Title

Quantitative Constraints on the Reionization History from the IGM Damping Wing Signature in Two Quasars at z > 7

Journal

Astrophysical Journal, 864(2)

ISSN

0004-637X

Authors

Davies, FB
Hennawi, JF
Bañados, E
et al.

Publication Date

2018-09-10

DOI

10.3847/1538-4357/aad6dc

Peer reviewed

Draft version February 19, 2018
Typeset using LaTeX twocolumn style in AASTeX62

Quantitative Constraints on the Reionization History from the IGM Damping Wing Signature in Two Quasars at $z > 7$

Frederick B. Davies,[1] Joseph F. Hennawi,[1] Eduardo Bañados,[2,*] Zarija Lukić,[3] Roberto Decarli,[4,5] Xiaohui Fan,[6] Emanuele P. Farina,[1] Chiara Mazzucchelli,[4] Hans-Walter Rix,[4] Bram P. Venemans,[4] Fabian Walter,[4] Feige Wang,[1,7,8] and Jinyi Yang[6,7,8]

[1] Department of Physics, University of California, Santa Barbara, CA 93106-9530, USA
[2] The Observatories of the Carnegie Institution for Science, 813 Santa Barbara Street, Pasadena, California 91101, USA
[3] Lawrence Berkeley National Laboratory, CA 94720-8139, USA
[4] Max Planck Institut für Astronomie, Königstuhl 17, D-69117 Heidelberg, Germany
[5] INAF–Osservatorio Astronomico di Bologna, via Gobetti 93/3, 40129 Bologna, Italy
[6] Steward Observatory, The University of Arizona, 933 North Cherry Avenue, Tucson, Arizona 85721-0065, USA
[7] Department of Astronomy, School of Physics, Peking University, Beijing 100871, China
[8] Kavli Institute for Astronomy and Astrophysics, Peking University, Beijing 100871, China

ABSTRACT

During reionization, neutral hydrogen in the intergalactic medium (IGM) imprints a damping wing absorption feature on the spectrum of high-redshift quasars. A detection of this signature provides compelling evidence for a significantly neutral Universe, and enables measurements of the hydrogen neutral fraction $x_{\rm HI}(z)$ at that epoch. Obtaining reliable quantitative constraints from this technique, however, is challenging due stochasticity induced by the patchy inside-out topology of reionization, degeneracies with quasar lifetime, and the unknown unabsorbed quasar spectrum close to rest-frame Lyα. We combine a large-volume semi-numerical simulation of reionization topology with 1D radiative transfer through high-resolution hydrodynamical simulations of the high-redshift Universe to construct models of quasar transmission spectra during reionization. Our state-of-the-art approach captures the distribution of damping wing strengths in biased quasar halos that should have reionized earlier, as well as the erosion of neutral gas in the quasar environment caused by its own ionizing radiation. Combining this detailed model with our new technique for predicting the quasar continuum and its associated uncertainty, we introduce a Bayesian statistical method to jointly constrain the neutral fraction of the Universe and the quasar lifetime from individual quasar spectra. We apply this methodology to the spectra of the two highest redshift quasars known, ULAS J1120+0641 and ULAS J1342+0928, and measured volume-averaged neutral fractions $\langle x_{\rm HI}\rangle(z = 7.09) = 0.48^{+0.26}_{-0.26}$ and $\langle x_{\rm HI}\rangle(z = 7.54) = 0.60^{+0.20}_{-0.23}$ (posterior medians and 68% credible intervals) when marginalized over quasar lifetimes of $10^3 \le t_{\rm q} \le 10^8$ years.

1. INTRODUCTION

The epoch of reionization was a landmark event in the history of the Universe when the cumulative number of ionizing photons escaping from the first stars, galaxies, and quasars surpassed the number of hydrogen atoms in the intergalactic medium (IGM). Our knowledge of reionization is bounded by the presence of transmission in the Lyα forest at $z \lesssim 6$ (Fan et al. 2006), and an integral constraint from the electron scattering optical depth of the cosmic microwave background (CMB) which constrains the volume of ionized IGM between the present day and $z \sim 1100$ (Planck Collaboration et al. 2016a) that suggests a characteristic reionization redshift of $z_{\rm re} = 6.4$–9.7 (95% credible interval, Planck Collaboration et al. 2016b). With only these constraints, the detailed reionization history – reflecting the nature and evolution of sources of ionizing photons – is still highly uncertain and model-dependent.

The discovery and deep follow-up spectroscopy of quasars with redshifts greater than six provided the first look at the IGM approaching the epoch of reionization (e.g. Fan et al. 2001, 2003; Becker et al. 2001; White et al. 2003). While Gunn-Peterson troughs (Gunn & Peterson 1965) in the Lyα and Lyβ forests of these quasars due to the presence of neutral hydrogen in the IGM may be signatures of ongoing reionization, they can only place lower limits on the volume-averaged hy-

* Carnegie-Princeton Fellow

drogen neutral fraction of $\langle x_{\rm H\,I}\rangle \gtrsim 10^{-4}$ (e.g. Fan et al. 2006). The sizes of the transparent proximity zones of these quasars have also been analyzed in the context of expanding Strömgren spheres in a (partially-)neutral IGM (Cen & Haiman 2000; Wyithe et al. 2005; Mesinger & Haiman 2007; Schroeder et al. 2013), but as recently demonstrated by Eilers et al. (2017), the sizes alone may be insensitive to the ionization state of the IGM. A more sensitive, and perhaps definitive, probe of neutral gas in the IGM is the Lyα damping wing (Miralda-Escudé 1998) which suppresses the quasar continuum redward of rest-frame Lyα.

The first quasar with a claimed damping wing signal was ULAS J1120+0641 (Mortlock et al. 2011) at $z = 7.09$, but the inferred constraints on $\langle x_{\rm H\,I}\rangle$ vary between different analyses. These differences are in part due to differences in physical models for the proximity zone and/or damping wing. One approach to constrain $\langle x_{\rm HI}\rangle$ is to fit the Lyα transmission spectrum with the analytic model of Miralda-Escudé (1998), as performed by Mortlock et al. (2011), but this formula does not include the substantial resonant Lyα absorption by residual H I inside of the proximity zone (see, e.g., Bolton et al. 2011; Keating et al. 2015), and the IGM outside the proximity zone is typically assumed to have a completely uniform ionization state instead of a more realistic patchy topology of ionized bubbles (Furlanetto et al. 2004). Greig et al. (2017a) constrained $\langle x_{\rm HI}\rangle$ from the damping wing of ULAS J1120+0641 using large-volume semi-numerical simulations of the reionization topology (Mesinger et al. 2016) to predict the distribution of damping wing strengths as a function of $\langle x_{\rm HI}\rangle$, but they only considered wavelengths redward of Lyα. The size of the proximity zone, and the strength of the damping wing, are sensitive to the quasar lifetime (Bolton et al. 2011; Keating et al. 2015) which has an uncertainty of several orders of magnitude (Martini 2004; Eilers et al. 2017).

Another complication is that uncertainties and differences between methodologies for estimating the intrinsic quasar continuum (Kramer & Haiman 2009; Greig et al. 2017b) can be similar in strength to the damping wing signal itself. Further exacerbating this challenge is the fact that the spectral properties of quasars at $z \gtrsim 6.5$, in particular their C IV emission line blueshifts, are often extreme outliers of the distribution of lower redshift quasars (Mazzucchelli et al. 2017). In light of its large C IV blueshift, Mortlock et al. (2011) estimated the intrinsic spectrum of ULAS J1120+0641 via stacking a sample of SDSS quasar spectra with matched C IV emission line properties, but did not test the accuracy of this continuum estimation method on quasars with known continua. Bosman & Becker (2015) found that a more closely matched sample of lower-redshift quasars (i.e., without any damping wing signal) with carefully selected C IV emission line strengths and blueshifts had Lyα spectral shapes consistent with the observed ULAS J1120+0641 spectrum. In contrast, the predictive continuum model of Greig et al. (2017b), as demonstrated by Greig et al. (2017a), appears to prefer a much stronger intrinsic Lyα profile, suggesting instead that the damping wing signal is quite strong.

Analysis of the recently discovered quasar ULAS J1342+0928 (Bañados et al. 2018) at $z = 7.54$ suggests that it exhibits a much more prominent damping wing absorption signal than ULAS J1120+0641, consistent with a predominantly neutral IGM. However, its C IV line exhibits a blueshift more than twice that of ULAS J1120+0641, and so only a very small number of similar quasars exist in lower redshift samples. Bañados et al. (2018) estimated the intrinsic spectrum of ULAS J1342+0928 by constructing a composite spectrum from 46 SDSS/BOSS quasars with similar C IV emission line properties and estimated the uncertainty via measuring the residuals between the composite and its constituent quasar spectra. They then derived their fiducial constraints on $\langle x_{\rm HI}\rangle$ in the surrounding IGM using the Miralda-Escudé (1998) model for the damping wing shape. In this work we describe one of the alternative models ("Model B") from Bañados et al. (2018) in more detail.

A complete model of the proximity zone and damping wing region of quasar spectra requires an estimate of the intrinsic quasar continuum, the uncertainty in the quasar continuum model, a model for the small-scale density fluctuations in the IGM, a realistic description of patchy reionization topology surrounding the massive dark matter halos that host luminous quasars, and time-dependent radiative transfer of ionizing photons from the quasar along the line of sight. The goal of this work is to put all of these pieces together for the first time to forward model mock quasar spectra and develop a statistical method to constrain the volume-averaged IGM neutral fraction and quasar lifetime from individual quasar spectra.

In Davies et al. (2018, henceforth Paper I), we developed a Principal Component Analysis (PCA)-based approach with a training set of $> 10,000$ quasar spectra from the SDSS/BOSS DR12Q catalog (Pâris et al. 2017) to predict the "blue-side" quasar continuum, at rest-frame wavelengths $1175 < \lambda_{\rm rest} < 1280$ Å, from the "red-side" spectrum, covering $1280 < \lambda_{\rm rest} < 2850$ Å. We quantified the covariant uncertainties by testing the method on the training set, finding that for a typical quasar the relative error of our predicted continua is $\sim$ 6–12% at rest-frame wavelengths most sensitive to damping wing absorption. Finally, we demonstrated the applicability of our method on the two known $z > 7$ quasars: ULAS J1120+0641, and ULAS J1342+0928. While these quasars represent outliers from the distribution of typical quasars in SDSS/BOSS, we have calibrated the uncertainty on the blue-side predicted continua from custom subsets of "nearest-neighbor" quasars in the training set that have similar red-side spectra to each quasar separately.

In this work, we present a hybrid model for quasar proximity zone and damping wing structures during reionization, and a statistical method to perform Bayesian parameter inference on high-redshift quasar spectra in conjunction with the PCA quasar continuum model from Paper I. We apply these new methods to constrain $\langle x_{\rm HI} \rangle$ from the spectra of ULAS J1342+0928 at $z = 7.54$ and ULAS J1120+0641 at $z = 7.09$. We find strong evidence for a substantially neutral IGM at $z > 7$, especially at $z = 7.54$, consistent with the latest constraints from the CMB (Planck Collaboration et al. 2016b).

The rest of the paper is structured as follows. In § 2, we briefly summarize the Principal Component Analysis (PCA) method for predicting the intrinsic blue-side quasar continuum from the red-side spectrum from Paper I. In § 3 we describe our hybrid model for quasar proximity zones and the Lyα damping wing, combining ionizing radiative transfer simulations (Davies et al. 2016) through density field skewers from high-resolution hydrodynamical simulations with semi-numerical simulations of the inside-out reionization topology around massive halos. In § 4 we describe our methodology for performing Bayesian parameter inference from millions of forward-modeled mock spectra. In § 5 we show the results of our analysis on the two quasars known at $z > 7$: ULAS J1342+0928 and ULAS J1120+0641. Finally, in § 6 we conclude with a discussion of the implications of the neutral fraction constraints from the two quasars on the reionization history of the Universe, and describe avenues for future investigation of existing quasar samples.

In this work we assume a flat ΛCDM cosmology with $h = 0.685$, $\Omega_b = 0.047$, $\Omega_m = 0.3$, $\Omega_\Lambda = 0.7$, and $\sigma_8 = 0.8$.

2. PCA CONTINUUM MODEL

We adopt the method for predicting the intrinsic blue-side quasar continuum ($1175 < \lambda_{\rm rest} < 1280$ Å) from the observed red-side spectrum ($1280 < \lambda_{\rm rest} < 2850$ Å) from Paper I, which we briefly summarize below.

