LINE-EMITTING GALAXIES BEYOND A REDSHIFT OF 7: AN IMPROVED METHOD FOR ESTIMATING THE EVOLVING NEUTRALITY OF THE INTERGALACTIC MEDIUM

Matthew A. Schenker[1], Richard S. Ellis[1], Nick P. Konidaris[1], and Daniel P. Stark[2]

[1] Department of Astrophysics, California Institute of Technology, MC 249-17, Pasadena, CA 91125, USA; schenker@astro.caltech.edu
[2] Department of Astronomy and Steward Observatory, University of Arizona, Tucson, AZ 85721, USA

ABSTRACT

The redshift-dependent fraction of color-selected galaxies revealing Lyman alpha (Lyα) emission, $x_{\mathrm{Ly}\alpha}$ has become the most valuable constraint on the evolving neutrality of the early intergalactic medium. However, in addition to resonant scattering by neutral gas, the visibility of Lyα is also dependent on the intrinsic properties of the host galaxy, including its stellar population, dust content, and the nature of outflowing gas. Taking advantage of significant progress we have made in determining the line-emitting properties of $z \simeq$ 4–6 galaxies, we propose an improved method, based on using the measured slopes of the rest-frame ultraviolet continua of galaxies, to interpret the growing body of near-infrared spectra of $z > 7$ galaxies in order to take into account these host galaxy dependencies. In a first application of our new method, we demonstrate its potential via a new spectroscopic survey of $7 < z < 8$ galaxies undertaken with the Keck MOSFIRE spectrograph. Together with earlier published data, our data provide improved estimates of the evolving visibility of Lyα, particularly at redshift $z \simeq 8$. As a by-product, we also present a promising new line-emitting galaxy candidate, detected at 4.0σ at redshift $z = 7.62$. We discuss the improving constraints on the evolving neutral fraction over $6 < z < 8$ and the implications for cosmic reionization.

Key words: dark ages, reionization, first stars – galaxies: high-redshift

Online-only material: color figures

1. INTRODUCTION

The transition from a neutral intergalactic medium (IGM) to one that is ionized, and therefore transparent to ultraviolet photons, represents the latest frontier in our overall understanding of cosmic history. In addition to determining when this cosmic reionization occurred, a key question is the role of star-forming galaxies in governing the process. Structure in the polarization of the cosmic microwave background suggests that the reionization process occurred within the redshift interval $6 < z < 20$ (Hinshaw et al. 2013), and deep infrared imaging with *Hubble Space Telescope (HST)* has provided the first opportunity to conduct a census of galaxies during the latter half of this period (Ellis et al. 2013; Oesch et al. 2014). Recent progress in this area has been reviewed by Robertson et al. (2013) and Bromm (2013).

In the absence of significant numbers of high-redshift QSOs or gamma-ray bursts, the most immediately available probe of the evolving neutrality of the IGM beyond $z \simeq 6$–7 is the visibility of the Lyman alpha (Lyα) emission line in controlled samples of color-selected galaxies. Although a prominent line in star-forming galaxies at $z \leqslant 6$, as Lyα is a resonant transition, it is readily suppressed by neutral gas, both in the host galaxy and, if present, in the surrounding IGM. First proposed as a practical experiment using Lyman break galaxies (LBGs) by Stark et al. (2010), the idea followed earlier theoretical work by Miralda-Escudé et al. (2000), Santos (2004), and others.

Ground-based near-infrared spectroscopic surveys have now targeted various samples of color-selected LBGs over $6 < z < 8$, allowing the construction of a redshift-dependent Lyα fraction, $x_{\mathrm{Ly}\alpha,25}$, calculated as the fraction of LBGs that display Lyα with an equivalent width (EW) in emission larger than 25 Å. This fraction falls sharply from a value of $\simeq 50\%$ at $z \simeq 6$ (Stark et al. 2010) to less than 10% at and beyond $z \simeq 7$ (Pentericci et al. 2011, 2014; Schenker et al. 2012; Ono et al.

2012; Treu et al. 2012). Although converting this downturn in the visibility of the line into the volume fraction of neutral hydrogen, $x_{\mathrm{H\,I}}$, is uncertain (Bolton & Haehnelt 2013), the prospects for improving the statistics of this test are promising given the arrival of multi-object instruments such as MOSFIRE on the Keck I telescope (McLean et al. 2012).

So far, this important measure of late reionization has been applied by adopting an empirical description of the demographics of Lyα emission in LBGs, parameterized according to the EW distribution for various UV luminosities over the redshift range $3 < z < 6$ when the universe is fully ionized. The trend is then extrapolated to higher redshifts in the form of a no evolution prediction, with the aim of rejecting this prediction at some level of significance (e.g., Schenker et al. 2012). As we discuss here, this method, now widely used, has several disadvantages. Recognizing these and noting that the spectroscopic and optical and near-infrared imaging data of LBGs over $3 < z < 6$ have improved in scope and quality, in this paper we adopt a more physically based approach to the visibility of Lyα in the vicinity of the host galaxy. Our new approach aims to predict its visibility in a high-redshift galaxy on the basis of its measured UV continuum slope, which, in turn, contains information on the dust content and stellar population, which both directly influence the strength of any Lyα emission. This approach has the distinct advantage that, for the new $z > 7$ samples being studied with MOSFIRE and other spectrographs, composite UV slopes for the population are usually available so that unnecessary extrapolation can be avoided.

The present paper is concerned with describing this improved Lyα fraction test and applying it to the first comprehensive set of spectroscopic data emerging from MOSFIRE. In addition to incorporating the earlier surveys conducted with Keck (Schenker et al. 2012; Ono et al. 2012; Treu et al. 2012) and FORS2 on the Very Large Telescope (VLT; Pentericci et al. 2011), we present the first results from a survey of high-quality Ultra Deep Field

(UDF) targets, which provides a valuable extension of the aforementioned studies. As part of this survey, we demonstrate a new Lyα-emitting galaxy at a redshift of $z = 7.62$, extending once again the frontier of spectroscopically confirmed *HST* sources.

A plan of the paper follows. In Section 2 we introduce our new method for the Lyα fraction test. Section 3 introduces the new compilation of $3 < z < 6$ data drawn from our now-completed Keck spectroscopic survey (D. P. Stark et al. 2014, in preparation), and Section 4 contains an analysis of these data in the context of our new method. In Section 5 we introduce our new MOSFIRE data and apply our new method to both these data and those obtained earlier.

2. Lyα FRACTION TEST—A NEW APPROACH

Although the traditional Lyα fraction test as first proposed by Stark et al. (2010) has already provided meaningful constraints on the evolution of the IGM beyond $z \sim 6.5$, there are two limitations in the current methodology. Firstly, as inferring the presence of neutral gas in the IGM represents a differential measurement, it is necessary to assume a form of the distribution of the EWs of Lyα emission unprocessed by the IGM for the galaxies at $z \geqslant 7$. Comparing this to the observed distribution allows the extinction imposed by the IGM and, through the application of theoretical models, the IGM neutrality to be derived. The current methodology splits the sample into UV-luminous and UV-faint bins and tracks the Lyα fraction in each bin as a function of redshift. It is necessary to assume that the intrinsic galaxy $EW_{Ly\alpha}$ distribution at $z \sim 7$ is either the same as observed at $z \sim 6$, while the universe is fully ionized, or a linear extrapolation of the $3 < z < 6$ data. However, as we cannot observe the sample at $z \geqslant 7$ in the absence of a neutral IGM, we can never know which, if either, of these assumptions is correct.

Secondly, the EW distribution used to predict the observable Lyα distribution has been characterized in many different ways, including an exponential (Dijkstra et al. 2011), a direct histogram (Schenker et al. 2012), a half Gaussian (Treu et al. 2012), and a half Gaussian with a constant probability tail (Pentericci et al. 2011). Though the distributions are largely similar, no detailed comparison has been performed to determine which one optimally represents the $3 < z < 6$ data. We perform this in the context of assembling our model in Section 4.

The most fundamental question, however, is whether the rest-frame UV luminosity is the optimum parameter to predict the visibility of Lyα in the absence of any suppression by a partly neutral IGM. The approach, based on correlations first noted by Shapley et al. (2003), was adopted by Stark et al. (2010) as M_{UV} can be readily determined from the available photometry of distant galaxies together with a photometric redshift. However, M_{UV} is likely to be a coarse predictor of the Lyα EW as it ignores second-order parameters such as metallicity, the stellar initial mass function, and dust content.

The UV continuum slope is a more natural choice as a basic variable as it encodes each of these physical properties (Meurer et al. 1999). Lower-metallicity and hotter stars produce more ionizing photons per unit UV continuum, thus driving $EW_{Ly\alpha}$ upward. Dust very efficiently absorbs Lyα photons given their large effective path lengths from the many scatterings required to escape an H II region. These changes also result in a bluer or redder UV continuum slope, respectively. Thus, as the UV slope reflects more of the parameters that likely govern $EW_{Ly\alpha}$ compared to M_{UV}, we should expect it to be a more robust predictor of the visibility of the line in high-redshift samples.

Until recently, determining the UV continuum slope was only possible for a restricted subset of $z < 4$ *B*-dropouts. Stark et al. (2010) showed that within this subset, strong Lyα emitting galaxies have bluer UV continuum slopes than their non-emitting counterparts, but as there existed no high-quality infrared photometry in the GOODS fields at this time, it was necessary to parameterize distributions at higher redshift by their absolute magnitude. However, in addition to the now-completed Keck spectroscopic survey of LBGs over $3 < z < 6$ (Stark et al. 2010; D. P. Stark et al. 2014, in preparation), the CANDELS *HST* imaging program (GO 12444-5; PI: Ferguson/Riess/Faber) provides the necessary data to explore the potential of the UV continuum slope as a predictor for $EW_{Ly\alpha}$. The addition of Y_{105}, J_{125}, and H_{160} photometric data enables the derivation of accurate UV continuum slopes for catalogued galaxies, given that for each source there are a minimum of three broadband filters longward of the Lyman break. As such UV continuum slopes are now available for the growing body of $z > 7$ photometric galaxy samples (e.g., Dunlop et al. 2013; Bouwens et al. 2013), we can realize an Lyα fraction test that overcomes several of the issues associated with the earlier approach.

In the following, we discuss the new data for the Keck $3 < z < 6$ spectroscopic sample (D. P. Stark et al. 2014, in preparation) and analyze them in the context of a distribution function based on the observed UV continuum slopes of the population. We then apply the method to an updated sample of near-infrared spectroscopic data beyond $z \simeq 6$.

3. IMPROVED POST-REIONIZATION DATA

3.1. DEIMOS/FORS2 Spectroscopy

As discussed in Stark et al. (2010, 2011), the $3 < z < 6$ LBG candidates that form the basis of the post-reionization sample were targeted in the GOODS-N and GOODS-S fields using the DEIMOS spectrograph on the Keck II telescope. The final catalog is being prepared for release by D. P. Stark et al. (2014, in preparation). By retroactively applying the same photometric selection criteria, a spectroscopic sample in the GOODS-S field using the FORS2 spectrograph on the ESO VLT was added (Vanzella et al. 2009 and references therein). Full details of these spectroscopic campaigns can be found in the above-referenced articles.

