



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

Frederick B. Davies¹, Joseph F. Hennawi^{1,2}, Eduardo Bañados^{3,9}, Zarija Lukic⁴, Roberto Decarli^{2,5}, Xiaohui Fan⁶, Emanuele P. Farina¹, Chiara Mazzucchelli², Hans-Walter Rix², Bram P. Venemans², Fabian Walter², 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 Max Planck Institut für Astronomie, Königstuhl 17, D-69117 Heidelberg, Germany

³ The Observatories of the Carnegie Institution for Science, 813 Santa Barbara Street, Pasadena, CA 91101, USA

⁴ Lawrence Berkeley National Laboratory, CA 94720-8139, USA

⁵ INAF–Osservatorio Astronomico di Bologna, via Gobetti 93/3, I-40129 Bologna, Italy

⁶ Steward Observatory, The University of Arizona, 933 North Cherry Avenue, Tucson, AZ 85721-0065, USA

⁷ Department of Astronomy, School of Physics, Peking University, Beijing 100871, People’s Republic of China

⁸ Kavli Institute for Astronomy and Astrophysics, Peking University, Beijing 100871, People’s Republic of China

Received 2018 February 15; revised 2018 July 10; accepted 2018 July 25; published 2018 September 11

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 $x_{\text{H I}}(z)$ at that epoch. Obtaining reliable quantitative constraints from this technique, however, is challenging due to stochasticity induced by the patchy inside-out topology of reionization, degeneracies with quasar lifetime, and the unknown unabsorbed quasar spectrum close to rest-frame Ly α . 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 quasars with the highest redshifts known, ULAS J1120+0641 and ULAS J1342+0928, and measure volume-averaged neutral fractions $\langle x_{\text{H I}} \rangle(z = 7.09) = 0.48^{+0.26}_{-0.26}$ and $\langle x_{\text{H I}} \rangle(z = 7.54) = 0.60^{+0.20}_{-0.23}$ (posterior medians and 68% credible intervals) when marginalized over quasar lifetimes of $10^3 \leq t_{\text{q}} \leq 10^8$ yr.

Key words: cosmology: observations – cosmology: theory – dark ages, reionization, first stars – intergalactic medium – quasars: absorption lines

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 Ly α 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) that constrains the volume of ionized IGM between the present day and $z \sim 1100$ (Planck Collaboration et al. 2016b), which suggests a characteristic reionization redshift of $z_{\text{re}} = 6.4\text{--}9.7$ (95% credible interval, Planck Collaboration et al. 2016a). 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., Becker et al. 2001; Fan et al. 2001, 2003; White et al. 2003). While

Gunn–Peterson troughs (Gunn & Peterson 1965) in the Ly α and Ly β 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 hydrogen neutral fraction of $\langle x_{\text{H I}} \rangle \gtrsim 10^{-4}$ (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 Strömberg 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 Ly α damping wing (Miralda-Escudé 1998), which suppresses the quasar continuum redward of rest-frame Ly α .

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 $\langle x_{\text{H I}} \rangle$ 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 $\langle x_{\text{H I}} \rangle$ is to fit the Ly α transmission spectrum with the analytic model of Miralda-Escudé (1998), as performed by Mortlock et al. (2011), but this formula does not

⁹ Carnegie-Princeton Fellow.

include the substantial resonant Ly α absorption by residual H I inside 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 $\langle x_{\text{H I}} \rangle$ 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 $\langle x_{\text{H I}} \rangle$, but they only considered wavelengths redward of Ly α . 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 $z \gtrsim 6.5$, particularly the blueshifts of their C IV emission lines, are often extreme outliers of the distribution of lower-redshift quasars (Mazzuchelli 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 quasar spectra from the Sloan Digital Sky Survey (SDSS) 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 Ly α spectral shapes consistent with the observed spectrum of ULAS J1120+0641. 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 Ly α 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 (Baryon Oscillation Spectroscopic Survey) 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 $\langle x_{\text{H I}} \rangle$ in the surrounding IGM using the model of Miralda-Escudé (1998) 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. (2018a, 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 $1175 < \lambda_{\text{rest}} < 1280 \text{ \AA}$, from the “red-side” spectrum, covering $1280 < \lambda_{\text{rest}} < 2850 \text{ \AA}$. 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 the 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 $\langle x_{\text{H I}} \rangle$ 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. 2016a).

The rest of the paper is structured as follows. In Section 2, we briefly summarize the PCA method for predicting the intrinsic blue-side quasar continuum from the red-side spectrum from Paper I. In Section 3 we describe our hybrid model for quasar proximity zones and the Ly α 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 Section 4 we describe our methodology for performing Bayesian parameter inference from millions of forward-modeled mock spectra. In Section 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 Section 6 we conclude with a discussion of the implications of the constraints on the neutral fraction 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 Λ CDM 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 ($1175 < \lambda_{\text{rest}} < 1280 \text{ \AA}$) from the observed red-side spectrum ($1280 < \lambda_{\text{rest}} < 2850 \text{ \AA}$) 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 $2.09 < z_{\text{pipe}} < 2.51$ with a signal-to-noise ratio > 7 at $\lambda_{\text{rest}} = 1290 \text{ \AA}$, 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 Ly α 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 A_i with corresponding weights a_i ,

