The Astrophysical Journal, 872:101 (35pp), 2019 February 10 https://doi.org/10.3847/1538-4357/aafee4
© 2019. The American Astronomical Society. All rights reserved.

Revealing Reionization with the Thermal History of the Intergalactic Medium: New Constraints from the Lyα Flux Power Spectrum

Elisa Boera[1], George D. Becker[1], James S. Bolton[2], and Fahad Nasir[1]
[1] Department of Physics and Astronomy, University of California, Riverside, CA 92521, USA; elisa.boera@gmail.com
[2] School of Physics and Astronomy, University of Nottingham, University Park, Nottingham, NG7 2RD, UK
Received 2018 September 17; revised 2018 December 21; accepted 2019 January 9; published 2019 February 14

Abstract

We present a new investigation of the thermal history of the intergalactic medium (IGM) during and after reionization using the Lyα forest flux power spectrum at $4.0 \lesssim z \lesssim 5.2$. Using a sample of 15 high-resolution spectra, we measure the flux power down to the smallest scales ever probed at these redshifts ($-1 \lesssim \log(k/\mathrm{km}^{-1}\,\mathrm{s}) \lesssim -0.7$). These scales are highly sensitive to both the instantaneous temperature of the IGM and the total energy injected per unit mass during and after reionization. We measure temperatures at the mean density of $T_0 \sim 7000$–8000 K, consistent with no significant temperature evolution for redshifts $4.2 \lesssim z \lesssim 5.0$. We also present the first observational constraints on the integrated IGM thermal history, finding that the total energy input per unit mass increases from $u_0 \sim 4.6$ to 7.3 eV m_p^{-1} from $z \sim 6$ to 4.2 assuming a ΛCDM cosmology. We show how these results can be used simultaneously to obtain information on the timing and the sources of the reionization process. Our first proof of concept using simplistic models of instantaneous reionization produces results comparable to and consistent with the recent Planck constraints, favoring models with $z_\mathrm{rei} \sim 8.5^{+1.1}_{-0.8}$.

Key words: cosmology: observations – early universe – methods: observational – methods: statistical – quasars: absorption lines

Supporting material: tar.gz file

1. Introduction

The epoch of hydrogen reionization represents one of the most dramatic phases of evolution of the universe. During this period, the UV radiation from the first luminous sources reionized the neutral hydrogen (and He I) atoms in the diffuse intergalactic medium (IGM), driving the transition from a neutral to a highly ionized universe. Understanding sources and timing of this transformation can reveal crucial information on the properties of the first objects and the environment in which they were formed. When and how reionization happened therefore remains a primary subject of interest in extragalactic astrophysics (for a review, see Becker et al. 2015a).

The most direct probes of the highly ionized IGM have been obtained from observations of intergalactic Lyα absorption along the lines of sight to high-redshift quasars. Measurements of Lyα transmission along some lines of sight suggest that reionization was largely complete by $z \sim 6$ (e.g., McGreer et al. 2015). On the other hand, large fluctuations in IGM opacity remain at $z \lesssim 6$, suggesting that lingering evidence of reionization may remain in the IGM to somewhat lower redshifts (Fan et al. 2006; Becker et al. 2015b; Bosman et al. 2018; Eilers et al. 2018). While current constraints from cosmic microwave background (CMB) observations are consistent with a rapid reionization at redshift $z_\mathrm{rei} \simeq 7.7 \pm 0.7$ (Planck Collaboration et al. 2018), measurements of the fraction of neutral hydrogen at high redshift have also been obtained from the presence of Lyα damping wings (Mortlock et al. 2011; Simcoe et al. 2012; Greig et al. 2017; Davies et al. 2018) and from the weakening of Lyα emission lines in $z \sim 6$–8 galaxies (e.g., Caruana et al. 2014; Schmidt et al. 2016; Sadoun et al. 2017; Mason et al. 2018). The available data seem to generally support a late reionization scenario (with the bulk of reionization happening at $z \sim 6$–8) but are still consistent with a relatively broad range of reionization histories.

The sources responsible for reionization also remain uncertain. Star-forming galaxies have commonly been considered the most likely candidate (e.g., Finkelstein 2016; Bouwens et al. 2015, 2016). Scenarios in which active galactic nuclei (AGNs) make a substantial contribution, however, continue to be considered (e.g., Giallongo et al. 2015; Madau & Haardt 2015; D'Aloisio et al. 2017; Parsa et al. 2018).

Further insight may be gained from using the IGM thermal history to constrain the reionization process. The temperature of the IGM should increase significantly via photoionization heating during hydrogen reionization, and because its cooling time is long, the low-density gas retains some useful memory of when and how it was reionized (Miralda-Escudé & Rees 1994; Abel & Haehnelt 1999; Upton Sanderbeck et al. 2016). At the mean density of the IGM the characteristic signature of reionization is expected to be an increase in temperature of tens of thousands of kelvin as an ionization front sweeps through (e.g., D'Aloisio et al. 2018), followed by cooling (over $\Delta z \sim 1$–2) toward a thermal asymptote set primarily by the balance between photoheating by the UV background (UVB) and adiabatic cooling due to the expansion of the universe (e.g., McQuinn et al. 2009). The interplay among these effects is expected to lead to a power-law temperature–density (T–ρ) relation for the low-density gas ($\Delta = \rho/\bar{\rho} \lesssim 10$) of the form

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

where T_0 is the temperature at the mean density and $\gamma - 1$ is the slope of the relation (Hui & Gnedin 1997; Puchwein et al. 2015; McQuinn & Upton Sanderbeck 2016).

Following reionization, the increase in gas pressure due to the boost in temperature smooths out of the gas on small scales (e.g., Gnedin & Hui 1998; Rorai et al. 2013; Kulkarni et al. 2015). The degree of "Jeans smoothing" in the IGM prior to a given redshift is sensitive to timing and the total heat injection during and after reionization. Measurements of both the gas temperature evolution and the Jeans smoothing at redshifts approaching reionization ($z \gtrsim 4$) can therefore constrain the timing of this process and potentially provide information on the nature of the ionizing sources.

In the past two decades the Lyα forest in quasar spectra has been the main laboratory for the study of the thermal state of the IGM. In combination with cosmological hydrodynamical simulations, previous efforts have used a variety of statistical approaches to measure the IGM temperature–density relation at $1.5 \lesssim z \lesssim 5.4$ from the shapes of the Lyα absorption lines (e.g., Ricotti et al. 2000; Schaye et al. 2000; McDonald et al. 2001; Theuns et al. 2002; Bolton et al. 2008, 2014; Lidz et al. 2010; Becker et al. 2011; Rudie et al. 2012; Boera et al. 2014, 2016; Hiss et al. 2018; Rorai et al. 2017a, 2018). However, the widths of these features are sensitive to both the instantaneous temperature of the gas (thermal broadening) and Jeans smoothing (which increases the Hubble broadening) due to the heat injection at higher redshifts. In previous works the impact of pressure smoothing either has generally not been included or has been considered a source of systematic error. For example, Viel et al. (2013a) and Iršič et al. (2017a) account for this effect by adding the redshift of reionization as a nuisance parameter for their warm dark matter constraints.

On the other hand, the first direct measurement of the characteristic filtering scale over which the gas is pressure smoothed (λ_P) has recently been obtained from the analysis of the Lyα absorption correlations using close quasar pairs at $z \sim 2$–3 (Rorai et al. 2017b). This method largely disentangles the impacts of thermal broadening and pressure smoothing; however, the lack of known quasar pairs at higher redshifts prevents it from being used at redshifts closer to hydrogen reionization.

An alternative means of simultaneously constraining temperature and Jeans smoothing is presented by the Lyα flux power spectrum (Puchwein et al. 2015; Nasir et al. 2016, hereafter N16; Walther et al. 2019). N16 demonstrated, using hydrodynamical simulations, that the Lyα flux power spectrum exhibits different scale dependences for the temperature and Jeans smoothing. In particular, probing small scales (wavenumber $\log(k/\mathrm{km}^{-1}\,\mathrm{s}) \gtrsim -1$) increases the sensitivity to different reionization scenarios (see also Oñorbe et al. 2017 for an independent analysis). Although the 1D flux power spectrum statistic has been already explored in several works (e.g., Kim et al. 2004; McDonald et al. 2006; Viel et al. 2013a; Palanque-Delabrouille et al. 2015; Iršič et al. 2017a; Yèche et al. 2017), the lack of high-resolution, high signal-to-noise ratio (S/N) Lyα forest spectra has so far prevented these small scales from being measured at redshifts approaching reionization (but see Walther et al. 2018 for an analysis at $z < 4$).

In this paper we present a new measurement of the Lyα flux power spectrum at $z \sim 4$–5.2 obtained from a sample of high-resolution, high-S/N spectra. We extend the measurement to previously unexplored small scales ($\log(k/\mathrm{km}^{-1}\,\mathrm{s}) = -0.7$). By comparing the data to predictions from a suite of hydrodynamical simulations, we investigate the IGM temperature evolution and, simultaneously, its integrated thermal history. We then demonstrate how the combined constraints offer new insights on the timing and sources of the hydrogen reionization process.

Table 1
List of Quasars Used for This Analysis

Name	z_{em}	$z_{\mathrm{Ly}\alpha}$		C/N	Instrument
		z_{start}	z_{end}		
J2111−0156	4.89	3.99	4.79	20	HIRES
J0011+1446	4.94	4.03	4.84	33	HIRES
J1425+0827	4.95	4.04	4.85	40	UVES
J1008−0212	5.04	4.11	4.94	22	UVES
J1101+0531	5.05	4.12	4.94	23	UVES
J0025−0145	5.07	4.12	4.95	26	HIRES
J1204−0021	5.09	4.16	4.99	15	HIRES
J0131−0321	5.12	4.23	5.08	20	HIRES
J0957+0610	5.17	4.22	5.07	27	UVES
J0741+2520	5.19	4.24	5.09	17	HIRES
J0915+4924	5.20	4.25	5.10	16	HIRES
J0747+1153	5.26	4.30	5.16	18	HIRES
J1659+2709	5.32	4.34	5.21	25	HIRES
J0306+1853	5.36	4.37	5.22	42	HIRES
J0231−0728	5.42	4.43	5.31	31	HIRES

Note. For each object we report the name (Column (1)) based on the J2000 coordinates of the object and the emission redshift (Column (2)). The redshift intervals associated with the Lyα absorption used for this analysis are also reported with the corresponding median C/N level per pixel (Columns (3), (4), and (5)). Finally, the instrument with which the spectrum was taken is listed in Column (6).

For this work we have adopted the parameterization of the Jeans smoothing effect described in N16. We characterize the integrated thermal history of the IGM using the cumulative energy per unit mass, u_0, injected into the gas at the mean cosmic density during and after the reionization process. As we demonstrate, this quantity can be directly used to constrain reionization models.

This paper is organized as follows. In Section 2 we introduce the observational sample of high-resolution spectra. An overview of the simulations used to test and interpret the measurements is presented in Section 3. In Section 4 we introduce the power spectrum method, discussing the effect of the most relevant thermal parameters. In Section 5 we present the observational power spectrum results and discuss the strategies applied to take into account and reduce systematic uncertainties. The calibration and analysis of the synthetic power spectrum models are described in Section 6. The Markov chain Monte Carlo (MCMC) analysis, comparing models with the observational measurements, is described in Section 7, where we also present our main results for the IGM temperature at the mean density and the integrated thermal history. As an example of how our thermal constraints can be used to test reionization histories, we apply our results to instantaneous reionization models in Section 8. We summarize our findings and conclude in Section 9. Tests for various systematic effects are described in the appendices.