The GOODS-N sample consists of 393 LBG candidates targeted with DEIMOS observed over the course of 2008-2010. The targets include *B*-, *v*-, and *i*-drops, and the spectroscopically confirmed sample spans a redshift range of $3.33 < z < 5.99$. Typical 10σ limiting Lyα fluxes for these targets were in the range $(1.0–1.5) \times 10^{-17}$ erg cm^{-2} s^{-1}.

The complementary FORS2 campaign (Vanzella et al. 2009) targeted 214 LBG candidates in GOODS-S between 2002 and 2006. These targets were, on average, brighter than those studied at Keck (see Stark et al. 2010, Figure 2), and the confirmed galaxies span a redshift range $3.19 < z < 6.28$. In total, the sample comprises 607 galaxies, 269 of which are spectroscopically confirmed. We direct the reader to D. P. Stark et al. (2014, in preparation) for further details.

3.2. Photometry

The primary advance in our analysis of the EW distribution of Lyα in the above spectroscopic sample relates to the combination of the earlier *HST* Advanced Camera for Surveys

(ACS) optical imaging data with new, deep WFC3/IR near-infrared data critical to assessing how Lyα emission correlates with the measured UV continuum slope. To reliably bring together the various imaging data sets, it is necessary to account for the significantly different point-spread functions (PSFs) between the ACS (FWHM $\sim 0\rlap{.}{''}09$) and WFC3/IR (FWHM $\sim 0\rlap{.}{''}16$) instruments. In the GOODS-S field, we utilized the published, PSF-matched catalog of Guo et al. (2013), which uses the publicly released v2.0 ACS and v1.0 WFC3 images, constructs stellar profiles to derive the PSFs in order to convolve the higher-resolution, shorter-wavelength data, and performs isophotal photometry on the smoothed images.

For GOODS-N, we also utilized the publicly released v2.0 ACS images, but as no CANDELS WFC3/IR mosaic was released at the time of this analysis, we constructed our own. The first 13 (out of 18) epochs of GOODS-N observations, released as individual v0.5 mosaics, were combined using the routine SWARP (Bertin et al. 2002), with individual weights assigned according to the exposure time. PSF matching was implemented using the ColorPro program (Coe et al. 2006). A PSF was constructed for each filter by shifting and stacking $\sim$20 bright, unsaturated stars. All objects were detected using the H_{160} image, and colors were determined using matched isophotal apertures after degrading the resolutions of all other images to that of the H_{160} image. Colors were corrected to total by using the offset between the isophotal magnitude and that derived from MAG AUTO using SExtractor (Bertin & Arnouts 1996).

4. ANALYSIS

4.1. Lyα and the UV Continuum

The basis of our analysis relies on accurate determinations of both the Lyα EWs (EW$_{Ly\alpha}$) and ultraviolet slopes of our sample. Thus, we detail here the methodology used in determining both these quantities for use in our analysis.

In order to measure the UV slope for each object in our sample, we first used a custom photometric redshift code to determine the approximate redshifts of those galaxies without spectroscopic confirmation. The code fits a suite of synthetic spectra from the Bruzual–Charlot (BC03, Bruzual & Charlot 2003) models to the ACS and WFC3 photometry, evaluating the likelihood of each by computing the χ^2 statistic between the observed and synthetic fluxes. To simplify the fitting procedure, we used a restricted set of templates, choosing only $Z = 0.2 \, Z_\odot$, constant-star-formation models without nebular emission. We use the reddening prescription of Calzetti et al. (2000), varying $E(B - V)$ from 0.0 to 1.5. To determine the best-fit redshift, we marginalized across all other parameters (mass, dust extinction, and age) and used the maximum likelihood value. Although we defer a full discussion of how well our code performs at recovering those galaxies with spectroscopic redshifts to a future paper (D. P. Stark et al., in preparation), we note that after removing nine outliers from the sample, we find excellent agreement, with $\sigma_z/(1 + z) = 0.04$.

We measure the UV continuum slope using the β formalism first introduced by Calzetti (1994), where the flux is parameterized as $f_\lambda \propto \lambda^\beta$. In our fitting process, we include photometric filters with central rest wavelengths within the range defined by Calzetti (1994), $1350 < \lambda/\text{Å} < 2600$, determined by either the spectroscopic redshift, if available, or the best-fit photometric redshift. This range is also similar to that used previously in the literature (Bouwens et al. 2013; Rogers et al. 2014) and ensures

that passbands contaminated by Lyα or IGM absorption will be excluded. As in Bouwens et al. (2013), we use the effective filter wavelengths appropriate for a $\beta = -2$ spectrum, since the measured UV slopes in our sample generally populate the $-2.5 < \beta < -1.5$ range, and an error floor of 0.05 mag, or 5%, for all filters. We note that this power-law-fitting method is one of two techniques commonly adopted at high redshift, with the other being a direct measurement of β from the best-fitting spectral energy distribution (SED), first advocated by Finkelstein et al. (2012). Given that detailed tests performed in both Finkelstein et al. (2012) and Rogers et al. (2014) find both these methods to perform comparably, we have adopted the power-law method for transparency and ease of reproduction without being tied to a specific set of spectral models. We compute the total UV absolute magnitude for each object using either the best-fit or, if available, spectroscopic redshift to determine the filter with rest-frame wavelength closest to 1400 Å.

A grid of power-law slopes with $-3.5 < \beta < 0.5$ and $\Delta\beta = 0.01$ was fit to the observed photometry, and the relative likelihood of each was computed using $p(\beta_i) \propto \exp(-\chi^2/2)$, appropriate for Gaussian-distributed errors. This allows us to construct a likelihood curve for the UV slope of each galaxy, central to the fitting method we describe later. After the fitting, the image cutouts, photometry, and resulting $p(\beta)$ were manually inspected for each galaxy, flagging and removing objects with clearly deviant photometry or incorrect solutions. After this process, 297 of our sample of 393 galaxies in GOODS-N and 154 of our 214 galaxies in GOODS-S remained. Galaxies removed were typically those adjacent to bright objects, for which accurate photometry could not be assured, and faint, distant $z \geqslant 5$ galaxies that appeared in the shallower CANDELS-Wide field, for which accurate UV slopes could not be determined.

Measurements of the Lyα EW were taken from D. P. Stark et al. (2014, in preparation), with errors computed using both the 1σ flux errors in the spectrum and errors in M_{UV}, added in quadrature to produce a likelihood curve for $p(\text{EW}_{Ly\alpha})$. For cases where Lyα was not spectroscopically detected, we assume one of three cases: (1) the line flux falls below the 10σ limit (typically $(1.0\text{--}1.5) \times 1.0^{-17}$ erg cm^{-2} s^{-1}); (2) the line emission, though brighter than the 10σ limit, is missed, owing to poor sky subtraction or obscuration by skylines; or (3) the object is a contaminant outside the expected redshift range. With the exception of skylines, which are discussed along with the rest of the implementation of our approach in Section 5.1.3, our limiting sensitivity varies only slowly as a function of wavelength. Thus, in the case that a galaxy is not spectroscopically confirmed, we compute the EW limit from the M_{UV} and limiting flux determined at the best-fit photometric redshift.

With these results in hand, we now provide the basic evidence that the UV continuum slope of a galaxy, β, is a reliable predictor of its Lyα EW. In Figure 1 we plot the best-fit β for each galaxy in our final sample against either its EW$_{Ly\alpha}$ or its 10σ upper limit, if a line is not detected. The red crosses denote the mean EW$_{Ly\alpha}$ in bins spanning $\Delta\beta = 0.25$, where undetected objects are set to have a value of 0. To provide a more realistic estimate of the mean, the blue squares are calculated by assigning undetected objects an EW using a Monte Carlo simulation, where each object has a constant probability density between 0 and its upper limit. Error bars for each bin are calculated by bootstrap sampling.
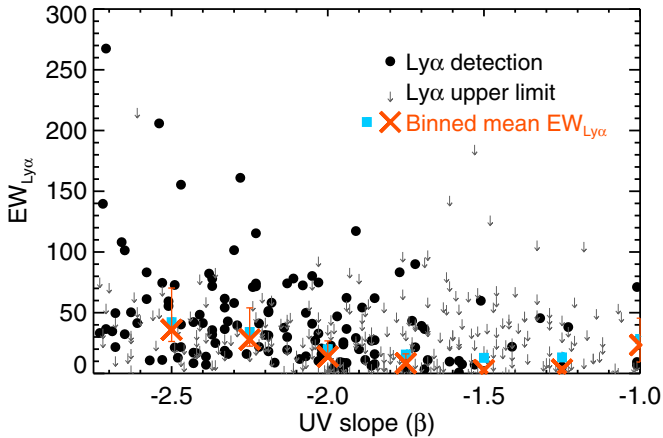
Figure 1. Compilation of our entire GOODS catalog of Lyα EWs as a function of UV slope, β. Red triangles and blue squares show the average EW, binned in steps of 0.25 in β, displaying a strong increase toward bluer slopes. Red triangles are computed by setting the EW of all undetected objects to 0, while blue squares use a Monte Carlo simulation where the average EW is uniformly distributed between the measured upper limit for each galaxy and 0. Thus, the true average value should lie between these two estimates. This data set forms the basis of our predictive model for Lyα emission incidence as a function of UV slope, but we undertake a full Bayesian analysis to deal with nondetections.

(A color version of this figure is available in the online journal.)

As shown from earlier work by Shapley et al. (2003) and Stark et al. (2010), a clear trend of increasing mean $EW_{Ly\alpha}$ with bluer (more negative) UV slopes is visible. With our large spectroscopic sample, this can be seen directly through the measures of individual objects, rather than via stacked spectra or consideration of average β values. Encouraged by this trend, we now develop a model that can predict the probability distribution of $EW_{Ly\alpha}$ given a measured value of β, for example, for a $z > 7$ galaxy. At the end of this section, we will also use this model to show that the UV slope is a more reliable predictor of $EW_{Ly\alpha}$ than the absolute UV magnitude.

4.2. The UV Slope-dependent Equivalent Width Distribution

We now seek to establish a formalism to predict the probability distribution for $EW_{Ly\alpha}$ given a particular measured UV slope. As our goal is to improve the accuracy of the Lyα fraction test, we first consider those galaxies with blue UV slopes, likely to have strong Lyα emission before processing through the IGM.

4.2.1. Equivalent Width Distributions for a Fixed UV Continuum Slope

As an illustration of our new method, we begin by examining the distribution of the Lyα EW for a fixed UV continuum slope, β. A natural choice is $\beta = -2.3$ given that Bouwens et al. (2013) find that faint ($M_{UV} \geqslant -19$), high-redshift ($z \geqslant 6$) galaxies have slopes that asymptote to this value. Near-IR spectrographs such as MOSFIRE have begun to target these galaxies in earnest (this work; Finkelstein et al. 2013; Treu et al. 2013), and it is becoming increasingly important to characterize their expected IGM unprocessed Lyα emission. To construct a sample of galaxies for this task, we limit our overall sample of 468 galaxies with $3.19 < z < 6.28$ to those galaxies with a best-fit value within $\Delta\beta = 0.25$ of -2.3, resulting in a total of 131 objects.

