PHYSICAL REVIEW D **89,** 083010 (2014)

Scale-dependent bias in the baryonic-acoustic-oscillation-scale intergalactic neutral hydrogen

Andrew Pontzen*

Department of Physics and Astronomy, University College London, Gower Street, London WC1E 6BT, United Kingdom

(Received 3 February 2014; published 30 April 2014)

I discuss fluctuations in the neutral hydrogen density of the $z \approx 2.3$ intergalactic medium and show that their relation to cosmic overdensity is strongly scale dependent. This behavior arises from a linearized version of the well-known "proximity effect," in which bright sources suppress atomic hydrogen density. Using a novel, systematic and detailed linear-theory radiative-transfer calculation, I demonstrate how HI density consequently anticorrelates with total matter density when averaged on scales exceeding the Lyman-limit mean-free path. The radiative-transfer thumbprint is highly distinctive and should be measurable in the Lyman-α forest. Effects extend to sufficiently small scales to generate significant distortion of the correlation function shape around the baryon acoustic oscillation peak, although the peak location shifts only by 1.2 percent for a mean source bias of $b_j = 3$. The distortion changes significantly with b_j and other astrophysical parameters; measuring it should provide a helpful observational constraint on the nature of ionizing photon sources in the near future.

DOI: 10.1103/PhysRevD.89.083010 PACS numbers: 98.62.Ra, 98.80.−k

I. INTRODUCTION

The Lyman-α forest [1] is the imprint of the intergalactic medium (IGM)—specifically, neutral hydrogen—on the spectra of distant quasars. At high redshift the rapidly changing forest probes hydrogen reionization [2–4]; at lower redshift, the forest has a steadier ionization state and is used to trace overall matter density fluctuations [5–7]. Correlating Lyman-α fluctuations over small scales therefore places a strong constraint on modifications to the standard cold dark matter picture of structure formation [8–10]. More recently attention has turned to the large-scale forest's ability to constrain the baryon acoustic oscillation peak, providing an independent distance measurement for constraining dark energy [11–13]. In addition to probing the power spectrum in these ways, the observed $z < 5$ forest constrains the thermal state of the intergalactic medium [14,15], allowing various interesting processes to be studied (such as helium reionization [16]).

When considering the forest after reionization, it is standard practice [6,17,18] to model the IGM ionization state in the presence of a uniform background of ultraviolet (UV) photons. Direct constraints on the Lyman-α cloud temperatures [19] dictate that collisional ionization is unimportant except in systems that are dense enough to be substantially self-shielded from the radiation.

However the assumption that the UV background is uniform is known to be incorrect, since the constituent photons are actually generated by galaxies and quasars. One can distinguish two limits in which the approximation fails. First, on small scales, quasars are rare; depending on the fraction of photons they contribute (likely around 50% for $2 < z < 3$ [20,21]) they can add significant shot noise on small scales. Further fluctuations are imprinted by intrinsic variability in the IGM opacity [22]. This and related astrophysical effects have been widely investigated elsewhere [9,11,23–29] with the conclusion that, if properly accounted for, the added noise is not problematic for observational cosmology at $z < 5$. Measurements at higher redshift, during the epoch of reionization, will be affected more strongly [25,27] as the UV undulation amplitude increases.

In this paper I will consider the post-reionization IGM and place more emphasis on a second failure of the uniform-radiation assumption. This appears only when source clustering is taken into account on scales around the mean-free path of an ionizing photon. By definition, regions separated by greater distances cannot efficiently exchange UV radiation. Ionization equilibrium will therefore depend on the density of sources in the local region; the bias of the forest on the largest scales will depend on the clustering of UV sources [28,30–32].

This effect has received less attention to date, probably because the relevant scale is seemingly extremely large (the mean-free path is of order 500 Mpc in comoving units [33] at $z = 2.4$). In fact, once redshifting and volume dilution are accounted for, the transition scale is somewhat smaller (more like 350 Mpc comoving; see Sec.n II A). To fully model such scales would require exceptionally large radiative-transfer simulations, with box sizes exceeding a gigaparsec to properly probe long-wavelength fluctuations.

To achieve this, previous work has employed a combination of large dark-matter-only boxes and smaller hydrodynamic simulations [30] or semianalytic prescriptions

*a.pontzen@ucl.ac.uk

[31]. In the former case the author reported a significant drop in large-scale flux power out to scales of $k^{-1} \sim 70h^{-1}$ Mpc relative to the homogeneous-radiation control case. Even so, the result has not received widespread attention. This is likely because extending such state-of-the-art work consumes a great deal of computer time and, furthermore, appropriate empirical constraints for such large separations have seemed out of reach.

The observational situation has now been changed radically by the BOSS (Baryonic Oscillation Spectroscopic Survey) project [34]. The team have released results demonstrating the viability of measuring the correlation function of Lyman-alpha clouds on large scales [7,12,35]. The major goal of BOSS is to measure the baryonic acoustic oscillation (BAO) feature in the correlation function at $100h^{-1}$ Mpc comoving. This is not so far off the reduced mean-free-path scale discussed above and derived in Sec. II A. It is timely, therefore, to reconsider the impact of large-scale fluctuations in the UV source density on the Lyman-α forest.

The remainder of this paper proceeds systematically from first principles to a detailed linear-theory calculation of these effects. This should be highly complementary to numerical studies, and motivate further work in the area. I will ignore observational questions such as the transformation from HI to flux power spectrum, redshift-space distortions, redshift evolution and flux calibration biases—since these require major computational machinery in themselves [12,35]—and focus on the bias of the physical intergalactic HI density at a single, fixed redshift. The quantitative results will be presented for $z = 2.3$, around the mean redshift of observed Lyman-α clouds [35].

The approximations that allow this calculation to be completed are (i) that the spatial variations in the UV spectrum are less important for HI than the changes in intensity (a "monochromatic approximation"); (ii) that the hydrogen can be split into a diffuse intergalactic component in photoionization equilibrium and a small population of self-shielded, collisionally ionized clumps (i.e. the highest-column-density Lyman-limit systems [36]); (iii) that nonlinear corrections (including quasar duty cycles) can be ignored on sufficiently large scales [11,27,28], although I will include shot noise from the rarity of sources; (iv) that sources averaged on large scales radiate isotropically. These seem reasonable to obtain a good estimate of the effects but in future they should be checked against numerical simulations and more complicated analytic treatments that allow for departure from equilibrium [37].

Section II develops the inhomogeneous, monochromatic radiative-transfer equations; Sec. III discusses the application of these equations to the large-scale, linear behavior of intergalactic HI. Section IV presents the main results, showing how various parameters change the distinctive imprint of radiative transfer on the intergalactic neutral hydrogen. Further discussion is given in Sec. V, especially in relation to observations of the Lyman-α forest. Two subsidiary issues are considered in Appendices. In Appendix A, I re-derive all equations using general relativity, so including peculiar velocities and inhomogeneous gravitational redshifting and elucidating the gauge dependence of the results (all of which considerations turn out to impact only on scales larger than those of interest here). Appendix B discusses the calibration of a particular parameter (the intergalactic HI bias in the absence of radiation transfer) from analytic arguments and numerical simulations.

There are a few notational matters worth settling before starting the calculation. It is helpful to be able to decompose any quantity X into its spatial mean value X_0 and fractional perturbations δ_X defined by

$$\delta_X = \frac{X - X_0}{X_0}. \tag{1}$$

Later I will mainly deal with the Fourier transform $\tilde{\delta}_X$ of these fractional variations; any quantity can be rewritten

$$\delta_X(\mathbf{x}) = \frac{1}{(2\pi)^{3/2}} \int \mathrm{d}^3 k e^{i\mathbf{k}\cdot\mathbf{x}} \tilde{\delta}_X(\mathbf{k}), \tag{2}$$

where $\mathbf{k}$ is the comoving wave vector. Finally, the power spectrum $P_X(k)$ is defined by

$$\langle \tilde{\delta}_X(\mathbf{k}')^* \tilde{\delta}_X(\mathbf{k}) \rangle = P_X(k) \delta(\mathbf{k} - \mathbf{k}'), \tag{3}$$

where angle brackets denote an ensemble average and, by an unfortunate quirk of conventional notation, the δ on the right-hand side represents the Dirac delta function. It follows from these definitions that the power spectrum for any quantity has units of a comoving volume. The expression above assumes statistical isotropy so that P_X is a function of $k = |\mathbf{k}|$ alone.

Numerical results will be derived assuming a fiducial Planck temperature-only [38] cosmology $(h, \Omega_{M0}, \Omega_{\Lambda 0}) = (0.6711, 0.3175, 0.6825)$, where Ω_{M0} and $\Omega_{\Lambda 0}$ are the present day matter and cosmological constant densities relative to critical and $h = H_0/(100\ \mathrm{km\,s^{-1}\,Mpc^{-1}})$ is the dimensionless Hubble parameter today. The main role of these quantities will be to fix the Hubble expansion rate at $z = 2.3$; any uncertainties are easily small enough to be ignored for the present study.

II. RADIATIVE TRANSFER

In this section, I will derive a monochromatic approximation to the radiative-transfer equation; this involves systematically integrating over frequency dependence. Because the scales of interest remain strongly subhorizon, relativistic corrections will be subdominant and are excluded. For the interested reader, they are reintroduced

in Appendix A which shows explicitly that they constitute a small correction.

To start, let $f(\mathbf{x},\mathbf{n},\nu)$ denote the physical number density of photons at comoving position $\mathbf{x}$ traveling in direction $\mathbf{n}$ with frequency ν. In the absence of collisional effects, the total number of photons is conserved. However the Lagrangian phase volume that those photons occupy changes over time: the spatial volume increases as a^3 while the frequency interval decreases as a, giving an overall expansion rate[1] of a^2. Overall, this implies the following Boltzmann equation:

$$\frac{\partial f}{\partial t}+\frac{c}{a}(\mathbf{n}\cdot\nabla)f+\frac{\partial f}{\partial \nu}\frac{\mathrm{d}\nu}{\mathrm{d}t}+2Hf=C[f], \tag{4}$$

where c is the speed of light, a is the universe scale factor and $H=\dot{a}/a$ is the usual Hubble expansion rate. $C[f]$ contains the collisional terms (i.e. those that alter the photon number) and will be expanded in a moment. In order, the terms on the left-hand side denote the Eulerian rate of change of photon density; the free-streaming of photons; the redshifting of the photons; and the volume dilution discussed above. The gradient operator ∇ is taken with respect to the comoving position $\mathbf{x}$ throughout this work. The term $\partial f/\partial t$ will now be set to zero, meaning I am approximating the radiation and ionization to be in equilibrium as noted in the Introduction. At the background level, this is a good approximation at $z=2.3$—the evolution of the photoionization rate Γ_0 is slow, $\mathrm{d}\ln\Gamma_0/\mathrm{d}\ln a\approx -0.04$ from the tabulations of Ref. [21]—but the implications of time dependence for perturbations should certainly be explored further in future work.

To formulate the collisional term, consider first the emission of radiation. There are two distinct relevant aspects: first, galaxies and quasars generate energy from stars and black holes; second, the intergalactic HI regenerates a fraction of photons it previously absorbed when the electron and proton recombine. I will treat these two terms separately in what follows.

Now consider absorption processes. A large effect will come from the IGM, corresponding to the low-column-density Lyman-α forest. The density of the neutral hydrogen $n_{\mathrm{HI}}(\mathbf{x})$ in this phase will be a key quantity. However, some portion (to be quantified later) of absorption comes from small, dense clumps which are strongly self-shielded against the UV radiation that is being modeled. At least three characteristics distinguish the clumped phase: first, the density of HI is determined by collisional ionization and hence essentially unaffected by variations in the radiation. Second, the majority of recombination radiation produced is reabsorbed internally within a clump. Third, the amount of radiation absorbed by such a population does not scale with the mass of HI in the population, but rather with the geometrical size and number density of the objects. For all three reasons, this population requires separate treatment.

Following the above discussion, the emission and absorption of photons can be expressed by

$$C_\nu[f]=j_\nu(\mathbf{x})+n_{\mathrm{HI}}(\mathbf{x})\left(\frac{\Gamma(\mathbf{x})}{4\pi}f_{\mathrm{rec}}(\nu,T)-c\sigma_{\mathrm{HI}}(\nu)f\right)-c\kappa_{\mathrm{clump}}(\mathbf{x},\nu)f, \tag{5}$$

where

(i) $j_\nu(\mathbf{x},\mathbf{n})$ is the emissivity per unit physical volume per frequency interval from sources other than the IGM itself;
(ii) $n_{\mathrm{HI}}(\mathbf{x})$ is the number of ground-state hydrogen atoms per unit physical volume in the IGM (excluding the shielded clumps);
(iii) $\kappa_{\mathrm{clump}}(\mathbf{x},\nu)$ is the opacity from collisionally ionized clumps;
(iv) $f_{\mathrm{rec}}(\nu,T)$ is the IGM recombination spectrum, which depends on the temperature T of the free electrons;
(v) $\Gamma(\mathbf{x})$ is the rate of ionization per HI atom (and therefore also the recombination rate, assuming photoionization equilibrium); and
(vi) $\sigma_{\mathrm{HI}}(\nu)$ is the cross section of a HI atom to ionization by a frequency ν photon.

In principle j_ν is a function of angle $\mathbf{n}$ as well as of position $\mathbf{x}$ but, in accordance with approximation (iv) above, the $\mathbf{n}$ dependence is now to be dropped (meaning that sources averaged over large scales radiate isotropically). I will also assume throughout that only HI can absorb photons in the frequency range of interest.