2. Observational Spectra

We obtained high-resolution spectra of a sample of 15 quasars spanning emission redshifts $4.8 \lesssim z_{\mathrm{em}} \lesssim 5.4$. The quasars and their basic properties are listed in Table 1. The spectra for 11 of the objects were obtained with the Keck High

Table 2
List of Hydrodynamical Simulations Used in This Work

Model	$L(h^{-1}$ cMpc)	$M_{\rm gas}(h^{-1}M_{\odot})$	$z_{\rm OT}$	ζ	$T_0^{z=5}$ (K)	$\gamma^{z=5}$	$u_0^{z=5}$(eV $m_{\rm p}^{-1}$)
S10_0.3z7	10	9.97×10^4	7	0.3	3162	1.52	1.3
S10_0.3z9	10	9.97×10^4	9	0.3	3388	1.49	2.3
S10_0.3z15	10	9.97×10^4	15	0.3	3388	1.51	5.0
S10_0.55z7	10	9.97×10^4	7	0.55	4553	1.51	1.9
S10_0.55z9	10	9.97×10^4	9	0.55	5086	1.48	3.3
S10_0.55z12	10	9.97×10^4	12	0.55	5110	1.51	5.2
S10_0.55z15	10	9.97×10^4	15	0.55	5093	1.52	7.4
S10_0.55z19	10	9.97×10^4	19	0.55	5074	1.52	10.0
S10_1z7	10	9.97×10^4	7	1.0	6607	1.50	2.7
S10_1z9	10	9.97×10^4	9	1.0	7413	1.51	4.7
S10_1z12	10	9.97×10^4	12	1.0	7510	1.51	7.6
S10_1z15 (10–512)	**10**	$\mathbf{9.97 \times 10^4}$	**15**	**1.0**	**7413**	**1.50**	**10.6**
S10_1z19	10	9.97×10^4	19	1.0	7457	1.52	14.7
S10_1.8z7	10	9.97×10^4	7	1.8	9725	1.49	3.9
S10_1.8z9	10	9.97×10^4	9	1.8	10866	1.50	6.8
S10_1.8z12	10	9.97×10^4	12	1.8	10900	1.51	10.9
S10_1.8z15	10	9.97×10^4	15	1.8	10865	1.51	15.5
S10_1.8z19	10	9.97×10^4	19	1.8	10827	1.52	21.4
S10_3.3z7	10	9.97×10^4	7	3.3	13803	1.48	5.5
S10_3.3z9	10	9.97×10^4	9	3.3	15488	1.50	9.9
S10_3.3z12	10	9.97×10^4	12	3.3	15821	1.48	16.3
S10_3.3z15	10	9.97×10^4	15	3.3	15488	1.52	23.2
S10_1z9_g1	10	9.97×10^4	9	1.0	7413	1.00	4.7
S20_1z15 (20–1024)	**20**	$\mathbf{9.97 \times 10^4}$	**15**	**1.0**	**7413**	**1.50**	**10.6**
S40_1z15 (40–2048)	**40**	$\mathbf{9.97 \times 10^4}$	**15**	**1.0**	**7413**	**1.50**	**10.6**
S40_1z9 (40–2048–zr9)	**40**	$\mathbf{9.97 \times 10^4}$	**9**	**1.0**	**7413**	**1.51**	**4.7**
S10_1z15_256	10	7.97×10^5	15	1.0	7413	1.50	10.6
S10_1z15_768	10	2.95×10^4	15	1.0	7413	1.50	10.6

Note. Entries in bold correspond to the Sherwood simulations first introduced in Bolton et al. (2017); the model names used in that work are given in brackets. For each simulation we report the name (Column (1)), box size (Column (2)), mass resolution (Column (3)), redshift at which the gas becomes optically thin (Column (4)), and constant factor used to rescale the photoheating rates for different thermal histories (Column (5)). The thermal parameters that describe the T–ρ relation at $z = 5$ are also listed: the temperature of the gas at the mean density (Column (6)) and the power-law index γ (Column (7)). Finally, the cumulative energy per unit mass deposited into the IGM at the mean density by $z = 5$ is given in Column (8) (see text for details). Further details on the simulation methodology are presented in Bolton et al. (2017).

Resolution Echelle Spectrometer (HIRES; Vogt et al. 1994), while the remaining four were taken with the Ultraviolet and Visual Echelle Spectrograph (UVES; Dekker et al. 2000) on the Very Large Telescope (VLT).

The spectra were reduced using a custom set of IDL routines that include optimal sky subtraction (Kelson 2003) and extraction techniques (Horne 1986). For each object a single 1D spectrum was extracted simultaneously from all exposures after individually applying telluric absorption corrections and relative flux calibration to the 2D frames. Telluric corrections were modeled based on the ESO SKYCALC Cerro Paranal Advanced Sky Model (Noll et al. 2012; Jones et al. 2013). For the UVES data we found that flux calibration derived from standard stars yielded sufficiently accurate agreement between overlapping orders. For HIRES, however, this approach produced well-known moderate (∼10%) interorder flux discrepancies. For all except one of our HIRES quasars, therefore, we used lower-resolution spectra from Keck/ESI, VLT/X-Shooter, or Gemini/GMOS to derive a custom response function for each exposure. The remaining object, J2111−0156, was calibrated using a response function from a standard star. We verified that our final flux power spectra remained essentially unchanged if standard-star flux calibration was used for every object. We therefore do not expect this aspect of the reduction to significantly impact our results.

The HIRES objects were observed using a $0\farcs86$ slit, giving a nominal resolution FWHM of $\sim$6 km s^{-1}. The UVES spectra were taken with a $1\farcs0$ slit, giving a nominal resolution of $\sim$7 km s^{-1}. The telluric models for the UVES data, however, indicated somewhat higher resolution consistent with a typical seeing of $0\farcs8$. Consequently, we adopt a resolution of 6 km s^{-1} for the full data set, which is sufficient to resolve small-scale features in the Lyα forest. We therefore expect that even the smallest scale of the flux power spectrum measured in this work ($\log(k/{\rm km}^{-1}\,{\rm s}) = -0.7$, or $\Delta v \sim 30$ km s^{-1}) will not be strongly affected by the finite spectroscopic resolution (but see Section 5.6.3). For all the quasars, the echelle orders were redispersed onto a common wavelength scale with a dispersion of 2.5 km s^{-1} per pixel.

According to the analysis presented in N16 using mock observations with a redshift path $\Delta z = 4$, a continuum-to-noise ratio (C/N) of $\sim$15 per 3 km s^{-1} pixel is necessary to break the degeneracy between thermal broadening and pressure smoothing and measure the thermal parameters with a statistical uncertainty of ∼20%. Conservatively, we have chosen our sample imposing this minimum threshold inside the Lyα forest region.

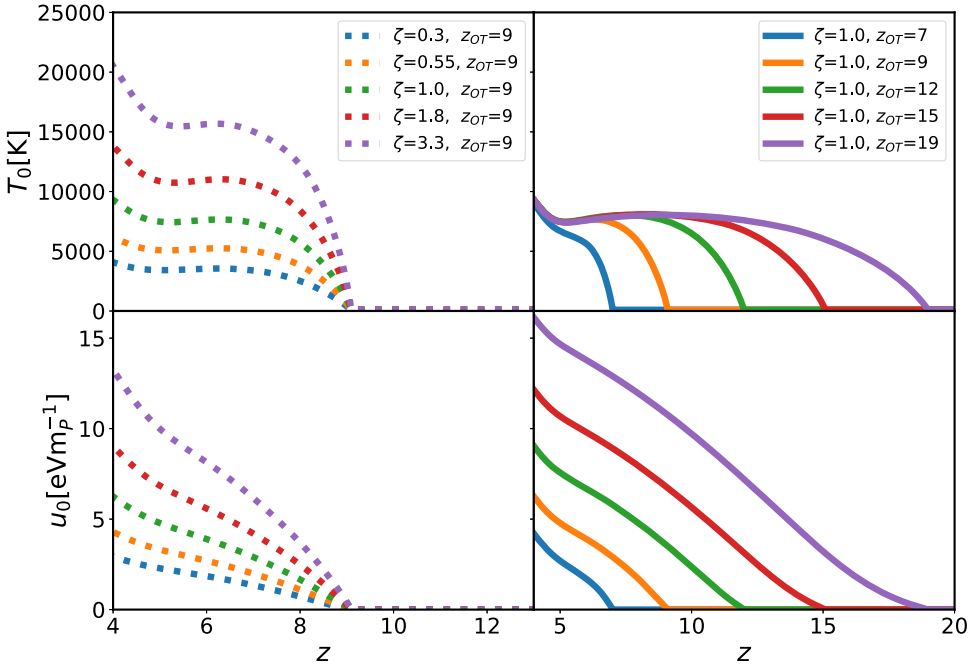

Figure 1. Examples of the evolution of parameters governing the thermal state of the IGM in our simulations. Left panels: evolution as a function of redshift of the temperature (top) and the cumulative energy per unit mass at the mean density (bottom) for models in which heating begins at the same $z_{\rm OT}$ but the photoheating rates are changed. Right panels: models with the same photoheating rates and different $z_{\rm OT}$. While increasing the photoheating rate produces larger values in both the temperature and u_0, models with different $z_{\rm OT}$ converge to the same value of T_0 provided that sufficient time has elapsed after the onset of heating. The values of these quantities at $z = 5$ for all our simulations are also listed in Table 2. The full suite of thermal histories is plotted in Figure 34.

We have fitted the continuum in our spectra using spline fits guided by power-law extrapolations of the continuum redward of the Lyα emission line. Given the high levels of absorption at $z \gtrsim 4$, the continuum measurements are necessarily characterized by large uncertainties (∼10%–20%). We therefore use these estimations only to derive a rough estimate of the C/N level. In measuring the power spectrum, as described in Section 5.1, we adopt an approach that does not require a priori knowledge of the continuum. The redshift coverage and the median C/N for the Lyα forest region of our sample are reported in Table 1. The majority of the spectra have larger C/N than our cut, with a typical value per pixel in the range of 20–30. This high C/N assures that the power spectrum measurement at small scales will not be strongly affected by uncertainties in the noise modeling.

3. The Simulations

To test systematics associated with the observed power spectrum and to interpret our observational results, we used synthetic spectra derived from hydrodynamical simulations and processed to closely match the characteristics of the real data. We ran a large set of hydrodynamical simulations that span a range of thermal histories at $z > 4$. The simulations run following the Sherwood simulations suite (Bolton et al. 2017), which uses a modified version of the parallel smoothed particle hydrodynamics code P-GADGET-3, an updated and extended version of GADGET-2 (Springel 2005). The models adopt the cosmological parameters $\Omega_m = 0.308$, $\Omega_\Lambda = 0.692$, $h = 0.678$, $\Omega_b = 0.0482$, $\sigma_8 = 0.829$, and $n_s = 0.961$, consistent with the CMB constraints of Planck Collaboration et al. (2014). Initial conditions were obtained using transfer functions generated by CAMB (Lewis et al. 2000). Because the vast majority of the absorption systems probed by the Lyα forest at $z > 4$ correspond to overdensities $\Delta = \rho/\bar{\rho} \lesssim 10$, our analysis will not be affected by the star formation prescription (Viel et al. 2013b). Therefore, to increase the computational speed, gas particles with temperature $T < 10^5$ K and overdensity $\Delta > 10^3$ are converted to collisionless particles (Viel et al. 2004).

The bulk of our simulations use a box size of 10 h^{-1} cMpc and 2×512^3 gas and dark matter particles, corresponding to a gas particle mass of $9.97 \times 10^4\, h^{-1}\, M_\odot$. In addition, we use runs with larger box size and different mass resolution to test numerical convergence (see Appendix F).

We note that our simulations are not intended to self-consistently model reionization. Instead, we employ models with a wide variety of thermal histories so that our ultimate constraints on the temperature and integrated heating of the IGM are as general as possible. The gas in our models becomes optically thin at a redshift $z_{\rm OT}$, after which it is photoionized and heated by a uniform UVB, which is a scaled version of the background from Haardt & Madau (2012). The thermal history of a given simulation is therefore determined by the choice of $z_{\rm OT}$ and UVB scaling factor.

The photoheating rates from Haardt & Madau (2012) ($\epsilon_i^{\rm HM12}$) for the different species (i = [H I, He I, He II]) have been rescaled proportionally by a constant factor ζ using the relation $\epsilon_i = \zeta \epsilon_i^{\rm HM12}$ (see Table 2). The combination of $z_{\rm OT}$ and ζ will determine both the instantaneous temperature and the total integrated heating per unit mass at the epoch where the power spectrum is measured. Models with larger $z_{\rm OT}$ and/or ζ will tend to have higher values of u_0.

A summary of the simulations used in this work is listed in Table 2. For each model we selected the simulation outputs between $4.0 \lesssim z \lesssim 5.4$ with a redshift step $\Delta z = 0.1$. At each redshift, synthetic spectra of the Lyα forest were produced by choosing 5000 "lines of sight" throughout the simulation box. In Section 6.2 we describe how these lines of sight were combined to create realistic mock spectra.

