negligible X-ray heating at $z \sim 8$–9 or very strong radio backgrounds (Ghara et al. 2020, 2021; Mondal et al. 2020; Greig et al. 2021a, b; The HERA Collaboration 2021a). Of course, the recent report of an absorption signal in the sky-averaged spectrum at $z \sim 17$ from EDGES (Bowman et al. 2018) requires an even colder IGM (Barkana 2018; Boddy et al. 2018; Fialkov, Barkana & Cohen 2018; Kovetz et al. 2018; Muñoz & Loeb 2018) or a brighter background (Ewall-Wice et al. 2018; Feng & Holder 2018; Fialkov & Barkana 2019; Mirocha &

★ E-mail: jordan.mirocha@mcgill.ca

© 2022 The Author(s)
Published by Oxford University Press on behalf of Royal Astronomical Society

Downloaded from https://academic.oup.com/mnras/article/514/2/2010/6595978 by Scuola Normale Superiore. Biblioteca user on 03 December 2022

Furlanetto 2019) than models in ΛCDM cosmologies generally predict. However, the most stringent power spectrum upper limits from The HERA Collaboration (2021b) are derived at sufficiently low redshifts relative to EDGES ($z \lesssim 10$ versus $z \simeq 18$) that they cannot yet directly address the EDGES controversy (Hills et al. 2018; Bradley et al. 2019; Singh & Subrahmanyan 2019; Sims & Pober 2020; Tauscher, Rapetti & Burns 2020; Singh et al. 2021).

The theoretical interpretations of 21-cm measurements have thus far been guided mostly by semi-numerical models of reionization (Santos et al. 2010; Mesinger, Furlanetto & Cen 2011; Fialkov et al. 2014; Hutter 2018) and other approximate techniques (Thomas et al. 2009; Ghara, Choudhury & Datta 2015) designed to avoid more accurate, but expensive, radiative transfer simulations (see e.g. Gnedin 2014; Rosdahl et al. 2018; Ocvirk et al. 2020; Kannan et al. 2022). Though several inference frameworks have emerged to jointly fit 21 cm and other constraints on reionization (e.g. Mirocha, Harker & Burns 2015; Greig et al. 2016; Ghara et al. 2018; Mondal et al. 2020), current power spectrum limits are still quite weak, and MCMC fits have thus yet to provide a strong constraint on any individual astrophysical parameter, with the exception of the ratio of X-ray luminosity to SFR in high-z galaxies (The HERA Collaboration 2021a), and instead reveal the ∼2–3 dimensional corners of parameter space that are most strongly disfavoured by the data. Of course, as upper limits improve and eventually become detections, constraints on astrophysical and cosmological parameters are expected to be exquisite (e.g. McQuinn et al. 2006; Greig & Mesinger 2015; Ewall-Wice et al. 2016; Liu et al. 2016; Liu & Parsons 2016; Kern et al. 2017; Muñoz 2019b). In the meantime, however, it may be prudent to focus also on simpler derived quantities, e.g. the mean neutral fraction and temperature of the IGM (or ratio of background temperature to 21-cm spin temperature, T_R/T_S), which are more directly probed by 21-cm experiments, and may thus be easier to constrain.

In this paper, we present a simple, phenomenological model for 21-cm fluctuations during reionization that abstracts away galaxies and instead works directly in terms of the mean properties of the IGM and the size distribution of ionized regions. The goal is first and foremost to build intuition for the results of more sophisticated semi-numerical models in use today. In addition, our phenomenological approach lets us break key assumptions built-in to physical models, and may thus help gauge the extent to which constraints on the IGM derived from 21-cm power spectra are model dependent. Our formalism is similar to that of Furlanetto, Zaldarriaga & Hernquist (2004) and extensions (Paranjape & Choudhury 2014; Paranjape, Choudhury & Padmanabhan 2016), though we do not attempt to model the size distribution of ionized regions using excursion set arguments. Instead, we parametrize it directly, which offers more flexibility than physical models. Our approach is similar in spirit to other recent efforts aimed at building intuition for, and providing a cross-check of, more detailed numerical simulations of reionization (e.g. Kaurov 2016; McQuinn & D'Aloisio 2018; Raste & Sethi 2018; Schneider, Giri & Mirocha 2021), but conceptually simpler than each. It is also complementary to efforts to constrain the IGM while abstracting away astrophysical parameters as much as possible (e.g. Mirocha, Harker & Burns 2013; Cohen et al. 2017; Mason et al. 2019).

In order to gain ground analytically, we will assume fully ionized and spherical bubbles with infinitely crisp edges, whose abundance is well described by a bubble size distribution (BSD) function. In reality, the ionization field is more complicated; as bubbles merge with neighbours their morphologies become complex, resulting in an interconnected network of ionized regions when the global ionized fraction is just ∼10 per cent (Furlanetto & Oh 2016). Partial ionization due to small-scale clumping – which we neglect here – can affect the BSD and 21-cm power spectrum (e.g. Sobacchi & Mesinger 2014; Bianco et al. 2021), as could a strong X-ray background, if sources with soft spectra dominate. Despite these known shortcomings of the approach, we forge ahead none the less, in order to thoroughly assess the accuracy of analytical models and the prospects for using them to derive meaningful constraints on the high-z IGM. To our knowledge, there has yet to be such an attempt to push any analytical model of 21-cm fluctuations through to parameter inference, and compare its results to those of a semi-numerical model (though see e.g. Santos et al. 2008; Schneider et al. 2021, for more general comparisons of analytical and seminumeric models).

The structure of the paper is as follows. In Section 2, we introduce our phenomenological approach to 21-cm fluctuations and present its basic predictions. Then, in Section 3, we compare various components of the phenomenological model to two illustrative 21cmFAST models, in order to gauge its accuracy and motivate different modelling choices. We present a forecast in Section 4, conducted by fitting our model to mock signals created both by the phenomenological model itself as well as 21cmFAST. We conclude in Section 5.

2 PHENOMENOLOGICAL MODELLING FRAMEWORK

In this paper, we attempt to remain as agnostic as possible about the source of 21-cm fluctuations at high redshifts. We assume only that the 21-cm field is composed of discrete bubbles embedded in a medium of uniform temperature. We make no effort to model the size distribution of these bubbles via forward modelling, nor do we attempt to evolve the properties of the 'bulk IGM' beyond bubbles. Instead, we absorb all the astrophysics of reionization and reheating into the size distribution of bubbles, the volume of space they occupy, their mean density, and the mean temperature of the bulk IGM. This section describes the core components of the model, which is implemented in the publicly available MICRO21CM package.[1] The key predictions of the model are summarized in Figs 1, 2, and 3, with comparisons to 21cmFAST and forecasts to follow in Sections 3 and 4.

2.1 Preliminaries

The brightness temperature of the 21-cm field at an arbitrary location in space is given by (see e.g. Madau et al. 1997; Furlanetto et al. 2006)

$$\boldsymbol{\delta T_b} \simeq 27\ \mathrm{mK} \left(\frac{\Omega_{\mathrm{b},0}h^2}{0.023}\right)\left(\frac{0.15}{\Omega_{\mathrm{m},0}h^2}\frac{1+z}{10}\right)^{1/2} \times \boldsymbol{x}_{\mathrm{H\,I}}(1+\boldsymbol{\delta})\left(1-\frac{T_{\mathrm{R}}}{\boldsymbol{T}_{\mathrm{S}}}\right)\left(\frac{H}{\partial_r v_r}\right), \quad (1)$$

where $x_{\mathrm{H\,I}} = 1 - x_{\mathrm{H\,II}}$ is the neutral hydrogen fraction, δ is the baryon density relative to the cosmic mean, $\partial_r v_r$ is the line-of-sight gradient of the velocity, T_{R} is the temperature of the background (assumed here to be spatially uniform), generally assumed to be the cosmic microwave background T_γ, and T_{S} is the spin temperature, which quantifies the relative abundance of hydrogen atoms in the ground hyperfine triplet and singlet states. Each of these quantities carries an implicit redshift dependence, while bolded quantities are those that vary spatially as well, e.g. $\boldsymbol{\delta T_b} = \boldsymbol{\delta T_b}(\boldsymbol{x})$, which we will discuss in detail momentarily. In general, T_{S} depends the hydrogen and

[1] https://github.com/mirochaj/micro21cm

Downloaded from https://academic.oup.com/mnras/article/514/2/2010/6595978 by Scuola Normale Superiore. Biblioteca user on 03 December 2022

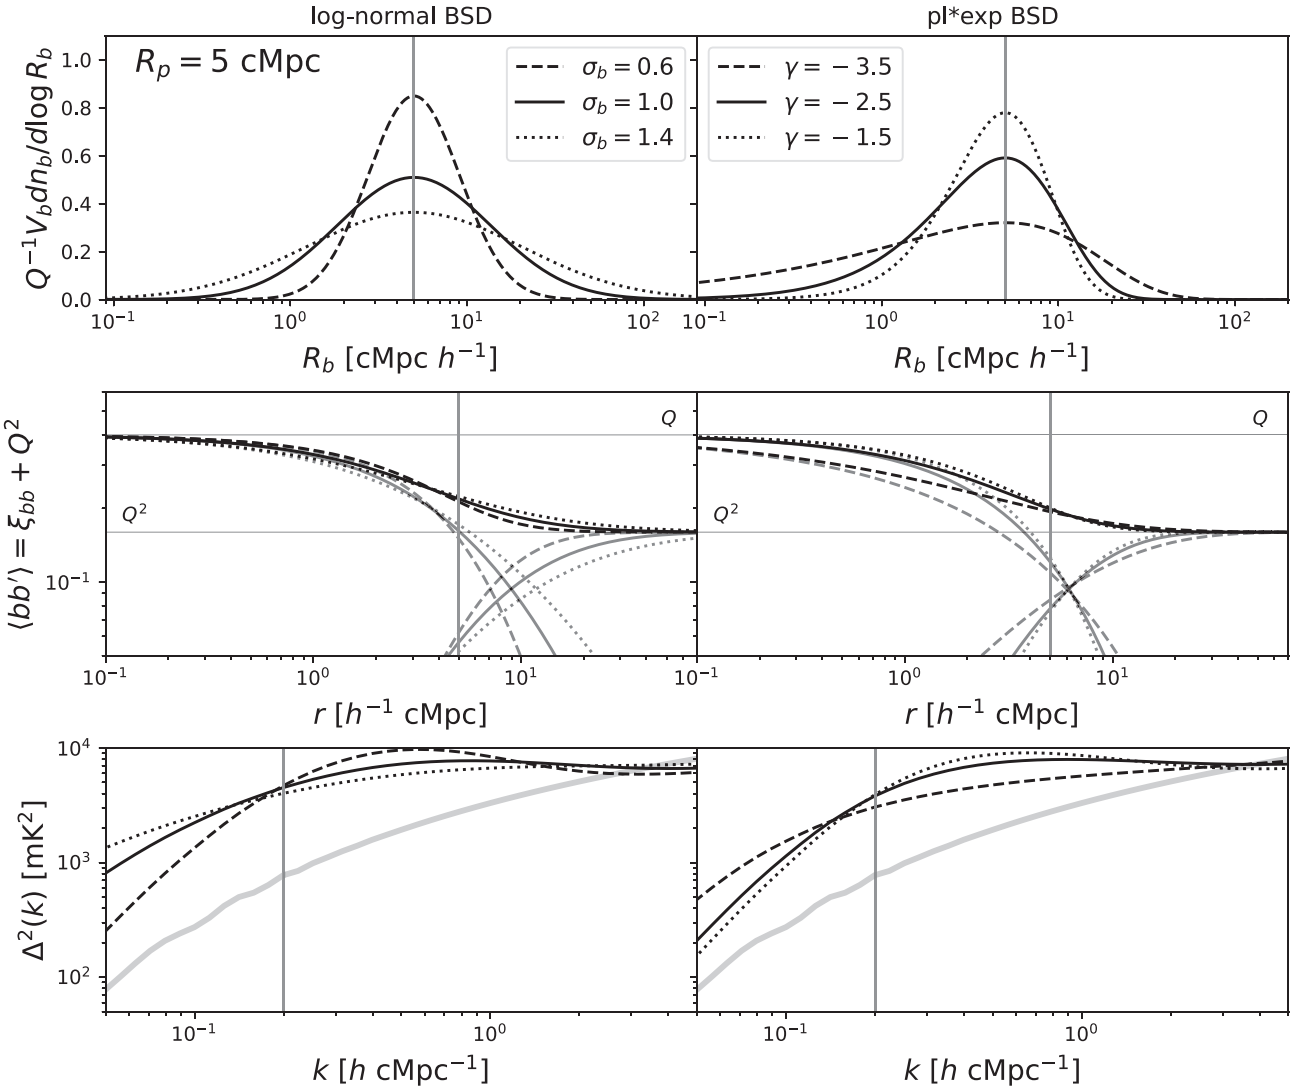

Figure 1. Connection between bubble size distribution (top), bubble correlation function (middle), and 21-cm power spectrum (bottom), both for lognormal BSDs (left) and BSDs modelled as a power law with an exponential cut-off (right). Each model assumes the same mean ionized fraction, $\mathcal{Q} = 0.2$, and adopts $T_S = 1.8$ K, as is appropriate for a uniform, adiabatically cooled $z = 8$ IGM. Vertical grey bars in the top and middle rows indicate the location of the peak in the volume-weighted BSD, while in the bottom row we instead show the scale $k_{\rm peak} = R_{\rm peak}^{-1}$. Horizontal lines in the middle row indicate the limiting behaviour of the one- and two-bubble terms, which are shown individually in grey. We describe the many terms contributing to 21-cm fluctuations in Section 2 and Fig. 2.

electron densities, gas kinetic temperature, and Lyman α background intensity, though in this work we abstract away all the physics embedded in T_S and simply treat it as a homogeneous free parameter. We take the same approach to T_R, though in general, e.g. radio emission from galaxies may drive non-trivial fluctuations in T_R as well, and leave interesting signatures in the 21-cm background in some scenarios (Reis, Fialkov & Barkana 2020).

Given our assumption of a field composed of bubbles, it will be convenient in what follows to rewrite equation (1) as

$$\boldsymbol{\delta T_b} = T_0(z)(1+\boldsymbol{\delta})(1+\alpha \boldsymbol{b})(1+\boldsymbol{\delta_\varphi}), \tag{2}$$

where T_0 is a redshift- and cosmology-dependent normalization,

$$T_0 \equiv 27 \left(\frac{\Omega_{\rm b,0}h^2}{0.023}\right)\left(\frac{0.15}{\Omega_{\rm m,0}h^2}\frac{1+z}{10}\right)^{1/2}\left(1-\frac{T_{\rm R}}{T_{\rm S}}\right)\ {\rm mK} \tag{3}$$

and δ_φ is a fractional perturbation in the temperature 'contrast,'

$$\boldsymbol{\varphi} = \frac{\boldsymbol{T}_{\rm S} - T_{\rm R}}{\boldsymbol{T}_{\rm S}}, \tag{4}$$

which is related to fractional perturbations in the spin temperature via

$$\boldsymbol{\delta_\varphi} = \left(\frac{T_{\rm R}}{T_{\rm S}-T_{\rm R}}\right)\boldsymbol{\delta}_{T_S} = \varphi^{-1}\left(\frac{T_{\rm R}}{T_{\rm S}}\right)\boldsymbol{\delta}_{T_S}. \tag{5}$$

The variable $\boldsymbol{b}$ in equation (2) represents a binary field of bubbles, and thus takes on values of 0 or 1 only, with 1 indicative of a fully ionized or fully heated bubble. In this work, we will focus entirely on ionized bubbles, noting here the possibility of heated bubbles for completeness – one can easily switch from one to the other with a

Downloaded from https://academic.oup.com/mnras/article/514/2/2010/6595978 by Scuola Normale Superiore. Biblioteca user on 03 December 2022

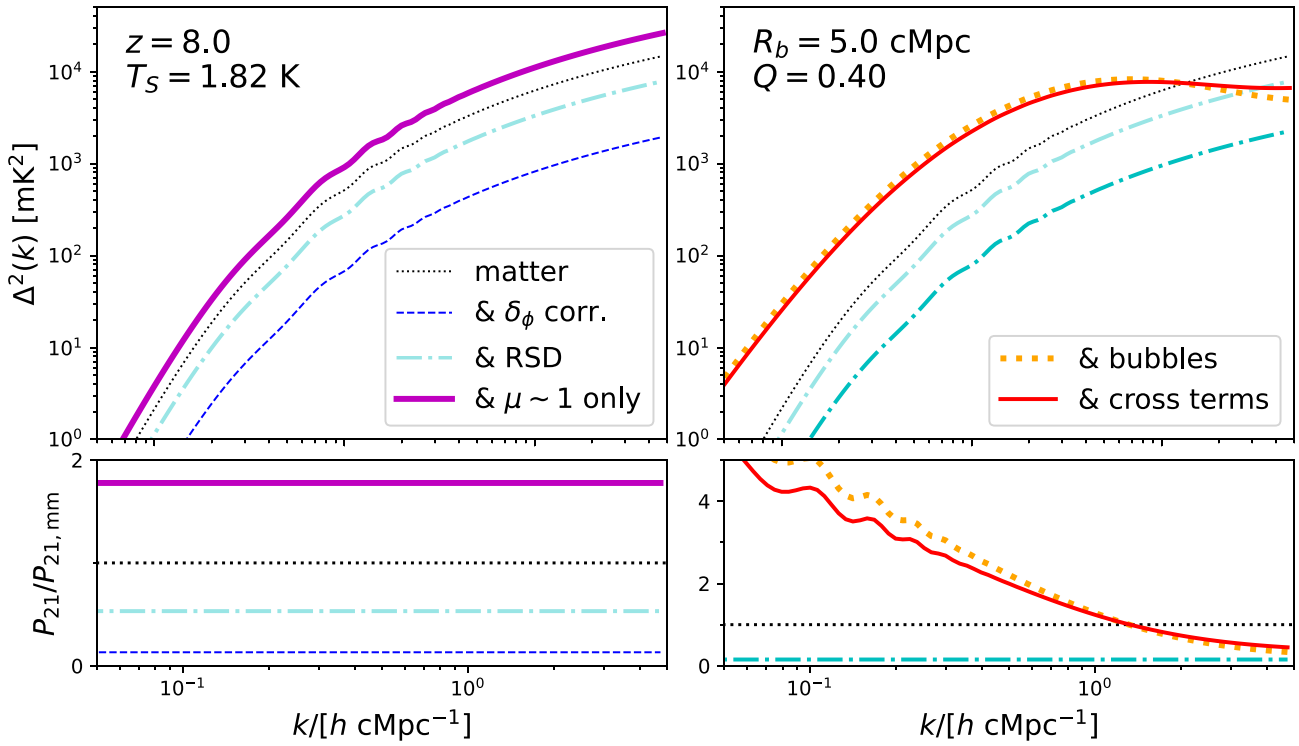

Figure 2. Main contributions to the 21-cm power spectrum. Starting from a universe with only linear density fluctuations (dotted black; left column), we add corrections for adiabatic compression/expansion (dotted blue), and redshift-space distortions, with spherical averaging (cyan) as well as pure line-of-sight modes (magenta). Then, in the right-hand column, we add the autocorrelation contributions from bubbles ($\langle bb' \rangle$; dotted orange) and the cross-correlation terms involving bubbles and the density field (solid red). The bottom row in each column shows the relative amplitude of the 21-cm power spectrum relative to a model with matter fluctuations only, which we indicate as $P_{21,\mathrm{mm}}$.