$$\log F \approx \langle \log F \rangle + \sum_i^n a_i A_i, \quad (1)$$

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

$$F \approx e^{\langle \log F \rangle} \prod_i^n e^{a_i A_i}. \quad (2)$$

This log-space decomposition naturally accounts for the variations in continuum slope between quasars, which dominate the total variance in flux space. For our analysis, we decomposed the red-side and blue-side spectra independently, keeping 10 red-side (R_i) and 6 blue-side (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_{temp} , allowing us to place each quasar onto a consistently defined rest frame. The blue-side coefficients b_i for each quasar in the training set were then found by fitting the blue-side (in the z_{temp} frame) spline fit continua assuming constant noise. From the sets of r_i and b_i for all quasars in the training set, \mathbf{r} and \mathbf{b} , we follow Suzuki et al. (2005) and P aris et al. (2011) and compute the *projection matrix* \mathbf{X} by finding the least-squares solution to the linear equation

$$\mathbf{b} = \mathbf{r} \cdot \mathbf{X}. \quad (3)$$

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 analyses of the proximity zone and damping wing ($1210 < \lambda_{\text{rest}} < 1240 \text{ \AA}$), 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; Ba˜nados 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

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

where $N_{\text{PCA},r}$ 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 spectra in the training set 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 for 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 quasars, i.e., the 1% nearest neighbors, as described above. We then use draws from these custom error distributions to generate forward-modeled mock spectra, described in more detail in Section 4. We note that, as shown in Appendix B of Paper I, additional uncertainties in the continuum estimation related to near-infrared echelle calibration artifacts may be present that we do not explicitly account for in this work.

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 must now 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; F. B. Davies & S. R. Furlanetto 2018, in preparation).
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 (Z. Lukić et al. 2018, in preparation), which is described in Sorini et al. (2018). We extract 1200 axis-aligned skewers from the centers of the 200 most massive halos, corresponding to halo masses $M_h \gtrsim 2 \times 10^{11} M_\odot$. The simulation, optimized for studying the Ly α 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 $\sim 1\text{--}2$ proper Mpc scales covered by the proximity zones of the $z > 7$ quasars in our analysis (compared to the halo virial radius, ~ 50 proper kpc).

We rescale 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 21cmFAST¹⁰ (Mesinger et al. 2011), to be presented in further detail in F. B. Davies & S. R. Furlanetto (2018, in preparation). The 21cmFAST code computes the fraction of material that has collapsed into dark matter halos, f_{coll} , following conditional Press–Schechter theory (Lacey & Cole 1993) applied to a nonlinear density field computed using the Zel’dovich approximation (Zel’dovich 1970). A region is considered ionized if $f_{\text{coll}} > \zeta^{-1}$ 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, recomputing f_{coll} at each filter scale. Our modified algorithm assigns collapsed mass to each cell according to the nonlinear density field, and it is this collapsed mass field that 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 $\langle x_{\text{HI}} \rangle$, 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




Figure 1. One-pixel (~ 800 comoving kpc) slice through the semi-numerical ionization field at $z = 7.5$ with $\langle x_{\text{HI}} \rangle = 0.5$ (grayscale; a linear stretch with white corresponding to neutral and black corresponding to fully ionized). The locations of $M_h > 10^{11} M_\odot$ 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 large regions that have already been ionized.

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 $M_{\text{min}} = 10^8 M_\odot$, and mean free path of ionizing photons $\lambda_{\text{mfp}} = 60$ cMpc. We then tuned ζ to produce ionization fields with global volume-averaged neutral fractions of $\langle x_{\text{HI}} \rangle = 0.05\text{--}0.95$ in steps of $\Delta x_{\text{HI}} = 0.05$.

The minimum halo mass was chosen to roughly correspond to a virial temperature of 10^4 K, the atomic cooling limit for gas of primordial composition. The relatively long λ_{mfp} was chosen to limit the impact of the new treatment of ionizing photon absorption to be presented in F. B. Davies & S. R. Furlanetto (2018, in preparation), and we note that a shorter value more consistent with extrapolation from observations at $z \lesssim 5$ (~ 30 cMpc, Worseck et al. 2014) would result in somewhat smaller initially ionized regions around massive halos.

Massive dark matter halos reside within larger-scale overdensities that 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 $\langle x_{\text{HI}} \rangle = 0.5$ 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 ionized regions (black). As seen by the subtle gray shading, the

¹⁰ <https://github.com/andreimesinger/21cmFAST>




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

ionization field is not entirely a binary neutral (white) versus 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_{HI} from the locations of the 500 most massive halos, corresponding to $M_{\text{h}} \gtrsim 3 \times 10^{11} M_{\odot}$. We show the distribution of distances from these massive halos to the nearest patch of neutral hydrogen ($x_{\text{HI}} > 0.01$) as a function of $\langle x_{\text{HI}} \rangle$ 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 $\langle x_{\text{HI}} \rangle$, 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_{min} and λ_{mfp}).

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 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:

$$\frac{dn_{\text{H II}}}{dt} = n_{\text{H I}}(\Gamma_{\text{H I}}^{\gamma} + n_e \Gamma_{\text{H I}}^e) - n_{\text{H II}} n_e \alpha_{\text{H II}}^A, \quad (5)$$

$$\begin{aligned} \frac{dn_{\text{He II}}}{dt} = & n_{\text{He I}}(\Gamma_{\text{He I}}^{\gamma} + n_e \Gamma_{\text{He I}}^e) + n_{\text{He III}} n_e \alpha_{\text{He III}}^A \\ & - n_{\text{He II}}(\Gamma_{\text{He II}}^{\gamma} + n_e \Gamma_{\text{He II}}^e - n_e \alpha_{\text{He II}}^A), \end{aligned} \quad (6)$$

$$\frac{dn_{\text{He III}}}{dt} = n_{\text{He II}}(\Gamma_{\text{He II}}^{\gamma} + n_e \Gamma_{\text{He II}}^e) - n_{\text{He III}} n_e \alpha_{\text{He III}}^A, \quad (7)$$

where n_i are the number densities for each species, Γ_i^{γ} are the photoionization rates, Γ_i^e are the collisional ionization rates, and α_i^A are the Case A recombination rate coefficients. In Γ_i^{γ} we include the effect of *secondary* ionizations, as tabulated by Furlanetto & Johnson Stoeber (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