The cross section σ_{HI} is sharply peaked at the Lyman limit ($\nu_{\mathrm{LL}}\approx 3\times 10^{15}$ Hz), which allows for a monochromatic approach. The key quantity will be an effective number density of Lyman-limit photons, f_{LL}, defined by

$$f_{\mathrm{LL}}(\mathbf{x},\mathbf{n})=\int f(\mathbf{x},\mathbf{n},\nu)\sigma_{\mathrm{HI}}(\nu)\mathrm{d}\nu. \tag{6}$$

If desired, one can divide through by a fixed cross section [e.g. $\sigma_{\mathrm{HI}}(\nu_{\mathrm{LL}})$] to "correct" the units of f_{LL}, leading to cosmetic differences. Either way, σ_{HI} defines a single particular broadband filter that we choose to focus on; the whole framework could be formulated in terms of another band if desired. This particular choice of filter is uniquely motivated because the ionization rate per HI atom—a critical quantity of interest—is given exactly by integrating f_{LL} over all angles:

[1]Some works, e.g. Refs. [21,39], choose to use the energy density per unit volume, which leads to an a^3 volume factor and accordingly a few cosmetic differences.

$$\Gamma(\mathbf{x}) = c \iint f_{\rm LL}(\mathbf{x}, \mathbf{n}) \mathrm{d}^2 n. \quad (7)$$

To obtain the Boltzmann equation for $f_{\rm LL}$, multiply Eq. (4) by $\sigma_{\rm HI}$ and integrate with respect to ν, giving

$$a^{-1}(\mathbf{n}\cdot\nabla) f_{\rm LL} + \kappa_{\rm tot} f_{\rm LL} = c^{-1}\left(j + \bar{\sigma}_{\rm HI} n_{\rm HI} \frac{\Gamma}{4\pi} \beta_{\rm r}(T)\right), \quad (8)$$

where $\kappa_{\rm tot}$ is an effective opacity, j is an effective source emissivity and $\beta_{\rm r}(T)$ is a dimensionless, temperature-dependent fraction of recombination radiation which lies in our Lyman-limit waveband. The formal definitions of these terms arise directly from the frequency integration, and will be given and discussed below in turn.

A. Absorption

First consider the effective opacity $\kappa_{\rm tot}$ which has been composed from separate diffuse IGM opacity, clump opacity, redshifting and volume-dilution contributions:

$$\kappa_{\rm tot} = \bar{\sigma}_{\rm HI} n_{\rm HI} + \bar{\kappa}_{\rm clump} + \alpha_z \frac{H}{c} + 3\frac{H}{c}. \quad (9)$$

I have written the volume term as $3H/c$ to directly associate it with comoving volume dilution; α_z, a dimensionless number to be defined below, will contain a compensating term to return the $2H/c$ of the original formulation (4). The quantities $\bar{\sigma}_{\rm HI}$ and α_z are dependent on the spectrum, but not on the normalization of the spectrum; the monochromatic approach therefore assumes them independent of position. Their values can be estimated by using tabulated mean UV background estimates [21] at $z = 2.3$:

$$\bar{\sigma}_{\rm HI} = \frac{1}{f_{\rm LL}} \int \sigma_{\rm HI}^2 f \mathrm{d}\nu \approx 3.87 \times 10^{-18}\ \mathrm{cm}^{-2}; \quad (10)$$

$$\alpha_z = -\frac{1}{f_{\rm LL}} \int \sigma_{\rm HI} \frac{\partial f}{\partial \ln \nu} \mathrm{d}\nu - 1 \approx 1.57. \quad (11)$$

Meanwhile I have defined the monochromatic clump opacity

$$\bar{\kappa}_{\rm clump}(\mathbf{x}) = \int \kappa_{\rm clump}(\mathbf{x}, \nu) \sigma_{\rm HI}(\nu) \mathrm{d}\nu. \quad (12)$$

It will be convenient later to write the fraction of effective opacity from the respective terms as

$$\beta_{\rm HI} = \frac{\bar{\sigma}_{\rm HI} n_{\rm HI}}{\kappa_{\rm tot}}; \qquad \beta_{\rm clump} = \frac{\bar{\kappa}_{\rm clump}}{\kappa_{\rm tot}};$$
$$\beta_z = \frac{\alpha_z H}{c\kappa_{\rm tot}}; \qquad \beta_V = \frac{3H}{c\kappa_{\rm tot}}. \quad (13)$$

By definition these obey $\beta_{\rm HI} + \beta_{\rm clump} + \beta_z + \beta_V = 1$. We can estimate their values by referring to the observational constraints on Lyman-limit opacity; for instance Ref. [33] quoted a mean-free path (MFP) of $\kappa_{\rm HI}^{-1} \equiv (\bar{\sigma}_{\rm HI} n_{\rm HI} + \bar{\kappa}_{\rm clump})^{-1} \approx 150$ Mpc in physical units at $z \approx 2.4$. Their value takes into account the intergalactic medium absorption alone (it excludes volume and redshifting effects, as well as circumgalactic absorption immediately around the emitting object). Correcting to $z = 2.3$ using [33] $\lambda_{\rm MFP} \propto (1+z)^{-4.5}$ and converting to comoving units gives a helpful reference value:

$$a^{-1}\kappa_{\rm HI}^{-1} \approx 570\ \mathrm{Mpc} \quad \text{comoving at } z = 2.3. \quad (14)$$

There are uncertainties in the analysis of the observational data [40] which could imply that the correct mean-free path is somewhat longer than this value; results for different $\kappa_{\rm HI}$ will be investigated at the end of the work.

With the Planck cosmology defined in the Introduction one has $H(z=2.3)/c \approx (1280\ \mathrm{Mpc})^{-1}$ and therefore, taking the reference value of $\kappa_{\rm HI}$ above,

$$\beta_{\rm HI} + \beta_{\rm clump} = 0.62; \qquad \beta_z = 0.13; \qquad \beta_V = 0.25. \quad (15)$$

One immediate implication of these calculations is that, compared against physical opacity, redshifting and volume dilution are subdominant but important factors in lowering the cosmological density of Lyman-limit photons. This implies that the relevant scale at which scale-dependent effects are centered is smaller than the quoted HI-only mean-free path; we now have

$$a^{-1}\kappa_{\rm tot}^{-1} \approx 350\ \mathrm{Mpc} \quad \text{comoving at } z = 2.3. \quad (16)$$

This is closer to the range measurable by BOSS. In fact when solving the equations in detail below, this characteristic path length will turn out to be sufficiently short that radiation transfer can have a significant impact on the forest at the BAO scale.

Finally, one needs to decide how to assign opacity between the intergalactic HI and clumps. Sadly there is no way to do this unambiguously so I will further parametrize:

$$p_{\rm clump} = \frac{\beta_{\rm clump}}{\beta_{\rm clump} + \beta_{\rm HI}}. \quad (17)$$

Reference [33] details the fraction of opacity from systems of differing column density, allowing an estimate of $p_{\rm clump}$. For a parcel of gas to count as clumped, the definition made above Eq. (5) requires it to be in collisional ionization equilibrium. The boundary will therefore be somewhat higher than the traditional Lyman-limit system threshold because reaching the collisionally ionized state requires a reduction in photoionization rate by a substantial fraction throughout the cloud.

The details depend strongly on the temperature, density and geometry of the system itself; here I will make a very rough order-of-magnitude estimate for a cutoff point. From the radiative-transfer simulations in Ref. [41], I determined that a typical Lyman-limit system has an electron density around 10^{-2} cm^{-3} and a temperature $T \approx 2.4 \times 10^4$ K. This gives a collisional ionization rate of approximately [42] 1.2×10^{-13} s^{-1}, compared to the photoionization rate [21] of around 1.0×10^{-12} s^{-1}. One therefore needs to suppress the intergalactic flux by a factor of around 10 to reach collisional ionization.

Then, making a simple uniform-density one-dimensional (1D) model of a clump irradiated by the intergalactic flux, the mean flux inside as a function of total column density N is given by $\Gamma_0(1 - e^{-\bar{\sigma}_{\rm HI}N})/(\bar{\sigma}_{\rm HI}N)$. Note therefore that, although the central photoionization rate falls exponentially with N, the mean falls only approximately linearly with N. Accordingly for the mean rate to drop to $\Gamma_0/10$ requires $N \approx 10\bar{\sigma}_{\rm HI}^{-1} \approx (2.6 \times 10^{18})$ cm^2. The cumulative effect of column densities greater than these limits constitute only around 10% of the IGM opacity (see Ref. [33], Fig. 10). For that reason I will adopt $p_{\rm clump} = 0.10$, showing the effect of varying the value at the end of the paper.

B. Recombination emission

Now let us turn attention to the recombination radiation term. This arises automatically from the integration discussed above Eq. (8), with the definition

$$\beta_{\rm r}(T) = \frac{1}{\bar{\sigma}_{\rm HI}} \int_0^\infty \sigma_{\rm HI}(\nu) f_{\rm rec}(\nu, T) {\rm d}\nu, \tag{18}$$

which provides a dimensionless measure of the amount of recombination radiation that ends up in the monochromatic waveband under consideration. Over the frequencies of interest, $f_{\rm rec}$ can be approximated as an offset Maxwell-Boltzmann distribution corresponding to the electron temperature T, scaled by the fraction of recombinations that occur directly to the ground state:

$$f_{\rm rec}(\nu, T) \approx 2h f_{\infty\to 1} \sqrt{\frac{h(\nu - \nu_{\rm LL})}{\pi (kT)^3}} e^{-h(\nu-\nu_{\rm LL})/kT}, \tag{19}$$

where $f_{\infty\to 1} \approx 0.40$ is the fraction of recombinations direct to the ground state [43], h is Planck's constant and k is Boltzmann's constant. At $z \approx 2.4$ (close enough to our fiducial redshift), a typical forest temperature is [16,19] $T = 2.5 \times 10^4$ K; evaluating Eq. (18) then gives $\beta_{\rm r} = 0.39$. In principle, we could keep track of how variations in the mean temperature correlate with variations in the mean density; note, however, that when considering the averaged effects on linear scales this may not be the same as the equation of state measured for individual clouds [19]. Worse, large-scale spatial temperature correlations could well be generated by unmodeled, nonequilibrium phenomena such as helium reionization [16,44–46]. Luckily, the final effect of these on the photoionization equilibrium will be subdominant because $\beta_{\rm r}$ changes quite slowly with temperature ($d\beta_{\rm r}/d\ln T = -0.15$). The impact of thermal fluctuations on the recombination spectral shape is thus small compared to their effect on the recombination rate (which scales approximately as $T^{-0.7}$ in the intergalactic regime). Even in the latter case, within our monochromatic approximation the temperature variations do not depend strongly on the local ionizing field strength [47]. Incorporating multiwavelength, time-dependent radiative transfer could introduce qualitatively important corrections to the temperature field and should be prioritized in future work (see Sec. V).

C. Other sources

There is one remaining term in Eq. (8) that as yet has not been discussed: $j(\mathbf{x})$. Recall that (8) is obtained by integrating (4) over frequency; accordingly $j(\mathbf{x})$ is defined by

$$j(\mathbf{x}) = \int j_\nu(\mathbf{x}) \sigma_{\rm HI}(\nu) {\rm d}\nu, \tag{20}$$

and specifically excludes the recombination emission which was treated separately above. Looking ahead in the calculation, we will need to understand the statistical properties of the $j(\mathbf{x})$ field. It is widely believed that, at $z \approx 2.3$, quasars and galaxies both contribute significantly to the UV emission [21]. For both populations, systematic fluctuations δ_j are thought to be proportional to a constant (the "bias," b_j) times the matter density fluctuations δ_ρ [48] when averaged over suitably large scales. However we will also need to consider the shot noise: because quasars are rare, even a uniform distribution would have a significant Poisson fluctuation in density from place to place. In the limit of large scales (and therefore large numbers), these Poisson fluctuations can be modeled as an additive Gaussian noise, giving the total large-scale emissivity variations:

$$\delta_j(\mathbf{x}) = b_j \delta_\rho(\mathbf{x}) + \delta_{\rm SN}(\mathbf{x}), \tag{21}$$

where, in accordance with definition (1), $\delta_\rho(\mathbf{x})$ is the fractional matter overdensity determined by the cosmology and $\delta_{\rm SN}(\mathbf{x})$ is the uncorrelated shot-noise Gaussian random field. If multiple source populations contribute to the emissivity, they add linearly from which it follows that

$$\delta_j(\mathbf{x}) = \sum_i \frac{j_{0,i}}{j_0} (b_{j,i} \delta_\rho(\mathbf{x}) + \delta_{{\rm SN},i}(\mathbf{x})), \tag{22}$$

where the sum extends over the different source categories i and the total mean emissivity is $j_0 = \sum_i j_{0,i}$. Comparing

Eqs. (21) and (22) shows that the net large-scale effect is the same as that of a single population with suitably averaged parameters as I will discuss below.

Consider first the bias b_j, which is established by estimating the correlation strength of the emitting objects. Quasars are found to be strongly biased ($b_q \approx 4$) with respect to the matter density field [49,50]. Galaxies are substantially less strongly correlated, and therefore less biased; Ref. [51] quotes a correlation length of $r_0 = 4.3h^{-1}\,\mathrm{Mpc}$ for a sample of bright ($23.5 < \mathcal{R} < 25.5$) galaxies at $z \approx 2.2$. This translates into a bias of $b_g \approx 2.4$ with the Planck cosmology described above, assuming the underlying matter fluctuations to be normalized [38] to $\sigma_8 = 0.834$ at $z = 0$.