To construct the PCA model, we selected a sample of 12,764 quasars from the BOSS DR12Q catalog (Pâris et al. 2017) at $2.09 < z_{\rm pipe} < 2.51$ with S/N > 7 at $\lambda_{\rm rest} = 1290$ Å, and fit each spectrum with an automated, piecewise spline fitting method designed to recover smooth quasar continua in the presence of absorption lines (Young et al. 1979; Carswell et al. 1982; Dall'Aglio et al. 2008). In this redshift range, the BOSS spectra cover the entire spectral range from Lyα to Mg II. We further processed the splined spectra by median stacking each one with its 40 nearest neighbors to clean up residual artifacts such as strong associated absorption. We then computed principal component spectra (or "basis spectra") from these median stacks of these spline fit spectra with the standard PCA approach using `scikit-learn` (Pedregosa et al. 2011), albeit in *log-space*. That is, the logarithm of each quasar spectrum is represented by a sum of basis spectra $\mathbf{A}_i$ with corresponding weights a_i,

$$\log \mathbf{F} \approx \langle \log \mathbf{F} \rangle + \sum_i^n a_i \mathbf{A}_i, \tag{1}$$

which in linear space becomes a *product* of basis spectra raised to powers,

$$\mathbf{F} \approx \mathrm{e}^{\langle \log \mathbf{F} \rangle} \prod_i^n \mathrm{e}^{a_i \mathbf{A}_i}. \tag{2}$$

This log-space decomposition naturally accounts for the continuum slope variations between quasars, which dominates the total variance in flux space. For our analysis, we decomposed the red-side and blue-side spectra independently, keeping 10 red-side ($\mathbf{R}_i$) and 6 blue-side ($\mathbf{B}_i$) basis spectra.

We then found the best-fit red-side coefficients r_i for each (original, not spline fit) quasar spectrum in the training set via χ^2 minimization while simultaneously fitting for a *template redshift* $z_{\rm temp}$, allowing us to place each quasar onto a consistently defined rest-frame. The blue-side coefficients b_i for each training set quasar were then found by fitting the blue-side (in the $z_{\rm temp}$ frame) spline fit continua assuming constant noise. From the sets of r_i and b_i for all training set quasars, $\mathbf{r}$ and $\mathbf{b}$, we follow Suzuki et al. (2005) and Pâris et al. (2011) and compute the *projection matrix* $\mathbf{X}$ by finding the least-squares solution to the linear equation,

$$\mathbf{b} = \mathbf{r} \cdot \mathbf{X}. \tag{3}$$

After fitting the 10 r_i of an arbitrary quasar spectrum, we can "project" to the corresponding 6 b_i (and thus reconstruct the blue-side spectrum) via a dot product with the 10×6 projection matrix $\mathbf{X}$.

By testing our PCA procedure on the training set, we found that the relative error in the projected blue-side continua (which we refer to as the "blue-side prediction") is $\sim$ 6–12% in the region of the spectrum most useful for proximity zone and damping wing analyses ($1210 < \lambda_{\rm rest} < 1240$), with a mean bias $\lesssim 1\%$. The continuum error was found to be highly covariant across wide regions of the spectrum corresponding to regions associated with broad emission lines. However, as mentioned above, the spectra of $z \gtrsim 6.5$ quasars are known to be irregular compared to typical quasar spectra at lower redshift – in particular, they exhibit large C IV blueshifts relative to lower ionization lines in the spectrum such as Mg II (Mortlock et al. 2011; Mazzucchelli et al. 2017; Bañados et al. 2018). The uncertainty of the predicted continua for these atypical spectra may not be properly represented by the average uncertainty for all spectra.

For individual quasars, we can estimate a more accurate uncertainty by measuring the continuum errors for quasars with similar red-side spectra. We defined a distance D_r in the space of red-side PCA coefficients r_i by

$$D_r \equiv \sqrt{\sum_i^{N_{\rm PCA,r}} \left(\frac{\Delta r_i}{\sigma(r_i)}\right)^2}, \qquad (4)$$

where $N_{\rm PCA,r}$ is the number of red-side PCA basis vectors, Δr_i is the difference between r_i values, and $\sigma(r_i)$ is the standard deviation of r_i values in the training set. In Paper I, we measured the predicted continuum errors for the 1% of training set spectra with the lowest D_r to each of the $z > 7$ quasars, allowing us to estimate a custom continuum uncertainty for each $z > 7$ quasar. The predictions for these "similar" quasars tended to be slightly less uncertain and somewhat more biased than the training set as a whole.

For the statistical analysis of quasar spectra that follows, we require the ability to generate mock realizations of the continuum prediction error, which we denote as ϵ_C following Paper I. We assume a multivariate Gaussian distribution for the relative continuum error, with the mean and covariance determined from the prediction errors measured for similar, i.e. the 1% nearest neighbor, quasars as described above. We then use draws from these custom error distributions to generate forward-modeled mock spectra, described in more detail in § 4.

3. HYBRID MODEL OF QUASAR PROXIMITY ZONES AND DAMPING WINGS DURING REIONIZATION

With our predictive model for intrinsic quasar continua and their errors in place, we now must develop a physical model for quasar proximity zones and damping wings. We construct a hybrid model with three parts:

1. High-resolution density field from a large-volume hydrodynamical simulation (Lukić et al. 2015).

2. Semi-numerical simulations of reionization morphology (Mesinger et al. 2011; Davies & Furlanetto, in prep.).

3. One-dimensional ionizing radiative transfer of hydrogen- and helium-ionizing photons emitted by the quasar (Davies et al. 2016).

In this section, we describe these model components in detail.

3.1. *Nyx Hydrodynamical Simulation*

The first ingredient of our hybrid model is the small-scale structure of the IGM, which determines the absorption features inside the quasar proximity zone. We use density, velocity, and temperature fields from the $z = 7.0$ output of a `Nyx` hydrodynamical simulation (Almgren et al. 2013), 100 Mpc/h (comoving) on a side with 4096^3 dark matter particles and 4096^3 baryon grid cells (see also Lukić et al. 2015). Dark matter halos were selected via an algorithm that finds topologically-connected regions above 138 times the mean density (Lukič et al., in prep.), which is described in Sorini et al. (2017). We extract 1200 axis-aligned skewers from the centers of the 200 most massive halos, corresponding to halo masses $M_{\rm h} \gtrsim 2 \times 10^{11}$ M$_\odot$. The simulation, optimized for studying the Lyα forest, was run on a fixed, Eulerian grid, and lacks prescriptions for star formation or feedback (Lukić et al. 2015) which are required to characterize the circumgalactic medium of massive dark matter halos. Nevertheless, they should be adequate for our purposes, because our primary goal is to capture the larger-scale overdensity surrounding these halos on the relatively large $\sim$ 1–2 proper Mpc scales covered by the proximity zones of the $z > 7$ quasars in our analysis (compared to the halo virial radius, $\sim$ 50 proper kpc).

We re-scale the gas density of the skewers by $(1+z)^3$ depending on which quasar we are simulating. We leave the computation of custom-redshift outputs (i.e. matched to the quasar redshifts) of large hydrodynamical simulations to future work, but note that between $z = 7$ and the redshifts of the two quasars we focus on here ($z = 7.09, 7.54$) the evolution of the overdensity field should be relatively unimportant.

3.2. *Semi-Numerical Reionization Simulations with 21cmFAST*

The second ingredient of our hybrid model is the large-scale morphology of reionization around massive quasar-hosting halos. To compute realistic ionization fields on large scales, we adopt a modified version of the semi-numerical reionization code `21cmFAST`[1] (Mesinger et al. 2011), to be presented in further detail in Davies & Furlanetto (in prep.). The `21cmFAST` code computes the fraction of material that has collapsed into dark matter halos, $f_{\rm coll}$, following conditional Press-Schechter (Lacey & Cole 1993) applied to a non-linear density field computed using the Zel'dovich approximation (Zel'dovich 1970). A region is considered ionized if $f_{\rm coll} > \zeta^{-1}$ on *any* scale, where ζ is the "ionizing efficiency," combining a series of assumptions about the efficiency of star formation and the production (and escape) of ionizing photons from galaxies into a single parameter that corresponds to the total number of ionizing photons emitted per collapsed baryon. In standard `21cmFAST`, this criterion is assessed by filtering the density field from large to small scales, re-computing $f_{\rm coll}$ at each filter scale. Our modified algorithm assigns collapsed mass to each cell according to the non-linear density field, and it is this collapsed mass field which is filtered to determine whether a given region is ionized, similar to `DexM` (Mesinger & Furlanetto 2007) but without explicitly generating a distribution of halos. In this way, the small-scale clustering of halos is better reflected on large scales, and the new algorithm produces ionization fields that are very similar to `DexM` at a very small fraction of the computation time. We have also implemented a novel approach to treat the mean free path of ionizing photons as a smooth attenuation rather than a sharp cutoff. As shown by Greig et al. (2017a), we do not expect the exact choice of model for reionization topology to make a substantial difference in our inference of $\langle x_{\rm HI}\rangle$, so we leave an exploration of different model assumptions to future work.

The hydrodynamical simulation is likely to be too small to fully characterize the distribution of ionized regions around rare, massive halos, so we compute the ionization fields in an independent larger volume, 400 comoving Mpc (cMpc) on a side. The resolution of the cosmological initial conditions was 2048^3, while the evolved density field and ionization fields were output at a lower resolution, 512^3. We assume a mass-independent ionizing efficiency ζ, a minimum halo mass of $M_{\rm min} = 10^8$ $M_\odot$, and mean free path of ionizing photons $\lambda_{\rm mfp} = 60$ cMpc. We tuned ζ to produce ionization fields with global volume-averaged neutral fractions of $\langle x_{\rm HI}\rangle = 0.05$–$0.95$ in steps of $\Delta x_{\rm HI} = 0.05$.

Massive dark matter halos reside within larger-scale overdensities which are reionized early (Alvarez & Abel 2007), leading to an important bias in the distribution of distances to the nearest patch of neutral gas (Lidz et al. 2007; Mesinger & Furlanetto 2008). We constructed dark matter halos directly from the initial conditions following the method of Mesinger & Furlanetto (2007) as now implemented in the public release of `21cmFAST`. In Figure 1, we show a 0.78 Mpc-thick slice through an ionization field at $z = 7.5$ with $\langle x_{\rm HI}\rangle = 0.5$ and the locations of massive halos within ± 2 Mpc of the slice. As expected for the inside-out progression of reionization (Furlanetto et al. 2004), halos preferentially (in fact, exclusively, for the massive halos shown here) lie inside of ionized regions (black). As seen by the subtle grey shading, the ionization field is not entirely a binary neutral (white) vs. ionized (black) field – `21cmFAST` by default includes a prescription for a (typically small) degree of partial ionization due to "unresolved" ionized bubbles within each cell, but this is unlikely to have any impact on our results.

We then extracted randomly-oriented sightlines of $x_{\rm HI}$ from the locations of the 500 most massive halos, corresponding to $M_{\rm h} \gtrsim 3 \times 10^{11}$ $M_\odot$. We show the distribution of distances from these massive halos to the nearest patch of neutral hydrogen ($x_{\rm HI} > 0.01$) as a function of $\langle x_{\rm HI}\rangle$ in Figure 2. This distance is what determines the initial strength of the damping wing feature when the quasar first turns on, and the distribution of distances leads to large variations in the damping wing profiles between different sightlines at the same global neutral fraction. More massive halos tend to sit in larger ionized bubbles (Alvarez & Abel 2007; Lidz et al. 2007; Mesinger & Furlanetto 2008), so assuming a lower (higher) halo mass than the actual quasar host halo would result in shorter (longer) distances to the nearest neutral patch at fixed $\langle x_{\rm HI}\rangle$, leading to an overestimate (underestimate) of the damping wing strength. In this work we ignore this potential source of bias, and note that it is likely degenerate with other assumptions in our model for the reionization topology (e.g. $M_{\rm min}$).

3.3. *Ionizing radiative transfer*

We use an updated version of the one-dimensional radiative transfer implementation in (Davies et al. 2016) to compute the effect of quasar radiation on the surrounding IGM, which we briefly summarize below. The radiative transfer code computes the time-dependent evolution of six species (e^-, H I, H II, He I, He II, He III) and

[1] https://github.com/andreimesinger/21cmFAST

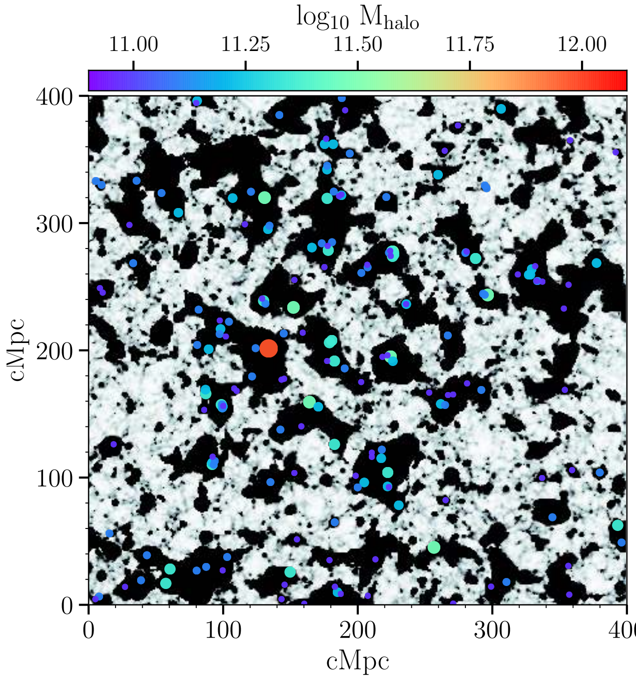

Figure 1. One pixel (~ 800 comoving kpc) slice through the semi-numerical ionization field at $z = 7.5$ with $\langle x_{\rm HI} \rangle = 0.5$ (greyscale; a linear stretch with white corresponding to neutral and black corresponding to fully ionized). The locations of $M_{\rm h} > 10^{11}$ $M_\odot$ halos are shown from a 4 comoving Mpc-thick slice centered on the ionization field slice, color- and size-coded by halo mass. Massive halos tend to lie inside of large regions that have already been ionized.