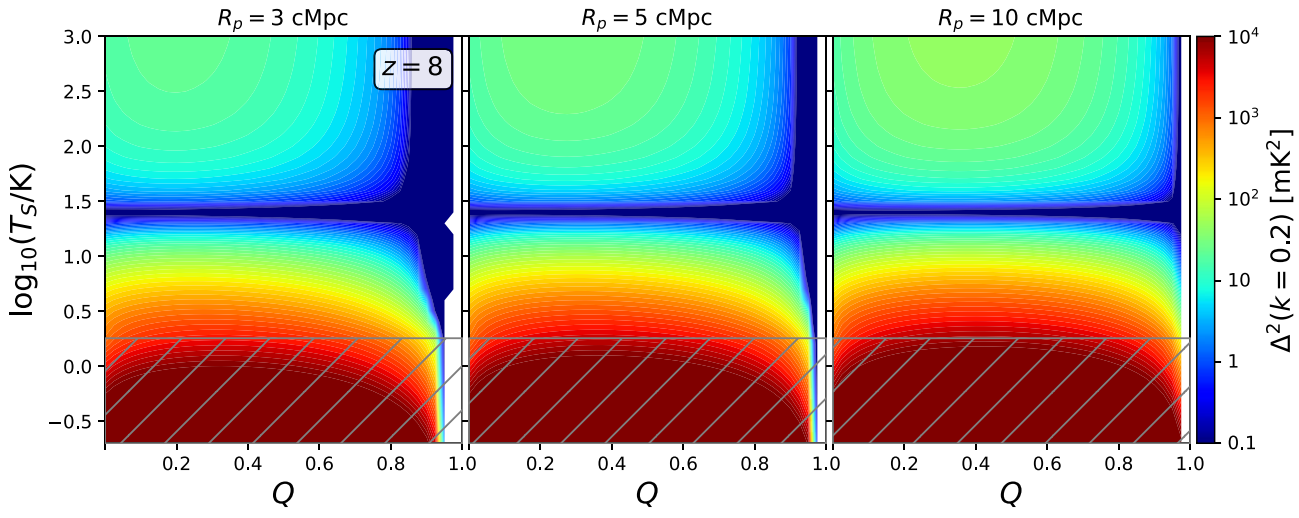

Figure 3. 2D parameter study showing the effects of ionized fraction Q and spin temperature T_S for lognormal bubble size distributions. The colour-scale indicates the dimensionless power, Δ^2_{21}, at $z = 8$ and $k = 0.2\,h\,\mathrm{Mpc}^{-1}$. We assume $\sigma_b = 1$ for characteristic bubble sizes of 1, 5, and 10 cMpc (left to right). The horizontal band in each panel corresponds to $T_\mathrm{S} = T_\gamma$ at $z = 8$, where the brightness temperature goes to zero, while the cross-hatched region in the bottom indicates spin temperatures below the adiabatic limit for a uniform medium at $z = 8$, $T_\mathrm{adi} = 1.82$ K.

Downloaded from https://academic.oup.com/mnras/article/514/2/2010/6595978 by Scuola Normale Superiore. Biblioteca user on 03 December 2022

suitable choice of α, i.e.

$$\alpha = \begin{cases} 0 & \text{no bubbles} \\ -1 & \text{ionized bubbles} \\ T_{\rm R}/(T_{\rm S} - T_{\rm R}) & \text{heated bubbles} \end{cases} \quad (6)$$

This expression assumes that heated 'bubbles' are fully saturated, $T_{\rm S} \gg T_{\rm R}$, in which case plugging $\alpha = T_{\rm R}/(T_{\rm S} - T_{\rm R})$ into equation (2) reduces to the brightness temperature of a saturated patch of the IGM. One could alternatively leave the temperature of heated bubbles as a free parameter.

From this point onward, we will discontinue the use of bold-faced variables, meaning any occurrence of $T_{\rm S}$ or φ refers to the mean, while δ's and b's carry all spatial information.

If we assume for the moment that the spin temperature field is spatially homogeneous, leaving $\delta_\varphi = 0$, we can write the correlation function of the 21-cm field, $\xi_{21} \equiv \langle \delta T_b \delta T_b' \rangle - \langle \delta T_b \rangle^2$, relatively compactly as

$$\begin{aligned} T_0^{-2}\xi_{21} = &\langle \delta\delta' \rangle + \alpha^2 \langle bb' \rangle - \alpha^2 \langle b \rangle^2 \\ &+ 2\alpha \langle b\delta' \rangle + 2\alpha^2 \langle bb'\delta \rangle + 2\alpha \langle b\delta\delta' \rangle + \alpha^2 \langle bb'\delta\delta' \rangle \\ &- 2\alpha^2 \langle b \rangle \langle b\delta \rangle - \alpha^2 \langle b\delta \rangle^2. \end{aligned} \quad (7)$$

Here, angular brackets indicate ensemble averages and primed quantities indicate points a distance r from unprimed points. So far, our results are exact, as only terms proportional to $\langle \delta \rangle$ have been dropped (the overdensity field has mean zero by definition). The ensemble average of the bubble field is equivalent to the volume-filling fraction in this framework, so in what follows we will use $Q \equiv \langle b \rangle$.

In the next few sections, we will discuss methods for modelling the various terms in equation (7), as well as 'correction terms' that arise when $\delta_\varphi > 0$ and peculiar velocities (i.e. redshift-space distortions) are included. We will adopt the following convention for correction factors that relate an arbitrary fluctuation in δ_X to the density field, i.e.

$$\delta_X \equiv \beta_X \delta. \quad (8)$$

In all that follows, we will plot only the dimensionless power spectrum

$$\Delta^2_{21}(k) = k^3 P_{21}(k)/2\pi^2, \quad (9)$$

where P_{21} is related to the 21-cm correlation function by the transform[2]

$$P_{21}(k) = \frac{1}{(2\pi)^3} \int 4\pi R^2 \frac{\sin(kR)}{kR} \xi_{21}(R) {\rm d}R. \quad (10)$$

Note that it is common in the literature to plot $\overline{\delta T_b}^2 \Delta^2_{21}(k)$, i.e. the global 21-cm signal squared times Δ^2_{21}. The difference is one of definitions and notation. If one works in terms of the fractional perturbation in the 21-cm brightness temperature, such that $\delta \mathbf{T_b}(\mathbf{x}) = \overline{\delta T_b}(1 + \delta_{21})$, and defines the 21-cm correlation function as $\xi_{21} \equiv \langle \delta_{21}\delta_{21}' \rangle - \langle \delta_{21} \rangle^2$, then the appropriate quantity to plot is indeed $\overline{\delta T_b}^2 \Delta^2_{21}(k)$. However, we work directly in terms of the 21-cm signal's constituent terms, i.e. δ, b, δ_ϕ, etc., which leaves a normalization factor of T_0^2 only, where $T_0 = \overline{\delta T_b}(1 - Q)^{-1}$ (see also equation 3). We absorb this factor of T_0^2 into the 21-cm correlation function, as in equation (7), which implicitly lends equations (9) and (10) units of mK^2. Hence, the absence of $\overline{\delta T_b}^2$ as a multiplicative pre-factor applied to $\Delta 21^2(k)$ in the y labels of our figures.

[2]We discuss numerical solutions to this integral in Section 2.6.

2.2 Adiabatic expansion and compression

The first correction we consider addresses the correlation between the density and temperature due to adiabatic expansion or compression alone (i.e. no X-ray heating from astrophysical sources). For clarity, we allow $\delta_\varphi > 0$ but hold $b = 0$, in which case the correlation function of the 21-cm background can be written as

$$\begin{aligned} T_0^{-2}\xi_{21,b=0} = &\langle \delta\delta' \rangle + \langle \delta_\varphi \delta_\varphi' \rangle \\ &+ 2\langle \delta_\varphi \delta' \rangle + 2\langle \delta_\varphi \delta_\varphi' \delta \rangle + 2\langle \delta_\varphi \delta\delta' \rangle + \langle \delta_\varphi \delta_\varphi' \delta\delta' \rangle \\ &- 2\langle \delta_\varphi \rangle \langle \delta_\varphi \delta \rangle - \langle \delta_\varphi \delta \rangle^2. \end{aligned} \quad (11)$$

We further assume that density and temperature fluctuations are small, which allows us to eliminate all 3- and 4-pt terms from the above expression, and note that the ensemble average of δ_φ is zero by construction and can thus be removed as well.

Perturbations in the density, temperature, and ionization can be evolved numerically (e.g. Barkana & Loeb 2005b; Naoz & Barkana 2005; Pritchard & Furlanetto 2007). However, Muñoz, Ali-Haïmoud & Kamionkowski (2015) found that fluctuations in the kinetic temperature can be more straightforwardly related to δ from recombination and throughout the cosmic dark ages up to whenever inhomogeneous X-ray heating occurs, via

$$\delta_T \equiv \frac{\delta T_{\rm K}}{T_{\rm K}} = \beta_T(z)\delta. \quad (12)$$

The coefficient β_T, usually termed the adiabatic index, would be exactly 2/3 for pure adiabatic cooling, though the Compton scattering with the CMB produces a deviation from this factor. For the redshift range of interest, the result in Muñoz et al. (2015) can be approximated as $\beta_T(z) = c_0 - c_1(z - 10)$, with $c_0 = 0.58$ and $c_1 = 0.005$ to within 3 per cent precision (for $z = 6$–50).

Using this linear relationship between δ and δ_T, and assuming saturated Wouthuysen–Field coupling (so that $Ts \approx T_{\rm K}$.), we have

$$\begin{aligned} T_0^{-2}\xi_{21,b=0} &= \langle \delta\delta' \rangle + \beta_\varphi^2 \langle \delta\delta' \rangle + 2\beta_\varphi \langle \delta\delta' \rangle \\ &= \langle \delta\delta' \rangle \left[1 + \beta_\varphi\right]^2, \end{aligned} \quad (13)$$

where we have made use of equation (5) and defined $\beta_\varphi \equiv \beta_T \varphi^{-1}\left(\frac{T_{\rm R}}{T_{\rm S}}\right)$.

The overall effect of adiabatic expansion and compression is to reduce the amplitude of the 21-cm emission in a given patch of the IGM relative to the uniform density case. To see this, we now write equation (1) with a correction term $\beta_\delta(z) \equiv 1 + \beta_\varphi$ applied to the density, i.e.

$$\delta T_b \simeq \overline{T_0}(1 + \beta_\delta \delta), \quad (14)$$

which yields

$$\delta T_b \simeq \frac{\overline{T_0}}{1 - T_\gamma/T_{\rm S}} \left[1 - \frac{T_\gamma}{T_S}(1 - \beta_T)\right] \delta. \quad (15)$$

As expected, $\beta_T > 0$ reduces the amplitude of the signal, holding all other quantities fixed.

For scenarios where the gas has a temperature above the adiabatic prediction, we ought to account for how heating changes this picture. We can always write that at every point and z,

$$T_g = T_g^{\rm ad} + \Delta T_g^X, \quad (16)$$

where $T_g^{\rm ad}$ is the adiabatic prediction, and ΔT_g^X is the heating term, which we assume to be approximately homogeneous, as predicted

Downloaded from https://academic.oup.com/mnras/article/514/2/2010/6595978 by Scuola Normale Superiore. Biblioteca user on 03 December 2022

for instance for hard X-rays. Therefore, $\delta T_g = \delta T_g^{\rm ad} = \beta_T \delta$, and thus

$$\delta_T = \delta_T^{\rm ad} \frac{T_g^{\rm ad}}{T_g} = \beta_T \frac{T_g^{\rm ad}}{T_g} \delta, \tag{17}$$

or equivalently, we can use our result in equation (14) with the correction

$$\beta_T \rightarrow \beta_T \min\left(1, \frac{T_g^{\rm ad}}{T_g}\right). \tag{18}$$

We note that for $T_g < T_g^{\rm ad}$, as predicted by models of DM-induced cooling (see e.g. Muñoz, Kovetz & Ali-Haïmoud 2015; Barkana 2018; Berlin et al. 2018; Muñoz, Dvorkin & Loeb 2018), our formula predicts an increase in the β_T term. To remain conservative, we cap β_T at its adiabatic value of 2/3.

2.3 Statistics of bubbles

We now turn our attention to 21-cm fluctuations sourced by fluctuations in the ionization field, starting with the autocorrelation term $\langle bb' \rangle$ (see equation 7). Much of this follows from Furlanetto et al. (2004) (hereafter FZH04), but we review it here none the less for completeness. Given that b is binary, the ensemble average greatly simplifies

$$\langle bb' \rangle \equiv \int {\rm d}b \int {\rm d}b' bb' f(b, b') = \mathcal{P}_{bb} \tag{19}$$

since the integrand is only non-zero when both b and b' are unity. In this case, the double integral over the joint distribution $f(b', b')$ is simply the probability that two points are both in bubbles, hence our use of the notation $\mathcal{P}_{bb}$. This probability will be determined entirely by the BSD, which we parametrize flexibly rather than model from physical arguments (see Furlanetto et al. 2004; Paranjape & Choudhury 2014, for excursion set models). Several example BSDs are shown in Fig. 1 and discussed in more detail in Section 2.3.2.

For a field composed of discrete bubbles, $\mathcal{P}_{bb}$ can be worked out for an assumed size distribution of bubbles, $n_{\rm b}(R_{\rm b})$, following FZH04. Drawing inspiration from the halo model (Cooray & Sheth 2002), only two configurations are possible: either two points are in the same bubble or they are in different bubbles, i.e. $\mathcal{P}_{bb} = \mathcal{P}_1 + \mathcal{P}_2$.

The probability $\mathcal{P}_1$ that two points are in the same bubble depends on the fraction of the volume in which a single bubble can enclose two points separated by a distance r. This amounts to an integral over the bubble size distribution weighted by the 'overlap volume,' V_o,

$$\mathcal{P}_1'(r) = \int {\rm d}R_b n_b(R_b) V_o(r, R_b). \tag{20}$$

The overlap volume is the intersectional volume of two spheres of radius R_b separated by r,

$$V_o = \begin{cases} \frac{4\pi}{3} R_b^3 - \pi r \left[R_b^2 - r^2/12\right] & r < 2R_b \\ 0 & \text{otherwise} \end{cases}. \tag{21}$$

A single bubble can only engulf two points separated by r if it is centred within this region. Note that the probabilities in equation (20) (as well as equation 22 below) are primed to indicate that they are not the final probabilities used in the model, as corrections are in order (see Section 2.3.1).

The other possibility is that two points reside in different bubbles. In this case, we need the probability that a single source can ionize one point but not the other, which is proportional to $V(R_b) - V_o$, with $V(R_b) = 4\pi R_b^3/3$:

$$\mathcal{P}_2'(r) = (1 - \mathcal{P}_1') \times \left(\int {\rm d}R_b n_b(R_b) [V(R_b) - V_o(r, R_b)] \right)^2. \tag{22}$$

The leading factor of $(1 - \mathcal{P}_1')$ ensures that the two points under consideration do not reside in a single bubble.

The one- and two-bubble terms are shown individually in the middle row of Fig. 1; the former dominates on small scales, and asymptotes to a value of Q, while the latter dominates on large scales, and asymptotes to Q^2. On intermediate scales comparable to the typical bubble size, both terms are comparable. The smoothness of the transition from the one-bubble regime to the two-bubble regime is governed by the details of the BSD, which we discuss more in Section 2.3.2.

We currently neglect the clustering of bubbles – a more sophisticated approach may be warranted at early times and/or intermediate scales where clustering of bubbles is important, in which case a factor of $1 + \epsilon$ within one of the integrands of equation (22) would be necessary, indicating an excess probability that a second source ionizes point 2 given that the first point is ionized (see e.g. section 3.3 in FZH04).

2.3.1 *Bubble overlap*

Perhaps the most obvious shortcoming of the treatment so far is that it neglects the potential for overlap between bubbles. In reality, overlapping bubbles cease to be distinct entities, and will instead form a single larger bubble – indeed, (Furlanetto & Oh 2016) showed that throughout most of reionization the vast majority of the ionized volume is contained in a single 'percolating cluster' with a very complex shape (although simulations do suggest that the percolating cluster is composed of subunits with a finite size; e.g. Lin et al. 2016; Busch et al. 2020). The percolation process qualitatively changes the meaning of the BSD, so that we cannot treat these stages self-consistently. However, motivated by the existence of a characteristic scale in simulations of reionization, we can modify our probabilities slightly to account for overlap in a statistical sense.

To assess the importance of overlap, it is useful to first consider the total volume contained in bubbles

$$Q_{\rm tot} = \int {\rm d}R_b n_b(R_b) V(R_b). \tag{23}$$

This quantity is *not* equivalent to the volume-filling fraction of ionized gas, $Q_{\rm tot} \neq Q$, since there is nothing stopping two (or more) bubbles from co-occupying the same space in our model. It is more accurate to consider Q as the probability that a single point is ionized, which is a sum over all possible configurations, e.g. that a point is engulfed by a single bubble of radius R_1, or that a point is instead engulfed by a bubble of radius R_2 (but not R_1), etc., i.e.

$$\mathcal{P}_b = \mathcal{P}_1 + (1 - \mathcal{P}_2) + (1 - \mathcal{P}_1)(1 - \mathcal{P}_2)\mathcal{P}_3 + \ldots \tag{24}$$

Note that each $\mathcal{P}_i$ term in the above sum requires an integral, and neglects the possibility that a single point resides within multiple bubbles.