To add complication, these biases are measured for bright objects; especially in the case of galaxies, a significant fraction of photons are emitted from a large population of individually under-luminous objects [52]. To understand how the bias scales with luminosity one can assume it arises from the underlying dark matter halo. In that case the bias implies a halo mass [53]; for instance, with $b_g = 2.4$ we obtain a characteristic mass[2] of $M = 4 \times 10^{11}\mathrm{M}_\odot$. Suppose we wish to consider galaxies a factor of 10 fainter than the sample of Ref. [51]; then a number of arguments point to the dark matter halos being approximately a factor of $\sqrt{10}$ less massive [41,55]. This can be translated back into a bias of 1.9. Similarly, dropping another factor of 10 in luminosity yields halos with bias 1.5.

So the appropriate "source bias" is sensitive to details of the underlying population generating the UV photons. For a combination of different sources, Eq. (22) shows that the net bias is exactly the average of the two individual biases, weighted by the emissivity of the populations. As a default value in this work, I will assume an average source bias of $b_j = 3$, representing the average between highly biased quasars and a range of galaxy luminosities. (Recall that the value does not need to be further reduced for recombination emission, since that is included elsewhere in the calculation.) Reflecting the uncertainty, I will also show results for a range of b_j from 1.5 to 4.0. A great attraction of future measurements of the effects in this paper is that they should in principle constrain b_j and so shed light on the origin of UV photons.

Now consider the shot-noise term δ_{SN} for a single population. This represents random variations in the density of sources. For a mean density of $\bar{n}$, the number of sources in a volume V is given by

$$N(V) = \bar{n}\left(V + \frac{1}{(2\pi)^{3/2}}\int_V \mathrm{d}^3x \int \mathrm{d}^3k e^{i\mathbf{k}\cdot\mathbf{x}}\tilde{\delta}_{\mathrm{SN}}(\mathbf{k})\right). \tag{23}$$

[2]These results have been calculated using the prescriptions of Ref. [53] as implemented by Ref. [54].

Using this expression to demand that $\langle N^2\rangle - \langle N\rangle^2 = N$ for any volume V (the Gaussian, large-N limit of Poisson noise) dictates the power spectrum

$$P_{\mathrm{SN}}(k) = \bar{n}^{-1}, \tag{24}$$

independent of scale, where $P_{\mathrm{SN}}(k)$ is defined by Eq. (3) and correctly has units of comoving volume as explained earlier. The shot noise is modeled as stationary; in fact quasars likely have a finite duty cycle, causing the realization of the shot noise to change over time. The effect of this cannot be analyzed rigorously with the time-stationary approach I have adopted, but it could plausibly change the impact of noise on large scales. It should therefore be investigated in future work.

Just as for the bias b_j, choosing an appropriate value of $\bar{n}$ is tricky. One can start by parametrizing the quasar luminosity function $\Phi(L)$ using a double power-law fit e.g [56,57]; following the consequences of Eq. (21) for a series of infinitesimal bins in luminosity, assuming an independent shot-noise realization for each bin, one is led to an L^2-weighted [27,32] effective quasar number density defined by

$$\bar{n}_q = \frac{(\int \Phi(L)L\mathrm{d}L)^2}{\int \Phi(L)L^2\mathrm{d}L}. \tag{25}$$

Since the intrinsic luminosity function in the ionizing radiation is unknown (being completely obscured by Lyman-limit absorption) I assume that the ionizing radiation of a given quasar scales linearly with its bolometric luminosity. Evaluating Eq. (25) using estimates for the bolometric population parameters [56] then gives $\bar{n}_q$ between 1.5×10^{-6} and $10^{-5}\ \mathrm{Mpc}^{-3}$ comoving over the parameter range quoted by Ref. [56].

As with b_j, the δ_{SN} appearing in Eq. (21) can be seen from Eq. (22) to be a photon-weighted average of the δ_{SN} appropriate to the two populations. The density of galaxies is so much higher that one can essentially ignore their shot-noise contribution compared to that of the quasars. Tracing this through, assuming again a 50% contribution from both populations, one has to multiply $\bar{n}$ by 4 to account for the galaxy part of the emission (since δ_{SN} scales with $\bar{n}^{-1/2}$). This yields the approximate upper limit $\bar{n} \approx 4 \times 10^{-5}\ \mathrm{Mpc}^{-3} \approx 10^{-4}h^3\,\mathrm{Mpc}^{-3}$, with a lower limit of $2 \times 10^{-5}h^3\,\mathrm{Mpc}^{-3}$. To be clear, the $\bar{n}$ derived in this way is not the density of any particular population—it is a weighted average which accounts for the very different densities of two populations.

These estimates neglect any effects of time variability and anisotropy which will introduce qualitative corrections and possibly lead to an increase in the effective $\bar{n}$ by pushing observed variation to smaller scales. I will therefore adopt the upper end of the naive uncertainty for the present. Results will later be shown for a full range of

possible $\bar{n}$. Many of the same physical considerations bear on the value of both $\bar{n}$ and b_j, but I will consider them to be separate parameters for the sake of clarity.

The total power spectrum for the sources $P_j(k)$ is just

$$P_j(k) = b_j^2 P_\rho(k) + \bar{n}^{-1}, \tag{26}$$

which follows because, by assumption, $\langle \tilde{\delta}_\rho(\mathbf{k}) \tilde{\delta}_{\rm SN}(\mathbf{k}') \rangle = 0$. A similar expression was given by Ref. [32]. However, despite appearing on an equal footing in Eq. (26), the shot noise and correlated components behave differently in terms of their effect on the HI [28], as we will see below.

III. LINEARIZATION

The preceding section concluded by discussing the behavior of sources averaged on large scales in terms of spatial perturbations δ defined by (1), leading to a number of dimensionless physical parameters which are summarized in Table I. The plan now is to rewrite the radiative-transfer equations in terms of δ's. Ignoring all terms of order δ yields the homogeneous or "zero-order" approximation; the Boltzmann equation (8) becomes (integrating over all angles without loss of information)

$$\kappa_{\rm tot,0}(1 - \beta_{\rm HI}\beta_{\rm r})\Gamma_0 = 4\pi j_0, \tag{27}$$

which expresses the equilibrium condition that the overall rate of photon production is balanced by the effective absorption from redshifting, dilution and ionization.

Expanding Eq. (8) to linear order and simplifying using the background solution (27), one obtains an expression for $\tilde{\delta}_{f_{\rm LL}}$:

$$\tilde{\delta}_{f_{\rm LL}} = \frac{(1 - \beta_{\rm HI}\beta_{\rm r})\tilde{\delta}_j + \beta_{\rm HI}\beta_{\rm r}[\tilde{\delta}_{n_{\rm HI}} + \tilde{\delta}_\Gamma] - \tilde{\delta}_{\kappa_{\rm tot}}}{i(a\kappa_{\rm tot,0})^{-1}(\mathbf{n}\cdot\mathbf{k}) + 1}, \tag{28}$$

where I have suppressed functional $\mathbf{k}$ and $\mathbf{n}$ dependencies for brevity. Excepting small gravitational effects (Appendix A), the effective opacity fluctuations $\delta_{\kappa_{\rm tot}}$ are linked solely to variations in the diffuse and clumped neutral hydrogen:

$$\delta_{\kappa_{\rm tot}} = \beta_{\rm HI}\delta_{n_{\rm HI}} + \beta_{\rm clump}\delta_{\bar{\kappa}_{\rm clump}}. \tag{29}$$

Integrating out the remaining angular dependence in Eq. (28), noting that $\int {\rm d}^2 n \tilde{\delta}_{f_{\rm LL}} = 4\pi\tilde{\delta}_\Gamma$, one obtains an implicit equation for $\tilde{\delta}_\Gamma$:

$$\begin{aligned}\tilde{\delta}_\Gamma(\mathbf{k}) &= [(1 - \beta_{\rm HI}\beta_{\rm r})\tilde{\delta}_j - \beta_{\rm HI}(1 - \beta_{\rm r})\tilde{\delta}_{n_{\rm HI}} \\ &\quad - \beta_{\rm clump}\tilde{\delta}_{\bar{\kappa}_{\rm clump}} + \beta_{\rm HI}\beta_{\rm r}\tilde{\delta}_\Gamma]S(k), \\ S(k) &= \frac{a\kappa_{\rm tot,0}}{k}\arctan\frac{k}{a\kappa_{\rm tot,0}},\end{aligned} \tag{30}$$

showing the characteristic scale dependence arising in the radiation field (Fig. 1, upper panel). Performing an inverse Fourier transform on the kernel $S(k)$ returns the radial function $e^{-\kappa_{\rm tot,0} r}/r^2$ (Fig. 1, lower panel). The systematic approach has therefore recovered something like the heuristic equations of Refs. [27,32], where sources are convolved with a similar kernel. However, those works do not take into account the shortened effective mean-free path from redshifting and volume-dilution contributions (they use $\kappa_{\rm HI}$ where $\kappa_{\rm tot,0}$ is more appropriate). Moreover the appearance of $\tilde{\delta}_{n_{\rm HI}}$ (from inhomogeneous absorption) and $\tilde{\delta}_\Gamma$ (from recombination radiation) on the right-hand side of Eq. (30) means that the convolution kernel is modified from this simple form once absorption fluctuations, as well as emission fluctuations, are included.

On large scales $S(k \to 0)$ is ≈ 1, meaning fluctuations in the effective source function are tracked by fluctuations in the number density of ionizing photons. This agrees with the intuitive picture outlined earlier, in which regions separated by more than the mean-free path must arrive

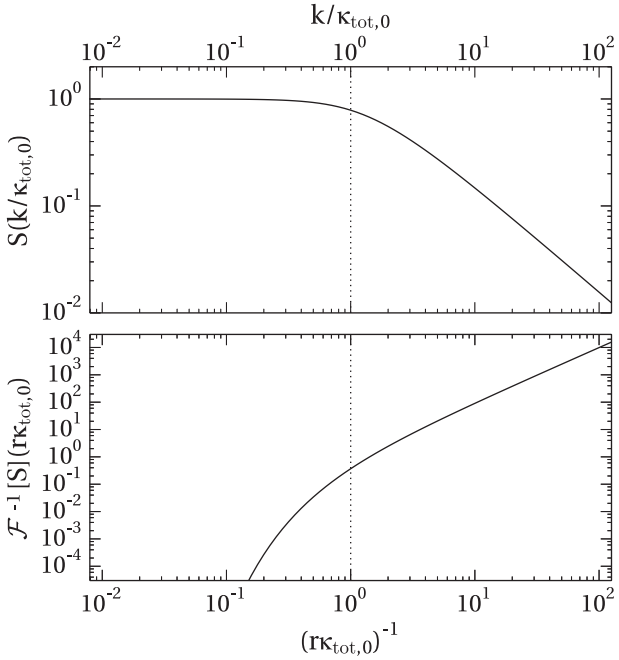

FIG. 1. In the linear approximation, the radiative transfer consists of convolving an effective source function (including emission, absorption and reradiation terms) with a kernel S, Eq. (30). The kernel is shown here in Fourier space (upper panel) as a function of wave number divided by $\kappa_{\rm tot,0}$. On large scales (toward the left) $S(k)$ is ≈ 1, meaning fluctuations in the effective source function are tracked by fluctuations in the number density of ionizing photons. On small scales (toward the right), $S(k)$ decays towards zero; fluctuations in the photon density are suppressed and the uniform UV approximation will apply. The lower panel shows the same kernel transformed into real space, $\mathcal{F}^{-1}[S]$; the horizontal axis is an inverse distance, so the two panels read in the same direction.

TABLE I. Dimensionless quantities used in this work, with a brief explanation and the default value calculated or estimated at $z = 2.3$.

	Spectrum-dependent coefficients	
α_z	Coefficient for background redshifting (11)	1.57
β_r	Fraction of HI recombinations to LL photons (18)	0.39
	Estimated origin of effective opacity $\kappa_{\mathrm{tot},0}$, Eq. (13)	
β_HI	Fraction from HI in photoionization equilibrium	0.56
β_clump	Fraction from collisional-equilibrium clumps	0.06
β_z	Fraction from redshifting	0.13
β_V	Fraction from dilution	0.25
	Input biases relative to the linear overdensity field	
$b_\mathrm{HI,u}$	Bias of HI in homogeneous-radiation limit	1.5
b_j	Bias of photon source objects	3
$b_{j,\mathrm{eff}}$	Effective bias of sources including recombination (33)	2.6

at independent ionization equilibria. On small scales $S(k \to \infty)$ decays towards zero; fluctuations in the photon density are suppressed and the uniform UV approximation is recovered. (On sufficiently small scales one will, however, have enhanced nonlinear shot noise; as discussed in the Introduction I will consider only the linear regime in the present work.)