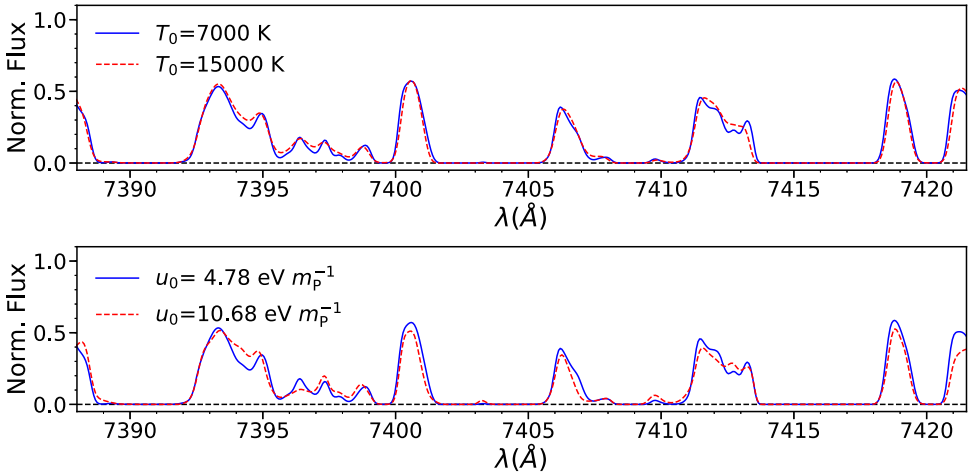

Figure 2. Effect of thermal broadening and integrated heating on simulated Lyα forest spectra at $z = 5$. Top panel: effect of thermal broadening on the absorption features in models characterized by the same integrated heating ($u_0 = 4.78$ eV m_p^{-1}) but post-processed to different instantaneous temperatures. Bottom panel: effect of pressure smoothing on the Lyα absorption for models with the same temperature ($T_0 = 7000$ K) but different thermal histories.

Following N16, the integrated thermal history in our simulations is parameterized using u_0, the cumulative energy deposited per unit mass into the gas at the mean density. At each redshift u_0 is defined as

$$u_0(z) = \int_z^{z_{\rm OT}} \frac{\sum_i n_i \epsilon_i}{\bar{\rho}} \frac{dz}{H(z)(1+z)}, \quad (2)$$

where $\bar{\rho}$ is the mean mass density and n_i and ϵ_i represent, respectively, the number density and the photoheating rates for the species $i = [\mathrm{H\,I}, \mathrm{He\,I}, \mathrm{He\,II}]$. As shown in N16 (see their Figure 4), this parameter correlates with the density power spectrum of the cosmic gas in the simulations, with larger u_0 corresponding to a smoother distribution of gas for overdensities $\Delta < 10$. These are the overdensities at which the Lyα forest is sensitive at $z > 4$ (e.g., Becker et al. 2011), suggesting that, at these redshifts, u_0 serves as a useful parameterization for the prior IGM thermal history. In Section 6.3.3 we will further consider the redshift range of integration over which u_0 optimally correlates with the flux power spectrum.

Examples of the evolution of u_0 in our models are presented in Figure 1, along with the corresponding evolution of the temperature at the mean density, T_0 (for the complete set of models, see Appendix I). The left panels show how increasing the photoheating rate in the simulations produces larger values in both the temperature and u_0. The right panels show models with the same photoheating rate but different $z_{\rm OT}$. These converge to the same value of T_0 provided that sufficient time has elapsed after the onset of heating ($\Delta z \sim 1$–2; e.g., McQuinn & Upton Sanderbeck 2016); however, they remain distinct in terms of u_0 values, reflecting differences in the total integrated thermal history and therefore in the amount of pressure smoothing.

4. The Lyα Flux Power Spectrum

Both thermal broadening and pressure smoothing tend to reduce the amount of small-scale structure in the forest. Figure 2 shows the effect of thermal broadening (top panel) and pressure smoothing (bottom panel) on simulated Lyα forest spectra at $z = 5$. While the impacts of T_0 and u_0 are visually similar, the scale dependences of these effects make it possible to break the degeneracy (e.g., N16; Oñorbe et al. 2017).

The top row of Figure 3 demonstrates how the shape of the 1D Lyα flux power spectrum at $z = 5$ varies for models with different instantaneous temperature (left panel) and integrated thermal histories (right panel). Similar results were shown in N16 but are expanded here to include a broader range of thermal histories. As described in Section 6.3.1, we use post-processing to vary the T_0 for a fixed u_0 (top left) or to impose the same T_0 for models with different u_0 (top right). We also demonstrate the impact of varying γ (bottom left) and the effective optical depth, $\tau_{\rm eff}$ (bottom right). As noted by N16, the scale dependences of T_0 and u_0 differ somewhat. While the impact of pure thermal broadening increases continuously toward smaller scales, the effect of changing u_0 peaks near $\log(k/\mathrm{km^{-1}\,s}) \sim -0.9$ to -0.8.

Comparing the two panels of the first row of Figure 3, it is clear that in order to distinguish models characterized by an early reionization (large u_0) from those with high T_0 values, it is necessary to probe the power spectrum down to $\log(k/\mathrm{km^{-1}\,s}) \sim -0.7$. Our effort to measure the power spectrum down to these scales is described in the following section.

5. Data Analysis

In this section we describe our procedure for measuring the flux power spectrum from the observed spectra. The following strategies have been tested using synthetic spectra for two reasons: first, to detect and quantify systematic effects in the calculation of the power spectrum, and second, to guarantee a fair comparison between simulated models and the observed data.

5.1. Rolling Mean

We performed the power spectrum measurement on the flux contrast estimator

$$\delta_F = \frac{F - \bar{F}}{\bar{F}}, \quad (3)$$

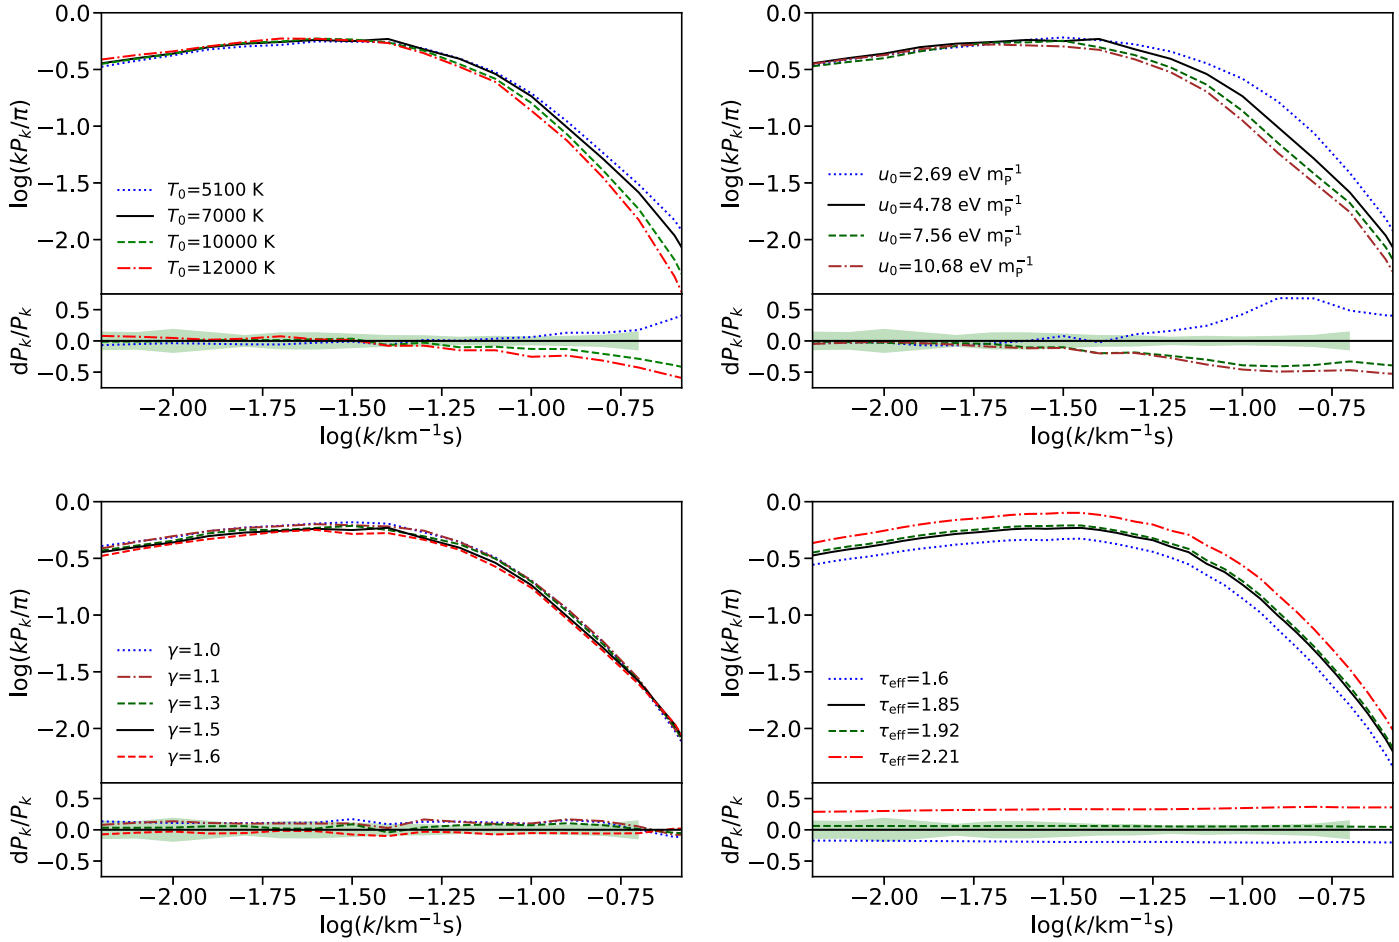

Figure 3. Effects of varying our model parameters on the 1D flux power spectrum at $z = 5.0$. In all panels we plot a fiducial model with $T_0 = 7000$ K, $\gamma = 1.5$, $u_0 = 4.78$ eV m_p^{-1}, and $\tau_{\rm eff} = 1.85$ using a black solid line. The four parameters are varied separately as indicated in each panel. Residuals dP_k/P_k relative to the fiducial value are displayed for each scale. For comparison, the 68% errors relative to the observational power spectrum computed in this work at $z = 5$ are also shown (shaded green region). Models with higher temperature show decreasing power toward smaller scales with the most prominent effect at scales $\log(k/{\rm km^{-1}\ s}) > -1$ (top left). Changes in the integrated thermal history (top right) produce variations in the pressure smoothing experienced by the gas. This effect has a somewhat different scale dependence than pure thermal broadening. The power spectrum at this redshift is not highly sensitive to variations in γ (bottom left), although decreasing γ tends to increase the power at $\log(k/{\rm km^{-1}\ s}) \lesssim -0.8$. Differences in the effective optical depth (i.e., mean flux) create changes in the normalization of the power spectrum (bottom right).

where F is the transmission in the Lyα forest and $\bar{F}$ is the mean flux. When computing δ_F, we need to first divide out the intrinsic shape of the quasar spectrum, which can impact the power spectrum at large scales ($\log(k/{\rm km^{-1}\ s}) \lesssim -2$; e.g., Kim et al. 2004; Viel et al. 2013a; Iršič et al. 2017a). However, directly estimating the continuum is difficult at $z \gtrsim 4$ owing to the high levels of absorption in the Lyα forest. We therefore used a rolling mean approach, wherein $\bar{F}$ is estimated locally by smoothing the observed spectrum using a boxcar average. We used a boxcar window of 40 h^{-1} cMpc, which was chosen to be short enough to roughly capture relevant features in the quasar continua over the forest (see Appendix A.1 for details). Examples of this approach are presented in Figure 4.

5.2. Proximity Regions

Regions near a quasar are subjected to the local influence of its UV radiation field and are therefore expected to show lower Lyα absorption with respect to the cosmic mean. To avoid any proximity effect bias, we excluded these regions from the analysis. The UV flux of a bright quasar is thought to affect regions $\lesssim$10 proper Mpc along its line of sight (e.g., Scott et al. 2000; Worseck & Wisotzki 2006). We conservatively masked 30 proper Mpc blueward of the quasar redshift Lyα emission line. Moreover, to exclude possible blueshifted Lyβ absorption, we also masked a velocity interval corresponding to 10 proper Mpc redward of the Lyβ emission line. Excluding the proximity regions moderately changes the power (by $\gtrsim$5%) only for the highest redshift bin at $z = 5$, although the correction is always well within the statistical error.

5.3. DLAs

We excluded damped Lyα (DLA) systems from our spectra. DLAs were identified visually and masked prior to computing the power spectrum. This step changes the power up to $\sim$5%–10%, which is within the statistical uncertainties at all scales.

