

## IGM CONSTRAINTS FROM THE SDSS-III/BOSS DR9 Ly $\alpha$ FOREST TRANSMISSION PROBABILITY DISTRIBUTION FUNCTION

KHEE-GAN LEE<sup>1,2</sup>, JOSEPH F. HENNAWI<sup>1</sup>, DAVID N. SPERTEL<sup>2</sup>, DAVID H. WEINBERG<sup>3</sup>, DAVID W. HOGG<sup>4</sup>,  
MATTEO VIEL<sup>5,6</sup>, JAMES S. BOLTON<sup>7</sup>, STEPHEN BAILEY<sup>8</sup>, MATTHEW M. PIERI<sup>9</sup>, WILLIAM CARITHERS<sup>8</sup>,  
DAVID J. SCHLEGEL<sup>8</sup>, BRITT LUNDGREN<sup>10</sup>, NATHALIE PALANQUE-DELABROUILLE<sup>11</sup>, NAO SUZUKI<sup>12</sup>,  
DONALD P. SCHNEIDER<sup>13,14</sup>, AND CHRISTOPHE YÈCHE<sup>11</sup>

<sup>1</sup> Max Planck Institute for Astronomy, Königstuhl 17, D-69117 Heidelberg, Germany; lee@mpia.de

<sup>2</sup> Department of Astrophysical Sciences, Princeton University, Princeton, NJ 08544, USA

<sup>3</sup> Department of Astronomy and Center for Cosmology and Astro-Particle Physics, Ohio State University, Columbus, OH 43210, USA

<sup>4</sup> Center for Cosmology and Particle Physics, New York University, 4 Washington Place, Meyer Hall of Physics, New York, NY 10003, USA

<sup>5</sup> INAF, Osservatorio Astronomico di Trieste, Via G. B. Tiepolo 11, I-34131 Trieste, Italy

<sup>6</sup> INFN/National Institute for Nuclear Physics, Via Valerio 2, I-34127 Trieste, Italy

<sup>7</sup> School of Physics and Astronomy, University of Nottingham, University Park, Nottingham NG7 2RD, UK

<sup>8</sup> E.O. Lawrence Berkeley National Lab, 1 Cyclotron Road, Berkeley, CA 94720, USA

<sup>9</sup> Institute of Cosmology & Gravitation, University of Portsmouth, Dennis Sciama Building, Portsmouth PO1 3FX, UK

<sup>10</sup> Department of Astronomy, University of Wisconsin, Madison, WI 53706, USA

<sup>11</sup> CEA, Centre de Saclay, Irfu/SPP, F-91191 Gif-sur-Yvette, France

<sup>12</sup> Kavli Institute for the Physics and Mathematics of the Universe (IPMU), The University of Tokyo, Kashiwano-ha 5-1-5, Kashiwa-shi, Chiba, Japan

<sup>13</sup> Department of Astronomy and Astrophysics, The Pennsylvania State University, University Park, PA 16802, USA

<sup>14</sup> Institute for Gravitation and the Cosmos, The Pennsylvania State University, University Park, PA 16802, USA

Received 2014 May 5; accepted 2014 December 4; published 2015 January 29

### ABSTRACT

The Ly $\alpha$  forest transmission probability distribution function (PDF) is an established probe of the intergalactic medium (IGM) astrophysics, especially the temperature–density relationship of the IGM. We measure the transmission PDF from 3393 Baryon Oscillations Spectroscopic Survey (BOSS) quasars from Sloan Digital Sky Survey Data Release 9, and compare with mock spectra that include careful modeling of the noise, continuum, and astrophysical uncertainties. The BOSS transmission PDFs, measured at  $\langle z \rangle = [2.3, 2.6, 3.0]$ , are compared with PDFs created from mock spectra drawn from a suite of hydrodynamical simulations that sample the IGM temperature–density relationship,  $\gamma$ , and temperature at mean density,  $T_0$ , where  $T(\Delta) = T_0 \Delta^{\gamma-1}$ . We find that a significant population of partial Lyman-limit systems (LLSs) with a column-density distribution slope of  $\beta_{\text{pLLS}} \sim -2$  are required to explain the data at the low-transmission end of transmission PDF, while uncertainties in the mean Ly $\alpha$  forest transmission affect the high-transmission end. After modeling the LLSs and marginalizing over mean transmission uncertainties, we find that  $\gamma = 1.6$  best describes the data over our entire redshift range, although constraints on  $T_0$  are affected by systematic uncertainties. Within our model framework, isothermal or inverted temperature–density relationships ( $\gamma \leq 1$ ) are disfavored at a significance of over  $4\sigma$ , although this could be somewhat weakened by cosmological and astrophysical uncertainties that we did not model.

*Key words:* intergalactic medium – large-scale structure of universe – methods: data analysis – quasars: absorption lines – quasars: emission lines – techniques: spectroscopic

### 1. INTRODUCTION

Remarkably soon after the discovery of the first high-redshift ( $z_{\text{qso}} \gtrsim 2$ ) quasars (Schmidt 1965), Gunn & Peterson (1965) realized that the amount of resonant Lyman- $\alpha$  (Ly $\alpha$ ) scattering off neutral hydrogen structures observed in the spectra of these quasars could be used to constrain the state of the intergalactic medium (IGM) at high-redshifts: they deduced that the hydrogen in the IGM had to be highly photo-ionized (neutral fractions of  $n_{\text{H}1}/n_{\text{H}} < 10^{-4}$ ) and hot (temperatures,  $T > 10^4$  K).

Lynds (1971) then discovered that this Ly $\alpha$  absorption could be separated into discrete absorption lines, i.e., the Ly $\alpha$  “forest.” Over the next two decades, it was recognized that the individual Ly $\alpha$  forest lines have Voigt absorption profiles corresponding to Doppler-broadened systems with  $T \sim 1 - 3 \times 10^4$  K (see, e.g., Rauch et al. 1992; Ricotti et al. 2000; Schaye et al. 2000; McDonald et al. 2001; Tytler et al. 2004; Lidz et al. 2010; Becker et al. 2011) and neutral column densities of  $N \sim 10^{13} - 10^{17} \text{ cm}^{-2}$  (Petitjean et al. 1993; Penton et al. 2000; Janknecht et al. 2006; Rudie et al. 2013), and increasingly precise measurements of mean Ly $\alpha$  forest transmission have been carried out (Theuns

et al. 2002; Bernardi et al. 2003; Faucher-Giguère et al. 2008; Becker et al. 2013). However, the exact physical nature of these absorbers was unclear for many years (see Rauch 1998, for a historical review of the field).

Beginning in the 1990s, detailed hydrodynamical simulations of the IGM led to the current physical picture of the Ly $\alpha$  forest arising from baryons in the IGM, which trace fluctuations in the dark matter (DM) field induced by gravitational collapse, in ionization balance with a uniform ultraviolet ionizing background (see, e.g., Cen et al. 1994; Miralda-Escudé et al. 1996; Croft et al. 1998; Davé et al. 1999; Theuns et al. 1998). A physically motivated analytic description of this picture is the fluctuating Gunn–Peterson approximation (Croft et al. 1998; Hui et al. 1997), in which the Ly $\alpha$  optical depth,  $\tau$ , scales with underlying matter density,  $\rho$ , through a polynomial relationship:

$$\tau \propto \frac{T^{-0.7}}{\Gamma} \Delta^2 \propto \frac{T_0^{-0.7}}{\Gamma} \Delta^{2-0.7(\gamma-1)}, \quad (1)$$

where  $\Gamma$  is the background photoionization rate, and  $\Delta \equiv \rho/\langle \rho \rangle$  is the matter density relative to the mean density of the universe

at the given epoch. In the second proportionality above, we have made the assumption that the local temperature of the gas has a polynomial relationship with the local density,

$$T = T_0 \Delta^{\gamma-1}, \quad (2)$$

where  $T_0$  is the gas temperature at mean density and  $\gamma$  parameterizes the temperature–density relation, which encodes the thermal history of the IGM (e.g., Hui & Gnedin 1997; Schaye et al. 1999; Ricotti et al. 2000; McDonald et al. 2001; Hui & Haiman 2003; see Meiksin 2009 for a detailed overview on the relevant physics).

Over the past decade and a half, the 2000–2008 Sloan Digital Sky Survey (SDSS-I and -II; York et al. 2000; Stoughton et al. 2002, <http://www.sdss.org>) spectroscopic data has represented a dramatic improvement in the statistical power available to Ly $\alpha$  forest studies: McDonald et al. (2006) measured the one-dimensional (1D) Ly $\alpha$  forest transmission power spectrum from  $\approx 3000$  SDSS quasar sightlines. This measurement was used to place significant constraints on cosmological parameters and large-scale structure (see, e.g., McDonald et al. 2005b; Seljak et al. 2005; Viel & Haehnelt 2006).

The McDonald et al. (2006) quasar sample, which in its time represented a  $\sim 100$  increase in sample size over previous data sets, is superseded by the Baryon Oscillations Sky Survey (BOSS; part of SDSS-III; Eisenstein et al. 2011; Dawson et al. 2013) quasar survey. This spectroscopic survey, which operated between fall 2009 and spring 2014, is aimed at taking spectra of  $\sim 150,000$   $z_{\text{qso}} \gtrsim 2.2$  quasars (Dawson et al. 2013) with the goal of constraining dark energy at  $z > 2$  using transverse correlations of Ly $\alpha$  forest absorption (see, e.g., Slosar et al. 2011) to measure the baryon acoustic oscillation (BAO) scale.<sup>15</sup> At time of writing, the full BOSS survey is complete, with  $\sim 170,000$  high-redshift quasars observed, although this paper is based on the earlier sample of  $\sim 50,000$  BOSS quasars from SDSS Data Release 9 (DR9; Ahn et al. 2012; Pâris et al. 2012; Lee et al. 2013).

The quality of the individual BOSS Ly $\alpha$  forest spectra might appear at first glance inadequate for studying the astrophysics of the IGM, that have to date been carried out largely with high-resolution, high-signal-to-noise ratio (S/N) spectra: the typical BOSS spectrum has  $S/N \sim 2 \text{ pixel}^{-1}$ ,<sup>16</sup> since the BAO analysis is optimized with large numbers of low S/N sightlines, densely sampled on the sky (McDonald & Eisenstein 2007; McQuinn & White 2011). It is therefore interesting to ask whether it is possible to model the various instrumental and astrophysical effects seen in the BOSS Ly $\alpha$  forest spectra, to sufficient accuracy level to exploit the unprecedented statistical power.

In this paper, we will measure the probability distribution function (PDF) of the Ly $\alpha$  forest transmission,  $F \equiv \exp(-\tau)$ , from BOSS. This one-point statistic, which was first studied by Jenkins & Ostriker (1991), is sensitive to astrophysical parameters such as the amplitude of matter fluctuations and the thermal history of the IGM. However, the transmission<sup>17</sup> PDF is also highly sensitive to effects such as pixel noise level, resolution of the spectra, and systematic uncertainties

in the placement of the quasar continuum level, especially in moderate resolution spectra such as SDSS or BOSS. Desjacques et al. (2007) studied the transmission PDF from a sample of  $\sim 3500$  Ly $\alpha$  forest spectra from SDSS Data Release 3 (Abazajian et al. 2005). Using mock spectra generated from a log-normal model of the Ly $\alpha$  forest with parameters tuned to reproduce high-resolution, high-S/N spectra, they fitted for the estimated pipeline noise level and continuum-fitting errors in the SDSS spectra. They concluded that the noise levels reported by the SDSS pipeline were underestimated by  $\sim 10\%$ , consistent with the findings of McDonald et al. (2006). They also found that the quasar continuum-level was systematically lower by  $\sim 10\%$  in comparison with a power law extrapolated from redward of the quasar Ly $\alpha$  line, with a rms variance of  $\sim 20\%$ , although certain aspects of their study, e.g., the noise modeling and quasar continuum model, were rather crude.

We intend to take an approach distinct from that of Desjacques et al. (2007): instead of treating the noise and continuum as free parameters, we will attempt to measure the BOSS Ly $\alpha$  forest transmission PDF using a rigorous treatment of the noise and continuum-fitting, and then adopt a “forward-modeling” approach of trying to model the various instrumental effects as accurately as possible in mock spectra generated from detailed hydrodynamical simulations. Using the raw individual exposures and calibration data from BOSS, we will first implement a novel probabilistic method for co-adding the exposures, which will yield more accurate noise estimates and also enable self-consistent noise modeling in mock spectra. Similarly, we will use a new method for continuum estimation called mean flux regulated/principal component analysis (MF-PCA; Lee et al. 2012). This technique provides unprecedented continuum accuracy for noisy Ly $\alpha$  forest spectra:  $< 10\%$  rms errors for  $S/N \sim 2$  and  $< 5\%$  rms errors for  $S/N \gtrsim 5$  spectra.

On the modeling side, we will use the detailed hydrodynamical IGM simulations of Viel et al. (2013a) as a basis. The mock spectra are then smoothed to BOSS resolution, and have Lyman-limit systems (LLSs) and metal contamination added, followed by the introduction of pixel noise based on our improved noise estimates. We will then self-consistently introduce continuum errors by applying our continuum-estimation procedure to the mock spectra.

With the increase in statistical power from the sheer number of BOSS spectra, and our improved modeling of the noise and continuum, we expect to significantly reduce the errors on the measured transmission PDF in comparison with Desjacques et al. (2007). This should enable us to place independent constraints on the shape of the underlying transmission PDF, and the thermal history of the IGM as parameterized by the power-law temperature–density relation,  $\gamma$  and  $T_0$ .

The IGM temperature–density relationship is a topic of recent interest, as Bolton et al. (2008) and Viel et al. (2009) have found evidence of an inverted temperature–density relation,  $\gamma < 1$ , at  $z \sim 2-3$  from the transmission PDF from high-resolution, high-S/N Ly $\alpha$  forest spectra (Kim et al. 2007), implying that IGM voids are hotter than overdensities. This result is in contrast with theoretical expectations of  $\gamma \approx 1.6$  (Miralda-Escudé & Rees 1994; Hui & Gnedin 1997; Theuns et al. 1998; Hui & Haiman 2003), which arises from the balance between adiabatic cooling in the lower-density IGM and photoheating in the higher-density regions. Even inhomogeneous He II reionization, which is expected to flatten the IGM temperature–density relation (see, e.g., Furlanetto & Oh 2008; Bolton et al. 2009; McQuinn et al. 2009), is insufficient to account for the extremely low values

<sup>15</sup> There is also a simultaneous effort to observe  $\sim 1.5$  million luminous red galaxies, to measure the BAO at  $z \sim 0.5$ . See, e.g., Anderson et al. (2014).

<sup>16</sup> All spectral signal-to-noise ratios quoted in this paper are per  $69 \text{ km s}^{-1}$  SDSS/BOSS pixel unless noted otherwise.

<sup>17</sup> The Ly $\alpha$  forest transmitted flux fraction is sometimes also referred to as “flux” in the literature; but we do however use the variable  $F$  to refer to this quantity.

of  $\gamma \sim 0.5$  estimated by the aforementioned authors (although inversions could occur at higher densities, see, e.g., Meiksin & Tittle 2012).

Indeed, earlier papers studying the temperature–density relationship using either the transmission PDF (McDonald et al. 2001) or by measuring the Doppler parameters and hydrogen column densities of individual forest absorbers (the so-called  $b - N_{\text{H I}}$  relation, e.g., Schaye et al. 1999; Ricotti et al. 2000; Rudie et al. 2012) have found no evidence of an inverted  $\gamma$ . In recent years, the decay of blazar gamma rays via plasma instabilities (Broderick et al. 2012; Chang et al. 2012; although see Sironi & Giannios 2014) has been invoked as a possible mechanism for supplying the heat necessary to flatten  $\gamma$  to the observed levels (Puchwein et al. 2012).

It would be desirable to perform an independent re-analysis of high-resolution data taking into account continuum-fitting bias (Lee 2012), to place these claims on a firmer footing. However, Lee & Spergel (2011) have argued that the complete SDSS DR7 (Abazajian et al. 2009) Ly $\alpha$  forest data set could have sufficient statistical power to place interesting constraints on  $\gamma$ , even assuming continuum-fitting errors at the  $\sim 10\%$  rms level. Therefore, with the current BOSS data, we hope to model noise and resolution, as well as astrophysical systematics, at a sufficient precision to place interesting constraints on the IGM thermal history.

This paper is organized as follows: we first give a broad overview of the BOSS Ly $\alpha$  forest data set, followed by our measurement of the BOSS transmission PDF with detailed descriptions of our method of combining multiple raw exposures and continuum estimation. We then discuss how we include various instrumental and astrophysical effects into our modeling of the transmission PDF, starting with hydrodynamical simulations. The model transmission PDF is then compared with the observed PDF to obtain constraints on the thermal parameters governing the IGM.

## 2. DATA

### 2.1. Summary of BOSS

BOSS (Dawson et al. 2013) is part of SDSS-III (Eisenstein et al. 2011; the other surveys are SEGUE-2, MARVELS, and APOGEE). The primary goal of the survey is to carry out precision BAO measurements at  $z \sim 0.5$  and  $z \sim 2.5$ , from the luminous red galaxy distribution and Ly $\alpha$  forest absorption field, respectively (see, e.g., Anderson et al. 2014; Busca et al. 2013; Slosar et al. 2013). Its eventual goal is to obtain spectra of  $\sim 1.5$  million luminous red galaxies and  $\sim 170,000$   $z > 2.15$  quasars over 4.5 yr of operation.

BOSS is conducted on upgraded versions of the twin SDSS spectrographs (Smee et al. 2013) mounted on the 2.5 m Sloan telescope (Gunn et al. 2006) at the Apache Point Observatory, NM. One thousand optical fibers mounted on a plug-plate at the focal plane (spanning a  $3^\circ$  field of view) feed the incoming flux to the two identical spectrographs, of which 160–200 fibers per plate are allocated to quasar targets (see Ross et al. 2012; Bovy et al. 2011 for a detailed description of the quasar target selection). Both spectrographs split the light into blue and red cameras that cover 3610–10140 Å, with the dichroic overlap region occurring at around 6000 Å. The resolving power  $R \equiv \lambda/\Delta\lambda$  ranges from 1300 at the blue end to 2600 at the red end.

Each plate is observed for sufficiently long to achieve the S/N requirements set by the survey goals; typically, 5 individual exposures of 15 minutes are taken. The data are processed,

calibrated, and combined into co-added spectra by the “idl-spec2d” pipeline, followed by a pipeline that operates on the 1D spectra to classify objects and assign redshifts (Bolton et al. 2012). However, as described later in this paper, we will generate our own co-added spectra from the individual exposures and other intermediate data products.

### 2.2. Data Cuts

In this paper we use data from the publicly available SDSS Data Release 9 (DR9; Ahn et al. 2012). This includes 87,822 quasars at all redshifts that have been confirmed by visual inspection as described in Pâris et al. (2012). In Lee et al. (2013), we have defined a further subset of 54,468 quasars with  $z_{\text{qso}} \geq 2.15$  that are suitable for Ly $\alpha$  forest analysis, and have provided in individual FITS files for each quasar various products such as sky masks, masks for damped Ly $\alpha$  absorbers (DLAs), noise corrections, and continua; these are designed to ameliorate systematics in the BOSS spectra and aid in Ly $\alpha$  forest analysis (see Table 1 in Lee et al. 2013 for a full listing). While we use this Lee et al. (2013) catalog as a starting point, in this paper we will generate our own custom co-added spectra and noise estimates.

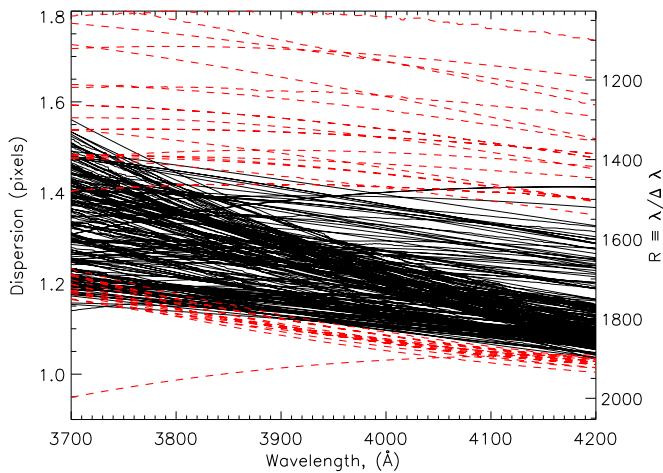
The typical S/N of the BOSS Ly $\alpha$  forest quasars is low:  $\langle S/N \rangle \approx 2$  pixel $^{-1}$  within the Ly $\alpha$  forest; this criterion is driven by a strategy to ensure a large number of sightlines over a large area in order to optimize the 3D Ly $\alpha$  forest BAO analysis (McDonald & Eisenstein 2007; McQuinn & White 2011), rather than increasing the S/N in individual spectra. However, for our analysis we wish to select a subset of BOSS Ly $\alpha$  forest sightlines with reasonably high S/N in order to reduce the sensitivity of our PDF measurement to inaccuracies in our modeling of the noise and continuum of the BOSS spectra. We therefore make a cut on the S/N, including only sightlines that have a median of  $\langle S/N \rangle \geq 6$  pixel $^{-1}$  within the Ly $\alpha$  forest,<sup>18</sup> defined with respect to the pipeline noise estimate (see Lee et al. 2013)—this selects only  $\sim 10\%$  of the spectra with the highest S/N. The 1041–1185 Å Ly $\alpha$  forest region of each quasar must also include at least 30 pixels ( $\Delta v = 2071$  km s $^{-1}$ ) within one of our absorption redshift bins of  $\langle z \rangle = 2.3$ ,  $\langle z \rangle = 2.6$ , and  $\langle z \rangle = 3.0$ , with bin widths of  $\Delta z = 0.3$  (see Section 3).

We discard spectra with identified DLAs in the sightline, as listed in the “DLA Concordance Catalog” used in the Lee et al. (2013) sample. This DLA catalog (W. Carithers 2014, in preparation) includes objects with column densities  $N_{\text{H I}} > 10^{20}$  cm $^{-2}$ ; however, the completeness of this catalog is uncertain below  $N_{\text{H I}} = 10^{20.3}$  cm $^{-2}$ . We therefore discard only sightlines containing DLAs with  $N_{\text{H I}} \geq 10^{20.3}$  cm $^{-2}$ , and take into account lower column-density absorbers in our subsequent modeling of mock spectra. At the relatively high S/N that we will work with (see below), the detection efficiency of DLAs is essentially 100% (see, e.g., Prochaska et al. 2005; Noterdaeme et al. 2012) and thus we expect our rejection of  $N_{\text{H I}} \geq 10^{20.3}$  cm $^{-2}$  DLAs to be quite thorough.

Measurements of the Ly $\alpha$  forest transmission PDF are known to be sensitive to the continuum estimate (Lee 2012), but in this paper we use an automated continuum-fitter, MF-PCA (Lee 2012), that is less susceptible to biases introduced by manual continuum estimation. Moreover, unlike the laborious process of manually fitting continua on high-resolution spectra, the automated continuum estimation can be used to explore various biases in continuum estimation. For this purpose, we will use the

<sup>18</sup> Defined as the 1041–1185 Å region in the quasar rest-frame.






**Figure 1.** Wavelength dispersions,  $\sigma_{\text{disp}}$ , for 236 BOSS quasar spectra randomly selected from the  $\langle z \rangle = 2.3$ ,  $6 < S/N < 8$  PDF bin. The ordinate axis on the right shows the equivalent spectral resolution,  $R \equiv \lambda/\Delta\lambda$ . The dashed red lines are objects that have been discarded from the analysis because they are outliers in spectral dispersion.

same MF-PCA continuum estimation used in Lee et al. (2013), albeit with minor modifications as described in Section 3.2. We select only quasars that appear to be well described by the continuum basis templates, based on the goodness-of-fit to the quasar spectrum redward of Ly $\alpha$ . This is flagged by the variable CONT\_FLAG = 1 as listed in the Lee et al. (2013) catalog (see Table 3 in that paper). Broad absorption line (BAL) quasars, the continua of which are difficult to estimate due to broad intrinsic absorption troughs, have already been discarded from the Lee et al. (2013) sample.

Another consideration is that the shape of the transmission PDF is affected by the resolution of the spectrum, especially since the BOSS spectrographs do not resolve the Ly $\alpha$  forest. The exact spectral resolution of a BOSS spectrum at a given wavelength varies as a function of both observing conditions and row position on the BOSS CCDs. The BOSS pipeline reports the wavelength dispersion at each pixel,  $\sigma_{\text{disp}}$ , in units of the co-added wavelength pixel size (binned such that  $\ln(10) \Delta(\lambda)/\lambda = 10^{-4}$ ). This is related to the resolving power by  $R \approx (2.35 \times 1 \times 10^{-4} \ln 10 \sigma_{\text{disp}})^{-1}$ . Palanque-Delabrouille et al. (2013) have recently found, using their own analysis of the width of the arc-lamp lines and bright sky emission lines, that the spectral dispersion reported by the pipeline had a bias that depended on the CCD row and increased with wavelength, up to 10% at  $\lambda \approx 6000 \text{ \AA}$ . We will correct for this bias when creating mock spectra to compare with the data, as described in Section 4. Figure 1 shows the (uncorrected) pixel dispersions from 236 BOSS quasars from the  $\langle z \rangle = 2.3$ ,  $S/N = 6\text{--}8$  bin, as a function of wavelength at the blue end ( $\lambda = 3700\text{--}4200 \text{ \AA}$ ) of the spectrograph. At a fixed wavelength, there are outliers that contribute to the large spread in  $\sigma_{\text{disp}}$ , e.g., ranging from  $\sigma_{\text{disp}} \approx 0.9\text{--}1.8$  at  $3700 \text{ \AA}$ . We therefore discard spectra with outlying values of  $\sigma_{\text{disp}}$  based on the following criterion: we first rank-order the spectra based on their  $\sigma_{\text{disp}}$  value evaluated at the central wavelength of each PDF bin (i.e.,  $\lambda = [4012, 4377, 4863] \text{ \AA}$  at  $\langle z \rangle = [2.3, 2.6, 3.0]$ ), and then discard spectra below the 5th percentile and above the 90th percentile. This is illustrated by the dashed red lines in Figure 1.