We are now in a position to understand the mean ionization state. Recalling that the effects of shielded clumps have already been dealt with thee n_HI field refers specifically to the intergalactic HI alone—we can write

$$\delta_{n_\mathrm{HI}} = \delta_{n_\mathrm{HI},\mathrm{u}} - \delta_\Gamma, \tag{31}$$

where $\delta_{n_\mathrm{HI},\mathrm{u}}$ describes the HI field in the case of a completely uniform ionizing background; the given relationship is a consequence of linearizing the photoionization equilibrium equation $n_\mathrm{HI} \propto 1/\Gamma$. In the absence of any radiative fluctuations, by definition $\delta_{n_\mathrm{HI}} = \delta_{n_\mathrm{HI},\mathrm{u}}$. Combining Eqs. (30) and (31) gives a solution for $\tilde{\delta}_{n_\mathrm{HI}}$ in terms of $\tilde{\delta}_j$, $\tilde{\delta}_{n_\mathrm{HI},\mathrm{u}}$ and $\tilde{\delta}_{\bar{\kappa}_\mathrm{clump}}$:

$$\tilde{\delta}_{n_\mathrm{HI}} = \frac{\tilde{\delta}_{n_\mathrm{HI},\mathrm{u}} - [(1-\beta_\mathrm{HI}\beta_\mathrm{r})\tilde{\delta}_j - \beta_\mathrm{clump}\tilde{\delta}_{\bar{\kappa}_\mathrm{clump}} + \beta_\mathrm{HI}\beta_\mathrm{r}\tilde{\delta}_{n_\mathrm{HI},\mathrm{u}}]S(k)}{1-\beta_\mathrm{HI}S(k)}. \tag{32}$$

Now assume that $\tilde{\delta}_{n_\mathrm{HI},\mathrm{u}}$ (the HI density fluctuations in a completely uniform UV field), $\tilde{\delta}_j$ (the source density fluctuations) and $\tilde{\delta}_{\bar{\kappa}_\mathrm{clump}}$ (the self-shielded clump opacity fluctuations) can be written as a bias (respectively $b_\mathrm{HI,u}$, b_j and b_clump) times the fiducial cosmic density field, and further define an effective source bias

$$b_{j,\mathrm{eff}} = (1-\beta_\mathrm{HI}\beta_\mathrm{r})b_j - \beta_\mathrm{clump}b_\mathrm{clump} + \beta_\mathrm{HI}\beta_\mathrm{r}b_\mathrm{HI,u}, \tag{33}$$

which takes into account the recombination emission from the IGM and absorption from the clumps.[3] The intergalactic HI density then follows immediately,

$$\tilde{\delta}_{n_\mathrm{HI}} = \frac{[b_\mathrm{HI,u} - b_{j,\mathrm{eff}}S(k)]\tilde{\delta}_\rho - [1-\beta_\mathrm{HI}\beta_\mathrm{r}]S(k)\tilde{\delta}_\mathrm{SN}}{1-\beta_\mathrm{HI}S(k)}, \tag{34}$$

using Eq. (21). The HI density perturbation $\tilde{\delta}_{n_\mathrm{HI}}$ can be split into two terms, corresponding to the correlated and shot-noise components respectively. The correlated part obeys $\tilde{\delta}_{n_\mathrm{HI}} = b_\mathrm{HI}\tilde{\delta}_\rho$ where

$$b_\mathrm{HI}(k) = \frac{b_\mathrm{HI,u} - b_{j,\mathrm{eff}}S(k)}{1-\beta_\mathrm{HI}S(k)}, \tag{35}$$

showing that the HI density traces the cosmological density in a scale-dependent way. This is the main result of the present work. Its implications will be discussed in the next section.

IV. RESULTS

A. Bias and power spectrum

The preceding section used a systematic linearization of first-principles radiative transfer to derive equations governing the IGM HI density on large scales. I will now explore what this implies for the bias b_HI (35) and total power spectrum. As previously discussed, a number of uncertain parameters enter the calculation, namely: the IGM bias in the uniform-radiation limit, $b_\mathrm{HI,u}$; the effective

[3] I will assume that $b_\mathrm{clump} = b_\mathrm{HI,u}$, since both unshielded and clumped HI are included in the estimate made in Appendix B; it could plausibly be the case that b_clump in reality differs from $b_\mathrm{HI,u}$ if the distinction between phases is made carefully—but since both b_j and b_clump only enter through $b_{j,\mathrm{eff}}$, any uncertainty in b_clump is degenerate with the uncertainty in b_j which will be explored later.

source density $\bar{n}$; the effective source bias b_j; the fraction of opacity in collisionally ionized clumps $p_{\rm clump}$; and the mean physical Lyman-limit opacity $\kappa_{\rm HI} = (\beta_{\rm HI} + \beta_{\rm clump})\kappa_{\rm tot,0}$. This section will explore the consequences of varying all of these except for $b_{\rm HI,u} = 1.5$ (see Appendix B for details). I will explore substantial variations around the default choices (justified earlier in the text) of $b_j = 3$, $\bar{n} = 10^{-4}h^3\,{\rm Mpc}^{-3}$, $p_{\rm clump} = 0.10$ and $(a\kappa_{\rm HI})^{-1} = 570$ Mpc comoving.

The top panel of Fig. 2 plots $b_{\rm HI}$ against comoving wave number k, Eq. (35). This represents the linear relationship between HI density and total density as a function of scale. The dashed and solid lines show respectively the assumed relationship when there are no effects of inhomogeneous radiation and the calculated relationship for the default parameters.

The basic functional form and its b_j dependence can be understood as follows. On small scales ($k \gg \kappa_{\rm tot,0}$), $S(k)$ asymptotes to zero, so

$$b_{\rm HI}(k \gg \kappa_{\rm tot,0}) = b_{\rm HI,u}, \tag{36}$$

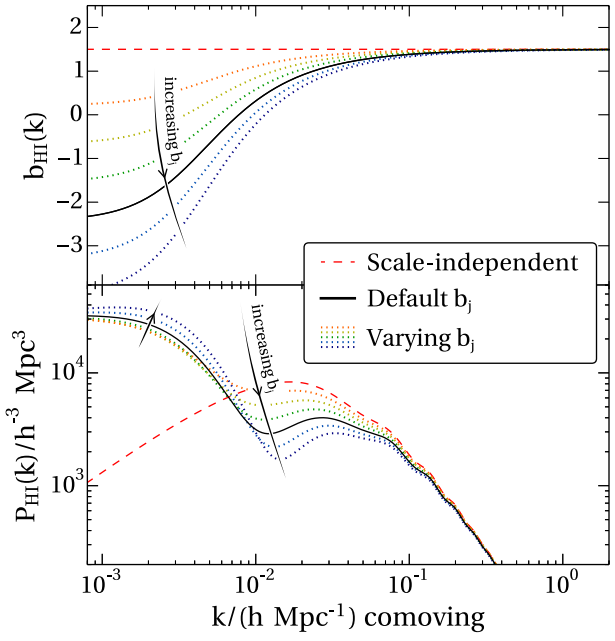

FIG. 2 (color online). (Upper panel) The calculated bias of intergalactic HI at $z = 2.3$ as function of comoving wave number. The dashed line shows the bias in the uniform-radiation case; the solid line shows the bias calculated with radiative transfer. On small scales (towards the right), the calculated bias agrees with that of the uniform case. On large scales, the HI is negatively biased because overdensities imply high emissivity, high radiation density and hence net HI under-density. Dotted lines show the effect of changing the source bias b_j; from top to bottom, $b_j = 1.5, 2.0, \cdots, 4.0$. (Lower panel) The corresponding power spectrum, $P_{\rm HI}(k)$, defined by Eq. (38), has a strong feature where $b_{\rm HI}$ passes through zero (near to $k = \kappa_{\rm tot,0}^{-1}$).

showing that the small modes are unaffected by radiative-transfer phenomena at the linear level. Conversely on large scales, $S(k)$ asymptotes to one, giving

$$b_{\rm HI}(k \ll \kappa_{\rm tot,0}) = \frac{b_{\rm HI,u} - b_{j,\rm eff}}{1 - \beta_{\rm HI}}. \tag{37}$$

For $b_{j,\rm eff} > b_{\rm HI,u}$, this makes the HI negatively biased on large scales, i.e. anticorrelated with the total density. The intensity of the radiation in dense regions overcompensates for the clustering of hydrogen, causing a net deficit in neutral hydrogen—a direct analogue of the proximity effect but averaged over many sources on large scales.

This has profound consequences for the power spectrum of HI fluctuations. Recall that $\tilde{\delta}_{\rm SN}$ and $\tilde{\delta}_\rho$ are uncorrelated, so we have

$$P_{\rm HI}(k) = b_{\rm HI}(k)^2 P(k) + \left[\frac{(1 - \beta_{\rm HI}\beta_{\rm r})S(k)}{1 - \beta_{\rm HI}S(k)}\right]^2 \frac{1}{\bar{n}}. \tag{38}$$

This power spectrum is plotted in the lower panel of Fig. 2, with the default value $\bar{n} = 5 \times 10^{-4}h^3\,{\rm Mpc}^{-3}$ (Sec. II C) and a fiducial $P(k)$ for the Planck cosmology calculated using CAMB [58]. The strong feature arises because $b_{\rm HI}^2$ touches zero at $k^{-1} \approx 125h^{-1}$ Mpc comoving (for the default parameters). Accordingly there is a sharp dip in $P(k)$ around that wave number. On larger scales still, at the far left of Fig. 2, the HI fluctuations become stronger than predicted in the scale-independent model. This arises from a mixture of shot noise (discussed in more detail below) and the large magnitude[4] of the limiting bias (37).

Dotted lines in Fig. 2 explore the impact of changing b_j over a wide range; from top to bottom, $b_j = 1.5, 2.0, \cdots, 4.0$. Recall that, as discussed in Sec. II C, the source bias b_j is composed of a photon-weighted average of different populations (excluding recombination emission, which is accounted for elsewhere within the calculation). As the source bias increases, the effects at a given wave number typically become stronger. Consequently the zero in $b_{\rm HI}$ moves to larger wave numbers (shorter distances), making the radiation thumbprint more observationally accessible. Even for small biases, however ($b_j = 1.5$) the effects are significant on scales of tens to hundreds of megaparsecs comoving.

On sufficiently small scales, the HI power spectrum is unaffected by radiative transfer, regardless of the value of b_j. In particular, 1D measurements of the Lyman-alpha forest are limited by the small path length that can be observed with an individual quasar. Only wave numbers greater than $\sim 0.02\,{\rm s\,km}^{-1}$, corresponding to $2h\,{\rm Mpc}^{-1}$

[4]Although $b_{\rm HI}$ turns negative on large scales, this is not directly seen in the power spectrum which is sensitive only to $b_{\rm HI}^2$. On the other hand $b_{\rm HI}$ appears linearly when the forest is cross correlated with another tracer population [29,59], so its sign is detectable in principle, a point explored a little more in a moment.

comoving at $z = 2.3$, are measured [60]; from Fig. 2 it is clear that the effects are minimal for such measurements. (The nonlinear, non-Gaussian contribution to the shot noise on those scales will become significant, but plausibly average out over many sightlines [11,28].)

B. Auto- and cross-correlation function

To test the large-scale predictions against observations one must turn to more recent three-dimensional (3D) analyses that take advantage of modern surveys with dense background sources [34]. In this context, it is more conventional to consider the correlation function $\xi(\mathbf{r}) = \langle\delta(\mathbf{x}+\mathbf{r})\delta(\mathbf{x})\rangle$. It has been widely used to constrain the BAO feature in Lyman-α and other large-scale structure tracers [12,35,61]. By isotropy ξ is actually a function of $r = |\mathbf{r}|$ alone; one can show it is related to the power spectrum via a Legendre transformation,

$$\xi_{\rm HI}(r) = \frac{1}{2\pi^2}\int {\rm d}k \frac{\sin kr}{kr} k^2 P_{\rm HI}(k). \tag{39}$$

The correlation function for the Lyman-α flux ξ_F is closely related to $\xi_{\rm HI}$ and can be measured from observations relatively directly. As explained in the Introduction, this paper will not go as far as calculating ξ_F, but a brief discussion of the relationship to $\xi_{\rm HI}$ is given in Sec. V.

Figure 3 shows $\xi_{\rm HI}$ (renormalized by r^2 to highlight the BAO structure) for the scale-free (dashed line) and default radiation model (solid line). Once again the dotted lines show the calculated correlation function for a range of different source biases b_j from 1.5 to 4.0. The mapping from power spectrum to correlation function causes a substantial mixing of information on different scales, so the new shape needs a little unpicking to understand. On scales smaller than $\sim 5h^{-1}$ Mpc, the scale-free predictions are barely altered; this corresponds to the small-scale limit $b_{\rm HI} \to b_{\rm HI,u}$ in the bias, Fig. 2. Moving to larger separations, the radiation-corrected correlation function falls rapidly compared to the scale-free counterpart, because the HI bias is declining and the power on these scales is suppressed. In fact, the new correlation function turns negative at around $55h^{-1}$ Mpc; this is an artifact of the constraint that $\int \xi(r) r^2 {\rm d}r = 0$ for a properly mean-calibrated sample, and the negativity in itself does not indicate anything physically special about these scales.

The BAO feature—a hump at around $r = 100h^{-1}$ Mpc—remains clearly visible in all cases, but the local maximum in $r^2\xi(r)$ shifts marginally. The local maximum can be found at $100.0h^{-1}$ Mpc in the homogeneous case (dashed line) but at $101.2h^{-1}$ Mpc in the fiducial $b_j = 3.0$ case (solid line). At distances exceeding $130h^{-1}$ Mpc, the new correlation function begins to rise because of contributions from the increased power on very large scales (far left of Fig. 2).

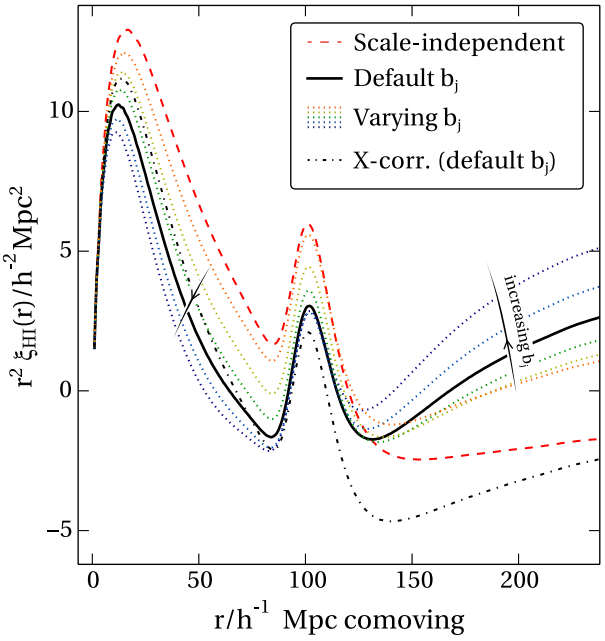

FIG. 3 (color online). The correlation function of intergalactic HI at $z = 2.3$, as defined by the Legendre transform (39) of the power spectrum shown in Fig. 2. As before, dashed lines show the constant-bias case, whereas the solid line shows the calculated bias for inhomogeneous radiation (in the default case, $b_j = 3.0$). The dotted lines show a series of different source biases ($b_j = 1.5$, nearest the dashed line; then 2.0, 2.5, 3.5 and 4.0). The result of cross correlating the HI against a tracer population with fixed bias is shown by the dash-dotted line.