We now require a model to represent the IGM unprocessed EW distribution. In Appendix A we review four such options

using the methodology outlined here, and we find that a lognormal distribution provides a significantly better fit than any of the others. In this case the natural logarithm of $EW_{Ly\alpha}$ obeys a normal distribution. The two relevant parameters of the distribution, μ and σ, are typically referred to as the location parameter and scale parameter, respectively. They denote the mean of the natural log of $EW_{Ly\alpha}$ and its variance. However, while the median of the distribution is given, as might be expected, by $\exp(\mu)$, the mean is slightly larger at $\exp(\mu+\sigma^2/2)$. A third parameter, A_{em}, determines the fraction of galaxies that have $EW_{Ly\alpha} > 0$, as there is no reason a priori to expect all galaxies to display Lyα in emission. The resulting distribution can be written as

$$p(EW) = A_{em} \times \frac{1}{\sqrt{2\pi}\,\sigma\,EW} \exp\left(-\frac{(\ln(EW) - \mu)^2}{2\sigma^2}\right) + (1.0 - A_{em}) \times \delta(EW). \qquad (1)$$

In Figure 2 we illustrate how the $EW_{Ly\alpha}$ probability distribution function (pdf), $p(EW)$, and the Lyα fraction, $x_{Ly\alpha}$, change as these parameters are varied.

We now describe the Bayesian formalism we developed to evaluate the likelihood of the underlying parameters for our lognormal distribution, and we determine which provides the best fit to the data. The entire set of spectroscopic Lyα observations is denoted as *Obs*; this contains the information for observations of each individual galaxy, Obs_i. We can then denote the parameters for the model being fit as $\theta \equiv [\mu, \sigma, A_{em}]$. Our overall goal is to determine the probability distributions for the underlying parameters of each model, given our observations, i.e., $p(\theta|Obs)$. Using Bayes's theorem, we can rewrite this as

$$p(\theta|Obs) \propto p(\theta) \times p(Obs|\theta). \qquad (2)$$

Here $p(\theta)$ represents our uniform priors for the underlying parameters, while the term on the right-hand side represents the probability of our observations given the model parameters. For any single object in which we measure a definite $EW_{Ly\alpha}$, this posterior probability can be expressed as

$$p(Obs_i|\theta) = \int_0^\infty p(EW_{Obs,i})p(EW|\theta)\,dEW. \qquad (3)$$

In the case of an object for which Lyα remains undetected above our (10σ) limit, we compute the posterior probability as

$$p(Obs_i|\theta) = p(EW < EW_{10\sigma}|\theta) + p(EW > EW_{10\sigma}|\theta) \times C_{1,i} + C_2. \qquad (4)$$

Here the first term represents the probability that the object intrinsically poses an $EW_{Ly\alpha}$ below our detection limits, while the $C_{1,i}$ term takes into account incompleteness in the sample caused by skyline residuals, particularly at longer wavelengths. We compute this term for each object as a function of wavelength by determining what fraction of the spectra has noise above $2.5\times$ the median level found in clean sky regions. This fraction is then weighted by the photometric redshift probability distribution. Thus, galaxies at $z \sim 4$ have a large completeness of $\sim$95%, while some at $z \sim 6$ can have completeness as low as $\sim$68%. Contamination by low-redshift sources is taken into account through the final term, C_2. We chose a modest value for our contamination terms of $C_2 = 0.05$, motivated by the lack of low-redshift interlopers found in spectroscopic follow-up surveys (Stark et al. 2010; Pentericci et al. 2011). We note
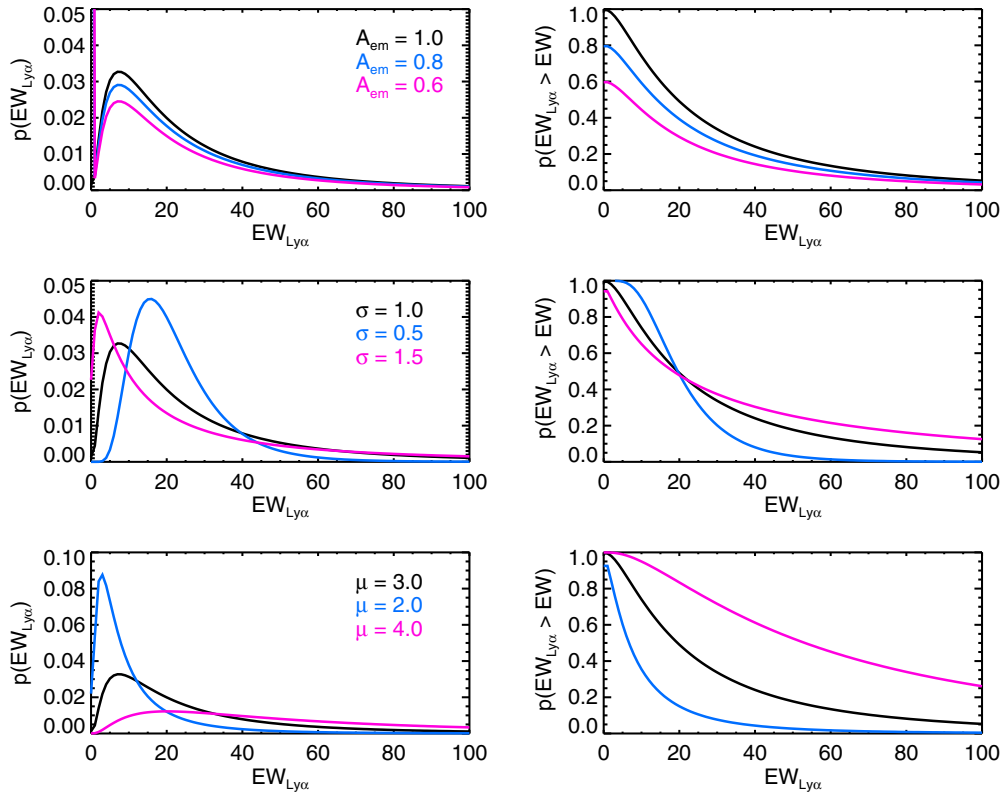
Figure 2. Example curves for how our lognormal model of $EW_{Ly\alpha}$ distribution varies with each parameter. Left: probability distributions for $EW_{Ly\alpha}$. In each panel, the black curve has the same parameter values: $\mu = 3.0$, $\sigma = 1.0$, $A_{em} = 1.0$. From top to bottom, the two colored curves each display the effect of a change in a single parameter on the distribution. Right: complementary cumulative distribution functions for the same parameters used in each left panel. This method of display is especially useful, as the Lyα fraction, $x_{Ly\alpha}$, for any $EW_{Ly\alpha}$ can simply be read off the plot by finding the value of the curve at the desired $EW_{Ly\alpha}$ along the x-axis.
(A color version of this figure is available in the online journal.)

that only a differential contamination between our $3 < z < 6$ sample and our $z > 6$ sample will alter our results. The full posterior distribution for the parameters can then be computed by multiplying the individual posterior probabilities for each object. This allows us to infer the most likely values on parameters, as well as their marginalized and unmarginalized errors.

We display the best-fit distribution for our sample of 131 galaxies with $3.19 < z < 6.28$, overplotted on a histogram of their $EW_{Ly\alpha}$ detections and upper limits in Figure 3. The best-fit parameters are $\mu = 2.7 \pm 0.7$, $\sigma = 1.4^{+0.9}_{-0.5}$, and $A_{em} = 1.0^{+0.0}_{-0.4}$. As we show in Appendix A, this formalism is the best fit to our post-reionization data.

4.2.2. A Generalized Approach

Now that we have determined the distribution that best fits the data at the key UV slope value of $\beta \sim -2.3$, we proceed with a more appropriate goal of determining the EW distribution across all values of β. Although faint galaxies at $z \geqslant 6$ may asymptote to $\beta \sim -2.3$, the UV-bright galaxies almost certainly do not (Bouwens et al. 2013). In order to fully leverage the $x_{Ly\alpha}$ test for the more luminous objects, we must use a model that determines the EW distribution across a wide range of β.

To achieve this goal, we extend the Bayesian formalism introduced above. The differences are twofold. Firstly, we now include our entire sample of spectroscopically observed galaxies when fitting, rather than just those narrowly clustered around a particular value of β. Secondly, we must reconsider the nature
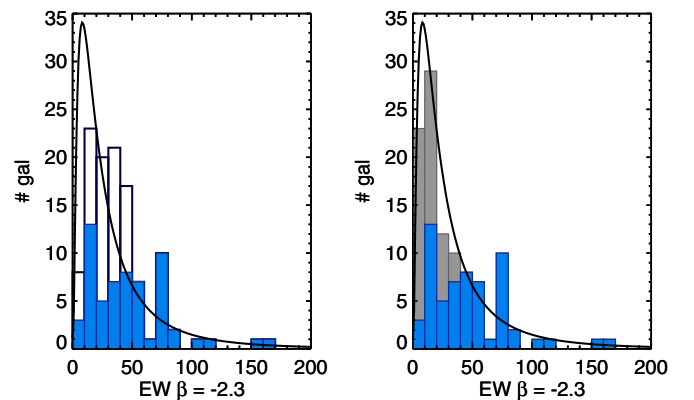

Figure 3. Left: histogram of our observed $EW_{Ly\alpha}$ detections (solid blue) and upper limits (unfilled black) for galaxies with best-fit $\beta = -2.3 \pm 0.25$. Overplotted in solid black is our best-fit lognormal model as described in Section 4.2.1. In this panel, undetected are added to the histogram as unfilled black simply at their measured upper limit, so the best-fit model should only trace the filled blue histogram at $EW > 50 \text{\AA}$ where incompleteness is not an issue. Right: same as on the left, but for this plot to demonstrate the model fit, we instead plot upper limits in gray probabilistically according to the best-fit model. For example, an upper limit for a single galaxy measured at $EW < 45 \text{\AA}$ will add counts of 0.30, 0.34, 0.22, and 0.14 to the bins centered at 5, 15, 25, and 35 $\text{\AA}$, respectively. In this plot, the solid curve should trace the total (filled gray+blue) histogram at all EWs.
(A color version of this figure is available in the online journal.)

of Equation (1) since it is clear that the EW distribution varies as a function of β from Figure 1.

It is most reasonable to consider that the location parameter, μ, varies with β since it is this parameter that governs the
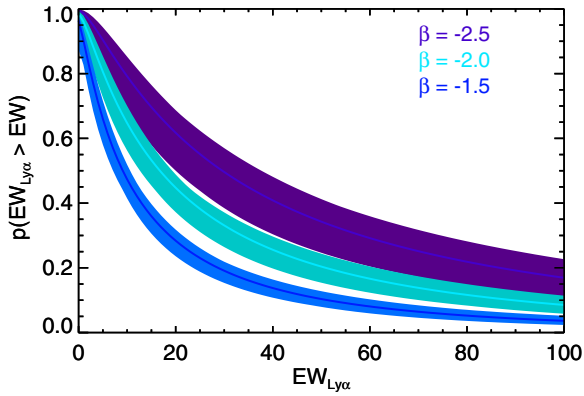
Figure 4. Inverse cumulative distribution functions of $EW_{Ly\alpha}$ for our best-fit model, plotted as a function of β. Shaded regions for a given β denote the spread of $EW_{Ly\alpha}$ allowed within the 68% highest posterior density confidence intervals of our full model. For a desired β, the Lyα fraction, $x_{Ly\alpha}$, for an arbitrary EW is defined by the y-value of the curve at the given $EW_{Ly\alpha}$.