5.4. Bad Pixels

We masked bad pixels characterized by negative or zero flux errors. We also masked discrete regions affected by sky

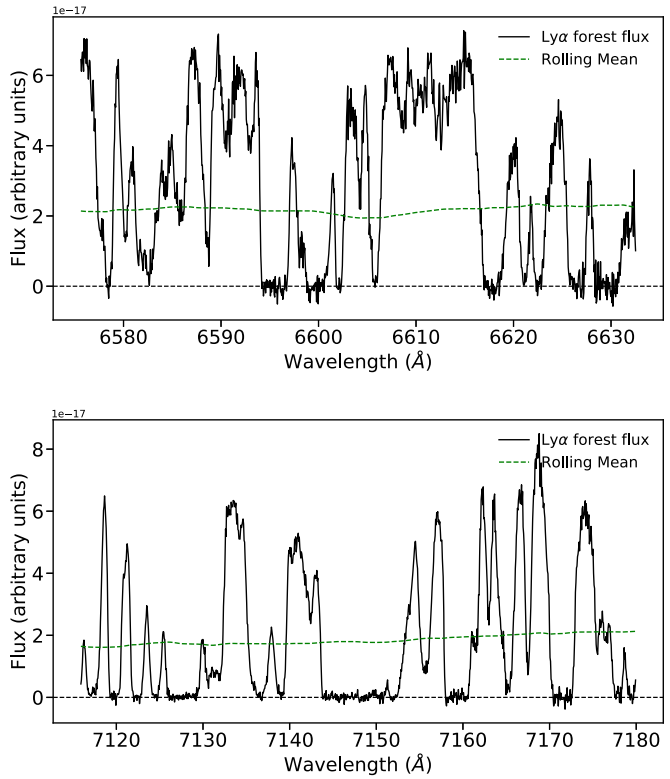

Figure 4. Examples of observed 20 h^{-1} cMpc Lyα forest sections extracted from our sample. Top panel: Lyα forest at $z \simeq 4.4$ extracted from the spectrum of J0741+2520 with a C/N per pixel of ∼17. Bottom panel: Lyα forest at $z \simeq 4.9$ extracted from the spectrum of J0306+1853 with a C/N per pixel of ∼40. The corresponding values of the boxcar rolling mean, measured within a 40 h^{-1} cMpc window, are also shown (green dashed line).

emission line residuals, which tend to be noisy. These features mainly impact smaller scales than the ones we want to compute ($\log(k/\mathrm{km}^{-1}\,\mathrm{s}) \gtrsim -0.5$), but they may affect the evaluation of the correct noise power (see Section 5.6.2) and therefore need to be removed.

5.5. Lyα Sections and Redshift Subsamples

We compute the flux power spectrum on sections of 20 h^{-1} cMpc (comoving distance). This scale was chosen to be small enough that we would have enough subsamples ($N > 30$) to evaluate the statistical uncertainty in the flux power via bootstrapping, yet large enough to avoid significant windowing affects (see Appendix A.2). Each of the Lyα sections has been examined by eye to avoid sections containing too many masked pixels. The power spectrum results from the useful sections are then collected and averaged in redshift bins of $\Delta z = 0.4$ centered at $z = 4.2$, 4.6, and 5.0.

5.6. Measuring the Power Spectrum

For each of the 20 h^{-1} cMpc forest regions we calculate the power spectrum from the flux contrast δ_F defined in Equation (3). Our spectra are unevenly sampled because they are masked, so we use a Lomb–Scargle algorithm (Lomb 1976; Scargle 1982) to compute the raw power of each region ($P_{\mathrm{masked}}(k)$). In all of our calculations we use k-bins logarithmically spaced with $\Delta \log k = 0.1$. To obtain the final power spectrum values, $P_{\mathrm{F}}(k)$, for each section, we first correct the raw $P_{\mathrm{masked}}(k)$ for the effect of masking. Second, we

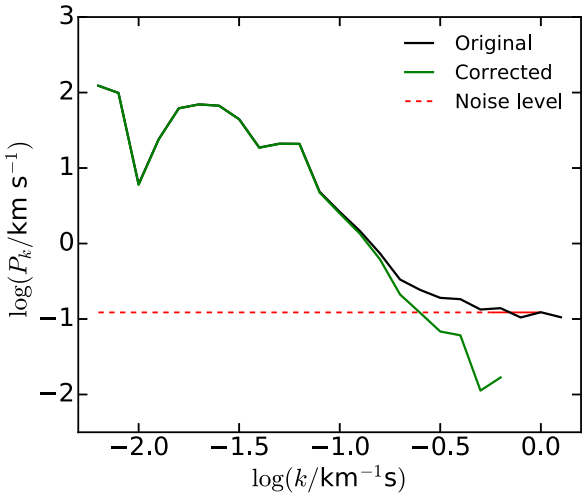

Figure 5. Demonstration of the subtraction of the noise power from the raw power spectrum of an observed 20 h^{-1} cMpc section of Lyα forest. The original power computed directly from the spectra (black solid line) shows a flattening toward the smallest scale becoming roughly constant for $\log(k/\mathrm{km}^{-1}\,\mathrm{s}) \gtrsim -0.2$ when the noise starts to dominate. Assuming white noise, we fit the noise power with a constant value at the smallest scale (red dashed line) and then subtract it from the total power spectrum obtaining the corrected, "noiseless" version (green solid line).

subtract from the corrected $P_{\mathrm{data}}(k)$ an estimate of the contribution to the power from noise, $P_{\mathrm{N}}(k)$. All these steps are described in the following sections.

5.6.1. Masking Correction Function

The masking procedure described in Sections 5.3 and 5.4, and in particular the masking of sky line residuals, impacts the power spectrum owing to the application of a complex window function. In order to correct for this, we apply a masking correction function, $C_m(k)$, to the raw power obtained from each of the 20 h^{-1} cMpc Lyα forest sections,

$$P_{\mathrm{data}}(k) = P_{\mathrm{masked}}(k) \times C_m(k), \tag{4}$$

where $P_{\mathrm{data}}(k)$ is the corrected quantity used to infer the final power and $P_{\mathrm{masked}}(k)$ is the raw power initially computed from masked spectra.

We determine the effect of masking for each of the Lyα forest sections contributing to the analysis using the following procedure. First, we create hundreds of synthetic spectra with the same characteristics (i.e., size, noise, redshift) of each of the real 20 h^{-1} cMpc sections with and without the same masking applied. The final correction is then obtained from the average of the ratio between the power of the unmasked (P_{sim}) and masked ($P_{\mathrm{sim}}^{\mathrm{mask}}$) simulated spectra,

$$C_m(k) = \left\langle \frac{P_{\mathrm{sim}}(k)}{P_{\mathrm{sim}}^{\mathrm{mask}}(k)} \right\rangle. \tag{5}$$

Because the impact of masking on the power spectrum in principle depends on its underlying shape, possible systematics may arise from choosing a particular simulation run to compute $C_m(k)$. We use the 20 h^{-1} cMpc run of Table 2 for the final correction; therefore, we quantify these possible uncertainties in Appendix A.5.

Table 3
List of Quasars Used for the Analysis of the Metal Power Spectrum at $z \sim 4.2$

Name	z_{em}
J010604−254651	3.36500
J162116−004250	3.70270
J132029−052335	3.70000
J124957−015928	3.63680
J014049−083942	3.71290
J115538+053050	3.47520
J014214+002324	3.37140
J123055−113909	3.52800
J110855+120953	3.67160
J005758−264314	3.65500

Note. For each object we report the name (Column (1)) based on the J2000 coordinates of the quasar and the emission redshift (Column (2)). All the spectra have been taken with the UVES spectrograph (see Boera et al. 2014 and Murphy et al. 2019 for details).

5.6.2. Noise Subtraction

In principle, the noise power can be directly computed from the flux error array output by the data reduction pipeline (e.g., Iršič et al. 2017a; Walther et al. 2018). This approach, however, relies on the precision of the pipeline; underestimating or overestimating these uncertainties could significantly impact the final power, especially at the small scales we are interested in. We therefore estimate the amount of noise for each forest section directly from the raw power spectrum of the data. At the smallest scales ($\log(k/\mathrm{km}^{-1}\,\mathrm{s}) \gtrsim -0.2$) the power is dominated by noise fluctuations and, assuming that the noise in adjacent wavelength bins is uncorrelated, can be fitted with a constant value. We then assume that P_N is constant over all scales and subtract it from the total power obtaining the "noiseless" power spectra. This method is illustrated in Figure 5 and has been tested on synthetic data after adding the observational noise arrays to the simulated lines of sight (see Appendix A.3).

5.6.3. Resolution Correction

In this work we forward-model the synthetic spectra generated from simulations to match the instrumental resolution and pixel size of the data. For reference, however, we include a version of the observed flux power spectrum that has been corrected for resolution as

$$P_F(k) = \frac{P_{data}(k) - P_N(k)}{W_R^2(k, R, dv_p)}, \quad (6)$$

using the window function adopted in Palanque-Delabrouille et al. (2015),

$$W_R(k, R, dv_p) = \exp\left(-\frac{1}{2}(kR)^2\right)\frac{\sin(kdv_p/2)}{(kdv_p/2)}. \quad (7)$$

Assuming our nominal resolution $R = 2.55\,\mathrm{km\,s}^{-1}$ (FWHM = $6\,\mathrm{km\,s}^{-1}$) and pixel size $dv_p = 2.5\,\mathrm{km\,s}^{-1}$, the correction for the smallest scale considered in this work ($\log(k/\mathrm{km}^{-1}\,\mathrm{s}) = -0.7$) is $W_R^2 \sim 0.76$.

We note that the actual spectral resolution of the data will depend on the seeing of the observation and may be different from the nominal one. A possible error in the power spectrum due to uncertainties in the spectral resolution, even when forward-modeling the simulations, must therefore be taken into account. We estimate an error of 10% in the spectral resolution, corresponding to an uncertainty in the power of $\lesssim$5% at $\log(k/\mathrm{km}^{-1}\,\mathrm{s}) \lesssim -0.7$. This correction is smaller than our statistical error, so we do not expect that uncertainties in the resolution will significantly affect the measurements (see Appendix A.4).

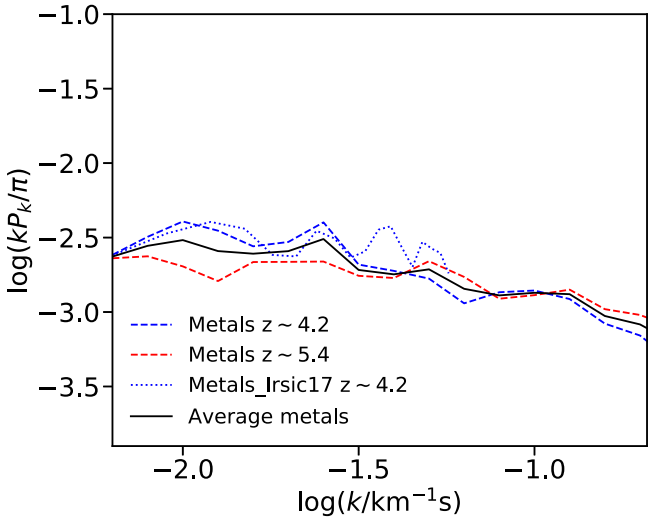

Figure 6. Metal power spectra measured redward of the Lyα emission line. The metal contribution at $z \sim 4.2$ (blue dashed line) computed from the quasars with emission redshifts $3.4 \lesssim z_{em} \lesssim 4.1$ of Table 3 is compared with the metal power measured at $z \sim 5.4$ (red dashed line) from the quasar subsample with $4.5 \lesssim z_{em} \lesssim 5.3$ of Table 1. The average is given by the black solid line. For comparison, the metal power spectrum computed from the XQ-100 data sample at $z \sim 4.2$ (Iršič et al. 2017a) is also shown (blue dotted line).

5.7. Metals

The flux power spectrum measured directly from the observational spectra contains both the power coming from the Lyα forest and a small contribution from intervening metal lines. These lines tend to show individual components significantly narrower than Lyα ($b \lesssim 15\,\mathrm{km\,s}^{-1}$), which will increase the power on small scales (e.g., Lidz et al. 2010). This section describes our approach to quantifying and removing the effect of metals on our final power spectrum measurements.

The high level of Lyα absorption at high redshift makes it very challenging to directly identify all metal lines in the forest. We therefore estimate the metal power spectrum directly from regions of quasar spectra redward of the Lyα emission line, where only metal absorption systems are present (e.g., McDonald et al. 2005; Palanque-Delabrouille et al. 2013). The metal power measured in this way will not take into account transitions with rest-frame wavelength shorter than the Lyα line. Correlation features like the one observed for Si III (λ1206) in McDonald et al. (2006), however, will tend to affect the power spectrum on scales larger than the ones considered in this work ($\log(k/\mathrm{km}^{-1}\,\mathrm{s}) \lesssim -2.5$).