In cross correlation, the signature looks slightly different. As an illustration, the dot-dashed line in Fig. 3 shows a hypothetical cross correlation against a fixed-bias population with $b = 1.5$ (this value has no significance except to scale the overall function similarly to the autocorrelation). In other words, I am plotting $\xi_{\rm HI\times} \equiv 1.5\langle\delta_{\rm HI}(\mathbf{x}+\mathbf{r})\delta_\rho(\mathbf{x})\rangle$. In simple cases this would return the geometric mean of the dashed and solid lines. However, there are a couple of more subtle effects here. First, the negativity of the HI bias on large scales reduces the large-distance cross correlation ($\xi_{\rm HI\times}$ probes $b_{\rm HI}$ whereas $\xi_{\rm HI}$ is sensitive only to $b_{\rm HI}^2$). Second, the plot assumes cross correlation against a population other than quasars so that the shot-noise term cancels. Overall this leads to a cross correlation that is suppressed more strongly on large scales than would be expected from an averaging argument.

C. Varying other parameters

So far I have only shown results for varying b_j. However there are other parametric dependencies which ought to be examined. The first is the physical HI opacity, $\kappa_{\rm HI}$, which is the inverse of the mean-free path of a photon in the absence of redshifting or volume dilution. The default value has been discussed extensively above; in Fig. 4 I have shown what happens when $\kappa_{\rm HI}$ is changed by a factor of 0.25, 0.5,

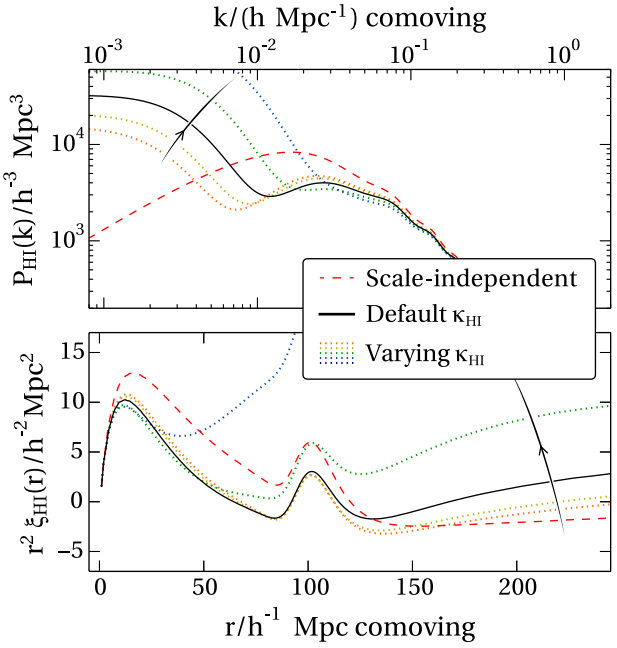

FIG. 4 (color online). The effect of varying the mean HI opacity on the power spectrum (upper panel) and correlation function (lower panel). Other parameters are held fixed. The dashed and solid lines show the uniform-radiation and reference cases respectively, so agreeing with the same lines in Figs. 2 and 3. Dotted lines show the results for an HI opacity 0.25, 0.5, 2 and 4 times that of the default case. As the opacity increases, the mean-free path decreases, meaning that the "dip" feature in the HI power spectrum moves to shorter wave numbers. Consequently small-scale power is increasingly suppressed, whereas large-scale power is enhanced.

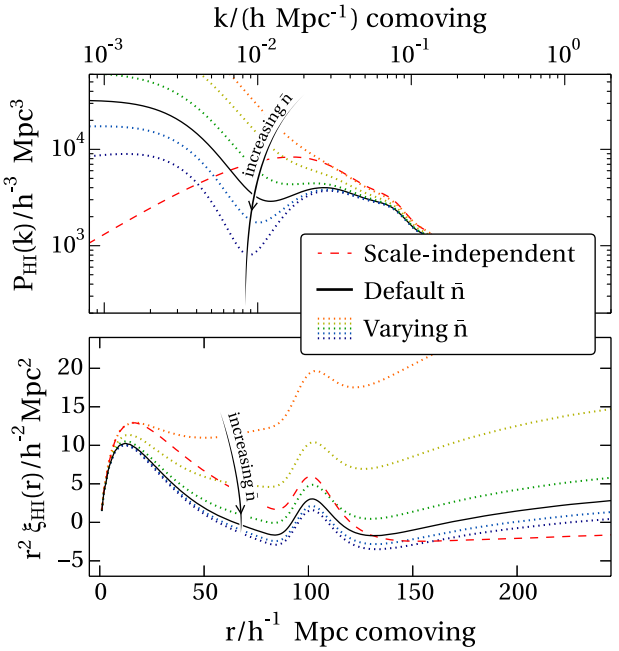

FIG. 5 (color online). The effect of source shot noise on the power spectrum (upper panel) and correlation function (lower panel). Other parameters are held fixed. The dashed and solid lines show the uniform-radiation and reference cases respectively, so agreeing with the same lines in Figs. 2 and 3. The dotted lines show the results for (top to bottom) $\bar{n} = 10^{-5}$, 5×10^{-5}, 10^{-4}, 5×10^{-4} (solid line), 10^{-3} and $5 \times 10^{-3} h^3\,\mathrm{Mpc}^{-3}$. As the effective source density $\bar{n}$ decreases, the amplitude of shot noise increases, confusing the HI signal. The changes are strongest on large scales for the reasons discussed in the text.

2 and 4. The actual uncertainty in the observationally constrained value [33,40] is more likely under a factor of 2. The upper and lower panels show the power spectrum and correlation function respectively. Solid and dashed lines therefore correspond exactly to those presented in Figs. 2 and 3; the dotted lines show the impact of changing κ_{HI}. As this opacity increases (i.e. the mean-free path decreases), the effects become more prominent on smaller scales. A slightly more subtle change occurs at long wavelengths: as the mean-free path decreases, the large-scale limiting bias increases, as does the noise contribution. Since β_{HI} increases when H is fixed but κ_{HI} increases, this behavior is in accordance with Eq. (37). Physically, photoionized HI amplifies fluctuations in radiation on large scales: an overdensity of radiation implies a lower HI fraction and therefore a deficit in opacity, in turn boosting the overdensity of radiation. This is why, as κ_{HI} increases, the fluctuations on large scales become more dramatic.

Next consider the effect of varying $\bar{n}$ from its fiducial value. Recall that this determines the large-scale shot-noise contribution and is related to the underlying source population densities (Sec. II C). Fixing the other parameters, Fig. 5 demonstrates the effect of $\bar{n}$ varying between $5 \times 10^{-3} h^3\,\mathrm{Mpc}^{-3}$ and $10^{-5} h^3\,\mathrm{Mpc}^{-3}$ on the power spectrum (upper panel) and autocorrelation function (lower panel). Smaller source densities lead to a stronger effect, with significant power added in the case of $\bar{n} = 10^{-5} h^3\,\mathrm{Mpc}^{-3}$. It may come as a surprise that, in all cases, the effects of low source density are most pronounced as r becomes large (or k small) rather than in the opposite limit. In the linear, averaged limit, however, this is correct. The shot-noise power spectrum is suppressed on small scales by $S(k)^2$ which declines steeply at increasing k (Fig. 1). The intuitive picture that shot noise matters more on small scales depends on the transition to the nonlinear, unaveraged regime which I have not attempted to model.

Finally let us return to the parameter p_{clump}, which controls the fraction of opacity arising from self-shielded, collisionally ionized clumps as opposed to diffuse, photoionized HI. As this parameter is increased, β_{HI} decreases and β_{clump} increases. The overall effects are shown in Fig. 6 for $p_{\mathrm{clump}} = 0.0, 0.1$ (the default), $0.2, \cdots, 0.5$. For scenarios with a greater fraction of opacity in clumps, the effect of radiation is slightly mitigated on very large scales. However the differences are minor.

For realistic observations, the effects of p_{clump} will be somewhat larger: here I am plotting the effect only on the

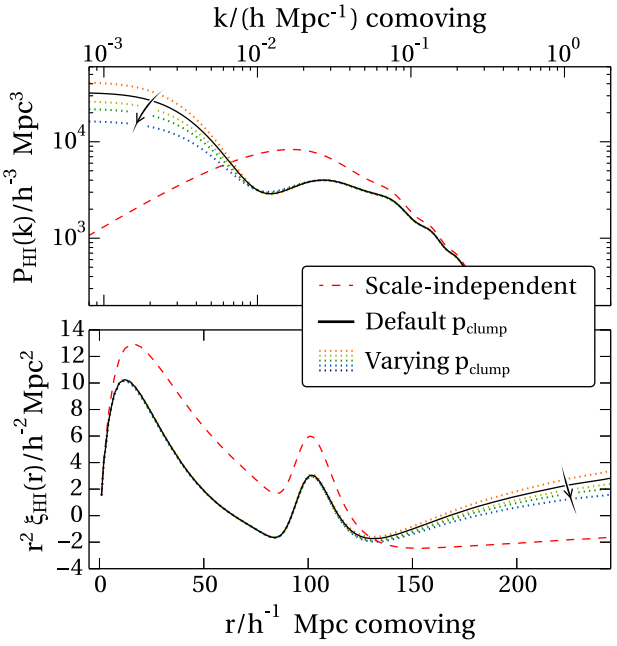

FIG. 6 (color online). The effect of changing the fraction of opacity from collisionally ionized clumps on the power spectrum (upper panel) and correlation function (lower panel). Other parameters are held fixed. The dashed and solid lines show the uniform-radiation and reference cases respectively, so agreeing with the same lines in Figs. 2 and 3. The dotted lines show the results for (top to bottom) $p_{\rm clump} = 0.0$, 0.1 (solid line), $0.2, \cdots, 0.5$. When more opacity arises from clumps (higher $p_{\rm clump}$), the effect of radiation is slightly weaker because the clumps are able to partially counterbalance the enhanced radiation in overdense regions. However, the effect is minor.

photoionized intergalactic medium. Increasing the fraction of clumps contributing to the Lyman-limit opacity will also increase the balance of such systems in the Lyman-α forest flux spectrum. Since they are collisionally ionized, they are not much affected by the inhomogeneous radiation field and therefore they dilute the scale-dependent effects roughly by a fraction $1 - p_{\rm clump}$. In other words, the leading-order effect of a large $p_{\rm clump}$ on observations will be different to, and more important than, the physical effect on the intergalactic HI which I have discussed here.

Nonetheless, whatever the value of any of these parameters, there are substantial changes to the intergalactic HI correlation function at all scales exceeding $5h^{-1}$ Mpc. It seems likely that these should be detectable with BOSS observations of the Lyman-α forest—even if observational complications lead to a substantial dilution. This prospect is considered further in the discussion below.

V. DISCUSSION

Radiative transfer imprints dramatic scale-dependent bias in the intergalactic HI and therefore the Lyman-α forest, even after reionization is complete. This follows because regions separated by distances larger than the UV photon mean-free path reach essentially independent photoionization equilibria. Source clustering is stronger than IGM clustering, leading to negative HI bias on large scales (i.e. the HI anticorrelates with large-scale overdensities).

This paper has presented a detailed calculation of these new effects by adopting a monochromatic, equilibrium, large-scale description, focusing on the large-scale, average correlations [30–32] rather than small-scale nonlinear fluctuations [9,11]. The systematic analytic treatment starts from first-principles radiative transfer and produces, with minimal computational effort, predictions for very large scales.

The calculation reveals, as expected from the argument above, a strong feature in the HI power spectrum and a corresponding distortion of its correlation function. According to the estimates here, this distortion should have an effect at the BAO scale ($\approx 100h^{-1}$ Mpc). The BAO bump position is slightly shifted—in Fig. 3, the local maximum is at 1.2% larger scales in the radiative-transfer case (solid line) compared to the constant-bias case (dashed line). That said, future cosmology constraints are unlikely to come from measuring the peak in such a simple way; so long as algorithms marginalize over possible broadband distortions to the correlation function, they will likely still recover an unbiased estimate of the BAO scale.

The most interesting conclusion is therefore that BAO-focused Lyman-α observational programmes will be able to recover helpful astrophysical constraints: the radiative-transfer distortions are strongly dependent on the mean bias of sources (Figs. 2 and 3), the HI opacity (Fig. 4) and the effective number density of sources $\bar{n}$ (Fig. 5). As b_j increases, the correlated component of the radiative fluctuations grows and the power on BAO scales decreases while the power on very large scales increases; as $\bar{n}$ decreases, the random component of the radiative fluctuations grows and the power on all scales increases. These trends seem to agree with numerical results where a comparison can be made [30–32].

