Lawrence Berkeley National Laboratory

Title
A New Method to Measure the Post-reionization Ionizing Background from the Joint Distribution of Lyα and Lyβ Forest Transmission∗ ∗ Some of the data presented herein were obtained at the W. M. Keck Observatory, which is operated as a scientific partn...

Authors
Davies, Frederick B
Hennawi, Joseph F
Eilers, Anna-Christina
et al.

A NEW METHOD TO MEASURE THE POST-REIONIZATION IONIZING BACKGROUND FROM THE JOINT DISTRIBUTION OF LYMAN-α AND LYMAN-β FOREST TRANSMISSION[*]

Frederick B. Davies[1], Joseph F. Hennawi[2,1], Anna-Christina Eilers[1,†], Zarija Lukić[3]

[1]Max-Planck-Institut für Astronomie, Königstuhl 17, D-69117 Heidelberg, Germany
[2]Department of Physics, University of California, Santa Barbara, CA 93106, USA
[3]Lawrence Berkeley National Laboratory, CA 94720-8139, USA

Draft version March 31, 2017

ABSTRACT

The amplitude of the ionizing background that pervades the intergalactic medium (IGM) at the end of the epoch of reionization provides a valuable constraint on the emissivity of the sources which reionized the Universe. While measurements of the ionizing background at lower redshifts rely on a simulation-calibrated mapping between the photoionization rate and the mean transmission of the Lyα forest, at $z \gtrsim 6$ the IGM becomes increasingly opaque, and transmission arises solely in narrow spikes separated by saturated Gunn-Peterson troughs. In this regime, the traditional approach of measuring the average transmission over large ~ 50 Mpc/h regions is less sensitive and sub-optimal. Additionally, the five times smaller oscillator strength of the Lyβ transition implies the Lyβ forest is considerably more transparent at $z \gtrsim 6$, even in the presence of contamination by foreground $z \sim 5$ Lyα forest absorption. In this work we present a novel statistical approach to analyze the joint distribution of transmission spikes in the co-spatial $z \sim 6$ Lyα and Lyβ forests. Our method relies on Approximate Bayesian Computation (ABC), which circumvents the necessity of computing the intractable likelihood function describing the highly correlated Lyα and Lyβ transmission. We apply ABC to mock data generated from a large-volume hydrodynamical simulation combined with a state-of-the-art model of ionizing background fluctuations in the post-reionization IGM, and show that it is sensitive to higher IGM neutral hydrogen fractions than previous techniques. As a proof of concept, we apply this methodology to a real spectrum of a $z = 6.54$ quasar and measure the ionizing background from $5.4 \leq z \leq 6.4$ along this sightline with ~ 0.2 dex statistical uncertainties.

1. INTRODUCTION

Following the recombination of the Universe and the release of the cosmic microwave background (CMB), the vast majority of baryonic matter in the intergalactic medium (IGM) consisted of neutral atoms, in stark contrast to the highly ionized IGM seen a few billion years later (Gunn & Peterson 1965) and at the present day (Field 1959). The first stars, galaxies, and black holes are believed to be responsible for the intervening phase transition known as the epoch of reionization (Loeb & Furlanetto 2013). Through study of the reionization process, we hope to understand the nature of the faintest and earliest collapsed structures.

The CMB itself provides an "integral" constraint on the total column density of ionized gas through the measured optical depth to electron scattering τ_e, and thus a characteristic reionization redshift $z_{\rm reion} \sim 8$ (Planck Collaboration et al. 2016a). However, a wide range of models for the reionization history are consistent with CMB measurements (Planck Collaboration et al. 2016b), and τ_e alone does not constrain the topology of the inho-

mogeneous reionization process (Furlanetto et al. 2004). An important boundary condition to the epoch of reionization is the ubiquitous transmission through the Lyα forest at $z \lesssim 5.5$ which is inconsistent with a predominantly neutral IGM (e.g. McGreer et al. 2011, 2015). Large stretches of neutral gas are still possible at higher redshift $z \gtrsim 6$ where the Lyα forest is almost entirely opaque.

The $z \sim 6$ Lyα forest has been studied for well over a decade to constrain the evolution of the ionization state of the Universe close to the reionization epoch (e.g. Becker et al. 2001; White et al. 2003; Songaila 2004; Fan et al. 2002, 2006). As shown by the data points in the second panel from the bottom of Figure 1, above $z \sim 5.5$ the typical opacity – and the typical variation between different lines of sight – increases rapidly. Eventually, the only transmission seen in the Lyα forest is in narrow transmission spikes, likely corresponding to small-scale regions of the universe with particularly low density (e.g. Oh & Furlanetto 2005). The large-scale transmission measurements then simply reflect the number and strength of these spikes that fall into wide bins in each spectrum, where the size is defined either in terms of redshift interval ($\Delta z \sim 0.15$; e.g. Fan et al. 2006) or by comoving distance (50 Mpc/h; Becker et al. 2015), both roughly corresponding to the scale of the first large-scale opaque regions, known as Gunn-Peterson troughs (after Gunn & Peterson 1965; henceforth GP troughs), observed in the spectra of $z > 6$ quasars (Fan et al. 2001; Becker et al. 2001; Fan et al. 2003; White et al. 2003). Indeed, some large-scale regions with formal limits on

their *mean* transmission contain high significance transmission spikes (Becker et al. 2015), which results in mean transmission below the significance threshold (2σ) when combined with surrounding GP troughs. Clearly, there is information in these small-scale transmission features which is lost when averaging on large scales, suggesting a "matched filtering" approach targeting transmission spikes would be more sensitive.

Another limiting factor for determining the ionization state of the $z \gtrsim 6$ Universe is that the Lyα forest becomes essentially completely saturated, even at small scales, with many lines of sight consistent with zero transmission. The red dotted line in the bottom panels of Figure 1 shows the 2σ limiting optical depth for 10 ks exposure time on a bright ($F_{\rm cont} = 10^{-17}$ erg s^{-1} cm^{-2}) high-redshift quasar with a 10-meter telescope. The solid curves of the panel second from the bottom show the evolution of the Lyα forest opacity computed from the hydrodynamical simulation described in § 2 assuming different toy models for the evolution in the hydrogen ionization rate $\Gamma_{\rm HI}$ and corresponding evolution in the volume-averaged hydrogen neutral fraction $\langle x_{\rm HI}\rangle_V$ shown in the top panels.

The Lyβ forest saturates later because of its weaker oscillator strength, $f_{\rm Ly\beta}\lambda_{\rm Ly\beta} \sim 0.16 f_{\rm Ly\alpha}\lambda_{\rm Ly\alpha}$ where f_i and λ_i are the oscillator strength and rest-frame wavelength of transition i, respectively. In practice the interpretation of Lyβ measurements is complicated by the presence of lower redshift Lyα absorption, which at $z_{\rm Ly\beta} \sim 6$ corresponds to $z_{\rm Ly\alpha} \sim 5$. Fortunately, the $z \sim 5$ Lyα forest has relatively small scatter and is well-understood compared to the strong fluctuations seen at $z > 5.5$ (e.g. Becker et al. 2015) and its statistics can be modeled with standard simulations of the IGM that employ a uniform ionizing background. The bottom panel of Figure 1 shows similar curves to the panel above it, but computed for the Lyβ forest including the presence of foreground Lyα, demonstrating that observable Lyβ transmission should persist to substantially higher redshift than Lyα at fixed observational cost.

The most striking example of the increased sensitivity provided by Lyβ is the ~ 110 Mpc/h GP trough observed in the $z = 5.98$ quasar ULAS J0148+0600 with an effective optical depth $\tau_{\rm eff} > 7.4$ (2σ limit) in the Lyα forest, but which shows many prominent transmission spikes in the Lyβ forest (see Figure 5 in Becker et al. 2015). The presence of Lyβ transmission proves the existence of a highly-ionized IGM in this region despite the complete saturation of Lyα, consistent with theoretical predictions for large-scale fluctuations in the ionizing background after the end of reionization (Davies & Furlanetto 2016) or relic temperature fluctuations due to inhomogeneous heating during reionization (D'Aloisio et al. 2015; Davies et al., in prep.).

In this work, we present theoretical expectations for the distribution of Lyα and Lyβ forest transmission on varying $\gtrsim$Mpc scales at $z \sim 6$, and demonstrate their ability to constrain the evolving ionizing background at high redshift. To account for sparseness, correlations, and extreme-value characteristics of potential data sets we apply a statistical method known as Approximate Bayesian Computation (ABC). ABC approximates the likelihood function through a massive Monte Carlo of mock data realizations, in contrast to more typical meth-
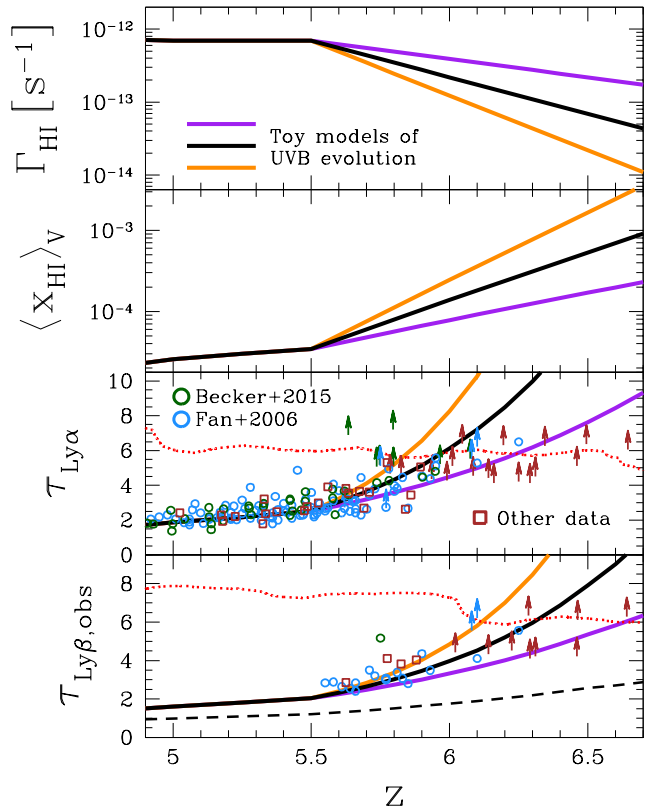

Figure 1. The top two panels show toy models for an evolving ionization rate $\Gamma_{\rm HI}$ and the corresponding volume-weighted IGM neutral fraction $\langle x_{\rm HI}\rangle_V$. The bottom two panels show the 50 Mpc/h Lyα and Lyβ forest $\tau_{\rm eff}$ measurements in Fan et al. (2006) and Becker et al. (2015) as the blue and green open circles, respectively, compared to simulated $\tau_{\rm eff}$ (solid curves) corresponding to the ionizing background evolution models in the top panel. The brown squares show quasar measurements not present in the other compilations (Willott et al. 2007; Mortlock et al. 2009; Goto et al. 2011; Tang et al. 2016; Barnett et al. 2017) and observations of high-redshift gamma ray bursts (Chornock et al. 2013, 2014). Upward arrows represent 2σ upper limits for measured transmission values that fall below reported 2σ uncertainties. The dotted curve in the bottom two panels shows the 2σ limiting effective optical depth on 50 Mpc/h scales for a simulated 10 ks Keck spectrum of a bright high-redshift quasar ($F_{\rm cont} = 10^{-17}$ erg s^{-1} cm^{-2} Å^{-1}). The long-dashed curve in the bottom panel shows the contribution of the foreground Lyα forest to the Lyβ forest $\tau_{\rm eff}$.

ods that assume a functional form (e.g. a multivariate Gaussian distribution). In § 2, we describe our numerical modeling of the Lyα and Lyβ forests and fluctuations in the ionizing background and present simulated distributions of Lyα and Lyβ forest transmission. In § 3, we describe our approach to measure the strength of the ionizing background using ABC. In § 4, we apply our statistical method to real data as a proof-of-concept test, providing (model-dependent) constraints on $\Gamma_{\rm HI}$ at $z \gtrsim 5.5$ from a single high-redshift quasar spectrum. In § 5, we conclude with a summary and a discussion of predicted constraints from future samples of $z > 6$ quasar spectra, where the addition of more quasars should both tighten the constraints on $\Gamma_{\rm HI}$ and allow for constraints on the topology of ionizing background fluctuations.