Finally, since our noise estimation procedure uses the individual BOSS exposures, we discard objects that have less than three individual exposures available.



**Figure 2.** Pixel distribution of Ly $\alpha$  absorber redshifts in the BOSS Ly $\alpha$  forest sample used in this paper, shown in bin sizes of  $\Delta z = 0.05$ . The different colors and line-styles denote the three redshift bins used in this paper. We have chosen these redshift bins—with the gap at  $2.75 < z < 2.85$ —to match the simulation redshifts (Section 4.1).

**Table 1**  
Binning of BOSS Ly $\alpha$  Forest Transmission PDFs

Ly $\alpha$ Forest Redshift	S/N <sup>a</sup> (pixel <sup>-1</sup> )	$N_{\text{spec}}^b$	$N_{\text{pix}}^c$	$\Delta v^d$ (km s <sup>-1</sup> )	$\Delta z^e$	$\Delta X^f$
2.15 < z < 2.45	6–8	1109	288442	$1.99 \times 10^7$	219	704
	8–10	501	129141	$8.90 \times 10^6$	97.9	315
	>10	561	146478	$1.01 \times 10^7$	111	357
2.45 < z < 2.75	6–8	1004	229898	$1.59 \times 10^7$	191	646
	8–10	490	107001	$7.38 \times 10^6$	88.6	300
	>10	604	140843	$9.71 \times 10^6$	117	396
2.85 < z < 3.15	6–8	511	108443	$7.48 \times 10^6$	99.7	358
	8–10	326	72448	$5.00 \times 10^6$	66.7	239
	>10	341	74284	$5.12 \times 10^6$	68.3	245

#### Notes.

<sup>a</sup> Median S/N within Ly $\alpha$  forest.

<sup>b</sup> Number of contributing spectra.

<sup>c</sup> Number of  $\Delta v = 69 \text{ km s}^{-1}$  pixels.

<sup>d</sup> Velocity path length.

<sup>e</sup> Redshift path length.

<sup>f</sup> Absorption distance, where  $dX/dz = (1+z)^2(\Omega_M(1+z)^3 + \Omega_\Lambda)^{-1/2}$ . For this conversion, we assume  $\Omega_M = 0.3$  and  $\Omega_\Lambda = 0.7$ .

Our final data set comprises 3373 unique quasars with redshifts ranging from  $z_{\text{qso}} = 2.255$  to  $z_{\text{qso}} = 3.811$ , and a median S/N of  $S/N = 8.08 \text{ pixel}^{-1}$ . This data set represents only a small subsample of the BOSS DR9 quasar spectra, but is over two orders of magnitude larger than high-resolution quasar samples previously used for transmission PDF analysis. Table 1 summarizes our data sample, and the statistics of the redshifts and S/N bins for which we measure the transmission PDF. Figure 2 shows histograms of the pixels used in our analysis, as a function of absorption redshift.

### 3. MEASURING THE TRANSMISSION PDF FROM BOSS

In this section, we will measure the Ly $\alpha$  forest transmission PDF from BOSS. In principle, the transmission PDF is simply the histogram of the transmitted flux in the Ly $\alpha$  forest after dividing by the quasar continuum. However, with the comparatively noisy BOSS data we need to ensure an accurate estimate of the pixel noise. We will therefore first describe a new

probabilistic method for co-adding the individual BOSS exposures that will enable us to have an accurate noise estimate. We will also describe the continuum-estimation method with which we normalize the forest transmission.

### 3.1. Co-addition of Multiple Exposures and Noise Estimation

Since we intend to model BOSS spectra with modest S/N, we need an accurate estimate of the pixel noise that also allows us to separate out the contributions from Poisson noise due to the background and sky as well as read noise from the detector. In this subsection, we will construct an accurate probabilistic model of the flux and noise of the BOSS spectrograph, based on the individual exposure data that BOSS delivers.

The basic BOSS spectral data consists of a spectrum of each raw exposure,  $f_{\lambda i}$  (inclusive of noise), an estimate of the sky  $s_{\lambda i}$ , and a calibration vector  $S_{\lambda i}$ , where  $i$  indicates the exposure of the  $n_{\text{exp}}$  exposures taken.<sup>19</sup> The quantity  $s_{\lambda i}$  is the actual sky model that was subtracted from the fiber spectra in the extraction. The calibration vector is defined as  $S_{\lambda i} \equiv f_{\lambda i}/f_{N i}$ , with  $f_{N i}$  being the flux of exposure  $i$  in units of photoelectrons. The idlspec2d pipeline then estimates the co-added spectrum of the true object flux,  $\mathcal{F}_{\lambda}$ , from the raw individual exposures, sky estimates, and calibration vectors.

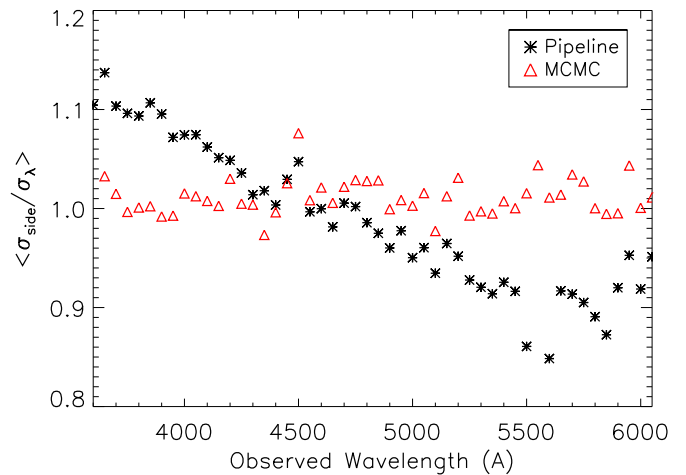
The BOSS data reduction pipeline also delivers noise estimates in the form of variance vectors, which are, however, known to be inaccurate (McDonald et al. 2006; Desjacques et al. 2007; Lee et al. 2013; Palanque-Delabrouille et al. 2013).

To quantify the fidelity of the BOSS noise estimate, we used the so-called ‘‘side-band’’ method described in Lee et al. (2014b) and Palanque-Delabrouille et al. (2013), which uses the variance in flat, absorption-free regions of the quasar spectra to quantify the fidelity of the noise estimate. First, we randomly selected 10,000 BOSS quasars (omitting BAL quasars) from the Pâris et al. (2012) catalog in the redshift range  $1.4 < z_{\text{qso}} < 3.4$ , evenly distributed into 20 redshift bins of width  $\Delta z_{\text{qso}} = 0.1$  (i.e., 500 objects per bin). We then consider the flat  $1460 \text{ \AA} < \lambda_{\text{rest}} < 1510 \text{ \AA}$  spectral region in the quasar rest-frame, which is dominated by the smooth power-law continuum and relatively unaffected by broad emission lines (e.g., Vanden Berk et al. 2001; Suzuki 2006) or absorption lines. The pixel variance in this flat portion of the spectrum should therefore be dominated by spectral noise, allowing us to examine whether the noise estimate provided by the pipeline is accurate. We then evaluate the ratio of  $\sigma_{\text{side}}$ , the pixel flux rms in the rest-frame  $1460 \text{ \AA} < \lambda_{\text{rest}} < 1510 \text{ \AA}$  region divided by the average pipeline noise estimate,  $\sigma_{\lambda}$ :

$$\left\langle \frac{\sigma_{\text{side}}}{\sigma_{\lambda}} \right\rangle = \frac{[\sum f_{\lambda}^2 - \bar{f}_{\lambda}^2]^{1/2}}{\sum \sigma_{\lambda}}, \quad (3)$$

where the summations and average flux is evaluated in the quasar rest-frame  $1460 \text{ \AA} < \lambda_{\text{rest}} < 1510 \text{ \AA}$ .

In Figure 3, this quantity is averaged over the 500 individual quasars per redshift bin and plotted as a function of the observed wavelength corresponding to  $\lambda = (1 + \langle z_{\text{qso}} \rangle)1485 \text{ \AA}$ . With a perfect noise estimate,  $\langle \sigma_{\text{side}}/\sigma_{\lambda} \rangle$  should be unity at all wavelengths, but we see that the BOSS pipeline underestimates the true noise in the spectra at  $\lambda \lesssim 5000 \text{ \AA}$ , by up to  $\sim 15\%$  at



**Figure 3.** Quantitative test of the noise estimation fidelity in the spectra. Each point shows the ratio of the pixel variance divided by the estimated noise variance, averaged over the rest-frame  $1460 \text{ \AA} < \lambda_{\text{rest}} < 1510 \text{ \AA}$  flat spectral region of 500 BOSS quasars within redshift bins of  $\Delta z_{\text{qso}} = 0.1$  and plotted as a function of the corresponding observed wavelength of the flat spectral region. If there is no bias in the noise estimation, this ratio should be unity. The black asterisks show this quantity estimated using the BOSS pipeline co-added spectra and noise estimates, while the red triangles show the results from the Markov Chain Monte Carlo (MCMC) co-addition and noise estimation procedure described in Section 3.1. The MCMC method clearly provides a better noise estimation than the BOSS pipeline.

the blue end of the spectra, with an overall tilt that changes over to an overestimate at  $\lambda \gtrsim 4500 \text{ \AA}$ . Lee et al. (2013) and Palanque-Delabrouille et al. (2013) provide a set of correction vectors that can be applied to the pipeline noise estimates to bring the latter to within several percent of the true noise level across the wavelength coverage of the blue spectrograph.

Unfortunately, these noise corrections are inadequate for our purposes, since we want to generate realistic mock spectra that have different realizations of the Ly $\alpha$  forest transmission field from the actual spectra, i.e., a different  $\mathcal{F}_{\lambda}$ . We therefore require a method that not only accurately estimates the noise in a given BOSS spectrum, but also separates out the photon-counting and CCD terms in the variance that results from applying the Horne (1986) optimal spectral extraction algorithm:

$$\sigma_{\lambda}^2 = S_{\lambda} (\mathcal{F}_{\lambda} + s_{\lambda}) + S_{\lambda}^2 \sigma_{\text{RN}}^2, \quad (4)$$

where  $\sigma_{\text{RN}}$  is the CCD read-noise.

To resolve this issue, we apply our own novel statistical method to the individual BOSS exposures to generate co-added spectra while simultaneously estimating the corresponding noise parameters for each individual spectrum. This procedure, which uses a Gibbs-sampled Markov-Chain Monte Carlo (MCMC) algorithm, is described in detail in the Appendix. Initially, we attempted to model the noise with just a single constant noise parameter that rescales the read-noise term of Equation (4), but this was found to be inadequate. This is likely because an optimal extraction algorithm weights by the product of the S/N and object profile, causing the corresponding variance to have a non-linear dependence on the flux and sky level. Furthermore, systematic errors in the reduction, sky-subtraction and calibration will result in additional noise contributions, which could depend on sky level, object flux, or wavelength, hence deviating from this simple model.

After considerable trial-and-error to find a model that best minimizes the bias illustrated in Figure 3, we settled

<sup>19</sup> Typically there are  $n_{\text{exp}} = 5$  exposures of 15 minutes each, although this can vary due to the requirements to achieve a given  $(\text{S/N})^2$  over each individual plug-plate, as determined by the overall BOSS survey strategy (see Dawson et al. 2013).

on the form:

$$\sigma_{\lambda_i}^2 = A_1 \hat{S}_{\lambda_i} (\mathcal{F}_\lambda + s_{\lambda_i}) + A_2 \hat{S}_{\lambda_i}^2 \sigma_{\text{RN,eff}}^2 \sigma_{\text{disp}}(\lambda), \quad (5)$$

where

$$\hat{S}_{\lambda_i} = S_{\lambda_i} (1 - \exp(-A_3 \lambda + A_4)), \quad (6)$$

where the  $A_j$  are free parameters in our noise model, while the  $\sigma_{\text{disp}}(\lambda)$  factor in the second term (the pixel dispersion) provides a rough approximation for the wavelength dependence of the spot size (i.e., the size of the raw CCD image in the spatial direction). Meanwhile,  $\sigma_{\text{disp}} = 12$  is the average CCD read-noise per wavelength bin in the BOSS spectra (D. J. Schlegel et al., in preparation). The quantities  $s_{\lambda,i}$ ,  $S_{\lambda,i}$ , and  $\sigma_{\text{disp}}(\lambda)$  (sky flux, calibration vector, and dispersion, respectively) are taken directly from the BOSS pipeline.

In addition, we assume that the pixel noise can be modeled as a Gaussian distribution with a variance given by Equation (5). The first, photon counting, term in the equation should formally be modeled as a Poisson distribution, but since the BOSS spectrograph always receives  $\gtrsim 30$ –40 counts even at the blue end of the spectrograph where the counts are the lowest, it is reasonable to use the Gaussian approximation because even in the limit of low S/N (i.e., when the spectrum is dominated by the sky flux), the moderate resolution ensures that there are at least several dozen sky photons per pixel in each exposure.

For each BOSS spectrum, we use the MCMC procedure described in the Appendix to combine the multiple exposures while simultaneously estimating the noise parameters  $A_j$  and true observed spectrum,  $\mathcal{F}_\lambda$ . With the optimal estimates of  $A_j$  and  $\mathcal{F}_\lambda$  for a given spectrum, the estimated noise variance is then simply Equation (5).

An important advantage of the form in Equation (5) is that the object photon noise  $\propto \mathcal{F}_\lambda$  is explicitly separated out. This facilitates the construction of a mock spectrum with the same noise characteristics as a true spectrum, *but with a different spectral flux*. For example, a mock spectrum of the Ly $\alpha$  forest will have a very different transmission field than the original data, and so the variance due to object photon counting noise can be added appropriately, in addition to contributions from the known sky, and the read noise term (Equation (5)). Our empirical determination of the parameters governing this noise model for each individual spectrum form a crucial ingredient in our forward model, which we will describe in Section 4.

Our MCMC procedure works for spectra from a single camera, either red or blue; we have not yet generalized it to combine blue and red spectra of each object. However, the spectral range of the blue camera alone ( $\approx 3600$ – $6400$  Å) covers the Ly $\alpha$  forest up to  $z \sim 5$ , i.e., most practical redshifts for Ly $\alpha$  forest analysis. For the purposes of this paper, we restrict ourselves to spectra from the blue camera alone.

In Figures 4 and 5, we show examples of co-added BOSS quasar spectra, using both the MCMC procedure and the standard BOSS pipeline. In the upper panels, the MCMC co-adds are not noticeably different from the BOSS pipeline, although the numerical values are different. In the lower panels, we show the estimated noise from both methods—the differences are larger than in the fluxes but still difficult to distinguish by eye.

We therefore return to the statistical analysis by calculating  $\langle \sigma_{\text{side}} / \sigma_\lambda \rangle$ , the ratio of the pixel variance against the estimated noise from the flat  $1460$  Å  $< \lambda_{\text{rest}} < 1510$  Å region of BOSS quasars; this ratio, computed for our MCMC coadds, is plotted in Figure 3. With these new co-adds, we see that this ratio is

within roughly  $\pm 3\%$  of unity across the entire  $\lambda \sim 3800$ – $5000$  Å wavelength range relevant to our subsequent analysis, with an overall bias of 1% (i.e., the noise is still underestimated by this level). Crucially, we have removed the strong wavelength dependence of  $\langle \sigma_{\text{side}} / \sigma_\lambda \rangle$  that was present in the standard pipeline, and we suspect most of the scatter about unity is caused by the limited number of quasars (500 per bin) available for this estimate, which will be mitigated by the larger number of quasars spectra available in the subsequent BOSS data releases. In principle, we could correct the remaining 1% noise bias, but since our selected spectra have S/N  $> 6$ , this remaining 1% noise bias would smooth the forest transmission PDF by an amount roughly 1/25 of the average PDF bin width ( $\Delta F = 0.05$ ). As we shall see, there are other systematic uncertainties in our modeling that have much larger effects than this, therefore we regard our noise estimates as adequate for the subsequent transmission PDF analysis, without requiring any further correction.

### 3.2. Mean-Flux Regulated Continuum Estimation

In order to obtain the transmitted flux  $F$  of the Ly $\alpha$  forest<sup>20</sup> we first need to divide the observed flux,  $\mathcal{F}_\lambda$ , by an estimate for the quasar continuum,  $c$ . We use the version of MF-PCA continuum fitting (Lee et al. 2012) described in Lee et al. (2013). Initially, PCA fitting with eight eigenvectors is performed on each quasar spectrum redward of the Ly $\alpha$  line ( $\lambda_{\text{rest}} = 1216$ – $1600$  Å) in order to obtain a prediction for the continuum shape in the  $\lambda_{\text{rest}} < 1216$  Å Ly $\alpha$  forest region (e.g., Suzuki et al. 2005). The slope and amplitude of this initial continuum estimate is then corrected to agree with the Ly $\alpha$  forest mean transmission,  $\langle F \rangle_{\text{cont}}(z)$ , at the corresponding absorber redshifts, using a linear correction function.

The only difference in our continuum-fitting with that in Lee et al. (2013) is that here we use the latest mean flux measurements of Becker et al. (2013) to constrain our continua. Their final result yielded the power-law redshift evolution of the effective optical depth in the unshielded Ly $\alpha$  forest, defined in their paper as  $N_{\text{HI}} \leq 10^{17.2} \text{ cm}^{-2}$  (although they only removed contributions from  $N_{\text{HI}} \geq 10^{19} \text{ cm}^{-2}$  absorbers). This is given by

$$\tau_{\text{Ly}\alpha, \text{B13}}(z) \equiv -\ln(\langle F \rangle(z)) = \tau_0 \left( \frac{1+z}{1+z_0} \right)^\beta + C, \quad (7)$$

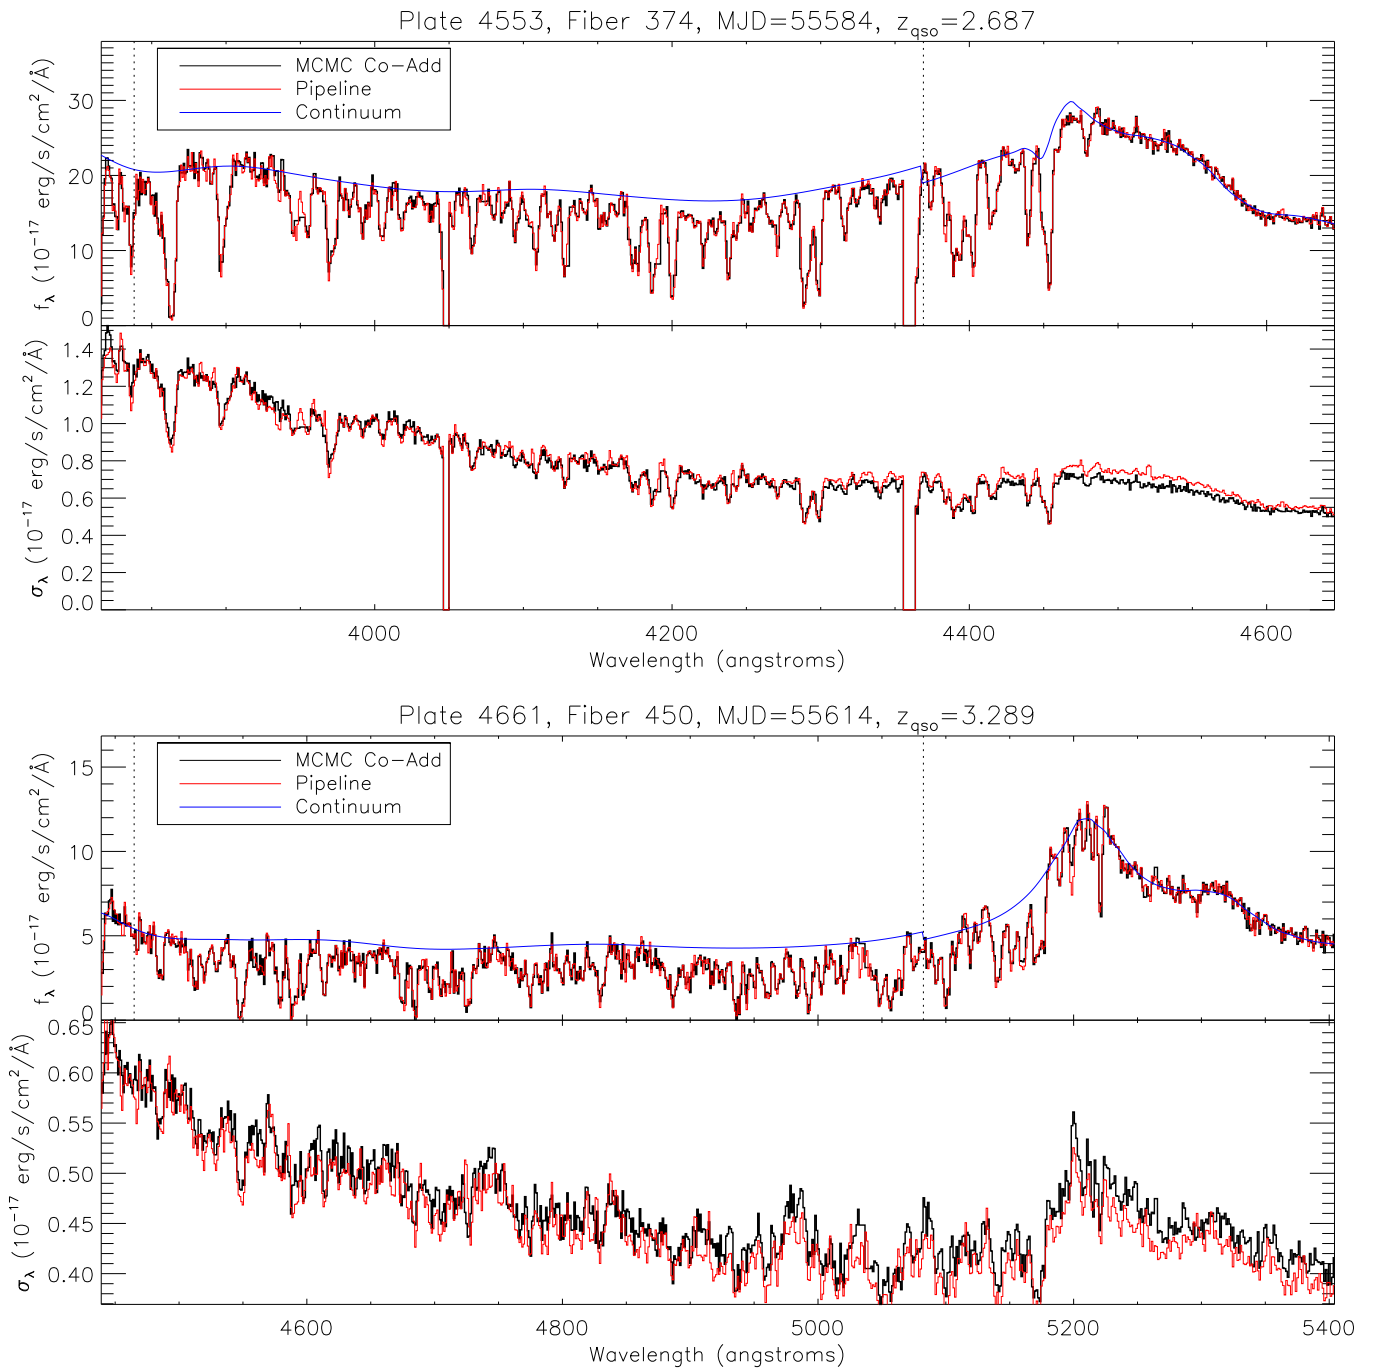
with best-fit values of  $[\tau_0, \beta, C] = [0.751, 2.90, -0.132]$  at  $z_0 = 3.5$ .

However, the actual raw measurement made by Becker et al. (2013) is the effective total absorption within the Ly $\alpha$  forest region of their quasars, which also contain contributions from metals and optically thick systems:

$$\tau_{\text{eff}}(z) \equiv \tau_{\text{Ly}\alpha, \text{B13}}(z) + \tau_{\text{metals}} + \tau_{\text{LLS}}(z), \quad (8)$$

where  $\tau_{\text{metals}}$  and  $\tau_{\text{LLS}}(z)$  denote the IGM optical depth contributions from metals and LLSs, respectively. For the purposes of our continuum-fitting, the quantity we require is  $\tau_{\text{eff}}(z)$ , since the  $\tau_{\text{metals}}$  and  $\tau_{\text{LLS}}(z)$  contributions are also present in our BOSS spectra. Becker et al. (2013) did not publish their raw  $\tau_{\text{eff}}(z)$ , therefore we must now “uncorrect” the metal and LLS

<sup>20</sup> Note that the ideal/model observed flux described in the noise modeling section,  $\mathcal{F}_\lambda$ , and the Ly $\alpha$  forest transmission  $F$ , are completely different quantities.



**Figure 4.** Examples of co-added BOSS spectra from the MCMC procedure described in Section 3.1 (red) and from the BOSS pipeline (black) are shown in the upper panels, in the rest-frame interval 1035–1260 Å. The corresponding pixel noise estimates are shown in the upper panels. The blue line shows the MF-PCA continuum used to extract the Ly $\alpha$  forest transmitted flux, while the vertical dotted lines delineate the 1041–1185 Å rest-frame interval, which we define as the Ly $\alpha$  forest. The continuum discontinuity at  $\lambda_{\text{rest}} = 1185$  Å is where we have applied the “mean flux regulation” correction to the Ly $\alpha$  forest. In the top figure, masked pixels have had their flux and noise set to zero. The signal-to-noise ratios for the two spectra are  $S/N \approx 11$  (top) and  $S/N \approx 6$  (bottom) within the Ly $\alpha$  forest.

contributions from the published  $\tau_{\text{Ly}\alpha, \text{B13}}(z)$ . The discussion below therefore attempts to retrace their footsteps and does not necessarily reflect our own beliefs regarding the actual level of these contributions.