We measured the metal power spectrum from two samples of high-resolution quasar spectra. First, we use a subset of the spectra listed in Table 1 with emission redshift $4.5 \lesssim z_{em} \lesssim 5.3$. Second, we use a sample of spectra of quasars with emissions redshifts $3.4 \lesssim z_{em} \lesssim 4.1$ from Boera et al. (2014, 2016) (Table 3). The latter sample allows us to measure

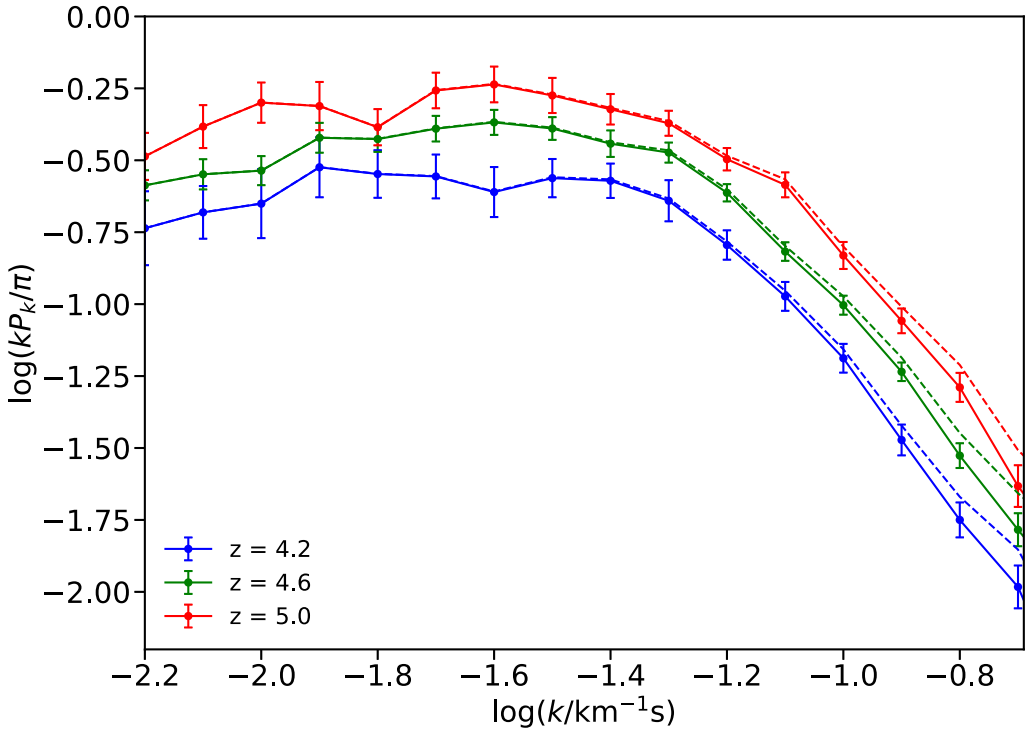

Figure 7. Our measurements of the Lyα flux power spectrum for the $\Delta z = 0.4$ bins centered at $z = 4.2$ (blue), 4.6 (green), and 5.0 (red). Higher effective optical depths determine an increase in the power toward higher redshifts. Values are obtained following the steps presented in Section 5 with (dashed lines) and without (solid line and data points) instrumental resolution correction. Vertical error bars are 1σ errors taken from the corrected and regularized covariance matrix (Section 5.9). Note that the larger number of spectra contributing to the $z = 4.6$ bin (roughly double the number of spectra of the other two bins) are reflected in the significantly smaller errors.

the metal power spectrum over observed wavelengths similar to those spanned by the Lyα forest at $4.0 \lesssim z \lesssim 4.4$. While metals redward of Lyα are not a perfect estimate of those that appear in the forest at higher redshifts, analyzing multiple samples allows us to check for redshift evolution in the metal power spectrum.

Figure 6 shows the comparison between our measurements of the metal power spectrum at $z \sim 4.2$ (blue dashed line) and at $z \sim 5.4$ (red dashed line), obtained using the same data analysis procedure described in the previous sections. The most significant difference in the metal power between these two redshifts is at large scales ($\log(k/\mathrm{km}^{-1}\,\mathrm{s}) \lesssim -1.5$), where the contribution of metals to the final flux power spectrum measurement is less relevant. Note that our measurement of the metal power spectrum at $z \sim 4.2$ is also in good agreement with the one computed from the XQ-100 data sample at the same redshift (Iršič et al. 2017a) (blue dotted line) even if the latter is slightly noisier.

Given the weak evolution in redshift of the metal power (already noted in previous works; e.g., Palanque-Delabrouille et al. 2013), we corrected our final power spectrum measurements assuming a metal contribution constant with redshift and equal to the average between our two measured metal power spectra (black solid line in Figure 6). The effect of the metal correction on the final flux power spectrum is shown in Appendix A.6.

5.8. The New Power Spectrum Measurements

The main observational results of this work are presented in Figure 7, where we plot the final Lyα flux power spectrum measured from our data. The values are tabulated in Appendix K and reported as a function of scale for the three redshift bins centered at $z = 4.2$, 4.6, and 5.0. Solid colored lines and data points represent the power spectrum results without the resolution correction described in Section 5.6.3, while the corresponding dashed lines are the measurements corrected for finite resolution and pixel size. The 1σ errors are estimated from the bootstrap covariance matrix of the data, corrected and regularized following the procedure described in Section 5.9.

Our measurements at scales $\log(k/\mathrm{km}^{-1}\,\mathrm{s}) \gtrsim -1$ are the first ones made at these redshifts (see Walther et al. 2018 for an analysis at $z < 4$). At larger scales, however, we can compare with the results derived from high-resolution spectra by Viel et al. (2013a) and medium-resolution data by Iršič et al. (2017a). These comparisons are presented in Appendices C and D.

5.9. Covariance Matrix

As demonstrated in previous works (e.g., Viel et al. 2013a; Iršič et al. 2017a), the covariance matrix obtained via bootstrapping of a limited data set is necessarily noisy. We therefore regularized the observed covariance matrix using the correlation coefficients estimated from the simulated spectra following an approach similar to the one used by Lidz et al. (2006). We first used the simulations to verify the ability of the bootstrapped errors to reproduce the real statistical variance. For this test, using the $40\,h^{-1}$ cMpc box simulation S40_1z15, we created hundreds of samples of simulated lines of sight that closely reproduce the characteristics of the observational data (see Section 6.2 for details), and we compare the variance computed directly from these realizations with the uncertainty obtained from the bootstrapping of only one synthetic sample randomly chosen. We verify that, as already shown by previous

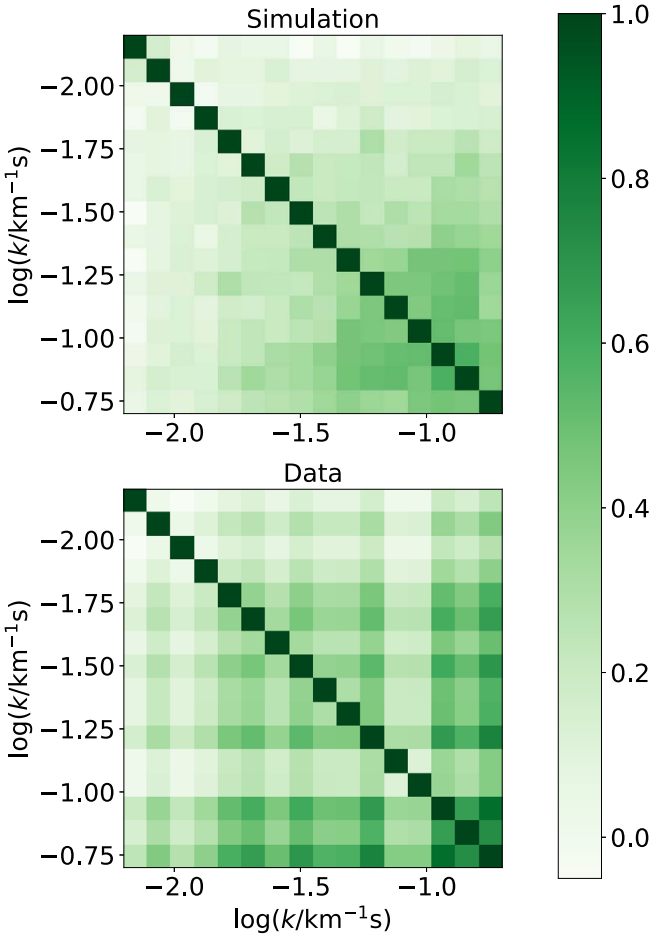

Figure 8. Example of the Lyα flux power spectrum covariance matrix computed by directly bootstrapping the data (bottom panel) and from many realizations of our 40 h^{-1} cMpc simulation model (top panel) for the redshift bin at $z = 5$. The bootstrapped matrix is visibly noisier than the simulated one owing to the limited sample size; however, the correlation structures are similar between the two panels, with stronger correlations toward smaller scales.

studies (e.g., Kim et al. 2004; Iršič et al. 2013; Palanque-Delabrouille et al. 2013), the bootstrapping technique underestimates the cosmic variance by up to ∼25%, with the discrepancy level increasing toward smaller scales. We therefore increased the observational bootstrapped error at all scales by ∼15%–25%, where the correction has been computed separately for each redshift.

The elements of the final covariance matrix, C_{ij}, are then computed as

$$C_{ij} = R_{ij}\sqrt{C_{ii}^{\rm data} C_{jj}^{\rm data}}, \tag{8}$$

with

$$R_{ij} = \frac{C_{ij}^{\rm sim}}{\sqrt{C_{ii}^{\rm sim} C_{jj}^{\rm sim}}}, \tag{9}$$

where $C_{ii}^{\rm data}$ are the diagonal elements of the bootstrapped observational covariance matrix, corrected as previously described, and $C^{\rm sim}$ are the elements of the simulated covariance matrix obtained from the multiple realizations of synthetic lines of sight.

For comparison, Figure 8 shows the simulated and observational covariance matrices for the redshift bin at $z = 5$. As expected, the bootstrapped matrix is noisier, but both the matrices show a similar structure, with correlations increasing toward the smallest scales, $\log(k/{\rm km}^{-1}\,{\rm s}) \gtrsim -1.5$. This similarity gives us confidence that the simulation model used for the covariance matrix regularization is reasonably capturing the data properties. We note that the off-diagonal correlation structure will depend somewhat on the precise shape of the power spectrum and therefore on the thermal parameters characterizing the model. In Appendix B we verify that the particular choice of the simulation S40_1z15 for this analysis is not significantly affecting our final constraints.

6. Simulation Analysis

In this section we describe how we calibrate and analyze synthetic spectra to create power spectrum models that will be used to fit the observational measurements in Section 7.

6.1. Constructing Mock Lines of Sight

To ensure the correct comparison between simulations and observational data, we need to produce mock lines of sight with the same resolution and redshift coverage as our observed sample. We first resample and smooth the synthetic Lyα spectra produced from the simulations in Table 2 to match the spectral resolution and the pixel size of the real spectra. We then progressively merge multiple synthetic sections randomly selected from the $\Delta z = 0.1$ simulation snapshot closest to the Lyα forest redshift that we want to cover. We choose an arbitrary starting point along each section, taking advantage of the periodicity of the simulation box. We take into account the mild redshift evolution of the mean flux along the line of sight by rescaling the optical depths such that the global effective Lyα optical depth ($\tau_{\rm eff} = -\ln(\bar{F})$) in the simulation box follows the relation

$$\tau_{\rm eff} = 1.56\left(\frac{1+z}{5.75}\right)^{4.0}. \tag{10}$$

See Appendix G for details. We note that while we use Equation (10) to calibrate the optical depth evolution within each simulated line of sight, the overall Lyα mean flux of each redshift bin will be treated as a free parameter in our models.

When required for testing purposes (see Appendix A), the spectral noise at the same level of the corresponding observational line of sight is added to the synthetic spectra, as well as the same pixel masking. Because the power spectrum computed from the real data is corrected for these systematics, the final power spectrum models used for the MCMC fit are computed without noise or pixel masking.