All distance units are comoving unless specified otherwise. We assume a ΛCDM cosmology with $\Omega_m = 0.3$,

$\Omega_\Lambda = 0.7$, $\Omega_b = 0.047$, $h = 0.685$, and $\sigma_8 = 0.8$.

2. MODELING THE Lyα AND Lyβ FORESTS

In this section, we describe our method for modeling the Lyα and Lyβ forests in the spectrum of $z \gtrsim 6$ quasars. Our physical model for the IGM is a cosmological hydrodynamical simulation using the Nyx code (Almgren et al. 2013; Lukić et al. 2015), 100 Mpc/h on a side, with 4096^3 dark matter particles and 4096^3 baryon fluid cells. We use outputs of the density, temperature, and velocity fields at $z = 6.0$ and $z = 5.0$ to model the $z \sim 6$ Lyα and Lyβ forest and the $z \sim 5$ Lyα contamination of the $z \sim 6$ Lyβ forest, respectively. We extract random skewers from these simulation outputs, compute the H I fraction in every cell assuming ionization equilibrium, and then calculate Lyα and Lyβ transmission along the line of sight including the effects of peculiar motions and thermal broadening (see Lukić et al. 2015 for more details). We re-scale the physical gas density along each skewer by $(1 + z)^3$ to account for cosmological expansion, making the simplifying assumption that within the redshifts of interest the evolution of the overdensity field is less important.

Our simulation meets the suggested mass resolution requirement to resolve the $z = 5$ Lyα forest proposed by Bolton & Becker (2009) ($M_{\rm gas} \lesssim 3 \times 10^5$ M$_\odot$), although convergence may be different for a grid code like Nyx versus the smoothed-particle hydrodynamics (SPH) simulations they studied. Convergence may also be more difficult for the extremely low transmission of the $z \gtrsim 6$ Lyα forest, suggested by the lack of convergence between the co-spatial Lyα and Lyβ optical depths at $z \sim 6$ of relatively high-resolution simulations shown in the Appendix of Becker et al. (2015). In future work we will quantify the convergence (or lack thereof) of the $z \gtrsim 6$ IGM in our simulation, extending the tests in Lukić et al. (2015) which were limited to $z \leq 4$, but preliminary tests suggest that the mean Lyα transmission is converged to better than $\sim 10\%$ (J. Oñorbe, private communication). We note that the convergence of the $z = 5$–6 Lyα forest is assisted by the non-adaptive spatial resolution of the Eulerian grid in the most underdense environments which lead to Lyα forest transmission (Oh & Furlanetto 2005) (as opposed to, e.g., nearest-neighbor smoothing kernels in SPH codes).

We include noise in our mock spectra using a noise vector extracted from a spectrum of the $z \sim 6.4$ quasar SDSS J1148+5251 taken with the Echellete Spectrograph and Imager (ESI; Sheinis et al. 2002) on the Keck II telescope with a total exposure time of 13,200s (PI: Cowie), re-scaled to an effective exposure time of 10,000s by multiplying the noise by $\sqrt{13200/10000} \sim 1.15$. The noise vector includes contributions from read noise, sky noise, and the varying sensitivity of ESI as a function of wavelength. We assume that the noise is independent and Gaussian-distributed in each pixel and that the presence of source signal does not affect the noise level, a good assumption for the faint $z \gtrsim 6$ quasars for which the sky background dominates the noise budget. The quasar spectrum is assumed to follow the composite spectrum of Lusso et al. (2015) assuming a flux density at 1450Å of $F_\lambda(1450\text{Å}) = 10^{-17}$ erg s^{-1} cm^{-2} Å^{-1}, corresponding to a fairly bright $z \sim 6$ quasar with $M_{1450} \sim -27$.
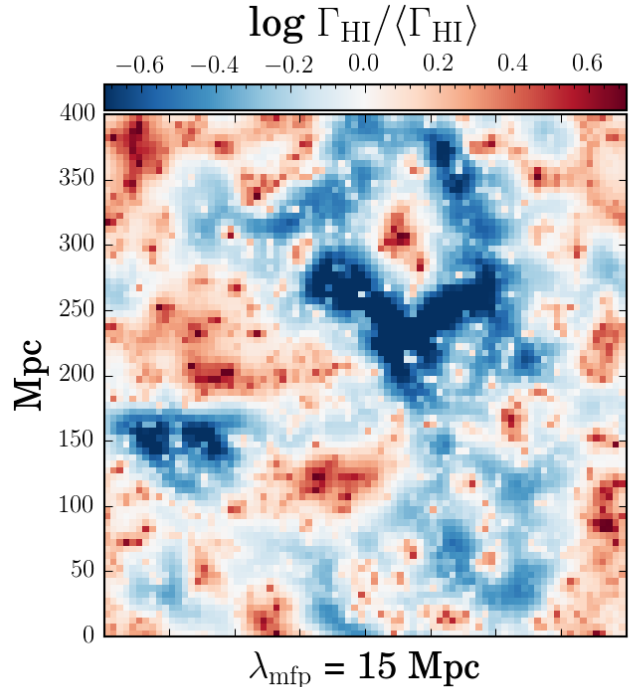

Figure 2. Slice through our simulation of the fluctuating ionizing background at $z = 6$ computed as in Davies & Furlanetto (2016) with the mean free path at the average strength of the ionizing background fixed to 15 Mpc. Strong large-scale fluctuations exist on ~ 100 Mpc scales as required by observations of Gunn-Peterson troughs at $z \gtrsim 5.5$ (Fan et al. 2006; Becker et al. 2015).

We model fluctuations in the post-reionization ionizing background in a separate $(400 \text{ Mpc})^3$ cosmological volume using the method of Davies & Furlanetto (2016, henceforth DF16). We defer the interested reader to DF16 for the details, but we briefly summarize the method here. We use the semi-numerical simulation code DEXM (Mesinger & Furlanetto 2007) to generate a realization of cosmological initial conditions in a $(400 \text{ Mpc})^3$ volume and populate it with dark matter halos down to $M_{\rm h,min} = 2 \times 10^9$ $M_\odot$ following the excursion set formalism. We then abundance match these halos to the measured Bouwens et al. (2015) $z \sim 6$ rest-frame UV luminosity function of galaxies down to absolute magnitudes of $M_{\rm UV,max} \sim -12.5$, and assume a constant ratio between the non-ionizing UV luminosity and the (escaping) ionizing luminosity of each galaxy. The ionizing background from these galaxies is computed on a coarse 80^3 grid allowing the mean free path of ionizing photons to vary spatially depending on the large-scale density field and local ionizing background, $\lambda \propto \Gamma_{\rm HI}^{2/3} \Delta^{-1}$, with the average mean free path in the volume normalized to $\lambda_{\rm mfp} = 15$ Mpc. In Figure 2 we show a slice through the ionizing background model, which exhibits strong large-scale fluctuations consistent with what is required to reproduce the distribution of Gunn-Peterson troughs in the $z \gtrsim 5.5$ Lyα forest (DF16).

We simulate the ionizing background in a separate volume from the hydrodynamical simulation because the latter is not large enough to fully model the effect of the fluctuating mean free path of ionizing photons, and thus would have much weaker ionizing background fluctuations on the ~ 100 Mpc scales we investigate in this
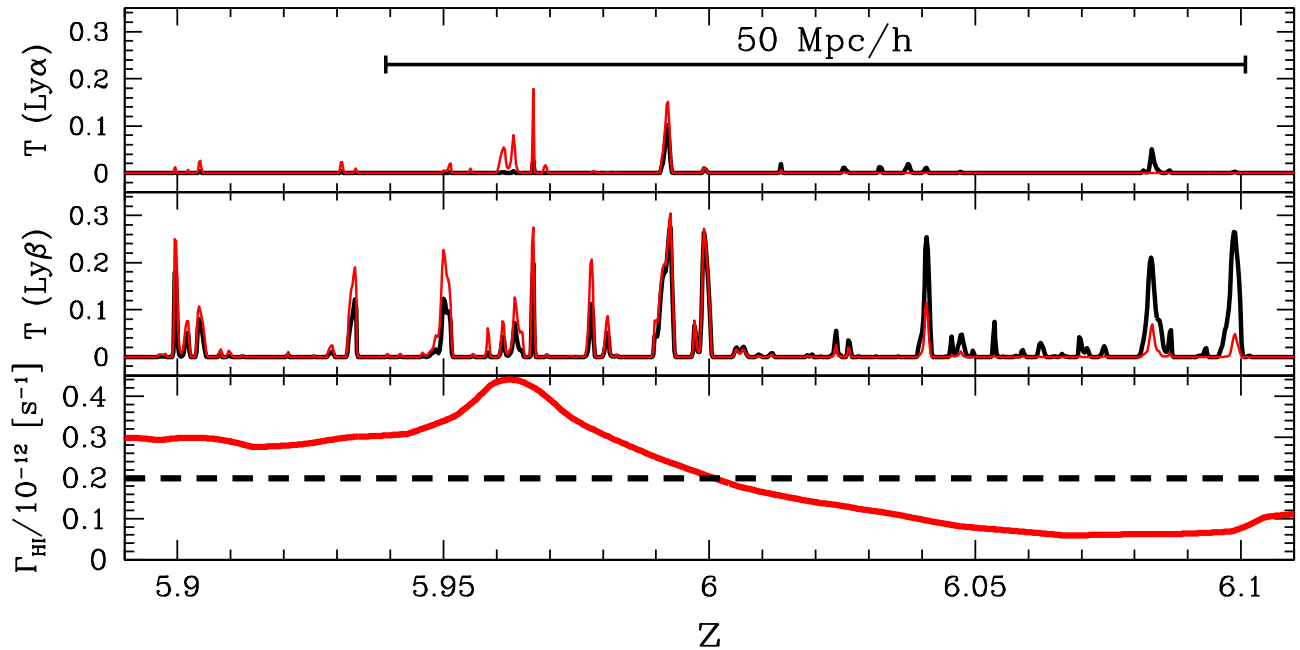
Figure 3. The top and middle panels show mock (noiseless) Lyα and Lyβ forest spectra, respectively, at $z \sim 6$. The black curve assumes a uniform ionization rate $\Gamma_{\mathrm{H\,I}} = 10^{-12.7}$ s^{-1}, the dashed curve in the bottom panel, while the red curve demonstrates the effect of a fluctuating ionizing background skewer shown by the solid red curve in the bottom panel. In the upper panel we show a bar corresponding to the typical scale of mean transmission measurements (50 Mpc/h) in the Lyα forest from the literature (e.g. Fan et al. 2006; Becker et al. 2015).

work (DF16). Because of this, we are explicitly ignoring the anti-correlation of the density and radiation fields on very large scales shown by DF16. In this work we assume that this large-scale anti-correlation does not strongly affect the statistics of small-scale transmission features in the IGM, an assumption we will test in future work.

In the top and middle panels of Figure 3, we show Lyα and Lyβ forest segments of a mock spectrum from $z = 5.9$–6.1. As expected, the Lyβ forest shows considerably more transmission, and despite contamination by foreground $z \sim 5$ Lyα forest absorption, transmission features in the Lyα forest tend to show counterparts in the Lyβ forest. The red curves show the effect of applying the fluctuating ionizing background skewer in the bottom panel. Small-scale features in the Lyα forest are sensitive to ionizing background fluctuations because the optical depth in each pixel changes as $\tau_\alpha \propto \Gamma_{\mathrm{H\,I}}^{-1}$. The Ly$\beta$ forest is less sensitive to these fluctuations because τ_β contains a constant factor from the foreground Lyα forest contamination which sets a floor to the opacity, i.e. $\tau_\beta = \tau_{\beta,z\sim6} + \tau_{\alpha,z\sim5}$ where only the first term responds to $z \sim 6$ background fluctuations[1].

2.1. *Lyα and Lyβ Transmission PDF on 2 & 20 Mpc Scales*

[1] In this work we assume that the ionizing background at $z \sim 5$ is uniform, a reasonable assumption given the relatively large measured mean free path of ionizing photons (Worseck et al. 2014, but see D'Aloisio et al. 2016) and the agreement between the distribution of large-scale Lyα forest optical depth with standard hydrodynamical simulations (Becker et al. 2015).