(A color version of this figure is available in the online journal.)

redistribution of EWs (see Figure 2). For convenience we assume that μ varies linearly with β, whence $\mu(\beta) = \mu_a + \mu_s \times (\beta + 2)$, where μ_a represents the location parameter at $\beta = -2.0$. Prior to selecting this model, we performed similar fits to a narrow range in β as in the above section, with different ranges in β, and found that while the best fit for μ varied strongly with the UV slope, both σ and A_{em} did not.

Figure 4 shows a representation of the product of the resulting formalism, while we detail the full results in Appendix B. Because our model treats the location parameter as a linear function of β, we can generate an EW pdf for any UV slope, although our sample only provides meaningful constraints in the observed range $-2.5 < \beta < -1.25$. In this figure, we plot three examples and their associated errors. Of course, a measured UV slope for a high-redshift galaxy has some error uncertainty and thus its own probability distribution, $p(\beta)$. To obtain $p(EW)$ given our model and an observation of β, we can marginalize over β for each galaxy in the sample:

$$p(EW_i) = \int p(EW|\beta) p(\beta_i) \, d\beta. \quad (5)$$

We also considered whether there remained any trends in the evolution of $EW_{Ly\alpha}$ with redshift, once the trend with UV slope had been accounted for. To do this, we split our sample into three bins in UV slope, centered on $\beta = -1.3, -2.0$, and -2.7, with objects assigned according to their best-fit value. Within these UV slope bins, we then split the sample into two bins at the median redshift of our survey, $z = 4.25$, computed the average $EW_{Ly\alpha}$ in each bin, and fit a slope to the relation between both bins. The only UV slope range for which evidence for evolution significantly greater than 1σ is seen is that centered at $\beta = -1.3$, with $(dEW/dz) = 8.4 \pm 3.4$. As this is much redder than the majority of UV slopes seen in $z \geqslant 7$ galaxies, we felt that adding a dependence on redshift to our model would needlessly complicate the results with little gain in predictive power.

4.3. UV Slope versus UV Luminosity

We now return to one of the assumptions motivating this paper. Does parameterizing the likelihood of Lyα emission via the UV slope represent a statistically better option than the combination of absolute UV magnitude and redshift used in previous high-redshift Lyα studies? Since we can measure β,

M_{UV}, $EW_{Ly\alpha}$, and a photometric or spectroscopic redshift for 451 objects in our spectroscopic sample, we can directly address this question.

In the widely used UV luminosity method, the dependence on M_{UV} is handled in a discrete manner, assigning galaxies into one of two bins, depending on whether they are greater or less than $M_{UV} = -20.25$. $x_{Ly\alpha}$ is also calculated in discrete bins at $z = 4, 5$, and 6. For purposes of comparison, we use the Bayesian formalism developed in this paper to generalize this to a continuous model. To do so, we alter the definition of μ from $\mu(\beta) = \mu_a + \mu_s \times (\beta + 2.0)$ to $\mu(M_{uv}, z) = \mu_a + \mu_{s,M_{uv}} \times (M_{uv} + 19.5) + \mu_{s,z} \times (z - 4.0)$. This thus replaces the linear dependence of the location parameter, μ, on UV slope, with a bivariate linear dependence on UV magnitude and redshift. We then use the same fitting process to determine the optimal values for all parameters in the model.

To compare how well each of these two models fits the available data, we use the Bayesian evidence ratio, or Bayes factor. The Bayes factor is a measure of the relative likelihood of a given model having generated the data, and it can be expressed as an integration of the likelihood function over all possible values for each parameter in the model

$$BF = E_1/E_2 = \int p(\theta_1|Obs) d\theta_1 / \int p(\theta_2|Obs) \, d\theta_2. \quad (6)$$

Evaluating this for both models yields a significant gain via a ratio of $E_\beta/E_{M_{uv},z} = 29$. Thus, the β formulation is 29 times more likely to have generated the data than the predictive model based on M_{UV} and z. Using the scale advocated by Kass & Raftery (1995) provides "strong" preference for the predictive model based on β.

5. FIRST APPLICATION TO DATA WITHIN THE REIONIZATION ERA

Although the body of spectroscopic data targeting galaxies beyond $z \simeq 6$ in the reionization era remains sparse, it is growing rapidly, particularly through the advent of multi-slit near-infrared spectrographs such as MOSFIRE. Thus, we are encouraged to apply our new method to such data. In addition to collating earlier relevant data available in the literature, we also present the first results from our new survey using MOSFIRE.

5.1. A New MOSFIRE Survey

As part of a long-term survey targeting $z > 7$ galaxies using the MOSFIRE spectrograph on the Keck I telescope, we have secured deep spectroscopic observations in both 2013 November and 2014 March targeting two different fields. One represents distant sources in a deep *HST* blank field with accurate photometric redshifts, and the other targeted gravitationally lensed sources with extensive multi-band photometry.

5.1.1. The GOODS-South/Ultra Deep Field

On the night of 2013 November 5, we secured a total of 3.5 hr of exposure in the Hubble UDF region of GOODS-South. Observations were taken through the MOSFIRE Y-band filter and grating, using 0''7 slits through intermittent high cirrus clouds and $\sim$0''8 seeing. Individual exposure times were 180 s, read out using MCDS-16 sampling, with an AB dither pattern separated by 2''5. Our observations reach a resolution of $R = 3270$.

A total of 16 $z > 7$ candidates were included on this mask selected from an initial list of $z > 6.5$ targets from our UDF12

Table 1
Summary of MOSFIRE Survey for Lyα

Name	R.A.	Decl.	J_{125}	z_{phot}	t_{exp} (hr)	5σ EW Limit
A611-0193	8:01:00.32	36:04:24.3	26.0	$7.9^{+0.2}_{-0.2}$	2.6	12.9
MACS 0647-1411	6:47:40.91	70:14:41.6	26.1	$7.6^{+0.3}_{-0.4}$	1.0	20.7
RX J1347-0943	13:47:33.90	−11:45:09.4	26.4	$7.5^{+0.4}_{-0.3}$	1.0	27.7
CANDY-2272447364	3:32:27.24	−27:47:36.4	27.1	$9.1^{+0.5}_{-1.1}$	2.35	60
CANDY-2243349150	3:32:24.33	−27:49:15.0	27.2	$8.5^{+0.4}_{-0.5}$	2.35	57
GS-2098-8535	3:32:20.98	−27:48:53.5	27.2	$9.7^{+0.3}_{-0.7}$	2.35	64
GS-2533-8541	3:32:25.33	−27:48:54.1	26.4	$7.6^{+0.2}_{-0.3}$	2.35	30
GS-2779-5141	3:32:27.79	−27:45:14.1	27.2	$9.6^{+0.4}_{-0.7}$	2.35	64
GS-4402-7273	3:32:44.02	−27:47:27.3	26.9	$7.4^{+0.3}_{-0.3}$	2.35	50
UDF12-3313-6545[a]	3:32:33.13	−27:46:54.5	28.5	$7.4^{+0.2}_{-0.3}$	2.35	200
UDF12-3722-8061	3:32:37.22	−27:48:06.1	27.8	$7.7^{+0.2}_{-0.3}$	2.35	110
UDF12-3846-7326	3:32:38.46	−27:47:32.6	29.0	$7.0^{+0.5}_{-0.5}$	2.35	330
UDF12-3880-7072	3:32:38.80	−27:47:07.2	26.9	$7.5^{+0.2}_{-0.2}$	2.35	220
UDF12-3939-7040	3:32:39.39	−27:47:04.0	28.6	$7.7^{+0.3}_{-0.7}$	2.35	180
UDF12-3762-6011	3:32:37.62	−27:46:01.1	28.4	$8.1^{+0.4}_{-0.5}$	2.35	150
UDF12-3813-5540	3:32:38.13	−27:45:54.0	28.2	$8.2^{+0.2}_{-0.4}$	2.35	70
UDF12-4256-7314	3:32:42.56	−27:47:31.4	27.3	$7.1^{+0.2}_{-0.2}$	2.35	50
UDF12-4308-6277	3:32:43.08	−27:46:27.7	28.5	$8.0^{+0.2}_{-0.1}$	2.35	200
UDF12-4470-6443	3:32:44.70	−27:46:44.3	27.3	$7.7^{+0.2}_{-0.3}$	2.35	70

Notes. MOSFIRE survey targets. 5σ limiting sensitivities are calculated using the median limiting flux limit between skylines and assuming a spectroscopic redshift of $z = 7.7$.
[a] Lyα detected at 4.0σ, as discussed in Section 5.2.

campaign and the GOODS-S field (Schenker et al. 2013b; McLure et al. 2013), augmented with two additional *Y*-drops outside the UDF proper from Oesch et al. (2012). We used the photometric redshift code, described in Section 4, to compute a redshift probability distribution for each object.

The UDF 2012 data set (Koekemoer et al. 2013; GO 12498; PI: Ellis) offers many distinct advantages for this first application of our method. Foremost, by virtue of the extraordinarily deep optical and F105W data available, contamination by foreground objects, as determined by the photometric scatter simulations in Schenker et al. (2013b), is ∼3%, down to the $J_{125} = 28.6$ limit of our targets. This contrasts with the ⩾10% contamination rate affecting galaxies in the CANDELS fields (Oesch et al. 2012). Secondly, as a result of a strategic deployment of near-infrared filters, our UDF 2012 candidates have better-defined redshift probability distributions, allowing us to more confidently exclude the possibility of Lyα emission in the event of a nondetection. The median 68% confidence interval in photometric redshift for the UDF 2012 objects on our MOSFIRE mask is smaller by $\Delta z = 0.2$ compared to our GOODS-S targets (and most likely to earlier published blank field surveys; see Section 5.3).

The final target selection for this aspect of our survey was arranged to formally maximize the expected number of detected lines and thus our leverage in calculating $x_{Ly\alpha}$. As a first attempt, we used the $z \sim 6$ histograms of Stark et al. (2011) to predict the distribution of $EW_{Ly\alpha}$ for each object, as a function of its UV magnitude. The fractional number of expected detections was calculated for each object, taking into account the photometric redshift distribution (as our spectral coverage is incomplete), UV magnitude, and expected limiting flux for a likely MOSFIRE exposure. This exercise resulted in the final list of 16 candidates presented in Table 1.

5.1.2. CLASH Lensing Sample

Over the course of our November and March observations, we also targeted three candidates with a photometric redshift $z \geqslant 7.5$ from the CLASH *HST* survey (GO 12065-12791; PI: Postman) as collated in Bradley et al. (2014). Although these targets can only be surveyed individually, limiting our efficiency, as each is sampled with eight *HST* filters at or longward of $\lambda = 7750\,\text{Å}$ they have sharp redshift probability distributions and well-determined UV slopes, making them ideally suited for our new method.

We first targeted the $z \sim 7.9$ candidate in A611 on 2013 November 5, securing 1.1 hr of on-source integration in 0″.8 seeing. A further 1.5 hr of integration were possible on 2014 March 5. In our 2014 March run, we also targeted candidates in RX J1347 and MACS 0647 for 1.0 hr each. Typical seeing conditions for the March nights were 0″.60–0″.65. Observations were undertaken in the same manner as described above, with the exception that all slits were 0″.8, in anticipation of poor seeing given in the forecast, resulting in a slightly decreased resolution of $R = 2860$. Full details of the three targets are presented along with our GOODS-S/UDF 2012 sample in Table 1.