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

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

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

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

The gas temperature is evolved taking into account photoionization heating and cooling¹¹ from recombinations, collisional excitation, the expansion of the universe, and inverse Compton scattering off 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 template of Lusso et al. (2015) to convert from the measured M_{1450} to L_{ν} at the ionizing edge of hydrogen ($E_{\text{H I}} \approx 13.6$ 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 transmission profile of the proximity zone (F. B. Davies et al. 2018, in preparation) and the strength of the damping wing. However, the size of the ionized bubble around the quasar R_{ion} (assuming a fully neutral universe) is only weakly dependent on the ionizing photon output \dot{N}_{ion} ($R_{\text{ion}} \propto \dot{N}_{\text{ion}}^{1/3}$; 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.

¹¹ We assume that the gas is of primordial composition, i.e., there is no cooling due to elements heavier than helium.




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 $\langle x_{\text{HI}} \rangle = 0.5$ 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 $\text{Ly}\alpha$ transmission, the middle panels show x_{HI} , and the bottom panels show the gas temperature.

We assume a “lightbulb” model for quasar emission: the quasar turned on at some point t_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_q as the “quasar lifetime.”

As mentioned in Section 3.1, we rescale the density of the gas in the hydrodynamical skewers by $(1+z)^3$. However, we do not account for cosmological evolution during the radiative transfer calculation, i.e., we assume a fixed $z = z_q$ even for the longest quasar lifetimes that we consider, which correspond to $\Delta z \sim 1$. We expect these evolutionary effects to be subtle, and as shown later, the small proximity zones of the two $z > 7$ quasars appear to rule out such long lifetimes.

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.¹² For regions with $x_{\text{HI}} = 0$ in the reionization simulations, we assume a uniform ionizing background such that $\langle x_{\text{HI}} \rangle$ inside ionized regions is $\sim 10^{-3}$ (corresponding to a hydrogen

photoionization rate $\Gamma_{\text{HI}} \sim 6 \times 10^{-14}$), consistent with the most stringent constraints from the $\text{Ly}\beta$ forest (Fan et al. 2006), 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 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 ($M_{1450} = -26.76$, $z_q = 7.5413$) using a skewer from the hydrodynamical simulation with initial x_{HI} given by a skewer from the $\langle x_{\text{HI}} \rangle = 0.5$ semi-numerical reionization simulation. The top panel shows the $\text{Ly}\alpha$ 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 ($\lambda_{\text{rest}} > \lambda_{\text{Ly}\alpha}$), 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 $T \sim (3-4) \times 10^4$ K. After 100 million

¹² We note that this stitching procedure breaks the correlation between large-scale IGM density and ionization state; in general, large-scale regions that are neutral should be less dense than large-scale regions that are ionized (Furlanetto et al. 2004). This means that the large-scale neutral patches of the IGM will be slightly more dense on average than they would be in a self-consistent simulation, and thus the effect of neutral gas on the spectra is biased slightly high.

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 Ly α forest absorption becomes too strong as the ionizing flux from the quasar decreases (e.g., Bolton & Haehnelt 2007; Eilers et al. 2017) and x_{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 leads 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 model with the longest lifetime (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 $0 \leq \langle x_{\text{HI}} \rangle \leq 1$ in 21 steps of $\Delta x_{\text{HI}} = 0.05$, 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 $\Delta \log[t_q/\text{yr}] = 0.5$ in 11 steps from 10^3 to 10^8 yr covering a velocity range $-10,000 \leq v - v_{\text{sys}} \leq +10,000 \text{ km s}^{-1}$, where v_{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_q at fixed $\langle x_{\text{HI}} \rangle = 0.5$ (top), and as a function of $\langle x_{\text{HI}} \rangle$ at fixed $t_q = 10^6$ yr (bottom). There is a clear trend toward stronger damping 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 uniform ionization rate in the entire IGM instead of our more realistic patchy reionization 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, $t_q \gtrsim 10^8$ yr, the damping wing almost entirely disappears, even for $\langle x_{\text{HI}} \rangle \sim 1$.

At even shorter timescales, $t_q \lesssim 10^5$ yr, 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, F. B. Davies et al. 2018, in preparation).




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

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 $(\langle x_{\text{HI}} \rangle, t_q)$. 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 (Section 2) introduces an additional multiplicative error that is strongly covariant. The combination of these processes cannot be simply described 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^{-1} binned

pixels,¹³

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

which is equivalent to the likelihood function of the (500 km s⁻¹ binned) transmission spectrum in the absence of correlations between pixels. By binning over the equivalent of several resolution elements of the observed spectra, we can safely ignore instrumental resolution, and thus do not convolve our mock spectra prior to binning. 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.¹⁴ 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, θ_{MLE} , 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 loss of information in this compression).