In the upper left set of panels in Figure 4 we show the transmission probability distribution function (PDF) on 2 Mpc scales (d$z \sim 0.005$) of the Lyα and Lyβ forests from simulations of $z \sim 6$ spectra with varying uniform $\Gamma_{\mathrm{H\,I}}$ of $10^{-12.4}$ s^{-1} (top, blue), $10^{-12.7}$ s^{-1} (middle, black), and $10^{-13.0}$ s^{-1} (bottom, orange). These $\Gamma_{\mathrm{H\,I}}$ values roughly correspond to the mean and $+/- 1\sigma$ uncertainty range of current measurements from the Lyα forest at $z \sim 6$ (Wyithe & Bolton 2011). The PDFs are characterized by a Gaussian peak around zero – corresponding to the noise floor, shown as the dotted line in every panel – and a tail towards higher transmission whose steepness is a function of ionizing background strength. The shape of the PDF is markedly different from the distributions seen at $z \sim 2$–3 (e.g. Lee et al. 2015) because the $z \gtrsim 6$ Lyα and Lyβ forests are so opaque that they almost never come close to the continuum (see also Becker et al. 2007). In the bottom panels the red curve corresponds to a model with the same average $\Gamma_{\mathrm{H\,I}} = 10^{-12.7}$ s^{-1} that includes fluctuations in the ionizing background (as shown in Figures 2 and 3). In general the effect of background fluctuations on the 2 Mpc-binned transmission PDF is minor, although it does lead to an extended tail of rare transmission spikes. The effect is larger on the Lyα PDF than the Lyβ PDF due to the contamination of Lyβ by foreground $z \sim 5$ Lyα forest absorption, i.e., even if $\Gamma_{\mathrm{H\,I}}$ fluctuates high, the transmission spikes can only increase up to the absorption level of the $z \sim 5$ Lyα forest (§ 2).

In the right set of panels in Figure 4 we show the PDFs on 20 Mpc scales. This is analogous to previous studies
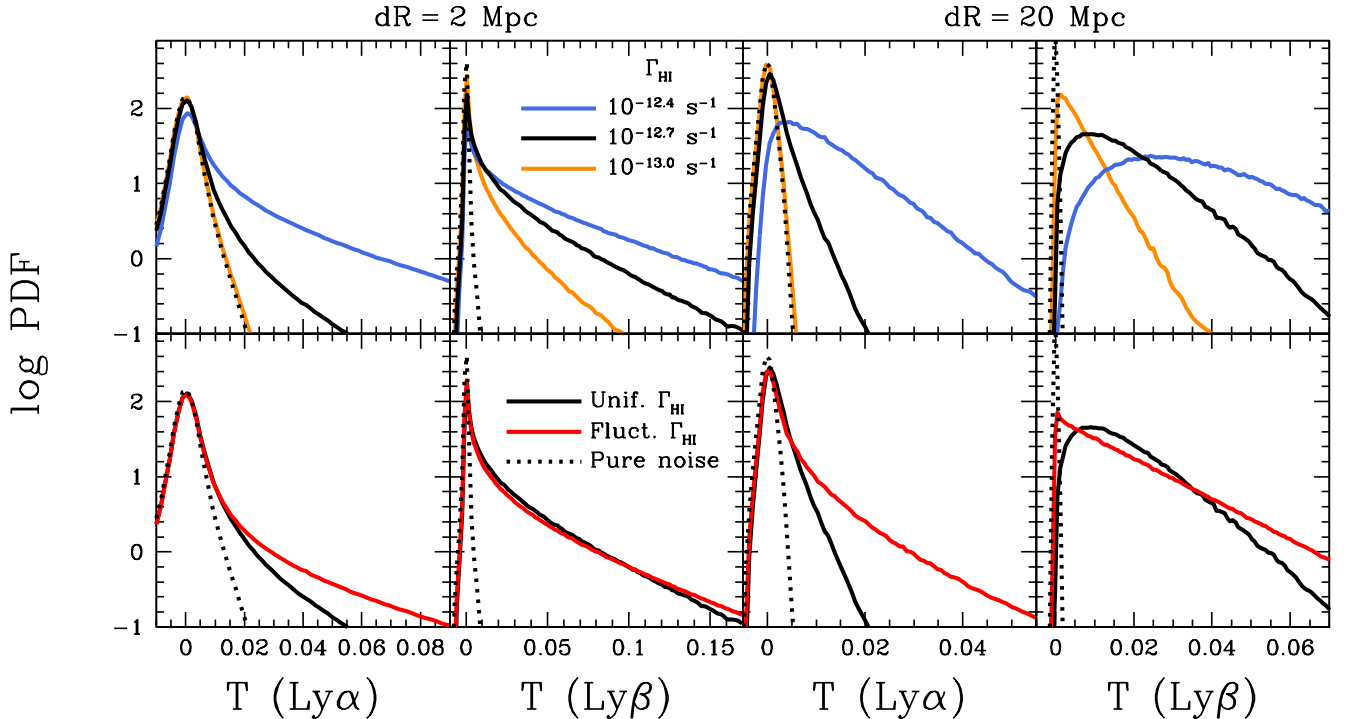
Figure 4. Simulated PDFs of 2 Mpc (left four panels) and 20 Mpc (right four panels) binned segments of the Lyα (left panels inside each dR) and Lyβ (right panels inside each dR) forests at $z \sim 6$. The upper panels show the variation in the PDFs as a function of the (uniform) ionization rate $\Gamma_{HI} = 10^{-13.0}$, $10^{-12.7}$, and $10^{-12.4}$ s^{-1} as the orange, black, and blue curves, respectively. The lower panels show the effect of including fluctuations in the ionizing background (red curve) with an average $\Gamma_{HI} = 10^{-12.7}$ s^{-1}. The dotted curve in every panel shows the PDF of the noise model, which assumes independent Gaussian-distributed noise according to a noise vector from a real Keck/ESI spectrum.

of large bins in the Lyα forest (Fan et al. 2006; Becker et al. 2015) but at a scale smaller by around a factor of three. The distributions are narrower (note the difference in transmission scale on the horizontal axis), reflecting not only a decrease in sky noise (S/N $\propto N_{pix}^{-1/2}$), but also the tendency for several transmission spikes to fall into a single bin leading to a somewhat more Gaussian shape to the distribution (as per the central limit theorem). For moderate ionization rates $\Gamma_{HI} \gtrsim 10^{-12.7}$ s^{-1}, the model PDFs suggest that high signal-to-noise transmission should nearly always be measured in the Lyβ forest at this scale. The effect of ionizing background fluctuations on the PDFs is considerably stronger on large scales, consistent with the $\gtrsim$ 10s of Mpc typical scales of background fluctuations seen in the simulations (e.g. Figure 3). In particular the large-scale Lyβ PDF is strongly affected by ionizing background fluctuations through an excess in 20 Mpc-scale Lyβ-dark regions (GP troughs) over the uniform background model, which are otherwise consistent with the noise floor in the Lyα forest, suggesting that it may have considerable constraining power for the parameters of models for strong fluctuations in IGM opacity (e.g. DF16, D'Aloisio et al. 2015). The 20 Mpc-scale PDFs shown here are analogous to the 50 Mpc/h-scale (cumulative) PDFs studied by Becker et al. (2015) – they similarly allow one to characterize fluctuations in the Lyα and Lyβ forests on the large-scales expected from an inhomogeneous reioniza-

tion process (e.g. Furlanetto et al. 2004).

From these PDFs we should be able to assess the probability for any given model to reproduce rare transmission spikes in the Lyα and Lyβ forests, and from their detailed shape, study the patchy nature of the post-reionization IGM. In practice, performing parameter inference from the PDFs is difficult because much of their constraining power lies in a handful of transmission spikes in the tail of the distribution, and the features in the Lyα and Lyβ forests should be correlated because they trace the same physical structures.

3. STATISTICAL METHODS: APPROXIMATE BAYESIAN COMPUTATION

The standard approach to IGM parameter inference from the transmission PDF of the Lyα forest at $z \sim 2$–3 (e.g. Lidz et al. 2006) is to choose a set of coarse bins in the range of $0 < F < 1$, compute the covariance matrix of the bins from either jackknife sampling of real data (Lidz et al. 2006; Lee et al. 2015) or forward-modeling of the Lyα forest in hydrodynamical simulations (Rollinde et al. 2013), and then approximate the likelihood function as a multivariate Gaussian. Because each PDF bin typically contains thousands of pixels, the central limit theorem ensures that a multivariate Gaussian likelihood is a good approximation.

The situation changes dramatically at higher redshifts where the IGM is more opaque and fewer quasars are available. At $z \gtrsim 6$, most pixels in the Lyα and Lyβ
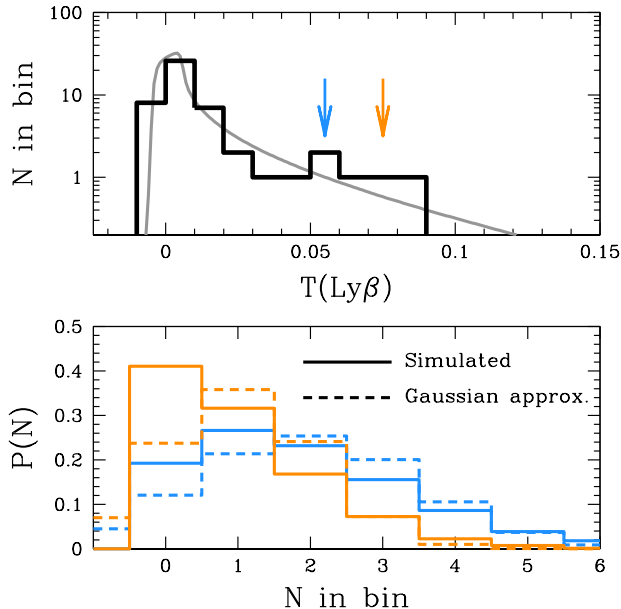
Figure 5. Top: The histogram shows the $\Delta T = 0.01$-binned transmission distribution of a single mock observation of the Lyβ forest covering 100 Mpc with transmission binned on 2 Mpc scales, with $\Gamma_{\rm H\,I} = 10^{-12.7}$ s^{-1} and $\lambda_{\rm mfp} = 15$ Mpc at $z \sim 5.9$–6.1. The curve shows the model PDF convolved with a $\Delta T = 0.01$ boxcar filter to be consistent with the binned PDF. Bottom: The solid and dashed histograms represent the true and Gaussian approximations to the distribution of PDF values, respectively. The two sets of histograms represent the two bins labeled (and color-coded) by the arrows in the top panel.

forests are consistent with zero transmission, with only a *sparse* tail of real transmission spikes. While the (co)variances of each bin can be estimated by forward modeling of mock spectra, the resulting Gaussian distributions (as defined by the mean and (co)variance of the mock samples) are not representative of the distributions of the number of pixels that fall into each bin. We show an example of this in Figure 5, where the bottom panel illustrates the true and Gaussian-approximated distributions of mock Lyβ PDF measurements in transmission bins highlighted in the top panel, with one mock observation shown.

The requirement of approximately Gaussian likelihood is more stringent for transmission measurements on larger bin scales, such as the 50 Mpc/h measurements by Becker et al. (2015), where the small number of measurements – and the tendency for measurements to be reported as limits – have typically led to rough "chi-by-eye" comparisons to simulations (e.g. D'Aloisio et al. 2015; Davies & Furlanetto 2016; D'Aloisio et al. 2016).