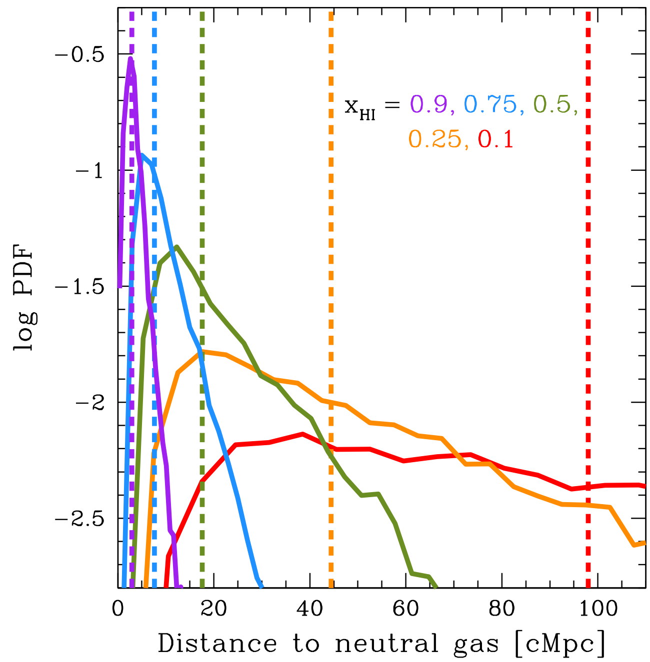

Figure 2. Distribution of distances from massive halos ($M_{\rm halo} \gtrsim 3 \times 10^{11}$ M$_\odot$) to the first patch of neutral gas in our semi-numerical simulations with global neutral fractions $\langle x_{\rm HI} \rangle = 0.1, 0.25, 0.5, 0.75, 0.9$ from right to left. The vertical dashed lines correspond to the median of the correspondingly-colored distribution.

the gas temperature, following the method described in the Appendix of Bolton & Haehnelt (2007). The abundances of ionic species are computed by integrating the following coupled system of equations,

$$\frac{dn_{\rm H\,II}}{dt} = n_{\rm H\,I}(\Gamma^{\gamma}_{\rm H\,I} + n_e\Gamma^{\rm e}_{\rm H\,I}) - n_{\rm H\,II}n_e\alpha^A_{\rm H\,II}, \quad (5)$$

$$\frac{dn_{\rm He\,II}}{dt} = n_{\rm He\,I}(\Gamma^{\gamma}_{\rm He\,I} + n_e\Gamma^{\rm e}_{\rm He\,I}) + n_{\rm He\,III}n_e\alpha^A_{\rm He\,III} - n_{\rm He\,II}(\Gamma^{\gamma}_{\rm He\,II} + n_e\Gamma^{\rm e}_{\rm He\,II} - n_e\alpha^A_{\rm He\,II}), \quad (6)$$

$$\frac{dn_{\rm He\,III}}{dt} = n_{\rm He\,II}(\Gamma^{\gamma}_{\rm He\,II} + n_e\Gamma^{\rm e}_{\rm He\,II}) - n_{\rm He\,III}n_e\alpha^A_{\rm He\,III}, \quad (7)$$

where n_i are the number densities for each species, Γ^γ_i are the photoionization rates, $\Gamma^{\rm e}_i$ are the collisional ionization rates, and α^A_i are the Case A recombination rate coefficients. In Γ^γ_i we include the effect of *secondary* ionizations, as tabulated by Furlanetto & Johnson Stoever (2010), whereby energetic photoelectrons (with kinetic energy greater than the ionization potential) lose energy by ionizing additional atoms rather than simply dumping the excess photoionization energy into the gas as heat. The remaining species are then solved for via the closure conditions

$$n_{\rm H\,I} = n_{\rm H} - n_{\rm H\,II}, \quad (8)$$

$$n_{\rm He\,I} = \frac{Y}{4(1-Y)}n_{\rm H} - n_{\rm He\,II} - n_{\rm He\,III}, \quad (9)$$

$$n_e = n_{\rm H\,II} + n_{\rm He\,II} + 2n_{\rm He\,III}, \quad (10)$$

where we have assumed $Y = 0.24$ for the mass fraction of helium.

The gas temperature is evolved taking into account photoionization heating and cooling[2] from recombinations, collisional excitation, the expansion of the Universe, and inverse Compton scattering off of CMB photons (see Davies et al. 2016 for more details). Adding to the model presented in Davies et al. (2016), we now include the prescription from Rahmati et al. (2013) to approximate self-shielding of the ionizing background in dense gas, and update this self-shielding at each time step to take into account the ionization of dense absorbers by the quasar.

For the ionizing spectrum of each quasar, we first use the Lusso et al. (2015) template to convert from the

[2] We assume that the gas is of primordial composition, i.e. there is no cooling due to elements heavier than helium.

measured M_{1450} to L_ν at the ionizing edge of hydrogen ($E_{\rm HI} \approx 13.6$ eV), and then extrapolate to higher frequencies by assuming $L_\nu \propto \nu^{-1.7}$, in agreement with the best-fit power-law spectrum from Lusso et al. (2015). Assuming a different average quasar template (e.g. Telfer et al. 2002; Stevans et al. 2014) could change the output of ionizing photons by tens of percent, which would have implications for the shape of the proximity zone transmission profile (Davies et al., in prep.) and strength of the damping wing. However, the size of the ionized bubble around the quasar $R_{\rm ion}$ (assuming a fully neutral universe) is only weakly dependent on the ionizing photon output $\dot{N}_{\rm ion}$ ($R_{\rm ion} \propto \dot{N}_{\rm ion}^{1/3}$; Cen & Haiman 2000), and this dependence is completely degenerate with the lifetime of the quasar. When computing the photoionization and photoheating rates, we integrate the ionizing spectrum over frequencies from the ionizing frequency ν_i to $40\nu_i$ for each (partially-)neutral species i, separately, with 25 logarithmic frequency bins.

We assume a "lightbulb" model for quasar emission: the quasar turned on at some point $t_{\rm q}$ in the past, and has been shining at a constant luminosity since then. In the rest of the paper we will refer to $t_{\rm q}$ as the "quasar lifetime."

3.4. *Hybrid model*

We synthesize these three model components by computing ionizing radiative transfer along hydrodynamical simulation skewers, with the initial neutral fraction along each sightline set by skewers from the semi-numerical reionization simulations. For regions with $x_{\rm HI} = 0$ in the reionization simulations, we assume a uniform ionizing background such that $\langle x_{\rm HI} \rangle$ inside of ionized regions is $\sim 10^{-3}$ (corresponding to a hydrogen photoionization rate $\Gamma_{\rm HI} \sim 6 \times 10^{-14}$), although we find that our results are insensitive to this choice. We initialize the IGM temperature in ionized regions to the values from the hydrodynamical simulation skewer, and assume that the IGM is initially cold (2000 K) inside of neutral regions (e.g. Furlanetto et al. 2006). The density field from the hydrodynamical simulation does not reflect the additional clumpiness of such cold gas, i.e. the gas in the simulation has been "pressure smoothed" to some extent (Gnedin & Hui 1998; Rorai et al. 2013; Kulkarni et al. 2015), but we do not expect this to have a large effect on the transmission profile.

In the left column of Figure 3, we show a typical example of the output from our full model, a simulated sightline assuming the luminosity and redshift of J1342+0928 ($M_{1450} = -26.76, z_{\rm q} = 7.5413$) using a skewer from the hydrodynamical simulation with initial $x_{\rm HI}$ given by a skewer from the $\langle x_{\rm HI} \rangle = 0.5$ semi-numerical reionization simulation. The top panel shows the Lyα transmission in the quasar spectrum, the middle panel shows the neutral fraction, and the bottom panel shows the gas temperature. The different colors correspond to quasar lifetimes of $10^{4.5}$ (orange), $10^{6.5}$ (green), and 10^8 (purple) years. The damping wing signal, shown by the absorption at negative distance ($\lambda_{\rm rest} > \lambda_{{\rm Ly}\alpha}$), is very strong soon after the quasar turns on (orange), with the first patch of neutral gas encountered at ~ 1.3 proper Mpc along the line of sight. This first neutral patch is ionized within a few million years (green), weakening the damping wing considerably and photoheating the gas to $T \sim 3$–4×10^4 K. After 100 million years (purple), the quasar has carved out a large enough ionized region (~ 7 proper Mpc) to completely wipe out the damping wing signal, and initially-neutral photoheated regions near the quasar have cooled substantially. At this stage, the proximity zone is no longer cut off by the onset of fully neutral gas in the IGM; instead, the Lyα forest absorption becomes too strong as the ionizing flux from the quasar decreases (e.g. Bolton & Haehnelt 2007; Eilers et al. 2017) and $x_{\rm HI}$ reaches $\sim 10^{-4}$, as shown by the complete disappearance of the purple transmission curve at a shorter distance ($\lesssim 6$ proper Mpc) than the location of the ionization front (6.7 proper Mpc).

Cosmic variance in the density field and in the reionization morphology lead to a wide variety of proximity zone and damping wing spectra at the same neutral fraction – we show another example sightline in the right column of Figure 3 which initially resides in a very large ionized region, and thus *never* shows a strong damping wing signal. For the longest lifetime model (purple), a modest amount of extra transmission appears at $R \sim 5$ proper Mpc due to heating from the reionization of He II by the quasar (the thermal proximity effect, e.g. Bolton et al. 2012; Khrykin et al. 2017).

For each semi-numerical reionization box, corresponding to $0 \leq \langle x_{\rm HI} \rangle \leq 1$ in 21 steps of $\Delta x_{\rm HI} = 0.05$, we ran 2400 radiative transfer simulations on 2400 different random skewers, using each of our 1200 hydrodynamical simulation skewers twice. From these simulations, we computed transmission spectra every $\Delta \log{[t_{\rm q}/yr]} = 0.5$ in 11 steps from 10^3 to 10^8 years covering a velocity range $-10,000 \leq v - v_{\rm sys} \leq +10,000$ km/s, where $v_{\rm sys}$ is the systemic velocity of the halo center. Our final set of transmission spectrum models is then 21×11, with 2400 spectra for each point in the coarse 2D grid. In Figure 4 we show the mean transmission profiles from our simulations as a function of $t_{\rm q}$ at fixed $\langle x_{\rm HI} \rangle = 0.5$ (top), and as a function of $\langle x_{\rm HI} \rangle$ at fixed $t_{\rm q} = 10^6$ years (bottom). There is a clear trend towards stronger damp-