5.1.3. Data Reduction

The data was reduced using the publicly available MOSFIRE data reduction pipeline.[3] This pipeline first creates a median, cosmic-ray-subtracted flat-field image for each mask. Wavelength solutions for each slit are fit interactively for the central pixel in each slit and then propagated outward to the slit edges to derive a full wavelength grid for each slit. Typical residual rms on the wavelength solution ranged from 0.3 to 0.6 Å. The sky

---

3 https://code.google.com/p/mosfire/

background is estimated as a function of wavelength and time using a series of B-splines and subtracted from each frame. The nodded A–B frames are differenced, stacked, rectified, and output for use along with inverse variance-weighted images used for error estimation.

Examination of our GOODS-S data revealed a gradual $0''6$ ($\sim$3 pixel) drift in the spatial direction over the course of our integration, which needed to be corrected for. The drift was tracked using a $J_{125} \sim 19$ star conveniently placed on one slit. The intensity of the star allowed us to follow the extinction for each frame, and eliminating those frames affected by significant extinction (measured by having less than 0.5 the maximum observed flux from the star on our mask), we secured 2.35 hr of useful exposure. As a result of this drift, the star on our original reduction displayed a somewhat greater FWHM, of $\sim 1''2$, than any of our individual exposures, which typically had an FWHM $\sim 0''8$. To correct this, we registered the relative positions of all frames by fitting a Gaussian profile to the star along the spatial axis. Given the drift over the entire exposure, we then arranged the frames into three separate groups, with the spatial positions in all frames consistent to within $\sim$1 pixel. Each of the three frame groups was reduced individually using the same procedure described above. To produce our final science stack, the three reductions were then shifted by the appropriate integer number of pixels and co-added, weighting by exposure time. Using this method, we were able to reduce the observed stellar FWHM to $\sim 1''0$.

In order to flux-calibrate our data, we observed an A0V star at twilight during each night using the long-slit mode of MOSFIRE and the same slit width used to observe each mask. To determine a relative flux calibration as a function of wavelength, we use a spectrum of Vega from the *HST* CALSPEC database,[4] scaled to the Y magnitude of our A0V standard. Then, to achieve an absolute calibration, we use the same method as in our previous work (Schenker et al. 2013a), including a star with $Y_{105} < 20.0$ on each mask in addition to our science targets. This star is then used to determine the normalization of our flux calibration by matching the median flux for the star in wavelengths probed by the MOSFIRE Y band to the value measured from *HST* photometry. This method has the advantage of intrinsically accounting for the effects of thin clouds and slit loss due to seeing and pointing error during our observations. We assume the same, conservative 15% error in flux calibration that was derived using this method in Schenker et al. (2013a).

Our final mask reaches a median 5σ limiting sensitivity between skylines of $\sim 7.0 \times 10^{-18}$ erg cm^{-2} s^{-1}. At the time of our reduction, there existed a known problem in the error spectrum output by the MOSFIRE data reduction pipeline, so we instead measured the limiting sensitivity directly from the final science frame. This is done by selecting rows on the mask that are free from sources, as determined from our *HST* imaging, and thus should contain no signal. To compute the limiting sensitivity, we then measure the variance in these 23 independent rows, using a 5×7 pixel aperture. This translates to $0''9$ in the spatial direction and 7 Å in the spectral direction. We note that approximately $\sim$33% of the Y-band spectral range is obscured by skylines at the MOSFIRE resolution of $R \sim 3380$ given our $0''7$ slit width.

5.2. A New z = 7.62 Lyα-emitting Galaxy

We inspected the reduced, two-dimensional spectra of all 16 objects by eye to search for Lyα emission. From this, we were

---

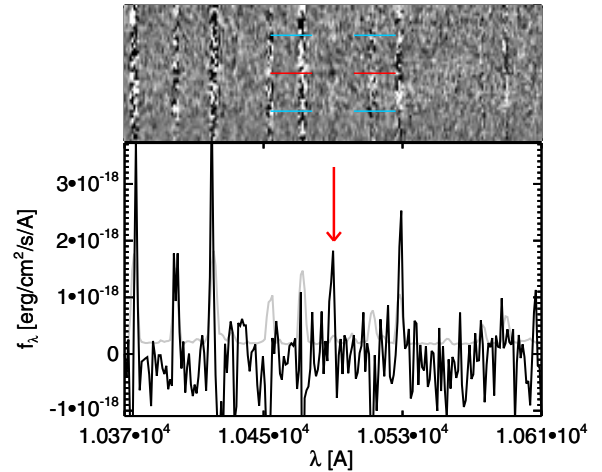[4] http://www.stsci.edu/hst/observatory/cdbs/calspec.html

Figure 5. MOSFIRE observations of our lone target with visible line emission, UDF12-3313-6545. The full two-dimensional spectrum is shown at the top, with the one-dimensional spectrum plotted at the bottom, along with the error spectrum in gray. Given our A–B reduction strategy (described in Section 2), our two-dimensional spectrum shows the expected positive signal (red line) flanked by two negative signals (blue lines), each separated by the amplitude of the dither pattern.

(A color version of this figure is available in the online journal.)

able to locate only a single candidate emission line. Surprisingly, this emission line is located in one of our faintest targets, UDF12-3313-6545 (first identified by McLure et al. 2010; Bouwens et al. 2011), with a measured flux of $5.6 \pm 1.4 \times 10^{-18}$ erg cm^{-2} s^{-1}. We display both the one- and two-dimensional spectra for this object in Figure 5. If the line is indeed Lyα, the galaxy lies at a spectroscopic redshift of $z = 7.62$, making it a promising candidate for the most distant spectroscopically confirmed galaxy to date. We present the relevant details and *HST* cutouts of this galaxy in Figure 6.

Given the faintness of the emission line (detected at 4.0σ), line asymmetry, commonly used to identify Lyα, is not detectable, so we present the arguments in favor of its association with our target and its reality. First, to search for lines, we predicted the expected y-pixel position of each object on the reduced image from its distance from the bottom slit edge, tying measurements to the rectified frame by measuring this offset for the star on our mask. For UDF12-3313-6545, the centroid of the identified line is consistent with our predicted position to within 1.5 pixels, or $0''27$. In our earlier investigations with MOSFIRE (Schenker et al. 2013a), we found that our code could accurately predict positions to an accuracy of roughly 1.0 pixels. Thus, given this information and our measured FWHM in the final reduction of $1''0$ (5.6 pixels), an offset of only $0''27$ is within the acceptable tolerance.

With regard to the reality of the line, we discuss two separate lines of evidence. Firstly, the line displays a clear positive signal flanked by two negative peaks, with signal-to-noise ratios of -2.3σ and -3.1σ. These are the result of the sky subtraction method, as our final A–B frame will have a negative signal contributed by the B images, and vice versa for the B–A frame. The presence of two negative signals indicates that the line was visible separately in both A and B exposures and is not due to an artifact such as a hot pixel or cosmic ray. Secondly, as part of our reduction we segmented the data into three separate blocks based on pixel position. Since each block contains both A and B frames, this represents an orthogonal split to the observations above. In the final stack of each block, we measure
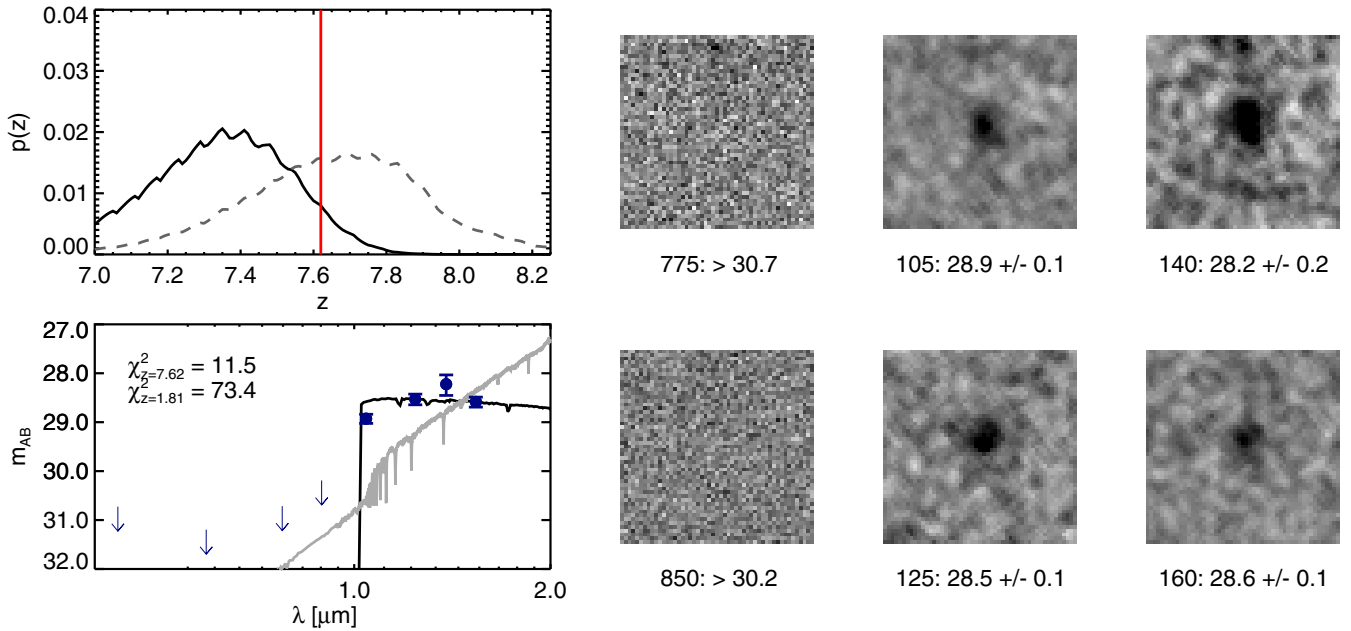
Figure 6. Right: 1″5 diameter cutouts and total magnitudes (see data section) of our MOSFIRE target, UDF12-3313-6545. As expected for a $z > 7$ candidate, the object is not formally detected in a stack of the optical data. Top left: photometric redshift probability distribution function (pdf) for our target, with the actual spectroscopic redshift, $z = 7.62$, denoted in red. The solid black curve displays the pdf from the raw photometry, while the dashed gray curve shows the pdf after the observed MOSFIRE line flux has been subtracted from the Y_{105} data point. Bottom left: *HST* WFC3 photometry along with best-fitting $z = 7.62$ SED for the galaxy in black (assuming that the detected line is Lyα), and best-fitting $z = 1.82$ SED for the galaxy in gray if the detected line is instead [O II] $\lambda3727$.
(A color version of this figure is available in the online journal.)

the signal-to-noise ratio of the positive peak and find significances of 2.2σ, 1.4σ, and 2.3σ, further instilling confidence in the reality of this feature. Although the line has a surprisingly large rest-frame EW of 160 ± 40 Å, this is comparable to some discovered in $z \sim 6$ Lyα emitting galaxies (Ouchi et al. 2010). Notably, the spectroscopic redshift lies well within the 1σ confidence interval of our derived photometric redshift distribution when line emission is accounted for.