To circumvent these challenges, we adopt the statistical approach known as ABC (e.g. Pritchard et al. 1999; see also Cameron & Pettitt 2012; Ishida et al. 2015; Hahn et al. 2016; Jennings & Madigan 2016 for other astronomical applications) which allows for parameter inference without any explicit analytic form for the likelihood function. We start with an abridged statement of Bayes' Theorem,

$$p(\theta|\mathbf{d}) \propto p(\mathbf{d}|\theta)p(\theta), \quad (1)$$

where θ represents the set of model parameters (e.g. $\Gamma_{\rm H\,I}$) and $\mathbf{d}$ represents the data (e.g. the transmission

measured in each spectral pixel), $p(\mathbf{d}|\theta)$ is the likelihood function, $p(\theta)$ is the prior distribution representing our previous knowledge of the model parameters, and $p(\theta|\mathbf{d})$ is the posterior distribution for θ that we hope to obtain. In many cases, the likelihood function can be approximated as a multivariate Gaussian distribution or can be straightforwardly computed numerically, allowing the posterior distribution to be computed through methods such as Markov Chain Monte Carlo (MCMC). In our case, as we have shown above, such an approximation fails to adequately describe the distribution of mock measurements in our models.

One could imagine a brute force computation of the likelihood function by computing mock data sets to determine the probability of exactly reproducing the data, but this is computationally intractable for many problems. The first and most fundamental approximation proposed by ABC is to slightly blur the definition of the likelihood function, instead describing the posterior distribution as (e.g. Marin et al. 2012)

$$p(\theta|\mathbf{d}) \approx p_{\rm ABC}(\theta|\mathbf{d}) \propto \int p(\rho(\mathbf{d},\mathbf{x}) < \epsilon, \mathbf{x}|\theta)p(\theta)d\mathbf{x}, \quad (2)$$

where $\mathbf{x}$ is a forward-modeled mock data set, $\rho(\mathbf{d},\mathbf{x})$ represents a distance measure between the data and mock data, ϵ is a small tolerance threshold, and the integral is performed over the space of forward-modeled mock datasets. That is, to approximate the likelihood function, ABC proposes that coming *close* to the measured data is "good enough," allowing for far easier computation than the brute force method mentioned above. In the process, the correlations between data elements are naturally accounted for through the forward-modeling procedure without any explicit computation of the covariance matrix.

The simplest application of ABC in practice is to draw model parameter samples θ^* from the prior, forward model a mock realization of the data given by those parameters $\mathbf{x}^*$, then measure the distance between the mock data and real data $\rho(\mathbf{d},\mathbf{x}^*)$. If this distance is "small enough", i.e. $\rho(\mathbf{d},\mathbf{x}^*) < \epsilon$, then θ^* is accepted as a sample of the posterior distribution. This procedure is called ABC rejection sampling.

The above methodology may still be computationally intractable if the data represent a highly stochastic process or high-dimensional data set. For example, in our case, producing a mock data set that nearly matches an observed Lyα and/or Lyβ forest spectrum at the pixel-to-pixel level (i.e. with small $\rho(\mathbf{d},\mathbf{x})$) requires nearly matching the peak heights and positions of every narrow transmission feature. In practice this would require matching not only the ionization state of the gas, but also the *entire set of cosmological initial conditions* that led to the structures observed along the sightline. To overcome this computational hurdle, the technique commonly employed in ABC is to use "summary statistics" to process the data (and mock data) into a form which allows them to be compared more easily, typically into a lower-dimensional state. This reformulation can be stated as

$$p(\theta|\mathbf{d}) \approx p(\theta|\mathbf{s}_{\rm obs})$$
$$\approx p_{\rm ABC}(\theta|\mathbf{s}_{\rm obs}) \propto \int p(\rho(\mathbf{s}_{\rm obs},\mathbf{s}) < \epsilon, \mathbf{x}|\theta)p(\theta)d\mathbf{x}, \quad (3)$$

where $\mathbf{s}_{\mathrm{obs}}$ is the summary statistic of the data and $\mathbf{s}$ is the summary statistic of mock data set $\mathbf{x}$. If the summary statistic is "sufficient" then no information is lost, but such summary statistics are only known analytically in very simple cases (e.g., for determining the mean of a normally-distributed quantity, the sample mean is a sufficient statistic). Thus, in most cases, the use of a summary statistic implies loss of information, leading to bias and/or loss of precision in the posterior distribution. In § 3.2, we perform inference tests on multiple summary statistics of the Lyα and Lyβ forests to investigate their potential biases.

While ABC may seem like a "silver bullet" to perform parameter inference with the least number of assumptions about the nature of the likelihood function, its primary weakness is the requirement of simulating an overwhelmingly large number of mock data sets in order to produce a reasonably converged approximation to the posterior PDF. Even with a one-dimensional summary statistic one may have to generate *millions* of mock data sets, limiting its potential application for practical reasons. Nevertheless, ABC allows for inference involving complicated likelihood functions, fully accounting for covariance (depending on the sufficiency of the chosen summary statistic) *implicitly* as they manifest in the assumed model for mock data.

3.1. *Application of ABC to the Lyα and Lyβ forests*

Application of ABC to inference of the Lyα and Lyβ forests is performed as follows:

1. Draw a value of Γ_{HI} from prior distribution, Γ_{HI}^*.

2. Compute mock spectrum of Lyα and Lyβ forest with $\Gamma_{\mathrm{HI}} = \Gamma_{\mathrm{HI}}^*$, including noise consistent with observations.

3. Compute summary statistic of mock and observed spectra ($\mathbf{s}$ and $\mathbf{s}_{\mathrm{obs}}$, respectively).

4. Compute distance between $\mathbf{s}_{\mathrm{obs}}$ and $\mathbf{s}$, $\rho(\mathbf{s}_{\mathrm{obs}}, \mathbf{s})$.

5. If $\rho(\mathbf{s}_{\mathrm{obs}}, \mathbf{s})$ is below the threshold ϵ, keep Γ_{HI}^* as a sample of the *approximate* posterior PDF, $p_{\mathrm{ABC}}(\Gamma_{\mathrm{HI}})$.

6. Repeat steps 1–5 until a predetermined number of posterior PDF samples have been drawn, e.g. 1000.

In practice, the Γ_{HI}^* can be drawn more efficiently, and in later sections we adopt a method based on importance sampling with an iteratively declining threshold ϵ set based on percentiles of the distribution of $\rho(\mathbf{s}_{\mathrm{obs}}, \mathbf{s})$ (Beaumont et al. 2009; see detailed implementations by e.g. Ishida et al. 2015; Hahn et al. 2016; Jennings & Madigan 2016). These procedures are simply methods to increase the efficiency of the basic rejection sampling method, and as such we will not discuss them in further detail.

For a given segment of the Lyα or Lyβ forest, perhaps the most basic summary statistic is the mean transmission: $\langle F \rangle = \sum_i^{N_{\mathrm{pix}}} F_i / N_{\mathrm{pix}}$. The use of this statistic reduces the dimensionality of the data from [# of Lyα forest pixels + # of Lyβ forest pixels] to 2: the mean transmission in Lyα and the mean transmission in Lyβ.
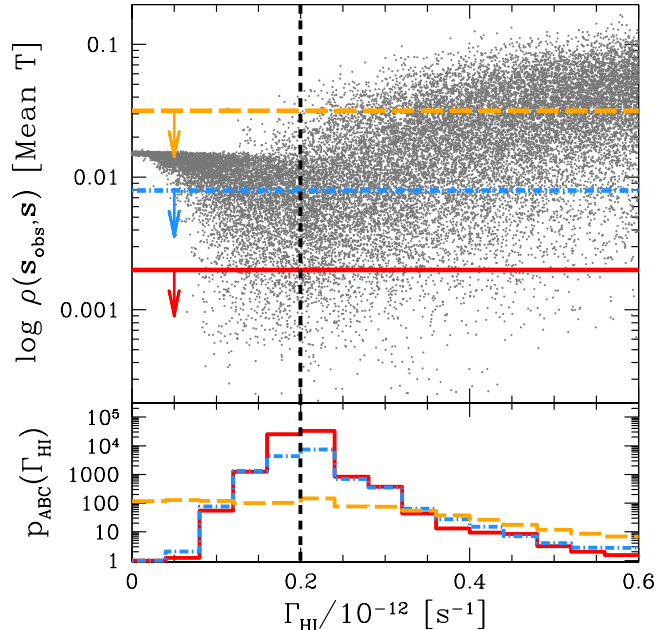

Figure 6. Demonstration of ABC applied to a mock observation of the Lyα and Lyβ forests. The "observation" consists of the mean transmissions of Lyα and Lyβ in a single mock spectrum covering 100 Mpc at $z \sim 6$ with a fluctuating ionizing background and $\Gamma_{\mathrm{HI}} = 10^{-12.7}$ s^{-1}. The points in the upper panel represent the distances between the "observed" mock spectrum and 25000 additional mock spectra with Γ_{HI} drawn from a uniform prior, where the distance has been defined using the mean transmission as the summary statistic (equation 4). The thick dashed orange, dot-dashed blue, and solid red histograms in the lower panel represent the ABC posterior distributions assuming thresholds of $\log \epsilon =$ -1.5, -2.1, and -2.7, respectively, corresponding to the horizontal lines in the upper panel.

That is, $\mathbf{s} = (\langle F_{\mathrm{Ly}\alpha} \rangle, \langle F_{\mathrm{Ly}\beta} \rangle)$. Using ABC, we do not have to make any assumptions about the nature of the intrinsic distributions of Lyα and Lyβ mean transmission, nor the details of their correlations. Instead, the process of comparison to forward modeled mock data naturally accounts for correlations under the assumption that our model for the IGM and the noise model of the spectrum is accurate.

We set up an example of the ABC procedure as follows. One mock data set with known Γ_{HI}, specifically a noisy realization of the Lyα and Lyβ forest with fluctuating Γ_{HI} shown in Figure 3, covering a path length of 100 Mpc, is chosen to be the "observed" sample. We then draw Γ_{HI}^* from a uniform prior covering 0–0.6 $\times 10^{-12}$ s^{-1}, forward model a mock data set using that Γ_{HI}^*, then compute the distance between the observed and mock data sets,

$$\rho(\mathbf{s}_{\mathrm{obs}}, \mathbf{s}) = \sqrt{(\langle F_{\mathrm{obs}}^\alpha \rangle - \langle F_{\mathrm{mock}}^\alpha \rangle)^2 + (\langle F_{\mathrm{obs}}^\beta \rangle - \langle F_{\mathrm{mock}}^\beta \rangle)^2}, \quad (4)$$

where we have adopted the L_2 norm for the distance ρ.

The result of such a test with 25000 mock data sets is shown in Figure 6. Each point in the top panel represents the distance defined above for a single mock spectrum corresponding to a draw of Γ_{HI}^*, a random IGM skewer from the hydrodynamical simulation, and a random fluctuating ionizing background skewer. The bottom panel

shows binned samples from the posterior distribution of $\Gamma_{\rm HI}$ obtained by rejecting samples with distance greater than the ϵ values given by the horizontal lines in the upper panel (i.e. the samples below the lines are accepted). As the tolerance ϵ is decreased, the ABC posterior distribution narrows, representing successive improvements to the approximated likelihood as it converges to the true distribution. Through the ABC process, this posterior distribution explicitly includes cosmic variance (as modeled by the cosmological simulation), measurement uncertainty, and the correlation between the Lyα and Lyβ forest mean transmissions.

However, the mean transmissions in the Lyα and Lyβ forests clearly do not tell the entire story. Due to sky noise, a large region of forest may have a mean transmission consistent with zero but still contain a significant positive spike (e.g. Figure 2 in Becker et al. 2015). The presence of any (real) transmission spike requires an ionizing background greater than zero, but the mean transmission alone may not contain this information because of limited signal-to-noise. A summary statistic which retains this information but disregards the locations of transmission spikes is the *rank-order distribution* of transmission on small scales. That is, a re-ordering of the transmission values such that they are in order from lowest to highest transmission. The scale over which the transmission is binned can be chosen to maximize the signal-to-noise ratio of transmission spikes, i.e. matched filtering. We also account for the spatially-varying noise of the spectrum (e.g. due to sky lines) by first dividing each transmission value by the local noise before re-ordering the measurements. An example of rank-order distributions of two mock Lyα forest spectra with the same $\Gamma_{\rm HI}$ is shown in Figure 7. The distance between observed and mock data is then

$$\rho(\mathbf{s}_{\rm obs}, \mathbf{s}) =$$
$$\sqrt{\sum_i^{N_\alpha}(S_{i,\rm obs}^\alpha - S_{i,\rm mock}^\alpha)^2 + \sum_i^{N_\beta}(S_{i,\rm obs}^\beta - S_{i,\rm mock}^\beta)^2}, \quad (5)$$

where S_i is the signal-to-noise ratio of the transmission in the bin with rank i (where the Lyα and Lyβ forests are ranked independently), and N_α and N_β are the number of bins in the Lyα and Lyβ forests, respectively.