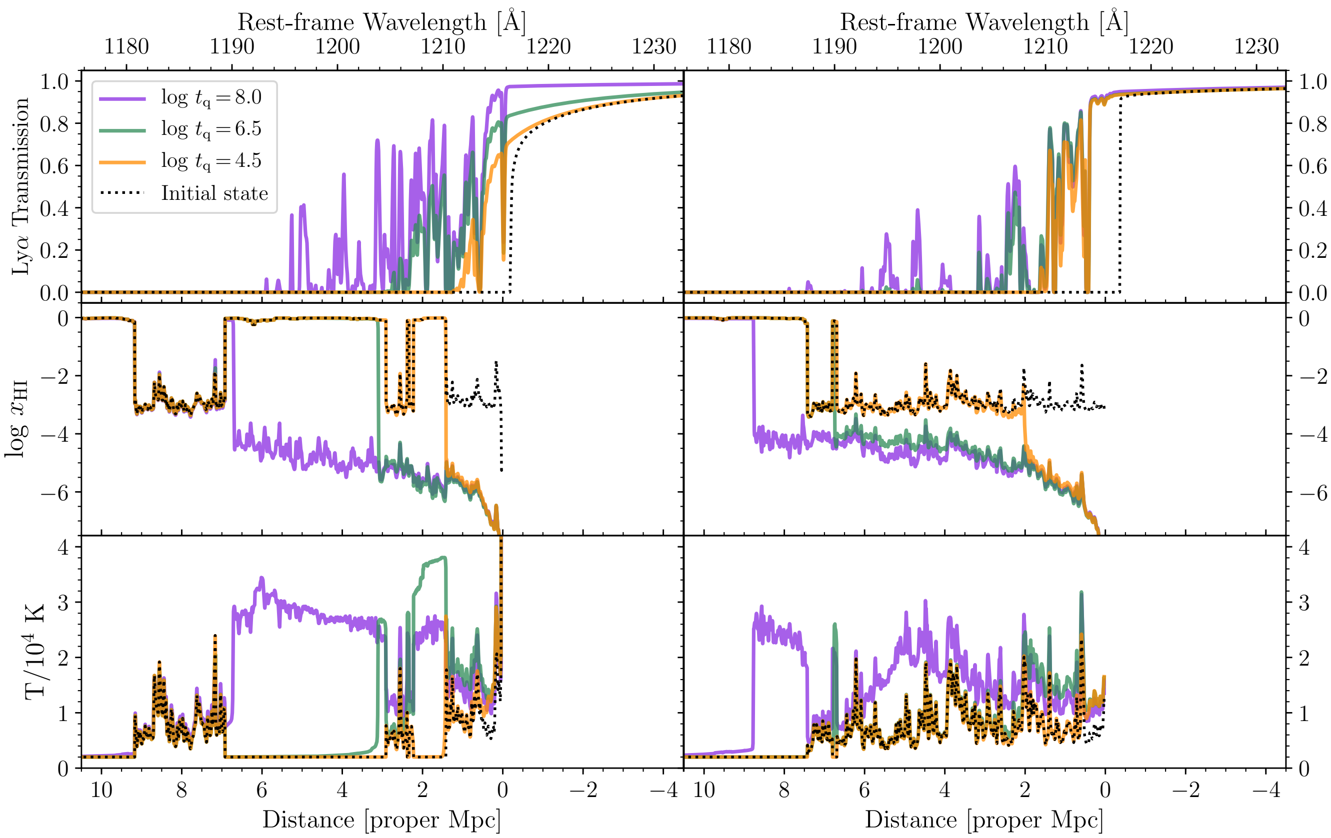

Figure 3. Example outputs from the hybrid model of quasar proximity zones at $\log t_q = 4.5$ (orange), 6.5 (green), and 8.0 (purple) for two skewers (left and right) through the $\langle x_{HI} \rangle = 0.5$ simulation at $z = 7.54$. The black dotted curves show the initial state of the skewer prior to the quasar turning on. The top panels show the Lyα transmission, the middle panels show x_{HI}, and the bottom panels show the gas temperature.

ing wings and smaller proximity zones for high neutral fractions and short quasar lifetimes.

As noted by Bolton et al. (2011), a degeneracy exists between IGM neutral fraction and quasar lifetime in determining the shape of the proximity zone and damping wing profile, wherein short quasar lifetime and small neutral fraction appears similar to long quasar lifetime and large neutral fraction (although these models assumed a constant x_{HI} in the IGM instead of our more realistic patchy topology). A similar degeneracy arises in our hybrid model because even at large neutral fraction, the quasar can carve out a large ionized region that greatly increases the distance to the nearest neutral patch, decreasing the strength of the damping wing feature and increasing the size of the proximity zone (see the purple curves in Figure 3). Another consequence of this is that at relatively long quasar lifetimes, $t_q \gtrsim 10^8$ years, the damping wing almost entirely disappears, even for $\langle x_{HI} \rangle \sim 1$.

At even shorter timescales, $t_q \lesssim 10^5$ years, the inner parts of the proximity zone start to disappear entirely (see the orange curves in Figure 3). This occurs because the gas has not been illuminated long enough to respond to the increased ionizing flux from the quasar (e.g. Khrykin et al. 2016), and such short lifetimes may explain the handful of very small proximity zones observed at $z \sim 6$ (Eilers et al. 2017, Davies et al. in prep.).

4. STATISTICAL METHOD FOR JOINTLY INFERRING THE NEUTRAL FRACTION AND QUASAR LIFETIME

The measured quasar proximity zone and damping wing signals are a highly covariant heteroskedastic process, with large sightline-to-sightline variance for any particular set of parameters $(\langle x_{HI} \rangle, t_q)$. In addition to the uncorrelated photon noise in the spectrum, additional sources of variance are the IGM density field, which leads to the absorption inside the proximity zone, and the distance to the nearest neutral patch of the IGM, which has strong covariant effects across the whole spectrum. The uncertainty in our prediction for the quasar continuum (§ 2) introduces an additional multiplicative error which is strongly covariant. The combi-

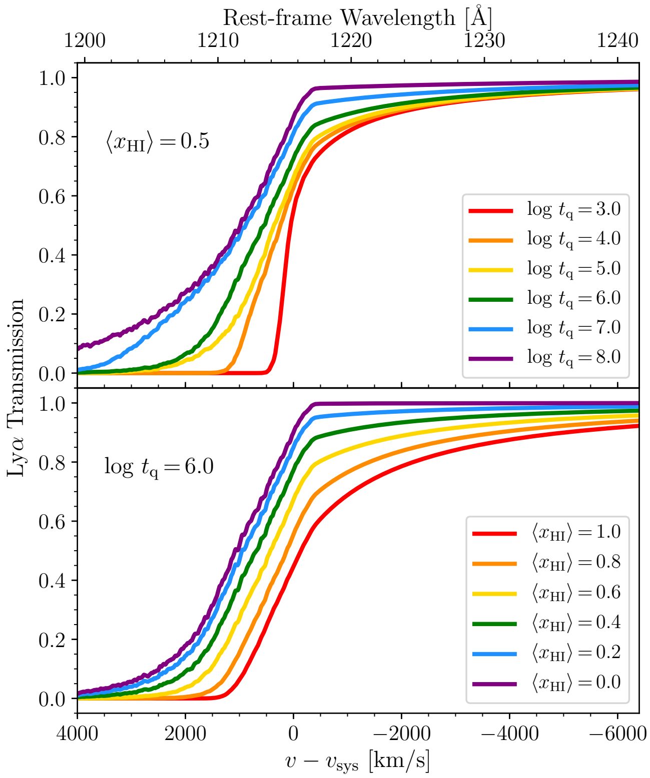

Figure 4. Mean Lyα transmission profiles from the 1D radiative transfer simulations. The top panel shows the variation in the mean profile for varying quasar lifetime ($\log[t_q/\text{yr}]$=3.0–8.0, $\Delta\log t_q$=1.0) at a fixed global neutral fraction $\langle x_{\rm HI}\rangle$=0.5. The bottom panel shows the variation in the mean profile for varying neutral fraction ($\langle x_{\rm HI}\rangle$=0.0–1.0, $\Delta x_{\rm HI}$=0.2) at a fixed quasar lifetime of $\log[t_q/\text{yr}] = 6.0$.

nation of these processes cannot be simply described by by a multivariate Gaussian likelihood, suggesting that inference via standard likelihood-based methods (e.g. Markov Chain Monte Carlo) may be difficult to interpret correctly.

Instead, we adopt an approach following principles of Bayesian indirect inference (Gourieroux et al. 1993; Drovandi et al. 2015), wherein the likelihood for auxiliary parameters or an auxiliary likelihood of the true parameters is used in place of an intractable true likelihood for the true parameters. We define a "pseudo-likelihood" $\tilde{L}$ as the product of flux probability distribution functions (PDFs) $P(F_i)$ of 500 km/s binned pixels,

$$\tilde{L}(\theta) = \prod_i P(F_i|\theta), \tag{11}$$

which is equivalent to the likelihood function of the (500 km/s binned) transmission spectrum in the absence of correlations between pixels. For computational simplicity, and to limit the impact of our finite number of simulated sightlines, we approximate the flux PDFs $P(F_i|\theta)$ of each bin i by fitting them with mixtures of three Gaussians[3]. While direct parameter inference from this likelihood would be formally incorrect due to the neglected correlations, one can still determine a set of maximum pseudo-likelihood model parameters, $\theta_{\rm M\tilde{L}E}$, which should be closely related to the true maximum likelihood parameters. This procedure reduces the dimensionality of our data from the number of transmitted flux bins in the spectrum to the number of model parameters, allowing for a full Bayesian treatment with modest computational expense (albeit likely with slightly less constraining power than the original data due to information lost in this compression).

We treat $\theta_{\rm M\tilde{L}E}$ as a summary statistic and compute the posterior PDF of the "true" model parameters θ following Bayes' theorem,

$$p(\theta|\theta_{\rm M\tilde{L}E}) = \frac{p(\theta_{\rm M\tilde{L}E}|\theta)p(\theta)}{p(\theta_{\rm M\tilde{L}E})}, \tag{12}$$

where $p(\theta|\theta_{\rm M\tilde{L}E})$ is the posterior PDF of θ, $p(\theta_{\rm M\tilde{L}E}|\theta)$ is the likelihood of $\theta_{\rm M\tilde{L}E}$ given the model θ, $p(\theta)$ is the prior on θ, and $p(\theta_{\rm M\tilde{L}E})$ is the evidence. We compute the likelihood function *directly* by measuring the distribution of $\theta_{\rm M\tilde{L}E}$ for forward-modeled mock data on a coarse grid of $\theta = (\langle x_{\rm HI}\rangle, t_q)$ and explicitly computing the evidence,

$$p(\theta_{\rm M\tilde{L}E}) = \int p(\theta_{\rm M\tilde{L}E}|\theta)p(\theta)d\theta. \tag{13}$$

We denote as $\hat{F}$ a forward-modeled transmission spectrum, which results from taking a draw F from the set of 2400 transmission spectra computed for the parameter set θ, multiplication by a random draw from a multivariate Gaussian distribution describing the relative continuum error (ϵ_C; Paper I), then a draw of additive Gaussian noise following the continuum-normalized noise vector of the spectrum (N):

$$\hat{F} = F \times (1 + \epsilon_C) + N. \tag{14}$$

We find the $\theta_{\rm M\tilde{L}E}$ for each mock spectrum via a simple brute force approach, computing $\tilde{L}$ (equation 11) for each of the 21×11 models in our coarse grid. For priors, we assume a flat *linear* prior on $\langle x_{\rm HI}\rangle$ from 0 to 1, and flat *log* priors on t_q from 10^3 to 10^8 years. The linear prior on $\langle x_{\rm HI}\rangle$ reflects our expectation that $z > 7$ is in the midst of the reionization epoch, while the log prior

[3] The exact form of the approximation to the individual flux PDFs appears to have only a minor effect on our analysis – similar, albeit somewhat less constraining, posterior PDFs can be obtained with single Gaussian fits.

on t_q reflects our broad uncertainty on the lifetime of the luminous quasar phase, incorporating recently discovered evidence for very short lifetimes (Eilers et al. 2017). We will also quote constraints for a stronger prior on t_q which only allows lifetimes as short as 10^5 years. We will discuss the impact of this choice of priors in § 6.3.

In this method, the data have been compressed into the measurement of $\theta_{\tilde{\mathrm{MLE}}}$, so any spectrum with a particular $\theta_{\tilde{\mathrm{MLE}}}$ will result in an identical 2D posterior PDF. Given our coarse grid in parameter space, there are only $21 \times 11 = 231$ possible posterior PDFs for any given quasar spectrum.

5. RESULTS

In Paper I, we estimated the intrinsic continua (and their uncertainties) of the two quasars known at $z > 7$: ULAS J1120+0641 ($z = 7.0851$; Mortlock et al. 2011; Venemans et al. 2017b) and ULAS J1342+0928 ($z = 7.5413$; Bañados et al. 2018; Venemans et al. 2017a). Here we apply the statistical method from the previous section to jointly constrain the neutral fraction at $z > 7$ and the quasar lifetimes by comparing the resulting transmission profiles to our simulated spectra.

For the purposes of modeling the physical state of the IGM along the line of sight, we adopt the precise systemic redshifts above as the true locations of the quasar host halos. However, these systemic redshifts have little relevance to the PCA continua, given that the PCA model was trained on quasars with imprecise redshifts. In Paper I we resolved this ambiguity by fitting for a "template redshift" simultaneously with the red-side PCA coefficients, resulting in an independent (but physically irrelevant) redshift estimate that can be applied to any quasar spectrum. This template redshift is what then defines the rest-frame wavelengths for the continuum prediction.

5.1. *ULAS J1120+0641*

In Figure 5, we show the red-side fit to the VLT/FORS2 + Gemini/GNIRS spectrum of ULAS J1120+0641 (Mortlock et al. 2011; top panel) and the predicted blue-side continuum (bottom panel) from Paper I. The predicted continuum has been corrected for the mean bias of predicted continua for similar quasars in the training set, as discussed in §2 and shown in Figure 12 of Paper I. We find a best-fit template redshift of $z = 7.0834$, a very small blueshift of $\Delta v = 63$ km/s from the systemic frame defined by the centroid of the [CII] emission line of the host galaxy ($z = 7.0851$, Venemans et al. 2017b). The blue-side profile shows a hint of absorption redward of Lyα and a relatively small proximity zone ($R_\mathrm{p} = 1.72$ proper Mpc, following the definition in Eilers et al. 2017) compared to the trend seen at $z \sim 5.7$–6.5 (Eilers et al. 2017; Mazzucchelli et al. 2017). While the damping wing signal appears to be fairly strong at rest-frame Lyα, similar to what was found by previous works (Mortlock et al. 2011; Simcoe et al. 2012; Greig et al. 2017a), when our large, covariant continuum uncertainty is taken into account the spectrum does not appear to definitively indicate a neutral IGM.