We find  $\tau_{\text{metals}} = 0.02525$  by simply averaging over the Schaye et al. (2003) metal correction tabulated by Faucher-Giguère et al. (2008) (i.e., the  $2.2 \leq z \leq 2.5$  values in  $\Delta z = 0.1$  bins from their Table 4), that were used by Becker et al. (2013) to normalize their relative mean flux measurements. Note that there is no redshift dependence on  $\tau_{\text{metals}}$  in this context, because Becker et al. (2013) argued that the metal contribution does not

vary significantly over their redshift range. Whether or not this is really true is unimportant to us at the moment, since we are merely “uncorrecting” their measurement.

The LLS contribution to the optical depth is re-introduced by integrating over  $f(N_{\text{H1}}, b, z)$ , the column-density distribution of neutral hydrogen absorbers:

$$\tau_{\text{LLS}}(z) \approx \frac{1+z}{\lambda_{\text{Ly}\alpha}} \int_{N_{\text{min}}}^{N_{\text{max}}} dN_{\text{H1}} \int db \times f(N_{\text{H1}}, b, z) W_0(N_{\text{H1}}, b), \quad (9)$$





**Figure 5.** Same as Figure 4, but the  $1050 \text{ \AA} < \lambda_{\text{rest}} < 1090 \text{ \AA}$  rest-frame region is expanded to better illustrate the differences between the MCMC and pipeline co-added spectra.

where  $b$  is the Doppler parameter and  $W_0(N_{\text{H I}}, b)$  is the rest-frame equivalent width (EW; we use the analytic approximation given by Draine 2011, valid in the saturated regime).


Following Becker et al. (2013), we adopted a fixed value of  $b = 20 \text{ km s}^{-1}$  and assumed that  $f(N_{\text{H I}}, z) = f(N_{\text{H I}})dn/dz$ , where  $f(N_{\text{H I}})$  is given by the  $z = 3.7$  broken power-law column density distribution of Prochaska et al. (2010) and  $dn/dz \propto (1+z)^2$ . Becker et al. (2013) had corrected for super-LLSs and DLAs in the column-density range  $[N_{\text{min}}, N_{\text{max}}] = [10^{19} \text{ cm}^{-2}, 10^{22} \text{ cm}^{-2}]$ , but as discussed above we have discarded all sightlines that include  $N_{\text{H I}} \geq 10^{20.3} \text{ cm}^{-2}$  DLAs; therefore, we reintroduce the optical depth contribution for super-LLSs, i.e.,  $[N_{\text{min}}, N_{\text{max}}] = [10^{19} \text{ cm}^{-2}, 10^{20.3} \text{ cm}^{-2}]$ . We

find  $\tau_{\text{LLS}}(z) = 0.0022 \times [(1+z)/3]^3$ . This is a small correction, giving rise to only a 0.5% change in  $\langle F \rangle$  at  $z = 3.0$ .

This estimate of the raw absorption,  $\langle F \rangle_{\text{eff}}(z) = \exp[-\tau_{\text{eff}}(z)]$ , is now the constraint used to fit the continua of the BOSS quasars, i.e., we set  $\langle F \rangle_{\text{cont}} = \langle F \rangle_{\text{eff}}(z)$ . Note that in our subsequent modeling of the data, we will use the same  $\langle F \rangle_{\text{cont}}(z)$  to fit the mock spectra to ensure an equal treatment between data and mocks. Since  $\langle F \rangle_{\text{cont}}(z)$  includes a contribution from  $N_{\text{H I}} < 10^{20.3} \text{ cm}^{-2}$  optically thick systems, our mock spectra will need to account for these systems as we shall describe in Section 4.2.

The MF-PCA technique requires spectral coverage in the quasar rest-frame interval  $1000\text{--}1600 \text{ \AA}$ . However, as noted





**Figure 6.**  $\text{Ly}\alpha$  forest transmission PDFs,  $p(F)$ , measured from different subsamples of our BOSS sample, at various redshift (with  $\Delta z = 0.3$ ) and S/N. Both the upper and lower panels show the PDF, but with linear and logarithmic ordinate axes, respectively. The different colors and line styles denote our different S/N subsamples at each redshift. The error bars are estimated from bootstrap resampling over  $\Delta v = 2 \times 10^4 \text{ km s}^{-1}$  segments from the contributing spectra. Table 1 summarizes the number of spectra and pixels that contribute to each bin.

in the previous section, we work with co-added BOSS spectra from only the blue cameras covering  $\lambda \lesssim 6400 \text{ \AA}$ ; this covers the full 1000–1600  $\text{\AA}$  interval required for the PCA fitting only for  $z \lesssim 3$  quasars. However, the differences in the fluxes between our MCMC co-adds and the BOSS pipeline co-adds are relatively small, and we do not expect the relative shape of the quasar spectrum to vary significantly. We can thus carry out PCA fitting on the BOSS pipeline co-adds, which cover the full observed range (3700–10000  $\text{\AA}$ ), to predict the overall quasar continuum shape. This initial prediction is then used to perform mean flux regulation using the MCMC co-adds and noise estimates, to fine-tune the amplitude of the continuum fits.

The observed flux,  $f_\lambda$ , is divided by the continuum estimate,  $c$ , to derive the  $\text{Ly}\alpha$  forest transmission,  $F = f_\lambda/c$ . For each quasar, we define the  $\text{Ly}\alpha$  forest as the rest wavelength interval 1041–1185  $\text{\AA}$ . This wavelength range conservatively avoids the quasar’s  $\text{Ly}\beta/\text{O VI}$  emission line blend by  $\Delta v \sim 3000 \text{ km s}^{-1}$  on the blue end, as well as the proximity zone close to the quasar redshift by staying  $\Delta v \sim 10,000 \text{ km s}^{-1}$  from the nominal quasar systemic redshift. We are now in a position to measure the transmission PDF, which is simply the histogram of pixel transmissions  $F \equiv \exp(-\tau)$ .

### 3.3. Observed Transmission PDF from BOSS

Since the  $\text{Ly}\alpha$  forest evolves as a function of redshift, we measure the BOSS  $\text{Ly}\alpha$  forest transmission PDF in three bins with mean redshifts of  $\langle z \rangle = 2.3$ ,  $\langle z \rangle = 2.6$ , and  $\langle z \rangle = 3.0$ , and bin sizes of  $\Delta z = 0.3$ . These redshift bins were chosen to match the simulations outputs (Section 4.1) that we will later use to make mock spectra to compare with the observed PDF; this choice of binning leads to the gap at  $2.75 < z < 2.85$  as seen in Figure 2. In this paper, we restrict ourselves to  $z \lesssim 3$  since the primary purpose is to develop the machinery to model the BOSS spectra. In subsequent papers, we will apply these techniques to analyze the transmission PDF in the full  $2 \lesssim z \lesssim 4$  range using the larger samples of subsequent BOSS data releases (DR10; Ahn et al. 2014).


Another consideration is that the transmission PDF is strongly affected by the noise in the data. While we will model this effect in detail (Section 4), there is a large distribution of S/N within our subsample ranging from  $S/N = 6 \text{ pixel}^{-1}$  to

$S/N \sim 20 \text{ pixel}^{-1}$ . We therefore further divide the sample into three bins depending on the median S/N per pixel within the  $\text{Ly}\alpha$  forest:  $6 < S/N < 8$ ,  $8 < S/N < 10$ ,  $S/N > 10$ . The consistency of our results across the S/N bins will act as an important check for the robustness of our noise model (Section 3.1).

We now have nine redshift and S/N bins in which we evaluate the transmission PDF from BOSS; the sample sizes are summarized in Table 1. For each bin, we have selected quasars that have at least 30  $\text{Ly}\alpha$  forest pixels within the required redshift range, and which occupy the quasar rest-frame interval 1041–1185  $\text{\AA}$ . The co-added spectrum is divided with its MF-PCA continuum estimate (described in the previous section) to obtain the transmitted flux,  $F$ , in the desired pixels. We then compute the transmission PDF from these pixels.

Physically, the possible values of the  $\text{Ly}\alpha$  forest transmission range from  $F = 0$  (full absorption) to  $F = 1$  (no absorption). However, the noise in the BOSS  $\text{Ly}\alpha$  forest pixels, as well as continuum fitting errors, leads to pixels with  $F < 0$  and  $F > 1$ . We therefore measure the transmission PDF in the range  $-0.2 < F < 1.5$ , in 35 bins with width  $\Delta(F) = 0.05$ , and normalized such that the area under the curve is unity. The statistical errors on the transmission PDF are estimated by the following method: we concatenate all the individual  $\text{Ly}\alpha$  forest segments that contribute to each PDF, and then carry out bootstrap resampling over  $\Delta v = 2 \times 10^4 \text{ km s}^{-1}$  segments with 200 iterations. This choice of  $\Delta v$  corresponds to  $\sim 250\text{--}300 \text{ \AA}$  in the observed frame at  $z \sim 2\text{--}3$ —according to Rollinde et al. (2013), this choice of  $\Delta v$  and number of iterations should be sufficient for the errors to converge (see also Appendix B in McDonald et al. 2000).

In Figure 6, we show the  $\text{Ly}\alpha$  forest transmission PDF measured from the various redshift and S/N subsamples in our BOSS sample. At fixed redshift, the PDFs from the lower S/N data have a broader shape as expected from increased noise variance. With increasing redshift, there are more absorbed pixels, causing the transmission PDFs to shift toward lower  $F$  values. As discussed previously, there is a significant portion of  $F > 1$  pixels due to a combination of pixel noise and continuum errors, with a greater proportion of  $F > 1$  pixels in the lower-S/N subsamples as expected. Unlike the high-resolution transmission PDF, at  $\langle z \rangle \lesssim 3$  there are few pixels that reach  $F = 0$ . This effect is due to the resolution of

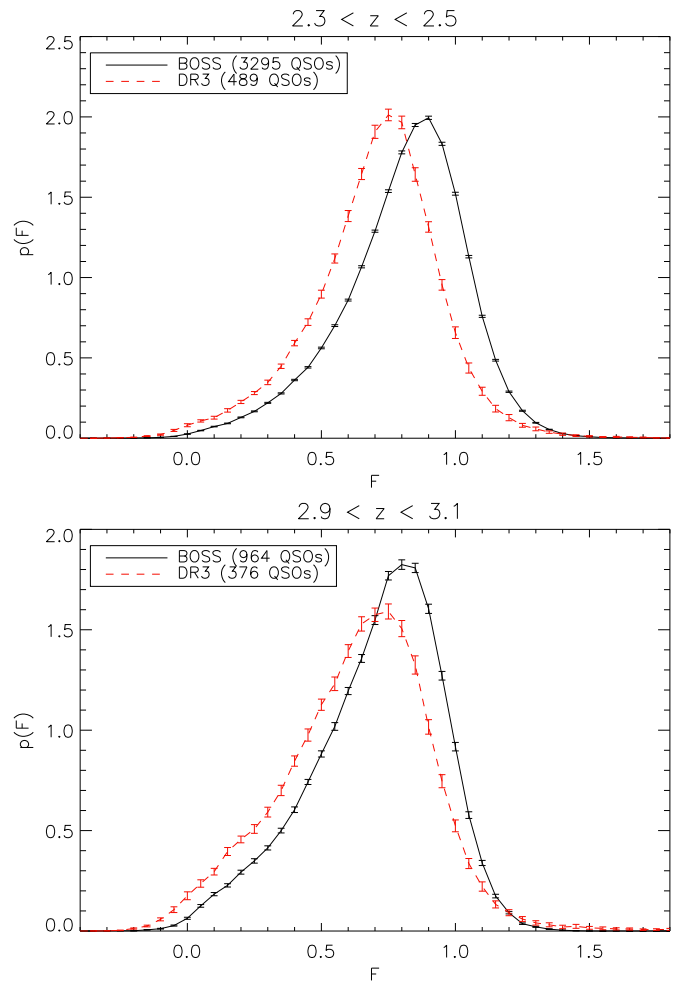


**Figure 7.** Top: 2D density plot of the error covariance matrix for the Ly $\alpha$  forest transmission PDF from the  $\langle z \rangle = 2.6$ ,  $S/N = 8\text{--}10$  BOSS subsample as a function of transmission bins, along with (bottom) the corresponding correlation function. The covariance matrix was estimated through bootstrap resampling, and the values have been multiplied by  $10^4$  for clarity. The covariances are largely diagonal, except for some cross-correlations between neighboring bins.

the BOSS spectrograph, which smooths over the observed Ly $\alpha$  forest such that even saturated Ly $\alpha$  forest absorbers with  $N_{\text{H I}} \gtrsim 10^{14}\text{--}10^{16} \text{ cm}^{-2}$  rarely reach transmission values of  $F \lesssim 0.3$ . The pixels with  $F \lesssim 0.3$  are usually contributed either by blends of absorbers or optically thick LLSs (see also Pieri et al. 2014).

An advantage of our large sample size is that it is also able to directly estimate the error covariances,  $\mathbf{C}_{\text{boot}}$ , via bootstrap resampling—an example is shown in Figure 7. In contrast to the Ly $\alpha$  forest transmission PDF from high-resolution data which have significant off-diagonal covariances (Bolton et al. 2008), the error covariance from the BOSS transmission PDF is nearly diagonal with just some small correlations between neighboring bins, although we also see some anti-correlation between transmission bins at  $F \sim 0.8$  and  $F \sim 1$ .

It is interesting to compare the transmission PDF from our data with that measured by Desjacques et al. (2007) from SDSS DR3. This comparison is shown in Figure 8, in which



**Figure 8.** Comparison between the Ly $\alpha$  forest transmission PDFs measured from our BOSS DR9 sample (solid black lines), and the SDSS DR3 sample from Desjacques et al. (2007; dashed red lines). Only sightlines with  $S/N > 4$  were used in evaluating these PDFs. The lower average transmission of the DR3 PDFs is because Desjacques et al. (2007) had directly extrapolated a power law from  $\lambda_{\text{rest}} > 1216 \text{ \AA}$  for continuum estimates, which does not take into account a flattening of the quasar continuum that occurs at  $\lambda_{\text{rest}} \sim 1200 \text{ \AA}$ ; our BOSS spectra, in contrast, have been normalized to mean transmission values in agreement with the latest measurements and takes this effect into account.

the transmission PDFs calculated from SDSS DR3 Ly $\alpha$  forest spectra with  $S/N > 4$  (kindly provided by Dr. V. Desjacques) are shown for two redshift bins, juxtaposed with the BOSS transmission PDFs calculated from spectra with the same redshift and  $S/N$  cuts.

While there is some resemblance between the two PDFs, the most immediate difference is that the Desjacques et al. (2007) PDFs are shifted to lower transmission values, i.e., the mean transmission,  $\langle F \rangle$ , is considerably smaller than that from our BOSS data:  $\langle F \rangle(z = 2.4) = 0.73$  and  $\langle F \rangle(z = 3.0) = 0.64$  from their measurement, whereas the BOSS PDFs have  $\langle F \rangle(z = 2.4) = 0.80$  and  $\langle F \rangle(z = 3.0) = 0.70$ . This difference arises because the Desjacques et al. (2007) used a power-law continuum (albeit with corrections for the weak emission lines in the quasar continuum) extrapolated from  $\lambda_{\text{rest}} > 1216 \text{ \AA}$  in the quasar rest-frame; this does not take into account the power-law break that appears to occur in low-redshift quasar spectra at  $\lambda_{\text{rest}} \approx 1200 \text{ \AA}$  (Telfer et al. 2002; Suzuki 2006). Later in their paper, Desjacques et al. (2007) indeed conclude that this must be the case in order to be consistent with other  $\langle F \rangle(z)$

measurements. Our continua, in contrast, have been constrained to match existing measurements of  $\langle F \rangle(z)$ , for which there is good agreement between different authors at  $z \lesssim 3$  (e.g., Faucher-Giguère et al. 2008; Becker et al. 2013).

Another point of interest in Figure 8 is that the error bars of the BOSS sample are considerably smaller than those of the earlier measurement. This difference is largely due to the significantly larger sample size of BOSS. The proportion of pixels with  $F \lesssim 0$  appears to be smaller in the BOSS PDFs compared with the older data set, but this is because Desjacques et al. (2007) did not remove DLAs from their data.

We next describe the creation of mock Ly $\alpha$  absorption spectra designed to match the properties of the BOSS data.

#### 4. MODELING OF THE BOSS TRANSMISSION PDF

In this section, we will describe simulated Ly $\alpha$  forest mock spectra designed, through a “forward-modeling” process, to have the same characteristics as the BOSS spectra, for comparison with the observed transmission PDFs described in the previous section. For each BOSS spectrum that contributed to our transmission PDFs in the previous section, we will take the Ly $\alpha$  absorption from randomly selected simulation sightlines, then introduce the characteristics of the observed spectrum using auxiliary information returned by our pipeline.

Starting with simulated spectra from a set of detailed hydrodynamical IGM simulations, we carry out the following steps, which we will describe in turn in the subsequent subsections:

1. Introduce LLS absorbers.
2. Smooth the spectrum to BOSS resolution.
3. Add metal absorption via an empirical method using lower-redshift SDSS/BOSS quasars.
4. Add pixel noise, based on the noise properties of the real BOSS spectrum using parameters estimated by our MCMC noise estimation technique.
5. Simulate continuum errors by refitting the noisy mock spectrum.

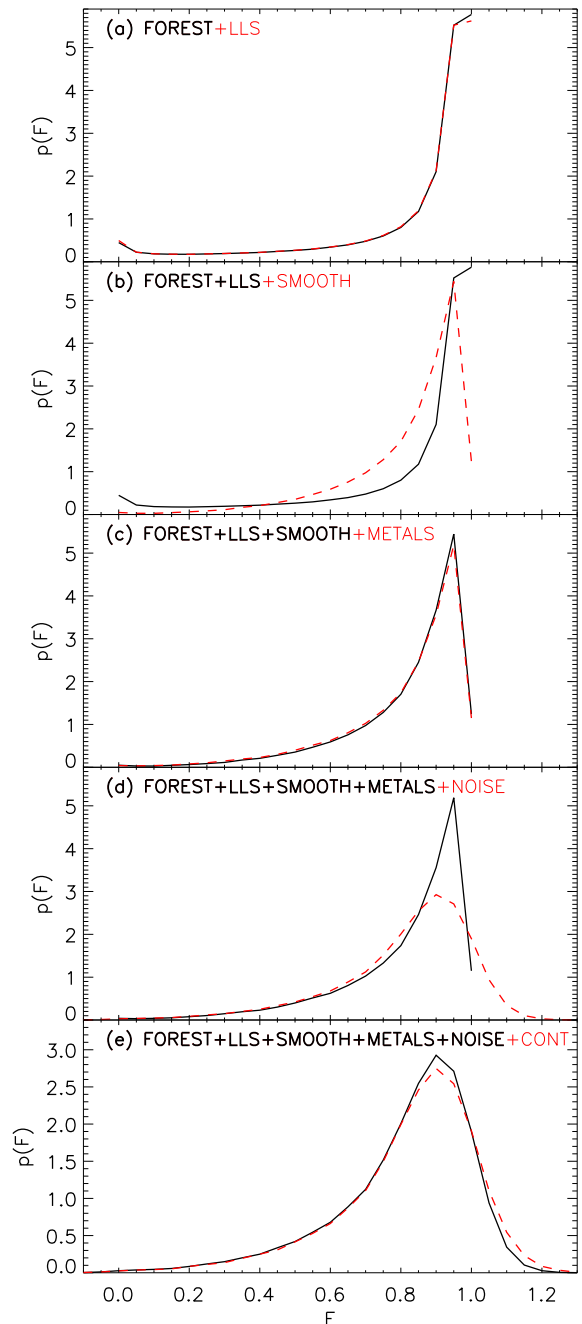
In the subsequent subsections, we will describe each step in detail. The effect of each step on the observed transmission PDF is illustrated in Figure 9.

##### 4.1. Hydrodynamical Simulations

As the basis for our mock spectra, we use hydrodynamic simulations run with a modification of the publicly available GADGET-2 code. This code implements a simplified star formation criterion (Springel et al. 2005) that converts all gas particles that have an overdensity above 1000 and a temperature below  $10^5$  K into star particles (see Viel et al. 2004). The simulations used are described in detail in Becker et al. (2011) and in Viel et al. (2013a).


The reference model that we use is a box of length  $20 h^{-1}$  comoving Mpc with  $2 \times 512^3$  gas and cold DM particles (with a gravitational softening length of  $1.3 h^{-1}$  kpc) in a flat  $\Lambda$ CDM universe with cosmological parameters  $\Omega_m = 0.274$ ,  $\Omega_b = 0.0457$ ,  $n_s = 0.968$ ,  $H_0 = 70.2 \text{ km s}^{-1} \text{ Mpc}^{-1}$  and  $\sigma_8 = 0.816$ , in agreement both with nine-year *WMAP* (Komatsu et al. 2011) and *Planck* data (Planck Collaboration et al. 2014). The initial condition power spectra are generated with CAMB (Lewis et al. 2000). For the boxes considered in this work, we have verified that the transmission PDF has converged in terms of box size and resolution.

We explore the impact of different thermal histories on the Ly $\alpha$  forest by modifying the ultraviolet (UV) background



**Figure 9.** Cumulative effect of various aspects of our forward model that attempts to reproduce the Ly $\alpha$  forest transmission PDF from BOSS. Starting with the “raw” transmission PDF from the simulations (top), the black curve in each panel shows the PDF from the prior panel, while the red curve shows the effect from (a) the addition of LLS, (b) smoothing from the finite spectrograph resolution, (c) contamination from lower-redshift metals, (d) pixel noise, and (e) continuum fitting errors. The transmission PDF modeled in this figure is from the  $(z) = 2.3, 8 < S/N < 10$  bin.

photo-heating rates in the simulations as done in, e.g., Bolton et al. (2008). A power-law temperature–density relation,  $T = T_0 \Delta^\gamma$ , arises in the low-density IGM ( $\Delta < 10$ ) as a natural consequence of the interplay between photo-heating and adiabatic cooling (Hui et al. 1997; Gnedin & Hui 1998). The value of  $\gamma$  within a simulation can be modified by varying a density-dependent heating term (see, e.g., Bolton et al. 2008). We consider a range of values for the temperature at mean density,  $T_0$ , and the power-law index of the temperature–density relation,  $\gamma$ , based on the observational measurements presented



**Figure 10.** Simulated  $\langle z \rangle = 2.3$  Ly $\alpha$  forest skewer from our hydrodynamical simulations, without smoothing (top panel) and smoothed to BOSS resolution (bottom panel). The black curve is the simulated transmission directly extracted from the simulations, while the red curve is the same transmission field but with an LLS added at  $\lambda = 4057 \text{ \AA}$  or  $z = 2.337$ . The blue curve in the bottom panel shows the effect of the metal absorbers added using our empirical method. For illustrative purposes, we have specifically chosen to this simulated sightline to have significant LLS and metal absorption; it is possible for a sightline to have neither. The dashed horizontal line denotes  $F = 0.3$ , below which our fiducial transmission PDF model disagrees with BOSS (see Section 5).

**Table 2**  
Evolution of  $T_0$  in Hydrodynamical Simulations

$\langle z \rangle$	T_COLD	T_REF	T_HOT
2.3	13000 K	18000 K	23000 K
2.6	11000 K	16000 K	21500 K
3.0	9000 K	14000 K	19000 K

recently by Becker et al. (2011). These consist of a set of three different indices for the temperature–density relation,  $\gamma(z = 2.5) \sim 1.0, 1.3, 1.6$ , that are kept roughly constant over the redshift range  $z = [2-6]$  and three different temperatures at mean density,  $T_0(z = 2.5) \sim [11000, 16000, 21500] \text{ K}$ , which evolve with redshift, yielding a total of nine different thermal histories. Between  $z = 2$  and  $z = 3$  there is some temperature evolution and the IGM becomes hotter at low redshift; at  $z = 2.3$ , the models have  $T_0 \sim [13000, 18000, 23000] \text{ K}$ . We refer to the intermediate temperature model as our “reference” model, or T\_REF, while the hot and cold models are referred to as T\_HOT and T\_COLD, respectively. The values of  $T_0$  of our simulations at the various redshifts are summarized in Table 2.

Approximately 4000 core hours were required for each simulation run to reach  $z = 2$ . The physical properties of the Ly $\alpha$  forest obtained from the TREE PM/SPH code GADGET-2 are in agreement at the percent level with those inferred from the moving-mesh code AREPO (Bird et al. 2013) and with the Eulerian code ENZO (O’Shea et al. 2004).

For this study, the simulation outputs were saved at  $z = [2.3, 2.6, 3.0]$ , from which we extract 5000 optical depth sightlines binned to 2048 pixels each. To convert these to transmission spectra, the optical depths were rescaled such that the skewers collectively yielded a desired mean transmission,  $\langle F \rangle_{\text{Ly}\alpha} \equiv \exp(-\tau_{\text{Ly}\alpha})$ . For our fiducial models, we would like to use the mean transmission values estimated by Becker et al. (2013), which we denote as  $\langle F \rangle_{\text{Ly}\alpha, \text{B13}} \equiv \exp(-\tau_{\text{Ly}\alpha, \text{B13}})$ . However, their estimates assume certain corrections from optically thick systems and metal absorption. We therefore add back in the corrections they made (see discussion in Section 3.2) to get their “raw” measurement for  $\langle F \rangle$  that now includes all optically thick systems and metals, and then remove these contributions assuming our own LLS and metal absorption models (see below).

Later in the paper, we will argue that our PDF analysis in fact places independent constraints on  $\langle F \rangle_{\text{Ly}\alpha}$ .

#### 4.2. Lyman-limit Systems

In principle, all optically thick Ly $\alpha$  absorbers such as LLSs and DLAs should be discarded from Ly $\alpha$  forest analyses, since they do not trace the underlying matter density field in the same way as the optically thin forest (Equation (1)), and require radiative transfer simulations to accurately capture their properties (e.g., McQuinn et al. 2011; Rahmati et al. 2013).