In Figure 8 we show a 2 Mpc-binned mock Lyα and Lyβ forest spectrum at $z \sim 6.5$ with a fluctuating ionizing background and true $\Gamma_{\rm HI} = 10^{-12.95}$ s^{-1}. The mean transmissions in Lyα and Lyβ are consistent with zero, but there is a highly significant ($S/N > 5$) transmission spike in the Lyβ forest. In Figure 9 we show posterior PDFs for $\Gamma_{\rm HI}$ using several different summary statistics. The blue and red curves shows the posterior PDF using the mean transmission and the 20 Mpc-binned rank-order distribution as summary statistics, both of which fail to measure $\Gamma_{\rm HI}$, i.e. the posterior PDF does not rule out $\Gamma_{\rm HI} = 0$. In contrast, using the 2 Mpc-binned rank-order distribution as the summary statistic (black curve) allows for an actual measurement of $\Gamma_{\rm HI}$, demonstrating the enhanced ability of small-scale transmission features to constrain $\Gamma_{\rm HI}$ relative to the large-scale mean.

As an example of the flexibility of the ABC procedure, the dotted curve in Figure 9 shows the posterior PDF from adopting a summary statistic that *com-*
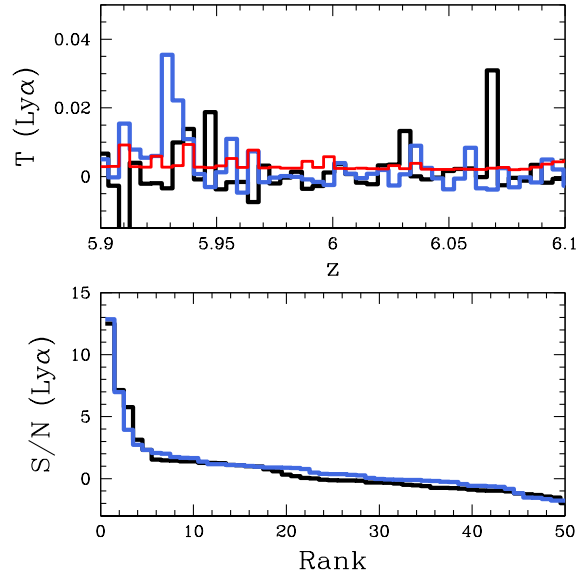

Figure 7. Top: Two mock Lyα forest spectra (black, blue) at $z \sim 6$ computed with the same ionization rate $\Gamma_{\rm HI} = 10^{-12.6}$ s^{-1} and binned to 2 Mpc. The noise vector is shown as the red curve. Bottom: Rank-order distribution of the 2 Mpc binned signal-to-noise of the spectra shown in the upper panel. The "distance" between the two spectra, qualitatively the amount of "space" between the black and blue curves, is much smaller when arranged in rank order.
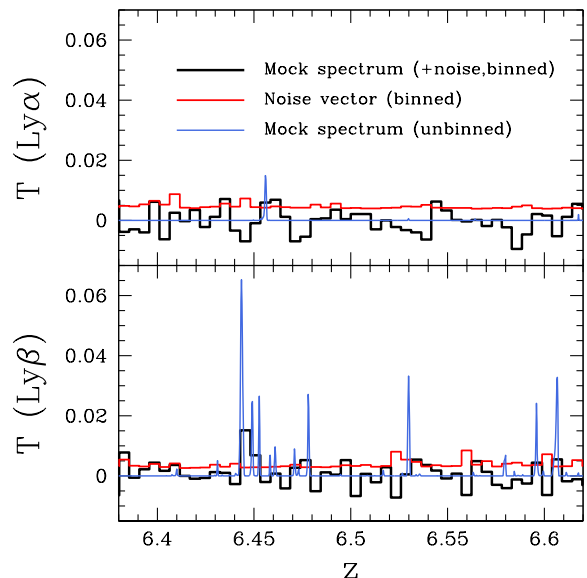

Figure 8. Top panels: Mock Lyα and Lyβ forest spectra (2 Mpc binned: black, noiseless unbinned: blue) and 1σ noise (2 Mpc binned, red) for $\Gamma_{\rm HI} = 10^{-12.95}$ s^{-1} at $z \sim 6.5$. The mean transmissions of Lyα and Lyβ are consistent with zero, but there is a highly significant transmission spike in the Lyβ forest which should definitively rule out $\Gamma_{\rm HI} = 0$.

bines the 2 Mpc and 20 Mpc rank-order distributions by adding their distances in quadrature, i.e. $\rho_{\rm 2Mpc+20Mpc} = \sqrt{\rho_{\rm 2Mpc}^2 + \rho_{\rm 20Mpc}^2}$. In Figure 10, we show the posterior PDFs using the 2 Mpc rank-order distribution summary statistic of the Lyα forest alone (orange) and the Lyβ forest alone (purple) compared to the joint constraints
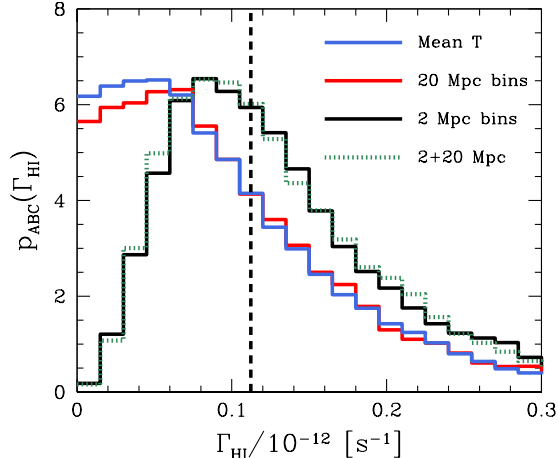
Figure 9. Posterior PDFs of $\Gamma_{\rm H\,I}$ for the Lyα and Lyβ forest spectra in Figure 8 using the mean transmission (blue) versus the rank-order distribution of 2 Mpc and 20 Mpc binned transmission (red and black, respectively) as the ABC summary statistic, compared to the true value (dashed line). While the mean transmission and 20 Mpc bins only provide an upper limit to $\Gamma_{\rm H\,I}$, the S/N ~ 5.5 2 Mpc binned spike in Lyβ is highly constraining. The green dotted curve shows the posterior PDF when the 2 and 20 Mpc rank-order distribution summary statistics are combined.
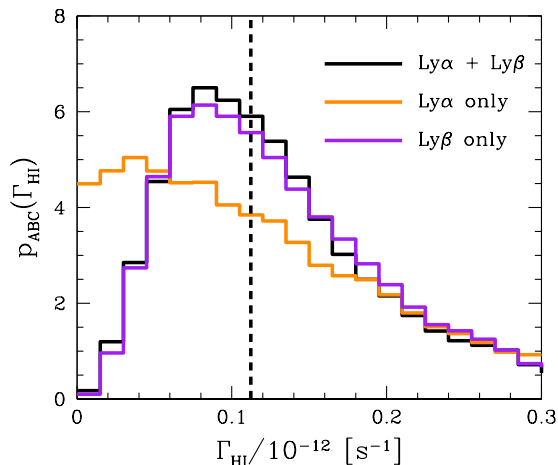

Figure 10. Posterior PDFs of $\Gamma_{\rm H\,I}$ for the Lyα and Lyβ forest spectra in Figure 8 using the rank-order distribution of 2 Mpc binned transmission in Lyα alone (orange), Lyβ alone (purple), and both (black) as the ABC summary statistic. The dashed line indicates the true value of $\Gamma_{\rm H\,I}$.

(black, same as Figure 9). While the Lyβ forest has the most constraining power due to the presence of the significant transmission spike, including the *non-detection* of transmission in Lyα modestly improves the posterior PDF density at the location of the true $\Gamma_{\rm H\,I}$.

As an alternative to the rank-order distribution, we have also devised a one-dimensional summary statistic that is sensitive to rare transmission spikes which we call the "pseudo-likelihood" of the spectrum. Through forward modeling many mock data sets, we can compute the transmission PDF in our simulations to arbitrarily high precision (e.g. Fig. 4) with any set of model parameters, and then compute the "pseudo-likelihood" of the spectrum under the assumption that each bin of the Lyα and Lyβ forests is entirely independent: $p(\mathbf{d}|\theta) = \prod_i P(F_i|\theta)$.

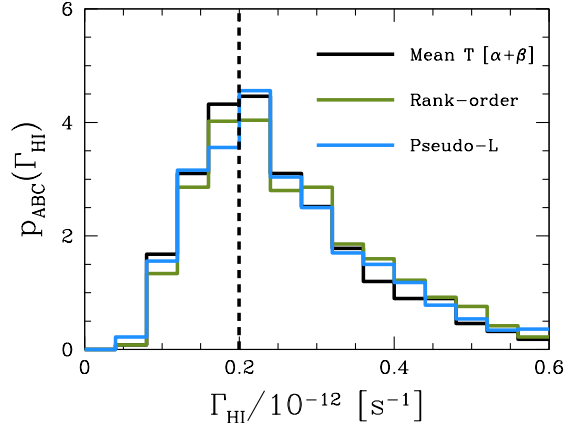**Figure 11.** Comparison of the ABC posterior distributions of $\Gamma_{\rm H\,I}$ for the same mock spectrum as in Figure 6 using three different summary statistics: mean transmission (black), rank-order distribution (green), and pseudo-likelihood (blue). The true $\Gamma_{\rm H\,I}$ is shown as the vertical dashed line.

If there were no spatial correlations in (or between) the Lyα and Lyβ forests, this likelihood would be exact. After finding the maximum pseudo-likelihood model parameters $\theta_{\rm ML}$ of the observed data, we then use the absolute difference of pseudo-likelihood values in log space as the distance metric between the observed and mock data,

$$\rho(\mathbf{s}_{\rm obs}, \mathbf{s}) = \left| \sum_i^{N_{\alpha+\beta}} \ln P(F_{i,\rm obs}|\theta_{\rm ML}) - \sum_i^{N_{\alpha+\beta}} \ln P(F_{i,\rm mock}|\theta_{\rm ML}) \right|, (6)$$

where i encompasses all of the bins in the Lyα and Lyβ forests and $N_{\alpha+\beta} = N_\alpha + N_\beta$. The pseudo-likelihood summary statistic described above is similar to "indirect inference" ABC methodologies (e.g. Drovandi et al. 2011; Gleim & Pigorsch 2013). We find that this summary statistic is equally powerful as the rank-order distribution in terms of its ability to constrain $\Gamma_{\rm H\,I}$ from sparse transmission spikes.