6.2. Modeling the Power Spectrum

To retrieve the final flux power spectrum for each of the models in Table 2, we average the power computed from hundreds of mock data samples. Each measurement has been obtained following a similar procedure to the one described in Section 5, with a few necessary expedients:

1. Lyα sections: As for the real data, we compute the flux contrast estimator (Equation (3)) using a 40 h^{-1} cMpc boxcar rolling mean over the reconstructed line of sight. After the rolling mean is applied, we redivide the line of sight into the original 10 h^{-1} cMpc sections and use

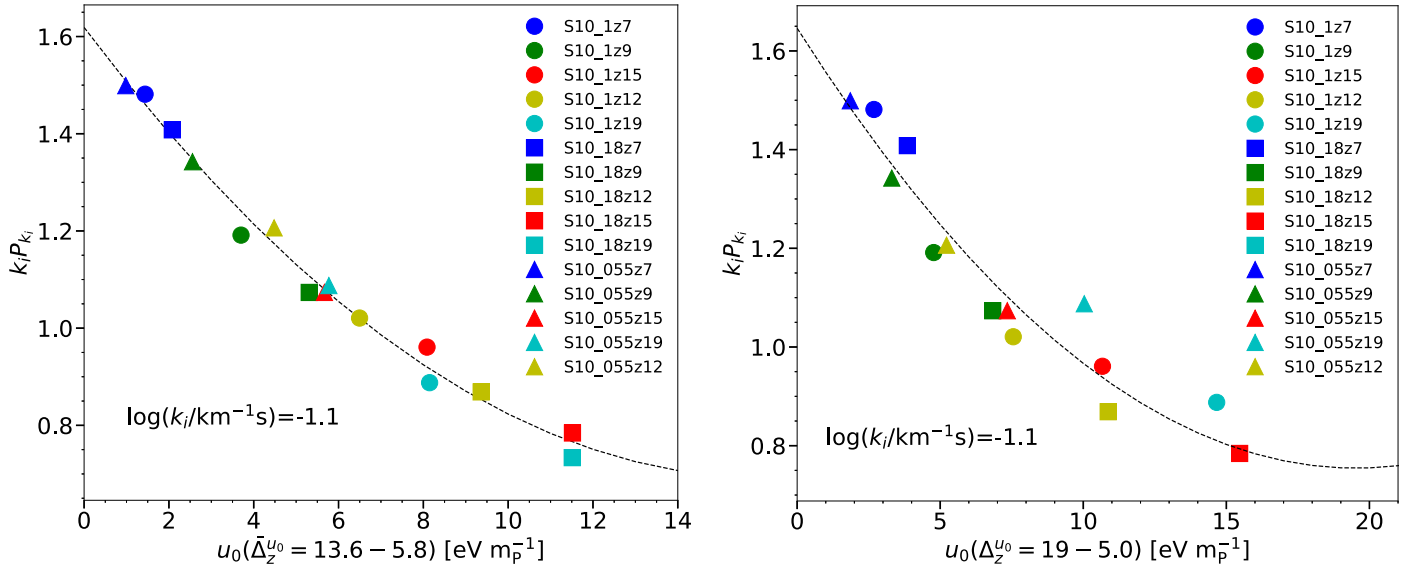

Figure 9. Examples of the relationship between the power spectrum at $z = 5.0$ computed for $\log(k/\mathrm{km}^{-1}\,\mathrm{s}) = -1.1$ and u_0 obtained integrating over the characteristic redshift range $\Delta\bar{z}^{u_0} = 13.6 - 5.8$ (left panel) and integrating over a nonoptimal redshift range $\Delta z^{u_0} = 19 - 5.0$ (right panel). Different colors correspond to different simulations post-processed to the same values of $T_0 = 10{,}000$ K, $\gamma = 1.5$, and $\tau_{\rm eff} = 1.85$.

these to compute the power spectra. We verified that discontinuities in the flux on the border between individual sections do not substantially affect the rolling mean.

2. Mass resolution and box size corrections: Our 10 h^{-1} cMpc models with 2×512^3 gas and dark matter particles represent a necessary compromise in terms of computational resources (Bolton & Becker 2009) and need to be corrected for small errors due to box size and resolution convergence. Therefore, we rescaled our models by factors obtained from reference simulations with larger box size (40 h^{-1} cMpc box) and higher mass resolution (2×768^3 particles) in the convergence tests presented in Appendix F.

6.3. Varying Model Parameters

To be able to fit the power spectrum measurement of Section 5.8, we need a grid of models that cover the parameter space that we want to explore. In each redshift bin we consider four parameters to describe the power spectrum: the thermal parameters T_0, u_0, and γ and the effective Lyα optical depth, $\tau_{\rm eff}$. While the large set of simulations listed in Table 2 spans a wide range of thermal histories, by themselves they are not sufficient to evaluate the power spectrum in all the possible combinations of thermal parameters. We therefore use the interpolation scheme described below.

6.3.1. Varying T_0 and γ

In order to separate the impact of thermal broadening and Jeans smoothing in the power spectrum models, we applied a simple post-processing procedure to the simulated spectra. This is achieved by translating and rotating the entire T–ρ plane of the simulations to match the new T_0 and γ values. We recompute the optical depths in each of our models over an extended range of power-law T–ρ relationships, with $T_0 = 3000$–$15{,}000$ K in steps of 1000 K and $\gamma = 0.7$–1.7 in steps of 0.1.

We note that at $z \gtrsim 4$ the Lyα forest is mainly sensitive to gas close to the mean density (e.g., Becker et al. 2011). For this reason we do not expect to place strong constraints on γ. We nevertheless treat γ as a free parameter in our fitting code. The impacts of T_0 and γ on the flux power spectrum are demonstrated in Figure 3. Note that scales $\log(k/\mathrm{km}^{-1}\,\mathrm{s}) \gtrsim -0.8$ seem to be insensitive to variations in γ, while considerable changes in this parameter create minor shifts in the power for scales $\log(k/\mathrm{km}^{-1}\,\mathrm{s}) \lesssim -0.8$.

6.3.2. Varying τ_{eff}

We rescaled the optical depths in our models to span a wide range of $\tau_{\rm eff}$ values. At each redshift the reference value has been obtained from Equation (10), while the entire range of optical depths covered by our models is $\tau_{\rm eff} = 0.6$–2.2 in steps of 0.1. Different mean fluxes within each redshift bin have been obtained by multiplying Equation (10) by single, fine-tuned scalars when calibrating the simulated lines of sight. The impact of varying $\tau_{\rm eff}$ on the power spectrum is shown in the bottom right panel of Figure 3.

6.3.3. Varying u_0

By post-processing our simulations to a common set of thermal parameters, we can isolate how the power spectrum depends on the integrated heating. N16 demonstrated that the flux power spectrum at $z = 5.0$ (averaged over scales $-1.5 \lesssim \log(k/\mathrm{km}^{-1}\,\mathrm{s}) \lesssim -0.8$) correlates with u_0. They further argued that the correlation is strongest when u_0 is integrated between $z = 12$ and 5, reflecting the timescales over which the Jeans smoothing is sensitive to heat injection. Here we reevaluate this redshift dependence using our more extended suite of models. For each of the redshift bins at which we compute the power spectrum and each of the scales sensitive to u_0 ($-1.4 \lesssim \log(k/\mathrm{km}^{-1}\,\mathrm{s}) \lesssim -0.8$) we empirically determine the "characteristic" redshift range of integration ($\Delta\bar{z}^{u_0}$) for which the power is closest to a one-to-one function of u_0. The method is demonstrated in Figure 9. We first post-process all of the 10 h^{-1} cMpc simulations of Table 2 to the

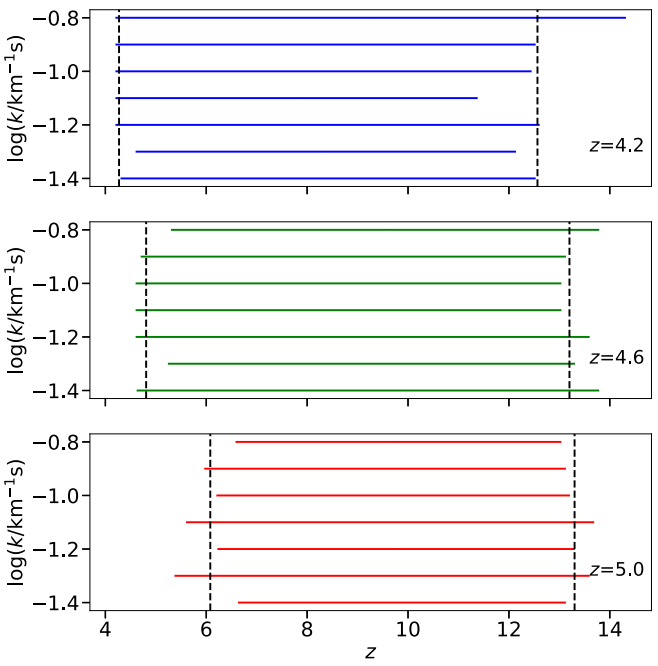

Figure 10. Characteristic $\Delta\bar{z}^{u_0}$ computed for $-1.4 \lesssim \log(k/\mathrm{km}^{-1}\,\mathrm{s}) \lesssim -0.8$. For each $\log(k/\mathrm{km}^{-1}\,\mathrm{s})$ on the y-axis, the colored horizontal line covers the redshift interval that produces the best fit between the power spectrum and the integrated heating. Each panel shows the results for a different redshift bin. While there is a mild evolution with redshift, the $\Delta\bar{z}^{u_0}$ values are generally constant within the same redshift bin. We use the average $\Delta\bar{z}^{u_0}$ indicated by the black dashed vertical lines to compute the fiducial u_0 parameters at each redshift.

same values of T_0, γ, and $\tau_{\rm eff}$. We then fit a power law to kP_k versus u_0, where u_0 is integrated using Equation (2) over a redshift interval Δz^{u_0}. The preferred interval, $\Delta\bar{z}^{u_0}$, is the one that minimizes the χ^2 for this fit.

The characteristic $\Delta\bar{z}^{u_0}$ computed for the scales $-1.4 \lesssim \log(k/\mathrm{km}^{-1}\,\mathrm{s}) \lesssim -0.8$ are reported in Figure 10 for the different redshift bins. As expected, the sensitivity of the power spectrum to the previous thermal history varies slightly with the redshift at which the power spectrum is measured. The Lyα structures observed at progressively lower redshifts seem to slowly lose sensitivity to earlier epochs; while the power spectrum measured at $z = 5.0$ still maintains sensitivity up to $z \gtrsim 13$, at $z = 4.2$ the forest traces the thermal history of the gas mainly for $z \lesssim 12$. Interestingly, we find that the power spectrum at $z = 5.0$ is less sensitive to heating happening at $z \lesssim 6$, even though the power spectra at $z = 4.6$ and 4.2 retain sensitivity all the way down to their respective redshifts. We generally expect that the gas density distribution will exhibit some delay in responding to changes in gas pressure. Further investigation revealed that a delay did appear for all three redshifts when peculiar velocities were turned off. The delay increased with the redshift at which the power spectrum was measured, with a delay at $z = 5.0$ that was larger than the one found with peculiar velocities turned on. This suggests that peculiar velocities may play a role by decreasing the delay between heat injection and a change in the power spectrum. Presumably this occurs because, as the gas is heated, redshift distortions created by accelerating the gas precede changes in the density field. This effect may partly explain the lack of a gap at $z = 4.2$ and 4.6. For now we adopt these relations as empirical and leave more detailed physical insights to future work.

Table 4
Fiducial Redshift Ranges Used to Compute the u_0 Parameters for Fitting the Flux Power Spectrum at Different Redshifts

z	$\Delta z_{\rm fid}^{u_0}$
4.2	4.2–12
4.6	4.6–13
5.0	6.0–13

Note. Redshift bins are indicated in Column (1), with the corresponding fiducial u_0 redshift intervals in Column (2).

Because the $\Delta\bar{z}^{u_0}$ values are generally constant among different scales within the same redshift bin, we adopt average values (black dashed vertical lines in Figure 10) as integration bounds in Equation (2). The fiducial redshift range over which u_0 is integrated is given in Table 4. We note that our u_0–kP_k relationship, while remarkably tight over scales sensitive to u_0, does exhibit scatter. In the final MCMC analysis, therefore, the amount of scatter about the u_0–kP_k fit at each scale has been included as systematic uncertainty. We note that while we chose our fiducial redshift ranges to maximize the sensitivity of the power spectrum to u_0, in principle we could constrain this parameter integrated within any reasonable redshift range if properly accounting for the systematic uncertainty in the u_0 versus kP_k fit.

7. Thermal State Constraints

To obtain constraints on the IGM temperature and integrated thermal history from the observational power spectrum measurements obtained in Section 5, we adopted a Bayesian MCMC approach to measure T_0, u_0, γ, and $\tau_{\rm eff}$ for each of the three redshift bins independently. In this section we present the method and the main findings of this analysis.

7.1. The MCMC Analysis

We constructed a grid of power spectrum models following the post-processing approach given above, where for a given choice of T_0, γ, and $\tau_{\rm eff}$ the dependence of the power spectrum on u_0 is derived from the fits described in Section 6.3.3. We then perform a multilinear interpolation among the grid points of the four-dimensional parameter space. We implemented the interpolation scheme using a Bayesian MCMC approach. At each redshift, applying flat priors for all variables, we obtain the set of parameters that maximize a Gaussian multivariate likelihood function:

$$\ln \mathcal{L} = -\frac{1}{2}\mathbf{\Delta}^T C^{-1}\mathbf{\Delta} - \frac{1}{2}\ln\det(C) - \frac{N}{2}\ln 2\pi, \qquad (11)$$

where $\mathbf{\Delta}$ is the residual vector between the power spectrum values of the data and the model and C is the $N \times N$ data covariance matrix (where N is the number of data points).

We tested the interpolation scheme by removing one of the models used for the interpolation and using it to generate mock data (see Appendix H). We found that the parameters u_0 and T_0 are recovered accurately. Small biases (within the 68% uncertainties) appear in the recovered values of γ and $\tau_{\rm eff}$ owing to their intrinsic degeneracy at large scales and the poor sensitivity of the high-redshift power spectrum to γ.