While DLAs are straightforward to identify through their saturated absorption and broad damping wings even in noisy BOSS data (see, e.g., Noterdaeme et al. 2012), the detection completeness of optically thick systems through their Ly $\alpha$  absorption drops rapidly at  $N_{\text{H I}} \lesssim 10^{20} \text{ cm}^{-2}$ . Even in high-S/N, high-resolution spectra, optically thick systems can only be reliably detected through their Ly $\alpha$  absorption at  $N_{\text{H I}} \gtrsim 10^{19} \text{ cm}^{-2}$  (“super-LLS”). Below these column densities, optically thick systems can be identified either through their rest-frame 912  $\text{\AA}$  Lyman-limit (albeit only one per spectrum) or using higher-order Lyman-series lines (e.g., Rudie et al. 2013). Neither of these approaches have been applied in previous Ly $\alpha$  forest transmission PDF analyses (McDonald et al. 2000; Kim et al. 2007; Calara et al. 2012; Rollinde et al. 2013), so arguably all these analyses are contaminated by LLSs.

Instead of attempting to remove LLSs from our observed spectra, we incorporate them into our mock spectra through the following procedure. For each PDF bin, we evaluate the total redshift pathlength of the contributing BOSS spectra (and corresponding mocks)—this quantity is summarized in Table 1. This is multiplied by  $l_{\text{LLS}}(z)$ , the number of LLS per unit redshift, to give the total number of LLS expected within our sample. We used the published estimates of this quantity by Ribaudo et al. (2011)<sup>21</sup> which is valid over  $0.24 < z < 4.9$ :

$$l_{\text{LLS}}(z) = l_{z_0}(1+z)^{\gamma_{\text{LLS}}}, \quad (10)$$

where  $l_{z_0} = 0.1157$  and  $\gamma_{\text{LLS}} = 1.83$ .

<sup>21</sup> Note that the value  $l_{z_0} = 0.30$  given in Table 6 of Ribaudo et al. (2011) is actually erroneous, and the correct normalization is in fact  $l_{z_0} = 0.1157$ , consistent with the data in their paper, which is used in Equation (10). Dr. J. Ribaudo 2014, private communication, has concurred with this conclusion.



After estimating the total number of LLSs in our mock spectra,  $l_{\text{LLS}}(z)\Delta z$ , we add them at random points within our set of simulated optical depth skewers. We also experimented with adding LLSs such that they are correlated with regions that already have high column density (e.g., Font-Ribera & Miralda-Escudé 2012), but we found little significant changes to the transmission PDF and therefore stick to the less computationally intensive random LLSs.

For each model LLS, we then draw a column density using the published LLS column density distribution,  $f(N_{\text{HI}})$ , from Prochaska et al. (2010). This distribution is measured at  $z \approx 3.7$ , so we make the assumption that  $f(N_{\text{HI}})$  does not evolve with redshift between  $2 \lesssim z \lesssim 3.7$ . For our column densities of interest, this distribution is represented by the broken power laws:

$$f(N_{\text{HI}}) = \begin{cases} k_1 N_{\text{HI}}^{-0.8} & \text{if } 10^{17.5} < N_{\text{HI}} < 10^{19.0} \\ k_2 N_{\text{HI}}^{-1.2} & \text{if } 10^{19.0} < N_{\text{HI}} < 10^{20.3} \end{cases} \quad (11)$$

For the normalizations  $k_1$  and  $k_2$ , we demand that

$$\int_{10^{17.5}}^{10^{19.0}} k_1 N_{\text{HI}}^{-0.8} dN_{\text{HI}} + \int_{10^{19.0}}^{10^{20.3}} k_2 N_{\text{HI}}^{-1.2} dN_{\text{HI}} = 1, \quad (12)$$

and require both power laws to be continuous at  $N_{\text{HI}} = 10^{19.0} \text{ cm}^{-2}$ . These constraints produce  $k_1 = 10^{-4.505}$  and  $k_2 = 10^{3.095}$ . After drawing a random value for the column density of each LLS, we add the corresponding Voigt profile to the optical depth in the simulated skewer.

In addition to the LLS with column densities of  $10^{17.5} \text{ cm}^{-2} < N_{\text{HI}} < 10^{20.3} \text{ cm}^{-2}$  that are defined to have  $\tau_{\text{HI}} \geq 2$ , there is also a population of partial Lyman-limit systems (pLLSs) that are not well-captured in our hydrodynamical simulations since they have column densities ( $10^{16.5} \text{ cm}^{-2} \lesssim N_{\text{HI}} < 10^{17.5} \text{ cm}^{-2}$ ) at which radiative transfer effects become significant ( $\tau_{\text{HI}} \gtrsim 0.1$ ). However, the incidence rates and column-density distribution of pLLSs are ill-constrained since they are difficult to detect in normal LLS searches. We therefore account for the pLLS by extrapolating the low end of the power-law distribution in Equation (11) down to  $N_{\text{HI}} = 10^{16.5} \text{ cm}^{-2}$ , i.e.,

$$f(10^{16.5} \text{ cm}^{-2} < N_{\text{HI}} < 10^{17.5} \text{ cm}^{-2}) = k_1 N_{\text{HI}}^{-0.8}. \quad (13)$$

This simple extrapolation does not take into account constraints from the mean free path of ionizing photons (e.g., Prochaska et al. 2010), which predicts a steeper slope for the pLLS distribution, but we will explore this later in Section 5.2.

Comparing the integral of this extrapolated pLLS distribution with Equation (12) leads us to conclude that

$$l_{\text{pLLS}}(z) = 0.197 l_{\text{LLS}}(z), \quad (14)$$

and we proceed to randomly add pLLSs to our mock spectra in the same way as LLSs.

The other free parameter in our LLS model is their effective  $b$ -parameter distribution. However, due to the observational difficulty in identifying  $N_{\text{HI}} \lesssim 18.5 \text{ cm}^{-2}$  LLSs, the  $b$ -parameter distribution of this distribution has, to our knowledge, never been quantified. Due to this lack of knowledge, it is common to simply adopt a single  $b$ -value when attempting to model LLSs (e.g., Font-Ribera & Miralda-Escudé 2012; Becker et al. 2013). We therefore assume that all our pLLSs and LLSs have a  $b$ -parameter of  $b = 70 \text{ km s}^{-1}$  similar to DLAs (Prochaska & Wolfe 1997), an “effective” value meant to capture the blending

of multiple Ly $\alpha$  components. However, the  $b$ -parameter for this population of absorbers is a highly uncertain quantity and as we shall see, it will need to be modified to provide a satisfactory fit to the data although it will turn out to not strongly affect our conclusions regarding the IGM temperature–density relationship. Figure 10 illustrates the effect of adding a LLS to a mock spectrum, and its subsequent smoothing with the spectrograph resolution kernel (next section).

### 4.3. Spectral Resolution

The spectral resolution of SDSS/BOSS spectra is  $R \equiv \lambda/\Delta\lambda \approx 1500\text{--}2500$  (Smee et al. 2013). The exact value varies significantly both as a function of wavelength, and across different fibers and plates depending on observing conditions (Figure 1).

For each spectrum, the BOSS pipeline provides an estimate of the  $1\sigma$  wavelength dispersion at each pixel,  $\sigma_{\text{disp}}$ , in units of the co-added wavelength grid size ( $\Delta \log_{10} \lambda = 10^{-4}$ ). The spectral resolution at that pixel can then be obtained from the dispersion, through the following conversion:  $R \approx (2.35 \times 1 \times 10^{-4} \ln 10 \sigma_{\text{disp}})^{-1}$ . Figure 1 shows the pixel dispersions from 236 randomly selected BOSS quasar as a function of wavelength at the blue end of the spectrograph. Even at fixed wavelength, there is a considerable spread in the dispersion, e.g., ranging from  $\sigma_{\text{disp}} \approx 0.9\text{--}1.8$  at  $3700 \text{ \AA}$ . The value of  $\sigma_{\text{disp}}$  typically decreases with wavelength (i.e., the resolution increases).


In their analysis of the Ly $\alpha$  forest 1D transmission power spectrum, Palanque-Delabrouille et al. (2013) made their own study of the BOSS spectral resolution by directly analyzing the line profiles of the mercury and cadmium arc lamps used in the wavelength calibration. They found that the pipeline underestimates the spectral resolution as a function of fiber position (i.e., CCD row) and wavelength: the discrepancy is  $< 1\%$  at blue wavelengths and near the CCD edges, but increases to as much as  $10\%$  at  $\lambda \sim 6000 \text{ \AA}$  near the center of the blue CCD (compare with Figure 4 in Palanque-Delabrouille et al. 2013). Our analysis is limited to  $\lambda \leq 5045 \text{ \AA}$ , i.e.,  $z \leq 3.15$ , where the discrepancy is under  $4\%$ . Nevertheless, we implement these corrections to the BOSS resolution estimate to ensure that we model the spectral resolution to an accuracy of  $< 1\%$ .

For each BOSS Ly $\alpha$  forest segment that contributes to the observed transmission PDFs discussed in Section 3, we concatenate randomly selected transmission skewers from the simulations described in the previous section. This is because the simulation box size of  $L = 20 h^{-1} \text{ Mpc}$  ( $\Delta v \sim 2000 \text{ km s}^{-1}$ ) is significantly shorter than the path length of our redshift bins ( $\Delta z = 0.3$ , or  $\Delta v \approx 27,000 \text{ km s}^{-1}$ ). This ensures that each BOSS spectrum in our sample has a mock spectrum that is exactly matched in pathlength.

We then directly convolve the simulated skewers by a Gaussian kernel with a standard deviation that varies with wavelength, using the estimated resolution from the real spectrum, multiplied by the Palanque-Delabrouille et al. (2013) resolution corrections. The effect of smoothing on the transmission PDF is illustrated by the dashed red curve in Figure 9(b). Smoothing decreases the proportion of pixels with high transmission ( $F \approx 1$ ) and with high absorption ( $F \approx 0$ ), and increases the number of pixels with intermediate transmission values.

### 4.4. Metal Contamination

Metal absorption along our observed Ly $\alpha$  forest sightlines acts as a contaminant since their presence alters the observed



**Figure 11.** Illustration of our empirical “sideband” model of metal contamination in our mock Ly $\alpha$  forest spectra. The lower panel shows the  $z_{\text{qso}} = 2.7$  quasar along with its Ly $\alpha$  forest region (red) that we wish to model. To its corresponding mock spectrum, we add metals observed in the  $\lambda_{\text{rest}} \approx 1260\text{--}1390$  Å region of a lower-redshift ( $z_{\text{qso}} = 2.0$ ) quasar (blue region in top panel).

statistics of the Ly $\alpha$  forest. In high-resolution data, this contamination is usually treated by directly identifying and masking the metal absorbers, although in the presence of line blending it is unclear how thorough this approach can be.

With the lower S/N and moderate resolution of the BOSS data, direct metal identification and masking is not a viable approach. Furthermore, most of the weak metal absorbers seen in high-resolution spectra are not resolved in the BOSS data.


Rather than removing metals from the BOSS Ly $\alpha$  forest spectra, we instead add metals as observed in lower-redshift quasar spectra. In other words, we add absorbers observed in the rest-frame  $\lambda_{\text{rest}} \approx 1260\text{--}1390$  Å region of lower-redshift quasars with  $1 + z_{\text{qso}} \approx (1216 \text{ Å}/1300 \text{ Å})(1 + \langle z \rangle)$ , such that the observed wavelengths are matched to the Ly $\alpha$  forest segment with average

redshift ( $\langle z \rangle$ ). Figure 11 is an illustration that illustrates this concept. This method makes no assumption about the nature of the metal absorption in the Ly $\alpha$  forest, and includes all resolved metal absorption spanning the whole range of redshifts down to  $z \sim 0$ . The disadvantage of this method is that it does not include metals with intrinsic wavelengths  $\lambda \lesssim 1300$  Å, but the relative contribution of such metal species toward the transmission PDF should be small<sup>22</sup> since most of the metal contamination comes from low-redshift ( $z \lesssim 2$ ) C IV and Mg II.

We use a metal catalog generated by B. Lundgren et al. (in preparation; see also Lundgren et al. 2009), which lists absorbers in SDSS (Schneider et al. 2010) and BOSS quasar spectra (Pâris et al. 2012)—the SDSS spectra were included in order to increase the number of  $z_{\text{qso}} \approx 1.9\text{--}2.0$  quasars needed to introduce metals into the  $\langle z \rangle = 2.3$  Ly $\alpha$  forest mock spectra, which are not well sampled by the BOSS target selection (Ross et al. 2012). We emphasize that we work with the “raw” absorber catalog, i.e., the individual absorption lines have not been identified in terms of metal species or redshift. For each quasar, the catalog provides a line list with the observed wavelength, EW ( $W_r$ ), FWHM, and detection S/N,  $W_r/\sigma_{W_r}$ . To ensure a clean catalog, we use only  $W_r/\sigma_{W_r} \geq 3.5$  absorbers in the catalog that were identified from quasar spectra with S/N > 15 per angstrom redward of Ly $\alpha$ . The latter criterion ensures that even relatively weak lines (with EW  $\gtrsim 0.5$  Å) are accounted for in our catalog. Figure 12 shows an example of the lower-redshift quasar spectra that we use for the metal modeling.

However, we want to add a smooth model of the metal-line absorption to add to our mock spectra, rather than adding in a noisy spectrum. We therefore use a simple model as follows: For each Ly $\alpha$  forest segment we wish to model at redshift  $\langle z \rangle$ , we select an absorber line-list from a random quasar with  $1 + z_{\text{qso}} \approx (1216 \text{ Å}/1300 \text{ Å})(1 + \langle z \rangle)$ . We next assume that all resolved metals in the SDSS/BOSS spectra are saturated and

<sup>22</sup> Si III is an obvious exception, although we will later account for this omission in our error bars (Section 5.3).



**Figure 12.** Continuum-normalized spectrum of a BOSS quasar showing the metal absorbers in the  $1300 \text{ Å} < \lambda_{\text{rest}} < 1390 \text{ Å}$  “sideband” region, which would be used to add metals to  $\langle z \rangle = 2.6$  mock Ly $\alpha$  forest spectra. The red curve shows our metal model for this spectrum, generated from the observed wavelengths and equivalent widths in the absorber catalog generated by the automatic algorithm of Lundgren et al. (2009). We also assume that the absorbers all lie on the saturated portion of the curve-of-growth and have  $\tau_0 = 3$ , with the equivalent width (labeled above each absorption line) proportional to the  $b$ -parameter. The model absorption profiles represented by the red curve would be added to our mock Ly $\alpha$  forest spectra. We have chosen to plot this particular “sideband” because it has more absorbers than average—the typical spectrum has less metal absorption than this.

thus in the flat regime of the curve-of-growth. The EW is then given by

$$W_r \approx \left(\frac{2b}{c}\right) \sqrt{\ln(\tau_0/\ln 2)}, \quad (15)$$

where  $\tau_0$  is the optical depth at line center,  $b$  is the velocity width and  $c$  is the speed of light. In the saturated regime,  $W_r$  is mostly sensitive to changes in  $b$  while being highly insensitive to changes in  $\tau_0$ . We can thus adopt  $\tau_0$  as a global constant and solve for  $b$ , given the  $W_r$  of each listed absorber in the selected “sideband” quasar. We have found that  $\tau_0 = 3$  provides a good fit for most of the absorbers.

We then add the Gaussian profile into our simulated optical depth skewers:

$$\tau = \tau_0 \exp \left[ - \left(\frac{c}{b}\right) \left(\frac{\Delta\lambda}{\lambda}\right)^2 \right] \quad (16)$$

centered at the same observed wavelength,  $\lambda$ , as the real absorber. The red curve in Figure 12 shows our model for the observed absorbers, using just the observed wavelength,  $\lambda$ , and EW,  $W_r$ , from the absorber catalog.

Our method for incorporating metals is somewhat crude since one should, in principle, first deconvolve the spectrograph resolution from the input absorbers, and then add the metal absorbers into the mock spectra prior to convolving with the BOSS spectral resolution. In contrast, we fit  $b$ -parameters to the absorber catalog without spectral deconvolution, and therefore these  $b$ -parameters can be thought of as combinations of the true absorber width,  $b_{\text{abs}}$ , and the spectral dispersion,  $\sigma_{\text{disp}}$ , i.e.,  $b^2 \sim b_{\text{abs}}^2 + \sigma_{\text{disp}}^2$ . While technically incorrect, this seems reasonable since the template quasar spectra and forest spectra that we are attempting to model both have approximately the same resolution, and in practical terms this *ad hoc* approach does seem to be able to reproduce the observed metals in the lower-redshift quasar spectra (Figure 12). The other possible criticism of our approach is that it does not incorporate weak metal absorbers, although we attempted to mitigate this by setting a very high S/N threshold on the template quasars for the metals. However, we have checked that such weak metals do not significantly change the forest PDF (and indeed metals in general do not seriously affect the PDF, compare with Figure 9(c)).

We also tried adding metals with similar redshifts to—and correlated with—forest absorbers (e.g., absorption by Si II and Si III) measured in Pieri et al. (2010, 2014) using a method described in the Appendix of Slosar et al. (2011). We found a negligible impact on the transmission PDF owing mainly to the fact that these correlated metals contribute only  $\sim 0.3\%$  to the overall flux decrement, so we neglect this contribution in our subsequent analysis.

#### 4.5. Pixel Noise

It is non-trivial to introduce the correct noise to a simulated Ly $\alpha$  forest spectrum: given a noise estimate from the observed spectrum, one needs to first ensure that the mock spectrum has approximately the same flux normalization as the data. This is challenging, as the Ly $\alpha$  forest transmission at any given pixel, which ranges from 0 to 1, will vary considerably between the simulated spectrum and the real data.

The simplest method of adding noise to a mock spectrum is simply to introduce Gaussian deviates using the pipeline noise estimate for each spectrum—this was essentially the method

used by Desjacques et al. (2007) and the BOSS mocks described in Font-Ribera et al. (2012). However, with the MCMC co-addition procedure described in Section 3.1, we are in a position to model the noise in a more robust and self-consistent fashion.

Recall that the MCMC procedure returns posterior probabilities for two quantities: the true underlying spectral flux density,  $\mathcal{F}_\lambda$ , and the four free parameters  $A_j$ , which parameterize the noise in each spectrum. This estimate of the  $A_j$  from each quasar spectrum allows us to accurately model the pixel noise using Equation (5).

The MF-PCA method (Section 3.2) produces an estimate of the quasar continuum,  $c$ , providing approximately the correct flux level at each point in the spectrum. We can now multiply  $c$  with the simulated Ly $\alpha$  forest transmission spectra,  $F$ , which had already been smoothed to the same dispersion as its real counterpart (the estimated quasar continuum is already at approximately the correct smoothing, since it was fitted to the observed spectrum).

This procedure produces a noiseless mock spectrum with the correct flux normalization and smoothing. We can now generate noisy spectra corresponding to a given BOSS quasar, using our MCMC noise estimation described in Section 3.1. First, we substitute our mock spectrum as  $\mathcal{F}_\lambda$  into Equation (5), and then combine the  $A_j$  noise parameters (estimated through our MCMC procedure) as well as the calibration vectors  $S_{\lambda,i}$  and sky estimates  $s_{\lambda,i}$ . This lets us generate self-consistent noise vectors corresponding to each individual exposure that make up the mock quasar spectrum,  $\sigma_{\lambda,i}$ . The noise vectors are then used to draw random Gaussian deviates that are added to the mock spectrum, on a per-pixel basis, to create the mock spectral flux density,  $f_{\lambda,i}$ . Finally, we combine these individual mock exposures into the optimal spectral flux density for the mock spectrum, through the expression (see the Appendix):

$$f_{\text{opt},\lambda} \equiv \frac{1}{\sigma_{\text{opt},\lambda}^2} \sum_i \frac{f_{\lambda,i}}{\sigma_{\lambda,i}^2}, \quad (17)$$

where


$$\frac{1}{\sigma_{\text{opt},\lambda}^2} \equiv \sum_i \frac{1}{\sigma_{\lambda,i}^2}. \quad (18)$$

Figure 9(c) illustrates the effect of adding pixel noise to the smoothed Ly $\alpha$  forest transmission PDF. As expected, this scatters a significant fraction of pixels to  $F > 1$ , and also to  $F < 0$  to a smaller extent.

#### 4.6. Continuum Errors

With the noisy mock spectrum in hand (see, e.g., bottom panel of Figure 13), we can self-consistently include the effect of continuum errors into our model transmission PDFs by simply carrying out our MF-PCA continuum-fitting procedure on the individual noisy mock spectra. Dividing out the mock spectra with the new continuum fits then incorporates an estimate of the continuum errors (estimated by Lee et al. 2012 to be at the  $\sim 4\%$ – $5\%$  rms level) into the evaluated model transmission PDF. This estimated error includes uncertainties stemming from the estimation of the quasar continuum shape due to pixel noise, as well as the random variance in the mean Ly $\alpha$  forest absorption in individual lines-of-sight.

Note that regardless of the overall mean absorption in the mock spectra (i.e., inclusive of our models for metals, LLSs, and mean forest absorption—see Section 5.4), we always use  $\langle F \rangle_{\text{cont}}(z)$ , the same input mean transmission derived from



**Figure 13.** Simulating the noise properties and continuum errors of a BOSS quasar. The top panel shows the observed spectrum of a BOSS quasar, and its associated continuum fit,  $c$ , in blue. The middle panel shows the simulated transmission spectra (after adding LLS, smoothing, and adding metals) multiplied by the quasar continuum fitted to the true spectrum. In the lower panel, we have added noise to the mock spectrum using the noise parameters estimated from the true spectrum (see Section 3.1). A new continuum,  $c'$ , (red) is re-fitted to the noisy mock spectrum. The difference between new continuum  $c'$  and “true” continuum,  $c$ , of the mock (blue) introduces continuum errors to our model. The vertical dotted lines indicate the range of pixels that contribute to the  $\langle z \rangle = 3.0$  subsample in our transmission PDF; a small segment between  $(1 + z_{\text{qso}})1040 \text{ \AA} = 4461 \text{ \AA}$  and  $(1 + 2.75)1216 \text{ \AA} = 4560 \text{ \AA}$  also contributes to the  $\langle z \rangle = 2.6$  bin.

Becker et al. (2013; described in Section 3.2) to fit the continua in both the data and mock spectra. While the overall absorption in our fiducial model is consistent with that from Becker et al. (2013), as we shall see later, the shape of the transmission PDF retains information on the true underlying mean transmission even if fitted with a mean flux regulated continuum with a wrong input  $\langle F \rangle(z)$ .

The effect of continuum errors on the transmission PDF is shown in Figure 9(e): like pixel noise, it degrades the peak of the PDF, but only near  $F \sim 1$ .

## 5. MODEL REFINEMENT

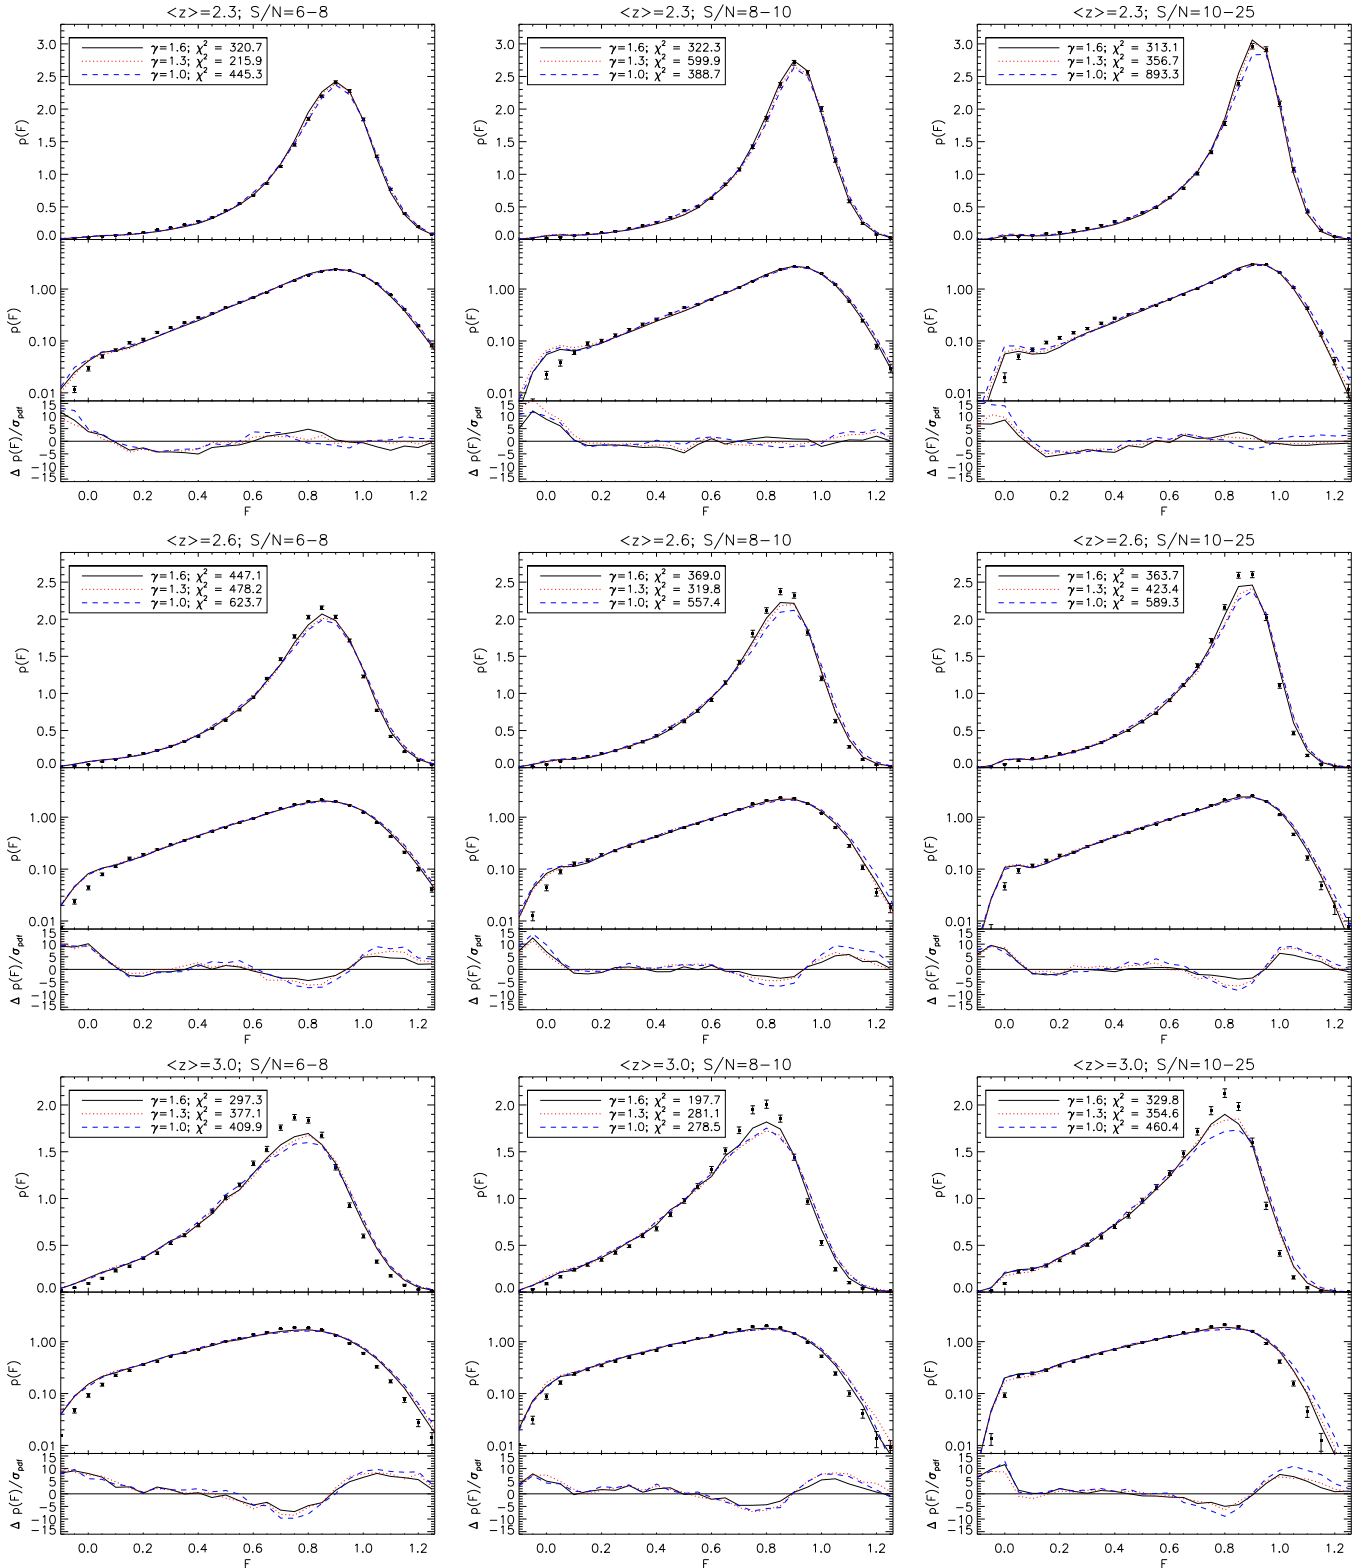
In an ideal world, one would like to do a blind analysis by generating the transmission PDF model (Section 4) in isolation from the data, before “unblinding” to compare with data—this would then in principle yield results free from psychological bias in the model building. However, as we shall see in Section 5.1, this does not give acceptable fits to the data so we have to instead modify our model to yield a better agreement, in particular our LLS model (Section 5.2) and assumed mean transmission (Section 5.4).