Outside of examples like Figure 8, the mean transmission, rank-order distribution, and pseudo-likelihood summary statistics are comparable in their ability to constrain $\Gamma_{\rm H\,I}$. We demonstrate this in Figure 11 where we show ABC posterior distributions for the three statistics computed for the same mock data as Figure 6, and we have chosen the tolerance ϵ such that we accept 5% of the 25,000 mock datasets. The rank-order distribution and pseudo-likelihood statistics should be more sensitive to additional parameters that affect the shape of the transmission PDF such as the thermal state of the IGM (e.g. Lee et al. 2015) and the strength of ionizing background fluctuations (e.g. Meiksin & White 2003). We restrict ourselves to inferring $\Gamma_{\rm H\,I}$ alone in this work for simplicity in this initial presentation of our methodology, and we note that this choice limits the additional information gathered by the distribution of Lyα and Lyβ forest transmission. That is, in the models we present in this work there is a one-to-one relationship between the mean transmission and $\Gamma_{\rm H\,I}$, and the scale-dependent distribution of transmission does not contain much additional
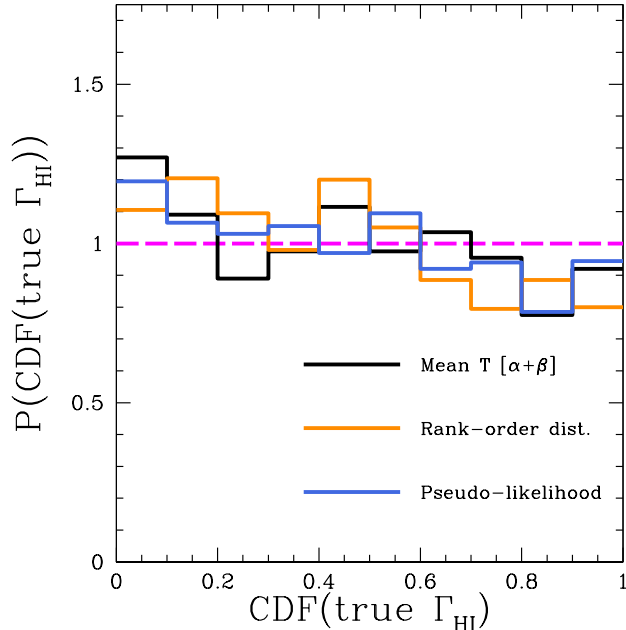
Figure 12. Inference test of the ABC methodology using the different summary statistics described in § 3.1. A total of 2000 mock measurements of Γ_{HI} were performed on 2000 different mock data sets (consisting of one noisy Lyα+β forest spectrum covering 100 Mpc) with $\Gamma_{\mathrm{HI}} = 10^{-12.7}$ s^{-1}. The solid curves show the distribution of the posterior CDF of each mock measurement evaluated at the true Γ_{HI} (i.e. CDF(true Γ_{HI}) as defined in the text) using the mean transmission (black), rank-order PDF (orange), and pseudo-likelihood (blue) as summary statistics.

information *unless* the mean transmission is close to the noise limit of the spectrum. The methodology in principle allows for measurements of other parameters, such as the strength and scale of Γ_{HI} fluctuations and the thermal state of the IGM, which we will pursue in future work.

3.2. *Validation of Statistical Methods: Inference Tests*

We rigorously test the capability of ABC to perform statistical inference of Γ_{HI} as follows. We conduct 2000 mock Γ_{HI} measurements similar to the examples in the previous section – for each measurement, we draw one mock dataset (100 Mpc of the Lyα and Lyβ forests at $z \sim 6$ with $\Gamma_{\mathrm{HI}} = 10^{-12.7}$ s^{-1}, e.g. Figure 3) to be the "data", then draw 500 samples from the ABC approximate posterior distribution of Γ_{HI}. For each of the 2000 mock measurements, we record the value of the cumulative distribution function (CDF), determined from the 500 posterior samples, at the true Γ_{HI}, CDF(true Γ_{HI}). The procedure then results in 2000 values of CDF(true Γ_{HI}).

The simplest inference test would be to test what fraction of the time, out of the 2000 mock measurements, $0.16 < \mathrm{CDF(true\ }\Gamma_{\mathrm{HI}}) < 0.84$ (central 68%), or $0.025 < \mathrm{CDF(true\ }\Gamma_{\mathrm{HI}}) < 0.975$ (central 95%), etc., but we opt for a somewhat stricter test covering the entire CDF. If the ABC posterior distribution can be treated as a true probability distribution, then the probability that the true Γ_{HI} lies within *any* interval of the CDF should be equal to the size of that probability interval, i.e.

$$P(A < \mathrm{CDF(true\ }\Gamma_{\mathrm{HI}}) < B) = B - A. \qquad (7)$$
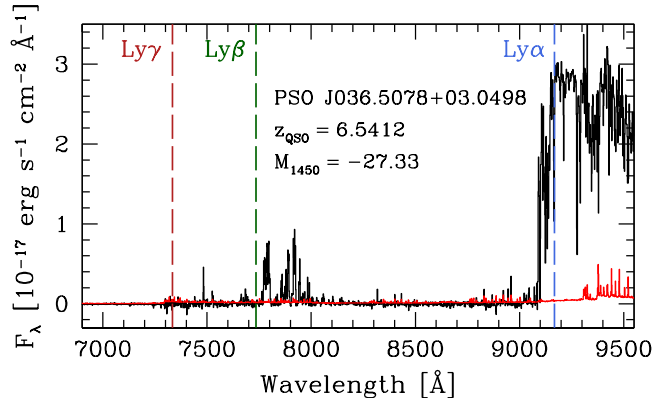

Figure 13. Keck/ESI spectrum of PSO J036.5078+03.0498 ($z_{\mathrm{QSO}} = 6.5412$). The spectrum (black) and noise vector (red) have been binned by six pixels ($\approx$ 1 resolution element). The wavelengths of Lyα, Lyβ, and Lyγ at z_{QSO} are denoted by vertical lines.

In other words, the distribution of CDF(true Γ_{HI}) should follow the uniform distribution over the interval [0,1].

We demonstrate this "CDF(true Γ_{HI})" test for the mean transmission, rank order distribution, and pseudo-likelihood summary statistics in Figure 12 for CDF(true Γ_{HI}) bins of width ΔCDF = 0.1. The dashed line shows the expected distribution of CDF(true Γ_{HI}) (i.e. a uniform distribution), while the solid curves show the distribution of CDF(true Γ_{HI}) of the 2000 mock measurements of Lyα and Lyβ forest spectra described above. In general, the ABC posterior distributions behave similar to the expectations for a true probability distribution, but there is a small tilt such that CDF(true Γ_{HI}) tends to be < 0.5 more often than it should (i.e. P(CDF(true Γ_{HI}) $< 0.5 \sim 0.53$–0.56). This deviation from a uniform distribution is likely due to a combination of factors. First, the summary statistics may be insufficient, leading to bias in the posterior distribution. Second, the posterior PDFs from each test may not be fully converged due to the fact that $\epsilon > 0$, i.e. the tilt demonstrates the inherent "approximate" nature of the ABC posterior. Finally, we drew only 500 posterior samples for each test, and this may not be enough to adequately populate the tails of the posterior PDF. Nevertheless, this rigorous test shows that the probabilities from the ABC posterior distribution can be reasonably interpreted as "true" probabilities, at least in the context of our theoretical model for the IGM.

4. PROOF OF CONCEPT Γ_{HI} MEASUREMENT

As a proof of concept, we have applied our statistical methodology to a Keck/ESI spectrum of the recently discovered quasar PSO J036.5078+03.0498 ($z = 6.5412$, Venemans et al. 2015b; Bañados et al. 2015, henceforth P036+03), shown in Figure 13. The quasar was observed on January 11th and 12th 2016 for a combined exposure time of 11,700s with a slit width of 1", resulting in a signal-to-noise ratio of $\sim$ 20–45 per pixel between sky lines in the extrapolated continuum level in the Lyα and Lyβ forests and a spectral resolution of $R \sim 4000$. In the following analysis, and for the purposes of generating mock spectra, we use the noise vector from this observation instead of the one described in § 2. Further details of the data reduction and continuum fitting of
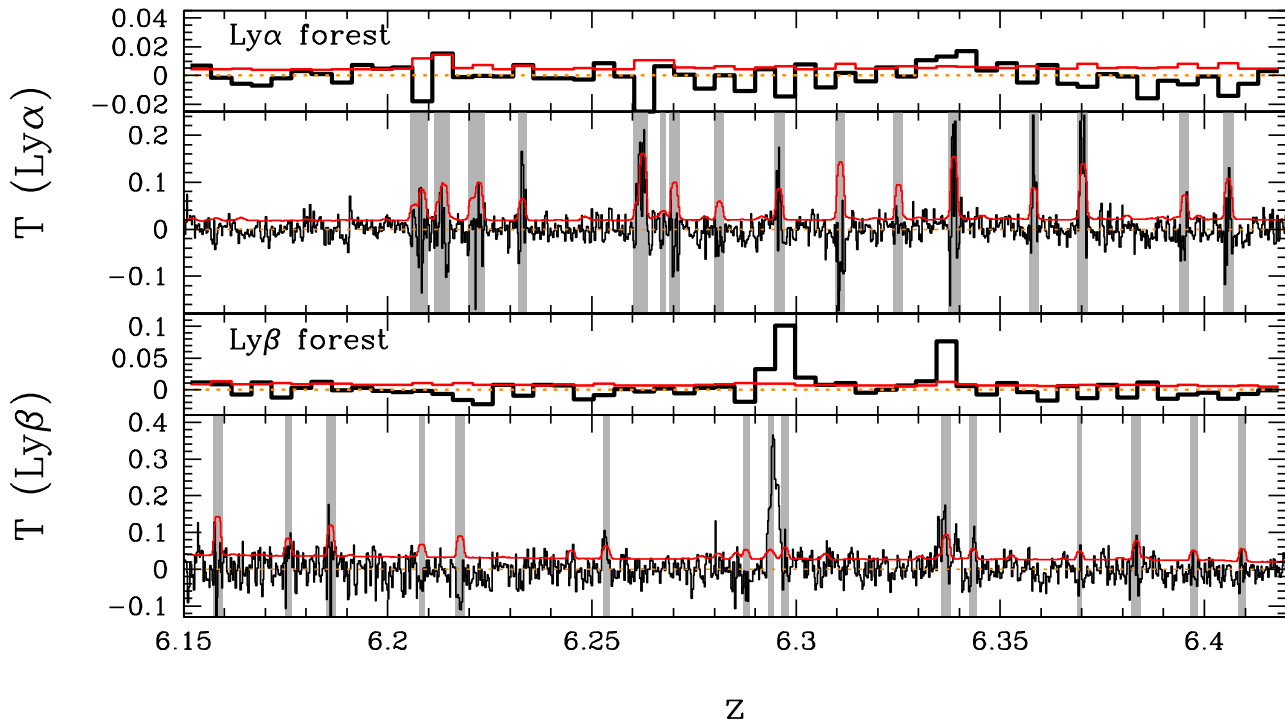
Figure 14. Overlapping Lyα and Lyβ forest in the highest redshift bin studied in our proof-of-concept test on P036+03. The top two panels show the Lyα forest, while the bottom two panels show the Lyβ forest. The lower panel in each pair shows the un-binned spectrum (black) and noise vector (red), with grey shading indicating regions with strong sky lines that we have masked in our current analysis with an automated threshold above the intra-line noise level. The upper panel in each pair shows the 2 Mpc binned spectrum (black) and noise vector (red) only including pixels outside of the masking in the lower panels.

this spectrum will be discussed in detail by Eilers et al. (in prep.). The usable Lyα forest of this quasar spans $z \sim 5.36$–6.42, bounded by the onset of Lyβ absorption at $(1+z) = (1+z_{\mathrm{QSO}})(\lambda_{\mathrm{Ly}\beta}/\lambda_{\mathrm{Ly}\alpha})$ and by the $\Delta z \sim 0.1$ extent of the quasar proximity zone (as defined by the declining visual excess in transmission spikes, cf. Fan et al. 2006). The usable Lyβ forest spans $z \sim 6.15$–6.42, similarly cut off by the onset of Lyγ absorption at $(1+z) = (1+z_{\mathrm{QSO}})(\lambda_{\mathrm{Ly}\gamma}/\lambda_{\mathrm{Ly}\gamma}) = 6.15$ and by the quasar proximity zone at $z = 6.42$.

In Figure 14, we show a zoom-in on the overlap between the Lyα and Lyβ forests at $z \sim 6.3$. Due to the presence of correlated noise in sky lines that we currently ignore in our modeling procedure, we have automatically masked regions which have noise greater than twice the smoothly varying inter-line noise, shown as the grey vertical bands. The Lyα forest is nearly consistent with zero transmission, but when binned to 2 Mpc (~ 20 ESI pixels) there is a region with S/N ~ 3 at $z \approx 6.34$. The Lyβ forest shows considerably more transmission, mostly notably a strong Lyβ spike at $z \approx 6.295$ which is partially masked due to its overlap with sky lines. A weaker Lyβ feature also appears at $z \approx 6.335$ which is similarly contaminated by sky lines, and may be due to the same physical structure (i.e. void in the IGM) as the Lyα feature at the same redshift.