Having investigated the reality of the line, we also consider alternative identifications for the line besides Lyα. Given the robust measurement of a >1.3 (>2.0) magnitude break between the z_{850} and Y_{105} filters at 2σ (1σ), the only reasonable alternative explanation would be a detection of the [O II] $\lambda\lambda3727$–3729 doublet, given its proximity to the Balmer/4000 Å break. As such, we fit an SED to the available photometry at the redshift this would imply, $z = 1.82$, with the results displayed in Figure 6. The fitting procedure including the template set used is the same as described previously in Section 4.1, with the χ^2 statistic computed uniformly across all filters, including those with formal nondetections. Clearly, the best SED at $z = 7.62$ is strongly statistically preferred, with a χ^2 value lower by 62, partially driven by the fact that the galaxy is not detected in any of the four optical filters. As such, we consider Lyα to be the most likely explanation for our observed line.

5.3. Additional Data from Published Surveys

In order to achieve the most up-to-date and precise measurement of the Lyα opacity at $z \geqslant 6.5$, we have compiled a comprehensive sample of other near-infrared surveys for Lyα at high redshift, which we will utilize in our analysis. This includes our own prior work with Keck's NIRSPEC (Schenker et al. 2012), as well as a number of other surveys, using red-sensitive optical detectors on the VLT and Keck (Fontana et al.

2010; Pentericci et al. 2011; Ono et al. 2012; Pentericci et al. 2014), as well as independent work by Treu and collaborators using NIRSPEC (Treu et al. 2012) and MOSFIRE (Treu et al. 2013).

In total, this literature sample comprises 83 $z \geqslant 6.5$ galaxies for which follow-up spectroscopy at various depths has been attempted, plus an additional 19 from this work. To apply our method, we split this sample into two redshift bins centered at $z \sim 7$ and $z \sim 8$. The manner in which targets were assigned to each bin required careful consideration given the limited wavelength response of each instrument with respect to the photometric redshift likelihood distribution $P(z)$. Rather than binning on photometric redshift alone, we carefully considered the redshift range within which a null detection could be determined. If the median redshift for which a null detection could be determined was less than (greater than) $z = 7.5$, we placed the object in the $z \sim 7(8)$ bin.

5.3.1. Monte Carlo Simulation

To predict the number of detections expected in an IGM with no additional opacity to Lyα, we use a similar Monte Carlo method to that developed in Schenker et al. (2012). This simulation has three key inputs for each object: flux limits as a function of wavelength from the spectroscopic observations, which also take into account the night-sky emission; a photometric redshift probability distribution; and a prediction for the IGM unprocessed EW$_{\mathrm{Ly}\alpha}$ (and thus $f_{\mathrm{Ly}\alpha}$) distribution.

For objects observed in either this paper or Schenker et al. (2012), flux limits as a function of wavelength were calculated directly from the reduced spectra by computing the variance in an aperture of 7 Å spectral extent. For the data in Treu et al. (2012, 2013) we rescaled our flux limits from NIRSPEC and MOSFIRE, respectively, to match the quoted limits in the paper for each object. For Pentericci et al. (2011, 2014) and Ono et al.

Table 2
Summary of MOSFIRE Survey for Lyα

Survey	Observed	5σ Detections	Transmission Fraction (f)	$x_{\mathrm{H\,I}}$
This work MOSFIRE[a] M_{UV}	19	0	<0.47	>0.43
This work MOSFIRE[a] β	19	0	<0.38	>0.49
Composite $z \sim 7$	72	11	$0.52^{+0.12}_{-0.13}$	$0.39^{+0.08}_{-0.09}$
Composite $z \sim 8$	27	0	<0.20	>0.64
Composite $z \sim 7$[b]	72	11	$0.38^{+0.11}_{-0.09}$	$0.49^{+0.07}_{-0.07}$
Composite $z \sim 8$[b]	27	0	<0.15	>0.68

Notes. Monte Carlo results for transmission fraction, f, and $x_{\mathrm{H\,I}}$.
[a] The median redshift our sample is $z = 7.66$, and the median absolute magnitude is $M_{\mathrm{UV}} = -19.6$.
[b] For these results, where individual UV slopes are not available, we instead use individual values of M_{UV} to predict a value of β, which in turn is used to generate the IGM unprocessed EW$_{\mathrm{Ly\alpha}}$ distribution.

(2012) we did the same with our LRIS limits, as presented in Schenker et al. (2012), and DEIMOS limits, respectively. In the case of the Pentericci et al. (2011, 2014) samples, which used FORS2, we fit a polynomial to the limiting sensitivity curve presented in Fontana et al. (2010) and apply it to our LRIS limits, such that the sensitivity as a function of wavelength is well matched. We assume that the $z \geqslant 6$ sample contains the same modest 5% contamination as our $3 < z < 6$ control sample.

We used the published photometry from each paper (and our own here) in conjunction with our photometric redshift code described previously to determine a photometric redshift distribution for each object. The only exceptions are for the 10 and 4 objects from Ono et al. (2012) and Pentericci et al. (2011), respectively, for which either photometry or coordinates were not available. For these objects, we used the photometric redshift distribution for ground-based z-drops from Ouchi et al. (2009), from which the targets of Ono et al. (2012) were selected. Additionally, the photometric redshift distribution for both the four Pentericci et al. (2011) objects and the Ouchi et al. (2009) objects peak at $z \sim 6.85$, assuaging concerns about any systematics this could create.

Finally, for the objects in our new MOSFIRE survey, we generated the prediction for the IGM unprocessed EW$_{\mathrm{Ly\alpha}}$ distribution using the observed UV slope, as described in Section 4. Ideally, we would prefer to use this new method for all objects in the combined sample, in order to eliminate the potential bias of simply using M_{UV} as a predictor. However, with the exception of galaxies in the UDF 2012 field and CLASH lensed sample, the requisite three infrared photometric data points longward of the Lyman break essential for achieving an accurate measure of β are not available. Thus, for all other objects, we must predict EW$_{\mathrm{Ly\alpha}}$ as a function of M_{UV} from Treu et al. (2012), using the data presented in Stark et al. (2011). As an illustrative exercise, we also generated a prediction for the IGM unprocessed EW$_{\mathrm{Ly\alpha}}$ distribution of these objects using their M_{UV} to calculate a β derived from the M_{UV}–β relation at $z \sim 7$ from Bouwens et al. (2013). Though not plotted, these results are available in Table 2.

With these inputs, we conduct a Monte Carlo simulation. In each trial for a given object, we draw a redshift from the photometric redshift distribution and an Lyα EW from the EW$_{\mathrm{Ly\alpha}}$ distribution. From the observed photometry, this EW$_{\mathrm{Ly\alpha}}$ is converted to a flux. We then sample the spectroscopic flux limit at the redshift drawn to determine whether the emission line would be observed at $\geqslant 5\sigma$. This process is then repeated with $N = 10{,}000$ trials for each object.

5.3.2. Comparison between UV Slope and UV Luminosity Predictions

Before considering the total sample (i.e., including previous data from the literature), we compare the difference in the expected Lyα statistics for the high-redshift sample using either M_{UV} or β as the basis for predicting the IGM unprocessed EW$_{\mathrm{Ly\alpha}}$ distribution. Since only the UDF 2012 and CLASH surveys currently have accurate individual measurements of β, this comparison can only be done for the 13 targets from our recent MOSFIRE survey.

The number of expected detections is compared in Figure 7. Although hampered by limited statistics, the difference is still significant. Using β as a basis, we predict an average of 1.4 more detections than using M_{UV}. This difference represents an important correction of a systematic error in the prior $x_{\mathrm{Ly\alpha}}$ tests. Our new results show that the M_{UV} method, for this specific sample of 13 targets, underpredicts the incidence of IGM unprocessed Lyα emission, which results in an overestimate of the IGM transmission. The difference in predicted detections is dependent on the properties of the sample considered, but, as the objects probed from the both the UDF and CLASH are intrinsically faint, with blue UV slopes, it is not surprising that the difference is so great.

This change in predicted Lyα emission has important consequences for the transmission fraction of the IGM, implying a lower limit on the neutral fraction that is a factor of 0.16 larger. Clearly, for a given survey, our new β method for predicting the IGM unprocessed EW$_{\mathrm{Ly\alpha}}$ distribution has significant advantages and reduces a key systematic error. This will be especially relevant for spectroscopic follow-up of the *HST* Frontier Fields (GO:13496; PI: Lotz) and their parallels, given that these fields will have full coverage with the same four WFC3/IR filters used for the UDF.

5.3.3. Analyzing the Entire Sample

We can now combine the various subsamples, our own MOSFIRE survey, and that from diverse sources in the literature. The net result is a histogram of the number of 5σ detections overall. These histograms are displayed, for both the $z \sim 7$ and $z \sim 8$ samples in Figure 8; our new MOSFIRE survey is, by design, more effective in constraining the higher-redshift limits. We see evidence for a moderate decline in the Lyα fraction at $z \sim 7$ and a continued sharper decline at higher redshift. Note that our new spectroscopic confirmation at $z = 7.62$ is not included as a detection as it lies below the a priori 5σ flux limit.

As an illustration, we can convert these observational-based results to an IGM extinction of Lyα by adopting the model used
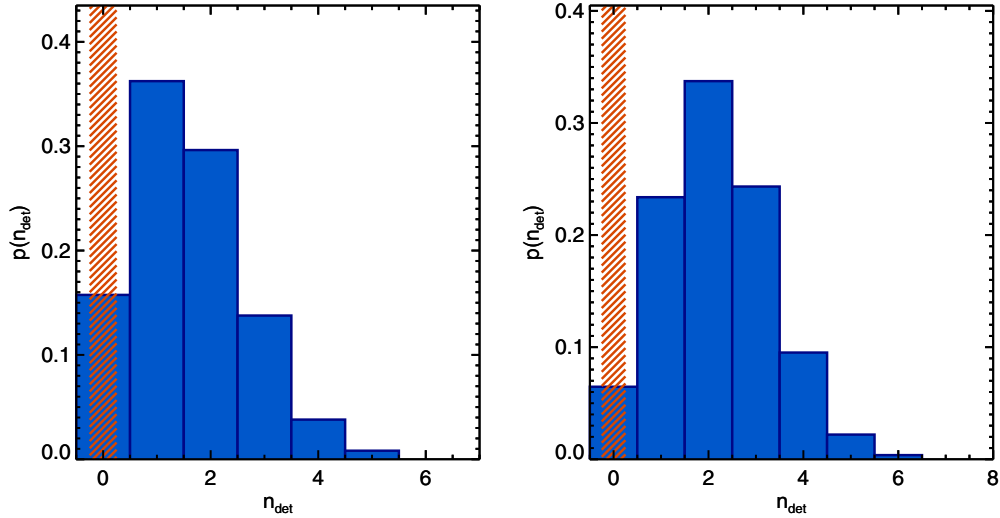
Figure 7. Left: predicted number of detections for our MOSFIRE survey using the $EW_{Ly\alpha}$ probability distribution from Treu et al. (2012), which uses M_{UV} as the predictor. The observed number of 0 detections is indicated by the red cross-hatch. Right: predicted number of detections for same survey, but using β as the predictor for $EW_{Ly\alpha}$, as outlined in Section 4. In this case, the average number of expected detections is increased by a factor of 0.4, highlighting the importance of using a model that accurately predicts the IGM unprocessed EW distribution. The equivalent upper limit on the transmission fraction is also decreased by a factor of 0.23.