The overall picture gives rise to a large number of questions. The most obvious is whether existing BOSS observations of the Lyman-α forest are compatible with the expected thumbprint. A variety of subtle observational issues must be taken into account before this can be answered. First, converting an HI correlation function into a flux correlation function is a nonlinear process that needs at a minimum to be calibrated by suitable numerical simulations [6]. Redshift-space distortions will mix the dynamical growth of structure with the tracer statistics into a final observed correlation function [6,62,63]. Dependent on the exact survey design, angular binning and data cuts, these effects could easily dilute the scale dependence, making the observed correlation function closer to the homogeneous-radiation result. However the changes in the

underlying intergalactic HI bias are sufficiently dramatic for it to seem implausible that the radiation-transfer signature would be obscured completely in forthcoming precision data. To be sure we will have to understand how the data processing and parameter degeneracies impact on our ability to measure the effects. Although the distortion is large, it is also a very smooth function of scale and therefore one needs to accurately calibrate the normalization of the correlation function over a wide range of scales to make a definitive detection; otherwise the effects are degenerate with a renormalization of the uniform-limit bias $b_{\mathrm{HI,u}}$.

The BOSS team has emphasized that their 3D Lyman-α forest pipeline is presently designed to pick out localized correlation-function features—i.e. the BAO bump—rather than reconstruct the entire function accurately [12,35]. Nonetheless an attempt to measure scale dependence in cross correlation against quasars has been made; none was found [29]. Conversely some scale dependence in the cross correlation between damped Lyman-alpha systems and the forest can be seen in the plots of Ref. [59]. It is unclear whether and how these results can be reconciled with the present work; observational difficulties such as continuum determination cause severely correlated errors in correlation functions and dealing with these leads to certain large-scale modes being projected out [12,35]. Overall, the task of determining whether the effects of radiative transfer are present in existing data is considerable. However, I hope that the present calculation has underlined the rewards of such an effort. The scale-dependent radiative transfer contains a rich, valuable source of information on the nature of UV sources.

21 cm emission studies will not be affected by these considerations because the HI they probe is largely in collisional- rather than photoionization equilibrium [41]. The 21 cm absorption forest would be affected in just the same way as the Lyman-α forest; but this phenomena is of most promise at high redshift before or during reionization [64]—so the present calculation does not apply. One way to tackle the larger fluctuations at high redshift is to use a halo-model-based calculation [65]; alternatively, a linear-theory approach has been taken to the problem by Refs. [37,66,67]. In these cases, an explicit time integration needs to be performed to follow the growth of ionized bubbles whereas in the present case the integration is absent because I have assumed equilibrium, making the present paper's calculations considerably simpler but more restricted in scope.

Depending on one's assumptions (for instance on the relative importance of quasars to the UV background, and on the quasar luminosity function), the rarity of sources also has a substantial impact on very large scales. Here I have modeled the resulting shot noise by a Gaussian approximation similar to that of Ref. [32]; in that work, noise was considered to be so large that the correlated component of the radiation fluctuations was thrown out of the calculation. In the present work the effects of shot noise seem milder, which reflects that I work at lower redshift (where the comoving density of quasars has increased) and make greater allowance for a UV contribution from star-forming galaxies. Crucially, the correlated component has a qualitatively different signature to the noise component of the radiation field: the former reduces the power of HI fluctuations on large scales, whereas the latter can only ever add power (at any scale). One effect that is absent from the present work concerns scales below 10 Mpc or so—here the noise would be significantly amplified [27] by 2-halo and other nonlinear effects. Another missing aspect from my analysis is that of time dependence which could, for example, add further confusion from quasar duty cycles.

With all this in mind it would be of great interest to supplement the linear-theory calculations of this work by revisiting the BAO-scale correlation function of the Lyman-α forest using nonlinear numerical simulations of gigaparsec chunks of the IGM with correlated sources, incorporating radiative transfer—along the lines of work described by Refs. [30–32]. Hints of the anticorrelation discussed at length in the present paper have been seen before in such efforts [28,30]. It would be helpful to include large-scale temperature fluctuations arising from helium reionization [46]. Or, one might be able to tackle temperature fluctuations analytically by relaxing the monochromatic assumption; it is worth reemphasizing that the current work includes the zero-order effects of hard photons [the spectral shape enters through Eq. (11)], but not first-order changes from local fluctuations in spectral shape. At this level of approximation, the gas thermal equilibrium is nearly unaffected by radiation intensity fluctuations [47] because the radiative heating rate can be approximately rewritten as a function of density and temperature (via the ionization equilibrium condition). To answer the important question of how thermal fluctuations change the large-scale signal one therefore needs either to incorporate multi-wavelength radiative transfer or go beyond an equilibrium approximation—or, preferably, both [47].

Further work is required to reach a unified view of how radiation changes the observational prospects for cosmology and astrophysics with the Lyman-α forest. The present investigation forms a first guide to the effects that will dominate on the largest scales.

ACKNOWLEDGMENTS

I am grateful to the anonymous referee for insightful comments and suggestions; to Jamie Bolton, Pedro Ferreira, Andreu Font, Martin Haehnelt, David Marsh, Jordi Miralda Escudé, Philip Hopkins, Anže Slosar, Hiranya Peiris and Matteo Viel for discussions; and to the Royal Society for financial support. Some numerical results in this paper were derived with the help of the pynbody framework [54].

Note added.—After circulating a draft of this work, I was made aware of an independent study by Gontcho A Gontcho, Miralda-Escudè and Busca [73]; at present it seems these authors reach many similar conclusions using a different calculation framework. This is encouraging, and it will be helpful to compare our approaches in due course.

APPENDIX A: ONCE MORE WITH GRAVITY

The plan for this appendix is to regenerate Eq. (8) but now including peculiar velocity and gravitational inhomogeneities in accordance with general relativity. In fact, all the effects turn out to be minor on scales of interest: dimensionless perturbations to the metric ϕ are going to be small compared to the dimensionless perturbations δ to the matter except on scales comparable to or larger than the horizon:

$$\phi(\mathbf{k}) \sim \frac{3H_0^2\Omega_{m,0}(1+z)}{2c^2k^2}\tilde{\delta}(\mathbf{k}) \approx \frac{4.2\ \mathrm{Gpc}^{-2}}{k^2}\tilde{\delta}(\mathbf{k}), \quad \text{(A1)}$$

where, as in the main text, k is the comoving wave number. If you are convinced by this argument, there is no need to read any further.

On the other hand, factors arising from spectral integrations could outweigh the scale contrast and make the effects relevant. To be sure either way one needs to press ahead with the calculation. I work in conformal-Newtonian gauge to make the geodesic equations relatively simple, but will briefly discuss the effect of gauge changes at the end of this Appendix. The formal derivation starts with a suitably perturbed flat Friedmann-Robertson-Walker universe described by the metric [68],

$$\mathrm{d}s^2 = -(1+2\psi)c^2\mathrm{d}t^2 + a(t)^2(1-2\phi)\mathrm{d}\mathbf{x}^2, \quad \text{(A2)}$$

where t is coordinate time, ψ and ϕ are the scalar potentials, and $\mathbf{x}$ are the comoving position coordinates.

Consider a photon with wave vector k_μ traveling through this perturbed metric. The null condition $k_\mu k^\mu = 0$ implies that (working throughout at first order in the potentials)

$$\frac{\mathrm{d}\mathbf{x}}{\mathrm{d}t} = \frac{c}{a}(1+\psi+\phi)\mathbf{n}, \quad \text{(A3)}$$

where $\mathbf{n}$ is the unit vector in the spatial propagation direction. The observed frequency of the photon in the coordinate frame is $\nu = ck^0(1+\psi)$; combining this with the geodesic equation for k^μ one finds that

$$\frac{\mathrm{d}\nu}{\mathrm{d}t} = \nu\left(-H - \frac{c}{a}\mathbf{n}\cdot\nabla\psi + \dot{\phi}\right). \quad \text{(A4)}$$

I will assume that on large scales all absorber and emitter streaming velocities follow that of the pressureless dark matter. The tangent 4-vector u^μ with $u_\mu u^\mu = -1$ can be related to the peculiar velocity $\mathbf{v}$ as

$$u^\mu = c^{-1}\begin{pmatrix} 1-\psi \\ \mathbf{v}/a \end{pmatrix}, \quad \text{(A5)}$$

where I have used a new assumption that $|\mathbf{v}|/c$ is small (the same order as the potentials). The frequency of the photon as seen by an absorber is then

$$\nu' \equiv -u^\mu k_\mu \approx \nu\left[1 - \mathbf{n}\cdot\frac{\mathbf{v}}{c}\right], \quad \text{(A6)}$$

again at first order. Finally, the physical 3-volume of a fixed $\mathbf{x}$ coordinate patch is proportional to $a^3(1-3\phi)$.

To define what is meant by an equilibrium solution to the Boltzmann equation in the relativistic setting, consider the rate of change $\dot{f}$ of the distribution function along a dark matter worldline. We have

$$\dot{f} = cu^\mu f_{,\mu} \approx (1-\psi)\frac{\partial f}{\partial t} + \frac{\mathbf{v}\cdot\nabla f}{a}, \quad \text{(A7)}$$

and since the gradient term is overall second order, we can again adopt the simple assumption that $\partial f/\partial t = 0$ to obtain a well-defined equilibrium at first order, independent of gauge.

Putting this together, the underlying number density f of photons satisfies the collisional Boltzmann equation,

$$\frac{c}{a}(1+\phi+\psi)(\mathbf{n}\cdot\nabla)f + \frac{\partial f}{\partial\nu}\frac{\mathrm{d}\nu}{\mathrm{d}t} + \left(\frac{\mathrm{d}\ln\Delta V\Delta\nu}{\mathrm{d}t}\right)f = C_{\nu'}[f], \quad \text{(A8)}$$

where $C_{\nu'}[f]$ is calculated according to Eq. (5) as before, but evaluated at the Doppler-shifted frequency ν' according to Eq. (A6); ∇ again means the ordinary derivative with respect to comoving spatial coordinates $\mathbf{x}$, and I have taken $\partial f/\partial t = 0$ as explained above. The quantity $\Delta V\Delta\nu$ appears because f is expressed in photons per unit physical volume per unit frequency. A bundle of photons which occupies a volume $\Delta V\Delta\nu$ in this space at one moment will occupy a different volume at the next. The evolution of ΔV can be calculated by setting up an initially cubic volume with edges $\Delta\mathbf{x}^{(\parallel)}$, $\Delta\mathbf{x}^{\perp 1}$ and $\Delta\mathbf{x}^{\perp 2}$ such that $\Delta\mathbf{x}^\parallel$ is parallel to $\nabla(\psi+\phi)$ and the others are perpendicular. Then

$$\Delta V = a^3(1-3\phi)\Delta\mathbf{x}^\parallel\cdot(\Delta\mathbf{x}^{\perp 1}\times\Delta\mathbf{x}^{\perp 2}) \quad \text{(A9)}$$

$$\Rightarrow \frac{\mathrm{d}\Delta V}{\mathrm{d}t} = \frac{\mathrm{d}\Delta\mathbf{x}^\parallel}{\mathrm{d}t}\cdot(\Delta\mathbf{x}^{\perp 1}\times\Delta\mathbf{x}^{\perp 2})a^3(1-3\phi) + 3\Delta V(H-\dot{\phi})$$
$$= \Delta V\left(\frac{c}{a}\mathbf{n}\cdot\nabla(\psi+\phi) + 3H - 3\dot{\phi}\right) \quad \text{(A10)}$$

at first order, with other terms canceling from the choice of $\Delta\mathbf{x}$ vectors. Along with the frequency factor, which follows

immediately from Eq. (A4), the overall volume term then reads

$$\frac{\mathrm{d}\ln \Delta V \Delta\nu}{\mathrm{d}t} = 2(H-\dot{\phi}) + \frac{c}{a}\mathbf{n}\cdot\nabla\phi. \tag{A11}$$

To link the gravitational effects to the perturbed density field ρ, we require the Einstein equations. The energy-momentum tensor for the pressureless, matter-dominated universe is

$$T_{\mu\nu} = \rho_0(1+\delta_\rho)u_\mu u_\nu. \tag{A12}$$

We will work at sufficiently high redshift that we can take $\Lambda = 0$. The zero-order Einstein equations recover the Friedmann and acceleration equations for the pressureless fluid universe; the linear-order equations reduce to

$$\phi = \psi \tag{A13}$$

$$c^2\nabla(\dot{\phi} + H\phi) = -4\pi G\rho_0 a v(\mathbf{x}) \tag{A14}$$

$$c^2\nabla^2\phi = 4\pi G\rho_0 a^2\delta_\rho + 3a^2H(\dot{\phi}+H\phi) \tag{A15}$$

$$0 = \ddot{\phi} + 4H\dot{\phi}. \tag{A16}$$

Equation (A16) is solved by $\dot{\phi} = 0$ (the potential is frozen, which corresponds to putting the matter perturbation δ_ρ in the growing mode); substituting also the zero-order Friedmann equation $3H^2/c^2 = 8\pi G\rho$ we have

$$(-2c^2\nabla^2 + 6a^2H^2)\phi = -3a^2H^2\delta_\rho \tag{A17}$$

$$\text{and} \quad 3aH\mathbf{v} = -2c^2\nabla\phi. \tag{A18}$$

Let us now follow exactly the same procedure as in Secs. II and III to obtain our previous approximation but with the relativistic terms present. Multiply Eq. (A8) by σ_{HI} and integrate over all frequencies; then, comparing against Eq. (8) at linear order, only two extra terms survive. In particular, the perturbation to $\mathrm{d}\mathbf{x}/\mathrm{d}t$ is irrelevant because there are no spatial gradients in the background. Of the two remaining terms, first consider how the peculiar velocities induce an extra term on the right-hand side:

$$\int C_\nu[f]\sigma_{\mathrm{HI}}\mathrm{d}\nu \approx \int C_\nu[f]\sigma_{\mathrm{HI}}\mathrm{d}\nu - \frac{\mathbf{n}\cdot\mathbf{v}}{c}\alpha_v H f_{\mathrm{LL}}, \tag{A19}$$

$$\text{where} \quad \alpha_v = \frac{1}{Hf_{\mathrm{LL}}}\int \frac{\partial C_\nu[f]}{\partial\ln\nu}\sigma_{\mathrm{HI}}\mathrm{d}\nu; \tag{A20}$$

the first term is the same as in our original calculation. The integral in the second term needs be evaluated only at zero-order, for which we can use the background (zero-order) Boltzmann equation (A8) in the form

$$\frac{1}{H}\frac{\partial C_\nu[f_0]}{\partial\ln\nu} = 2\frac{\partial f_0}{\partial\ln\nu} - \frac{\partial^2 f_0}{\partial(\ln\nu)^2}. \tag{A21}$$

With the above, again using the $z = 2.3$ spectrum from Ref. [21], I obtain $\alpha_v = -12.0$.