We show the resulting transmission spectrum (i.e. observed spectrum divided by the continuum model) as the grey curve in Figure 6 and the 500 km/s-binned spectrum in black. For the statistical analysis, we only use pixels at $v - v_\mathrm{sys} > -4,400$ km/s ($\lambda_\mathrm{rest} \lesssim 1233$ Å) to avoid the strong (and unresolved) associated N V absorption, and we choose to end the blue-side coverage at $v - v_\mathrm{sys} = +6,400$ km/s ($\lambda_\mathrm{rest} \sim 1190$ Å) because all of the proximity zone models have no detectable signal beyond that distance. We find maximum pseudo-likelihood parameter values of $\theta_{\tilde{\mathrm{MLE}}} = (\langle x_\mathrm{HI} \rangle = 0.65, \log t_\mathrm{q} = 6.5)$, and we show the median transmission profile of the $\tilde{\mathrm{MLE}}$ model (blue solid) and the expected 16–84th percentile scatter (blue shaded) from forward modeled spectra with $\theta = \theta_{\tilde{\mathrm{MLE}}}$ in Figure 6. From the $\tilde{\mathrm{MLE}}$ parameter values we infer the 2D posterior PDF $p(\theta|\theta_{\tilde{\mathrm{MLE}}})$ shown in Figure 7. The posterior PDF is relatively flat across a wide swathe of $(\langle x_\mathrm{HI} \rangle, \log t_\mathrm{q})$ parameter space, with a trend towards higher $\langle x_\mathrm{HI} \rangle$ for longer t_q, reflectingv the degeneracy between these two parameters discussed in § 3.4 and shown in Figure 4. The non-zero size of the proximity zone rules out quasar lifetimes shorter than $\sim 10^{4.5}$ years, while the combination of damping wing strength and small proximity zone rule out quasar lifetimes longer than $\sim 10^7$ years.

5.2. *ULAS J1342+0928*

In Figure 8, we show the red-side fit to the Magellan/FIRE + Gemini/GNIRS spectrum of ULAS J1342+0928 and the predicted blue-side (bias-corrected) continuum from Paper I. We find a best-fit template redshift of $z = 7.4438$, a blueshift of $\Delta v = 3422$ km/s from the systemic frame ($z = 7.5413$, Venemans et al. 2017a). The red-side spectrum is very different from a typical quasar, however, similar examples do exist in our PCA training set (Paper I), and the PCA model is capable of broadly reproducing the spectrum. In fact, the uncertainty in the continuum derived from nearest-neighbor quasars in the training set is somewhat lower than for typical quasars due to the relatively weak broad emission lines. The blue-side profile shows a strong damping wing redward of Lyα, and a very small proximity zone

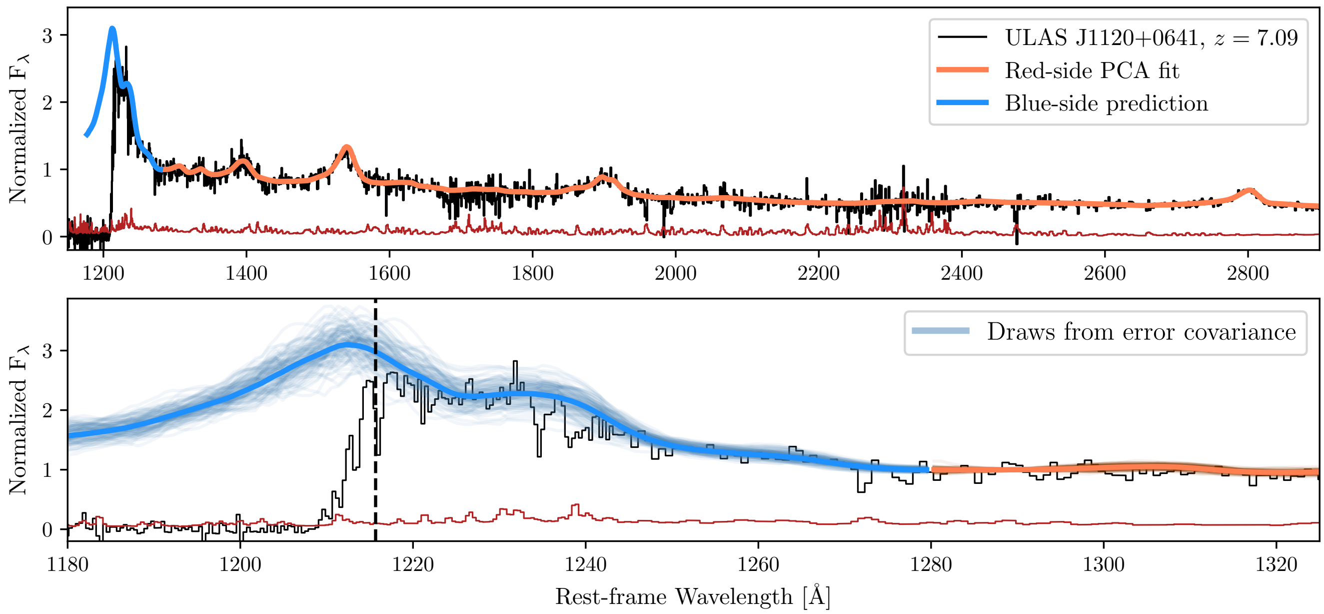

Figure 5. Top: VLT/FORS2 + Gemini/GNIRS spectrum of ULAS J1120+0641 (Mortlock et al. 2011, black) and its noise vector (red). The red-side PCA fit and blue-side prediction are shown as the orange and blue curves, respectively. Bottom: Zoom in of the Lyα region of the spectrum, where the vertical dashed line shows rest-frame Lyα ($\lambda_{\rm rest} = 1215.67$ Å). The transparent curves show 100 draws from the covariant blue-side prediction error calibrated from the 1% most similar quasars in the training set. This quasar shows modest evidence for a damping wing and has a relatively small proximity zone.

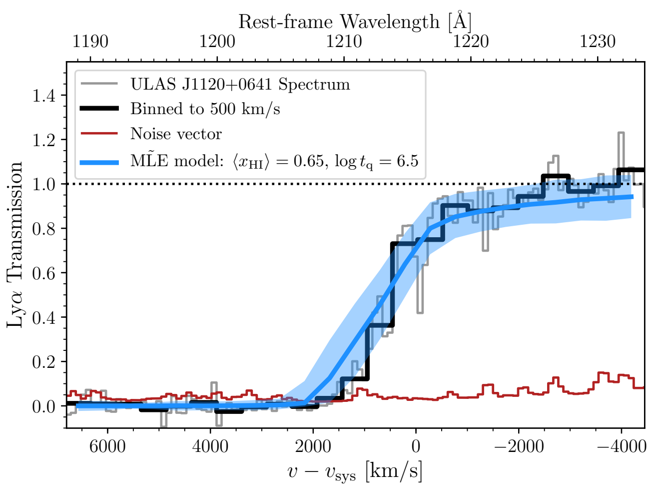

Figure 6. Continuum-divided spectrum of ULAS J1120+0641 (grey) and its noise vector (red). The black histogram shows the spectrum rebinned to ~ 500 km/s in the region we use for model comparison. The blue solid curve shows the median ~ 500 km/s-binned transmission spectrum of mock spectra with the ML̃E parameter values $\theta_{\rm M\tilde{L}E} = (\langle x_{\rm HI} \rangle = 0.65, \log t_{\rm q} = 6.5)$, while the associated blue shaded region shows the 16th–84th percentile range for mock spectra with $\theta = \theta_{\rm M\tilde{L}E}$

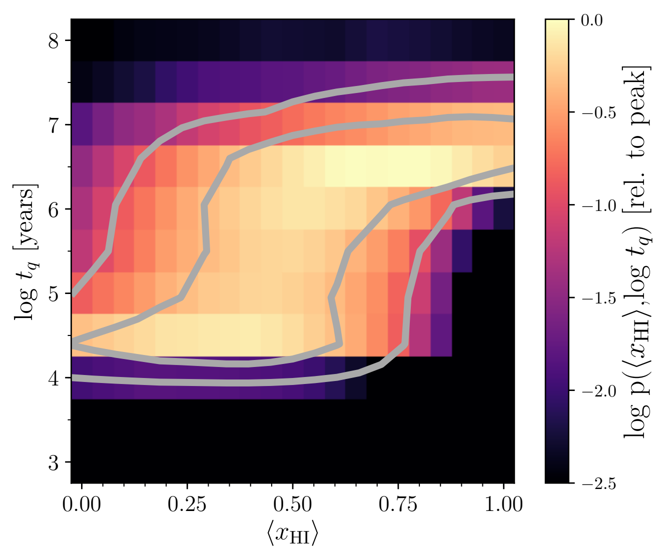

Figure 7. 2D posterior PDF of $\langle x_{\rm HI} \rangle$ and $\log t_{\rm q}$ resulting from the ML̃E parameter values $\theta_{\rm M\tilde{L}E} = (\langle x_{\rm HI} \rangle = 0.65, \log t_{\rm q} = 6.5)$ derived from the ULAS J1120+0641 spectrum. The contours enclose 68% and 95% of the total probability.

($R_p = 1.20$ pMpc). The damping wing signal is clearly stronger and the proximity zone is even smaller than ULAS J1120+0641, despite the slightly higher luminosity of ULAS J1342+0928. Both of these properties point towards a substantially neutral IGM surrounding ULAS J1342+0928.

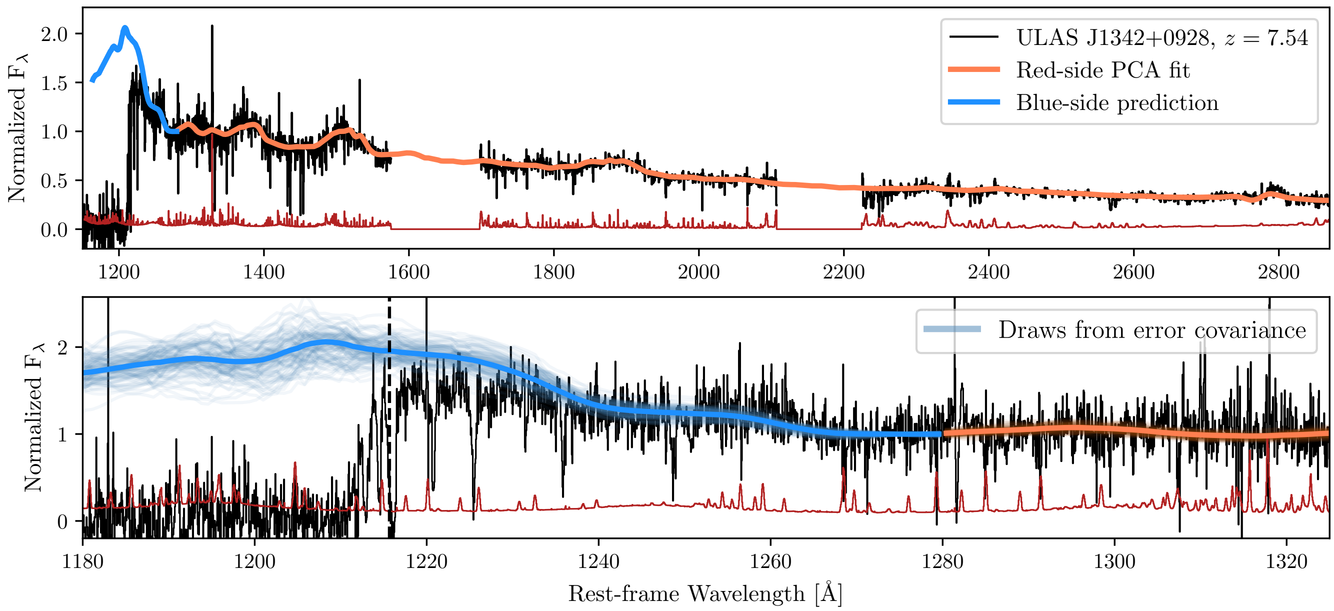

Figure 8. Similar to Figure 5 but for the Magellan/FIRE + Gemini/GNIRS spectrum of ULAS J1342+0928 (black), including its noise vector (red), red-side PCA fit (orange), and blue-side prediction (blue). The FIRE spectrum in the top panel has been re-binned to match the resolution of the GNIRS data used in the K-band, while the bottom panel is shown at the higher FIRE resolution. This quasar shows strong evidence for a damping wing and has a very small proximity zone.