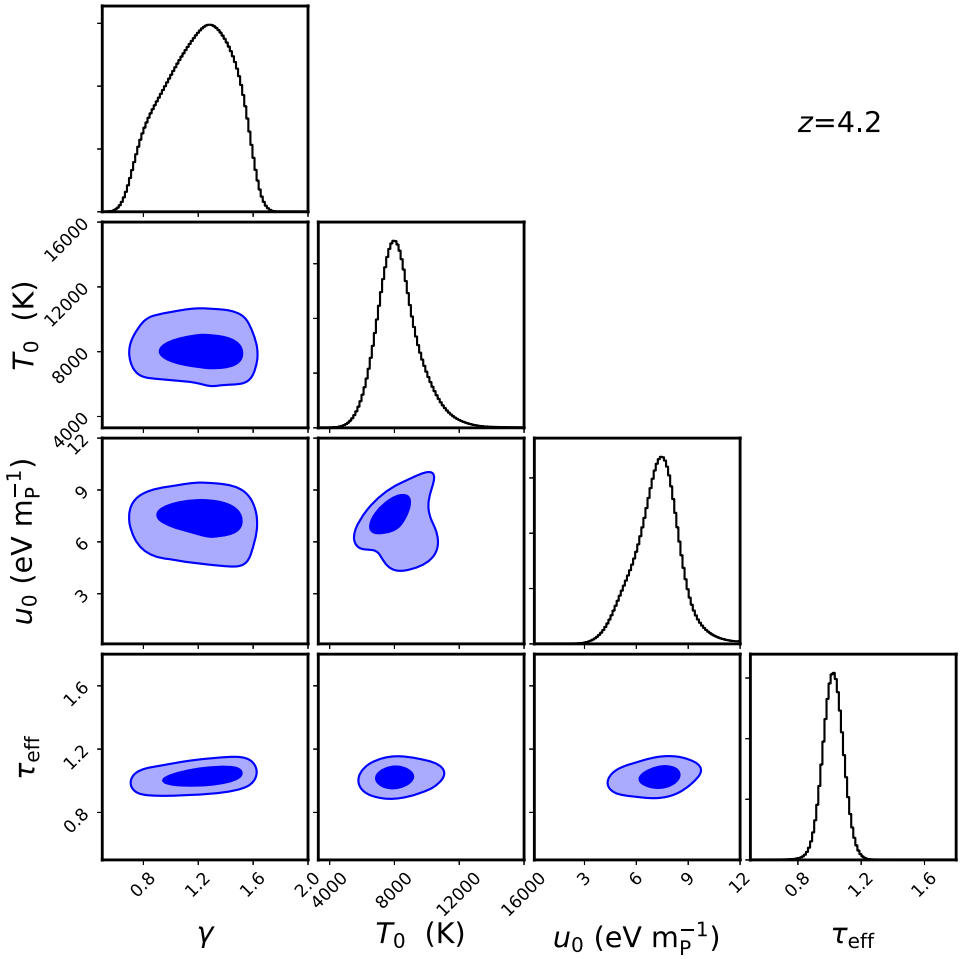

Figure 11. Probability distributions for the parameters T_0, u_0, γ, and $\tau_{\rm eff}$ obtained from the MCMC analysis of the power spectrum at $z = 4.2$. Contours show the 68% and 95% marginalized 2D probability distributions, while the black histograms display the 1D marginalized posterior distributions for each parameter.

Fortunately, however, the relatively weak constraints on these parameters do not bias our results for T_0 and u_0.

To test the reliability of the best-fitting values, for each redshift we ran three independent chains of 2×10^5 iterations (half of which are discarded as burn-in) from different randomly chosen initial parameters. We verify that all the chains were converged by comparing the between-chain and within-chain variances for each parameter using the Gelman–Rubin test.

7.2. Results

Figures 11–13 display the posterior likelihood distributions for the parameters T_0, u_0, γ, and $\tau_{\rm eff}$ at redshifts 4.2, 4.6, and 5.0, respectively. While the inclusion of small scales ($\log(k/{\rm km}^{-1}\,{\rm s}) \gtrsim -1.0$) in the power spectrum allows relatively tight constraints on both T_0 and u_0, some degeneracy between these two variables is still noticeable at all redshifts. This is expected since both of these parameters act on intermediate scales in a similar way. Degeneracies between γ and $\tau_{\rm eff}$ increase with redshift, with slightly weaker constraints on $\tau_{\rm eff}$ obtained toward higher redshifts. As expected, γ shows broad bounds at all redshift (with the 1σ contours covering almost the entire parameter space), reaffirming that the Lyα forest at high redshifts mainly probes gas around the mean density and is not highly sensitive to the slope of the T–ρ relation. Figures 11–13 demonstrate that our measurements of u_0 and T_0 are not highly affected by degeneracies with γ and $\tau_{\rm eff}$.

The final results of the MCMC analysis are summarized in Table 5. The temperatures are constrained with ∼15% uncertainties at all redshifts, while the error on u_0 varies from ∼18% for the $z = 4.2$ and $z = 4.6$ redshift bins up to ∼30% at the highest redshift, in good agreement with the forecast presented by N16. Our results for $\tau_{\rm eff}$ are highly consistent with the measurements of Becker et al. (2013) at $z = 4.2$. We are somewhat higher at $z = 4.6$, but altogether our constraints appear to bridge the evolution of $\tau_{\rm eff}$ at $z \lesssim 4$ measured by Becker et al. (2013) and at $z \gtrsim 5$ from Bosman et al. (2018) (see Appendix G).

In Figure 14 we show the best power spectrum models compared with the measurement. Visually there is good agreement with the data at all redshifts.

The final temperature constraints at the IGM mean density are presented in Figure 15. Our new results (green points) are compared with the previous measurements of Becker et al. (2011) (gray triangles) at $z > 3.5$ obtained with the curvature method. We have added the systematic uncertainty for Jeans smoothing estimated by Becker et al. to those data. See also Appendix E for a more comprehensive comparison of recent temperature measurements at the redshifts covered by our analysis. Our temperature measurements show good agreements with this previous work in the overlapping redshift bins.

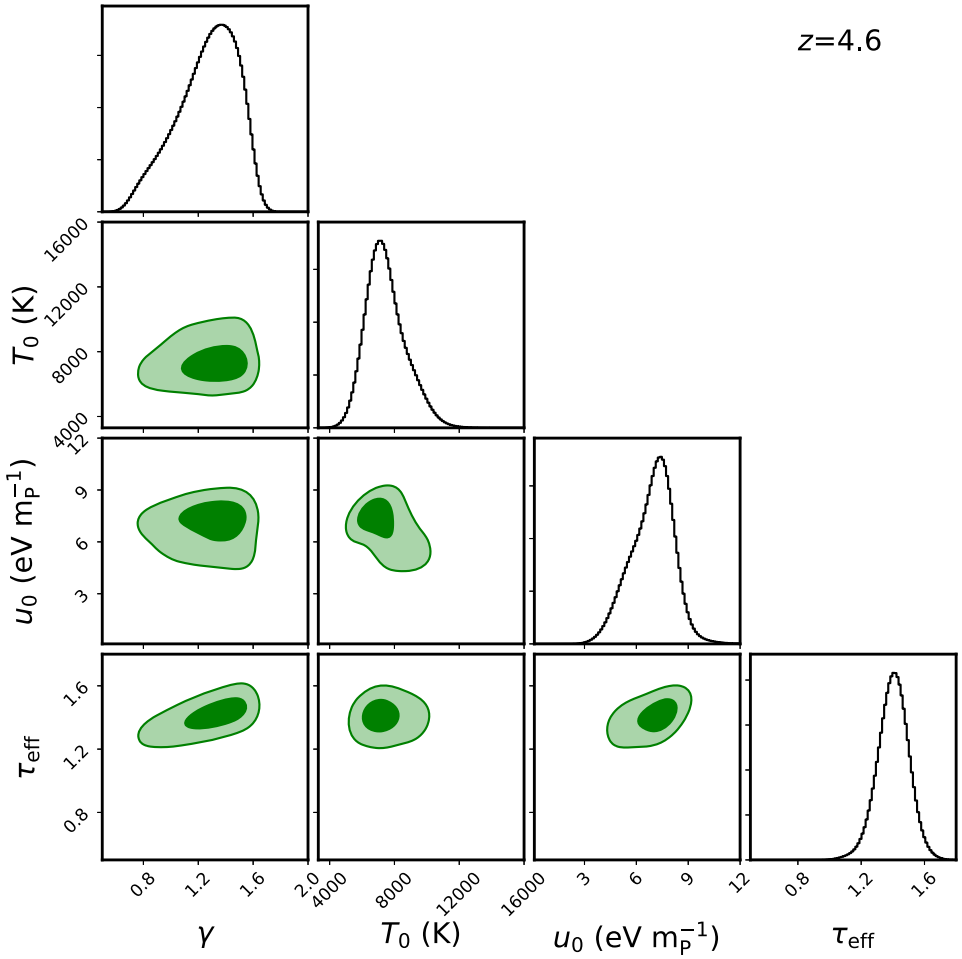

Figure 12. Same as Figure 11, but for the power spectrum at $z = 4.6$.

This accord is significant because we analyzed a largely independent set of quasar spectra with a different method and using a new suite of hydrodynamical simulations. Most significantly, we now explicitly fit for u_0, removing the systematic uncertainty in T_0 related to Jeans smoothing. Overall our T_0 values are consistent with little evolution over $4.2 \lesssim z \lesssim 5.0$. Given the known trend of increasing temperatures at $z \lesssim 4$ (e.g., Becker et al. 2011; Boera et al. 2014; Walther et al. 2019), our measurement at $z = 4.2$ may include a contribution from the initial phase of IGM reheating due to the He II reionization (e.g., Worseck et al. 2011; Syphers & Shull 2014). We consider this possibility in the final part of our analysis, when we use the new thermal constraints to evaluate hydrogen reionization scenarios.

Finally, in Figure 16 we present our first constraints on the integrated thermal history of the IGM. Our u_0 measurements are plotted at the minimum redshift of the fiducial ranges given in Table 5. As expected, u_0 increases from $z = 6$ to $z = 4.2$, reflecting ongoing heat injection after reionization.

8. Reionization Constraints

Our observational constraints on T_0 and u_0 can, in principle, be used to test any reionization model for which a thermal history can be calculated. While we leave for future work the analysis of extended and more realistic reionization scenarios, in this section we demonstrate the potential of this approach using semianalytical models of instantaneous reionization.

8.1. Modeling Instantaneous Reionization

We model the thermal history of instantaneous hydrogen reionization using a semianalytical approach similar to the one adopted in Upton Sanderbeck et al. (2016). To obtain T_0 as a function of redshift, we solve the equation describing the temperature evolution of a Lagrangian fluid element at the cosmic mean density, i.e., with $\Delta = 1$ (e.g., Miralda-Escudé & Rees 1994; Hui & Gnedin 1997; McQuinn & Upton Sanderbeck 2016),

$$\frac{dT}{dt} = -2HT + \frac{2T}{3\Delta}\frac{d\Delta}{dt} + \frac{2}{3k_B n_{\rm tot}}\frac{dQ}{dt}, \qquad (12)$$

where H is the Hubble parameter and $n_{\rm tot}$ is the total number density of particles (electrons and ions). Equation (12) is valid in the approximation that the number of particles remains fixed, describing well the post-reionization gas. The first term on the right-hand side of Equation (12) takes into account the cooling due to adiabatic expansion, while the second term gives the adiabatic heating and cooling due to structure formation. The thermal history at the mean density has been shown to depend weakly on this second term (McQuinn & Upton Sanderbeck 2016); we will therefore ignore it in our calculation. The