The only other remaining gravitational term is the gradient term in (A4) and (A11); this and the velocity term discussed above appear in the effective opacity which now reads

$$\kappa_{\mathrm{tot}} = \bar{\sigma}_{\mathrm{HI}}n_{\mathrm{HI}} + \bar{\kappa}_{\mathrm{clump}} + \frac{H}{c}\left(\alpha_z + 2 + \frac{\mathbf{n}\cdot\mathbf{v}}{c}\alpha_v\right) + \frac{\alpha_z+1}{a}\mathbf{n}\cdot\nabla\phi. \tag{A22}$$

Using the Einstein constraint equation in the form (A18) we can update our expression (29) for the fractional variations in κ_{tot}:

$$\tilde{\delta}_{\kappa_{\mathrm{tot}}} = \beta_{\mathrm{HI}}\tilde{\delta}_{n_{\mathrm{HI}}} + \beta_{\mathrm{clump}}\tilde{\delta}_{\bar{\kappa}_{\mathrm{clump}}} + \beta_\phi\frac{i\mathbf{n}\cdot\mathbf{k}\phi}{a\kappa_{\mathrm{tot},0}}, \tag{A23}$$

$$\text{where} \quad \beta_\phi = \alpha_z + 1 - \frac{2}{3}\alpha_v \approx 10.6. \tag{A24}$$

With this updated definition, Eq. (28) remains valid. Integrating (28) over $\mathbf{n}$ to obtain the solution for $\tilde{\delta}_\Gamma$ is slightly more involved because angular dependence now appears on the numerator as well as denominator; I obtain

$$\tilde{\delta}_\Gamma = -\beta_\phi\phi + [(1-\beta_{\mathrm{HI}}\beta_{\mathrm{r}})j - \beta_{\mathrm{HI}}(1-\beta_{\mathrm{r}})\tilde{\delta}_{n_{\mathrm{HI}}} - \beta_{\mathrm{clump}}\tilde{\delta}_{\bar{\kappa}_{\mathrm{clump}}} + \beta_{\mathrm{HI}}\beta_{\mathrm{r}}\tilde{\delta}_\Gamma + \beta_\phi\phi]S(k). \tag{A25}$$

The relation between local HI density and radiation fluctuations is still specified by Eq. (31), which allows us to find the solution for the HI fluctuations:

$$\tilde{\delta}_{n_{\mathrm{HI}}} = \frac{\tilde{\delta}_{n_{\mathrm{HI,u}}} + \beta_\phi\phi(k) - [\tilde{\delta}_{j,\mathrm{eff}} + \beta_\phi\phi(k)]S(k)}{1-\beta_{\mathrm{HI}}S(k)}, \tag{A26}$$

where $\tilde{\delta}_{n_{\mathrm{HI,u}}}$ is the HI density fluctuations in the absence of radiation inhomogeneities. Writing the difference between Eqs. (A26) and (32) as $\Delta\tilde{\delta}_{n_{\mathrm{HI}}}$, one can define the change to the HI bias from the gravitational and Doppler effects:

$$\Delta b_{\mathrm{HI}} = \frac{\Delta\tilde{\delta}_{n_{\mathrm{HI}}}}{\tilde{\delta}_\rho} = \beta_\phi\frac{1-S(k)}{1-\beta_{\mathrm{HI}}S(k)}\frac{-3a^2H^2}{2c^2k^2+6a^2H^2}, \tag{A27}$$

where I have made use of Eq. (A17). This function is plotted in Fig. 7 (dashed line), where it can be seen that even on gigaparsec scales it reaches a maximum shift of around -0.05, a tiny change in the bias (compare to Fig. 2). Note that the shot-noise component is unaffected.

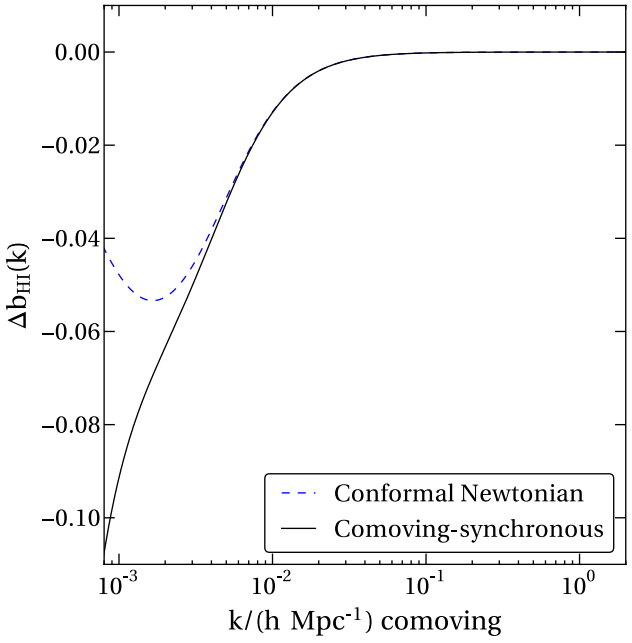

FIG. 7 (color online). The change in the HI bias arising from the velocity and potential terms in the radiative transfer. The solid and dashed lines show the conformal-Newtonian (A27) and comoving-synchronous expressions respectively. The effects are extremely minor as anticipated.

It may be more natural to think of the bias on large scales in another gauge—it is more plausible, in particular, to imagine that $b_{\rm HI,u}$ and b_j are scale invariant in the comoving-synchronous than in the conformal-Newtonian gauge [69]. (Ultimately one ought to derive gauge-invariant observables, but for the reasons outlined in the conclusions, that is beyond the scope of the current work.) The gauge transformation consists of a small coordinate transformation $(t,\mathbf{x})\to(t+T,\mathbf{x}+\mathbf{X})$; to reach the comoving-synchronous gauge one uses the freedom to eliminate the peculiar velocities in the coordinate frame. Following this through for any quantity Z, assuming $Z_0\propto a^{-q}$, one finds that

$$\delta_Z^N=\delta_Z^S-2q\phi/3, \tag{A28}$$

where superscripts N and S stand for perturbations in conformal-Newtonian and comoving-synchronous gauges respectively. Equation (A28) can be used to transform the conformal-Newtonian expression (A26) into the synchronous equivalent, giving

$$\tilde\delta^S_{n_{\rm HI}}=\frac{\tilde\delta_{n_{\rm HI},\rm u}{}^S+\beta_\phi\phi-[\tilde\delta_j^S+\beta_\phi\phi+\frac{2}{3}(\beta_{\rm HI}q_{\rm HI}-q_{j,\rm eff})\phi]S(k)}{1-\beta_{\rm HI}S(k)}. \tag{A29}$$

Then, in the synchronous gauge, we have

$$\Delta b^S_{\rm HI}=\frac{-3a^2H^2}{2c^2k^2}\frac{\beta_\phi-[\beta_\phi+\frac{2}{3}(\beta_{\rm HI}q_{\rm HI}-q_{j,\rm eff})]S(k)}{1-\beta_{\rm HI}S(k)}, \tag{A30}$$

where $\Delta b^S_{\rm HI}\equiv\Delta\tilde\delta^S_{n_{\rm HI}}/\tilde\delta^S_\rho$, and the result has been obtained using the relation between Newtonian potential and synchronous-gauge density,

$$\phi=\frac{-3a^2H^2}{2c^2k^2}\tilde\delta^S. \tag{A31}$$

To gain a quantitative picture we must estimate $q_{\rm HI}$ and q_j. Note that for any quantity Y composed of a linear sum of components, $Y=\sum_i Y_i$, one has

$$q_Y=\frac{\mathrm{d}\ln Y}{\mathrm{d}\ln a}=\frac{1}{Y}\sum_i\frac{\mathrm{d}Y_i}{\mathrm{d}\ln a}=\sum_i\beta_i q_{Y_i}, \tag{A32}$$

where $\beta_i=Y_i/Y$. It therefore follows from using the background equilibrium (27) that

$$q_j\approx[\beta_{\rm HI}(1-\beta_r)+\beta_{\rm clump}]q_{\rm HI}-\frac{3}{2}[\beta_z+\beta_V]\approx 2.7;$$
$$q_{j,\rm eff}\approx(1-\beta_{\rm HI}\beta_{\rm r})q_j+(\beta_{\rm HI}\beta_{\rm r}-\beta_{\rm clump})q_{\rm HI}\approx 2.3, \tag{A33}$$

where I have used $\mathrm{d}\ln H/\mathrm{d}\ln a\approx -3/2$ and $q_{\rm HI}\simeq 4.3$, the latter from Ref. [33].

Adopting these estimates, Eq. (A30) is plotted as a solid line in Fig. 7. The changes are larger than in the Newtonian gauge but still small. It is worth noting that the synchronous-gauge bias as derived above describes a different universe—it is not, in fact, related by a gauge transformation to the Newtonian case. This follows because I have formulated both descriptions assuming a constant large-scale bias as an input distribution; this assumption implies something physically different in the two different gauges. As previously stated, it is probably a more appropriate assumption in the synchronous than in the Newtonian gauge.

This, however, is a tangential question because the gravitational and Doppler effects are tiny in both gauges. The original decision to drop these terms is therefore shown to be strongly justified.

APPENDIX B: AN ESTIMATE OF $b_{\rm HI,u}$

To complete the calculation in the main text it was necessary to specify a value of the bias $b_{\rm HI,u}$ of HI in the absence of inhomogeneous radiative effects. One can estimate this from the photoionization equilibrium equations for a uniform field, coupled with a description of the temperature-density relation for the averaged IGM. Specifically, the uniform ionization equilibrium in the limit that only a trace of neutral HI survives specifies [70] that

$$n_{\rm HI} \propto \frac{\alpha(T)\rho^2}{\Gamma_0}, \tag{B1}$$

where I have assumed the local electron and proton densities are both proportional to the cosmic density ρ. Next write the equation of state $T \propto \rho^{\gamma-1}$ and approximate [42] $\alpha(T) \propto T^{-0.7}$; expanding both ρ and $n_{\rm HI}$ in terms of their background values and perturbations, one obtains

$$b_{\rm HI,u} = \frac{\delta_{n_{\rm HI,u}}}{\delta_\rho} = 2 - 0.7(\gamma - 1). \tag{B2}$$

For a value [47] $\gamma = 1.6$, this gives an estimate of $b_{\rm HI,u} \simeq 1.6$.

However, there is a slight inconsistency in the derivation above. The recombination rate actually depends on the strictly local value of the electron and proton densities, not on any linear-theory average on large scales. Depending on how small-scale clustering reflects large-scale density inhomogeneities, the assumption that the local density scales with the environmental density may fail. I therefore also estimated $b_{\rm HI,u}$ directly from a cosmological simulation with 256^3 dark matter and 256^3 gas particles in a 50 Mpc-side box. The code Gasoline [71] implements gravity, hydrodynamics, star formation feedback (which is likely of minor importance here) and a uniform UV field, the values for which I adopted from Ref. [21]. Much more careful work has been performed in simulating the forest by other authors [6,18,72] but they quote statistics on the flux field, which is related to the HI field by a nonlinear transformation and therefore does not directly tell us $b_{\rm HI,u}$.

Taking the output at $z = 2.3$, I interpolated the gas and dark matter particles back onto a 256^3 grid to obtain two 3D density maps, the first of HI and the second of total mass density with ~0.2 Mpc resolution. To study the behavior of the intergalactic medium, I flagged all cells with dark matter density less than ten times the cosmic mean. Only the flagged cells were subsequently used to produce a degraded map with 8^3 supercells, in which the mean dark matter and HI density of the flagged subcells was recorded.

This allows us to see the large-scale relationship between IGM overdensity and HI (Fig. 8). Each point represents an IGM supercell; the two axes correspond to dimensionless total mass overdensity and HI overdensity in the IGM, expressed as a fraction according to Eq. (1). The plots show

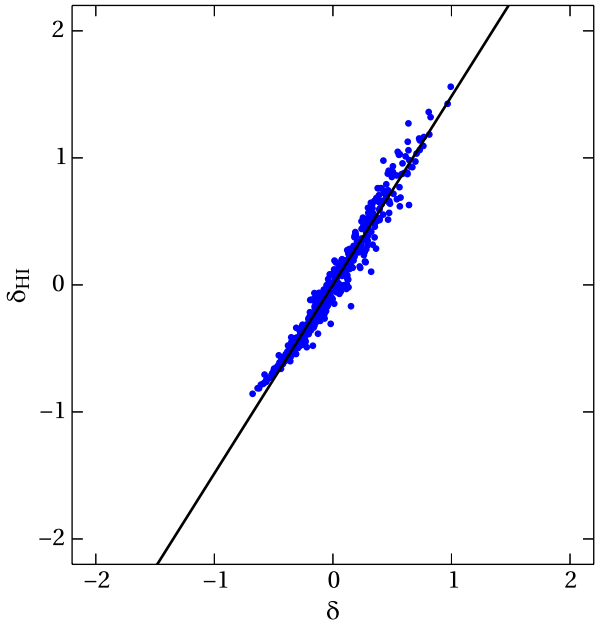

FIG. 8 (color online). The relationship between HI overdensity and total matter overdensity in the IGM; each point represents the values averaged in a cube of side length 6 Mpc comoving at $z = 2.3$. This derives from a simulation with uniform UV background. There is a very near-linear relationship between the two quantities, allowing the measurement of the slope (solid line), $b_{\rm HI,u} \approx 1.5$, which is a parameter entering the main calculation.

a very near-linear relationship between the total overdensity and the HI overdensity as expected. The slope of the line gives the bias, which is found to be $b_{\rm HI,u} = 1.48 \approx 1.5$. This is in fair agreement with the analytic estimate of 1.6 given above, given the multitude of uncertainties.