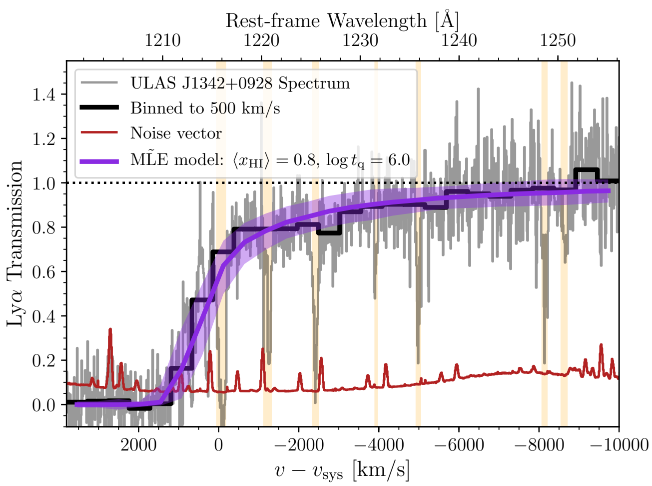

Figure 9. Similar to Figure 6 but for the continuum-divided spectrum of ULAS J1342+0928. The purple solid curve shows the median binned transmission spectrum in the mock spectra assuming the $\tilde{\rm MLE}$ parameter values $\theta_{\rm M\tilde{L}E} = (\langle x_{\rm HI}\rangle = 0.8, \log t_{\rm q} = 6.0)$, while the associated purple shaded region shows the 16th–84th percentile range for mock spectra with $\theta = \theta_{\rm M\tilde{L}E}$. The orange shaded regions highlight identified metal absorption systems that we have masked in our analysis.

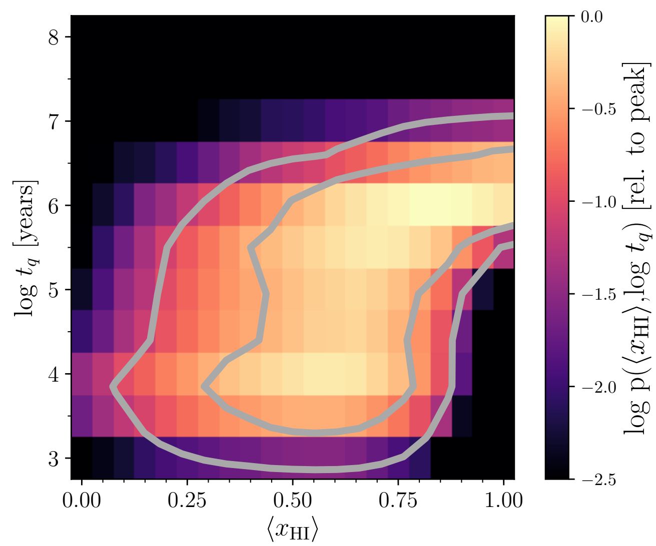

Figure 10. 2D posterior PDF of $\langle x_{\rm HI}\rangle$ and $\log t_{\rm q}$ resulting from the $\tilde{\rm MLE}$ parameter values $\theta_{\rm M\tilde{L}E} = (\langle x_{\rm HI}\rangle = 0.65, \log t_{\rm q} = 6.5)$ derived from the ULAS J1342+0928 spectrum. The contours enclose 68% and 95% of the total probability.

We show the resulting transmission spectrum as the grey curve in Figure 9 and the 500 km/s-binned spectrum in black. No strong associated absorption is visible in the spectrum, so we include all pixels redward of Lyα that are covered by our modeled transmission spectra ($v - v_{\rm sys} < 10,000$ km/s, $\lambda_{\rm rest} \lesssim 1255$ Å). We cut off the blue-side coverage at $+4000$ km/s (~ 1200 Å) due to the presence of non-Gaussian noise features (both positive and negative spikes) in the spectrum that are not included in our noise model, which could spuriously impact the statistical analysis[4]. The maximum pseudo-likelihood parameter values are $\theta_{\tilde{\mathrm{MLE}}} = (\langle x_{\mathrm{HI}} \rangle = 0.8, \log t_{\mathrm{q}} = 6.0)$, and we compare the binned transmission spectrum to the median transmission profile of the M$\tilde{\mathrm{L}}$E model (solid purple) and the expected 16–84th percentile scatter (dashed purple) from our forward modeling in Figure 9. From the M$\tilde{\mathrm{L}}$E parameter values we infer the 2D posterior PDF $p(\theta|\theta_{\tilde{\mathrm{MLE}}})$ shown in Figure 10. Due in large part to the strong damping wing, the posterior PDF has a clear preference for a significantly neutral universe, even for short quasar lifetimes, although there is still some degeneracy between $\langle x_{\mathrm{HI}} \rangle$ and t_{q}.

6. DISCUSSION

In the preceding sections, we have demonstrated a method for jointly constraining the global neutral fraction $\langle x_{\mathrm{HI}} \rangle$ and quasar lifetime t_{q} from analysis of the proximity zone and (presence or absence) damping wing, and applied it to the two quasars known at $z > 7$: ULAS J1120+0641 at $z = 7.09$, and ULAS J1342+0928 at $z = 7.54$. Here we present our constraints on $\langle x_{\mathrm{HI}} \rangle$ marginalized over quasar lifetime, compare our analysis of ULAS J1120+0641 to those from previous works, and discuss our choice of $\langle x_{\mathrm{HI}} \rangle$ and t_{q} priors.

6.1. *The History of Reionization: $\langle x_{\mathrm{HI}} \rangle(z)$*

In Figure 11, we show the posterior PDFs for $\langle x_{\mathrm{HI}} \rangle$ from each quasar, marginalized over quasar lifetime with a flat prior in log space from 10^3 to 10^8 years (solid curves) and a more restrictive prior from 10^5 to 10^8 years (dashed curves), and provide the 68% and 95% credible intervals in Table 1. We chose to include extremely short quasar lifetimes $\lesssim 10^5$ years in our fiducial analysis due to the existence of a surprisingly large fraction of small proximity zones at $z \sim 6$ ($\sim 10\%$, Eilers et al. 2017) which imply lifetimes shorter than 10^5 years. Including the possibility of the shortest quasar lifetimes shifts the $\langle x_{\mathrm{HI}} \rangle$ posterior PDF to slightly lower values, consistent with the degeneracy between $\langle x_{\mathrm{HI}} \rangle$ and t_{q} shown in Figure 4, but in general does not have a large effect. While our analysis of ULAS J1120+0641 suggests a neutral fraction of $\langle x_{\mathrm{HI}} \rangle \sim 0.5$, the posterior PDF is not particularly constraining, with significant probability density at the $\langle x_{\mathrm{HI}} \rangle = 0$ and $\langle x_{\mathrm{HI}} \rangle = 1$ boundaries. In contrast, the posterior PDF for ULAS J1342+0928 strongly indicates a significantly neutral IGM, ruling out $\langle x_{\mathrm{HI}} \rangle < 0.08$ (0.14) at 99% probability marginalized over quasar lifetime for our prior covering $10^3 \leq t_{\mathrm{q}} \leq 10^8$ ($10^5 \leq t_{\mathrm{q}} \leq 10^8$) years.

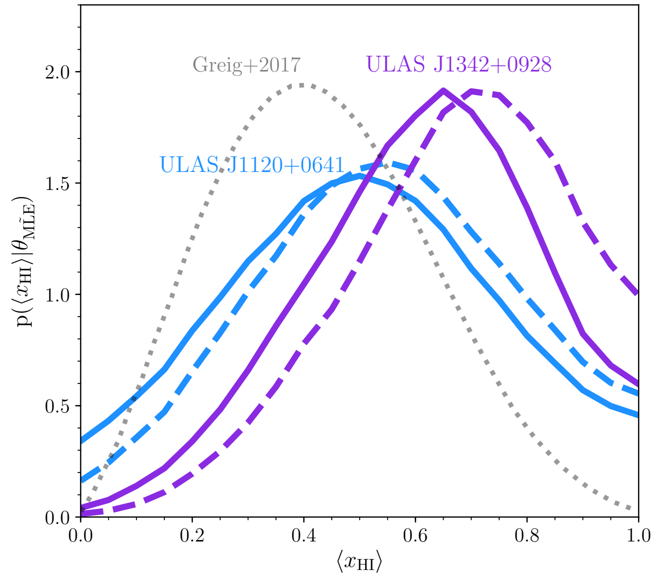

Figure 11. Posterior PDFs of $\langle x_{\mathrm{HI}} \rangle$ for ULAS J1120+0641 (blue) and ULAS J1342+0928 (purple) marginalized over quasar lifetime assuming a flat prior covering our entire model grid ($3.0 \leq \log t_{\mathrm{q}}/\mathrm{yr} \leq 8.0$; solid curves) or adopting a prior that excludes extremely short lifetimes ($5.0 \leq \log t_{\mathrm{q}}/\mathrm{yr} \leq 8.0$; dashed curves). The posterior PDF from the damping wing analysis of ULAS J1120+0641 in Greig et al. (2017a) is shown as the dotted grey curve.

In Figure 12, we compare the posterior PDFs for $\langle x_{\mathrm{HI}} \rangle$ from each quasar to the broad swathe of reionization histories consistent with the measured electron-scattering optical depth of the CMB (Planck Collaboration et al. 2016b)[5]. Under our most conservative prior, allowing quasar lifetimes from 10^3 to 10^8 years, we find 68% (95%) credible intervals of $\langle x_{\mathrm{HI}} \rangle(z = 7.09) = 0.48^{+0.26}_{-0.26}(^{+0.47}_{-0.46})$ and $\langle x_{\mathrm{HI}} \rangle(z = 7.54) = 0.60^{+0.20}_{-0.23}(^{+0.36}_{-0.45})$. These constraints are consistent with the CMB and are in broad agreement with recent calculations of the reionization history (e.g. Robertson et al. 2015; Bouwens et al. 2015; Khaire et al. 2016). The large cosmic variance between the damping wing profiles at fixed $\langle x_{\mathrm{HI}} \rangle$ (§ 3.4), the strong degeneracy with quasar lifetime, and the limited precision of our continuum reconstructions greatly limits the constraining power of any single $z > 7$

[4] This non-Gaussianity is likely also affecting the region of spectrum that we do analyze; however, we expect that even a modestly enhanced noise in the spectrum will still be subdominant compared to the other sources of variance that we consider (cosmic variance, continuum error).

[5] We compare to the combined Planck + ACT/SPT constraints on the reionization history that take into account upper limits on the strength of the kinetic Sunyaev-Zel'dovich effect and a prior from Lyα forest measurements that the end of reionization occurred before $z = 6$, see § 5.3 of Planck Collaboration et al. (2016b)

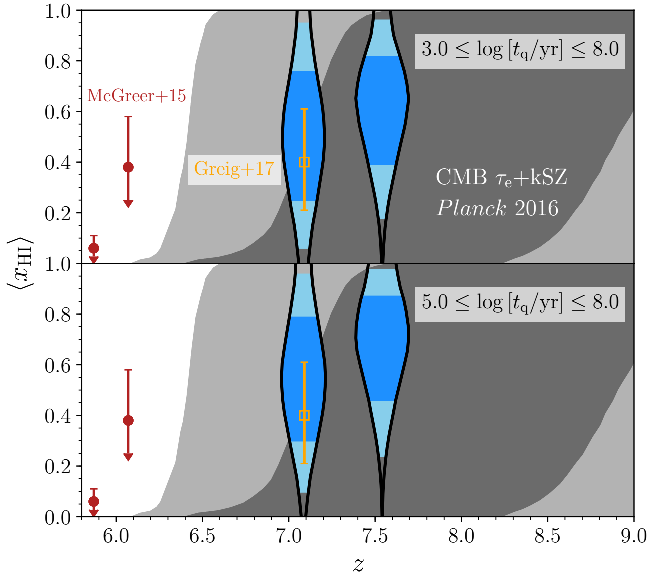

Figure 12. Violin plot comparing the posterior PDFs from our analysis with the reionization history constraints from Planck Collaboration et al. (2016b), with the dark and light grey shaded regions corresponding to the 68% and 95% credible intervals, respectively. Also shown are the Lyα+Lyβ forest dark pixel constraints from McGreer et al. (2015) (red crosses) and the damping wing analysis of ULAS J1120+0641 from Greig et al. (2017a) (orange square).