### 5.1. Initial Comparison with T\_REF Models

For each of our nine hydrodynamical simulations (sampling three points each in  $T_0$  and  $\gamma$ ), we determine the transmission PDF from the Ly $\alpha$  forest mock spectra that include the effects described in the previous section, for the various redshift and S/N subsamples in which we had measured the PDF in BOSS (Section 3.3). In Figure 14, we show the transmission PDFs for all our redshift and S/N subsamples in BOSS, compared with the corresponding simulated transmission PDFs from the T\_REF simulation with  $\gamma = [1.0, 1.3, 1.6]$ . Note that the error bars shown are the diagonal elements of the covariance matrix estimated through bootstrap resampling on the data.

At first glance, the model transmission PDFs seem to be a reasonable match for the data, especially considering we have carried out purely forward modeling without fitting for any parameters. However, when comparing the “pull,”  $(p_{\text{data},i} - p_{\text{model},i})/\sigma_{p,i}$ , between the data and model (bottom panels of Figure 14), we see significant discrepancies in part due to the extremely small bootstrap error bars. Nevertheless, it is gratifying to see that the shape of the residuals is relatively consistent across the different S/N subsamples at fixed redshift and  $\gamma$ , since this indicates that our spectral noise model is robust.





**Figure 14.** Initial comparison between the transmission PDFs observed from BOSS Ly $\alpha$  forest data (error bars) and simulated PDFs generated from the T\_REF hydrodynamical simulations (curves) with the method described in Section 4; each row is at the same redshift, while the different columns display the different S/N cut. The points with the error bars are the PDFs measured from the BOSS data (estimated from bootstrap resampling, while the black, dotted red, and dashed blue curves denote simulated PDFs with  $\gamma = [1.5, 1.3, 1.0]$ , respectively). The top and middle panels show the transmission PDFs with linear and logarithmic axes, while the lower panels show the pull, i.e., residuals between the simulated PDF and the data PDF, divided by the error. The  $\chi^2$  values indicated in these plots are for 24 dof, and clearly indicate unacceptable fits to the data—modifications to the model are required.

We proceed to quantify the differences between the simulated transmission PDFs,  $p_{\text{model}}$ , and observed transmission PDFs,  $p_{\text{data}}$ , with the  $\chi^2$  statistic:

$$\chi^2 = \sum_{ij} (p_{\text{model},i} - p_{\text{data},i})^T C_{ij}^{-1} (p_{\text{model},j} - p_{\text{data},j}), \quad (19)$$

where we use the bootstrap error covariance matrix,  $C_{\text{boot}}$ . Note that we also include a bootstrap error term that accounts for the sample variance in the model transmission PDFs, since our pipeline for generating mock spectra is too computationally expensive to include sufficiently large amounts of skewers to fully beat down the sample variance in the models.<sup>23</sup>

We limit our model comparison to the range  $-0.1 \leq F \leq 1.2$ , i.e., 27 transmission bins with bin width  $\Delta(F) = 0.05$ . This range covers pixels that have been scattered to “unphysical” values of  $F < 0$  or  $F > 1$  due to pixel noise, as is expected from the low-S/N of our BOSS data, and also captures  $>99.8\%$  of the pixels within each of our data subsets. In particular, it is important to retain the bins with  $F > 1$  because the  $F \sim 1$  transmission bins are highly sensitive to  $\gamma$  (Lee 2012) and therefore we want to fully sample that region of the PDF even if it will require careful modeling of pixel noise and continuum errors.

There are two constraints on all our transmission PDFs: the normalization convention

$$\int p(F) dF = 1 \quad (20)$$

and the imposition of the same mean transmission due to the mean flux regulated continuum-fitting

$$\int F p(F) dF = \langle F \rangle_{\text{cont}} \quad (21)$$

such that all the mock spectra have the same absorption,  $\langle F \rangle_{\text{cont}}(z)$ . This is because the mock spectra have been continuum-fitted (Section 4.6) in exactly the same way as the BOSS spectra, which assumes the same mean Ly $\alpha$  transmission inferred from the Becker et al. (2013) measurements (Section 3.2). The “true” optically thin mean transmission,  $\langle F \rangle_{\text{Ly}\alpha}$ , imposed on the simulation skewers is in principle a different quantity from  $\langle F \rangle_{\text{cont}}$ , since the latter includes contribution from metal contamination and optically thick LLSs.

This leaves us with  $\nu = 27 - 1 - 2 = 24$  degrees of freedom (dof) in our  $\chi^2$  comparison. The  $\chi^2$  for all the models shown in Figure 14 are shown in the corresponding figure legends.

In this initial comparison, the  $\chi^2$  values for the models in Figure 14 are clearly unacceptable: we find  $\chi^2 \gtrsim 200$  for 24 dof in all cases. However, it is interesting to note that the  $\gamma = 1.6$  or  $\gamma = 1.3$  models are preferred at all redshifts and S/N cuts. Note that the S/N = 8–10 subsamples (middle column in Figure 14) tends to have a slightly better agreement between model and data compared to the other S/N cuts at the same redshift: this simply reflects the smaller quantity of data of the subsample (compare with Table 1) and hence larger bootstrap errors.

A closer inspection of the residuals in Figure 14 indicate that there are two major sources of discrepancy between the models and data: first, at the low-transmission end, we underproduce pixels at  $0.1 \lesssim F \lesssim 0.4$  while simultaneously overproducing

$F \lesssim 0.1$  pixels, especially at  $\langle z \rangle = 2.3$  and  $\langle z \rangle = 2.6$ . This seems to affect all  $\gamma$  models equally. Pieri et al. (2014) found that at BOSS resolution, pixels with  $F \lesssim 0.3$  come predominantly from saturated Ly $\alpha$  absorption from LLS. We therefore investigate possible modifications to our LLS model in Section 5.2.

The other discrepancy in the model transmission PDFs manifests at the higher-transmission end in the  $\langle z \rangle = 2.6$  and  $\langle z \rangle = 3$  subsamples, where we see a sinusoidal shape in the residuals at  $F > 0.6$  that appears consistent across different S/N. This portion of the transmission PDF depends on both  $\gamma$  and, as we shall see, on the assumed mean transmission  $\langle F \rangle(z)$ , which we shall discuss in more detail in Section 5.4.

Finally, our transmission PDF model includes various uncertainties in the modeling of metals, LLSs, and continuum-fitting which have not yet been taken into account. In Section 5.3, we will estimate the contribution of these uncertainties, by means of a Monte Carlo method, in our error covariances.

## 5.2. Modifying the LLS Column Density Distribution

With the moderate spectral resolution of BOSS, there are few individual pixels in the optically thin Ly $\alpha$  forest that reach transmission values of  $F \lesssim 0.4$ . Such low-transmission pixels are typically due to either the blending of multiple absorbers (see, e.g., Figure 2 in Pieri et al. 2014), or optically thick systems (see Figure 10 in this paper).

As we have seen in Figure 14, at low-transmission values the discrepancy between data and model has a distinct shape, which is particularly clear at  $\langle z \rangle = 2.3$ : the models underproduce pixels at  $0.1 \lesssim F \lesssim 0.4$  while at the same time overproducing saturated pixels with  $F \approx 0$ .

To resolve this particular discrepancy would therefore require either drastically increasing the amount of clustering in the Ly $\alpha$  forest, or modifying our assumptions on the LLSs in our mock spectra. The first possibility seems rather unlikely since the Ly $\alpha$  forest power on relevant scales are well-constrained (Palanque-Delabrouille et al. 2013), and would in any case require new simulation suites to address—beyond the scope of this paper.


On the other hand, it is not altogether surprising that our fiducial column density distribution (Section 4.2)—which was measured at  $z \approx 3.7$  (Prochaska et al. 2010)—do not reproduce the BOSS data at  $\langle z \rangle = 2.3$ –2.6. We therefore search for an LLS model that better describes the low-transmission end of the BOSS Ly $\alpha$  forest. Looking at the  $\langle z \rangle = 2.3$  PDFs in Figure 14, we see that our fiducial model *overproduces* pixels at  $F = 0$ , yet is deficient at slightly higher  $F$ . This suggests that our model is overproducing super-LLS ( $N_{\text{H1}} > 10^{19} \text{ cm}^{-2}$ ) that contribute large absorption troughs with  $F = 0$ , while not providing sufficient lower-column density absorbers that can individually reach minima of  $0.1 \lesssim F \lesssim 0.4$  when smoothed to BOSS resolution. In other words, our fiducial model appears to have an excessively “top-heavy” LLS column density distribution.

For a change, we will try an LLS column density distribution with a more ample bottom end, using the steepest power laws within the  $1\sigma$  limits estimated by Prochaska et al. (2010):

$$f(N_{\text{H1}}) = \begin{cases} k_1 N_{\text{H1}}^{-1.2} & \text{if } 10^{17.5} < N_{\text{H1}} < 10^{19.0} \\ k_2 N_{\text{H1}}^{-1.4} & \text{if } 10^{19.0} < N_{\text{H1}} < 10^{20.3} \end{cases} \quad (22)$$

We use the same  $l_{\text{LLS}}(z)$  as before, and obey the integral constraints from Prochaska et al. (2010) that demand that the ratios of  $\int f(N_{\text{H1}}) dN_{\text{H1}}$  between the two column-density

<sup>23</sup> We aim for three to four times more mock spectra than in the corresponding data sample, but later when we have to compute large model grids we are limited to models with the same size as the data.




**Figure 15.** LLS and pLLS column-density power-law distributions used in our initial model (black; Section 4.2) and steeper modification (red; Section 5.2). The distributions are normalized assuming the overall LLS incidence rate at  $z = 2.25$  (compare with Equation 10). The vertical dashed lines denote the  $N_{\text{HI}} = 10^{17.5} \text{ cm}^{-2}$  boundary between pLLS and LLS, and  $N_{\text{HI}} = 10^{19} \text{ cm}^{-2}$  boundary between LLS and super-LLS. The shaded regions show the range of possible distributions as determined by Prochaska et al. (2010), but there are few robust constraints in the  $10^{16.5} \text{ cm}^{-2} \leq N_{\text{HI}} \leq 10^{17.5} \text{ cm}^{-2}$  pLLS regime. The “initial” distribution was used in the preliminary data comparisons in Section 5.1, but all subsequent analysis (after Section 5.2) assumes the “steep” distribution.

regimes be fixed. This gives us  $k_1 = 10^{2.819}$  and  $k_2 = 10^{7.039}$ , although the new distribution is no longer continuous at  $N_{\text{HI}} = 10^{19} \text{ cm}^{-2}$ . This new distribution is illustrated by the red power laws in Figure 15.

Another change we have made is to the partial LLS model, which was possibly too conservative in the fiducial model. Instead of extrapolating from the LLS distribution, we now adopt the pLLS power-law slope of  $\beta_{\text{pLLS}} = -2.0$  inferred from the total mean free path to ionizing photons by Prochaska et al. (2010). This dramatically increases the incidence of pLLS in our spectra relative to LLS: we now have  $l_{\text{pLLS}} = 1.8 l_{\text{LLS}}$ , where  $l_{\text{LLS}}$  is the same value we used previously (Equation (10)). This increase, while large, is not unreasonable in light of the large uncertainties in direct measurements on the H I column-density distribution from direct Ly $\alpha$  line-profile fitting (e.g., Janknecht et al. 2006; Rudie et al. 2013). Note also that even this increased pLLS incidence only amounts to, on average, less than one pLLS per quasar ( $\Delta(z) \sim 0.3\text{--}0.4$  per quasar at our redshifts).

We found that while increasing the number of pLLS relieves the tension between data and model at  $0.1 \lesssim F \lesssim 0.4$ , it does not resolve the excess at the fully absorbed  $F \approx 0$  pixels in the models. However, changing the  $b$ -parameter of the LLS and pLLS from our original fiducial value of  $b = 70 \text{ km s}^{-1}$  modifies the PDF in a way that improves the agreement. This is a reasonable step, since the effective  $b$ -parameter is otherwise observationally ill-constrained for the LLS and pLLS populations. This is because LLSs are typically complexes of multiple systems separated in velocity space, and while there have been analyses of the  $b$ -parameter in these individual components, the “effective”  $b$ -parameter for complete LLS systems has never been quantified to our knowledge.

We therefore search for the best-fit  $b$ -parameter with respect to the T\_REF,  $\gamma = 1.3$  model at  $(z) = 2.3$ , focusing primarily on the agreement in the  $0 \leq F \leq 0.4$  bins (Figure 16). Our choice of model for this purpose should not significantly affect our subsequent conclusions regarding the IGM temperature–density



**Figure 16.** Variation of the transmission PDF as a function of LLS  $b$ -parameter. All model transmission PDFs here are computed from the T\_REF,  $\gamma = 1.6$  model assuming the revised pLLS/LLS distribution described in Section 5.2 (curves), compared with the S/N = 6–8 BOSS transmission PDF at  $\langle z \rangle = 2.3$  (error bars). The quoted  $\chi^2$  values are for 24 dof, and evaluated using only bootstrap error covariances. We find that  $b = 45 \text{ km s}^{-1}$  gives the best fit to the data.


slope, since there is little sensitivity toward the latter in the relevant low-transmission bins (compare with Figure 14). However, there will be some degeneracy between the LLS  $b$ -parameter and  $T_0$  (Figure 17) since changing the latter does somewhat change the low-transmission portion of the PDF—we will come back to this point in Section 7.

As shown in Figure 16, a value of  $b = 45 \text{ km s}^{-1}$  gives the best agreement with the data at  $0 \leq F \leq 0.4$ . This yields  $\chi^2 = 116$  for 24 dof, which is dramatically improved over that quoted in Figure 14, but still not quite a good fit. In the subsequent results, we will adopt this steeper pLLS/LLS model and  $b$ -parameter as the fiducial model in our analysis, and will correspondingly decrease the dof in our  $\chi^2$  analysis to account for the fitting of  $b$ .

Note that while significantly improving the PDF fit, this new  $b$ -parameter still does not give a perfect fit to the low-transmission ( $F < 0.4$ ) end. This is probably due to the simplified nature of our LLS model, which neglects the finite distribution of  $b$ -parameters and internal velocity dispersion of individual components. These properties are currently not well known, and it seems likely that an improved model would allow a better fit to the low-transmission end of the PDF.

### 5.3. Estimation of Systematic Uncertainties

While we have estimated the sample variance of our BOSS transmission PDFs by bootstrap resampling on the spectra, there are significant uncertainties associated with each component of our transmission PDF model as described above, e.g., the LLS



**Figure 17.** Variation of the transmission PDF as a function of the IGM temperature at mean density,  $T_0$ . All model transmission PDFs here have the same temperature–density relationship,  $\gamma = 1.6$ , and are compared with the  $S/N = 6\text{--}8$  BOSS transmission PDF at  $\langle z \rangle = 2.3$  (error bars). The quoted  $\chi^2$  values are for 23 dof. Note that in these models we have already implemented the improved LLS/pLLS model described in Section 5.2, hence the much improved  $\chi^2$  values compared to those quoted in Figure 14.

incidence rate and level of continuum error. These uncertainties can be incorporated into a systematics covariance matrix,  $\mathbf{C}_{\text{sys}}$  that can then be added to the bootstrap covariance,  $\mathbf{C}_{\text{boot}}$ , when computing the model likelihoods. This requires assuming that  $\mathbf{C}_{\text{sys}}$  and  $\mathbf{C}_{\text{boot}}$  are uncorrelated, and that the errors are Gaussian distributed.

We adopt a Monte Carlo approach to estimate  $\mathbf{C}_{\text{sys}}$  by generating 200 model transmission PDFs that randomly vary the systematics. We then evaluate the covariance of the transmission PDFs,  $p_i$ , relative to the fiducial model,  $p_{\text{ref},i}$  at each transmission bin  $i$ . This allows us to construct a covariance matrix with the elements

$$C_{\text{sys},ij} = \langle (p_i - p_{\text{ref},i})(p_j - p_{\text{ref},j}) \rangle \quad (23)$$

that encompasses the errors from the uncertainties in the LLS model, metal absorption, and continuum scatter. Note that estimation of systematic uncertainties is typically a subjective process, and for most of these contributions we can only make educated guesses as to their uncertainty.

Our Monte Carlo iterations sample the various components of our model as follows:

1. *LLS incidence.* We sample the uncertainty in the power-law exponent  $\gamma_{\text{LLS}}$  of the redshift evolution in LLS incidence rate (Equation (10)), which is  $\sigma_{\gamma_{\text{LLS}}} \pm 0.21$  as reported by Ribaldo et al. (2011). We assume this uncertainty is Gaussian and draw  $l_{\text{LLS}}(z)$  accordingly. This primarily affects the low-flux regions  $-0.1 \lesssim F \lesssim 0.3$  of the PDF.


2. *Partial-LLS slope.* Our choice of slope for the distribution of partial LLS ( $N_{\text{HI}} < 10^{17.5} \text{ cm}^{-2}$  absorbers) is from an indirect constraint with significant uncertainty (Prochaska et al. 2010). We therefore vary the pLLS slope around the fiducial  $\beta_{\text{pLLS}} = -2.0$  by  $\pm 0.5$  assuming a flat prior in this range, which primarily alters the  $0 \lesssim F \lesssim 0.4$  portion of the PDF since pLLS typically do not saturate at BOSS resolution.
3. *LLS  $b$ -parameters.* Also in the previous section, we found that a global  $b$ -parameter of  $b = 45 \text{ km s}^{-1}$  gives the best agreement with the data, but this is an *ad hoc* approach with significant uncertainties. In our Monte Carlo sampling we therefore adopt a conservative  $b = 45 \text{ km s}^{-1} \pm 20 \text{ km s}^{-1}$  with a uniform prior. This primarily affects the PDF at  $-0.1 \lesssim F \lesssim 0.4$  as can be seen in Figure 16.
4. *Intervening metals.* Although we used an empirical method to model intervening metals (Section 4.4), we may have missed metals with rest wavelengths  $\lambda \lesssim 1300 \text{ \AA}$ . Furthermore, we have a relatively small set ( $\sim 300\text{--}400$ ) of “template” quasars from which our metal model is derived, which may contribute some sampling variance. We therefore guess at a Gaussian error of  $\pm 30\%$  for the metal incidence rate. This modulates the extent to which metals pulls the overall PDF toward lower  $F$ -values (compare with Figure 9(c)).
5. *Continuum errors.* The overall rms scatter in our continuum estimation also affect the flux PDF (Figure 9(e)). This can be varied in our model by rescaling the quantity  $c'(\lambda)/c(\lambda) - 1$ , where  $c$  is the “true” continuum used to generate the mock spectrum, while  $c'$  is the model continuum that we subsequently fit (Figure 13). For each iteration in our Monte Carlo systematics estimation, we dilate or reduce  $c'(\lambda)/c(\lambda) - 1$  by a Gaussian deviate assuming  $\pm 20\%$  scatter. This primarily affects the high-transmission ( $F > 0.8$ ) end of the PDF.

For these Monte Carlo iterations, we used the identical thermal model ( $\gamma = 1.6$ , T\_REF) as well as fixed the same random number seeds used for the selection of simulation skewers and generation of noise vectors in our spectra, in order to ensure that the only variation between the different iterations are from the randomly sampled systematics. Figure 18 shows 50 of these Monte Carlo iterations on the transmission PDF for the  $\langle z \rangle = 2.3$ ,  $S/N = 8\text{--}10$  subsample.

Figure 19 shows an example of the systematic contribution to the covariance matrix. The overall amplitude of the systematic contribution is considerably higher than that estimated from the bootstrap resampling (compare with Figure 7), indicating that we are in the systematics-limited regime. We also see significant anti-correlations at almost the same level as the positive correlations, which are due mostly to correlations between transmission bins on either side of “pivot points” as the transmission PDF varies from the systematics—these anti-correlations will somewhat counteract the increased size of the diagonal components. In the subsequent analysis, we will use an error covariance matrix,  $\mathbf{C} = \mathbf{C}_{\text{boot}} + \mathbf{C}_{\text{sys}}$ , in which the systematics covariance matrix estimated in this subsection is added to the bootstrap covariance matrix (described in Section 3) estimated from the BOSS transmission PDFs.

We have at this point yet to address one more parameter that can significantly change the shape of our model transmission PDFs, namely the Ly $\alpha$  forest mean transmission assumed in the mock spectra,  $\langle F \rangle_{\text{Ly}\alpha}$ . However, this is an important astrophysical parameter which we did not want to treat as a






**Figure 18.** Gray curves show 50 model transmission PDFs with a random sampling of different LLS incidence rates, metal absorption, and continuum scatter, evaluated for the  $\langle z \rangle = 2.3$ ,  $S/N = 8-10$  BOSS subsample and using the T\_REF simulation with  $\gamma = 1.6$ . The red curve shows the transmission PDF at our fiducial level of LLS incidence, metal absorption, and continuum scatter. The top panel has a linear abscissa, while the lower panel has a logarithmic abscissa.

“systematic,” so the next subsection will describe our treatment of  $\langle F \rangle_{\text{Ly}\alpha}$ .

#### 5.4. Modifying the Mean-transmission


In the initial comparison of the model transmission PDFs shown in Figure 14, the models show a discrepancy with the data at higher transmission bins  $F \gtrsim 0.6$ . Such differences can be alleviated by varying the mean transmission of the pure Ly $\alpha$  forest,  $\langle F \rangle_{\text{Ly}\alpha} \equiv \exp(-\tau_{\text{Ly}\alpha})$ , i.e., ignoring the contribution from metals and LLS. This quantity can be varied directly in the simulation skewers (Section 4.1). When we vary  $\langle F \rangle_{\text{Ly}\alpha}$  in the simulations, the quantity  $\langle F \rangle_{\text{cont}}$ , which is used to normalize the continuum level of the mock quasar spectrum, is always kept fixed to  $\langle F \rangle_{\text{eff}}(z) = \exp[-(\tau_{\text{Ly}\alpha} + \tau_{\text{metals}} + \tau_{\text{LLS}})]$  as derived from Becker et al. (2013; see Section 3.2). However, since we are applying the same  $\langle F \rangle_{\text{cont}}$  to both the real and mock spectra,  $\langle F \rangle_{\text{cont}}$  can be best thought of as a normalization that does not actually need to match  $\langle F \rangle_{\text{eff}}$ . Once both the real and mock spectra have been normalized by  $\langle F \rangle_{\text{cont}}$ , the transmission PDF retains information on the respective contributions from the Ly $\alpha$  forest, metals, and LLSs regardless of the assumed  $\langle F \rangle_{\text{cont}}$ , because these contributions affect the shape of the PDF in different ways. In principle, it is possible to vary these all components to infer their relative contributions, but due to the crudeness of our metal and LLS models, we choose to have only  $\langle F \rangle_{\text{Ly}\alpha}$  as a free parameter while keeping  $\langle F \rangle_{\text{metals}} = \exp(-\tau_{\text{metals}})$  and  $\langle F \rangle_{\text{LLS}} = \exp(-\tau_{\text{LLS}})$  fixed. The possible variation of these latter two components are instead incorporated into the systematic uncertainties determined in Section 5.3. The effect of varying  $\langle F \rangle_{\text{Ly}\alpha}$  is illustrated in Figure 20, where we plot the same IGM model with different underlying values of  $\langle F \rangle_{\text{Ly}\alpha}$  in the simulation skewers while keeping fixed the contribution from metals, LLSs, etc.