We treat θ_{MLE} as a summary statistic and compute the posterior PDF of the “true” model parameters θ following Bayes’ theorem,

$$p(\theta|\theta_{\text{MLE}}) = \frac{p(\theta_{\text{MLE}}|\theta)p(\theta)}{p(\theta_{\text{MLE}})}, \quad (12)$$

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

$$p(\theta_{\text{MLE}}) = \int p(\theta_{\text{MLE}}|\theta)p(\theta)d\theta. \quad (13)$$

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):

$$\hat{F} = F \times (1 + \epsilon_{\text{C}}) + N. \quad (14)$$

We find θ_{MLE} 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 $\langle x_{\text{HI}} \rangle$ from 0 to 1, and flat *log* priors on t_{q} from 10^3 to 10^8 yr. The linear prior on $\langle x_{\text{HI}} \rangle$ reflects our expectation that $z > 7$ is in the midst of the reionization epoch, while the log

prior 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 yr. We will discuss the impact of this choice of priors in Section 6.4.

In this method, the data have been compressed into the measurement of θ_{MLE} , so any spectrum with a particular θ_{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 ($z = 7.5413$; Venemans et al. 2017a; Bañados et al. 2018). 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 Section 2 and shown in Figure 12 of Paper I. We find a best-fit template redshift of $z = 7.0834$, and a very small blueshift of $\Delta v = 63$ km s⁻¹ from the systemic frame defined by the centroid of the [C II] emission line of the host galaxy ($z = 7.0851$, Venemans et al. 2017b). The blue-side profile shows a hint of absorption redward of Ly α and a relatively small proximity zone ($R_{\text{p}} = 1.72$ 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 Ly α , 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 gray curve in Figure 6 and the 500 km s⁻¹ binned spectrum in black. For the statistical analysis, we only use pixels at $v - v_{\text{sys}} > -4400$ km s⁻¹ ($\lambda_{\text{rest}} \lesssim 1233$ Å) to avoid the strong (and

¹³ Our results are insensitive to the exact choice of binning; we find nearly identical results from 300 and 800 km s⁻¹ bins.

¹⁴ 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.

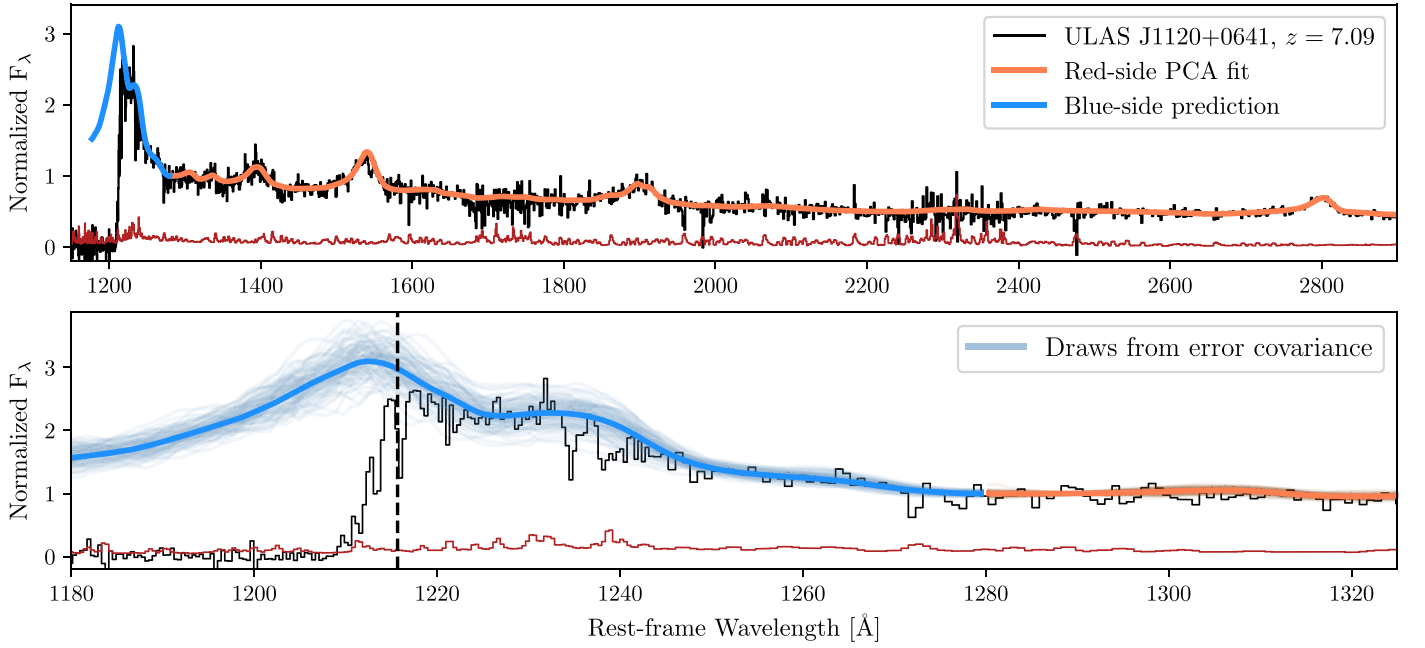


Figure 5. Top: VLT/FORS2 + Gemini/GNIRS spectrum of ULAS J1120+0641 (Mortlock et al. 2011, black) and its noise vector (red). The red-side PCA fit and blue-side prediction are shown as the orange and blue curves, respectively. Bottom: enlargement of the Ly α region of the spectrum, where the vertical dashed line shows rest-frame Ly α ($\lambda_{\text{rest}} = 1215.67 \text{ \AA}$). 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 (gray) and its noise vector (red). The black histogram shows the spectrum rebinned to $\sim 500 \text{ km s}^{-1}$ in the region we use for model comparison. The blue solid curve shows the median $\sim 500 \text{ km s}^{-1}$ binned transmission spectrum of mock spectra with the MLE parameter values $\theta_{\text{MLE}} = (\langle x_{\text{HI}} \rangle = 0.65, \log t_q = 6.5)$, while the associated blue shaded region shows the 16th–84th percentile range for mock spectra with $\theta = \theta_{\text{MLE}}$.