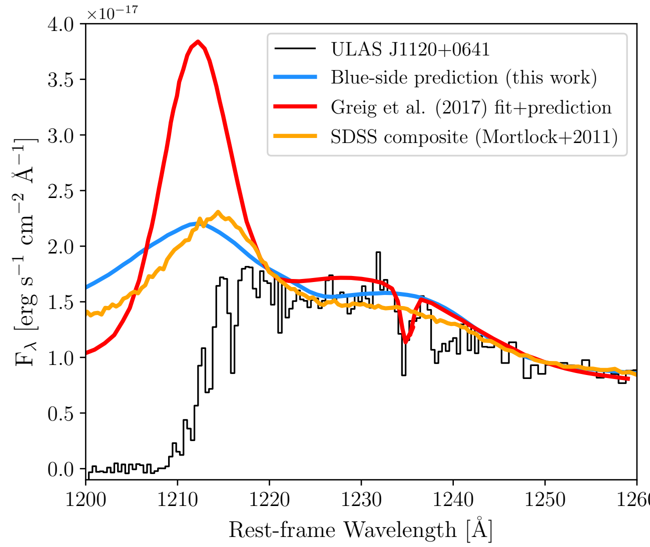

Figure 13. Comparison between our blue-side prediction for ULAS J1120+0641 (blue), the SDSS matched composite spectrum from Mortlock et al. (2011) (orange), and the model from Greig et al. (2017a). The Greig et al. (2017a) model has been renormalized to our blue-side prediction at $\lambda_{\rm rest} = 1245$ Å to correct for the different flux calibration of the Simcoe et al. (2012) spectrum used in their analysis.

quasar. However, with only a handful of additional quasars at $z > 7$, it may be possible to constrain $\langle x_{\rm HI} \rangle(z)$ to $\sim 10\%$. That said, despite the substantial uncertainties, our analysis of two $z > 7$ quasars already constrains the reionization history more than the integral constraint from the CMB.

6.2. *Previous studies of ULAS J1120+0641*

In the original discovery paper for ULAS J1120+0641, Mortlock et al. (2011) suggested that the spectrum showed signs of an IGM damping wing. They selected a sample of lower-redshift quasars with similar C IV blueshifts (relative to Mg II) and equivalent widths, and stacked their spectra to predict the shape of ULAS J1120+0641. The resulting composite spectrum was somewhat above the observed spectrum at wavelengths at and just redward of rest-frame Lyα, with a shape resembling the characteristic damping wing profile. However, the uncertainty in the stacked composite was not fully quantified, and the physical model was limited to the Miralda-Escudé (1998) expression for the damping wing. A followup work by Bolton et al. (2011) expanded upon the physical model with 1D radiative transfer simulations, and found that the combination of absorption at rest-frame Lyα and the small proximity zone were suggestive of neutral gas close to the quasar ($x_{\rm HI} \sim 0.1$), although they noted that an identical signal could potentially come from small-scale optically thick gas along the line of sight instead of a neutral IGM (see also Keating et al. 2015). With a higher resolution FIRE spectrum, Simcoe et al. (2012) found that any such dense gas would have to be extremely metal-poor ($[Z/H] < -4$), which would seem to favor the IGM interpretation.

The accuracy of the Mortlock et al. (2011) composite spectrum as a prediction for ULAS J1120+0641 was called into question by Bosman & Becker (2015), because the composite spectrum fails to match the C IV line and this may lead to an overestimate of the Lyα emission. Bosman & Becker (2015) selected a comparison sample of low-redshift quasars with more precisely-matched C IV emission line profiles. They found that the shape of the Lyα+N V region of these spectra was nearly identical to ULAS J1120+0641, suggesting that there may not be any damped absorption at all.

The most recent analysis of the ULAS J1120+0641 damping wing profile was undertaken by Greig et al. (2017a). Similar to this work, they trained a predictive model for the intrinsic blue-side continuum from a large sample of BOSS quasar spectra (Greig et al. 2017b). Their parametric model predicts Gaussian emission line parameters for Lyα (line width, amplitude, and velocity shift of two components) from fits to several broad emission lines on the red side of the spectrum. In ad-

Table 1. Neutral fraction constraints from the proximity zone and damping wing

Quasar	z	$\langle x_{\rm HI} \rangle$ $(10^3 \leq t_{\rm q}/{\rm yr} \leq 10^8)$	$\langle x_{\rm HI} \rangle$ $(10^5 \leq t_{\rm q}/{\rm yr} \leq 10^8)$
ULAS J1120+0641	7.0851	$0.48^{+0.26}_{-0.26}(^{+0.47}_{-0.46})$	$0.52^{+0.25}_{-0.25}(^{+0.44}_{-0.46})$
ULAS J1342+0928	7.5413	$0.60^{+0.20}_{-0.23}(^{+0.36}_{-0.45})$	$0.67^{+0.19}_{-0.23}(^{+0.31}_{-0.45})$

The tabulated constraints represent the median and 68% (95%) credible intervals obtained via linear interpolation of the $t_{\rm q}$-marginalized posterior PDFs in Figure 11.

dition, they separately fit the N V+Si II complex at $1230 < \lambda_{\rm rest} < 1275$ Å and introduce this fit as a strict prior to the model for the Lyα damping wing region ($1218 < \lambda_{\rm rest} < 1230$ Å). Again similar to our analysis, Greig et al. (2017a) employed a large-volume semi-numerical simulation of reionization (the Evolution Of 21 cm Structure simulation, Mesinger et al. 2016) to characterize the large-scale distribution of neutral gas around massive halos during the reionization epoch. By restricting their analysis to wavelengths redward of Lyα, they did not need to explicitly model the proximity zone of the quasar. To approximate the effect of the quasar ionizing radiation on neutral gas along the line of sight, they ionize the first 16 comoving Mpc (~ 2 proper Mpc) of every sightline to be consistent with the observed profile, but beyond this distance the quasar does not affect the ionization topology. Their final statistical constraints on $\langle x_{\rm HI} \rangle$ were derived from a χ^2-based likelihood analysis of intrinsic Lyα profiles drawn from their predictive continuum model multiplied by the damping wing absorption from sightlines through their reionization simulation. While Greig et al. (2017a) presented the most sophisticated study of a quasar damping wing at the time, their method has a handful of potential shortcomings which we describe below.

First, their fit of the N V and Si II complex and subsequent prior on the spectrum begins at $\lambda_{\rm rest} = 1230$ Å under the assumption that the damping wing absorption there is minimal. However, the absorption at $\lambda_{\rm rest} \sim 1230$ Å can still be significant at large $\langle x_{\rm HI} \rangle$, a fact which can be readily seen in the inset panel of Figure 2 in Greig et al. (2017a) where their best-fit damping wing model still shows $> 2\sigma$ absorption at $\lambda_{\rm rest} \sim 1230$ Å (see also Figure 4). This prior may then result in a bias towards lower $\langle x_{\rm HI} \rangle$, as smooth absorption redward of $\lambda_{\rm rest} \sim 1230$ Å may be fitted out, although it is unclear what the magnitude of this effect would be in practice.

Additionally, while they account for scatter in the Gaussian parameters for the Lyα line, they do not appear to account for any error in the Gaussian fits themselves. Indeed, as described in Greig et al. (2017b), they remove quasars from the training set whose spectra are "not well fit or characterized" by their double-Gaussian model for the Lyα line. While we have also excluded some discrepant quasars from our analysis, they were exclusively the most extreme cases of BALs and associated absorption which would be readily apparent in the spectra of $z > 7$ quasars. The distribution of possible continua for ULAS J1120+0641 used in Greig et al. (2017a) thus represents an underestimate of the true continuum error, lacking the additional error resulting from any deviations in the true continuum from the multiple Gaussian model.

Finally, by always treating the first ~ 2 proper Mpc of every sightline as ionized, Greig et al. (2017a) introduce a complicated prior on the quasar lifetime. For sightlines which originally intersect neutral gas within 2 proper Mpc, the quasar must have been on long enough to ionize material out to that distance. However, for sightlines where the first 2 proper Mpc are already ionized, the quasar then has no effect at all on neutral gas along the line of sight, implying a very short lifetime. As such the Greig et al. (2017a) posterior PDF cannot be considered fully marginalized over quasar lifetime, making a direct comparison between our $t_{\rm q}$-marginalized $\langle x_{\rm HI} \rangle$ constraints (Figure 11) and the results of the Greig et al. (2017a) analysis very difficult.

In Figure 13, we compare our blue-side prediction for ULAS J1120+0641 to the Mortlock et al. (2011) composite spectrum and the Greig et al. (2017a) model, where the latter has been renormalized to match our prediction at $\lambda_{\rm rest} = 1245$ Å. While our method predicts a very similar continuum to Mortlock et al. (2011), the intrinsic Lyα emission line strength predicted by Greig et al. (2017a) is dramatically higher (albeit with nearly identical Lyα centroids). Most importantly, however, we predict a substantially lower continuum (i.e. much closer to the observed spectrum) at $\lambda_{\rm rest} \sim 1225$ Å, the spectral region which contributed the most to the damping wing detection in Greig et al. (2017a). As a result, our measurement does not rule out $\langle x_{\rm HI} \rangle \sim 0$, although we nevertheless prefer somewhat higher $\langle x_{\rm HI} \rangle$ than Greig et al. (2017a) (see Figure 11), in large part due to the small proximity zone. While it is currently unclear why our two methods predict substantially different continua, we note that our predicted continuum is more consistent with the weak Lyα lines of the ULAS J1120+0641-analogs discussed in Bosman & Becker (2015), and with the nearest-neighbor quasars we used to calibrate the continuum uncertainty in Paper I.

6.3. *Choice of* $\langle x_{\rm HI}\rangle$ *and* t_q *Priors*

As mentioned above, we assume a flat prior on $\langle x_{\rm HI}\rangle$ from "0" (in truth, a model where reionization has finished with residual $\langle x_{\rm HI}\rangle \sim 10^{-3}$) to 1.0. If we were to instead assume a flat logarithmic prior (i.e. $p(\langle x_{\rm HI}\rangle) \propto 1/\langle x_{\rm HI}\rangle$) that extended down to $\langle x_{\rm HI}\rangle \sim 10^{-4}$, which would still be consistent with a completely opaque Lyα forest at $z \gtrsim 7$, our constraints on $\langle x_{\rm HI}\rangle$ would be dragged down to $\langle x_{\rm HI}\rangle \sim 0$, and there would be little evidence at all for ongoing reionization – the posterior PDF at $\langle x_{\rm HI}\rangle \sim 0$, which is small but non-negligible (Figure 11), would be boosted by a factor of $> 10^3$ relative to $\langle x_{\rm HI}\rangle > 0.1$. Because the damping wing signal is only detected at modest statistical significance in the context of the covariant continuum errors in our PCA method (even for ULAS J1342+0928), switching to a prior on $\langle x_{\rm HI}\rangle$ that is instead uniform in log space would shift the posterior PDF to peak close to $\langle x_{\rm HI}\rangle \sim 0$ for essentially *any* realistic quasar damping wing signal unless $\langle x_{\rm HI}\rangle \sim 1$ and $t_{\rm q}$ is very short. One could argue that our prior knowledge of $\langle x_{\rm HI}\rangle$ is not simply log-uniform, however, but can instead be thought of as bimodal. If reionization is complete, the Universe is "highly ionized," with $\langle x_{\rm HI}\rangle \lesssim 10^{-3}$ set by photoionization equilibrium with a metagalactic ionizing radiation field (e.g. Haardt & Madau 2012). If the Universe is undergoing the reionization phase transition, we instead have $\langle x_{\rm HI}\rangle$ of order unity.

While we are still starved for $z > 7$ quasars we will remain in a regime where quantitative constraints on $\langle x_{\rm HI}\rangle$ depend strongly on our choice of priors, but larger samples will greatly reduce this dependence: the prior enters the posterior PDF only once while each additional quasar contributes to the likelihood function. Assuming a Gaussian likelihood with $\langle x_{\rm HI}\rangle = 0.5$ and $\sigma_{\langle x_{\rm HI}\rangle} \sim 0.3$, similar to our constraint from the ULAS J1120+0641 spectrum, to counteract a factor of 10^4 prior advantage at $\langle x_{\rm HI}\rangle \sim 10^{-4}$ would require $\gtrsim 7$ quasars. Such a population of $z \gtrsim 7$ quasars is within reach of current programs exploiting wide-field optical and near-infrared surveys (e.g. Wang et al. 2017; Bañados et al. 2018).

Our fiducial prior on t_q is a flat, logarithmic prior from 10^3 to 10^8 years. The lower limit is motivated by the extremely small proximity zones in Eilers et al. (2017) whose sizes are consistent with such a short lifetime ($t_{\rm q} < 10^5$ years). The upper limit comes from the fact that, assuming accretion at the Eddington limit with 10% radiative efficiency, 10^8 years ago the quasar would have been roughly an order of magnitude fainter than currently observed. The quasar would then be effectively shut off at early times, so longer lifetimes would be largely irrelevant to the proximity zone and damping wing structure and can thus be excluded. Other estimates from the thermal proximity effect (Bolton et al. 2012) and He II transverse proximity effect (Schmidt et al. 2017b,a) suggest lifetimes of at least 10^7 years. The spectra of ULAS J1120+0641 and ULAS J1342+0928 both appear to exclude lifetimes at the upper and lower ends of our fiducial prior range (Figures 7 and 10), however, so expanding the bounds in either direction would make little difference to our posterior PDFs.