We have divided the Lyα forest into the five redshift bins listed in Table 1. The highest redshift bin covers 108 Mpc corresponding to the entire Lyβ forest range, while the remaining four bins split the forest into 88 Mpc re-

Table 1
Proof-of-concept $\Gamma_{\mathrm{H\,I}}$ measurements for P036+03

z	dR^a	Ly-series	$\log \Gamma_{\mathrm{H\,I}}^{\mathrm{unif}}$ (s^{-1})	$\log \Gamma_{\mathrm{H\,I}}^{\mathrm{fluct}}$ (s^{-1})
6.151–6.418	108	Lyα+Lyβ	$-12.82^{+0.11}_{-0.11}$	$-12.76^{+0.23}_{-0.22}$
5.944–6.150	88	Lyα	$(-12.83)^b$	$(-12.55)^b$
5.746–5.944	88	Lyα	$-12.70^{+0.09}_{-0.10}$	$-12.62^{+0.23}_{-0.20}$
5.556–5.746	88	Lyα	$-12.55^{+0.07}_{-0.07}$	$-12.44^{+0.21}_{-0.19}$
5.374–5.556	88	Lyα	$-12.18^{+0.06}_{-0.07}$	$-12.07^{+0.22}_{-0.19}$

The reported values of $\Gamma_{\mathrm{H\,I}}$ are the medians of the posterior PDFs, while the error bars represent the central 68% credible interval. The posterior PDFs were computed using ABC with the rank-order distribution summary statistic. Note that these uncertainties *do not* include systematic uncertainties due to the assumed thermal state of the intergalactic medium or the continuum fit of the quasar spectrum.
^{a} Size of redshift range in Mpc
^{b} 95% credible upper limits.

gions covering the rest of the Lyα forest. In each redshift bin, we perform separate ABC analyses using the three summary statistics described in § 3.1: the mean transmission, the rank-order distribution of 2 Mpc pixels, and the pseudo-likelihood of 2 Mpc pixels. The same pixels in the mock data sets were masked as in the observed spectra so as to keep the analysis of the mocks similar to what was applied to the real data. We assume a broad uniform prior on $\Gamma_{\mathrm{H\,I}}$ from 0–3×10^{-12} s^{-1} and draw 2500 samples from the (approximate) posterior PDF of $\Gamma_{\mathrm{H\,I}}$.
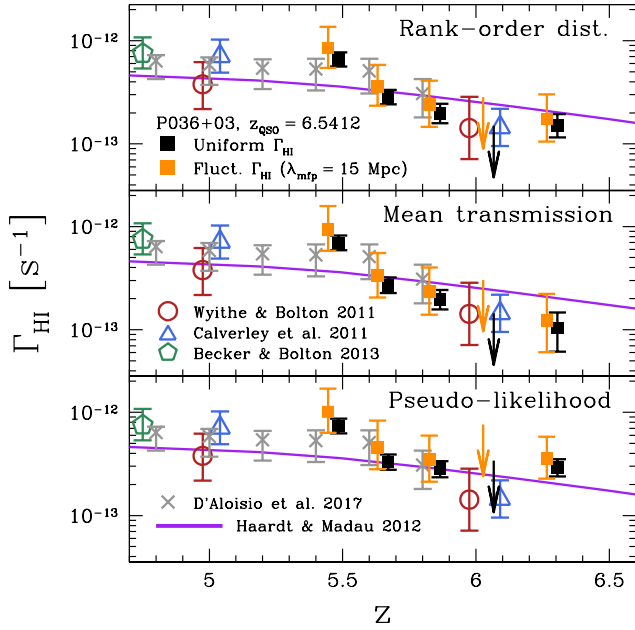
Figure 15. Proof of concept measurements of $\Gamma_{\mathrm{H\,I}}$ from the transmission spectrum of P036+03 shown in Figure 13 divided into the bins listed in Table 1, compared to previous measurements and the commonly assumed Haardt & Madau (2012) model. The black and orange points show the median $\Gamma_{\mathrm{H\,I}}$ of the posterior PDFs in each redshift bin assuming uniform and fluctuating ionizing background ($\lambda_{\mathrm{mfp}} = 15$ Mpc) models, respectively. Error bars denote the 16th-84th percentile range of the posterior PDFs. Where a measurement is consistent with zero, we show a downward arrow at the 95th percentile of the posterior PDF. The top, middle, and bottom panels apply the rank-order distribution, mean transmission, and pseudo-likelihood summary statistics, respectively.
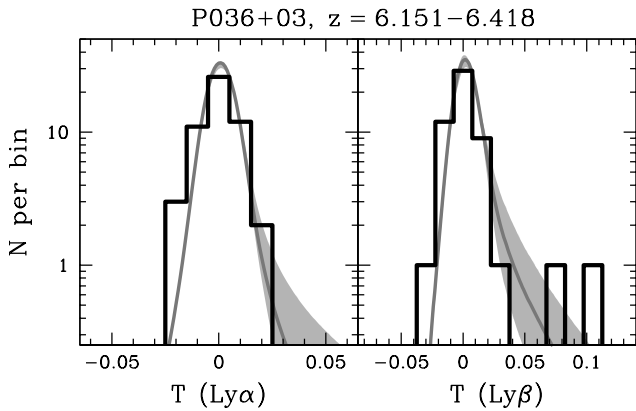

Figure 16. Transmission PDFs of the 2 Mpc-binned Lyα (left) and Lyβ (right) forest of P036+03 at $z = 6.151$–6.420. The black histograms show the measured distribution of transmission for P036+03, while the grey curve and shaded regions show (fluctuating background) model PDFs corresponding to the median and 68% credible $\Gamma_{\mathrm{H\,I}}$ values, respectively, as measured by the rank-order distribution summary statistic assuming a fluctuating ionizing background (Table 1).

Figure 15 shows the resulting $\Gamma_{\mathrm{H\,I}}$ measurements and uncertainties compared to previous measurements. We show two limiting cases of the ionizing radiation field – a uniform background (black points), and a highly fluctuating background as suggested by DF16 at $z \gtrsim 5.6$ (orange points). The fluctuating background models assume
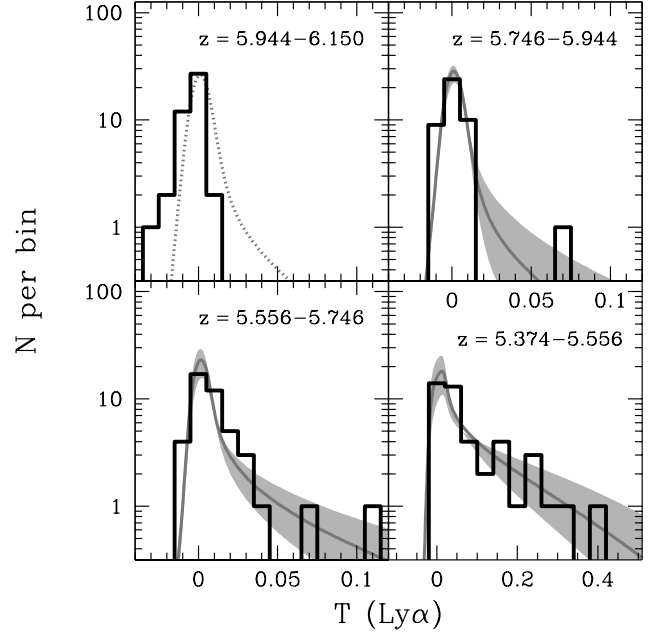

Figure 17. Transmission PDFs of the 2 Mpc-binned Lyα forest of P036+03 in redshift bins below the Lyβ forest overlap with corresponding (fluctuating background) model PDFs similar to Figure 16. The dotted curve in the upper left panel reflects the model PDF corresponding to the 95% credible upper limit for $\Gamma_{\mathrm{H\,I}}$ given the lack of significantly detected transmission. Note the extended horizontal axis scale on the lower right panel.

the $\Gamma_{\mathrm{H\,I}}/\langle\Gamma_{\mathrm{H\,I}}\rangle$ field computed as described in § 2 with $\lambda_{\mathrm{mfp}} = 15$ Mpc (e.g. bottom panel of Figure 3, drawn from the simulation shown in Figure 2), and the constraints in Figure 15 apply to $\langle\Gamma_{\mathrm{H\,I}}\rangle$. In general the measurements are consistent with previous work at $z \sim 5$–6 (large open points; Wyithe & Bolton 2011; Calverley et al. 2011; Becker & Bolton 2013; D'Aloisio et al. 2016). These works measured $\Gamma_{\mathrm{H\,I}}$ two different ways, either by calibrating hydrodynamical simulations to measurements of the mean transmission (Wyithe & Bolton 2011; Becker & Bolton 2013; D'Aloisio et al. 2016) or by analyzing the transmission profile of quasar proximity zones (Calverley et al. 2011). For those measurements based on the mean transmission we have adjusted their $\Gamma_{\mathrm{H\,I}}$ values slightly to account for their different assumed cosmologies following the scaling factors in Bolton & Haehnelt (2007). Each of these measurements represents samples of ≥ 7 quasars, while our comparably precise measurements are from a single quasar, and we more explicitly model large-scale cosmic variance of the density field and a strongly fluctuating ionizing background. However, our stated uncertainties do not include significant systematic uncertainties considered by previous authors, which we discuss further below.

In Figure 16 we show the measured PDFs of 2 Mpc transmission ($dT = 0.01$-binned) in the Lyα and Lyβ forests at $z = 6.151$–6.420 compared to fluctuating background model PDFs corresponding to the median and 68% credible posterior values of $\Gamma_{\mathrm{H\,I}}$ measured by the rank-order distribution (Table 1). The constraints demonstrate the capability of the ABC procedure to constrain $\Gamma_{\mathrm{H\,I}}$ from a sparsely sampled tail of significant transmission measurements. In Figure 17 we show simi-

lar plots for the remaining lower redshift bins which are constrained by the Lyα forest alone. The upper left panel shows the bin of $z = 5.943$–6.151 over which there are no significant transmission spikes detected, with the model PDF corresponding to the 95% credible upper limit on $\Gamma_{\rm HI}$ shown by the dotted curve. In this redshift range, corresponding to observed wavelengths 8440–8693Å, the excess of pixels with relatively large negative transmission values compared to the model PDF suggests that there are systematics in the data reduction that are not present in our simple Gaussian noise model. The other redshift bins show good agreement between the shape of the model PDFs and the data.

The constraints from different summary statistics are nearly identical, although the pseudo-likelihood statistic leads to somewhat elevated $\Gamma_{\rm HI}$ at all redshifts, most prominently in the highest redshift bin. This is likely due to our simplistic modeling of the noise in the spectrum as uncorrelated and Gaussian-distributed. In future work, we will use more realistic noise realizations extracted from real data. Both the pseudo-likelihood and mean transmission statistics are adversely affected by regions with an excess of negative transmission – which exist in the data but not in our noise model – leading to positive and negative biases in the $\Gamma_{\rm HI}$ measurements, respectively.

In detail, the pseudo-likelihood statistic cannot distinguish between negative and positive outliers because the negative side of the transmission PDF is only weakly dependent on $\Gamma_{\rm HI}$ (e.g. Figure 4), so excess negative transmission (wrongly) influences the posterior PDF the same way as a positive transmission spike, e.g. towards higher $\Gamma_{\rm HI}$. Conversely, excess negative transmission decreases the mean transmission more than otherwise expected by our assumed noise variations, so the inferred $\Gamma_{\rm HI}$ can be biased low. Despite the presence of small-scale transmission in the Lyα and Lyβ forests in the highest redshift bin (Figure 14), the mean transmission in Lyα is consistent with zero, and the mean transmission in Lyβ is only barely detected with S/N ≈ 2.