unresolved) associated N V absorption, and we choose to end the blue-side coverage at $v - v_{\text{sys}} = +6400 \text{ km s}^{-1}$ ($\lambda_{\text{rest}} \sim 1190 \text{ \AA}$) because all of the proximity zone models have no detectable signal beyond that distance. We find maximum pseudo-likelihood parameter values of $\theta_{\text{MLE}} = (\langle x_{\text{HI}} \rangle = 0.65, \log t_q = 6.5)$, and we show the median transmission profile of the MLE model (blue solid) and the expected 16th–84th percentile scatter (blue shaded) from forward-modeled spectra with $\theta = \theta_{\text{MLE}}$ in Figure 6. From the MLE parameter values we infer the 2D posterior PDF $p(\theta|\theta_{\text{MLE}})$ shown in Figure 7. The posterior PDF is relatively flat across a wide swathe of




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

$(\langle x_{\text{HI}} \rangle, \log t_q)$ parameter space, with a trend toward higher $\langle x_{\text{HI}} \rangle$ for longer t_q , reflecting the degeneracy between these two parameters discussed in Section 3.4 and shown in Figure 4. The non-zero size of the proximity zone rules out quasar lifetimes shorter than $\sim 10^{4.5} \text{ yr}$, while the combination of damping wing strength and small proximity zone rules out quasar lifetimes longer than $\sim 10^7 \text{ yr}$.

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




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 rebinned 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.

predicted blue-side (bias-corrected) continuum from Paper I. We find a best-fit template redshift of $z = 7.4438$, and a blueshift of $\Delta v = 3422 \text{ km s}^{-1}$ from the systemic frame ($z = 7.5413$, Venemans et al. 2017a). The red-side spectrum is very different from that of 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 $\text{Ly}\alpha$, and a very small proximity zone ($R_p = 1.20$ proper Mpc). The damping wing signal is clearly stronger and the proximity zone is even smaller than for ULAS J1120+0641, despite the slightly higher luminosity of ULAS J1342+0928. Both of these properties point toward a substantially neutral IGM surrounding ULAS J1342+0928.

We show the resulting transmission spectrum as the gray curve in Figure 9 and the 500 km s^{-1} binned spectrum in black. No strong associated absorption is visible in the spectrum, so we include all pixels redward of $\text{Ly}\alpha$ that are covered by our modeled transmission spectra ($v - v_{\text{sys}} < 10,000 \text{ km s}^{-1}$, $\lambda_{\text{rest}} \lesssim 1255 \text{ \AA}$). We cut off the blue-side coverage at $+4000 \text{ km s}^{-1}$ ($\sim 1200 \text{ \AA}$) 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.¹⁵ The maximum pseudo-likelihood parameter values are $\theta_{\text{MLE}} = (\langle x_{\text{HI}} \rangle = 0.8, \log t_q = 6.0)$, and we compare the binned transmission spectrum to the median transmission profile of the MLE model (solid purple) and the expected 16th–84th percentile scatter (dashed purple) from our forward-modeling in Figure 9. From

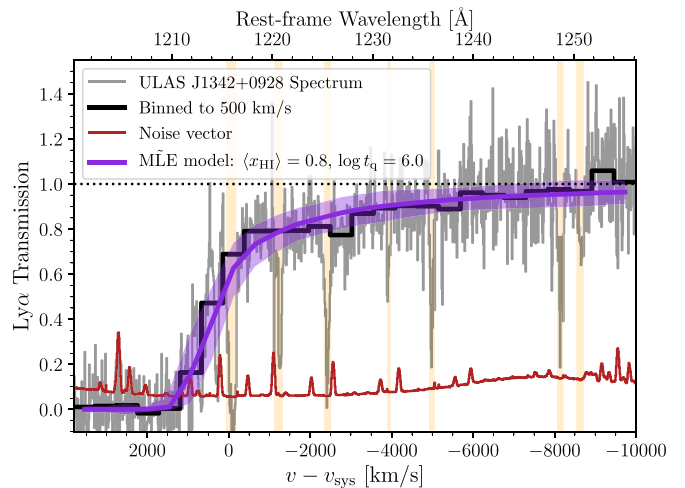


Figure 9. Similar to Figure 6 but for the continuum-divided spectrum of ULAS J1342+0928. The purple solid curve shows the median binned transmission spectrum in the mock spectra assuming the MLE parameter values $\theta_{\text{MLE}} = (\langle x_{\text{HI}} \rangle = 0.8, \log t_q = 6.0)$, while the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

the MLE parameter values we infer the 2D posterior PDF $p(\theta|\theta_{\text{MLE}})$ 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 $\langle x_{\text{HI}} \rangle$ and t_q .

6. Discussion

In the preceding sections, we have demonstrated a method for jointly constraining the global neutral fraction $\langle x_{\text{HI}} \rangle$ and quasar lifetime t_q from analysis of the proximity zone and (the presence or absence of) the damping wing, and applied it to the

¹⁵ 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).