One can dramatically simplify this computation by realizing that we only care if a point is ionized, regardless of how many bubbles contain it. The ionized fraction is simply the complement of the probability P_0 that *no* bubbles engulf a point, which we can write as a Poisson distribution (see also sections 3.1–3.2 in FZH04),

$$Q = 1 - \exp\left[- \int {\rm d}R_b n_b(R_b) V(R_b) \right]. \tag{25}$$

Downloaded from https://academic.oup.com/mnras/article/514/2/2010/6595978 by Scuola Normale Superiore. Biblioteca user on 03 December 2022

The same logic applies to the calculation of one- and two-bubble terms. As a result, we take our final, unprimed probabilities to be

$$\mathcal{P}_1(r) = 1 - \exp\left[-\int \mathrm{d}R_b n_b(R_b) V_o(r, R_b)\right] \tag{26}$$

and

$$\mathcal{P}_2(r) = (1-\mathcal{P}_1) \times \left(1 - \exp\left[-\int \mathrm{d}R_b n_b(R_b)\left(V(R_b) - V_o(r, R_b)\right)\right]\right)^2. \tag{27}$$

To gauge the importance of overlap, we can easily compute the difference of $Q_{\rm tot}$ and Q, i.e. the fraction of the volume composed of *more than one* bubble. We find that this 'global inter-sectional volume'

$$Q_{\rm int} \equiv Q_{\rm tot} - Q \tag{28}$$

is $\ll 1$ when Q is small, indicating that overlap is unimportant at early times, as expected. However, $Q_{\rm int}$ rises as Q grows, reaching a value of $Q_{\rm int} \simeq 0.2$ when $Q = 0.5$, and $Q_{\rm int} \simeq Q$ when $Q \simeq 0.8$, i.e. overlap is likely an order unity effect for the last $\simeq$20 per cent of reionization. Note that for this calculation we normalize n_b to satisfy a user-supplied value of Q using equation (25). Note also that $Q_{\rm tot}$ and $Q_{\rm int}$ can exceed unity, because these volumes include every occurrence of overlap at a given point, e.g. a single region will contribute three times its volume if three separate bubbles contain it.

2.3.2 Bubble size distributions

We have yet to explicitly define the BSD, instead leaving it as a generic function $n_b(R_b)$. Rather than model the BSD from a galaxy formation model and the excursion set approach, we parametrize it flexibly.[3] Our fiducial model adopts a lognormal form

$$\frac{\mathrm{d}n_b}{\mathrm{d}R_b} \propto (R_c\sigma)^{-1} \exp\left[-\frac{(R_b - R_c)^2}{2\sigma_b^2}\right], \tag{29}$$

peaked near a radius R_c, and with a width σ_b. We also consider a power-law times an exponential

$$\frac{\mathrm{d}n_b}{\mathrm{d}R_b} \propto \left(\frac{R_b}{R_c}\right)^{\gamma} \exp\left[-R_b/R_c\right], \tag{30}$$

with the index γ as a free parameter rather than σ. In each case, the BSD is normalized to preserve the mean ionized fraction defined in equation (25).

Note that the more relevant quantity in our analysis is actually the volume-weighted BSD, $V\mathrm{d}n_b/\mathrm{dlog}\,R_b$, where $V = 4\pi R_b^3/3$, since the probability that two points are ionized is related to volume more directly than bubble size. The peak in the volume-weighted BSD is *not* equal to the peak in $\mathrm{d}n_b/\mathrm{d}R_b$. As a result, in all that follows, we use the variable R_p or $R_{\rm peak}$ to indicate the location of the peak in $V\mathrm{d}n_b/\mathrm{dlog}\,R_b$. For both BSDs we explore, it is easy to translate R_c to R_p via

$$R_p = R_c \exp\{3\sigma^2\} \tag{31}$$

for the lognormal, and

$$R_p = R_c(\gamma + 4) \tag{32}$$

for the power-law-times-exponential model.

[3]The parametric forms we choose are known to mimic physically motivated models reasonably well (see e.g. Kakiichi et al. 2017)

2.3.3 Limiting behaviour checks

Before moving on to cross-terms, let us examine the limiting behaviour of the correlation function of the bubble field, $\xi_{bb} \equiv \langle bb' \rangle - Q^2$. We have only two obvious requirements: (i) fluctuations must vanish on large scales $r \gg R_b$, and (ii) fluctuations must vanish at $Q = 0$ and $Q = 1$. First, for an arbitrary value of Q, on large scales ($r \gg R_b$) we find $V_o \to 0$, so $\xi_{bb} \to \mathcal{P}_2 - Q^2 \simeq 0$, and fluctuations vanish as they must, since $\mathcal{P}_2$ tends towards Q^2 on large scales. Secondly, the $Q = 0$ limit is satisfied by construction in our framework, since the BSD is normalized by Q, thus forcing $\mathcal{P}_1 = \mathcal{P}_2 = 0$ when $Q = 0$. This leaves only the question of whether or not ξ_{bb} vanishes at $Q = 1$.

On large scales, fluctuations vanish regardless of Q, as shown above. On small scales, $r \ll R_b$, the overlap volume tends to the bubble volume, $V_o(r, R_b) \to V(R_b)$. As a result, $\mathcal{P}_1 \to Q$, and $\mathcal{P}_2 \to 0$. This leaves a bubble correlation function $\xi_{bb} \simeq Q - Q^2 = Q(1 - Q)$, which is indeed zero when $Q = 1$.

2.4 Cross-correlations between ionization and density

We have thus far neglected correlations between the density field and bubbles (see second and third rows of equation 7). In keeping with the phenomenological spirit of this paper, we take a simple approach to these terms that abstracts away assumptions about the sources (to the extent that this is possible). The busy reader may skip ahead to Fig. 2, which shows the effects of cross-terms.

To forge ahead analytically, we first assume that each of the two phases in our toy IGM also have different densities, each uniform. We indicate the average density of bubble material as $\langle\delta\rangle_i$, with the density in the bulk IGM $\langle\delta\rangle_n$, enforced by continuity such that

$$\langle\delta\rangle_n = -\langle\delta\rangle_i \frac{Q}{1-Q}. \tag{33}$$

The only unknown here is $\langle\delta\rangle_i$. One could treat it as a free parameter, or parametrize it flexibly as a function of redshift and/or bubble size. We discuss this possibility further in Section 2.4.2, where we introduce a simple model for $\langle\delta\rangle_i$.

Now, because we are assuming that neutral and ionized patches are of uniform (but redshift-dependent) densities, we can use a halo-model-like argument to write down the two-point terms involving b and δ. For example, $\langle b\delta' \rangle$ will be the sum of two terms: one in which the primed point is neutral, and one in which it is ionized, in each case weighted by the density of the relevant medium:

$$\langle b\delta' \rangle = \langle\delta\rangle_i \mathcal{P}_{bb} + \langle\delta\rangle_n \mathcal{P}_{bn}, \tag{34}$$

where we have used the previous notation to indicate the probability that two points are in bubbles, $\mathcal{P}_{bb}$, and a new term $\mathcal{P}_{bn}$ to indicate the probability that only one point is in a bubble. We can write the latter as

$$\mathcal{P}_{bn} = (1-\mathcal{P}_1)\exp\left[-\int \mathrm{d}R_b n_b(R_b)\left(V(R_b) - V_o(r, R_b)\right)\right] \times \left\{1 - \exp\left[-\int \mathrm{d}R_b n_b(R_b)\left(V(R_b) - V_o(r, R_b)\right)\right]\right\}. \tag{35}$$

In words, the above equation computes the probability that a single source can ionize one point but *not* the other (term in curly brackets), times the probability that the other point is *not* ionized by a different source that leaves the first point untouched (first exponential term), times the probability that a single source does not reside in the overlap volume and engulf both points (leading $1 - \mathcal{P}_1$ factor).

From equation (34), it is clear that if the density of ionized material is order unity ($\langle\delta\rangle_i \sim 1$), the cross-term $\langle b\delta' \rangle$ will be comparable to

Downloaded from https://academic.oup.com/mnras/article/514/2/2010/6595978 by Scuola Normale Superiore. Biblioteca user on 03 December 2022

the auto term $\langle bb' \rangle$, so long as $Q \ll 1$ and $P_{bn} \ll P_{bb}$. The first condition is plausible, and we will find that indeed $\langle \delta \rangle_i \approx 1$ when $Q \lesssim 0.2$. The second condition is less clear; certainly, on small scales, the P_{bn} term should be suppressed significantly by the $1 - P_1$ factor, since two points are increasingly likely to be in the same bubble on progressively smaller scales. However, on large scales, $P_{bn} \to Q(1 - Q)$, which means $P_{bn} \geq P_{bb}$ when $Q \leq 0.5$.

This is a curious feature of this model: for $\langle b\delta' \rangle \simeq \langle bb' \rangle$, it is possible that ξ_{21} becomes negative. Just comparing 2-pt terms, $\langle b\delta' \rangle = 2\langle bb' \rangle$, this will occur when

$$\frac{P_{bn}}{P_{bb}} < 2 \left(\frac{2\langle \delta \rangle_i - 1}{\langle \delta \rangle_i} \right) \left(\frac{1 - Q}{Q} \right). \tag{36}$$

Furthermore, if $\langle \delta \rangle_i < 0.5$, the RHS of equation (36) becomes negative, a condition that cannot be satisfied since P_{bn} and P_{bb} are both positive.

Such strong contributions from cross-terms involving ionization and density are not expected from more physically motivated models, but it is not surprising that they can become significant in our framework given the assumption of sharp, spherical bubbles and a perfect two-zone IGM. However, in detail, the amplitude and sign of ξ_{21} depend not only on the 2-pt contributions, but also on higher order terms, which we discuss next.

2.4.1 3- and 4-pt contributions to the power spectrum

The two most obvious ways to proceed with the remaining terms in equation (7) are to (i) neglect them, or (ii) fully embrace the binary framework and write down these terms following the logic applied to the $\langle b\delta' \rangle$ term above. Though higher order terms are likely to be smaller than the two-point terms, at least on scales $k \lesssim 1\,h\,\mathrm{Mpc}^{-1}$, they are not negligible in general (e.g. Lidz et al. 2007; Georgiev et al. 2021). Both options have some undesirable properties.

For example, option (i) must artificially set $\langle b\delta \rangle = 0$ in order to ensure that fluctuations vanish on large scales, despite the fact that correlations between bubbles and density imply that $\langle b\delta \rangle \neq 0$. However, the three and four-point terms $\langle b\delta\delta' \rangle$ and $\langle bb'\delta\delta' \rangle$ cannot be set to zero, otherwise 21-cm fluctuations will not vanish as $Q \to 1$ as they must. This is apparent from equation (7) – the leading factor of the matter fluctuations, $\langle \delta\delta' \rangle$, will persist regardless of Q, and so non-zero contributions from other terms involving the fluctuation $\delta\delta'$ are required in order for 21-cm fluctuations to vanish as $Q \to 1$. More on this momentarily.

Regarding option (ii), the binary field model predicts:

$$\langle bb'\delta \rangle = \langle \delta \rangle_i \mathcal{P}_{bb} \tag{37}$$

$$\langle b\delta\delta' \rangle = \langle \delta \rangle_i^2 \mathcal{P}_{bb} + \langle \delta \rangle_i \langle \delta \rangle_n \mathcal{P}_{bn} \tag{38}$$

$$\langle bb'\delta\delta' \rangle = \langle \delta \rangle_i^2 \mathcal{P}_{bb}. \tag{39}$$

A few observations about these terms:

(i) The only additional k-dependent suppression of 21-cm power beyond that caused by the two-point term $\langle b\delta' \rangle$ comes from the $\langle b\delta\delta' \rangle$ term above, which has a single leading factor of $\alpha = -1$ (see equation 7), in contrast to the $\langle bb'\delta \rangle$ and four-point terms which are both positive.

(ii) If we want to reduce the contribution from these higher order terms by, e.g. setting 39 to zero, we can only do so if we also set $\langle b\delta \rangle = 0$. Otherwise, fluctuations will not vanish on large scales (see last two terms of equation 7). However, this on its own will violate the requirement that fluctuations vanish as $Q \to 1$ (see above).

(iii) On large scales, the leading factor of $\mathcal{P}_{bb}$ means that the contribution of higher order terms will grow as reionization proceeds ($\mathcal{P}_{bb} \to Q^2$).

(iv) The power spectrum of a binary density field that traces the bubble field will exhibit a sharp feature on the typical bubble scale and no structure on smaller scales, at odds with the well-understood shape of the matter power spectrum.

Given these challenges, we employ a third option, which ensures that 21-cm fluctuations vanish on large scales and as $Q \to 1$. From equation (7), it is clear that the latter condition requires

$$-2\alpha\langle b\delta\delta' \rangle - \alpha^2 \langle bb'\delta\delta' \rangle \to \langle \delta\delta' \rangle \tag{40}$$

as $Q \to 1$. We take

$$\langle b\delta\delta' \rangle = Q\langle \delta\delta' \rangle \tag{41}$$

$$\langle bb'\delta\delta' \rangle = \langle bb' \rangle \langle \delta\delta' \rangle + \langle b\delta' \rangle^2 + \langle b\delta \rangle^2. \tag{42}$$

The second expression invokes Wick's theorem, common in the literature despite expectations that the 21-cm field is non-Gaussian, while the first is the simplest treatment that recovers the desired limiting behaviour. Together, these terms can be thought of as a correction factor applied to the matter fluctuations in equation (7), $\langle \delta\delta' \rangle \to (1 - Q)^2 \langle \delta\delta' \rangle$.

Finally, we set $\langle bb'\delta = 0$ to satisfy the requirement that $\xi_{21} \to 0$ on large scales.

2.4.2 Model for the density of bubble material

Critical to the simple cross-term treatment described above is knowledge of the density of bubble material, $\langle \delta \rangle_i$. To determine this bubble density, we make an argument similar to abundance matching in galaxy formation models: we assume that if a fraction Q of the IGM is in bubbles, then that volume is also the densest fraction Q of the IGM. Then, our task is to determine the density threshold above which Q per cent of the IGM resides.

To do this, we first define the variance of fluctuations at redshift z over a region of radius R to be

$$\sigma_R^2 = \int \frac{d^3k}{(2\pi)^3} P(k, z) |W_R(k)^2|, \tag{43}$$

where $P(k, z)$ is the matter power spectrum at z, and $W_R(k)$ is a window function encoding the shape of the region. We assume a spherical top-hat, which has the form

$$W_R(k) = \frac{3}{(kR)^3} \left[\sin(kR) - (kR)\cos(kR) \right]. \tag{44}$$

We further assume that the PDF of the density field, $\mathcal{P}(\delta_R)$, is lognormal (Coles & Jones 1991; Bi & Davidsen 1997). Then, if a fraction Q of the volume of the universe is ionized, we can associate that with a minimum density $\delta_R^{\min}$ through

$$\int_{\delta_R^{\min}}^{\infty} d\delta_R \mathcal{P}(\delta_R) = \frac{1}{2} \mathrm{erfc} \left(\frac{\delta_R^{\min}}{\sqrt{2}\sigma_R} \right) = Q, \tag{45}$$

or equivalently

$$\delta_R^{\min}(Q) = \sqrt{2}\sigma_R \, \mathrm{erfc}^{-1}(2Q). \tag{46}$$

Now it is easy to compute the density of ionized material

$$\langle \delta \rangle_i = \int_{\delta_R^{\min}}^{\infty} d\delta_R \mathcal{P}(\delta_R)\delta_R = \exp[-(\delta_R^{\min})^2 / \left(2\sigma_R^2\right)] \frac{\sigma_R}{\sqrt{2\pi}}. \tag{47}$$

Downloaded from https://academic.oup.com/mnras/article/514/2/2010/6595978 by Scuola Normale Superiore. Biblioteca user on 03 December 2022

The only free parameter of this model is the smoothing scale R employed to compute the variance in the density field, which we will hereafter refer to as $R_{\rm sm}$. There are only a few natural length-scales in our model thus far, all of which are related to the characteristic bubble size. We will explore several possibilities in Section 3, and compare to the mean density of ionized gas and ionization – density cross spectrum from 21cmFAST models for guidance.

Note that in this section, we effectively *have* made a physical argument about the nature of reionization, namely, that it occurs 'inside out.' Though this choice departs from our effort to avoid explicit astrophysical assumptions, it is appropriate for comparisons to 21cmFAST, and could be generalized in the future (see e.g. Pagano & Liu 2020, for one approach). Fortunately, scenarios with a strong degree of *anticorrelation* between ionization and density fields generate stronger 21-cm fluctuations than the alternative, and should thus be easier to rule out as upper limits become more stringent (Pagano & Liu 2020, 2021).

2.5 Redshift space distortions

We now add the effect of peculiar velocities (see e.g. Kaiser 1987; Barkana & Loeb 2005a). These give rise to redshift-space distortions (RSDs), which to linear order we can write as

$$\delta T_b = T_0(z)(1+\alpha b)(1+\delta-\delta_v)(1+\delta_\varphi), \tag{48}$$

where δ_v is the line-of-sight velocity-gradient anisotropy. This last quantity is given (again to linear order and during the epoch of interest) by $\delta_v = -\mu^2\delta$ in Fourier space, where $\mu = k_\parallel/k$, so we have

$$\delta T_b = T_0(z)(1+\alpha b)[1+(1+\mu^2+\beta_\varphi)\delta]. \tag{49}$$

We ignore non-linear RSDs (e.g. Mao et al. 2012; Jensen et al. 2013) – see Greig & Mesinger (2018) for the implementation in 21cmFAST. We also neglect the light cone effect (see e.g. Datta et al. 2012; La Plante et al. 2014; Chapman & Santos 2019, for detailed treatments). This is a conservative approach given that these effects boost the power spectrum, and so will drive lower limits on the IGM spin temperature upward once included.

In traditional galaxy surveys one can measure different modes $\vec{k}$, and thus probe the μ dependence of the RSDs. In 21-cm studies, however, the situation is different. Foregrounds impose an observational cutoff, as small wavenumbers along the line of sight are inaccessible for cosmology (e.g. Datta, Bowman & Carilli 2010; Morales et al. 2012; Parsons et al. 2012; Vedantham, Udaya Shankar & Subrahmanyan 2012; Thyagarajan et al. 2013; Liu, Parsons & Trott 2014). This 'foreground wedge' in fact occupies the majority of the Fourier plane, so in practice the modes observed at any fixed spherical k are chiefly along the line of sight, with $\mu \approx 1$.[4] We will therefore often simply fix $\mu = 1$. Alternatively, when comparing to 21cmFAST simulations we will manually set $\mu = 0.6$, which recovers $\langle(1+\mu^2)^2\rangle = 1.87$, the (linear) spherically averaged value of RSDs. Note that the use of a single average value of μ is likely to be overly simplistic due to the non-linear relationship between density and ionization (Pober 2015). We defer a more detailed treatment to future work.