(A color version of this figure is available in the online journal.)
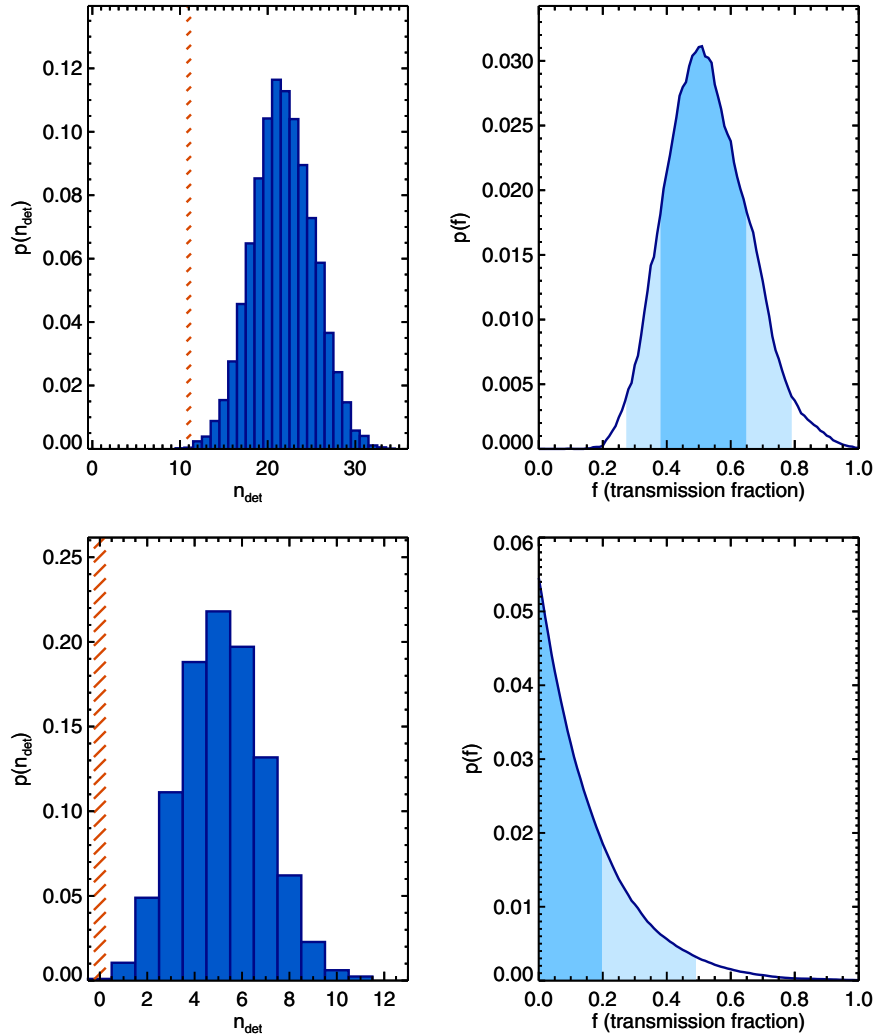
Figure 8. Results from our new MOSFIRE campaign, combined with data from the literature. Top left: histogram of expected 5σ detections of Lyα, computed using the Monte Carlo method described in the text for our $z \sim 7$ sample. The red cross-hatches denote the combined number of detections observed in all surveys. Top right: given our predicted and observed number of detections, the constraints on the average extinction fraction of Lyα, assuming a patchy opacity. Dark blue and light blue shadings encompass the 1σ and 2σ confidence intervals, respectively. Bottom left and right: same as above, but for our $z \sim 8$ sample.

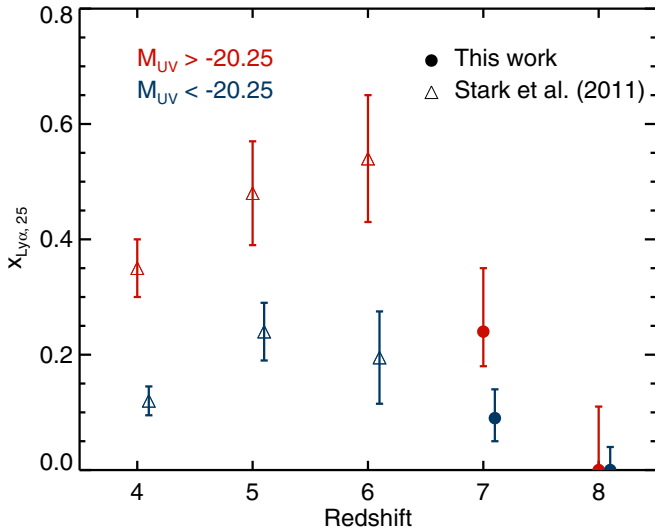(A color version of this figure is available in the online journal.)

Figure 9. Fraction of LBGs that display Lyα in emission at an EW $\geqslant$ 25 Å, plotted as a function of redshift. The values at $z = 7$ and 8 reflect differential measurements with the data at $z = 6$, as described in the text. Thus, these data points and errors are simply the convolution of the $x_{\mathrm{Ly}\alpha}$ pdf at $z = 6$ and the transmission fraction pdf at $z = 7$ and 8.

(A color version of this figure is available in the online journal.)

in Schenker et al. (2012) appropriate for patchy reionization. In this model, the IGM is partially opaque, such that Lyα escapes to the observer unattenuated from a fraction, f, of galaxies, while it is completely extinguished by the IGM in a fraction $1 - f$. Given the histogram of expected detections and the number actually observed, we infer a probability distribution for this transmission fraction, f. Formally, this involves the use of Bayes's theorem,

$$p(f|N_{\mathrm{obs}}) = \frac{p(N_{\mathrm{obs}}|f)p(f)}{\int_0^1 p(N_{\mathrm{obs}}|f)\,df}, \qquad (7)$$

where $p(N_{\mathrm{obs}}|f = 1)$ is derived from our Monte Carlo simulations and shown in Figure 8. In the case where f is not equal to one, we compute $p(N_{\mathrm{obs}}|f \neq 1)$ by modifying the Monte Carlo simulations such that a fraction of emission, $1 - f$, is extinguished. To compute $p(f|N_{\mathrm{obs}})$, the full probability distribution for f, we substitute the number of galaxies observed with Lyα emission at $>5\sigma$ and evaluate the right-hand side of the above equation. Our f and the ϵ_p used in Treu et al. (2012) represent the same quantity, though they use a more involved framework that takes into account the entire spectrum and observed line flux, while we concentrate only on the number of 5σ detections.

At $z \sim 7$, we find $f = 0.52^{+0.12}_{-0.13}$, and at $z \sim 8$, a 1σ upper limit of $f < 0.20$. The full results can be found in Table 2 and are plotted along with the lower redshift data on $x_{\mathrm{Ly}\alpha}$ in Figure 9.

Discussing the uncertainties in the transformation from transmission fraction to $x_{\mathrm{H\,I}}$ is beyond the scope of the present paper. However, clearly this conversion is dependent on a number of physical parameters, some internal to the galaxy, and others from the IGM state itself. These include the velocity offset of Lyα from the galaxy's systemic velocity (e.g., Hashimoto et al. 2013; Schenker et al. 2013a), the ionizing photon escape fraction (Dijkstra et al. 2014), and the possible presence of optically thick absorption systems (Bolton & Haehnelt 2013). Until the theoretical models converge and/or observations of these key quantities are available, absolute measures of the neutral fraction will still be subject to systematic errors. Nonetheless, we have demonstrated substantial observational progress with our new survey and improved methodology, reducing one of the key

systematic errors. We use the models of McQuinn et al. (2007), which calculated the visibility of Lyα emitters in a 186 Mpc side length cosmological simulation as a function of $x_{\mathrm{H\,I}}$. Applying these to our visibility data to provide an estimate of $x_{\mathrm{H\,I}}$, we find $x_{\mathrm{H\,I}} = 0.39^{+0.08}_{-0.09}$ at $z \sim 7$ and $x_{\mathrm{H\,I}} > 0.64$ at $z \sim 8$.

6. CONCLUSIONS

Using our sample of 451 $3 < z < 6$ spectroscopically followed-up LBGs, we demonstrate an improved correlation between the ultraviolet continuum slope of a galaxy, β, and its Lyα emission strength. Given the availability of deep WFC3 photometry for both the GOODS-N and GOODS-S fields, this progress follows measurements for many individual galaxies in this redshift range, rather than via stacked or averaged UV slopes, as in earlier work (Shapley et al. 2003; Stark et al. 2010).

We demonstrate that this correlation with the presence of Lyα is stronger and more physically motivated than that based on the UV luminosity and thus provides a natural basis for an improved model for the Lyα fraction test, now widely used to measure the evolving neutrality of the $z > 6.5$ IGM. We demonstrate the benefits of this new model using a new MOSFIRE spectroscopic survey of $7 < z < 8$ targets from the UDF 2012 catalog and CLASH lensing survey and combine this with data at these redshifts already published in the literature. As a result, we present the implications of the most comprehensive search for Lyα emission at $z \simeq 8$ to date, confirming once again important evidence that cosmic reionization ended at redshifts $z \simeq 6.5$.

As a by-product, we also present a promising 4.0σ detection of Lyα in a galaxy at $z = 7.62$, possibly the most distant spectroscopically confirmed galaxy.

APPENDIX A

MODELS FOR $p(\mathrm{EW}|\beta)$

The maximum likelihoods inferred from each of the four distributions are noted in Table 3. These results demonstrate that the lognormal distribution provides the best fit to the available data—its likelihood surpasses that of any other model by two orders of magnitude. Thus, we use this distribution as the basis for the more general form of $p(\mathrm{EW}|\beta)$ we consider next.