**Figure 19.** Top: 2D density plot of the error covariance matrix representing our systematic uncertainties in the LLS incidence rate, pLLS column-density distribution, LLS  $b$ -parameter, metal absorption, and continuum scatter, as estimated through the Monte Carlo method described in Section 5.3. The bottom plot shows the corresponding correlation function. This particular covariance matrix was estimated for the  $\langle z \rangle = 2.6$ ,  $S/N = 8-10$  subsample, and the values in the covariance have been multiplied by  $10^4$  for clarity.

We therefore explore a range of  $\langle F \rangle_{\text{Ly}\alpha}$  around the vicinity of that estimated by Becker et al. (2013),  $\langle F \rangle_{\text{Ly}\alpha, \text{B13}}$ , and at each value of  $\langle F \rangle_{\text{Ly}\alpha}$  evaluate the  $\chi^2$  summed over all the  $S/N$  subsamples for each  $\langle z \rangle$  and  $\gamma$  combination. In addition, we now adopt the updated LLS/pLLS model described in Section 5.2, while the  $\chi^2$  evaluation now uses the full covariance matrix including both the bootstrap and systematics (Section 5.3) uncertainties to compare with the transmission PDFs measured from the BOSS data.

The models are compared with the BOSS data as we vary  $\langle F \rangle_{\text{Ly}\alpha}$ , and for each  $\langle F \rangle_{\text{Ly}\alpha}$  we compute the total chi-squared summed over all three  $S/N$  subsamples, where each subsample contributes  $27 - 1 - 2 = 24$  dof (compare with Equations (20) and (21)) along with a further reduction of one dof since we have effectively fitted for the LLS  $b$ -parameters in Section 5.2, for a total of  $\nu = 71$  dof. The result of this exercise is shown in




**Figure 20.** Variation of the model transmission PDFs (curves) with respect to changing the mean transmission,  $\langle F \rangle_{\text{Ly}\alpha}$ , of the Ly $\alpha$  forest simulations. The model PDFs were generated from the  $\gamma = 1.6$ , T\_REF model, while the error bars show the corresponding transmission PDFs from BOSS data. In the bottom panel, the dashed horizontal lines indicate  $\pm 1\sigma$  discrepancies between models and data, although we caution against “chi-by-eye” due to the significantly non-diagonal covariances in the errors. The central  $\langle F \rangle_{\text{Ly}\alpha}$  value shown here corresponds to that estimated by Becker et al. (2013), while the other two are evaluated at  $\pm 1\sigma$  of their reported errors. The mean transmission value,  $\langle F \rangle_{\text{cont}}$ , assumed in the mean flux regulated continuum fitting is constant in all cases. Note that the  $\chi^2$  values, which are for 23 dof, are much improved over the previous data comparisons, since they now include the improved LLS/pLLS model as well as the full covariance matrix including systematic uncertainties.

Figure 21 which shows the  $\chi^2$  values for the T\_REF models with different  $\gamma$ —we only vary  $\gamma$  and not  $T_0$  because the  $F \gtrsim 0.6$  portions of the transmission PDF that change the most with  $\langle F \rangle_{\text{Ly}\alpha}$  do not vary as much with respect to changes in  $T_0$  (compare with Figure 17). Examples of the corresponding best-fit model PDFs in one S/N subsample are shown in Figure 22, where we see that varying  $\langle F \rangle_{\text{Ly}\alpha}$  can indeed change the shape of the  $F \gtrsim 0.6$  portion of the transmission PDF sufficiently, improving the fits in those transmission ranges compared to the fiducial models (Figure 14).

In all our redshift bins, the best-fitting models seen in Figure 21 are  $\gamma = 1.6$  with  $\chi^2 = [69, 67, 54]$  for 70 dof<sup>24</sup> at  $\langle z \rangle = [2.3, 2.6, 3.0]$  (for the combined data using all S/N bins), respectively, implying probabilities of  $P = [52\%, 59\%, \text{and } 92\%]$  of obtaining larger values.<sup>25</sup> At the higher redshifts best-fitting mean transmission for the  $\gamma = 1.6$  case is

<sup>24</sup> In this particular section, when we quote the  $\chi^2$  for the best-fitting  $\langle F \rangle_{\text{Ly}\alpha}$  the dof is further reduced by 1 compared to the other  $\chi^2$  summed over the S/N subsamples.


<sup>25</sup> These  $\chi^2$  values are very small for the dof, suggesting that we may have overestimated the size of our systematic errors, but as we shall see this does not affect our ability to place constraints on  $\gamma$  and merely makes our conclusions rather conservative.



**Figure 21.**  $\chi^2$  values for the T\_REF models (with different  $\gamma$ ) plotted as a function of Ly $\alpha$  forest mean transmission values,  $\langle F \rangle_{\text{Ly}\alpha}$ , used to normalize the simulation skewers. The quoted  $\chi^2$  values (with  $\nu = 71$  dof) were obtained by summing over the  $\chi^2$  for the different S/N subsamples at each redshift. The fiducial transmission values inferred from Becker et al. (2013) are shown as the solid vertical lines, while the dot-dashed vertical lines denote their  $1\sigma$  errors. The dashed lines in the  $\langle z \rangle = 3$  panel denote the inflated error bars we use to account for the quasar selection bias shown in Figure 23. In Section 6 we will marginalize over the uncertainties in  $\langle F \rangle_{\text{Ly}\alpha}$  to obtain our final results.

pushed to significantly discrepant values with respect to the fiducial Becker et al. (2013) values (Figure 21).

The  $\gamma = 1.3$  model also provides acceptable fits to the models, with  $\chi^2 = [71, 73, 58]$  for 70 dof ( $P = [43\%, 40\%, 84\%]$ ) at  $\langle z \rangle = [2.3, 2.6, 3.0]$ , but at the two higher redshift bins this requires  $\langle F \rangle_{\text{Ly}\alpha}$  values that are increasingly discrepant compared to Becker et al. (2013) ( $+2.3\sigma$  and  $+5\sigma$  respectively at  $\langle z \rangle = [2.6, 3.0]$ ). The isothermal  $\gamma = 1.0$  models are disfavored at the




**Figure 22.** Model transmission PDFs (curves) with the best-fit Ly $\alpha$  forest mean transmission  $\langle F \rangle_{\text{Ly}\alpha}$  for different  $\gamma$  values from the T\_REF family of models (using the improved LLS/pLLS model). These are for the S/N = 8–10 subsample and compare with the corresponding BOSS data transmission PDFs (error bars). The upper two panels in each plot show the transmission PDFs in linear and logarithmic ordinate axes, respectively, while the bottom panels show residuals divided by the errors, with dashed horizontal lines indicating the  $\pm 1\sigma$  region relative to the data. The best-fitting  $\langle F \rangle_{\text{Ly}\alpha}$  values correspond to the minima in Figure 21, but here we have labeled them relative to  $\langle F \rangle_{\text{Ly}\alpha, \text{B13}}$ , the fiducial Becker et al. (2013) values and errors. The  $\chi^2$  values quoted are for 23 dof (taking into account the fitting of the LLS  $b$ -parameter), and were computed using the full error covariances including both bootstrap and systematic terms.

two lower redshift bins, with best-fit values of  $\chi^2 = [98, 97]$  for 70 dof ( $P = [2\%, 2\%]$ ) at  $\langle z \rangle = [2.3, 2.6]$ , whereas at  $\langle z \rangle = 3$ , the error bars on the PDF are sufficiently large that acceptable fits are obtainable using  $\gamma = 1.0$ , with  $\chi^2 = 68$  for 70 dof ( $P = 54\%$ ). However, this requires a  $+5\sigma$  discrepancy in  $\langle F \rangle_{\text{Ly}\alpha}$  with respect to Becker et al. (2013). In Figure 22, one sees that fitting for  $\langle F \rangle_{\text{Ly}\alpha}$  allows the  $\gamma = 1.0$  models to be in good agreement with the data in the  $F > 0.7$  portion of the PDF, but gives rise to discrepancies in the  $0.4 \lesssim F \lesssim 0.7$  range, which limits the goodness-of-fit, and cannot easily be compensated by modifying the metals or LLS model.

From Figure 21, it is clear that as we move to higher redshifts, we require increasingly higher  $\langle F \rangle_{\text{Ly}\alpha}$  relative to the fiducial Becker et al. (2013) values in order to agree with the data: at  $\langle z \rangle = 2.3$ , our best-fit mean transmission for the  $\gamma = 1.6$  model agrees with Becker et al. (2013), but at  $\langle z \rangle = 3$  there is a significant deviation of  $+2\sigma$  with respect to the Becker et al. (2013) measurement. The same trend is true for the best-fit  $\gamma = 1.3$  and  $\gamma = 1.0$  models, but these require even greater discrepancies with respect to the fiducial  $\langle F \rangle_{\text{Ly}\alpha}$ .

One possible explanation for this discrepancy is the effect on the Becker et al. (2013) measurement of  $u$ -band selection bias in the SDSS quasars. This was first noted by Worseck & Prochaska (2011), who found that the color–color criteria used to select SDSS quasars preferentially selected quasars, specifically in the redshift range  $3 \lesssim z_{\text{qso}} \lesssim 3.5$ , that have intervening Lyman-breaks at  $\lambda_{\text{rest}} < 912 \text{ \AA}$ . The  $3 \lesssim z_{\text{qso}} \lesssim 3.5$  SDSS quasars are thus more likely to have intervening LLSs in their sightlines, yielding an additional contribution to the Ly $\alpha$  absorption and hence causing Becker et al. (2013) to possibly underestimate  $\langle F \rangle_{\text{Ly}\alpha}$  when stacking the impacted quasars. Becker et al. (2013) mentioned this effect in their paper but argued that it was much smaller than their estimated errors by referencing theoretical IGM transmission curves estimated by Worseck & Prochaska (2011; Figure 17 in the latter paper).

Dr. G. Worseck has kindly provided us with these transmission curves,  $T_{\text{IGM}}(\lambda)$ , which were generated for both the average IGM absorption and that extracted from SDSS quasars



**Figure 23.** Red and black curves show the excess Ly $\alpha$  absorption expected from sightlines of  $z_{\text{qso}} = 3.2$  and  $z_{\text{qso}} = 3.4$  quasars, respectively, relative to the mean IGM transmission. This is caused by the SDSS selection bias described in Worseck & Prochaska (2011), which yields above-average numbers of intervening LLSs. These are derived from the same curves shown in Figure 17 of Worseck & Prochaska (2011), but replotted as ratios smoothed by a boxcar function over 12 pixels for clarity. The top axis labels the Ly $\alpha$  absorption redshift corresponding to each wavelength, while the shaded region indicates the wavelength range of our  $\langle z \rangle = 3.0$  bin. The dashed line shows, for comparison, the relative errors on the Ly $\alpha$  forest mean transmission estimated by Becker et al. (2013). The discrepancy due to the SDSS bias is significant compared to the Becker et al. (2013) errors.

affected by the color–color selection bias. In Figure 23 we plot the relative difference between the biased Ly $\alpha$  transmission deduced from  $z_{\text{qso}} = 3.2$  and  $z_{\text{qso}} = 3.4$  quasars and the true mean IGM transmission, using the Worseck & Prochaska (2011) transmission curves. It is clear that at Ly $\alpha$  absorption redshifts of  $z_{\text{abs}} \approx 3$ , the excess LLSs picked up from such quasars contribute an additional  $\sim 1\%$  compared to the mean IGM decrement, a discrepancy that is of the same magnitude as the error bars in the Becker et al. (2013) measurement, indicated by the dashed line.

This could partially explain the higher  $\langle F \rangle_{\text{Ly}\alpha}$  required to make our  $\langle z \rangle = 3$  models fit the data in Figure 21. Note that we expect this UV color selection bias to be much less significant in our BOSS data, since we have selected bright quasars in the top 5th percentile of the S/N distribution. Given that such quasars have high-S/N photometry, their colors separate much more cleanly from stellar contaminants. Furthermore, such bright quasars are much more likely to have been selected with multi-wavelength data (e.g., including near-IR and radio in addition to optical photometry see Ross et al. 2012). For both of these reasons, we expect our quasars to be much less susceptible to biases in color-selection related to the presence of an LLS. A careful accounting of this bias is beyond the scope of this paper, but from now on we will inflate by a factor of two the corresponding errors on  $\langle F \rangle_{\text{Ly}\alpha}$  at  $\langle z \rangle = 3$  to account for this possible bias in the mean transmission measurements (dashed vertical lines in bottom panel of Figure 21).

Another possibility that could explain a bias in the  $\langle F \rangle_{\text{Ly}\alpha}$  measured by Becker et al. (2013) is their assumption that the metal contamination of the Ly $\alpha$  forest does not evolve with redshift. While there are few clear constraints on the aggregate metal contamination within the forest, assuming that the metals actually decrease with increasing redshift (e.g., in the case of CIV, Cooksey et al. 2013), the assumption of an unevolving metal contribution calibrated at  $z \approx 2.3$  would lead to an underestimate of  $\langle F \rangle_{\text{Ly}\alpha}$  at higher redshifts, which could explain the trend we seem to be seeing.

It is clear from the previous discussion that there is some degeneracy between  $\gamma$  and  $\langle F \rangle_{\text{Ly}\alpha}$  in our transmission PDFs. However, we are primarily interested in  $\gamma$ , while the  $\langle F \rangle_{\text{Ly}\alpha}$  has been extensively measured over the years, which allows the placement of strong priors. In the next section, we will therefore marginalize over  $\langle F \rangle_{\text{Ly}\alpha}$  in order to obtain our final results.

## 6. RESULTS

Due to the uncertainties in  $\langle F \rangle_{\text{Ly}\alpha}$  described in the previous subsection, for a better comparison between transmission PDFs,  $p$ , from models with different  $[\gamma, T_0]$  we will marginalize the model likelihoods,  $\mathcal{L} = \exp(-\chi^2/2)$ , over the Ly $\alpha$  forest mean transmission,  $\langle F \rangle$ :

$$\mathcal{L}(p | \gamma, T_0) = \int_{-\infty}^{\infty} \mathcal{L}(p | \gamma, T_0, \langle F \rangle) A(\langle F \rangle) dF, \quad (24)$$

where  $A(\langle F \rangle)$  is the prior on  $\langle F \rangle$  (for clarity in these equations,  $\langle F \rangle$  is used as a shorthand for  $\langle F \rangle_{\text{Ly}\alpha}$ ). We assume a Gaussian prior:

$$A(\langle F \rangle) = \frac{1}{\sigma_F \sqrt{2\pi}} \exp\left[-\frac{(\langle F \rangle - \langle F \rangle_{\text{B13}})^2}{2\sigma_F^2}\right], \quad (25)$$

where  $\langle F \rangle_{\text{B13}}$  and  $\sigma_F$  are the optically thin Ly $\alpha$  forest mean transmission and associated errors, respectively, estimated from Becker et al. (2013). Note that for  $\langle z \rangle = 3$ , we have decided to dilate the error bars by a factor of two to account for the suspected quasar selection bias discussed in the previous section.

For each model, we generate transmission PDFs with different  $\langle F \rangle_{\text{Ly}\alpha}$  (similar to Figure 21) and evaluate the combined  $\chi^2$  summed over different S/N. We interpolate the  $\chi^2$  over  $\langle F \rangle_{\text{Ly}\alpha}$  to obtain a finer grid, which then allows us to numerically integrate Equation (24) using five-point Newton–Coates quadrature.

At this stage, we also analyze models with different IGM temperatures at mean density,  $T_0$ . Hitherto, we have been working

**Table 3**  
Marginalized  $\chi^2$  for  $\nu = 71$  dof


$\langle z \rangle = 2.3$			
$\gamma$	T_COLD ( $T_0 = 13000$ K)	T_REF ( $T_0 = 18000$ K)	T_HOT ( $T_0 = 23000$ K)
1.6	87.7	72.9	79.5
1.3	103.4	76.0	71.8
1.0	174.2	105.5	88.4
$\langle z \rangle = 2.6$			
$\gamma$	T_COLD ( $T_0 = 11000$ K)	T_REF ( $T_0 = 16000$ K)	T_HOT ( $T_0 = 21500$ K)
1.6	88.8	72.0	71.4
1.3	118.0	82.6	91.8
1.0	203.3	127.3	111.1
$\langle z \rangle = 3.0$			
$\gamma$	T_COLD ( $T_0 = 9000$ K)	T_REF ( $T_0 = 14000$ K)	T_HOT ( $T_0 = 19000$ K)
1.6	61.7	65.1	62.5
1.3	77.6	72.7	63.8
1.0	119.5	77.7	85.8

only with the central T\_REF model ( $T_0(z = 2.5) \sim 16,000$  K), but we now also compare models from the T\_HOT and T\_COLD simulations, which have  $T_0(z = 2.5) \sim 11,000$  K and  $T_0(z = 2.5) \sim 21,500$  K, respectively. Each of these temperature models also sample temperature–density relationships of  $\gamma = [1.0, 1.3, 1.6]$  for a model grid of  $3 \times 3$  parameters at each redshift.

The marginalized  $\chi^2$  values for all the models are tabulated in Table 3, and plotted as a function of  $\gamma$  in Figure 24. In general, the T\_REF models with  $\gamma = 1.6$  provide the best agreements with the data at all redshifts with  $\chi^2 \approx 60$ –70 for 69 dof. The T\_HOT models (with higher IGM temperatures at mean density) provide fits of comparable quality, and indeed at  $\langle z \rangle = 2.3$  the T\_HOT model with  $\gamma = 1.3$  gives essentially the same goodness-of-fit as the  $\gamma = 1.6$  T\_REF model. The cooler T\_COLD models are less favored by the data, and at  $\langle z \rangle = 2.6$  give unreasonable fits to the data with  $\chi^2 = 89$  for 69 dof ( $P = 5\%$ ), but at other redshifts they are acceptable fits to the data. In other words, the transmission PDF does not show a strong sensitivity for  $T_0$ , which we shall show later is due to degeneracy with our LLS model in the low-transmission end of the transmission PDF.

The more important question to address is the possibility of isothermal or inverted temperature–density relationships ( $\gamma \leq 1$ ) as suggested by some studies on the transmission PDF of high-resolution, high-S/N echelle quasar spectra (e.g., Bolton et al. 2008; Viel et al. 2009; Calura et al. 2012). It is clear from Table 3 and Figure 24 that for all  $T_0$  models the isothermal,  $\gamma = 1.0$  models disagree strongly with the BOSS data. The closest match for an isothermal IGM is the T\_REF model at  $\langle z \rangle = 3.0$ , which yields  $\chi^2 = 78$  for 69 dof, or a probability of 21% of obtaining the data from this model. However, relative to the  $\gamma = 1.6$  model at  $\langle z \rangle = 3.0$ , which gives the minimum  $\chi^2$  at that redshift, we find  $\Delta\chi^2 \approx 16$  for the isothermal model, i.e., a  $\sqrt{\Delta\chi^2} = 4\sigma$  discrepancy from the best-fit model. The isothermal model is also strongly disfavored at the other redshifts, where we find  $\Delta\chi^2 \approx [15, 40]$  at  $\langle z \rangle = [2.3, 2.6]$  or  $\sqrt{(\Delta\chi^2)} \approx [3.9\sigma, 6.3\sigma]$ . Since the shape of the transmission PDF varies continuously as a function of  $\gamma$  (see, e.g., Bolton et al. 2008; Lee 2012), these results imply





**Figure 24.**  $\chi^2$  values (for 71 dof) from models with different  $\gamma$  and  $T_0$  at different redshifts, after marginalizing over uncertainties in the mean transmission ( $F$ ) of the Ly $\alpha$  forest. Models with  $\gamma = 1.6$  are generally favored, although  $\gamma = 1.3$  with the T\_HOT model is also acceptable at  $\langle z \rangle = 2.3$ . The same quantities are also tabulated in Table 3.

that inverted ( $\gamma < 1$ ) IGM temperature–density slopes are even more strongly ruled out.

## 7. DISCUSSION

In this paper, we have studied the  $\langle z \rangle = 2.3$ –3 Ly $\alpha$  forest transmission PDF from 3373 BOSS DR9 quasar spectra.

Although this is a relatively small subsample selected to be in the top 95th percentile in terms of S/N, it provides a two-orders-of-magnitude-larger Ly $\alpha$  forest path length than high-resolution, high-S/N data sets previously used for this purpose, providing unprecedented statistical power for transmission PDF analysis.

In order to ensure accurate characterization and allow subsequent modeling of the spectral noise, we have introduced a novel, probabilistic method of combining the multiple exposures that comprise each BOSS observation, using the raw sky and calibration data. This method significantly improves the accuracy of the noise estimation, and additionally allows us to generate mock spectra with noise properties tailored to each individual BOSS spectrum, but self-consistently for different Ly $\alpha$  forest realizations. We believe that our noise modeling—which yields noise estimates accurate to  $\sim 3\%$  across the relevant wavelength range—is the most careful treatment of spectral noise in multi-object fiber spectra to date, and we invite readers with similarly stringent requirements in understanding the BOSS spectral noise to contact the authors. In the future, the spectral extraction algorithm described by Bolton & Schlegel (2010) may solve some of the issues that affected us, but this has yet to be implemented.

For the continuum estimation, we used the MF-PCA method introduced in Lee (2012). This method, which reduces the uncertainty in the continuum estimation to  $\sigma_{\text{cont}} \lesssim 5\%$ , fits for a continuum such that the resulting Ly $\alpha$  forest has a mean transmission  $\langle F \rangle$  matched to external constraints, for which we use the precise measurements by Becker et al. (2013). While MF-PCA does require external constraints for  $\langle F \rangle$ , we argue that so long as both the real quasars and mock spectra are continuum-fitted in exactly the same way, the shape of the transmission PDF retains independent information on the Ly $\alpha$  forest mean transmission.

To compare with the data, we used the detailed hydrodynamical simulations of Viel et al. (2013a), which explore a range of IGM temperature–density slopes ( $\gamma \approx 1.0$ –1.6) and temperatures at mean density ( $T_0(z = 2.5) \approx [11000, 16000, 21500]$  K). We processed the simulated spectra to take into account the characteristics of the individual BOSS spectra in our sample, such as spectral resolution, pixel noise, and continuum fitting errors. We also incorporate the effects of astrophysical “nuisance” parameters such as LLSs and metal contamination. The LLSs are modeled by adding  $10^{16.5} \text{ cm}^{-2} \lesssim N_{\text{HI}} \lesssim 10^{20.3} \text{ cm}^{-2}$  absorbers into our mock spectra, based on published measurements of the observed incidence  $l_{\text{LLS}}(z)$  (Ribaldo et al. 2011) and HI column density distribution  $f(N_{\text{HI}})$  (Prochaska et al. 2010). Meanwhile, contamination from lower-redshift metals are modeled in an empirical fashion by inserting  $\lambda_{\text{rest}} > 1216 \text{ \AA}$  absorbers observed in lower-redshift SDSS/BOSS quasars into the same observed wavelengths of our mock spectra.

Our initial models did not provide satisfactory agreement with the transmission PDF measured from the BOSS spectra, with discrepancies at both the high-transmission and low-transmission bins. However, the differences between data and models were consistent across the different S/N subsamples, indicating that our noise modeling is robust. To resolve the discrepancies at the low-transmission end of the PDF, we explored various modifications to our LLS model. First, we steepened the column-density distribution slope of partial LLSs ( $16.5 < \log_{10}(N_{\text{HI}}) < 17.5$  systems) to  $\beta_{\text{LLS}} = -2$ , a value suggested from the mean free path of ionizing photons (Prochaska et al. 2010). This change relieved the tension between model and data in the  $F \approx 0.1$ –0.4 bins, but implies increasing the

number of pLLS by nearly an order of magnitude, but this is not unreasonable given the current uncertainties for this population (Janknecht et al. 2006; Prochaska et al. 2010). We believe that the necessity of a pLLS distribution with  $\beta_{\text{LLS}} \approx -2$  to fit the BOSS Ly $\alpha$  transmission PDF supports the claims of Prochaska et al. (2010) regarding the column-density distribution of this population.


However, after adding pLLSs a major discrepancy remained in the saturated  $F \approx 0$  bins, which we addressed by adjusting the effective  $b$ -parameter assumed in all the optically thick systems in our model. We found that an effective value of  $b = 45 \text{ km s}^{-1}$  provided the best fit to our model.<sup>26</sup>

At the high-transmission ( $F \gtrsim 0.6$ ) end of the model transmission PDFs, we found that modifying the Ly $\alpha$  forest mean transmission in the simulations,  $\langle F \rangle_{\text{Ly}\alpha}$ , allowed much better agreement with the BOSS data. At  $\langle z \rangle = [2.3, 2.6]$ , the  $\langle F \rangle_{\text{Ly}\alpha}$  that gave the best-fitting model PDFs were within  $1\sigma$  of the Becker et al. (2013) measurements, but at  $\langle z \rangle = 3$  we required a value that was  $\sim 2\sigma$  larger. We argue that this discrepancy could be due to a color-color selection bias in the  $3 \lesssim z_{\text{qso}} \lesssim 3.5$  SDSS quasars used by Becker et al. (2013), which preferentially selected sightlines with intervening LLSs, giving rise to additional Ly $\alpha$  absorption (and thus lower  $\langle F \rangle_{\text{Ly}\alpha}$ ) at a level comparable to the errors estimated by Becker et al. (2013). Our BOSS spectra, on the other hand, should be comparatively unaffected on account of being the brightest quasars in the survey, hence they separate more cleanly from the stellar locus in color-space, and were more likely to have been selected with additional criteria (radio, near-IR, variability, etc.) beyond color-color information (Ross et al. 2012).