Upon including RSDs, we take $\delta \to \delta - \delta_v$, resulting in the following modifications to cross-terms:

$$\langle b\delta\delta'\rangle \to \langle b\delta\delta'\rangle - \langle b\delta\delta'_v\rangle \tag{50}$$

$$\langle b\delta\delta'\rangle \to \langle b\delta\delta'\rangle - \langle b\delta\delta'_v\rangle - \langle b\delta'\delta_v\rangle + \langle b\delta_v\delta'_v\rangle \tag{51}$$

$$\langle bb'\delta\delta'\rangle \to \langle bb'\delta\delta'\rangle - \langle bb'\delta\delta'_v\rangle - \langle bb'\delta'\delta_v\rangle + \langle bb'\delta_v\delta'_v\rangle, \tag{52}$$

where we have shown for completeness the 3- and 4-pt terms, despite neglecting them in what follows.

2.6 Density fluctuations

Finally, $\langle\delta\delta'\rangle$ is equivalent to the matter correlation function, $\xi_{\delta\delta}$ – we compute the linear matter power spectrum using CAMB (Lewis, Challinor & Lasenby 2000) and Fourier transform to obtain $\xi_{\delta\delta}$, i.e.

$$P_{\delta\delta}(k) = \frac{1}{(2\pi)^3}\int 4\pi R^2 \frac{\sin(kR)}{kR}\xi_{\delta\delta}(R)\mathrm{d}R \tag{53}$$

which is the same operation used to convert ξ_{21} to the 21-cm power spectrum in equation (10). The highly-oscillatory nature of these integrals pose a challenge – we include options for Clenshaw–Curtis integration (as in, e.g. Diemer 2018) and FFTLog algorithms (Talman 1978; Hamilton 2000), which are implemented in SCIPY and MCFIT,[5] respectively. The MCFIT approach is generally faster by ∼2–3x, and so is the default in MICRO21CM.

2.7 Putting it all together

In Fig. 2, we assemble a representative 21-cm power spectrum term by term. In the left-hand panel, we start in a $z = 8$ IGM with uniform temperature $T_{\rm S} = 1.8$ K, and only matter fluctuations (dotted black). In reality, gas density and kinetic temperature are coupled due to adiabatic cooling – the blue dashed curve accounts for this correlation, which suppresses power on all scales since denser regions are also warmer than less dense regions (see Section 2.2). These first two cases ignore redshift-space distortions (see Section 2.5). The dotted–dashed cyan curve averages the power spectrum over all μ, providing a boost in power that nearly cancels out the suppression caused by the φ correction. Finally, we note that many current experiments almost exclusively probe line of sight modes, with $\mu \sim 1$, which we show in the solid magenta curve (also see fig. 3 in The HERA Collaboration 2021a). The ratio of each case to the matter-only case is shown in the bottom left panel.

Next, in the right-hand panel of Fig. 2, we start from the dotted–dashed cyan curve of the left-hand panel and add ionized bubbles with a lognormal size distribution (see Section 2.3). We further assume an ionized fraction $Q = 0.4$, and a typical bubble size of $R_b = 5$ cMpc, which are reasonable choices for $z \sim 8$. If reionization were spatially homogeneous, the power would be suppressed by a factor of $(1-Q)^2$ at all k, which we show in the opaque cyan curve. However, a model with discrete bubbles boosts power around the bubble scale, as seen in the orange-dotted curve, which adds the ionization autocorrelation term only (see equation 7). Finally, the red curve adds the cross-terms involving both ionization and density following the procedure of Sections 2.4 and 2.4.1. In the bottom row, we once again show the ratio between each power spectrum in the top panel with the matter-fluctuations-only case (dotted black).

Given that many current experiments probe large $k \simeq 0.2h\,\mathrm{Mpc}^{-1}$ scales, largely at frequencies $\nu \gtrsim 100$ MHz (e.g. Barry et al. 2019; Mertens et al. 2020; The HERA Collaboration 2021b) where the bulk of reionization is expected to occur (e.g. Bouwens et al. 2015; Mason, Trenti & Treu 2015; Robertson et al. 2015; Gorce et al. 2018;

[4] For HERA $\mu \gtrsim 0.97$, with values closer to one indicative of increasingly aggressive wedge cuts.

[5] https://github.com/eelregit/mcfit

Downloaded from https://academic.oup.com/mnras/article/514/2/2010/6595978 by Scuola Normale Superiore. Biblioteca user on 03 December 2022

Finkelstein et al. 2019), in Fig. 3 we show predictions for the large scale $k = 0.2\,h\,\mathrm{Mpc}^{-1}$ power over all of (Q, T_S) space. From left to right we increase the typical bubble size from 3 to 5 and 5 to 10 cMpc, and colour code by $\Delta^2(k = 0.2)$ from 1 to 10^4 mK^2 (blue to red colours). The cross-hatched region in the bottom of each panel indicates temperatures below 1.8 K, which is the expected minimum temperature of an adiabatically cooled $z = 8$ IGM, making it clear that 10^4 mK^2 signals (red) require a super-cooled IGM.

Focusing next on the 10^3 mK^2 range (orange), which is comparable to recent limits, we see that viable scenarios generally require $T_\mathrm{S} \simeq 3 - 10$ K. If bubbles are small (left-hand panel), sub-adiabatic temperatures may be required, but for larger bubbles, $R_p = 5$ or 10 cMpc (centre, right columns), 10^3 mK^2 fluctuations are possible without sub-adiabatic temperatures provided that reionization is not just beginning or just ending. Ionization fluctuations are maximized near the reionization midpoint, in which case, such fluctuations can be achieved if the IGM temperature is $\log_{10}(T_\mathrm{S}/\mathrm{K}) \simeq 0.5$, i.e. if $T_\mathrm{S} \simeq 3$ K, consistent with the interpretation of HERA's recent 10^3 mK^2 limits The HERA Collaboration (2021a).

In general, as the strength of ionization fluctuations grow, 21-cm fluctuations will also grow stronger if the temperature is held fixed. As ionization fluctuations decline in the latter half of reionization, holding the power constant demands that the spin temperature deviate more strongly from T_R. There are two exceptions to this behaviour. At early times, if bubbles are small (left column), the assumed positive correlation between ionization and density results in a decline in the power as Q grows (at fixed T_S). Second, if the spatial scale of interest is much larger than the typical bubble size, $k \lesssim R_p^{-1}$, the amplitude of fluctuations on that scale depend very little on the ionized fraction.

Fig. 3 provides a means of rough, by-eye inference. Provided our phenomenological model is reasonably accurate, one can simply 'read off' the Q, T_S, and R values that are consistent with new power spectrum upper limits. Though we have some indication already that the phenomenological model performs well compared to more sophisticated calculations, e.g. the association of 10^3 mK^2 fluctuations at $k = 0.2\,h\,\mathrm{Mpc}^{-1}$ with $T_\mathrm{S} \simeq 3$ K (see above; also The HERA Collaboration 2021a), in the next section, we provide a much more detailed comparison to 21cmFAST calculations.

3 COMPARISON TO 21CMFAST MODELS

Having outlined the various components of our phenomenological model, we now compare its predictions to two illustrative models generated with 21cmFAST (Mesinger & Furlanetto 2007; Mesinger et al. 2011; Murray et al. 2020). Our goal is to assess the accuracy of the model relative to more sophisticated calculations, test various modeling choices, and set expectations for interpreting fits to 21cmFAST mocks in Section 4.

We will compare to two benchmark 21cmFAST models. The first, model #1, has the same set of parameters as the `AllGalaxies` simulations of Muñoz et al. (2021b), and thus includes atomic-cooling galaxies forming Pop II stars (following Park et al. 2019) as well as molecular-cooling haloes forming Pop III (following Qin et al. 2020), with joint feedback from Lyman–Werner photons (Haiman, Rees & Loeb 1997; Machacek, Bryan & Abel 2001; Visbal et al. 2014) and streaming velocities (Tseliakhovich & Hirata 2010; Visbal et al. 2012; Muñoz 2019a). The second model, on the other hand, imposes a cutoff for star formation at $M_\mathrm{turn} = 10^9\,\mathrm{M}_\odot$, so haloes below that mass do not form stars. As a consequence, there are no Pop III stars in that model, and the evolution of the 21-cm signal is faster. The rest of galaxy properties are the same between the two models, with star formation parameters calibrated to high-z luminosity functions (from Finkelstein et al. 2015) and X-ray spectra representative of X-ray binaries hardened by neutral columns expected of low-mass galaxies at high-z (Das et al. 2017). The parameter values for each 21cmFAST model are summarized in Table 1.

Table 1. Summary of 21cmFAST models. We use two reionization models: slow [corresponding to the EOS21 parameters of Muñoz et al. (2021b)] and fast (which have the same PopII parameters but a nonzero turnover mass M_turn, and thus no PopIII stars). For each of those two models we vary the X-ray heating efficiency parameter L_X as indicated in this table. For the entire set of galaxy parameters, and how they fit all current EoR data, see Muñoz et al. (2021b).

ID	Model name	M_turn	$\log_{10}(L_X/\mathrm{SFR})$
1a	Slow/no heat	n/a	37.5
1b	Slow/cold	n/a	38.5
1c	Slow/warm	n/a	39.5
1d	**Slow/hot (EOS21)**	**n/a**	**40.5**
2a	Fast/no heat	$10^9\,\mathrm{M}_\odot$	37.5
2b	Fast/cold	$10^9\,\mathrm{M}_\odot$	38.5
2c	Fast/warm	$10^9\,\mathrm{M}_\odot$	39.5
2d	Fast/hot	$10^9\,\mathrm{M}_\odot$	40.5

The only parameter we will vary is the X-ray luminosity of the first galaxies, as it strongly affects the values of T_S during the epoch of interest. We will start with the fiducial choice of $\log_{10}(L_X/\mathrm{SFR}) = 40.5$, which is $\sim$10x higher than that generated by high-mass X-ray binaries in nearby star-forming galaxies (Mineo, Gilfanov & Sunyaev 2012), as expected of low-metallicity environments at high redshift (e.g. Fragos et al. 2013; Brorby et al. 2016). We explore this parameter in order-of-magnitude steps down to 37.5, so as to cover a broad range of possibilities.

We show the mean ionization histories (top) and ionization power spectra (bottom) for both our models in Fig. 4. Reionization occurs more gradually in model #1 than in model #2, so we dub them 'slow' and 'fast,' respectively, though they are both 'late reionization' scenarios, with neutral fractions of $\sim$20 per cent at $z \sim 6$, in accordance with recent constraints (Becker et al. 2015; Keating et al. 2020; Bosman et al. 2021; Qin et al. 2021). We will examine all four possibilities for the spin temperature evolution in Section 4.3 but for now focus only on the ionization field. In Fig. 5, we compare the ionization power spectrum computed by our model to those from 21cmFAST at the same mean ionized fraction, $1 - \bar{x}_{\mathrm{H\,I}}$, or equivalent volume filling fraction, Q in the phenomenological model.[6] At a series of ionized fractions spanning from 20 to 80 per cent, we show power spectra drawn from 21cmFAST model 1 (points), and a best-fitting representation using the phenomenological model. The parameters of a lognormal BSD (R and σ; solid lines) and power-law-with-exponential-cut-off model (R and γ; dashed lines) are calibrated to match the 21cmFAST power spectra at $0.1 \leq k/(h\,\mathrm{Mpc}^{-1}) \leq 0.8$. Overall, the shape of the ionization power spectrum can be well-modelled using either BSD parametrization. We exclude points at $k > 0.8h\,\mathrm{Mpc}^{-1}$ from the fit, as discrete sampling effects start to become apparent at smaller scales and so would artificially bias the calibration.

In Fig. 6, we show the best-fitting BSD parameters as a function of Q for both 21cmFAST scenarios (left and right columns). For both BSD models, the power spectra are consistent with rapidly growing

[6]Note that this is itself a potential source of uncertainty in our comparison, since the ionization field is not binary in 21cmFAST.

Downloaded from https://academic.oup.com/mnras/article/514/2/2010/6595978 by Scuola Normale Superiore. Biblioteca user on 03 December 2022

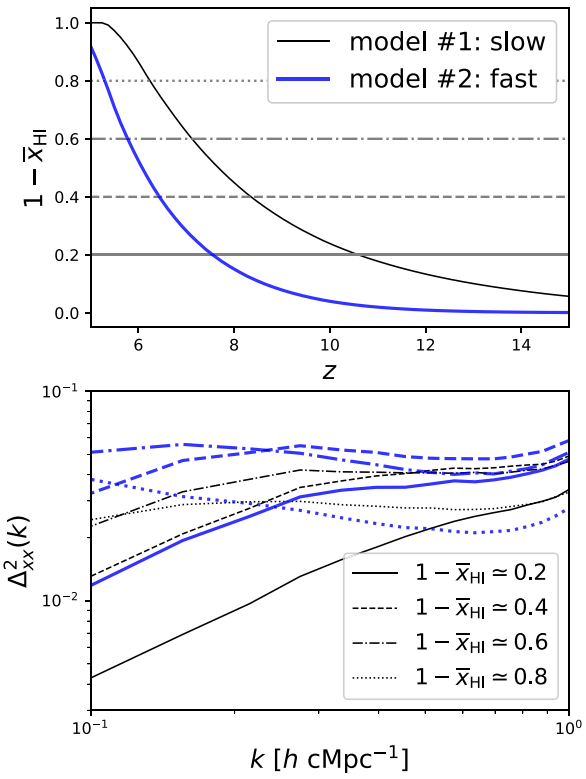

Figure 4. Example 21cmFAST models used for comparison throughout. Mean ionization histories (top) and ionization power spectra (bottom) at four different mean ionized fractions, Q. Both models are 'late reionization' scenarios, differing largely in the duration of reionization, with the 'slow' model (black), the transition from 20 to 80 per cent ionization takes $\Delta z \simeq 4$, while in the 'fast' scenario (blue), reionization begins later and the 20 to 80 per cent duration is $\Delta z \simeq 2$. The 'fast' scenario generates stronger fluctuations at fixed ionized fraction, as the sources of reionization are more biased.

bubble sizes (top row). For the lognormal BSD, the dispersion σ gradually increases from $\sigma \simeq 1.25$ to 2.25 (second row). The power-law-times-exponential BSD varies less with Q; generally γ lies between $\gamma \sim -3.75$ and $\gamma \sim -3.25$.

With calibrated $R(Q)$ and $\sigma(Q)$ values in hand, we now compare the phenomenological model's prediction for the density of bubble gas to the results extracted from 21cmFAST boxes. As discussed in Section 2.4.2, the key choice in the phenomenological model is the scale on which to smooth the density field when computing its variance, which sets the fraction of the volume above a given density contour. Here, we explore three options, which assume a smoothing scale equal to the scale of (i) the peak of the volume-weighted, logarithmic BSD, $V{\rm d}n/{\rm dlog}\,R$, (ii) the peak of the volume-weighted, linear BSD, $V{\rm d}n/{\rm d}R$, and (iii) the radius at which the joint probability of ionization, $\langle bb' \rangle$ is no longer equivalent to the one-bubble term. The final option requires choosing a threshold, e.g. the scale at which $P_1 = X\langle bb' \rangle$, with X a free parameter. It is not obvious how to choose this critical threshold, or if it is more or less meaningful than options (i) or (ii) – it merely serves as another approach to employ in comparisons with 21cmFAST. We show cases for a threshold of 0.97 ± 0.025, which result in densities that generally lie between the predictions of options (i) and (ii) described above.

In Fig. 7, we show the bubble density predictions for the lognormal BSD model compared to 21cmFAST. While each approach results in the correct behaviour qualitatively, none provide an accurate match at all Q or in both reionization scenarios. In general, the density of ionized gas evolves much more rapidly in the phenomenological model than in 21cmFAST models suggest at early times, $Q \lesssim 0.3$. While model #2 is well-matched by smoothing scale option (ii) described above (dashed curves; bottom panel), at least at $Q \gtrsim 0.2$, the same approach does not provide as good a match for model #1 (dashed curve; top panel). For a suitably chosen threshold, a smoothing scale linked to the decline of the one-bubble term does result in slightly more gradual evolution (dotted curves). Though the solid curves do not provide as good of a match in general, we adopt them in all that follows for reasons we discuss further momentarily.

In Fig. 8, we move on to the phenomenological model's predictions for the cross spectrum between the ionization and density fields. Line-style conventions are the same as in Fig. 7. As in the case of the ionization power spectrum, the phenomenological model provides a very reasonable prediction for the *shape* of $k^3P_{x\delta}$ as a function of k. However, there are systematic offsets from the 21cmFAST models that vary as a function of $R_{\rm sm}$. Interestingly, the best matches occur when $R_{\rm sm}$ is tied to the peak in $V{\rm d}n/{\rm d}R$ (dashed) or the decline of the one-bubble term (dotted).

Though setting $R_{\rm sm}$ to the scale where $V{\rm d}n/{\rm dlog}\,R$ peaks is not the obvious choice based on Figs 7 and 8, it is the only option that keeps the 21-cm power spectrum positive (see Section 2.4). Because our two-zone IGM model overestimates the strength of the fluctuations in the ionization field, as well as cross-terms involving ionization and density, *under*estimating the density of ionized gas acts as a countermeasure that keeps the amplitude of fluctuations in check. While this is far from an ideal solution, for our purposes it may not matter, as long as our predictions for the 21-cm fluctuations are reasonably accurate.

Having compared predictions for each component of our model to the statistics of the ionization field and its relation to the density field, we now do one last comparison to the 21-cm power spectrum. With R and σ calibrated to 21cmFAST models at a series of Q values, we can perform a one-dimensional fit to the 21-cm power spectrum produced by 21cmFAST varying $T_{\rm S}$ alone, and see if the $T_{\rm S}$ found in our fits agrees well with the mean $T_{\rm S}$ drawn from 21cmFAST. Because $T_{\rm S}$ in the phenomenological model refers to the mean temperature of the bulk IGM only i.e. neglecting ionized regions), we average over all voxels with $x_{\rm H\,I} \geq 0.95$ in the 21cmFAST boxes (as in, e.g. The HERA Collaboration 2021a).

The results of this test are shown in Figs 9 and 10, for slow and fast reionization scenarios, respectively. In each case, the model in the first row has minimal heating from astrophysical sources, while subsequent rows have increasingly efficient heating. Dashed lines in the first four columns show 21-cm power spectra obtained using the phenomenological model *when assuming the true underlying value for* T_S. Solid lines instead show power spectra obtained if T_S is allowed to vary freely in a 1D fit. The final column shows the temperature evolution, again including the true evolution (points and dashed line) and that recovered in the 1D fit to the power spectrum (solid lines).