The choice of flat prior in linear space vs. log space for $t_{\rm q}$ is more subtle. If one assumes that all quasars live for a fixed amount of time $t_{\rm q,max}$, then a random quasar will have been shining continuously for a time $t_{\rm q}$ (which is what we have defined as "lifetime" in this work) drawn from a uniform distribution between $0 < t_{\rm q} < t_{\rm q,max}$, and so a flat prior in linear space would be appropriate. However, the uncertainty on this maximum lifetime spans multiple orders of magnitude (e.g. Martini 2004), so we believe that a flat prior in log space is reasonably well justified.

7. CONCLUSION

In this work we have used the intrinsic quasar continuum models of ULAS J1120+0641 and ULAS J1342+0928 from Paper I, in combination with extensive forward modeling of the proximity zone and damping wing features in the context of patchy reionization, to jointly constrain the lifetimes of the two quasars and the volume-averaged neutral fraction of the Universe at $z > 7$.

Our hybrid model of quasar spectra combines large-scale semi-numerical reionization simulations, hydrodynamical simulations, and 1D radiative transfer of ionizing photons from the quasars. We computed 2400 transmission spectra covering the proximity zone and damping wing for each pair of $\langle x_{\rm HI}\rangle$ and $\log t_{\rm q}$ on a coarse 21×11 grid for both quasars. Accounting for the covariant intrinsic quasar continuum uncertainty from Paper I, we can then construct realistic forward modeled representations of quasar transmission spectra. Based on these mock spectra we developed a Bayesian statistical method for recovering the joint posterior PDF of $\langle x_{\rm HI}\rangle$ and $\log t_{\rm q}$ from an observed quasar spectrum.

Applying our statistical methodology to the spectra of ULAS J1120+0641 at $z = 7.09$ (Mortlock et al. 2011) and ULAS J1342+0928 at $z = 7.54$ (Bañados et al. 2018), we found that both quasars are consistent with an ongoing epoch of reionization at $z > 7$. When marginalized over quasar lifetimes from 10^3 to 10^8 years, the resulting medians and 68% credible intervals of the posterior PDFs are $\langle x_{\rm HI}\rangle(z = 7.09) = 0.48^{+0.26}_{-0.26}$ and $\langle x_{\rm HI}\rangle(z = 7.54) = 0.60^{+0.20}_{-0.23}$.

Using our method it should be possible to constrain $\langle x_{\rm HI}\rangle$ at lower redshifts $z \sim$ 6–7 where there are far more quasars known (e.g. Venemans et al. 2013, 2015; Bañados et al. 2016; Reed et al. 2017; Wang et al. 2017; Mazzucchelli et al. 2017). The most constraining measurements at lower redshift to date are the model-independent upper limits from McGreer et al. (2015) who measured the fraction of dark pixels in the co-spatial Lyα and Lyβ forests, shown as red points in Figure 12, but this method becomes less constraining as the Ly-series forests become almost entirely opaque at $z \gtrsim 6$, which may simply result from density evolution in the IGM or a mild decrease in the ionizing background towards higher redshift (e.g. Davies et al. 2017). We predict that we will be able to obtain stronger constraints than the $z \sim 6.1$ upper limit of McGreer et al. (2015) from the multitude of existing spectra of quasars at $z \sim 6 - 7$, an endeavor we leave for future work.

Large samples of $z \gtrsim 7$ quasars to be discovered in further follow-up of quasar candidates from ground-based surveys (e.g. ULAS, Lawrence et al. 2007; VIKING, Arnaboldi et al. 2007; VHS, McMahon et al. 2013; DECaLS[6]; UHS, Dye et al. 2018) and in future wide-field near-infrared surveys by Euclid and WFIRST, together with high signal-to-noise spectra from JWST, will allow for exquisitely precise constraints on $\langle x_{\rm HI}\rangle(z)$.

ACKNOWLEDGEMENTS

We would like to thank A. Price-Whelan and D. Hogg for consultation on statistical methods, and D. Stern for supporting the discovery of ULAS J1342+0928.

BPV and F. Walter acknowledge funding through the ERC grants “Cosmic Dawn” and “Cosmic Gas”.

REFERENCES

Almgren, A. S., Bell, J. B., Lijewski, M. J., Lukić, Z., & Van Andel, E. 2013, ApJ, 765, 39

Alvarez, M. A., & Abel, T. 2007, MNRAS, 380, L30

Arnaboldi, M., Neeser, M. J., Parker, L. C., et al. 2007, The Messenger, 127

Bañados, E., Venemans, B. P., Decarli, R., et al. 2016, ApJS, 227, 11

Bañados, E., Venemans, B. P., Mazzucchelli, C., et al. 2018, Nature, 553, 473

Becker, R. H., Fan, X., White, R. L., et al. 2001, AJ, 122, 2850

Bolton, J. S., Becker, G. D., Raskutti, S., et al. 2012, MNRAS, 419, 2880

Bolton, J. S., & Haehnelt, M. G. 2007, MNRAS, 374, 493

Bolton, J. S., Haehnelt, M. G., Warren, S. J., et al. 2011, MNRAS, 416, L70

Bosman, S. E. I., & Becker, G. D. 2015, MNRAS, 452, 1105

Bouwens, R. J., Illingworth, G. D., Oesch, P. A., et al. 2015, ApJ, 803, 34

Carswell, R. F., Whelan, J. A. J., Smith, M. G., Boksenberg, A., & Tytler, D. 1982, MNRAS, 198, 91

Cen, R., & Haiman, Z. 2000, ApJL, 542, L75

Dall’Aglio, A., Wisotzki, L., & Worseck, G. 2008, A&A, 491, 465

Davies, F. B., Furlanetto, S. R., & McQuinn, M. 2016, MNRAS, 457, 3006

[6] http://legacysurvey.org/

Davies, F. B., Hennawi, J. F., Eilers, A.-C., & Lukić, Z. 2017, ArXiv e-prints, arXiv:1703.10174

Davies, F. B., Hennawi, J. F., Bañados, E., et al. 2018, ArXiv e-prints, arXiv:1801.07679

Drovandi, C. C., Pettitt, A. N., & Lee, A. 2015, Statist. Sci., 30, 72

Dye, S., Lawrence, A., Read, M. A., et al. 2018, MNRAS, 473, 5113

Eilers, A.-C., Davies, F. B., Hennawi, J. F., et al. 2017, ApJ, 840, 24

Fan, X., Narayanan, V. K., Lupton, R. H., et al. 2001, AJ, 122, 2833

Fan, X., Strauss, M. A., Schneider, D. P., et al. 2003, AJ, 125, 1649

Fan, X., Strauss, M. A., Becker, R. H., et al. 2006, AJ, 132, 117

Furlanetto, S. R., & Johnson Stoever, S. 2010, MNRAS, 404, 1869

Furlanetto, S. R., Oh, S. P., & Briggs, F. H. 2006, PhR, 433, 181

Furlanetto, S. R., Zaldarriaga, M., & Hernquist, L. 2004, ApJ, 613, 1

Gnedin, N. Y., & Hui, L. 1998, MNRAS, 296, 44

Gourieroux, C., Monfort, A., & Renault, E. 1993, Journal of Applied Econometrics, 8, S85

Greig, B., Mesinger, A., Haiman, Z., & Simcoe, R. A. 2017a, MNRAS, 466, 4239

Greig, B., Mesinger, A., McGreer, I. D., Gallerani, S., & Haiman, Z. 2017b, MNRAS, 466, 1814

Gunn, J. E., & Peterson, B. A. 1965, ApJ, 142, 1633

Haardt, F., & Madau, P. 2012, ApJ, 746, 125

Keating, L. C., Haehnelt, M. G., Cantalupo, S., & Puchwein, E. 2015, MNRAS, 454, 681

Khaire, V., Srianand, R., Choudhury, T. R., & Gaikwad, P. 2016, MNRAS, 457, 4051

Khrykin, I. S., Hennawi, J. F., & McQuinn, M. 2017, ApJ, 838, 96

Khrykin, I. S., Hennawi, J. F., McQuinn, M., & Worseck, G. 2016, ApJ, 824, 133

Kramer, R. H., & Haiman, Z. 2009, MNRAS, 400, 1493

Kulkarni, G., Hennawi, J. F., Oñorbe, J., Rorai, A., & Springel, V. 2015, ApJ, 812, 30

Lacey, C., & Cole, S. 1993, MNRAS, 262, 627

Lawrence, A., Warren, S. J., Almaini, O., et al. 2007, MNRAS, 379, 1599

Lidz, A., McQuinn, M., Zaldarriaga, M., Hernquist, L., & Dutta, S. 2007, ApJ, 670, 39

Lukić, Z., Stark, C. W., Nugent, P., et al. 2015, MNRAS, 446, 3697

Lusso, E., Worseck, G., Hennawi, J. F., et al. 2015, MNRAS, 449, 4204

Martini, P. 2004, Coevolution of Black Holes and Galaxies, 169

Mazzucchelli, C., Bañados, E., Venemans, B. P., et al. 2017, ApJ, 849, 91

McGreer, I. D., Mesinger, A., & D'Odorico, V. 2015, MNRAS, 447, 499

McMahon, R. G., Banerji, M., Gonzalez, E., et al. 2013, The Messenger, 154, 35

Mesinger, A., & Furlanetto, S. 2007, ApJ, 669, 663

Mesinger, A., Furlanetto, S., & Cen, R. 2011, MNRAS, 411, 955

Mesinger, A., & Furlanetto, S. R. 2008, MNRAS, 385, 1348

Mesinger, A., Greig, B., & Sobacchi, E. 2016, MNRAS, 459, 2342

Mesinger, A., & Haiman, Z. 2007, ApJ, 660, 923

Miralda-Escudé, J. 1998, ApJ, 501, 15

Mortlock, D. J., Warren, S. J., Venemans, B. P., et al. 2011, Nature, 474, 616

Pâris, I., Petitjean, P., Rollinde, E., et al. 2011, A&A, 530, A50

Pâris, I., Petitjean, P., Ross, N. P., et al. 2017, A&A, 597, A79

Pedregosa, F., Varoquaux, G., Gramfort, A., et al. 2011, Journal of Machine Learning Research, 12, 2825

Planck Collaboration, Aghanim, N., Ashdown, M., et al. 2016a, A&A, 596, A107

Planck Collaboration, Adam, R., Aghanim, N., et al. 2016b, A&A, 596, A108

Rahmati, A., Pawlik, A. H., Raičević, M., & Schaye, J. 2013, MNRAS, 430, 2427

Reed, S. L., McMahon, R. G., Martini, P., et al. 2017, MNRAS, 468, 4702

Robertson, B. E., Ellis, R. S., Furlanetto, S. R., & Dunlop, J. S. 2015, ApJL, 802, L19

Rorai, A., Hennawi, J. F., & White, M. 2013, ApJ, 775, 81

Schmidt, T. M., Hennawi, J. F., Worseck, G., et al. 2017a, ArXiv e-prints, arXiv:1710.04527

Schmidt, T. M., Worseck, G., Hennawi, J. F., Prochaska, J. X., & Crighton, N. H. M. 2017b, ApJ, 847, 81

Schroeder, J., Mesinger, A., & Haiman, Z. 2013, MNRAS, 428, 3058

Simcoe, R. A., Sullivan, P. W., Cooksey, K. L., et al. 2012, Nature, 492, 79

Sorini, D., Oñorbe, J., Hennawi, J. F., & Lukić, Z. 2017, ArXiv e-prints, arXiv:1709.03988

Stevans, M. L., Shull, J. M., Danforth, C. W., & Tilton, E. M. 2014, ApJ, 794, 75

Suzuki, N., Tytler, D., Kirkman, D., O'Meara, J. M., & Lubin, D. 2005, ApJ, 618, 592

Telfer, R. C., Zheng, W., Kriss, G. A., & Davidsen, A. F. 2002, ApJ, 565, 773

Venemans, B. P., Findlay, J. R., Sutherland, W. J., et al. 2013, ApJ, 779, 24

Venemans, B. P., Bañados, E., Decarli, R., et al. 2015, ApJL, 801, L11

Venemans, B. P., Walter, F., Decarli, R., et al. 2017a, ApJL, 851, L8

—. 2017b, ApJ, 837, 146

Wang, F., Fan, X., Yang, J., et al. 2017, ApJ, 839, 27

White, R. L., Becker, R. H., Fan, X., & Strauss, M. A. 2003, AJ, 126, 1

Wyithe, J. S. B., Loeb, A., & Carilli, C. 2005, ApJ, 628, 575

Young, P. J., Sargent, W. L. W., Boksenberg, A., Carswell, R. F., & Whelan, J. A. J. 1979, ApJ, 229, 891

Zel'dovich, Y. B. 1970, A&A, 5, 84