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

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 $\langle x_{\text{HI}} \rangle$ marginalized over quasar lifetime, compare our analysis of ULAS J1120+0641 to analyses from previous works, and discuss our choice of $\langle x_{\text{HI}} \rangle$ and t_q priors.

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

In Figure 11, we show the posterior PDFs for $\langle x_{\text{HI}} \rangle$ from each quasar, marginalized over quasar lifetime with a flat prior in log space from 10^3 to 10^8 yr (solid curves) and a more restrictive prior from 10^5 to 10^8 yr (dashed curves), and provide the 68% and 95% credible intervals in Table 1. We chose to include extremely short quasar lifetimes of $\lesssim 10^5$ yr 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 yr. Including the possibility of the shortest quasar lifetimes shifts the $\langle x_{\text{HI}} \rangle$ posterior PDF to slightly lower values, consistent with the degeneracy between $\langle x_{\text{HI}} \rangle$ 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 $\langle x_{\text{HI}} \rangle \sim 0.5$, the posterior PDF is not particularly constraining, with significant probability density at the $\langle x_{\text{HI}} \rangle = 0$ and $\langle x_{\text{HI}} \rangle = 1$ boundaries. In contrast, the posterior PDF for ULAS J1342+0928 strongly indicates a significantly neutral IGM, ruling out $\langle x_{\text{HI}} \rangle < 0.08$ (0.14) at 99% probability marginalized over quasar lifetime for our prior covering $10^3 \leq t_q \leq 10^8$ yr ($10^5 \leq t_q \leq 10^8$ yr).

In Figure 12, we compare the posterior PDFs for $\langle x_{\text{HI}} \rangle$ 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. 2016a).¹⁶ Under our most conservative

¹⁶ 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 Ly α forest measurements that the end of reionization occurred before $z = 6$, see Section 5.3 of Planck Collaboration et al. (2016a).




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

prior, allowing quasar lifetimes from 10^3 to 10^8 yr, we find 68% (95%) credible intervals of $\langle x_{\text{HI}} \rangle(z = 7.09) = 0.48_{-0.26}^{+0.26}$ ($_{-0.46}^{+0.47}$) and $\langle x_{\text{HI}} \rangle(z = 7.54) = 0.60_{-0.23}^{+0.20}$ ($_{-0.45}^{+0.36}$). These constraints are consistent with the CMB and recent measurements based on Ly α emission properties of galaxies (Inoue et al. 2018; Mason et al. 2018), and are in broad agreement with recent calculations of the reionization history (e.g., Bouwens et al. 2015; Robertson et al. 2015; Khaire et al. 2016). The large cosmic variance between the damping wing profiles at fixed $\langle x_{\text{HI}} \rangle$ (Section 3.4), the strong degeneracy with quasar lifetime, and the limited precision of our continuum reconstructions greatly limit the constraining power of any single $z > 7$ quasar. However, with only a handful of additional quasars at $z > 7$, it may be possible to constrain $\langle x_{\text{HI}} \rangle(z)$ 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 Ly α , 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 expression of Miralda-Escudé (1998) for the damping wing. A follow-up 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 Ly α and the small proximity zone were suggestive of neutral gas close

Table 1
Constraints on the Neutral Fraction from the Proximity Zone and Damping Wing

Quasar	z	$\langle x_{\text{HI}} \rangle$ ($10^3 \leq t_q/\text{yr} \leq 10^8$)	$\langle x_{\text{HI}} \rangle$ ($10^5 \leq t_q/\text{yr} \leq 10^8$)
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	$0.60_{-0.23}^{+0.20+0.36}$ (-0.45)	$0.67_{-0.23}^{+0.19+0.33}$ (-0.45)

Note. The tabulated constraints represent the median and 68% (95%) credible intervals obtained via linear interpolation of the t_q -marginalized posterior PDFs in Figure 11.




Figure 12. Violin plot comparing the posterior PDFs from our analysis with the constraints on reionization history from Planck Collaboration et al. (2016a); the dark and light gray shaded regions correspond to the 68% and 95% credible intervals, respectively. Also shown are the dark pixel constraints on the Ly α +Ly β forest from McGreer et al. (2015) (red circles), the damping wing analysis of ULAS J1120+0641 from Greig et al. (2017a) (yellow square), statistics of Ly α emission in star-forming galaxies at $z \sim 7$ (Mason et al. 2018) (green pentagon), and the evolution of luminosity functions of Ly α emitters (Inoue et al. 2018) (brown hexagons and 1σ upper limits).

to the quasar ($x_{\text{HI}} \sim 0.1$), 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 composite spectrum of Mortlock et al. (2011) 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 Ly α 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 Ly α +NV region of these spectra was nearly identical to that of 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 Ly α (line width, amplitude, and

velocity shift of two components) from fits to several broad emission lines on the red side of the spectrum. In addition, they separately fit the NV+Si II complex at $1230 < \lambda_{\text{rest}} < 1275 \text{ \AA}$ and introduce this fit as a strict prior to the model for the Ly α damping wing region ($1218 < \lambda_{\text{rest}} < 1230 \text{ \AA}$). 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 Ly α , 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 $\langle x_{\text{HI}} \rangle$ were derived from a χ^2 -based likelihood analysis of intrinsic Ly α 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 NV and Si II complex and subsequent prior on the spectrum begins at $\lambda_{\text{rest}} = 1230 \text{ \AA}$ under the assumption that the damping wing absorption there is minimal. However, the absorption at $\lambda_{\text{rest}} \sim 1230 \text{ \AA}$ can still be significant at large $\langle x_{\text{HI}} \rangle$, a fact that 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 $\lambda_{\text{rest}} \sim 1230 \text{ \AA}$ (see also Figure 4). This prior may then result in a bias toward lower $\langle x_{\text{HI}} \rangle$, as smooth absorption redward of $\lambda_{\text{rest}} \sim 1230 \text{ \AA}$ 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 Ly α 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 Ly α line. While we have also excluded some discrepant quasars from our analysis, they were exclusively the most extreme cases of broad absorption line quasars and quasars with strong 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 that




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 model from Greig et al. (2017a) has been renormalized to our blue-side prediction at $\lambda_{\text{rest}} = 1245 \text{ \AA}$ to correct for the different flux calibration of the spectrum of Simcoe et al. (2012) used in their analysis.