There is much to unpack from Figs 9 and 10. Generally speaking, the phenomenological model performs best in the early stages of reionization, $Q \lesssim 0.6$ (first three columns), and begins to struggle later in reionization, particularly when the spin temperature is large. This result is not unexpected; the sharp spherical bubbles in our model will result in stronger ionization fluctuations at fixed Q compared to 21cmFAST, meaning we should expect to underestimate the contrast, $1 - T_\gamma/T_{\rm S}$. Indeed, this is generally what we find in Figs 9 and 10. The shape of our power spectra is a reasonably good

Downloaded from https://academic.oup.com/mnras/article/514/2/2010/6595978 by Scuola Normale Superiore. Biblioteca user on 03 December 2022

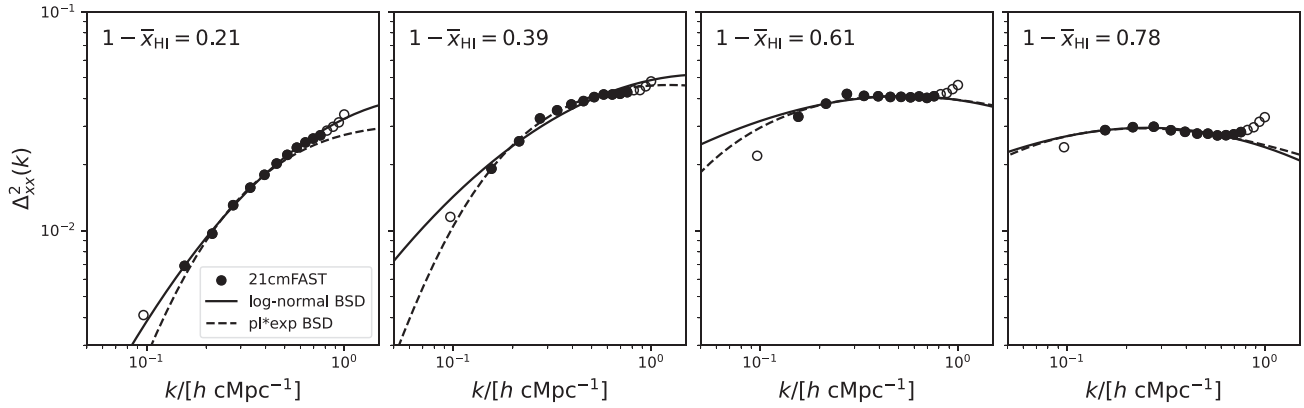

Figure 5. Comparison of phenomenological model (lines) to 21cmFAST ionization power spectra (points) for the 'slow' reionization scenario. From left to right, we compare at fixed mean ionized fraction from $Q \simeq 0.2$ to $Q \simeq 0.8$, as indicated in the upper left corner of each panel, assuming that $Q = 1 - \overline{x}_{\rm HI}$. We fit the 21cmFAST points to calibrate R and σ (or γ) for a fair comparison. Dashed lines correspond to the power-law-times-exponential BSD, while solid lines correspond to the lognormal BSD.

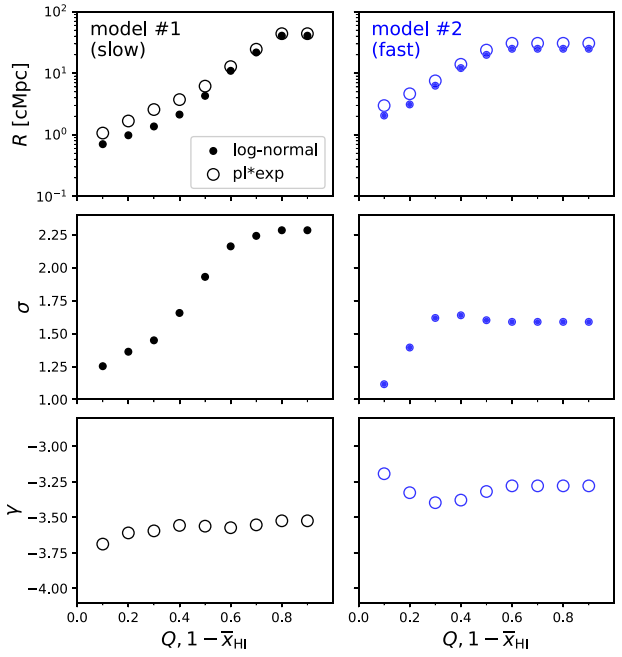

Figure 6. Evolution of BSD parameters inferred from fits to example 21cmFAST ionization power spectra. The typical bubble size evolution is well-captured by a power-law dependence on Q (top) for both 'slow' and 'fast' scenarios (left and right columns, respectively), while σ exhibits a gradual rise and eventual plateau (middle row). The power-law slope γ is relatively constant for both reionization scenarios (bottom row).

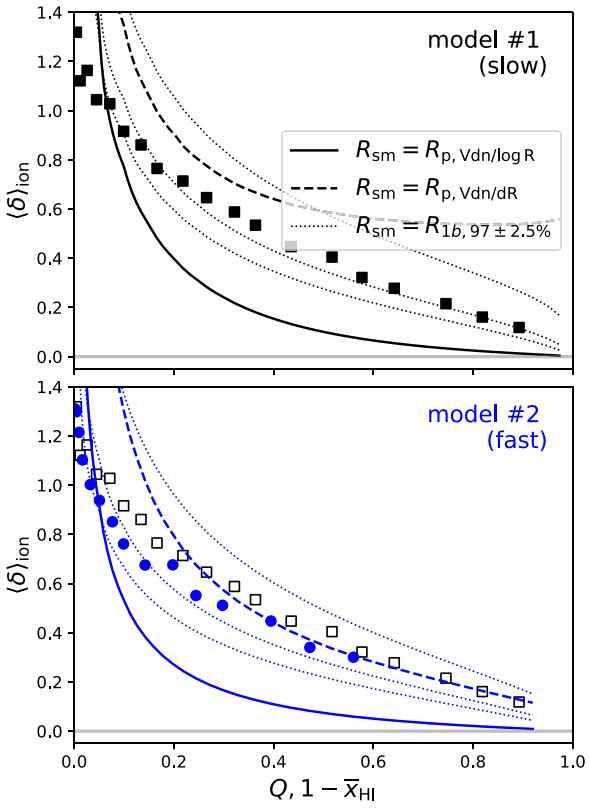

Figure 7. Mean density of ionized gas in 21cmFAST (points) and phenomenological models (lines) versus ionized fraction for the 'slow' (top) and 'fast' EoR scenarios (bottom). We show predictions for several different smoothing scales (see Section 2.4.2), which sets the variance in the density field and thus density contour containing Q per cent of the volume, as indicated in the legend in the top panel. We have assumed the lognormal BSD, with $R(z)$ and $\sigma(z)$ calibrations shown in Figs 5 and 6. Black points are repeated in both panels for ease of comparison.

match, becoming flatter near the midpoint of reionization. Similarly, the shape of our thermal histories are in good agreement with the 21cmFAST models, though exhibit small biases, as described above.

There are cases in which the 21-cm power on large scales departs from the phenomenological model's prediction even at early times. We have verified that this excess power on large scales $k \lesssim 0.1\,h\,{\rm Mpc}^{-1}$ and early times $Q \lesssim 0.2$ (see bottom two panels in the first column) is not due to the presence of temperature fluctuations

Downloaded from https://academic.oup.com/mnras/article/514/2/2010/6595978 by Scuola Normale Superiore. Biblioteca user on 03 December 2022

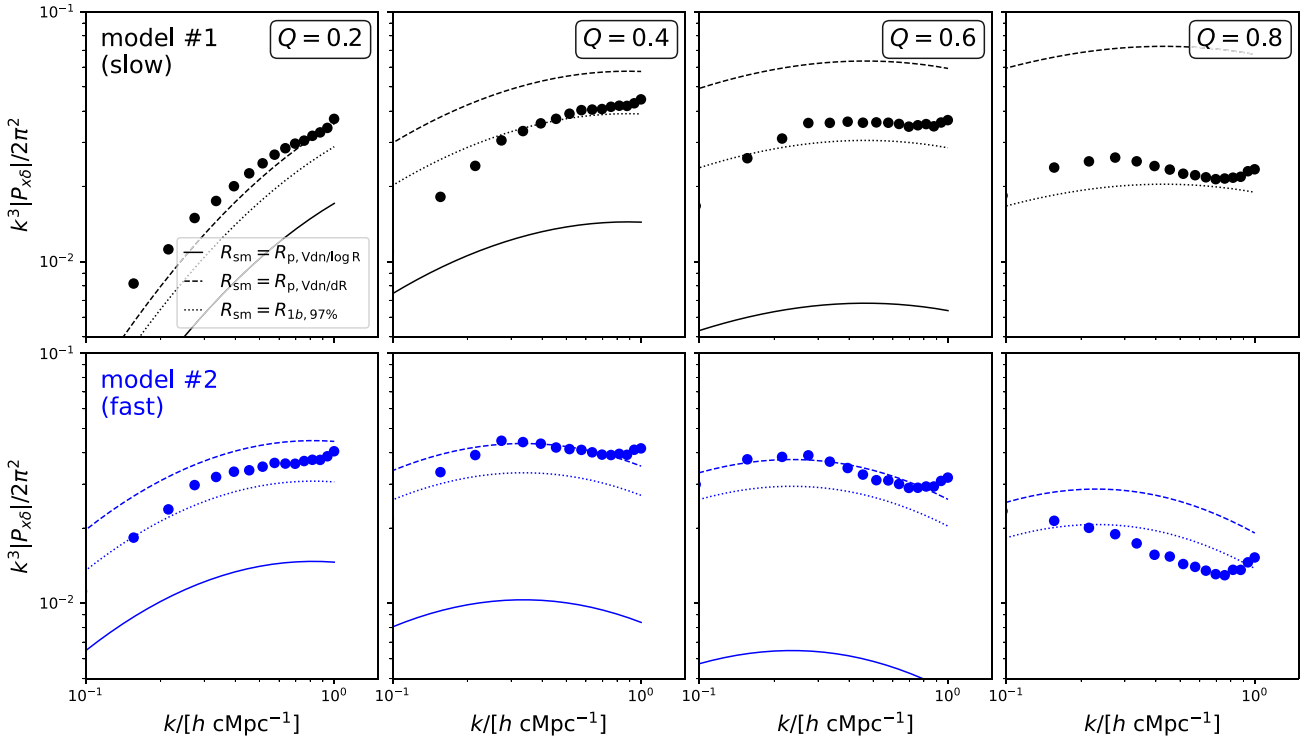

Figure 8. Cross spectrum between ionization and density in 21cmFAST (points) and phenomenological models (lines) for models 1 (top) and 2 (bottom). As in Fig. 7, we show predictions for several different smoothing scales (see Section 2.4.2), which sets the variance in the density field and thus density contour containing Q per cent of the volume. We have assumed the lognormal BSD, with $R(z)$ and $\sigma(z)$ calibrations shown in Figs 5 and 6.

by running an additional 21cmFAST simulation in which we set $T_S = \infty$ everywhere by hand. This large-scale power is likely due to the additional ionization from X-ray sources, which our phenomenological approach is not equipped to model. As we will see in Section 4.3, uncertainties on these k modes are large, and so unlikely to contribute significantly to biasing our fits.

Finally, notice that in some cases, T_S values just above and just below T_γ are both valid, because the power spectrum is proportional to $(1 - T_\gamma/T_S)^2$ (neglecting adiabatic corrections). When this is the case, we show two solid lines, one for each T_S solution. In reality, this degeneracy will be broken by fitting data at multiple redshifts simultaneously, unless the true temperature evolution of the IGM is very gradual, and hovers near $T_S \approx T_\gamma$. However, the phenomenological model is never driven to $T_S \gg T_\gamma$ – as seen in the bottom rows of Figs 9 and 10. Though the spin temperature saturates early in 21cmFAST in each case, the recovered T_S by the phenomenological model never exceeds $\sim$ few times T_γ. This behaviour is not unexpected. First, we are less and less sensitive to the spin temperature as T_S grows much larger than T_γ, and so should expect constraints to be poor. And second, given that our model overestimates the strength of ionization fluctuations at fixed Q due to the idealized assumption of sharp, spherical bubbles, we will then *underestimate* the contrast, $1 - T_\gamma/T_S$ when fitting to a given 21-cm power. As we will see in the next section, uncertainties on the inferred spin temperature history inflat dramatically in these regime, providing an indicator that a measurement lies in a region of parameter space captured poorly by our model in its current form.

In this section, we showed that the phenomenological model provides a reasonably accurate match to predictions from 21cmFAST simulations. There are differences, as expected, but they are largely systematic, with biases at the level of $\sim$20–40 per cent in the amplitude of the ionization power spectrum and ionization–density cross spectrum. These biases will of course affect our ability to recover the mean spin temperature and ionized fraction in fits to mock data sets generated with 21cmFAST, since these parameters act largely as normalization factors in the phenomenological model, whereas R and σ (or γ) carry more shape information. We will assess this possibility in the next section.

4 RECOVERY OF IGM PROPERTIES FROM 21-CM MOCK POWER SPECTRA

We now determine the extent to which the mean properties of the IGM and parameters describing the size distribution of bubbles can be recovered from observations using phenomenological models. We explore two different scenarios. First, we fit a mock signal generated with the phenomenological model itself (Section 4.1), and explore the power spectrum's sensitivity to the BSD in more detail in Section 4.2. Then, in Section 4.3, we fit mocks generated with 21cmFAST. We use EMCEE (Foreman-Mackey et al. 2013) (version 2.2.1), an implementation of the affine-invariant sampler of Goodman & Weare (2010), to perform all Markov Chain Monte Carlo (MCMC) fits to mock observations.

4.1 Fits to phenomenological model mocks

We begin with a simple exercise to make sure that the parameters of the phenomenological model can be recovered under idealized circumstances. The input model adopted for this calculation assumes $Q = 0.4$, $T_S = 3$ K, $R_b = 5$ cMpc, and $\sigma_b = 1$ (i.e. a lognormal BSD). We compute error-bars appropriate for HERA using 21CM-

Downloaded from https://academic.oup.com/mnras/article/514/2/2010/6595978 by Scuola Normale Superiore. Biblioteca user on 03 December 2022

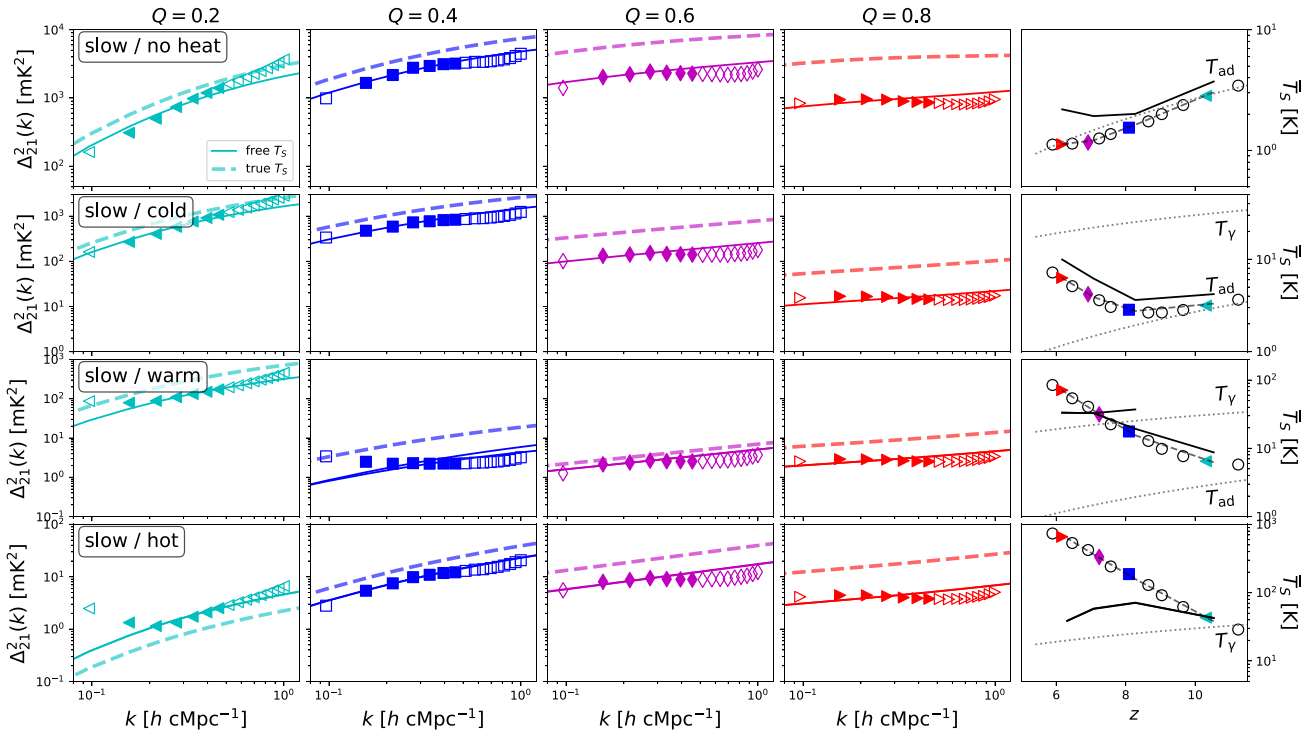

Figure 9. 21-cm power spectrum from 21cmFAST (points) and phenomenological models (lines) for the 'slow' reionization scenario, including several models with increasingly efficient heating (top to bottom). The first four columns show 21-cm power spectra at $Q = 0.2$, 0.4, 0.6, and 0.8, while the final column shows the redshift evolution of T_S. In each panel, we freeze $Q = 1 - \overline{x}_{HI}$, and take R and σ to their calibration values (derived from the ionization power spectrum alone; see Figs 5 and 6). We then fit the 21-cm power spectrum with a single free parameter, T_S. The results of this fit are shown as solid lines, while dashed lines adopt the true T_S from 21cmFAST. Note that because T_S is obtained at each snapshot independently, it can be double-valued (final column), since $\Delta^2_{21} \propto (1 - T_R/T_S)^2$. Such behaviour does not occur when jointly fitting measurements at different redshifts with a parametric form for $T_S(z)$ (see Section 4.3).