To deal with these uncertainties, we decided to marginalize over the mean transmission in our  $\chi^2$  analysis. At  $\langle z \rangle = 2.3$ , the preferred model is for a hot IGM with ( $T_0 = 23000 \text{ K}$ ) along with  $\gamma = 1.3$  ( $P \approx 45\%$ ), although the intermediate-temperature model ( $T_0 = 18000 \text{ K}$ ) with  $\gamma = 1.6$  is nearly as good a fit with  $P \approx 82\%$ . The preferred models at  $\langle z \rangle = [2.6, 3.0]$  are for  $\gamma = 1.6$  at temperatures at mean density of  $T_0 = [21500, 9000] \text{ K}$  ( $P = [46\%, 78\%]$ , respectively). We find that the isothermal ( $\gamma = 1$ ) temperature-density relationship is strongly disfavored at all redshifts regardless of  $T_0$ , with discrepancies of  $\sqrt{\Delta\chi^2} \sim 4\text{--}6\sigma$  compared to the best-fit models.

One might be skeptical of the results given the various assumptions we had to make in modeling astrophysical nuisance parameters. To test the robustness of our results to systematics, we generated 20 iterations of model transmission PDFs sampling all nine of our  $[T_0, \gamma]$  models (i.e., 180 PDFs in total) in the  $\langle z \rangle = 2.6$ , S/N = 8–10 bin, where each iteration has a random realization of the systematics (LLSs, metals, continuum errors, etc.) drawn in the same way as our Monte Carlo estimate of systematic uncertainty (Section 5.3). We then asked how many times each  $T_0$  or  $\gamma$  model gave the lowest  $\chi^2$  when compared with the data. For this test we only evaluated the  $\chi^2$  at the fiducial  $\langle F \rangle_{\text{Ly}\alpha}$  without marginalization.

The results of this test are shown in Figure 25. In the top panel, the T\_REF and T\_HOT models are favored  $\sim 40\%$  of the time but the T\_COLD has  $\sim 15\%$  of being favored depending on the (random) choice of systematics. In other words, there is significant degeneracy between our systematics model and  $T_0$ . We suspect this is driven largely by the choice of the LLS



**Figure 25.** Histogram indicating the fraction of times a given  $T_0$  (top) or  $\gamma$  (bottom) model is favored for the  $\langle z \rangle = 2.3$ , S/N = 8–10 transmission PDF when the systematics levels in the model are randomly sampled 20 times. While different systematics could lead to different best-fitting models for  $T_0$ , the models with  $\gamma = 1.6$  are always preferred. This indicates some degeneracy in our systematics model with  $T_0$ , but our conclusions on  $\gamma$  are robust.

$b$ -parameter, which changes the shape of the transmission PDF in a similar way to  $T_0$  (compare Figure 16 with Figure 17). In contrast, the bottom panel of Figure 25 shows that whatever systematics we choose,  $\gamma = 1.6$  is always favored, indicating a robust constraint.

There is, however, some degeneracy between  $\gamma$  and the Ly $\alpha$  forest mean transmission,  $\langle F \rangle_{\text{Ly}\alpha}$ . While we marginalize over the latter quantity, the choice of prior can, in principle, affect the results. However, at  $\langle z \rangle = [2.3, 2.6]$ , the chi-squared minimum of the  $\gamma = 1.0$  PDF model as a function of  $\langle F \rangle_{\text{Ly}\alpha}$  is  $\chi^2 \approx 100$  for 71 dof (Figure 21), which has a probability of  $P \approx 1\%$ . In other words, even if we fine-tuned  $\langle F \rangle_{\text{Ly}\alpha}$  in an attempt to force the isothermal model as the best-fit model at these redshifts, it would still be an unacceptable fit, and the  $\gamma = 1.3$  model would still be preferred over it. This is less clear cut at  $\langle z \rangle = 3$ , where the error bars are large enough to permit a reasonable minimum chi-squared of  $\chi^2 \approx 70$  for 71 dof using the  $\gamma = 1$  model, but this requires a value of  $\langle F \rangle_{\text{Ly}\alpha} = 0.71$ , which is  $5\sigma$  discrepant from the value reported by Becker et al. (2013). While this  $\langle F \rangle_{\text{Ly}\alpha}$  measurement is dependent on corrections for metals and LLS absorption (and indeed we argue that they have neglected a subtle bias related to SDSS quasar selection), they have attempted to incorporate these uncertainties into their errors and we have no particular reason to believe that they

<sup>26</sup> Note that we have quoted an *effective*  $b$ -parameter, which must not be confused with the  $b$  from individual kinematical components, which is often quoted by workers carrying out Voigt profile analysis of high-resolution spectra.

have underestimated this by a factor of  $>5$ . A quick survey of the available measurements on the forest mean transmission from the past decade yields  $\langle F \rangle_{\text{Ly}\alpha}(z=3) \approx 0.65\text{--}0.69$  (Kim et al. 2007; Faucher-Giguère et al. 2008; Dall’Aglio et al. 2008), albeit with larger errors. The use of any of these measurements as priors for our analysis would therefore disfavor an IGM with  $\gamma \leq 1$ , (which requires  $\langle F \rangle_{\text{Ly}\alpha}(z=3) \geq 0.71$ ), unless all the available literature in the field has significantly underestimated the mean transmission.

There are several cosmological and astrophysical effects that we did not model, which could in principle affect our conclusions on  $\gamma$ . Since the Ly $\alpha$  forest transmission PDF essentially measures the contrast between high-absorption and low-absorption regions of the IGM, this can be degenerate with the underlying amplitude of matter fluctuations, which is specified by a combination of  $\sigma_8$  and  $n_s$ , the matter fluctuation variance on  $8 h^{-1}$  Mpc scales and the slope of the amplitude power spectrum, respectively. While these parameters are increasingly well-constrained (e.g., Planck Collaboration et al. 2014), there is still some uncertainty regarding the level of the fluctuations on the sub-Mpc scales relevant to the Ly $\alpha$  forest, which could be degenerate with our  $\gamma$  measurement. Bolton et al. (2008) explored this degeneracy between  $\sigma_8$  and  $\gamma$  in the context of transmission PDF measurements from high-resolution spectra, and found that the PDF is less sensitive to plausible changes in  $\sigma_8$  compared to  $\gamma$ , e.g., modifying  $\sigma_8$  by  $\Delta\sigma_8 \pm 0.1$ , affected the shape of the PDF less than a modification of  $\Delta\gamma \pm 0.25$  (Figure 2 in their paper). This degeneracy is in fact further weakened when an MCMC analysis of the full parameter space is considered, as shown by the likelihood contours in Viel et al. (2009).

The astrophysical effects that could be degenerate with  $\gamma$  include galactic winds and inhomogeneities in the background UV ionizing field. The injection of gas into the IGM by strong galaxy outflows could in principle modify Ly $\alpha$  forest statistics at fixed  $\gamma$ ; this was studied using hydrodynamical simulations by Viel et al. (2013b), who concluded that the effect on the PDF is small compared to the uncertainties in high-resolution PDF measurements. Our BOSS measurement has roughly the same errors as those from high-resolution spectra once systematic uncertainties are taken into account, and it therefore seems unlikely that galactic winds could significantly bias our conclusions on  $\gamma$ . Meanwhile, fluctuations in the UV ionizing background,  $\Gamma$ , that are correlated with the overall density field could also be degenerate with the temperature–density relationship (compare with Equation (1)). This effect was studied by McDonald et al. (2005a) in simulations using an extreme model that considered only UV background contributions from highly biased active galactic nuclei, which maximizes the inhomogeneities. They concluded that while these UV fluctuations affected forest transmission statistics at  $z \sim 4$ , the effect was small at  $z \lesssim 3$ , the redshift range of our measurements.

Various observational and systematic effects could also, in principle, affect our constraints on  $\gamma$ . For example, our modeling of the BOSS spectral resolution assumes a Gaussian smoothing kernel which might affect our constraints if this were untrue. However, in their analysis of the 1D forest transmission power spectrum, Palanque-Delabrouille et al. (2013) examined the BOSS smoothing kernel and did not find significant deviations from Gaussianity. There are also possible systematics caused by our simplified modeling of LLS and metal contamination in the data, for example in our assumption of a single  $b$ -parameter for all LLSs and our neglect of very weak metal absorbers. However, we believe that the test performed in

Figure 25 samples larger differences in the transmission PDF than those caused by our model simplifications, e.g., it seems unlikely that going from a single LLS  $b$ -parameter to a finite  $b$ -distribution could cause greater differences in the flux PDF than varying the single  $b$ -parameter by  $\pm 50\%$  as was done in Figure 25. As for continuum-estimation, we carry out the exact same continuum-fitting procedure on the mock spectra as on the real quasar spectra, which leads no overall bias since in both cases the resulting forest transmission field is forced to have the same overall transmission,  $\langle F \rangle_{\text{cont}}$ . The only uncertainty then relates to the distribution of  $c'/c - 1$ , i.e., the per-pixel error of the estimated continuum,  $c'$ , relative to the true continuum,  $c$ . In reality the shape of this distribution could be different between the data and the mock versions, whereas within our mock framework we could only explore overall rescalings of the distribution width. Again, we find it unlikely that differences in the transmission PDF caused by the true shape of the  $c'/c - 1$  distribution could be so large as to be comparable to the effect caused by varying the width of the continuum error distribution, which we have examined.

While we do not think that the effects described in the previous few paragraphs qualitatively affect our conclusion that the BOSS data is inconsistent with isothermal or inverted IGM temperature–density relationships ( $\gamma \leq 1$ ), when taken in aggregate these systematic uncertainties do weaken our formal  $4\sigma\text{--}6\sigma$  limits against  $\gamma \leq 1$  and need to be explicitly considered in future analyses.

### 7.1. Astrophysical Implications

How does this compare with other results on the thermal state of the IGM? McDonald et al. (2001) analyzed the transmission PDF from eight high-resolution, high S/N spectra and compared with them now-obsolete hydrodynamical simulations. They found the data to be consistent with a temperature–density relationship with the expected values of  $\gamma \approx 1.5$  (Hui & Gnedin 1997). More recently, Bolton et al. (2008) and Viel et al. (2009) carried out analyses of the transmission PDF measured from a larger sample (18 spectra) of Ly $\alpha$  forest sightlines measured by Kim et al. (2007) and found evidence for an inverted temperature–density relationship ( $\gamma < 1$ ). Viel et al. (2009) found that at  $z \approx 3.0$ , the temperature–density relation was highly inverted ( $\gamma \approx 0.5$ ), and remained so as low as  $z \approx 2.0$  although at the lower redshifts the data was marginally consistent with an isothermal IGM. They suggested that the difference between their results and those of McDonald et al. (2001) was due to the now-obsolete cosmological parameters and less-detailed treatment of intervening metals in the earlier study. However, Lee (2012) then pointed out that there is a sensitivity of the measured values of  $\gamma$  from the transmission PDF on continuum-fitting. Since continuum-fitting of high-resolution data generally involves manually placing the continuum at Ly $\alpha$  forest transmission peaks that do not necessarily reach the true continuum, it is conceivable that continuum biases combined with underestimated jackknife errors bars (e.g., Rollinde et al. 2013) could have led Bolton et al. (2008) and Viel et al. (2009) to erroneously deduce an inverted temperature–density relationship (see Bolton et al. 2014 for a detailed discussion on this point). In our analysis we have fitted our continua using an automated process that is free from the same continuum-fitting bias, although it does require an assumption on the underlying Ly $\alpha$  forest transmission, which we have marginalized over in our analysis.



Most recent measurements of the transmission PDF from high-resolution data have continued to favor an isothermal or inverted  $\gamma$ —Calura et al. (2012) analyzed the transmission PDF from a sample of  $z \approx 3.3$ – $3.8$  quasars and also found an isothermal temperature–density relationship at  $z = 3$ , although combining with the Kim et al. (2007) data drove the estimated  $\gamma$  to inverted values at  $z < 3$ . However, Rollinde et al. (2013) carried out a re-analysis of the transmission PDF from various high-resolution echelle data sets, which included significant overlap with the Kim et al. (2007) data. They argue that previous analyses have underestimated the error on the transmission PDF, and found the observed transmission PDF to be consistent with simulations that have  $\gamma \approx 1.4$  over  $2 < z < 3$ —this discrepancy is probably also driven by a different continuum-estimation from the Kim et al. (2007) measurement.

The use of other statistics on high-resolution spectra have, however, tended to disfavor an isothermal or inverted temperature–density relationship. Rudie et al. (2012) analyzed the lower-end of the  $b$ – $N_{\text{H I}}$  cutoff from individual Ly $\alpha$  forest absorbers measured in a set of 15 very high-S/N quasar echelle spectra, and estimated  $\gamma \approx 1.5$  at  $z = 2.4$ . Bolton et al. (2014) compared the Rudie et al. (2012) measurements to hydrodynamical simulations and corroborated their determination of the temperature–density relationship slope.

Garzilli et al. (2012) analyzed the Kim et al. (2007) sample and found that while the transmission PDF supports an isothermal or inverted temperature–density relationship, a wavelet analysis favors  $\gamma > 1$ . Note, however, that the  $b$ – $N_{\text{H I}}$  cutoff and the transmission PDF are sensitive to different density ranges, with the PDF probing gas densities predominantly below the mean (e.g., Bolton et al. 2014).

Our result of  $\gamma \approx 1.6$  at  $\langle z \rangle = [2.3, 2.6, 3.0]$  is thus in rough agreement with measurements that do *not* involve the transmission PDF from high-resolution Ly $\alpha$  forest spectra (with the exception of Rollinde et al. 2013). Our value of  $\gamma$  at  $\langle z \rangle = 3$  is somewhat unexpected because one expects a flattening of the temperature–density relationship close to the He II reionization epoch at  $z \sim 3$  (Furlanetto & Oh 2008; McQuinn et al. 2009; but see Gleser et al. 2005; Meiksin & Tittley 2012), although  $\gamma = 1.3$  is not strongly disfavored ( $\sqrt{\Delta\chi^2} \sim 2.6$ ).

Taken at face value, the temperature–density relationship during He II reionization can be made steeper by a density-independent reionization and/or a lower heating rate in the IGM (Furlanetto & Oh 2008), which could be reconciled with an extended He II event (Shull et al. 2010; Worseck et al. 2011).

Our constraints on  $\gamma$  appear to be in conflict with the prediction of the theories of Broderick et al. (2012) and Chang et al. (2012), who elucidated a relativistic pair-beam channel for plasma-instability heating of the IGM from TeV gamma-rays produced by a population of luminous blazars. This mechanism provides a uniform volumetric heating rate, which would cause an inverted temperature–density relationship in the IGM (Puchwein et al. 2012) since voids would experience a higher specific heating rate compared with heating by He II reionization alone. This picture has been challenged by the recent study of Sironi & Giannios (2014), who dispute the amount of heating this mechanism could provide, since they found that the momentum dispersion of such relativistic pair beams allows  $\ll 10\%$  of the beam energies to be deposited into the IGM.

However, in this paper we have assumed relatively simple temperature–density relationships in which the bulk of the IGM in the density range  $0.1 \lesssim \Delta \lesssim 5$  follows a relatively tight power

law. We have therefore not studied more complicated  $T$ – $\Delta$  relationships, e.g., with a spread of temperatures at fixed density (e.g., Meiksin & Tittley 2012; Compostella et al. 2013) that might be caused by He II reionization or other phenomena. It is therefore possible that such complicated temperature–density relationships could result in Ly $\alpha$  forest transmission PDFs that mimic the  $\gamma \approx 1.6$  power law; this is something that needs to be examined in more detail in future work.

## 7.2. Future Prospects

Looking forward, the subsequent BOSS data releases will significantly enlarge our sample size, e.g., DR10 (Ahn et al. 2014) is nearly double the size of the DR9 sample used in this paper, while the final BOSS sample (DR12) should be three times as large as DR9. In particular, the newer data sets should be sufficiently large for us to analyze the transmission PDF and constrain  $\gamma$  during the epoch of He II reionization at  $z > 3$ . This would be a valuable measurement, since high-resolution spectra are particularly affected by continuum-fitting biases at these redshifts (Faucher-Giguère et al. 2008; Lee 2012).

The analysis of the optically thin Ly $\alpha$  forest transmission PDF from these expanded data sets will have vanishingly small sample errors, and the errors will be dominated by systematic and astrophysical uncertainties. At the high-transmission end, our uncertainties are dominated by the scatter of the continuum-fitting, which is dominated by the question of whether our quasar PCA templates, derived from low-luminosity low-redshift quasars (Suzuki et al. 2005), or high-luminosity SDSS quasars (Pâris et al. 2011), respectively, are an accurate representation of the BOSS quasars. This uncertainty should be eliminated in the near future by PCA templates derived self-consistently from the BOSS data (N. Suzuki et al. 2014, in preparation). The modeling of metal contamination could also be improved in the near future by advances in our understanding of how metals are distributed in the IGM (e.g., Zhu et al. 2014), although metals are a comparatively minor contribution to the uncertainty in our transmission PDF.

We also aim to improve on the rather ad hoc data analysis in this paper, in which we accounted for some uncertainties in our modeling by incorporating them into our error covariances (e.g., LLSs, metals, continuum errors), while  $\langle F \rangle_{\text{Ly}\alpha}$  was marginalized over a fixed grid. In future analyses, it would make sense to carry out a full MCMC treatment of all these parameters which would rigorously account for all the uncertainties and allow straightforward marginalization over nuisance parameters.

Since this paper was initially focused on modeling the BOSS spectra, for the model comparison we used only simulations sampling a very coarse  $3 \times 3$  grid in  $T_0$  and  $\gamma$  parameter space, and were unable to take account for uncertainties in other cosmological ( $\sigma_8$ ,  $n_s$ , etc) and astrophysical (e.g., Jeans scale, Rorai et al. 2013; or galactic winds, Viel et al. 2013b) parameters in our analysis. However, methods already exist to interpolate Ly $\alpha$  forest statistics from hydrodynamical simulations given a set of IGM and cosmological parameters (e.g., Viel & Haehnelt 2006; Borde et al. 2014; Rorai et al. 2013). In the near future we expect to perform joint analyses using other Ly $\alpha$  forest statistics in conjunction with the transmission PDF, such as new measurements of the small-scale ( $k \gtrsim 0.2 \text{ s km}^{-1}$ ) 1D transmission power spectrum (M. Walther et al. 2014, in preparation), moderate-scale ( $0.002 \text{ s km}^{-1} \lesssim k \lesssim 0.2 \text{ s km}^{-1}$ ) transmission power spectrum in both 1D (e.g., Palanque-Desabrouille et al. 2013) and 3D (from ultra-dense Ly $\alpha$  forest surveys using



high-redshift star-forming galaxies, Lee et al. 2014a, 2014b), the phase angle PDF determined from close quasar pair sight-lines (Rorai et al. 2013), and others. Such efforts would require a fine grid sampling the full set of cosmological and IGM thermal parameters in order to ensure that the interpolation errors are small compared to the uncertainties in the data (see e.g., Rorai et al. 2013). Efforts are underway to utilize massively parallel adaptive-mesh refinement codes (Almgren et al. 2013) to generate such parameter grids to study the IGM (Lukić et al. 2014). However, one of the findings of this paper is the importance of correct modeling of LLSs, in particular partial LLSs ( $10^{16.5} \text{ cm}^{-2} \lesssim N_{\text{H I}} \lesssim 10^{17.5} \text{ cm}^{-2}$ ), in accounting for the shape of the observed Ly $\alpha$  transmission PDF. Since our hydrodynamical simulations did not include radiative transfer and cannot accurately capture optically thick systems, we had to add these in an ad hoc manner based on observational constraints that are currently rather imprecise. In the near future, we would want to use hydrodynamical simulations with radiative transfer (even if only in post-processing, e.g., Altay et al. 2011; McQuinn et al. 2011; Altay et al. 2013; Rahmati et al. 2013) to self-consistently model the optically thick absorbers in the IGM. With the unprecedented statistical power of the full BOSS Ly $\alpha$  forest sample, this could provide the opportunity to place unique constraints on the column-density distribution function of partial LLSs.

## 8. SUMMARY/CONCLUSIONS

In this paper, we analyzed the PDF of the Ly $\alpha$  forest transmitted flux using 3393 BOSS quasar spectra (with  $\langle S/N \rangle \geq 6$ ) from Data Release 9 of the SDSS-III survey.

To rectify the inaccurate noise estimates in the standard pipeline, we first carried a custom co-addition of the individual exposures of each spectrum, using a probabilistic procedure that also separates out the signal and CCD contributions, allowing us to later create mock spectra with realistic noise properties. We then estimated the intrinsic quasar continuum using a mean flux regulated technique that reduces the scatter in the estimated continua by forcing the resultant Ly $\alpha$  forest mean transmission to match the precise estimates of Becker et al. (2013), although we had to make minor corrections on the latter to account for our different assumptions on optically thick systems in the data. This now allows us to measure the transmission PDF in the data, which we do so at  $\langle z \rangle = [2.3, 2.6, 3.0]$  (with bin widths of  $\Delta z = 0.3$ ), and split it into S/N subsamples of  $S/N = [6-8, 8-10, 10-25]$  at each redshift bin.

The second part of the paper describes finding a transmission PDF model that describes the data, based on detailed hydrodynamical simulations of the optically thin Ly $\alpha$  forest that sample different IGM temperature–density relationship slopes,  $\gamma$ , and temperatures at mean density,  $T_0$  (where  $T(\Delta) = T_0 \Delta^{\gamma-1}$ ). Using these simulations we generate mock spectra based on the real spectra. These take into account the following instrumental and astrophysical effects:

1. *Lyman-limit systems.* These are randomly added into our mock spectra based on published incidence rates (Ribaud et al. 2011) and column-density distributions (Prochaska et al. 2010), including a large population of partial LLSs ( $10^{16.5} \text{ cm}^{-2} \leq N_{\text{H I}} \leq 10^{17.5} \text{ cm}^{-2}$ ) with a power-law distribution of roughly  $f(N_{\text{H I}}) \propto N_{\text{H I}}^{-2}$ . We assumed an effective  $b = 45 \text{ km s}^{-1}$  for the velocity width of these absorbers.
2. *Metal contamination.* We measure metal absorption from the  $1260 \text{ \AA} \lesssim \lambda \lesssim 1390 \text{ \AA}$  rest-frame region of lower-

redshift quasars at the same observed wavelength, and then add these directly into our mock spectra.

3. *Spectral resolution and noise.* Each mock spectrum is smoothed by the dispersion vector of the corresponding real spectrum (determined by the BOSS pipeline), and we apply corrections that bring the spectral resolution modeling to within  $\sim 1\%$  accuracy. We then introduce pixel noise based on the noise parameters estimated by our probabilistic co-addition procedure on the real data, which also achieves percent-level accuracy on modeling the noise.
4. *Continuum errors.* Since we generate a full mock Ly $\alpha$  forest spectrum including the simulated quasar continuum (based on the continua fitted to the actual data), we can apply our continuum-estimation procedure on each mock to fit a new continuum. The difference between the new continuum and the underlying simulated quasar continuum yields an estimate of the continuum error.

We then compare the model transmission PDFs with the data, using an error covariance that includes both bootstrap errors and systematic uncertainties in the model components described above. At  $\langle z \rangle = 3.0$  we find a discrepancy in the assumed Ly $\alpha$  forest mean transmission,  $\langle F \rangle_{\text{Ly}\alpha}$ , between our data and that derived from Becker et al. (2013), which we argue is likely caused by a selection bias in the SDSS quasars used by the latter. We therefore marginalize out these uncertainties in  $\langle F \rangle_{\text{Ly}\alpha}$  to obtain our final results.

The models with an IGM temperature–density slope of  $\gamma = 1.6$  give the best-fit to the data at all our redshift bins ( $\langle z \rangle = [2.3, 2.6, 3.0]$ ). Models with an isothermal or inverted temperature–density relationship ( $\gamma \leq 1$ ) are disfavored at the  $\sqrt{(\Delta\chi^2)} = [3.9, 6.3, 4.0]\sigma$  at  $\langle z \rangle = [2.3, 2.6, 3.0]$ , respectively. Due to a degeneracy with our LLS model, we are unable to put robust constraints on  $T_0$ , but we have checked that our conclusions on  $\gamma$  are robust to such systematics as can be considered within our model framework. There are other possible systematics we did not consider that could in principle affect our measurement, such as cosmological parameters ( $\sigma_8$ ,  $n_s$ ) and astrophysical effects (galactic winds, inhomogeneous UV ionizing background), but we argue that these are unlikely to qualitatively affect our conclusions.

We thank Michael Strauss, J. Xavier Prochaska, Gabor Worseck, and Joop Schaye for useful comments and discussion. We also thank the members of the ENIGMA group (<http://www.mpia-hd.mpg.de/ENIGMA/>) at the Max Planck Institute for Astronomy (MPIA) for helpful discussions. J.F.H. acknowledges generous support from the Alexander von Humboldt foundation in the context of the Sofja Kovalevskaja Award. The Humboldt foundation is funded by the German Federal Ministry for Education and Research. The hydrodynamic simulations in this work were performed using the COSMOS Supercomputer in Cambridge (UK), which is sponsored by SGI, Intel, HEFCE and the Darwin Supercomputer of the University of Cambridge High Performance Computing Service (<http://www.hpc.cam.ac.uk/>), provided by Dell Inc. using Strategic Research Infrastructure Funding from the Higher Education Funding Council for England. COSMOS and DARWIN are part of the DIRAC high performance computing facility funded by STFC. M.V. is supported by the FP7 ERC grant ‘‘cosmoIGM’’ GA-257670, PRIN-MIUR and INFN/PD51 grants. J.S.B. acknowledges the support of a Royal Society University Research Fellowship. B.L. acknowledges support from the NSF Astronomy and Astrophysics Fellowship grant AST-1202963.

Funding for SDSS-III has been provided by the Alfred P. Sloan Foundation, the Participating Institutions, the National Science Foundation, and the U.S. Department of Energy Office of Science. The SDSS-III Web site is <http://www.sdss3.org/>.

SDSS-III is managed by the Astrophysical Research Consortium for the Participating Institutions of the SDSS-III Collaboration including the University of Arizona, the Brazilian Participation Group, Brookhaven National Laboratory, University of Cambridge, Carnegie Mellon University, University of Florida, the French Participation Group, the German Participation Group, Harvard University, the Instituto de Astrofísica de Canarias, the Michigan State/Notre Dame/JINA Participation Group, Johns Hopkins University, Lawrence Berkeley National Laboratory, Max Planck Institute for Astrophysics, Max Planck Institute for Extraterrestrial Physics, New Mexico State University, New York University, Ohio State University, Pennsylvania State University, University of Portsmouth, Princeton University, the Spanish Participation Group, University of Tokyo, University of Utah, Vanderbilt University, University of Virginia, University of Washington, and Yale University.