Note that, for simplicity, we have not included uncertainties in the thermal state of the IGM nor uncertainties in the quasar continuum. In the case of a uniform background, the former uncertainty dominates the error budget of $\Gamma_{\rm HI}$ (Becker & Bolton 2013). Preliminary tests marginalizing over a wide range in potential IGM thermal states and $\pm 20\%$ continuum uncertainty[2] suggest that the true error budget should be inflated by up to ~ 0.2 dex, but we leave a detailed multi-parameter analysis to future work. We have also ignored uncertainties in the mean transmission at $z \sim 5$ which will contribute uncertainty to the Lyβ measurement – however, the uncertainty in the mean transmission is likely to be smaller than the intrinsic scatter between sightlines due to the IGM density field, which we explicitly model. Thus the uncertainties shown in Figure 15 and

---

[2] Due to the weak transmission at $z \gtrsim 5.5$, uncertainties in the quasar continuum have a relatively small effect on the inferred ionizing background (see Eilers et al. 2017). In the simplest case of inference via the mean transmission, $\Gamma_{\rm HI} \propto 1/\tau_{\rm eff}^{\sim 2.0}$ (Becker & Bolton 2013). Propagating uncertainties through $\tau_{\rm eff} = -\ln \langle F \rangle$, a continuum error of 20% implies $\sigma_{\tau_{\rm eff}} \sim 0.2$, which then corresponds to $\sigma_{\Gamma_{\rm HI}}/\Gamma_{\rm HI} \sim 10\%$ for $\tau_{\rm eff} \sim 4$.
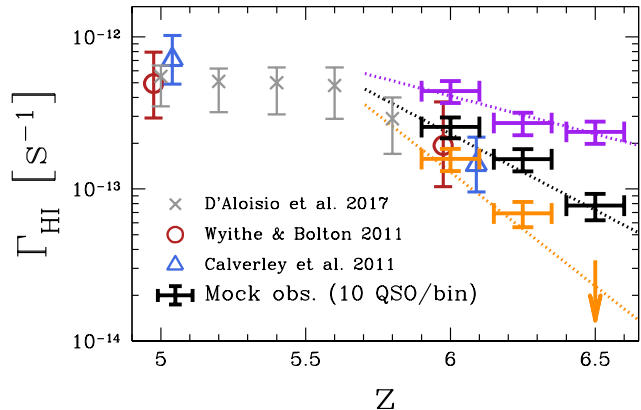

Figure 18. Purple, black, and orange error bars demonstrate the median $\Gamma_{\rm HI}$ of the posterior PDF from measurements performed on mock data sets (10 quasars per bin, using the overlapping Lyα and Lyβ forest region) following the evolving $\Gamma_{\rm HI}$ models shown by the corresponding dotted curves. The downward arrow in the highest redshift bin of the lowest $\Gamma_{\rm HI}$ model represents a 2σ upper limit from the lack of significant transmission spikes in the mock data.

listed in Table 1 should be understood as lower limits based on statistical uncertainty alone. Given our ability to measure $\Gamma_{\rm HI}$ from individual sightlines, in principle we can determine how realistic our uncertainties are by comparing our predicted dispersion (from the posterior PDFs) to the actual dispersion in measurements on multiple sightlines. In future work we will directly compare our estimated uncertainties in the mean $\Gamma_{\rm HI}$ to the dispersion in actual measurements from a large ensemble of quasars.

5. DISCUSSION AND CONCLUSION

The $\sim$Mpc-scale distribution of transmission in the Lyα and Lyβ forests at $z \gtrsim 6$ contains considerable statistical constraining power for the ionization state of the post-reionization IGM, but extracting this information from observed spectra is challenging. We have shown that an application of ABC should allow for precise measurements of the ionizing background at high redshift, even from ~ 1 significant transmission spike in a broad spectral region.

We have demonstrated above that our method, applied to a single high-quality quasar spectrum, allows for a measurement of $\Gamma_{\rm HI}$ to a precision of roughly a factor of two at $z > 6$ in the presence of a highly fluctuating ionizing background. This uncertainty. is dominated by cosmic variance, primarily due to the fluctuating ionizing background with a smaller contribution from density fluctuations In a fixed redshift interval, the constraining power of the Lyα and Lyβ forests increases roughly as $\sqrt{N_{\rm QSO}}$, so larger samples will not only be able to test the consistency of our uncertainty estimates but also allow for unprecedented precision in $\Gamma_{\rm HI}$ measurements. The population of known $z \gtrsim 6$ quasars has greatly increased in the last few years from all-sky surveys such as SDSS (e.g. Jiang et al. 2015) and Pan-STARRS (Venemans et al. 2015b; Bañados et al. 2016, Mazzucchelli et al., in prep.), and other ongoing searches for $z \gtrsim 6$ quasars (e.g. Reed et al. 2015, 2017; Venemans et al. 2015a) will likely increase the potential spectroscopic

sample size even further.

In Figure 18, we show mock measurements of $\Gamma_{\mathrm{H\,I}}$ for a hypothetical survey of three samples of ten quasars each with $z_{\mathrm{Q}} = 6.25, 6.50, 6.75$ to measure $\Gamma_{\mathrm{H\,I}}$ at $z_{\mathrm{IGM}} = 6.00, 6.25, 6.50$, respectively, using the rank-order distribution summary statistic on 100 Mpc of the overlapping Lyα and Lyβ forest regions. The three colors/dotted curves correspond to different toy evolution models that roughly bracket published measurements of the Lyα forest effective optical depth at $z \gtrsim 6$ (Figure 1) and we assume a strongly fluctuating ionizing background ($\lambda_{\mathrm{mfp}} = 15$ Mpc). While we do not include them in Figure 18, at redshifts below where the Lyβ forest is cut off by the onset of Lyγ absorption the constraints from the Lyα forest alone (down to $z \sim 5.25$ for quasars at $z_{\mathrm{QSO}} = 6$) would be of comparable precision.

Currently, there are ~ 20 $z \gtrsim 6$ QSOs with moderate-resolution (e.g. Keck/ESI, VLT/X-Shooter) spectra of comparable quality (or better) to our spectrum of P036+03 in public data archives. In future work, we will perform a joint statistical analysis of all publicly available spectra of $z \gtrsim 6$ quasars to measure $\Gamma_{\mathrm{H\,I}}$ from $z \sim 5.5$ to $z \gtrsim 6.5$ and constrain IGM thermal parameters. The lower panels of Figure 4 show that background fluctuations can imprint substantial signals in the distribution of transmission in the Lyα and Lyβ forests on small and large scales. With a large number of independent sightlines through the IGM, it may be possible to constrain not only the mean strength of the ionizing background, but also the strength and scale of ionizing background fluctuations through a multi-scale approach.

REFERENCES

Almgren, A. S., Bell, J. B., Lijewski, M. J., Lukić, Z., & Van Andel, E. 2013, ApJ, 765, 39
Bañados, E., Decarli, R., Walter, F., et al. 2015, ApJ, 805, L8
Bañados, E., Venemans, B. P., Decarli, R., et al. 2016, ArXiv e-prints, arXiv:1608.03279
Barnett, R., Warren, S. J., Becker, G. D., et al. 2017, ArXiv e-prints, arXiv:1702.03687
Beaumont, M. A., Cornuet, J.-M., Marin, J.-M., & Robert, C. P. 2009, Biometrika, doi:10.1093/biomet/asp052
Becker, G. D., & Bolton, J. S. 2013, MNRAS, 436, 1023
Becker, G. D., Bolton, J. S., Madau, P., et al. 2015, MNRAS, 447, 3402
Becker, G. D., Rauch, M., & Sargent, W. L. W. 2007, ApJ, 662, 72
Becker, R. H., Fan, X., White, R. L., et al. 2001, AJ, 122, 2850
Bolton, J. S., & Becker, G. D. 2009, MNRAS, 398, L26
Bolton, J. S., & Haehnelt, M. G. 2007, MNRAS, 382, 325
Calverley, A. P., Becker, G. D., Haehnelt, M. G., & Bolton, J. S. 2011, MNRAS, 412, 2543

Cameron, E., & Pettitt, A. N. 2012, MNRAS, 425, 44
Chornock, R., Berger, E., Fox, D. B., et al. 2014, ArXiv e-prints, arXiv:1405.7400
—. 2013, ApJ, 774, 26
D'Aloisio, A., McQuinn, M., Davies, F. B., & Furlanetto, S. R. 2016, ArXiv e-prints, arXiv:1611.02711
D'Aloisio, A., McQuinn, M., & Trac, H. 2015, ApJ, 813, L38
Davies, F. B., & Furlanetto, S. R. 2016, MNRAS, 460, 1328
Drovandi, C. C., Pettitt, A. N., & Faddy, M. J. 2011, Journal of the Royal Statistical Society: Series C (Applied Statistics), 60, 317
Eilers, A.-C., Hennawi, J. F., & Lee, K.-G. 2017, ApJ submitted
Fan, X., Narayanan, V. K., Strauss, M. A., et al. 2002, AJ, 123, 1247
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
Field, G. B. 1959, ApJ, 129, 525
Furlanetto, S. R., Zaldarriaga, M., & Hernquist, L. 2004, ApJ, 613, 1
Gleim, A., & Pigorsch, C. 2013, Draft paper: http://ect-pigorsch.mee. uni-bonn. de/data/research/papers
Goto, T., Utsumi, Y., Hattori, T., Miyazaki, S., & Yamauchi, C. 2011, MNRAS, 415, L1
Gunn, J. E., & Peterson, B. A. 1965, ApJ, 142, 1633
Hahn, C., Vakili, M., Walsh, K., et al. 2016, ArXiv e-prints, arXiv:1607.01782
Ishida, E. E. O., Vitenti, S. D. P., Penna-Lima, M., et al. 2015, Astronomy and Computing, 13, 1
Jennings, E., & Madigan, M. 2016, ArXiv e-prints, arXiv:1608.07606
Jiang, L., McGreer, I. D., Fan, X., et al. 2015, AJ, 149, 188
Lee, K.-G., Hennawi, J. F., Spergel, D. N., et al. 2015, ApJ, 799, 196
Lidz, A., Heitmann, K., Hui, L., et al. 2006, ApJ, 638, 27
Loeb, A., & Furlanetto, S. R. 2013, The First Galaxies in the Universe
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
Marin, J.-M., Pudlo, P., Robert, C. P., & Ryder, R. J. 2012, Statistics and Computing, 22, 1167
McGreer, I. D., Mesinger, A., & D'Odorico, V. 2015, MNRAS, 447, 499
McGreer, I. D., Mesinger, A., & Fan, X. 2011, MNRAS, 415, 3237
Meiksin, A., & White, M. 2003, MNRAS, 342, 1205
Mesinger, A., & Furlanetto, S. 2007, ApJ, 669, 663
Mortlock, D. J., Patel, M., Warren, S. J., et al. 2009, A&A, 505, 97
Oh, S. P., & Furlanetto, S. R. 2005, ApJ, 620, L9
Planck Collaboration, Aghanim, N., Ashdown, M., et al. 2016a, ArXiv e-prints, arXiv:1605.02985
Planck Collaboration, Adam, R., Aghanim, N., et al. 2016b, ArXiv e-prints, arXiv:1605.03507
Pritchard, J. K., Seielstad, M. T., Perez-Lezaun, A., & Feldman, M. W. 1999, Molecular Biology and Evolution, 16, 1791
Reed, S. L., McMahon, R. G., Banerji, M., et al. 2015, MNRAS, 454, 3952
Reed, S. L., McMahon, R. G., Martini, P., et al. 2017, ArXiv e-prints, arXiv:1701.04852
Rollinde, E., Theuns, T., Schaye, J., Pâris, I., & Petitjean, P. 2013, MNRAS, 428, 540
Sheinis, A. I., Bolte, M., Epps, H. W., et al. 2002, PASP, 114, 851
Songaila, A. 2004, AJ, 127, 2598
Tang, J.-J., Goto, T., Ohyama, Y., et al. 2016, ArXiv e-prints, arXiv:1612.06148
Venemans, B. P., Verdoes Kleijn, G. A., Mwebaze, J., et al. 2015a, MNRAS, 453, 2259
Venemans, B. P., Bañados, E., Decarli, R., et al. 2015b, ApJ, 801, L11
White, R. L., Becker, R. H., Fan, X., & Strauss, M. A. 2003, AJ, 126, 1
Willott, C. J., Delorme, P., Omont, A., et al. 2007, AJ, 134, 2435

Worseck, G., Prochaska, J. X., O'Meara, J. M., et al. 2014, MNRAS, 445, 1745

Wyithe, J. S. B., & Bolton, J. S. 2011, MNRAS, 412, 1926