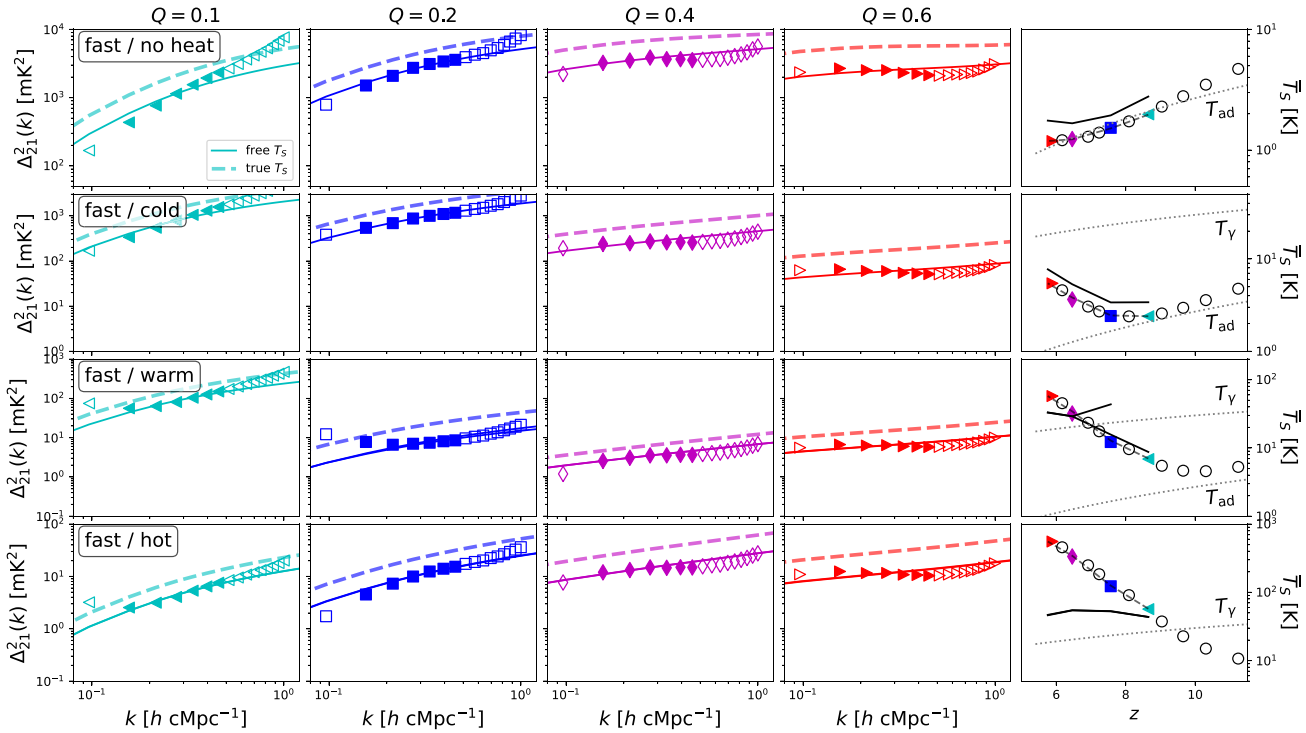

Figure 10. Same as Fig. 9, but for the 'fast' reionization scenario. The first four columns show 21-cm power spectra at $Q = 0.1$, 0.2, 0.4, and 0.6, while the final column shows the redshift evolution of T_S. The efficiency of X-ray production increases from top to bottom.

Downloaded from https://academic.oup.com/mnras/article/514/2/2010/6595978 by Scuola Normale Superiore. Biblioteca user on 03 December 2022

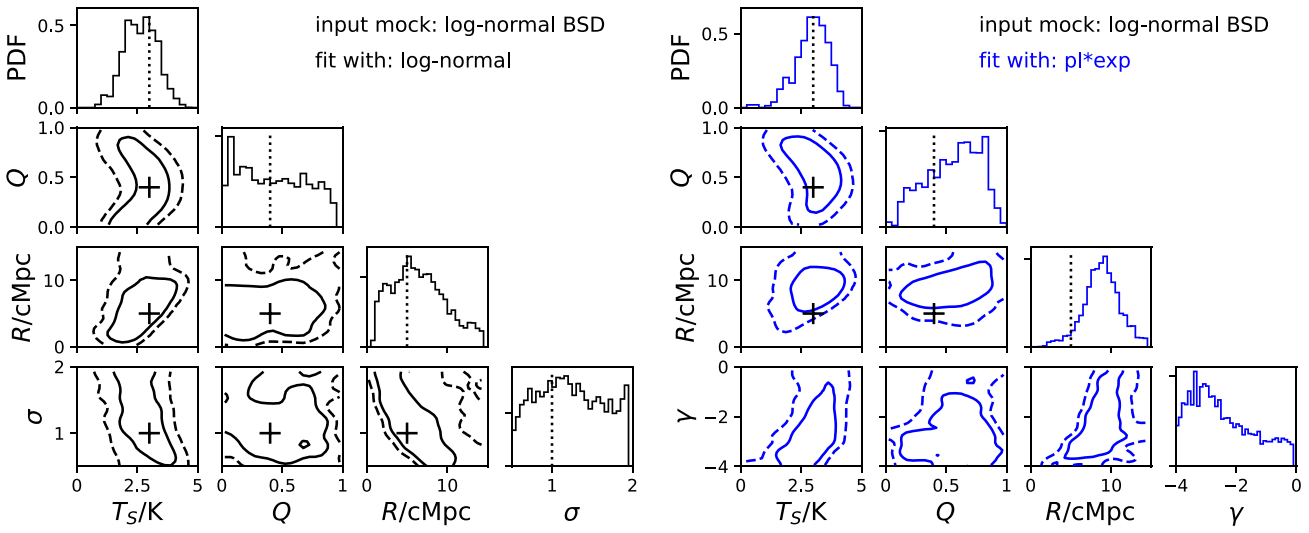

Figure 11. Parameter constraints obtained via fits to a mock phenomenological signal with $T_S = 3$ K, $Q = 0.4$, $R = 5$, and $\sigma = 1$ at $z = 8$ using both BSD parametrizations. *Left*: Posterior distributions when mock signal is fit with the same BSD as that used to generate the mock (lognormal). *Right*: Results when we assume a power-law-times-exponential BSD in the fit. Solid and dashed curves indicate 68 and 95 per cent confidence regions, respectively, while the input values are denoted by crosses in the interior panels and vertical dotted lines for the 1D posteriors. Only the spin temperature is constrained well in each case (first column), though the 2D relationship between T_S and R can also be reliably constrained.

SENSE[7] (Pober et al. 2013, 2014) under the assumption of 'moderate' foregrounds, and set a superhorizon buffer of $a = 0.05\,h\,\mathrm{Mpc}^{-1}$. We take a system temperature

$$T_{\rm sys} = 100\,\mathrm{K} + 120\,\mathrm{K} \times (\nu/150\mathrm{GHz})^{-2.55}, \tag{54}$$

following DeBoer et al. (2017), and assume 1 yr (1080 h) of observation time. We bin linearly in wavenumbers (with $\Delta k = 0.1\,h\,\mathrm{Mpc}^{-1}$), and frequency (with a bandwidth $\Delta\nu = 8$ MHz), from $z \approx 6$–24. This results in a signal-to-noise ratio SNR ~ 200 for the EOS21 model (1d in Table 1).

The noise on the 21-cm power spectrum has a cosmic-variance component that depends on the fiducial signal that we study. Rather than re-running `21cmSense` for each different model in our array, we have separated the thermal $\sigma_{\rm th}$ and cosmic-variance $\sigma_{\rm CV}$ contributions to the error, which can be added to find the full error as (Muñoz et al. 2021a)

$$\sigma_{\rm full} = \sigma_{\rm th} + \sigma_{\rm CV}, \tag{55}$$

where $\sigma_{\rm th}$ is a standard output of `21cmSense`, and

$$\sigma_{\rm CV} = a_{21}\,\Delta^2_{21} \tag{56}$$

is given by the k- and z-dependent coefficients a_{21}, which we find for our chosen bins. This relationship simply states that the total noise $\sigma_{\rm full}$ grows linearly with the size of the fiducial signal, due to cosmic variance.[8]

In Fig. 11, we show the result of this simple forecast. In the left triangle plot, we show the recovery of the lognormal BSD model parameters, and in the right-hand panel, we show the results of a fit that uses a power-law-times-exponential BSD, though the input mock remains that generated with a lognormal BSD. In the first case, the true input values are recovered well, though generally with large uncertainties. For example, Q can only be constrained to $0.04 \lesssim Q \lesssim 0.91$ at 2σ confidence. The typical bubble size is constrained to $3.5 \lesssim R/\mathrm{cMpc} \lesssim 12.8$ (1σ), with no real constraints on σ. The mean spin temperature of the IGM, however, is constrained well, to $T_S \simeq 3 \pm 1$ K. The constraint on the spin temperature largely holds even if fitting with the 'wrong' BSD, as shown in the right-hand panel of Fig. 11. The ionized fraction Q is similarly only weakly constrained, while the typical bubble size is biased to slightly larger sizes, $R \simeq 8$ cMpc. In principle, one could perform model selection and select the best-fitting parametrization with the Bayesian evidence, though we defer such an analysis to future work.

The weakness of a single-epoch fit is not unexpected. Our model has four parameters, and is being used to fit a curve that only mildly departs from a pure power law. We chose an input mock with a cold enough spin temperature that only solutions with $T_S \ll T_\gamma$ are viable, which significantly helps to constrain T_S. The results are largely consistent for each BSD, in that T_S is the best constrained parameter, with the others only weakly constrained. We take a deeper look at the power spectrum's sensitivity to the BSD next.

4.2 Sensitivity of PS to BSD

In Section 4.1, we emphasized potential constraints on Q and T_S, and discussed the BSD only as a nuisance and its potentially odd behaviour in the absence of physical priors. However, given that our phenomenological model takes BSDs as input for power-spectrum predictions, a complementary approach would be to examine the extent to which power spectrum measurements can place constraints on BSDs.

Consider a BSD that is given by our fiducial lognormal form plus an arbitrary perturbation in each $\ln(R_b)$ bin. Such a model possesses a large number of free parameters – one perturbation amplitude per discrete $\ln(R_b)$ bin – and thus the expectation is that a single power spectrum will not likely be able to place meaningful constraints on every degree of freedom of the BSD. Given this limitation, we

[7] https://github.com/j pober/21cmSense

[8] As a point of comparison, the variance of the CMB temperature power spectrum C_ℓ scales as $\sigma^2(C_\ell) \propto (C_\ell + N_\ell)^2$ Kamionkowski, Kosowsky & Stebbins (1997), scaling the same way as our result here.

Downloaded from https://academic.oup.com/mnras/article/514/2/2010/6595978 by Scuola Normale Superiore. Biblioteca user on 03 December 2022

instead perform a principal component analysis to quantify a few shapes in the BSD that can be easily probed using power spectrum observations. As our starting point, we compute a Fisher matrix to quantify the information content on each perturbation amplitude that is contained in the power spectrum. This is given by

$$F_{\alpha\beta} = \sum_i \frac{1}{\sigma^2_{\mathrm{full},i}} \frac{\partial \Delta^2(k_i)}{\partial \eta_\alpha} \frac{\partial \Delta^2(k_i)}{\partial \eta_\beta}, \tag{57}$$

where $\sigma_{\mathrm{full},i}$ is the error bar on a measurement of the ith bin of the power spectrum (as described in Section 4.1) and η_α denotes the value of $\mathrm{d}n_b/\mathrm{d}R_b$ in the αth $\ln(R_b)$ bin. It is understood that the derivatives are to be evaluated at a set of fiducial values for any free parameters in the model. Performing an eigen decomposition of this Fisher matrix provides a set of eigenvectors that serves as a series of orthonormal basis templates for the BSD. Fig. 12 provides some example eigenvectors and eigenvalues. The top panel shows the inverse square root of the eigenvalues, which quantify the error in a potential inference of the amplitude of each template from a power spectrum measurement. The templates are therefore ordered in the sense that they represent a set of modes ordered from most precisely measurable to least precisely measurable. The corresponding eigenvectors are shown in the bottom three panels, along with the fiducial BSD. Note that in this section, we plot $\mathrm{d}n_b/\mathrm{d}R_b$ rather than $Q^{-1}V_b\mathrm{d}n_b/\mathrm{dlog}\,R_b$, since it is the former that is an input in the MICRO21CM package. Since the eigenvectors are normalized, the errors in the top panel can be directly compared to the typical amplitude of $\mathrm{d}n_b/\mathrm{d}R_b$. The dotted horizontal line in the top panel indicates the peak value of the BSD, providing a rough sense for what modes might be measurable.

Several trends are immediately apparent. First, we note that all the eigenvectors are relatively localized to the high end of the BSD. Unsurprisingly, the scales over which these template modes have appreciable amplitude is in rough correspondence to the scales that we probe using the power spectrum. Secondly, we see that the modes are essentially tapered Fourier-like modes. Each mode measures finer and finer details of the upper end of the BSD. However, from the top panel we see that only the first few are likely to be measurable – far fewer than the 1000 bins in $\ln(R_b)$ that were perturbable in our analysis. We thus conclude that while the power spectrum certainly does have some sensitivity to the BSD, the information that can be extracted about it may be limited. Of course, these results do depend on the fiducial BSD used in the analysis, and we find from our experimentation that although our qualitative trends seem fairly robust, the details can vary substantially.

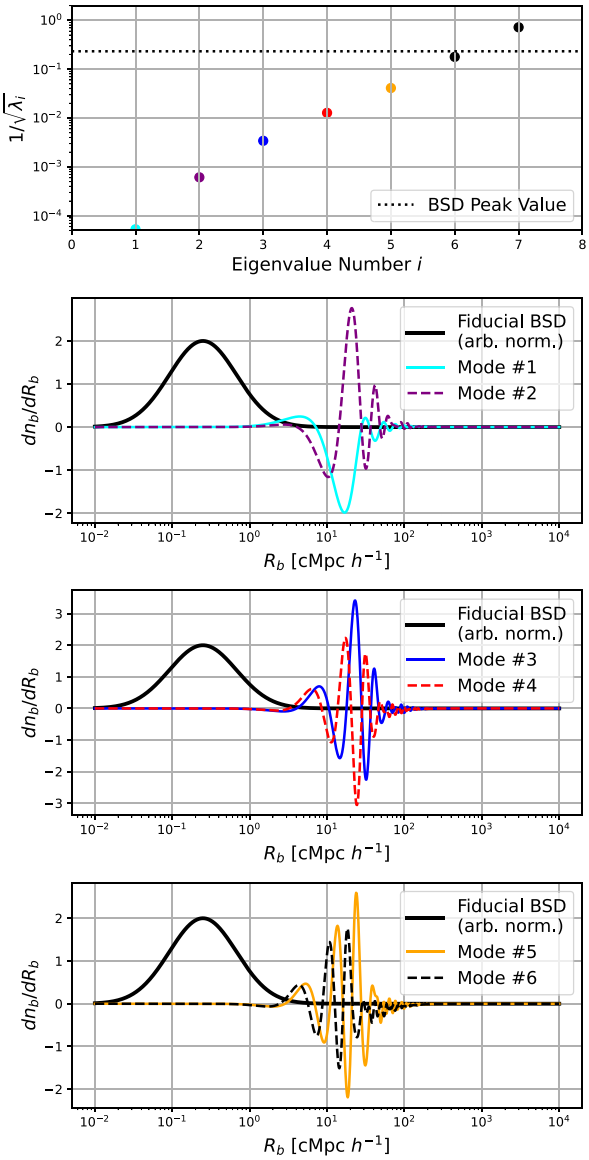

Figure 12. Principal components modes of accessible modes to power-spectrum constraints on the bubble size distribution (BSD). *Top*: Inverse square root of the eigenvalue spectrum of the principal components, which provides a sense for the measurability of each principal component. *Next three*: First six principal components of the BSD, representing the most measurable perturbations on a fiducial lognormal bubble BSD (solid black). Note that in contrast to Fig. 1, here we plot $\mathrm{d}n_b/\mathrm{d}R_b$ rather than $Q^{-1}V_b\mathrm{d}n_b/d\log R_b$.

4.3 Fits to 21cmFAST mocks

We now test the phenomenological model's ability to recover the mean IGM properties of mocks generated with 21cmFAST. While seminumerical models are known to produce 21-cm power spectra in reasonably good agreement with radiative transfer simulations, at least in the saturated limit (Zahn et al. 2011; Hutter 2018), to our knowledge, there has yet to be an attempt to compare analytical and seminumerical models of reionization or to recover the inputs of one model from the outputs of another. Once again, we expect the phenomenological model to struggle most when temperature fluctuations or partial ionization are important in the input 21cmFAST mocks, and/or late in reionization when the BSD is an increasingly poor descriptor of the ionization field. We plan to experiment with a non-uniform temperature field in the future, as there is evidence that semi-analytical and seminumerical techniques agree fairly well before reionization (see Schneider et al. 2021). However, for now we proceed with our model as-is, in order to establish a baseline for potential development in the future.

As in Sections 4.1 and 4.2, we use uncertainties generated with 21CMSENSE for HERA with moderate foregrounds, which provides a fiducial benchmark for the state of 21-cm observations over the next few years, and explore all eight 21cmFAST models introduced thus far (in Section 3 and Fig. 4): four different X-ray heating scenarios for each reionization scenario. We fit all redshift snapshots between $6 \lesssim z \lesssim 10$, including only modes in the range $0.1 \leq k/[h\,\mathrm{Mpc}^{-1}] \leq 1$.