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 posterior PDF of Greig et al. (2017a) cannot be considered fully marginalized over quasar lifetime, making a direct comparison between our t_q -marginalized $\langle x_{\text{H I}} \rangle$ constraints (Figure 11) and the results of the analysis of Greig et al. (2017a) very difficult.

In Figure 13, we compare our blue-side prediction for ULAS J1120+0641 to the composite spectrum of Mortlock et al. (2011) and the model of Greig et al. (2017a), where the latter has been renormalized to match our prediction at $\lambda_{\text{rest}} = 1245 \text{ \AA}$. While our method predicts a very similar continuum to Mortlock et al. (2011), the intrinsic Ly α emission line strength predicted by Greig et al. (2017a) is dramatically higher (albeit with nearly identical Ly α centroids). Most importantly, however, we predict a substantially lower continuum (i.e., much closer to the observed spectrum) at $\lambda_{\text{rest}} \sim 1225 \text{ \AA}$, the spectral region that contributed the most to the detection of the damping wing in Greig et al. (2017a). As a result, our measurement does not rule out $\langle x_{\text{H I}} \rangle \sim 0$, although we nevertheless prefer somewhat higher $\langle x_{\text{H I}} \rangle$ 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 Ly α 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. Comparison to Model A of Bañados et al. (2018)

We note here briefly the differences between our analysis of ULAS J1342+0928 and the default “Model A” presented in Bañados et al. (2018).

Bañados et al. (2018) modeled the intrinsic quasar continuum via a stack of 46 “analog” quasars with similar C IV emission line properties (equivalent width and blueshift relative to Mg II) identified in SDSS/BOSS spectra, and predicted a somewhat weaker blue-side continuum than our PCA analysis. We note that, by focusing only on the properties of C IV, their continuum prediction uses less information present in the red-side spectrum than our red-side PCA fit. In addition, the C IV properties within the selection region were not uniformly distributed, so the stacked spectrum is not “centered” on the properties of ULAS J1342+0928, whereas our PCA method effectively allows for smooth interpolation to any location in spectral parameter space.

Finally, Bañados et al. (2018) estimated $\langle x_{\text{H I}} \rangle$ via maximum likelihood estimation applied to the analytical model of Miralda-Escudé (1998) for IGM damping wings. While this model is effective for modeling damping wing absorption redward of Ly α , it does not include resonant Ly α absorption in the proximity zone, it assumes a uniform ionization state of the IGM, and it does not account for the local reionization bias due to the environment of massive quasar halos. In spite of these differences, the resulting constraints on $\langle x_{\text{H I}} \rangle$ are remarkably similar to those from our analysis.

6.4. Choice of $\langle x_{\text{H I}} \rangle$ and t_q Priors

As mentioned above, we assume a flat prior on $\langle x_{\text{H I}} \rangle$ from “0” (in truth, a model where reionization has finished with residual $\langle x_{\text{H I}} \rangle \sim 10^{-3}$) to 1.0. If we were to instead assume a flat logarithmic prior (i.e., $p(\langle x_{\text{H I}} \rangle) \propto 1/\langle x_{\text{H I}} \rangle$) that extended down to $\langle x_{\text{H I}} \rangle \sim 10^{-4}$, which would still be consistent with a completely opaque Ly α forest at $z \gtrsim 7$, our constraints on $\langle x_{\text{H I}} \rangle$ would be dragged down to $\langle x_{\text{H I}} \rangle \sim 0$, and there would be little evidence at all for ongoing reionization—the posterior PDF at $\langle x_{\text{H I}} \rangle \sim 0$, which is small but non-negligible (Figure 11), would be boosted by a factor of $>10^3$ relative to $\langle x_{\text{H I}} \rangle > 0.1$. 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 $\langle x_{\text{H I}} \rangle$ that is instead uniform in log space would shift the posterior PDF to peak close to $\langle x_{\text{H I}} \rangle \sim 0$ for essentially *any* realistic quasar damping wing signal unless $\langle x_{\text{H I}} \rangle \sim 1$ and t_q is very short. One could argue that our prior knowledge of $\langle x_{\text{H I}} \rangle$ is not simply log-uniform, however, but can instead be thought of as bimodal. If reionization is complete, the universe is “highly ionized,” with $\langle x_{\text{H I}} \rangle \lesssim 10^{-3}$ 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 $\langle x_{\text{H I}} \rangle$ of order unity.