APPENDIX B

RESULTS OF FULL MODELING PROCEDURE

For reference, and so that they are available for use in future work, we list here the final values for our generalized lognormal fit to the $\mathrm{EW}_{\mathrm{Ly}\alpha}$ distribution at $3 < z < 6$. They are $\mu_a = 2.875^{+0.125}_{-0.25}$, $\mu_s = -1.125 \pm 0.25$, $\sigma = 1.3 \pm 0.1$, and $A_{\mathrm{em}} = 1.0^{+0.0}_{-0.1}$. We also provide a plot of the posterior probability distribution in Figure 10, so the reader is able to appreciate the sometimes non-negligible covariances between parameters.
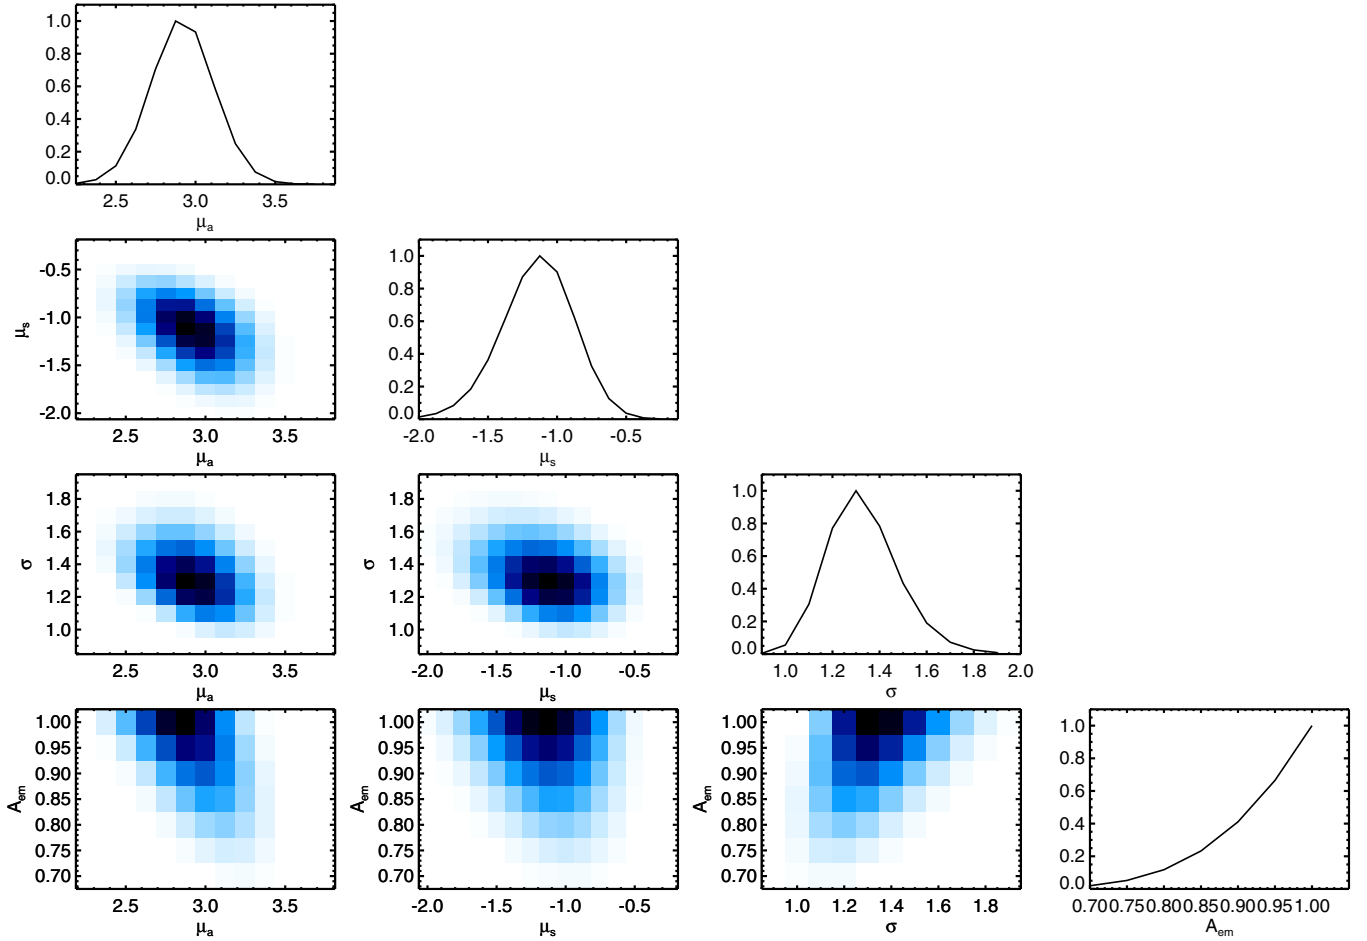
Figure 10. Posterior probability distribution for our full model, $p(\mathrm{EW}_{\mathrm{Ly}\alpha}|\beta)$. Shaded plots represent the posterior pdf marginalized over all but the two variables labeling the axes, while line plots are marginalized over all but one variable. Thus, the one-dimensional pdfs for each variable, from which we quote our error bars, can be read off along the diagonal.

(A color version of this figure is available in the online journal.)

Table 3
Lyα Functional Forms

Distribution Name	Equation	Free Parameters	Reference	$\log_{10}$ Max Likelihood
Lognormal	$A_{\mathrm{em}}\frac{1}{\sqrt{2\pi}\sigma\mathrm{EW}}\exp(-(\ln(\mathrm{EW})-\mu)^2/2\sigma^2)$	$A_{\mathrm{em}}, \mu, \sigma$	This work	-75.4
Half-Gaussian	$\frac{1}{\sqrt{2\pi}\sigma}\exp(-(\mathrm{EW}-\mu)^2/2\sigma^2)$	A_{em}, σ	Treu et al. (2012)	-80.7
w/high-EW tail	$A_{\mathrm{em},g}\frac{1}{\sqrt{2\pi}\sigma}\exp(-(\mathrm{EW}-\mu)^2/2\sigma^2)+A_{\mathrm{em},c}$	$A_{\mathrm{em},g}, A_{\mathrm{em},c}, \mu, \sigma$	Pentericci et al. (2011)	-80.8
Declining exponential	$A_{\mathrm{em}}\exp(-\mathrm{EW}/\mathrm{EW}_0)$	$A_{\mathrm{em}}, \mathrm{EW}_0$	Dijkstra et al. (2011)	-77.9

Notes. List of the mathematical distributions used to fit the $\mathrm{EW}_{\mathrm{Ly}\alpha}$ distribution at $\beta \sim -2.3$, and the calculated maximum likelihoods. Having a maximum likelihood two orders of magnitude greater than any other distribution considered demonstrates that the lognormal distribution provides the best fit.

REFERENCES

Bertin, E., & Arnouts, S. 1996, A&AS, 117, 393
Bertin, E., Mellier, Y., Radovich, M., et al. 2002, in ASP Conf. Ser. 281, Astronomical Data Analysis Software and Systems XI, ed. D. A. Bohlender, D. Durand, & T. H. Handley (San Francisco, CA: ASP), 228
Bolton, J. S., & Haehnelt, M. G. 2013, MNRAS, 429, 1695
Bouwens, R. J., Illingworth, G. D., Oesch, P. A., et al. 2011, ApJ, 737, 90
Bouwens, R. J., Illingworth, G. D., Oesch, P. A., et al. 2013, arXiv:1306.2950
Bradley, L. D., Zitrin, A., Coe, D., et al. 2014, ApJ, 792, 76
Bromm, V. 2013, RPPh, 76, 112901
Bruzual, G., & Charlot, S. 2003, MNRAS, 344, 1000
Calzetti, D., Armus, L., Bohlin, R. C., et al. 2000, ApJ, 533, 682
Calzetti, D., Kinney, A. L., & Storchi-Bergmann, T. 1994, ApJ, 429, 582
Castellano, M., Fontana, A., Paris, D., et al. 2010, A&A, 524, A28
Coe, D., Benítez, N., Sánchez, S. F., et al. 2006, AJ, 132, 926

Dijkstra, M., Mesinger, A., & Wyithe, J. S. B. 2011, MNRAS, 414, 2139
Dijkstra, M., Wyithe, S., Haiman, Z., Mesinger, A., & Pentericci, L. 2014, MNRAS, 440, 3309
Dunlop, J. S., Rogers, A. B., McLure, R. J., et al. 2013, MNRAS, 432, 3520
Ellis, R. S., McLure, R. J., Dunlop, J. S., et al. 2013, ApJL, 763, L7
Finkelstein, S. L., Papovich, C., Dickinson, M., et al. 2013, Natur, 502, 524
Finkelstein, S. L., Papovich, C., Salmon, B., et al. 2012, ApJ, 756, 164
Fontana, A., Vanzella, E., Pentericci, L., et al. 2010, ApJL, 725, L205
Guo, Y., Ferguson, H. C., Giavalisco, M., et al. 2013, ApJS, 207, 24
Hashimoto, T., Ouchi, M., Shimasaku, K., et al. 2013, ApJ, 765, 70
Hinshaw, G., Larson, D., Komatsu, E., et al. 2013, ApJS, 208, 19
Kass, R. E., & Raftery, A. E. 1995, J. Am. Stat. Assoc., 430, 773
Koekemoer, A. M., Ellis, R. S., McLure, R. J., et al. 2013, ApJS, 209, 3
McLean, I. S., Steidel, C. C., Epps, H. W., et al. 2012, Proc. SPIE, 8446, 84460J
McLure, R. J., Dunlop, J. S., Bowler, R. A. A., et al. 2013, MNRAS, 432, 2696
McLure, R. J., Dunlop, J. S., Cirasuolo, M., et al. 2010, MNRAS, 403, 960

McQuinn, M., Hernquist, L., Zaldarriaga, M., & Dutta, S. 2007, MNRAS, 381, 75
Meurer, G. R., Heckman, T. M., & Calzetti, D. 1999, ApJ, 521, 64
Miralda-Escudé, J., Haehnelt, M., & Rees, M. J. 2000, ApJ, 530, 1
Oesch, P. A., Bouwens, R. J., Illingworth, G. D., et al. 2012, ApJ, 759, 135
Oesch, P. A., Bouwens, R. J., Illingworth, G. D., et al. 2014, ApJ, 786, 108
Oke, J. B., & Gunn, J. E. 1983, ApJ, 266, 713
Ono, Y., Ouchi, M., Mobasher, B., et al. 2012, ApJ, 744, 83
Ouchi, M., Mobasher, B., Shimasaku, K., et al. 2009, ApJ, 706, 1136
Ouchi, M., Shimasaku, K., Furusawa, H., et al. 2010, ApJ, 723, 869
Pentericci, L., Fontana, A., Vanzella, E., et al. 2011, ApJ, 743, 132
Pentericci, L., Vanzella, E., Fontana, A., et al. 2014, arXiv:1403.5466
Robertson, B. E., Furlanetto, S. R., Schneider, E., et al. 2013, ApJ, 768, 71
Rogers, A. B., McLure, R. J., Dunlop, J. S., et al. 2014, MNRAS, 440, 3714
Santos, M. R. 2004, MNRAS, 349, 1137
Schenker, M. A., Ellis, R. S., Konidaris, N. P., & Stark, D. P. 2013a, ApJ, 777, 67
Schenker, M. A., Robertson, B. E., Ellis, R. S., et al. 2013b, ApJ, 768, 196
Schenker, M. A., Stark, D. P., Ellis, R. S., et al. 2012, ApJ, 744, 179
Shapley, A. E., Steidel, C. C., Pettini, M., & Adelberger, K. L. 2003, ApJ, 588, 65
Stark, D. P., Ellis, R. S., Bunker, A., et al. 2009, ApJ, 697, 1493
Stark, D. P., Ellis, R. S., Chiu, K., Ouchi, M., & Bunker, A. 2010, MNRAS, 408, 1628
Stark, D. P., Ellis, R. S., & Ouchi, M. 2011, ApJL, 728, L2
Treu, T., Schmidt, K. B., Trenti, M., Bradley, L. D., & Stiavelli, M. 2013, ApJL, 775, L29
Treu, T., Trenti, M., Stiavelli, M., Auger, M. W., & Bradley, L. D. 2012, ApJ, 747, 27
Vanzella, E., Giavalisco, M., Dickinson, M., et al. 2009, ApJ, 695, 1163