I tested that this result is reasonably insensitive to the size of the supercells and the IGM threshold density. Repeating the exercise with the IGM threshold at $\delta = 5$, for instance, gives $b_{\rm HI,u} = 1.45$; with the original threshold but 16^3 supercells, I retrieve $b_{\rm HI,u} = 1.43$, although the nonlinearity in the relation starts to become more prominent (as we are probing smaller scales) so the fit is less meaningful. For the illustrative purposes of this paper, adopting $b_{\rm HI,u} = 1.5$ seems to pin down the large-scale relationship to a sufficient accuracy.

[1] M. Rauch, Annu. Rev. Astron. Astrophys. **36**, 267 (1998).
[2] R. H. Becker, X. Fan, R. L. White, M. A. Strauss, V. K. Narayanan, R. H. Lupton, J. E. Gunn, J. Annis, N. A. Bahcall, J. Brinkmann *et al.*, Astrophys. J. **122**, 2850 (2001).
[3] X. Fan, V. K. Narayanan, M. A. Strauss, R. L. White, R. H. Becker, L. Pentericci, and H.-W. Rix, Astron. J. **123**, 1247 (2002).
[4] X. Fan, C. L. Carilli, and B. Keating, Annu. Rev. Astron. Astrophys. **44**, 415 (2006).

[5] R. A. C. Croft, D. H. Weinberg, M. Bolte, S. Burles, L. Hernquist, N. Katz, D. Kirkman, and D. Tytler, Astrophys. J. **581**, 20 (2002).
[6] P. McDonald, Astrophys. J. **585**, 34 (2003).
[7] A. Slosar, A. Font-Ribera, M. M. Pieri, J. Rich, J.-M. Le Goff, É. Aubourg, J. Brinkmann, N. Busca, B. Carithers, R. Charlassier *et al.*, J. Cosmol. Astropart. Phys. 09 (2011) 001.
[8] M. Viel, J. Lesgourgues, M. G. Haehnelt, S. Matarrese, and A. Riotto, Phys. Rev. D **71**, 063534 (2005).
[9] M. Viel, G. D. Becker, J. S. Bolton, M. G. Haehnelt, M. Rauch, and W. L. W. Sargent, Phys. Rev. Lett. **100**, 041304 (2008).
[10] A. Boyarsky, J. Lesgourgues, O. Ruchayskiy, and M. Viel, J. Cosmol. Astropart. Phys. 05 (2009) 012.
[11] A. Slosar, S. Ho, M. White, and T. Louis, J. Cosmol. Astropart. Phys. 10 (2009) 019.
[12] A. Slosar, V. Iršič, D. Kirkby, S. Bailey, N. G. Busca, T. Delubac, J. Rich, É. Aubourg, J. E. Bautista, V. Bhardwaj *et al.*, J. Cosmol. Astropart. Phys. 04 (2013) 026.
[13] A. Font-Ribera, D. Kirkby, N. Busca, J. Miralda-Escudé, N. P. Ross, A. Slosar, É. Aubourg, S. Bailey, V. Bhardwaj, J. Bautista *et al.*, arXiv:1311.1767.
[14] G. D. Becker, M. Rauch, and W. L. W. Sargent, Astrophys. J. **662**, 72 (2007).
[15] J. S. Bolton, M. Viel, T.-S. Kim, M. G. Haehnelt, and R. F. Carswell, Mon. Not. R. Astron. Soc. **386**, 1131 (2008).
[16] G. D. Becker, J. S. Bolton, M. G. Haehnelt, and W. L. W. Sargent, Mon. Not. R. Astron. Soc. **410**, 1096 (2011).
[17] R. A. C. Croft, D. H. Weinberg, N. Katz, and L. Hernquist, Astrophys. J. **495**, 44 (1998).
[18] P. McDonald, J. Miralda-Escudé, M. Rauch, W. L. W. Sargent, T. A. Barlow, R. Cen, and J. P. Ostriker, Astrophys. J. **543**, 1 (2000).
[19] J. S. Bolton, G. D. Becker, M. G. Haehnelt, and M. Viel, Mon. Not. R. Astron. Soc. **438**, 2499 (2014).
[20] C.-A. Faucher-Giguère, A. Lidz, M. Zaldarriaga, and L. Hernquist, Astrophys. J. **703**, 1416 (2009).
[21] F. Haardt and P. Madau, Astrophys. J. **746**, 125 (2012).
[22] A. Maselli and A. Ferrara, Mon. Not. R. Astron. Soc. **364**, 1429 (2005).
[23] L. Zuo, Mon. Not. R. Astron. Soc. **258**, 36 (1992).
[24] J. A. Kollmeier, D. H. Weinberg, R. Davé, and N. Katz, Astrophys. J. **594**, 75 (2003).
[25] A. Meiksin and M. White, Mon. Not. R. Astron. Soc. **350**, 1107 (2004).
[26] J. A. Kollmeier, J. Miralda-Escudé, R. Cen, and J. P. Ostriker, Astrophys. J. **638**, 52 (2006).
[27] A. Mesinger and S. Furlanetto, Mon. Not. R. Astron. Soc. **400**, 1461 (2009).
[28] M. White, A. Pope, J. Carlson, K. Heitmann, S. Habib, P. Fasel, D. Daniel, and Z. Lukic, Astrophys. J. **713**, 383 (2010).
[29] A. Font-Ribera, E. Arnau, J. Miralda-Escudé, E. Rollinde, J. Brinkmann, J. R. Brownstein, K.-G. Lee, A. D. Myers, N. Palanque-Delabrouille, I. Pâris *et al.*, J. Cosmol. Astropart. Phys. 05 (2013) 018.
[30] R. A. C. Croft, Astrophys. J. **610**, 642 (2004).
[31] P. McDonald, U. Seljak, R. Cen, P. Bode, and J. P. Ostriker, Mon. Not. R. Astron. Soc. **360**, 1471 (2005).
[32] M. McQuinn, L. Hernquist, A. Lidz, and M. Zaldarriaga, Mon. Not. R. Astron. Soc. **415**, 977 (2011).
[33] G. C. Rudie, C. C. Steidel, A. E. Shapley, and M. Pettini, Astrophys. J. **769**, 146 (2013).
[34] K. S. Dawson, D. J. Schlegel, C. P. Ahn, S. F. Anderson, É. Aubourg, S. Bailey, R. H. Barkhouser, J. E. Bautista, A. Beifiori, A. A. Berlind *et al.*, Astron. J. **145**, 10 (2013).
[35] N. G. Busca, T. Delubac, J. Rich, S. Bailey, A. Font-Ribera, D. Kirkby, J.-M. Le Goff, M. M. Pieri, A. Slosar, É. Aubourg *et al.*, Astron. Astrophys. **552**, A96 (2013).
[36] M. McQuinn, S. P. Oh, and C.-A. Faucher-Giguère, Astrophys. J. **743**, 82 (2011).
[37] J. Zhang, L. Hui, and Z. Haiman, Mon. Not. R. Astron. Soc. **375**, 324 (2007).
[38] P. A. R. Ade, N. Aghanim, C. Armitage-Caplan, M. Arnaud, M. Ashdown, F. Atrio-Barandela, J. Aumont, C. Baccigalupi, A. J. Banday *et al.* (Planck Collaboration), arXiv:1303.5076.
[39] F. Haardt and P. Madau, Astrophys. J. **461**, 20 (1996).
[40] J. X. Prochaska, P. Madau, J. M. O'Meara, and M. Fumagalli, Mon. Not. R. Astron. Soc. **438**, 476 (2014).
[41] A. Pontzen, F. Governato, M. Pettini, C. M. Booth, G. Stinson, J. Wadsley, A. Brooks, T. Quinn, and M. Haehnelt Mon. Not. R. Astron. Soc. **390**, 1349 (2008).
[42] J. H. Black, Mon. Not. R. Astron. Soc. **197**, 553 (1981).
[43] A. Loeb and S. Furlanetto, *The First Galaxies in the Universe*, Princeton Series in Astrophysics (Princeton University Press, Princeton, NJ, 2013).
[44] K. Lai, A. Lidz, L. Hernquist, and M. Zaldarriaga, Astrophys. J. **644**, 61 (2006).
[45] J. S. Bolton, M. G. Haehnelt, M. Viel, and R. F. Carswell, Mon. Not. R. Astron. Soc. **366**, 1378 (2006).
[46] S. R. Furlanetto, Astrophys. J. **703**, 702 (2009).
[47] L. Hui and N. Y. Gnedin, Mon. Not. R. Astron. Soc. **292**, 27 (1997).
[48] C. Bonvin and R. Durrer, Phys. Rev. D **84**, 063505 (2011).
[49] S. M. Croom, B. J. Boyle, T. Shanks, R. J. Smith, L. Miller, P. J. Outram, N. S. Loaring, F. Hoyle, and J. da Ângela, Mon. Not. R. Astron. Soc. **356**, 415 (2005).
[50] M. White, A. D. Myers, N. P. Ross, D. J. Schlegel, J. F. Hennawi, Y. Shen, I. McGreer, M. A. Strauss, A. S. Bolton, J. Bovy *et al.*, Mon. Not. R. Astron. Soc. **424**, 933 (2012).
[51] K. L. Adelberger, C. C. Steidel, M. Pettini, A. E. Shapley, N. A. Reddy, and D. K. Erb, Astrophys. J. **619**, 697 (2005).
[52] N. A. Reddy and C. C. Steidel, Astrophys. J. **692**, 778 (2009).
[53] S. Cole and N. Kaiser, Mon. Not. R. Astron. Soc. **237**, 1127 (1989).
[54] A. Pontzen, R. Roskar, G. Stinson, and R. Woods, pynbody: NBody /SPH analysis for python (May 2013), astrophysics Source Code Library, ascl:1305.002.
[55] B. P. Moster, T. Naab, and S. D. M. White, Mon. Not. R. Astron. Soc. **428**, 3121 (2012).
[56] P. F. Hopkins, G. T. Richards, and L. Hernquist, Astrophys. J. **654**, 731 (2007).
[57] N. P. Ross, I. D. McGreer, M. White, G. T. Richards, A. D. Myers, N. Palanque-Delabrouille, M. A. Strauss, S. F. Anderson, Y. Shen, W. N. Brandt *et al.*, Astrophys. J. **773**, 14 (2013).

[58] A. Lewis, A. Challinor, and A. Lasenby, Astrophys. J. **538**, 473 (2000).
[59] A. Font-Ribera, J. Miralda-Escudé, E. Arnau, B. Carithers, K.-G. Lee, P. Noterdaeme, I. Pâris, P. Petitjean, J. Rich, E. Rollinde *et al.*, J. Cosmol. Astropart. Phys. 11 (2012) 059.
[60] P. McDonald, U. Seljak, S. Burles, D. J. Schlegel, D. H. Weinberg, R. Cen, D. Shih, J. Schaye, D. P. Schneider, N. A. Bahcall *et al.*, Astrophys. J. Suppl. Ser. **163**, 80 (2006).
[61] L. Anderson, E. Aubourg, S. Bailey, F. Beutler, A. S. Bolton, J. Brinkmann, J. R. Brownstein, C.-H. Chuang, A. J. Cuesta, K. S. Dawson *et al.*, Mon. Not. R. Astron. Soc. **439**, 83 (2014).
[62] N. Kaiser, Mon. Not. R. Astron. Soc. **227**, 1 (1987).
[63] J. A. Peacock, *Cosmological Physics* (Cambridge University Press, Cambridge, 1999).
[64] J. R. Pritchard and A. Loeb, Rep. Prog. Phys. **75**, 086901 (2012).
[65] S. R. Furlanetto, M. Zaldarriaga, and L. Hernquist, Astrophys. J. **613**, 1 (2004).
[66] A. D'Aloisio, J. Zhang, P. R. Shapiro, and Y. Mao, Mon. Not. R. Astron. Soc. **433**, 2900 (2013).
[67] Y. Mao, A. D'Aloisio, J. Zhang, and P. R. Shapiro, Phys. Rev. D **88**, 081303 (2013).
[68] V. F. Mukhanov, H. A. Feldman, and R. H. Brandenberger, Phys. Rep. **215**, 203 (1992).
[69] D. Jeong, F. Schmidt, and C. M. Hirata, Phys. Rev. D **85**, 023504 (2012).
[70] D. E. Osterbrock and G. J. Ferland, *Astrophysics of Gaseous Nebulae and Active Galactic Nuclei* (University Science Books, Sausalito, CA, 2006).
[71] J. W. Wadsley, J. Stadel, and T. Quinn, New Astron. **9**, 137 (2004).
[72] M. Viel and M. G. Haehnelt, Mon. Not. R. Astron. Soc. **365**, 231 (2006).
[73] Gontcho A Gontcho, Miralda-Escudè, and Busca (to be published).