While we are still starved for $z > 7$ quasars we will remain in a regime where quantitative constraints on $\langle x_{\text{H I}} \rangle$ 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 $\langle x_{\text{H I}} \rangle = 0.5$ and $\sigma_{\langle x_{\text{H I}} \rangle} \sim 0.3$, similar to our constraint from the ULAS J1120+0641 spectrum, it would require $\gtrsim 7$ quasars to counteract a factor of 10^4 prior advantage at $\langle x_{\text{H I}} \rangle \sim 10^{-4}$. 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 yr. 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 ($t_q < 10^5$ yr). The upper limit comes from the fact that, assuming accretion at the Eddington limit with 10% radiative efficiency, 10^8 yr 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. 2017, 2018) suggest lifetimes of at least 10^7 yr. The spectra of ULAS J1120+0641 and ULAS J1342+0928 both appear to exclude lifetimes at the upper and lower ends of the range for our fiducial prior (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 versus log space for t_q is more subtle. If one assumes that all quasars live for a fixed amount of time $t_{q,\max}$, then a random quasar will have been shining continuously for a time t_q (which is what we have defined as “lifetime” in this work) drawn from a uniform distribution over the range $0 < t_q < t_{q,\max}$, 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 $\langle x_{\text{H I}} \rangle$ and $\log t_q$ 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 $\langle x_{\text{H I}} \rangle$ and $\log t_q$ 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 yr, the resulting medians and 68% credible intervals of the posterior PDFs are $\langle x_{\text{H I}} \rangle(z = 7.09) = 0.48^{+0.26}_{-0.26}$ and $\langle x_{\text{H I}} \rangle(z = 7.54) = 0.60^{+0.20}_{-0.23}$.

It should be possible to use our method to constrain $\langle x_{\text{H I}} \rangle$ 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; Mazzucchelli et al. 2017; Reed et al. 2017; Wang 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 Ly α and Ly β 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 toward higher redshift (e.g., Davies et al. 2018b). 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, Dey et al. 2018; UHS, Dye et al. 2018) and in future wide-field near-infrared surveys by Euclid and WFIRST, together with spectra of high signal-to-noise ratio from the *James Webb Space Telescope*, will allow for exquisitely precise constraints on $\langle x_{\text{H I}} \rangle(z)$.

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.

B.P.V. and F. Walter acknowledge funding through the ERC grants “Cosmic Dawn” and “Cosmic Gas.”

ORCID iDs

Frederick B. Davies  <https://orcid.org/0000-0003-0821-3644>
 Joseph F. Hennawi  <https://orcid.org/0000-0002-7054-4332>
 Eduardo Bañados  <https://orcid.org/0000-0002-2931-7824>
 Roberto Decarli  <https://orcid.org/0000-0002-2662-8803>
 Xiaohui Fan  <https://orcid.org/0000-0003-3310-0131>
 Emanuele P. Farina  <https://orcid.org/0000-0002-6822-2254>
 Chiara Mazzucchelli  <https://orcid.org/0000-0002-5941-5214>
 Hans-Walter Rix  <https://orcid.org/0000-0003-4996-9069>
 Bram P. Venemans  <https://orcid.org/0000-0001-9024-8322>
 Fabian Walter  <https://orcid.org/0000-0003-4793-7880>
 Feige Wang  <https://orcid.org/0000-0002-7633-431X>
 Jinyi Yang  <https://orcid.org/0000-0001-5287-4242>

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, *Msngr*, 127, 28
- Bañados, E., Venemans, B. P., Decarli, R., et al. 2016, *ApJS*, 227, 11
- Bañados, E., Venemans, B. P., Mazzucchelli, C., et al. 2018, *Natur*, 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
- Cent, 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
- Davies, F. B., Hennawi, J. F., Bañados, E., et al. 2018a, *ApJ*, 864, 143
- Davies, F. B., Hennawi, J. F., Eilers, A.-C., & Lukić, Z. 2018b, *ApJ*, 855, 106
- Dey, A., Schlegel, D. J., Lang, D., et al. 2018, arXiv:1804.08657
- 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., Becker, R. H., et al. 2006, *AJ*, **132**, 117
- Fan, X., Strauss, M. A., Schneider, D. P., et al. 2003, *AJ*, **125**, 1649
- Furlanetto, S. R., & Johnson Stoeber, 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, *J. Appl. Econ.*, **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
- Inoue, A. K., Hasegawa, K., Ishiyama, T., et al. 2018, *PASJ*, **70**, 55
- 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, in *Coevolution of Black Holes and Galaxies* (Cambridge: Cambridge Univ. Press), 169
- Mason, C. A., Treu, T., Dijkstra, M., et al. 2018, *ApJ*, **856**, 2
- 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, *Msngr*, **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, *Natur*, **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, *JMLR*, **12**, 2825
- Planck Collaboration, Adam, R., Aghanim, N., et al. 2016a, *A&A*, **596**, A108
- Planck Collaboration, Aghanim, N., Ashdown, M., et al. 2016b, *A&A*, **596**, A107
- 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. 2018, *ApJ*, **861**, 122
- Schmidt, T. M., Worseck, G., Hennawi, J. F., Prochaska, J. X., & Crighton, N. H. M. 2017, *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, *Natur*, **492**, 79
- Sorini, D., Oñorbe, J., Hennawi, J. F., & Lukić, Z. 2018, *ApJ*, **859**, 125
- Stevens, 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., Bañados, E., Decarli, R., et al. 2015, *ApJL*, **801**, L11
- Venemans, B. P., Findlay, J. R., Sutherland, W. J., et al. 2013, *ApJ*, **779**, 24
- Venemans, B. P., Walter, F., Decarli, R., et al. 2017a, *ApJL*, **851**, L8
- Venemans, B. P., Walter, F., Decarli, R., et al. 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
- Worseck, G., Prochaska, J. X., O'Meara, J. M., et al. 2014, *MNRAS*, **445**, 1745
- 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