One could in principle let R, σ, Q, and T_S vary independently at every redshift of a multi-epoch fit, though this will of course result

Downloaded from https://academic.oup.com/mnras/article/514/2/2010/6595978 by Scuola Normale Superiore. Biblioteca user on 03 December 2022

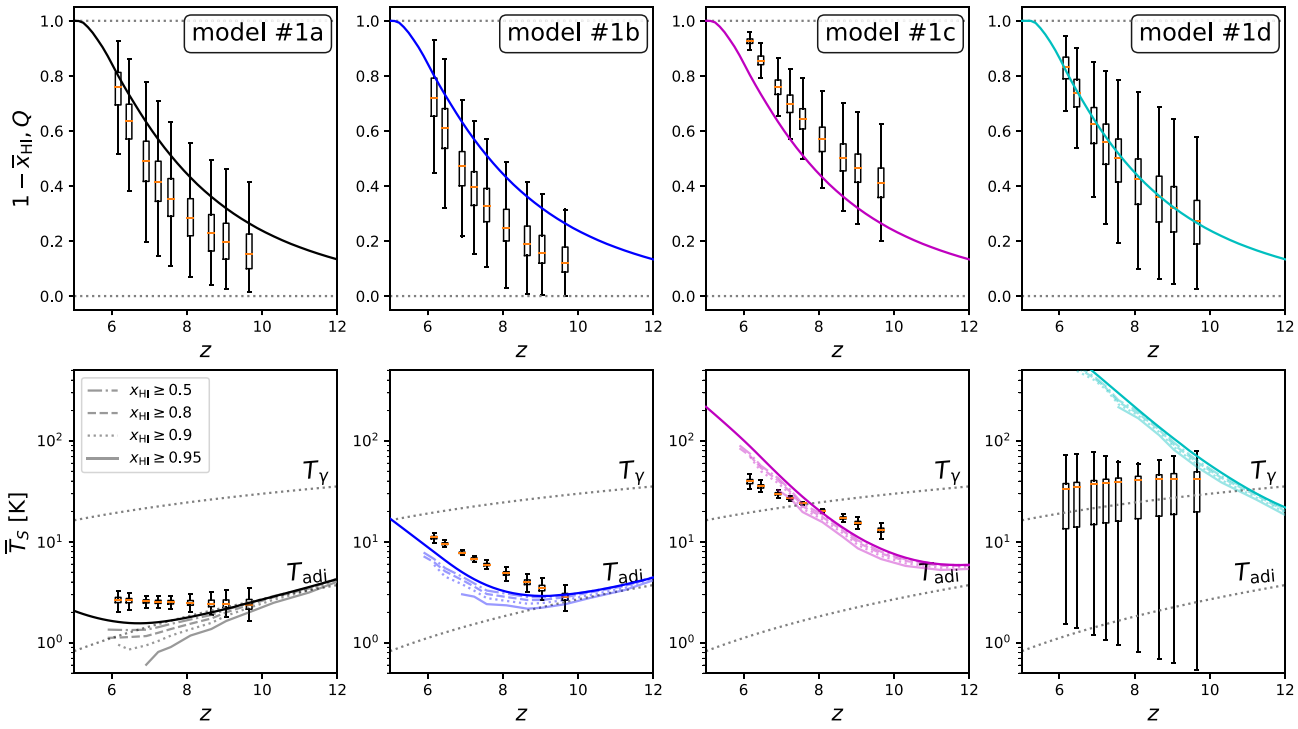

Figure 13. Recovery of mean ionization (top) and spin temperature (bottom) histories from 21cmFAST mocks, for each 'slow' reionization model. Lines show true evolution extracted from 21cmFAST directly, while recoveries are shown with error-bars, the full extent of which indicate the 95 per cent confidence region, and the boxes the 68 per cent confidence interval. Linestyles in the bottom row indicate the spin temperature averaged over voxels above different neutral fraction thresholds, as indicated in the legend. There are clear biases in both quantities, though the thermal histories are constrained well enough to associate each mock with the correct level of X-ray heating (see Section 4.3).

in a very high dimensional model. In addition, single-epoch fits only loosely constrain the ionized fraction (see Fig. 11), since they lack even simple priors, e.g. that Q increase from high redshift to low. Motivated by Fig. 6, we employ simple functions for the evolution of each parameter in subsequent MCMC fits. For R, we employ a power law in Q, which we found performed more efficiently than a power law in z. We further parametrize Q and $T_{\rm S}$ as power laws in redshift and assume σ is a constant, which results in a total of 7 free parameters. This allows us to compare to 21cmFAST models over a range of redshifts, and fit to multi-epoch mock data sets without dramatically increasing the dimensionality of the fit.

In the future, large databases of seminumeric models (e.g. Prelogovic et al., in preparation), could be used to map out a prior volume in (R, σ, Q) space in order to avoid unphysical regions and so reduce uncertainties on Q and $T_{\rm S}$. Here, however, we assume broad, uninformative priors on each parameter. We take:

(i) $Q \in [0, 1]$ and $\mathrm{dlog}\, Q/\mathrm{dlog}\, z = [-20, 0]$.
(ii) $\log_{10} T_{\rm S}/\mathrm{K} \in [-1, 3]$ and $\mathrm{dlog}\, T_{\rm S}/\mathrm{dlog}\, z \in [-20, 5]$.
(iii) $R(Q = 0.5)/\mathrm{cMpc} \in [0.5, 50]$ and $\mathrm{dlog}\, R/\mathrm{dlog}\, Q \in [0, 5]$.
(iv) $\sigma \in [0.25, 2.5]$.

We also enforce $Q \geq 0.99$ at $z = 5.3$, a conservative end-of-reionization prior consistent with the latest interpretation of Lyα forest constraints (Becker et al. 2015; Keating et al. 2020; Bosman et al. 2021).

Before we show the results of these fits, we note that each reionization scenario spans a different range in *ionized fraction* over the same interval in *redshift* used for fitting. For example, the IGM in the 'slow' reionization model is already $\simeq$30 per cent ionized by $z \sim 10$, the highest redshift used in the fit, while the 'fast' reionization scenario has $Q \lesssim 0.1$ at $z \simeq 10$ (see Fig. 4). As a result, we may see different outcomes in the fits given that our model fares poorly at the very end of reionization, and there are effectively more data points late in reionization for the 'slow' scenario.

In Figs 13 and 14, we show our recovery of the mean ionization history (top) and spin temperature history (bottom), for 'slow' and 'fast' reionization models, respectively. In each plot, we show results for increasingly efficient X-ray heating going from left to right.

Starting first with the 'slow' models, two things are clear immediately at a glance: (i) the recovered ionization histories have large uncertainties and are generally biased low, and (ii) the recovered thermal histories, while also slightly biased, are correct in order of L_X/SFR, and so are good enough to identify the appropriate heating scenario. However, for the most efficient X-ray heating scenario (model 1d; right-most columns), uncertainties on $T_{\rm S}$ grow much larger, as we are increasingly insensitive to $T_{\rm S}$ once it becomes large. This is in some sense a good thing, i.e. huge errors on $T_{\rm S}$ but *not* on Q may alone indicate $T_{\rm S} \gg T_\gamma$.

The 'fast' reionization scenario in Fig. 14 shows the same general trends in its recovery of the thermal history. However, here, the early stages of reionization are recovered much more accurately, while it is the later stages that suffer from more of a bias. This is at least in part caused by our decision to fit both models over the range $6 \lesssim z \lesssim 10$. For the 'fast' models, the full rise in 21-cm power from $Q \simeq 0$ to $Q \simeq 0.5$ lies within the fitted range, but the latter half of reionization does not. Because the mock contains the part of reionization that the

Downloaded from https://academic.oup.com/mnras/article/514/2/2010/6595978 by Scuola Normale Superiore. Biblioteca user on 03 December 2022

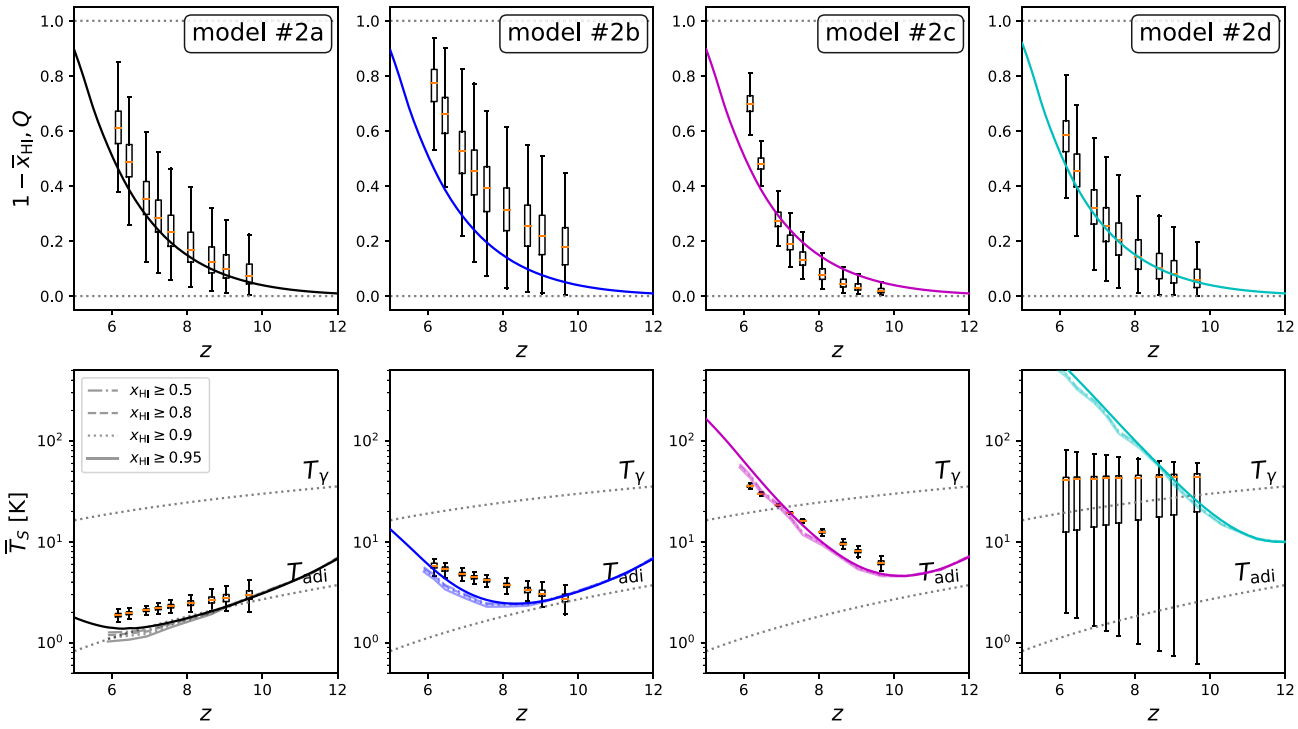

Figure 14. Recovery of mean ionization (top) and spin temperature (bottom) histories from 21cmFAST mocks, for each 'fast' reionization model. Lines show true evolution extracted from 21cmFAST directly, while recoveries are shown with error-bars, the full extent of which indicate the 95 per cent confidence region, and the boxes the 68 per cent confidence interval. Linestyles in the bottom row indicate the spin temperature averaged over voxels above different neutral fraction thresholds, as indicated in the legend. Though some biases are visible in both quantities, the early history of reionization is better constrained than that of the 'slow' scenario (Fig. 13), likely because the first half of reionization lies in the fit range, $6 \lesssim z \lesssim 10$. Once again, the thermal histories are biased, but are constrained well enough to associate each mock with the correct level of X-ray heating (see Section 4.3).

model predicts best, a more reliable fit is obtained. In contrast, the 'slow' scenario covers the middle ∼50 per cent of reionization, so the fit cannot leverage the model's accuracy at $Q \lesssim 0.3$.

There are other factors at play, aside from the fitted redshift range, that could explain the biases in Figs 13 and 14. For example, as discussed previously, because our model overestimates ionization fluctuations we should expect to underestimate the contrast. This is clearly the case in the scenarios with minimal X-ray heating – the recovered $T_{\rm S}$ is higher than the true $T_{\rm S}$ once reionization gets underway. In the final scenario, with very efficient X-ray heating, the exact temperature is constrained very poorly (last column). However, once the temperature is $T_{\rm S} \gg T_\gamma$, the 21-cm background is of course increasingly insensitive to $T_{\rm S}$. The bias in our recovery is in part caused once again by our overestimation of ionization fluctuations, except now, in the emission regime, $T_{\rm S}$ must be reduced to preserve large-scale 21-cm power as Q rises.

The other likely source of error is our neglect of temperature fluctuations and partial ionization. While our models a and d are nearly equivalent to cases with identically zero heating and full saturation, respectively, the intermediate cases b and c likely have residual temperature fluctuations to some extent during reionization. Similarly, cases c and d are most likely to have partial ionization caused by strong X-ray backgrounds. These factors could be causing the change in the recovered ionization histories as a function of L_X/SFR. We do not attempt to quantify the magnitude of these effects here, but a closer look may be warranted in future studies.

For completeness, we show example recovered 21-cm power spectra in Fig. 15. Now, the four different 'slow' reionization scenarios are shown from top to bottom, in order of increasing X-ray heating, at nine different mean ionized fractions from left to right. In most cases, the power spectrum is recovered very well, indicating that our 7-parameter model is sufficiently flexible to capture both the time evolution and shape of the 21-cm power spectrum.

5 DISCUSSION AND CONCLUSIONS

In this work, we have introduced a simple, phenomenological model for the 21-cm power spectrum during reionization that abstracts away assumptions about galaxy formation and instead works directly in terms of the bubble size distribution, mean ionized fraction, Q, and mean spin temperature, $T_{\rm S}$. The goal was to build intuition for the results of more sophisticated models like 21cmFAST, and determine if efficient, IGM-focused models like ours may be sufficiently accurate for inference in the next few years, as upper limits from, e.g. MWA, HERA, and LOFAR continue to improve. To our knowledge, this is the first attempt to recover the IGM properties from 21-cm power spectra generated with 21cmFAST using a completely different model (though see Zhou & La Plante 2021, for a similarly motivated analysis in the 21-cm imaging context).

We find that a binary ionization field, with a size distribution of bubbles given by simple lognormal and power-law-times-exponential functions, gives rise to 21-cm power spectra in qualitative agreement with those generated by seminumeric models (see Figs 1 and 2). We provide a grid of model predictions for the large-scale 21-cm power spectrum at $z = 8$ in Fig. 3, which show that 'cold reionization' models, with $T_{\rm S} \simeq 1.8$ K (the limit corresponding to pure adiabatic cooling), can be disfavoured by upper limits at the $\sim 10^3$ mK^2 level.

Downloaded from https://academic.oup.com/mnras/article/514/2/2010/6595978 by Scuola Normale Superiore. Biblioteca user on 03 December 2022

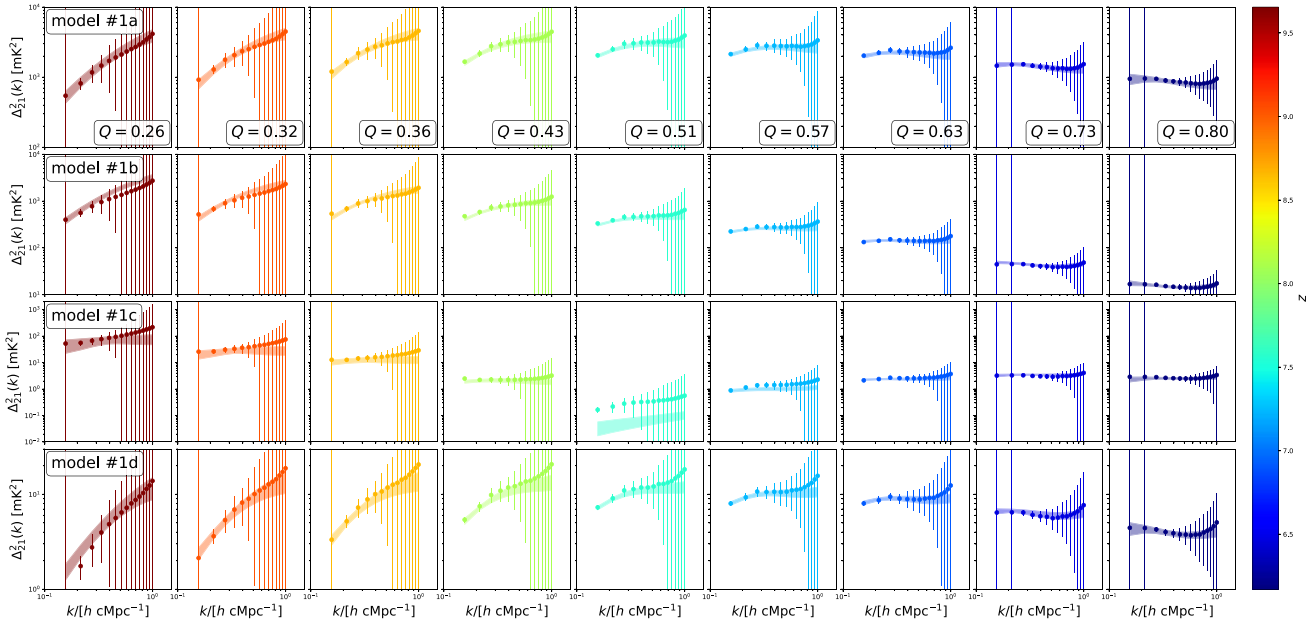

Figure 15. Recovered 21-cm power spectra from fits to 'slow' reionization 21cmFAST mocks. From top to bottom, we show models with increasing L_X/SFR, in order of increasing global ionized fraction from left to right. Shaded regions show 68 per cent confidence reconstructions, while solid lines indicate the maximum likelihood model.

Indeed, the recent The HERA Collaboration et al. (2021b) upper limits ($(30.76\ \mathrm{mK})^2$ at $z \sim 7.9$ and $k = 0.192\,h\,\mathrm{Mpc}^{-1}$) imply spin temperatures in excess of the adiabatic limit, as found in four independent analyses, including an analytical bias model, our phenomenological model, and two semi-numerical models (The HERA Collaboration 2021a). Our model predicts that an order-of-magnitude improvement, resulting in upper limits of $\sim 10^2$ mK^2 at $z \simeq 8$, would drive lower limits on the spin temperature securely to $T_S \gtrsim 10$ K. This is an important milestone, as $T_S \sim 10$ K is expected in models anchored to galaxy luminosity functions (Mirocha, Furlanetto & Sun 2017; Park et al. 2019), if one assumes there is no evolution in the efficiency of X-ray production (see also fig. 7–8 in The HERA Collaboration 2021a). As a result, ~ 100 mK^2 limits would substantiate expectations of redshift evolution in the L_X–SFR relation due to declining metallicities in galaxies at high redshift (Fragos et al. 2013; Brorby et al. 2016).

We also performed several forecasts, first a single-epoch fit to a mock 21-cm power spectrum generated with the phenomenological model (Section 4.1), followed by a multi-epoch fit to mocks generated with 21cmFAST (Section 4.3).

When fitting mock signals generated with the phenomenological model, we do recover the input model parameters, though uncertainties are generally large (see Fig. 11). Simultaneously fitting data at multiple redshifts is thus vital to obtaining tight constraints on model parameters, e.g. the mean ionized fraction. The IGM spin temperature, T_S, is an exception here, at least for strong signals. The detailed shape of the BSD is likely beyond the reach of current experiments, though, encouragingly, the power spectrum is sensitive to the typical bubble size and the distribution of sizes just above the peak (see Section 4.2).