## APPENDIX

In this Appendix, we describe our probabilistic procedure for combining the multiple BOSS exposures of each spectrum<sup>27</sup> while simultaneously estimating the noise variance in terms of a parameterized model. We assume that the noise in each pixel can be described by

$$\sigma_{\lambda i}^2 = A_1 \hat{S}_{\lambda i} (\mathcal{F}_\lambda + s_{\lambda i}) + A_2 \hat{S}_{\lambda i}^2 \sigma_{\text{RN,eff}}^2 \sigma_{\text{disp}}^2(\lambda), \quad (\text{A1})$$

where

$$\hat{S}_{\lambda i} = S_{\lambda i} (1 - \exp(-A_3 \lambda + A_4)). \quad (\text{A2})$$

The true object flux  $\mathcal{F}_\lambda$  and  $A_{j=1-4}$  are noise parameters that we will determine given the individual exposure spectra  $f_{\lambda,i}$ , sky flux estimates  $s_{\lambda,i}$ , and calibrations vectors  $S_{\lambda,i}$  (which convert between detector counts and photons).  $\sigma_{\text{RN,eff}}$  is the effective read noise that we fixed to  $\sigma_{\text{RN,eff}} = 12$ ; this can be thought of as an effective number of pixels times the true read noise of the CCD squared, which we multiplied by the spectrograph dispersion  $\sigma_{\text{disp}}(\lambda)$  to approximately account for the change in spot size as a function of wavelength. Equation (A2) parameterizes wavelength-dependent biases in the calibration vector.

We search for the model that best describes the multiple exposure spectra  $f_{\lambda,i}$ , where our model parameters are  $A_j$  from Equation (A1) and  $\mathcal{F}_\lambda$  is the true flux of the object. In what follows, we will outline a method for determining the posterior distribution  $P(A_j, \mathcal{F}_\lambda | f_{\lambda,i})$  using an MCMC method. From this distribution, we can obtain both an accurate model for the noise via Equation (A1), and our final combined spectrum. The estimates for  $A_j$  can also be used to self-consistently generate pixel noise in mock Ly $\alpha$  forest spectra.

The probability of the data given the model, or the likelihood, can be written as

$$\begin{aligned} L(A, \mathcal{F}_\lambda) &= P(f_{\lambda,i} | A_j, \mathcal{F}_\lambda) \\ &= \prod_{\lambda i} \frac{1}{\sqrt{2\pi} \sigma_{\lambda i}} \exp\left(-\frac{(f_{\lambda i} - \mathcal{F}_\lambda)^2}{2\sigma_{\lambda i}^2}\right). \end{aligned} \quad (\text{A3})$$

Note that individual exposure data  $f_{\lambda,i}$  are on the native wavelength grid of each CCD exposure, whereas the BOSS pipeline interpolates and then combines these individual spectra into a final co-added spectrum, defined on a wavelength grid with uniform spacing. Furthermore, flexure and other variations in the spectrograph wavelength solution will result in small (typically sub-pixel) shifts between the individual exposure wavelength grids. In Equation (A3) our model  $\mathcal{F}_\lambda$  must be computable at every wavelength  $f_{\lambda,i}$  of the individual exposures. We are free to choose the wavelengths at which  $\mathcal{F}_\lambda$  is represented, but this choice is a subtle issue for several reasons. First, note that we want to avoid interpolating the data,  $f_{\lambda,i}$ , onto the model wavelength grid, as this would correlate the data pixels, and require that we track covariances in the likelihood in Equation (A3), making it significantly more complicated and challenging to evaluate. Similarly, it is undesirable to interpolate our model  $\mathcal{F}_\lambda$ , as this would introduce correlations in the model parameters, making it much more difficult to sample them with our MCMC. Finally, note that  $\mathcal{F}_\lambda$  also represents our final co-added spectrum, so we might consider opting for a uniform wavelength grid, similar to what is done by the BOSS pipeline. Our approach is to simply determine the model flux  $\mathcal{F}_\lambda$  at each wavelength of the individual exposures  $f_{\lambda,i}$ . Shifts among the individual exposure wavelength grids result in a more finely sampled model grid. For the reasons explained above, we use nearest grid point (NGP) interpolation, so that the  $f_{\lambda,i}$  are evaluated on the  $\mathcal{F}_\lambda$  grid (and vice versa) by assigning the value from the single nearest pixel.

In our MCMC iterations, we use the standard Metropolis–Hastings criterion to sample the parameters  $A_j$ , with trials drawn from a uniform prior. For the  $\mathcal{F}_\lambda$ , we exploit an analogy with Gibbs sampling, which dramatically simplifies the MCMC for likelihood functions with a multivariate Gaussian form. Gibbs sampling exploits the fact that given a multivariate distribution, it is much simpler to sample from conditional distributions than to integrate over a joint distribution. To be more specific, the likelihood in Equation (A3) is proportional to the joint probability distribution of the noise parameters  $A_j$  and  $\mathcal{F}_\lambda$ , but it is also proportional to the conditional probability distribution of the  $\mathcal{F}_\lambda$  at fixed  $A_j$ . With  $A_j$  fixed the probability of  $\mathcal{F}_\lambda$  is then

$$P(\mathcal{F}_\lambda | A, f_{\lambda,i}) \propto \prod_{\lambda i} \frac{1}{\sqrt{2\pi} \sigma_{\lambda i}} \exp\left(-\frac{(f_{\lambda i} - \mathcal{F}_\lambda)^2}{2\sigma_{\lambda i}^2}\right), \quad (\text{A4})$$

which is very nearly a multivariate Gaussian distribution for  $\mathcal{F}_\lambda$  with a diagonal covariance matrix. The equation above slightly deviates from a Gaussian because the  $\sigma_{\lambda i}$  depends on  $\mathcal{F}_\lambda$  via Equation (A1). In what follows, we ignore this small deviation, and assume that the conditional PDF of the  $\mathcal{F}_\lambda$  (at fixed  $A_j$ ) is Gaussian.

Given that Equation (A4) is a multivariate Gaussian with diagonal covariance, the Gibbs sampling of the  $\mathcal{F}_\lambda$  becomes trivial. Since Equation (A4) can be factored into a product of individual Gaussians, we need not follow the standard Gibbs sampling algorithm, whereby each parameter is updated sequentially holding the others fixed. Instead we need only hold  $A_j$  fixed (since the likelihood is not Gaussian in these parameters), and we can sample all of the  $\mathcal{F}_\lambda$  simultaneously. This simplification, which dramatically speeds up the algorithm, is possible because the conditional distribution for  $\mathcal{F}_\lambda$  can be factored into a product of Gaussians for each pixel  $\mathcal{F}_\lambda$ , thus the conditional distribution at any wavelength is completely independent of all the others.

<sup>27</sup> Defined as unique combinations of plate number, fiber number, and MJD of observation.

Completing the square in Equation (A4) we can then write

$$P(\mathcal{F}_\lambda | A_j, f_{\lambda i}) \propto \prod_{\lambda} \exp\left(-\frac{(f_{\text{opt},\lambda} - \mathcal{F}_\lambda)^2}{2\sigma_{\text{opt},\lambda}^2}\right), \quad (\text{A5})$$

where

$$f_{\text{opt},\lambda} \equiv \frac{1}{\sigma_{\text{opt},\lambda}^2} \sum_i \frac{f_{\lambda i}}{\sigma_{\lambda i}^2} \quad \text{and} \quad \frac{1}{\sigma_{\text{opt},\lambda}^2} \equiv \sum_i \frac{1}{\sigma_{\lambda i}^2}. \quad (\text{A6})$$

The expressions above for  $f_{\text{opt},\lambda}$  and  $\sigma_{\text{opt},\lambda}^2$  simply represent the *optimally* combined flux estimator and the resulting variance. Thus one can think of our MCMC algorithm as performing an optimal combination of the individual exposure spectra  $f_{\lambda i}$ , whereby the noise is simultaneously determined via an iterative procedure.

Thus the basic steps of our algorithm can be summarized as follows.

1. Initialize, by creating a model  $\lambda$  grid from all unique wavelengths in the individual exposures, and use NGP interpolation to assign a  $f_{\lambda i}$  to this grid for each exposure.
2. Choose a starting guess for noise parameters  $A_j$ . For the starting  $\mathcal{F}_\lambda$  use  $\mathcal{F}_\lambda = f_{\text{opt},\lambda}$  from Equation (A6), but with the model  $\sigma_{\lambda i}$  replaced by the noise delivered by the pipeline.
3. Begin the MCMC loop:
  - (1) Use the current values of  $A_j$  and  $\mathcal{F}_\lambda$  to compute the variance  $\sigma_{\lambda i}^2$  for each exposure via Equation (A1).
  - (2) Compute  $f_{\text{opt},\lambda}$  and  $\sigma_{\text{opt},\lambda}^2$  from Equation (A6).
  - (3) Take a Gibbs step for each wavelength of  $\mathcal{F}_\lambda = f_{\text{opt},\lambda} + g_\lambda \sigma_{\text{opt},\lambda}$  simultaneously, where  $g_\lambda$  is a vector of unit variance Gaussian deviates.
  - (4) Use NGP to interpolate the model  $\mathcal{F}_\lambda$  onto each individual exposure  $f_{\lambda i}$  wavelength grid.
  - (5) Compute the likelihood  $L(A_j, \mathcal{F}_\lambda)$  according to Equation (A3).
  - (6) Take trial steps in the  $A_j$  according to  $A_{j,\text{try}} = A_j + g_j dA_j$ , where  $dA_j$  is a stepsize and  $g_j$  is a Gaussian deviating between zero and one, drawn for each individual noise parameter  $A_j$ .
  - (7) Compute the likelihood at  $L(A_{j,\text{try}}, \mathcal{F}_\lambda)$ .
  - (8) Apply the Metropolis–Hastings criteria to the likelihood difference. If it is satisfied then accept the values of  $A_j$  as part of the Markov chain. If not, then use the previous values. Note that the  $\mathcal{F}_\lambda$  are always accepted, because they are Gibbs sampled.
4. Use only the second half of the chain for the posterior distributions, as the first half is the burn in phase.

Our MCMC algorithm directly determines the posterior distribution  $P(A_j, \mathcal{F}_\lambda | f_{\lambda i})$ , which provides all the information we need to construct mock spectra using Equation (A1) as described in Section 4.5.

The distribution of  $P(\mathcal{F}_\lambda | f_{\lambda i})$ , on the other hand, contains everything we need to know about the combined spectrum. Namely, we can define

$$\bar{\mathcal{F}}_\lambda \equiv \int P(\mathcal{F}_\lambda | f_{\lambda i}) \mathcal{F}_\lambda d\mathcal{F}_\lambda \quad (\text{A7})$$

as the combined spectrum, and

$$\sigma_\lambda^2 \equiv \int P(\mathcal{F}_\lambda | f_{\lambda i}) (\mathcal{F}_\lambda - \bar{\mathcal{F}}_\lambda)^2 d\mathcal{F}_\lambda \quad (\text{A8})$$

as its variance. If the formal noise returned by BOSS pipeline were actually the true noise in the data, then our  $\bar{\mathcal{F}}_\lambda$  in Equation (A8) would be equivalent to the optimally combined noise and our variance the optimal variance, i.e., according to Equation (A6). In practice, the BOSS pipeline does not return the true noise and so our  $\bar{\mathcal{F}}_\lambda$  is optimal whereas the pipeline flux is sub-optimal, and our  $\sigma_\lambda^2$  is an empirical estimate of the actual noise in the data.

## REFERENCES

- Abazajian, K., Adelman-McCarthy, J. K., Agüeros, M. A., et al. 2005, *AJ*, **129**, 1755
- Abazajian, K. N., Adelman-McCarthy, J. K., Agüeros, M. A., et al. 2009, *ApJS*, **182**, 543
- Ahn, C. P., Alexandroff, R., Allende Prieto, C., et al. 2012, *ApJS*, **203**, 21
- Ahn, C. P., Alexandroff, R., Allende Prieto, C., et al. 2014, *ApJS*, **211**, 17
- Almgren, A. S., Bell, J. B., Lijewski, M. J., Lukić, Z., & Van Andel, E. 2013, *ApJ*, **765**, 39
- Altay, G., Theuns, T., Schaye, J., Booth, C. M., & Dalla Vecchia, C. 2013, *MNRAS*, **436**, 2689
- Altay, G., Theuns, T., Schaye, J., Crighton, N. H. M., & Dalla Vecchia, C. 2011, *ApJL*, **737**, L37
- Anderson, L., Aubourg, É., Bailey, S., et al. 2014, *MNRAS*, **441**, 24
- Becker, G. D., Bolton, J. S., Haehnelt, M. G., & Sargent, W. L. W. 2011, *MNRAS*, **410**, 1096
- Becker, G. D., Hewett, P. C., Worseck, G., & Prochaska, J. X. 2013, *MNRAS*, **430**, 2067
- Bernardi, M., Sheth, R. K., SubbaRao, M., et al. 2003, *AJ*, **125**, 32
- Bird, S., Vogelsberger, M., Sijacki, D., et al. 2013, *MNRAS*, **429**, 3341
- Bolton, A. S., & Schlegel, D. J. 2010, *PASP*, **122**, 248
- Bolton, A. S., Schlegel, D. J., Aubourg, É., et al. 2012, *AJ*, **144**, 144
- Bolton, J. S., Becker, G. D., Haehnelt, M. G., & Viel, M. 2014, *MNRAS*, **438**, 2499
- Bolton, J. S., Oh, S. P., & Furlanetto, S. R. 2009, *MNRAS*, **395**, 736
- Bolton, J. S., Viel, M., Kim, T., Haehnelt, M. G., & Carswell, R. F. 2008, *MNRAS*, **386**, 1131
- Borde, A., Palanque-Delabrouille, N., Rossi, G., et al. 2014, *JCAP*, **07**, 005
- Bovy, J., Hennawi, J. F., Hogg, D. W., et al. 2011, *ApJ*, **729**, 141
- Broderick, A. E., Chang, P., & Pfrommer, C. 2012, *ApJ*, **752**, 22
- Busca, N. G., Delubac, T., Rich, J., et al. 2013, *A&A*, **552**, A96
- Calura, F., Tescari, E., D'Odorico, V., et al. 2012, *MNRAS*, **422**, 3019
- Cen, R., Miralda-Escudé, J., Ostriker, J. P., & Rauch, M. 1994, *ApJL*, **437**, L9
- Chang, P., Broderick, A. E., & Pfrommer, C. 2012, *ApJ*, **752**, 23
- Compostella, M., Cantalupo, S., & Porciani, C. 2013, *MNRAS*, **435**, 3169
- Cooksey, K. L., Kao, M. M., Simcoe, R. A., O'Meara, J. M., & Prochaska, J. X. 2013, *ApJ*, **763**, 37
- Croft, R. A. C., Weinberg, D. H., Katz, N., & Hernquist, L. 1998, *ApJ*, **495**, 44
- Dall'Aglio, A., Wisotzki, L., & Worseck, G. 2008, *A&A*, **491**, 465
- Davé, R., Hernquist, L., Katz, N., & Weinberg, D. H. 1999, *ApJ*, **511**, 521
- Dawson, K. S., Schlegel, D. J., Ahn, C. P., et al. 2013, *AJ*, **145**, 10
- Desjacques, V., Nusser, A., & Sheth, R. K. 2007, *MNRAS*, **374**, 206
- Draine, B. T. 2011, *Physics of the Interstellar and Intergalactic Medium* (Princeton, NJ: Princeton Univ. Press)
- Eisenstein, D. J., Weinberg, D. H., Agol, E., et al. 2011, *AJ*, **142**, 72
- Faucher-Giguère, C., Prochaska, J. X., Lidz, A., Hernquist, L., & Zaldarriaga, M. 2008, *ApJ*, **681**, 831
- Font-Ribera, A., McDonald, P., & Miralda-Escudé, J. 2012, *JCAP*, **01**, 001
- Font-Ribera, A., & Miralda-Escudé, J. 2012, *JCAP*, **07**, 028
- Furlanetto, S. R., & Oh, S. P. 2008, *ApJ*, **682**, 14
- Garzilli, A., Bolton, J. S., Kim, T.-S., Leach, S., & Viel, M. 2012, *MNRAS*, **424**, 1723
- Gleser, L., Nusser, A., Benson, A. J., Ohno, H., & Sugiyama, N. 2005, *MNRAS*, **361**, 1399
- Gnedin, N. Y., & Hui, L. 1998, *MNRAS*, **296**, 44
- Gunn, J. E., & Peterson, B. A. 1965, *ApJ*, **142**, 1633
- Gunn, J. E., Siegmund, W. A., Mannery, E. J., et al. 2006, *AJ*, **131**, 2332
- Horne, K. 1986, *PASP*, **98**, 609
- Hui, L., & Gnedin, N. Y. 1997, *MNRAS*, **292**, 27
- Hui, L., Gnedin, N. Y., & Zhang, Y. 1997, *ApJ*, **486**, 599
- Hui, L., & Haïman, Z. 2003, *ApJ*, **596**, 9
- Janknecht, E., Reimers, D., Lopez, S., & Tytler, D. 2006, *A&A*, **458**, 427
- Jenkins, E. B., & Ostriker, J. P. 1991, *ApJ*, **376**, 33
- Kim, T., Bolton, J. S., Viel, M., Haehnelt, M. G., & Carswell, R. F. 2007, *MNRAS*, **382**, 1657



- Komatsu, E., Smith, K. M., Dunkley, J., et al. 2011, *ApJS*, 192, 18
- Lee, K.-G. 2012, *ApJ*, 753, 136
- Lee, K.-G., Bailey, S., Bartsch, L. E., et al. 2013, *AJ*, 145, 69
- Lee, K.-G., Hennawi, J. F., Stark, C., et al. 2014a, *ApJL*, 795, L12
- Lee, K.-G., Hennawi, J. F., White, M., Croft, R. A. C., & Ozbek, M. 2014b, *ApJ*, 788, 49
- Lee, K.-G., & Spergel, D. N. 2011, *ApJ*, 734, 21
- Lee, K.-G., Suzuki, N., & Spergel, D. N. 2012, *AJ*, 143, 51
- Lewis, A., Challinor, A., & Lasenby, A. 2000, *ApJ*, 538, 473
- Lidz, A., Faucher-Giguère, C.-A., Dall'Aglio, A., et al. 2010, *ApJ*, 718, 199
- Lukić, Z., Stark, C., Nugent, P., et al. 2014, *MNRAS*, 446, 3697
- Lundgren, B. F., Brunner, R. J., York, D. G., et al. 2009, *ApJ*, 698, 819
- Lynds, R. 1971, *ApJL*, 164, L73
- McDonald, P., & Eisenstein, D. J. 2007, *PhRvD*, 76, 063009
- McDonald, P., Miralda-Escudé, J., Rauch, M., et al. 2000, *ApJ*, 543, 1
- McDonald, P., Miralda-Escudé, J., Rauch, M., et al. 2001, *ApJ*, 562, 52
- McDonald, P., Seljak, U., Burles, S., et al. 2006, *ApJS*, 163, 80
- McDonald, P., Seljak, U., Cen, R., Bode, P., & Ostriker, J. P. 2005a, *MNRAS*, 360, 1471
- McDonald, P., Seljak, U., Cen, R., et al. 2005b, *ApJ*, 635, 761
- McQuinn, M., Lidz, A., Zaldarriaga, M., et al. 2009, *ApJ*, 694, 842
- McQuinn, M., Oh, S. P., & Faucher-Giguère, C.-A. 2011, *ApJ*, 743, 82
- McQuinn, M., & White, M. 2011, *MNRAS*, 415, 2257
- Meiksin, A., & Tittley, E. R. 2012, *MNRAS*, 423, 7
- Meiksin, A. A. 2009, *RvMP*, 81, 1405
- Miralda-Escudé, J., Cen, R., Ostriker, J. P., & Rauch, M. 1996, *ApJ*, 471, 582
- Miralda-Escudé, J., & Rees, M. J. 1994, *MNRAS*, 266, 343
- Noterdaeme, P., Petitjean, P., Carithers, W. C., et al. 2012, *A&A*, 547, L1
- O'Shea, B. W., Bryan, G., Bordner, J., et al. 2004, arXiv:astro-ph/0403044
- Palanque-Delabrouille, N., Yèche, C., Borde, A., et al. 2013, *A&A*, 559, A85
- Pâris, I., Petitjean, P., Aubourg, É., et al. 2012, *A&A*, 548, A66
- Pâris, I., Petitjean, P., Rollinde, E., et al. 2011, *A&A*, 530, A50
- Penton, S. V., Shull, J. M., & Stocke, J. T. 2000, *ApJ*, 544, 150
- Petitjean, P., Webb, J. K., Rauch, M., Carswell, R. F., & Lanzetta, K. 1993, *MNRAS*, 262, 499
- Pieri, M. M., Frank, S., Weinberg, D. H., Mathur, S., & York, D. G. 2010, *ApJL*, 724, L69
- Pieri, M. M., Mortonson, M. J., Frank, S., et al. 2014, *MNRAS*, 441, 1718
- Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2014, *A&A*, 571, A16
- Prochaska, J. X., Herbert-Fort, S., & Wolfe, A. M. 2005, *ApJ*, 635, 123
- Prochaska, J. X., O'Meara, J. M., & Worseck, G. 2010, *ApJ*, 718, 392
- Prochaska, J. X., & Wolfe, A. M. 1997, *ApJ*, 487, 73
- Puchwein, E., Pfrommer, C., Springel, V., Broderick, A. E., & Chang, P. 2012, *MNRAS*, 423, 149
- Rahmati, A., Pawlik, A. H., Raic'evic', M., & Schaye, J. 2013, *MNRAS*, 430, 2427
- Rauch, M. 1998, *ARA&A*, 36, 267
- Rauch, M., Carswell, R. F., Chaffee, F. H., et al. 1992, *ApJ*, 390, 387
- Ribaudo, J., Lehner, N., & Howk, J. C. 2011, *ApJ*, 736, 42
- Ricotti, M., Gnedin, N. Y., & Shull, J. M. 2000, *ApJ*, 534, 41
- Rollinde, E., Theuns, T., Schaye, J., Pâris, I., & Petitjean, P. 2013, *MNRAS*, 428, 540
- Rorai, A., Hennawi, J. F., & White, M. 2013, *ApJ*, 775, 81
- Ross, N. P., Myers, A. D., Sheldon, E. S., et al. 2012, *ApJS*, 199, 3
- Rudie, G. C., Steidel, C. C., & Pettini, M. 2012, *ApJL*, 757, L30
- Rudie, G. C., Steidel, C. C., Shapley, A. E., & Pettini, M. 2013, *ApJ*, 769, 146
- Schaye, J., Aguirre, A., Kim, T.-S., et al. 2003, *ApJ*, 596, 768
- Schaye, J., Theuns, T., Leonard, A., & Efstathiou, G. 1999, *MNRAS*, 310, 57
- Schaye, J., Theuns, T., Rauch, M., Efstathiou, G., & Sargent, W. L. W. 2000, *MNRAS*, 318, 817
- Schmidt, M. 1965, *ApJ*, 141, 1295
- Schneider, D. P., Richards, G. T., Hall, P. B., et al. 2010, *AJ*, 139, 2360
- Seljak, U., Makarov, A., McDonald, P., et al. 2005, *PhRvD*, 71, 103515
- Shull, J. M., France, K., Danforth, C. W., Smith, B., & Tumlinson, J. 2010, *ApJ*, 722, 1312
- Sironi, L., & Giannios, D. 2014, *ApJ*, 787, 49
- Slosar, A., Font-Ribera, A., Pieri, M. M., et al. 2011, *JCAP*, 09, 001
- Slosar, A., Iršič, V., Kirkby, D., et al. 2013, *JCAP*, 04, 026
- Smee, S. A., Gunn, J. E., Uomoto, A., et al. 2013, *AJ*, 146, 32
- Springel, V., Di Matteo, T., & Hernquist, L. 2005, *MNRAS*, 361, 776
- Stoughton, C., Lupton, R. H., Bernardi, M., et al. 2002, *AJ*, 123, 485
- Suzuki, N. 2006, *ApJS*, 163, 110
- Suzuki, N., Tytler, D., Kirkman, D., O'Meara, J. M., & Lubin, D. 2005, *ApJ*, 618, 592
- Telfer, R. C., Zheng, W., Kriss, G. A., & Davidsen, A. F. 2002, *ApJ*, 565, 773
- Theuns, T., Leonard, A., Efstathiou, G., Pearce, F. R., & Thomas, P. A. 1998, *MNRAS*, 301, 478
- Theuns, T., Schaye, J., Zaroubi, S., et al. 2002, *ApJL*, 567, L103
- Tytler, D., Kirkman, D., O'Meara, J. M., et al. 2004, *ApJ*, 617, 1
- Vanden Berk, D. E., Richards, G. T., Bauer, A., et al. 2001, *AJ*, 122, 549
- Viel, M., Becker, G. D., Bolton, J. S., & Haehnelt, M. G. 2013a, *PhRvD*, 88, 043502
- Viel, M., Bolton, J. S., & Haehnelt, M. G. 2009, *MNRAS*, 399, L39
- Viel, M., & Haehnelt, M. G. 2006, *MNRAS*, 365, 231
- Viel, M., Haehnelt, M. G., & Springel, V. 2004, *MNRAS*, 354, 684
- Viel, M., Schaye, J., & Booth, C. M. 2013b, *MNRAS*, 429, 1734
- Worseck, G., & Prochaska, J. X. 2011, *ApJ*, 728, 23
- Worseck, G., Prochaska, J. X., McQuinn, M., et al. 2011, *ApJL*, 733, L24
- York, D. G., Adelman, J., Anderson, J. E., Jr., et al. 2000, *AJ*, 120, 1579
- Zhu, G., Ménard, B., Bizyaev, D., et al. 2014, *MNRAS*, 439, 3139