We also fit a total of eight 21cmFAST mocks: a 'slow' and 'fast' reionization scenario with four different X-ray heating scenarios for each. Because our model *overestimates* the ionization fluctuations (see Figs 5–6), we generally *underestimate the contrast*, $1 - T_\gamma/T_S$. As a result, for cold reionization scenarios, the phenomenological model yields slightly higher temperatures than are assumed by the input mock, and if heating drives $T_S \gg T_\gamma$, we instead obtain T_S values that are biased low. In detail, the recovered ionization and thermal histories are underwhelming in some parts of parameter space (Figs 13 and 14). However, one can reliably place recovered thermal histories into broad categories (no heat and low/medium/high heat), a triumph for such a simple, phenomenological approach.

Finally, we note that our model adopts several key simplifying assumptions, improvements to which may bring our approach into closer agreement with semi-numeric models. For example, we assume:

(i) Bubbles are fully ionized, perfectly spherical, and have perfectly sharp edges, which results in an overestimate of fluctuations in the ionization field (see Section 3 and Figs 5–6).
(ii) The degree of bubble overlap is estimated (see Section 2.3.1) but we do not attempt to correct for overlap effects.
(iii) The density field is assumed to mirror the ionization field, which allows a simple approach to computing cross-terms (see Sections 2.4–2.4.1) and the mean density of ionized regions (see Section 2.4.2 and Fig. 7), but results in biases (see Fig. 8).
(iv) Though we focus entirely on ionized bubbles, one can instead treat bubbles of a fixed temperature with the same formalism by changing the parameter α (see Section 2). However, we do not allow both kinds of fluctuations to operate simultaneously in this work, and in some cases our ability to recover IGM properties from 21cmFAST models may suffer as a result.

With a run-time of less than a second per redshift, we can afford improvements to the model, even if they come with a non-trivial penalty in computational efficiency. The assumptions listed above may thus be prime targets for improving the model, in an attempt to reduce biases in IGM constraints derived from our phenomenological approach, and perhaps even tighten constraints over a broader range of parameter space (e.g. the $T_S \gg T_\gamma$ regime). We defer an exploration of potential improvements to future work, and in the

Downloaded from https://academic.oup.com/mnras/article/514/2/2010/6595978 by Scuola Normale Superiore. Biblioteca user on 03 December 2022

meantime welcome additions and/or improvements to the code at https://github.com/mirochaj/micro21cm.

Software: NUMPY (Van Der Walt, Colbert & Varoquaux 2011), SCIPY (Virtanen et al. 2020), MATPLOTLIB (Hunter 2007), EMCEE (Foreman-Mackey et al. 2013), CAMB (Lewis et al. 2000).

ACKNOWLEDGEMENTS

The authors acknowledge Adélie Gorce, Paul La Plante, Stefan Heimersheim, Anastasia Fialkov, and Yuxiang Qin for helpful feedback and encouragement throughout this effort. This material is based upon work supported by the National Science Foundation under Grant No. 1636646, the Gordon and Betty Moore Foundation, and institutional support from the HERA collaboration partners. 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 Technology. JBM is supported by a Clay fellowship at the Smithsonian Astrophysical Observatory. SRF was supported by the National Science Foundation through award AST-1812458. and was directly supported by the NASA Solar System Exploration Research Virtual Institute cooperative agreement number 80ARC017M0006. AL acknowledges support from the New Frontiers in Research Fund Exploration grant program, a Natural Sciences and Engineering Research Council of Canada (NSERC) Discovery Grant and a Discovery Launch Supplement, a Fonds de recherche Nature et echnologies Quebec New Academics grant, the Sloan Research Fellowship, the William Dawson Scholarship at McGill, as well as the Canadian Institute for Advanced Research (CIFAR) Azrieli Global Scholars program. Computations were made on the supercomputer Cedar at Simon Fraser University managed by Compute Canada. The operation of this supercomputer is funded by the Canada Foundation for Innovation (CFI).

DATA AVAILABILITY

The data underlying this article is available upon request.

REFERENCES

Barkana R., 2018, Nature, 555, 71
Barkana R., Loeb A., 2005a, ApJ, 624, L65
Barkana R., Loeb A., 2005b, MNRAS, 363, L36
Barry N. et al., 2019, ApJ, 884, 1
Becker G. D., Bolton J. S., Madau P., Pettini M., Ryan-Weber E. V., Venemans B. P., 2015, MNRAS, 447, 3402
Berlin A., Hooper D., Krnjaic G., McDermott S. D., 2018, Phys. Rev. Lett., 121, 011102
Bi H., Davidsen A. F., 1997, ApJ, 479, 523
Bianco M., Iliev I. T., Ahn K., Giri S. K., Mao Y., Park H., Shapiro P. R., 2021, MNRAS, 504, 2443
Boddy K. K., Gluscevic V., Poulin V., Kovetz E. D., Kamionkowski M., Barkana R., 2018, Phys. Rev. D, 98, 123506
Bosman S. E. I. et al., 2021, preprint(arXiv:2108.03699)
Bouwens R. J., Illingworth G. D., Oesch P. A., Caruana J., Holwerda B., Smit R., Wilkins S., 2015, ApJ, 811, 140
Bowman J. D., Rogers A. E. E., 2010, Nature, 468, 796
Bowman J. D., Rogers A. E. E., Monsalve R. A., Mozdzen T. J., Mahesh N., 2018, Nature, 555, 67
Bradley R. F., Tauscher K., Rapetti D., Burns J. O., 2019, ApJ, 874, 153
Brorby M., Kaaret P., Prestwich A., Mirabel I. F., 2016, MNRAS, 457, 4081
Burns J. O. et al., 2017, ApJ, 844, 33
Busch P., Eide M. B., Ciardi B., Kakiichi K., 2020, MNRAS, 498, 4533
Chapman E., Santos M. G., 2019, MNRAS, 490, 1255
Cohen A., Fialkov A., Barkana R., Lotem M., 2017, MNRAS, 472, 1915
Coles P., Jones B., 1991, MNRAS, 248, 1
Cooray A., Sheth R., 2002, Phys. Rep., 372, 1
Das A., Mesinger A., Pallottini A., Ferrara A., Wise J. H., 2017, MNRAS, 469, 1166
Datta A., Bowman J. D., Carilli C. L., 2010, ApJ, 724, 526
Datta K. K., Mellema G., Mao Y., Iliev I. T., Shapiro P. R., Ahn K., 2012, MNRAS, 424, 1877
DeBoer D. R. et al., 2017, PASP, 129, 045001
de Lera Acedo E., 2019, in 2019 International Conference on Electromagnetics in Advanced Applications (ICEAA). IEEE, p. 0626
Diemer B., 2018, ApJS, 239, 35
Eastwood M. W. et al., 2019, AJ, 158, 84
Ewall-Wice A., Hewitt J., Mesinger A., Dillon J. S., Liu A., Pober J., 2016, MNRAS, 458, 2710
Ewall-Wice A., Chang T.-C., Lazio J., Doré O., Seiffert M., Monsalve R. A., 2018, ApJ, 868, 63
Feng C., Holder G., 2018, ApJ, 858, L17
Fialkov A., Barkana R., 2019, MNRAS, 486, 1763
Fialkov A., Barkana R., Pinhas A., Visbal E., 2014, MNRAS, 437, L36
Fialkov A., Barkana R., Cohen A., 2018, Phys. Rev. Lett., 121, 011101
Finkelstein S. L. et al., 2015, ApJ, 810, 71
Finkelstein S. L. et al., 2019, ApJ, 879, 36
Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013, PASP, 125, 306
Fragos T. et al., 2013, ApJ, 764, 41
Fraser S. et al., 2018, Phys. Lett. B, 785, 159
Furlanetto S. R., Oh S. P., 2016, MNRAS, 457, 1813
Furlanetto S. R., Zaldarriaga M., Hernquist L., 2004, ApJ, 613, 1 (FZH04)
Furlanetto S. R., Oh S. P., Briggs F. H., 2006, Phys. Rep., 433, 181
Garsden H. et al., 2021, MNRAS, 506, 5802
Gehlot B. K. et al., 2019, MNRAS, 488, 4271
Gehlot B. K. et al., 2020, MNRAS, 499, 4158
Georgiev I., Mellema G., Giri S. K., Mondal R., 2022, MNRAS, 513, 5109
Ghara R., Choudhury T. R., Datta K. K., 2015, MNRAS, 447, 1806
Ghara R., Mellema G., Giri S. K., Choudhury T. R., Datta K. K., Majumdar S., 2018, MNRAS, 476, 1741
Ghara R. et al., 2020, MNRAS, 493, 4728
Ghara R., Giri S. K., Ciardi B., Mellema G., Zaroubi S., 2021, MNRAS, 503, 4551
Gnedin N. Y., 2014, ApJ, 793, 29
Goodman J., Weare J., 2010, Commun. Appl. Math. Comput. Sci., 5, 65
Gorce A., Douspis M., Aghanim N., Langer M., 2018, A&A, 616, A113
Greig B., Mesinger A., 2015, MNRAS, 449, 4246
Greig B., Mesinger A., 2018, MNRAS, 477, 3217
Greig B., Mesinger A., Pober J. C., 2016, MNRAS, 455, 4295
Greig B. et al., 2021a, MNRAS, 501, 1
Greig B. et al., 2021b, MNRAS, 500, 5322
Haiman Z., Rees M. J., Loeb A., 1997, ApJ, 476, 458
Hamilton A. J. S., 2000, MNRAS, 312, 257
Hills R., Kulkarni G., Meerburg P. D., Puchwein E., 2018, Nature, 564, E32
Hunter J. D., 2007, Comput. Sci. Eng., 9, 90
Hutter A., 2018, MNRAS, 477, 1549
Jensen H. et al., 2013, MNRAS, 435, 460
Kaiser N., 1987, MNRAS, 227, 1
Kakiichi K. et al., 2017, MNRAS, 471, 1936
Kamionkowski M., Kosowsky A., Stebbins A., 1997, Phys. Rev. D, 55, 7368
Kannan R., Garaldi E., Smith A., Pakmor R., Springel V., Vogelsberger M., Hernquist L., 2022, MNRAS, 511, 4005
Kaurov A. A., 2016, ApJ, 831, 198
Keating L. C., Weinberger L. H., Kulkarni G., Haehnelt M. G., Chardin J., Aubert D., 2020, MNRAS, 491, 1736
Kern N. S., Liu A., Parsons A. R., Mesinger A., Greig B., 2017, ApJ, 848, 23
Kovetz E. D., Poulin V., Gluscevic V., Boddy K. K., Barkana R., Kamionkowski M., 2018, Phys. Rev. D, 98, 103529
La Plante P., Battaglia N., Natarajan A., Peterson B. J., Trac H., Cen R., Loeb A., 2014, ApJ, 789, 31
Lewis A., Challinor A., Lasenby A., 2000, ApJ, 538, 473

Downloaded from https://academic.oup.com/mnras/article/514/2/2010/6595978 by Scuola Normale Superiore. Biblioteca user on 03 December 2022

Lidz A., Zahn O., McQuinn M., Matias Z., Dutta S., Hernquist L., 2007, ApJ, 659, 865
Lin Y., Peng Oh. S., Furlanetto S. R., Sutter M. P., 2016, MNRAS, 461, 3361
Liu A., Parsons A. R., 2016, MNRAS, 457, 1864
Liu A., Shaw J. R., 2020, PASP, 132, 062001
Liu A., Parsons A. R., Trott C. M., 2014, Phys. Rev. D, 90, 023018
Liu A., Pritchard J. R., Allison R., Parsons A. R., Seljak U., Sherwin B. D., 2016, Phys. Rev. D, 93, 043013
Machacek M. E., Bryan G. L., Abel T., 2001, ApJ, 548, 509
Madau P., Meiksin A., Rees M. J., 1997, ApJ, 475, 429
Mao Y., Shapiro P. R., Mellema G., Iliev I. T., Koda J., Ahn K., 2012, MNRAS, 422, 926
Mason C. A., Trenti M., Treu T., 2015, ApJ, 813, 21
Mason C. A., Naidu R. P., Tacchella S., Leja J., 2019, MNRAS, 489, 2669
McQuinn M., D'Aloisio A., 2018, J. Cosmol. Astropart. Phys., 2018, 016
McQuinn M., Zahn O., Zaldarriaga M., Hernquist L., Furlanetto S. R., 2006, ApJ, 653, 815
Mertens F. G. et al., 2020, MNRAS, 493, 1662
Mesinger A., Furlanetto S., 2007, ApJ, 669, 663
Mesinger A., Furlanetto S., Cen R., 2011, MNRAS, 411, 955
Mineo S., Gilfanov M., Sunyaev R., 2012, MNRAS, 419, 2095
Mirocha J., Furlanetto S. R., 2019, MNRAS, 483, 1980
Mirocha J., Harker G. J. A., Burns J. O., 2013, ApJ, 777, 118
Mirocha J., Harker G. J. A., Burns J. O., 2015, ApJ, 813, 11
Mirocha J., Furlanetto S. R., Sun G., 2017, MNRAS, 464, 1365
Mondal R. et al., 2020, MNRAS, 498, 4178
Morales M. F., Wyithe J. S. B., 2010, ARA&A, 48, 127
Morales M. F., Hazelton B., Sullivan I., Beardsley A., 2012, ApJ, 752, 137
Muñoz J. B., 2019a, Phys. Rev. D, 100, 063538
Muñoz J. B., 2019b, Phys. Rev. Lett., 123, 131301
Muñoz J. B., Loeb A., 2018, Nature, 557, 684
Muñoz J. B., Ali-Haïmoud Y., Kamionkowski M., 2015, Phys. Rev. D, 92, 083508
Muñoz J. B., Kovetz E. D., Ali-Haïmoud Y., 2015, Phys. Rev. D, 92, 083528
Muñoz J. B., Dvorkin C., Loeb A., 2018, Phys. Rev. Lett., 121, 121301
Muñoz J. B., Bohr S., Cyr-Racine F.-Y., Jesús Z., Vogelsberger M., 2021a, Phys. Rev. D, 103, 043512
Muñoz J. B., Qin Y., Mesinger A., Murray S. G., Greig B., Mason C., 2022, MNRAS, 511, 3657
Murray S., Greig B., Mesinger A., Muñoz J., Qin Y., Park J., Watkinson C., 2020, J. Open Source Softw., 5, 2582
Naoz S., Barkana R., 2005, MNRAS, 362, 1047
Ocvirk P. et al., 2020, MNRAS, 496, 4087
Paciga G. et al., 2013, MNRAS, 433, 639
Pagano M., Liu A., 2020, MNRAS, 498, 373
Pagano M., Liu A., 2021, MNRAS, 508, 1915
Paranjape A., Choudhury T. R., 2014, MNRAS, 442, 1470
Paranjape A., Choudhury T. R., Padmanabhan H., 2016, MNRAS, 460, 1801
Park J., Mesinger A., Greig B., Gillet N., 2019, MNRAS, 484, 933
Parsons A., Pober J., McQuinn M., Jacobs D., Aguirre J., 2012, ApJ, 753, 81
Parsons A. R. et al., 2014, ApJ, 788, 106
Patil A. H. et al., 2017, ApJ, 838, 65
Philip L. et al., 2019, J. Astron. Instrum., 8, 1950004
Pober J. C., 2015, MNRAS, 447, 1705
Pober J. C. et al., 2013, AJ, 145, 65
Pober J. C. et al., 2014, ApJ, 782, 66
Pober J. C. et al., 2015, ApJ, 809, 62
Pospelov M., Pradler J., Ruderman J. T., Urbano A., 2018, Phys. Rev. Lett., 121, 031103
Pritchard J. R., Furlanetto S. R., 2007, MNRAS, 376, 1680
Pritchard J. R., Loeb A., 2012, Rep. Prog. Phys., 75, 086901
Qin Y., Mesinger A., Park J., Greig B., Muñoz J. B., 2020, MNRAS, 495, 123
Qin Y., Mesinger A., Bosman S. E. I., Viel M., 2021, MNRAS, 506, 2390
Raste J., Sethi S., 2018, ApJ, 860, 55
Reis I., Fialkov A., Barkana R., 2020, MNRAS, 499, 5993
Robertson B. E., Ellis R. S., Furlanetto S. R., Dunlop J. S., 2015, ApJ, 802, L19
Rosdahl J. et al., 2018, MNRAS, 479, 994
Santos M. G., Amblard A., Pritchard J., Trac H., Cen R., Cooray A., 2008, ApJ, 689, 1
Santos M. G., Ferramacho L., Silva M. B., Amblard A., Cooray A., 2010, MNRAS, 406, 2421
Schneider A., Giri S. K., Mirocha J., 2021, Phys. Rev. D, 103, 083025
Shaver P. A., Windhorst R. A., Madau P., de Bruyn A. G., 1999, A&A, 345, 380
Sims P. H., Pober J. C., 2020, MNRAS, 492, 22
Singh S., Subrahmanyan R., 2019, ApJ, 880, 26
Singh S. et al., 2017, ApJ, 845, L12
Singh S. et al., 2022, Nature Astronomy, 6, 607
Sobacchi E., Mesinger A., 2014, MNRAS, 440, 1662
Talman J. D., 1978, J. Comput. Phys., 29, 35
Tauscher K., Rapetti D., Burns J. O., 2020, ApJ, 897, 132
The HERA Collaboration et al., 2022a, ApJ, 925, 221
The HERA Collaboration et al., 2022b, ApJ, 924, 51
Thomas R. M. et al., 2009, MNRAS, 393, 32
Thyagarajan N. et al., 2013, ApJ, 776, 6
Tingay S. J. et al., 2013, Publ. Astron. Soc. Aust., 30, e007
Tseliakhovich D., Hirata C., 2010, Phys. Rev. D, 82, 083520
van der Walt S., Colbert S. C., Varoquaux G., 2011, Comput. Sci. Eng., 13, 22
van Haarlem M. P. et al., 2013, A&A, 556, A2
Vedantham H., Udaya Shankar N., Subrahmanyan R., 2012, ApJ, 745, 176
Virtanen P. et al., 2020, Nat. Methods, 17, 261
Visbal E., Barkana R., Fialkov A., Tseliakhovich D., Hirata C. M., 2012, Nature, 487, 70
Visbal E., Haiman Z., Terrazas B., Bryan G. L., Barkana R., 2014, MNRAS, 445, 107
Yoshiura S. et al., 2021, MNRAS, 505, 4775
Zahn O., Mesinger A., McQuinn M., Trac H., Cen R., Hernquist L. E., 2011, MNRAS, 414, 727
Zhou Y., La Plante P., 2022, PASP, 134, 044001

This paper has been typeset from a TEX/LATEX file prepared by the author.

Downloaded from https://academic.oup.com/mnras/article/514/2/2010/6595978 by Scuola Normale Superiore. Biblioteca user on 03 December 2022