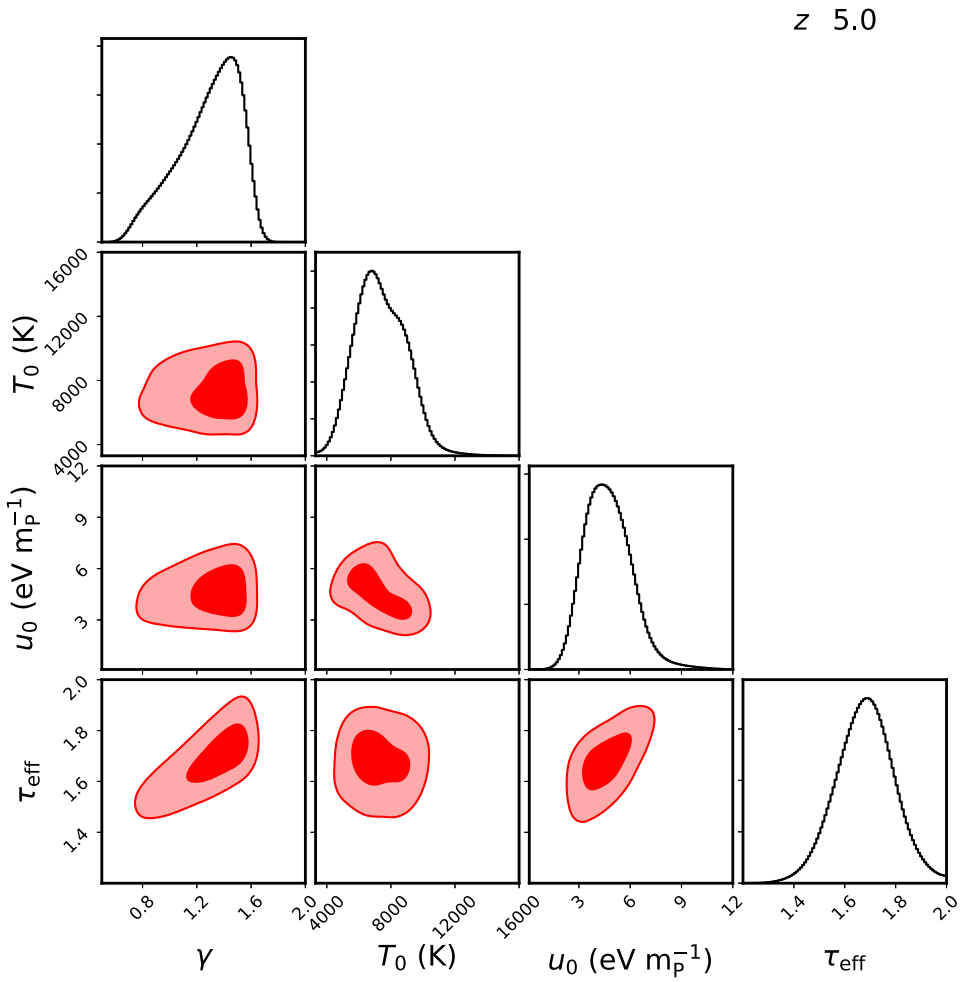

Figure 13. Same as Figure 11, but for the power spectrum at $z = 5.0$.

differences among the models will depend instead on the third term, which encodes photoheating (the only heating source considered in these calculations) and radiative cooling processes. We can expand this term as

$$\frac{dQ}{dt} = \sum_X \frac{dQ_{\mathrm{photo},X}}{dt} + \frac{dQ_{\mathrm{Compton}}}{dt} + \sum_i \sum_X R_{i,X} n_e n_X, \quad (13)$$

where $dQ_{\mathrm{photo},X}/dt$ is the photoheating rate of ion X, $dQ_{\mathrm{Compton}}dt$ is the Compton cooling rate, and $R_{i,X}$ is the cooling rate coefficient for the ion X and cooling mechanism i. Because we are modeling the temperature of the gas at the end of hydrogen reionization, in Equation (13) we will consider only the H I (dominant) and He I photoheating. As for the cooling term, we include Compton and H II recombination in our calculations. As discussed in McQuinn & Upton Sanderbeck (2016), these represent the relevant cooling processes that shape the temperature evolution. We compute these cooling terms using the rate coefficients provided by Hui & Gnedin (1997). The optically thin photoheating after reionization is modeled as (e.g., Upton Sanderbeck et al. 2016)

$$\frac{dQ_{\mathrm{photo},X}}{dt} \approx \frac{h\nu_X}{\gamma_X - 1 + \alpha_{bk}} \alpha_{A,X} n_{\tilde{X}} n_e, \quad (14)$$

where ν_X is the frequency associated with the ionization potential of species X, and γ_X is the corresponding approximate power-law index of the photoionization cross section, for which we assume $\gamma_X = 2.8$ for H I and $\gamma_X = 1.7$ for He I. Equation (14) is valid in the approximation of photoionization equilibrium with an ionizing background that has a power-law specific intensity of the form $J_\nu \propto \nu^{-\alpha_{\mathrm{bk}}}$. In detail, the photoheating rate will also depend on $\alpha_{A,X}$, the case A recombination coefficient associated with the transition from $X\,\mathrm{I} \rightarrow X$ for species $X \in$ [H I, He I]; on the number density of the species $\tilde{X} \in$ [H, He]; and on the electron number density n_e.

To compute the total energy deposited into the gas by photoheating, u_0, we just need to consider the third term of Equation (12) (and the first term of Equation (13)). The

Table 5
Best-fitting Values and Marginalized 68% Confidence Intervals for the Fits to Our Power Spectrum Measurements

z	$T_0/10^3$ (K)	u_0(eV m_p^{-1})	γ	τ_{eff}
4.2	$8.13^{+1.34}_{-0.97}$	$7.29^{+0.98}_{-1.35}$	$1.21^{+0.23}_{-0.28}$	$1.02^{+0.04}_{-0.04}$
4.6	$7.31^{+1.35}_{-0.88}$	$7.10^{+0.83}_{-1.45}$	$1.29^{+0.19}_{-0.26}$	$1.41^{+0.08}_{-0.09}$
5.0	$7.37^{+1.67}_{-1.39}$	$4.57^{+1.37}_{-1.16}$	$1.33^{+0.18}_{-0.27}$	$1.69^{+0.10}_{-0.11}$

Note. The power spectrum redshift (Column (1)) is reported along with the best-fitting values of T_0 (Column (2)), u_0 (Column (3)), γ (Column (4)), and τ_{eff} (Column (5)).

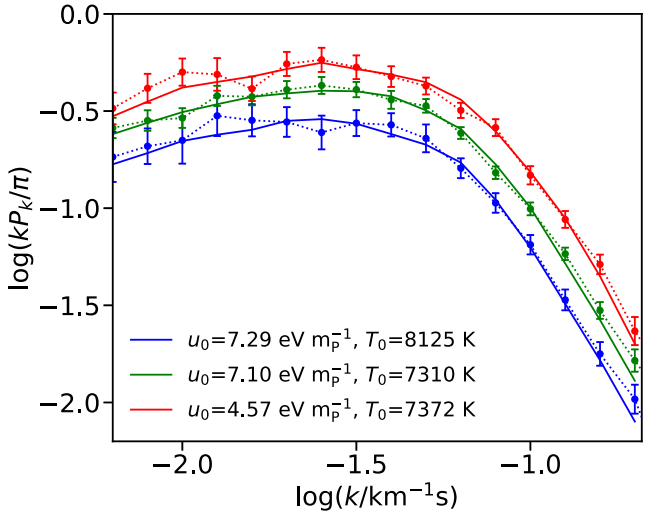

Figure 14. Best-fitting models for our high-resolution power spectrum measurements. The best-fit models at $z = 5.0$ (red solid line), $z = 4.6$ (green solid line), and $z = 4.2$ (blue solid line) are superimposed on the corresponding observational measurement (color-coded data points and dotted lines). The corresponding best-fitting parameters are also reported.

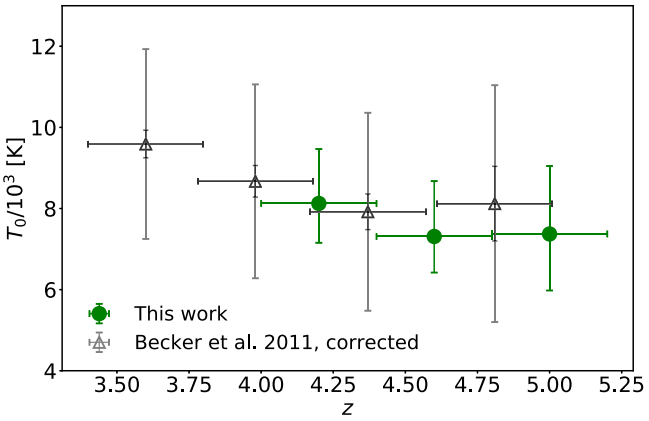

Figure 15. Temperature at the mean density of the IGM obtained in this work (green points) and from the curvature measurement of Becker et al. (2011) (gray triangles) at $z \gtrsim 3.5$. The Becker et al. T_0 values have been inferred assuming $\gamma \sim 1.5$. Vertical error bars are 68% confidence intervals for this work. For Becker et al. the small error bars are the 68% statistical uncertainties, while the extensions in lighter gray include the Jeans smoothing uncertainty estimated by those authors.

equation to solve for u_0 will then be

$$\frac{dT}{dt} = +\frac{2}{3k_B n_{\rm tot}} \sum_X \frac{dQ_{{\rm photo},X}}{dt}. \quad (15)$$

Because the specific internal energy can be expressed as $u = \frac{3}{2}k_B T \frac{1}{m}$, Equation (15) can be solved as

$$\frac{du_0}{dt} = \frac{1}{\bar{\rho}} \sum_X \frac{dQ_{{\rm photo},X}}{dt}, \quad (16)$$

where $\bar{\rho}$ is the mean mass density.

8.2. Instantaneous Reionization Parameters

We parameterize our models using three numbers: the redshift of instantaneous reionization, $z_{\rm rei}$, the temperature reached by the IGM during hydrogen reionization, $T_{\rm rei}$, and the spectral index of the post-reionization ionizing background, $\alpha_{\rm bk}$. Figure 17 presents the effects on the evolution of T_0 (top row) and u_0 (bottom row) of these three parameters. The first column shows models with the same $z_{\rm rei}$ and $\alpha_{\rm bk}$ but different reionization temperatures. Radiative transfer calculations suggest that temperatures during reionization should reach 17,000 K $\lesssim T_{\rm rei} \lesssim$ 25,000 K (e.g., Miralda-Escudé & Rees 1994; D'Aloisio et al. 2018); however, we explored $T_{\rm rei}$ down to 10,000 K and up to 30,000 K, where the upper range is similar to the hottest scenario of short and late reionization considered in D'Aloisio et al. (2018). The second column in Figure 17 demonstrates how changing the timing of reionization influences the histories of T_0 and u_0. We tested models spanning a range of redshifts from $z_{\rm rei} = 5.5$ up to $z_{\rm rei} = 12$. Finally, the third column shows the effect of changing the spectral index of the post-reionization ionizing background. The value of $\alpha_{\rm bk}$ can be connected to the intrinsic spectral index of the sources, α_s, via the expression $\alpha_{\rm bk} \approx \alpha_s - 3(\beta - 1)$, where β is the logarithmic slope of the column density distribution of intergalactic hydrogen absorbers (Upton Sanderbeck et al. 2016), which is valid at $z \gtrsim 3$ when the physical mean free path of 1 ryd photons $\lambda_{\rm MFP} \ll cH^{-1}$. The value of β may vary, but for this analysis we adopt $\beta = 1.3$ from Songaila & Cowie (2010).

Here we focus on two cases: reionization driven by star-forming galaxies with a soft $\alpha_{\rm bk} = 1.5$ ($\alpha_s \sim 2.4$, within the commonly adopted range between 1 and 3; e.g., Bolton & Haehnelt 2007; Kuhlen & Faucher-Giguère 2012), and models of quasar-driven reionization with $\alpha_{\rm bk} = 0.5$ (corresponding to $\alpha_s \sim 1.4$; e.g., Telfer et al. 2002; Shull et al. 2012). We note that radiative transfer calculations have shown that the temperature increase from the passage of an ionization front may not depend strongly on the spectrum of the ionizing sources (D'Aloisio et al. 2018). We therefore constrain $T_{\rm rei}$ independently from $\alpha_{\rm bk}$.

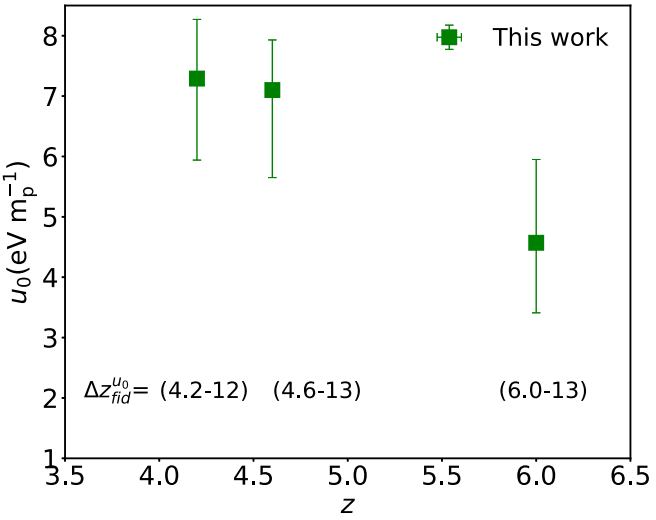

Figure 16. Our constraints on the integrated heat input into the IGM, computed for the fiducial redshift range indicated below each point. Error bars are marginalized 68% confidence intervals. Note that, for each of the observational power spectra, the fiducial redshift range used to compute u_0 has been chosen to maximize the sensitivity of the power spectrum to this parameter and therefore minimize the measurement's errors (see Section 6.3.3 for details).

8.3. Method

To be conservative, the constraints on instantaneous reionization models presented in this paper will be obtained

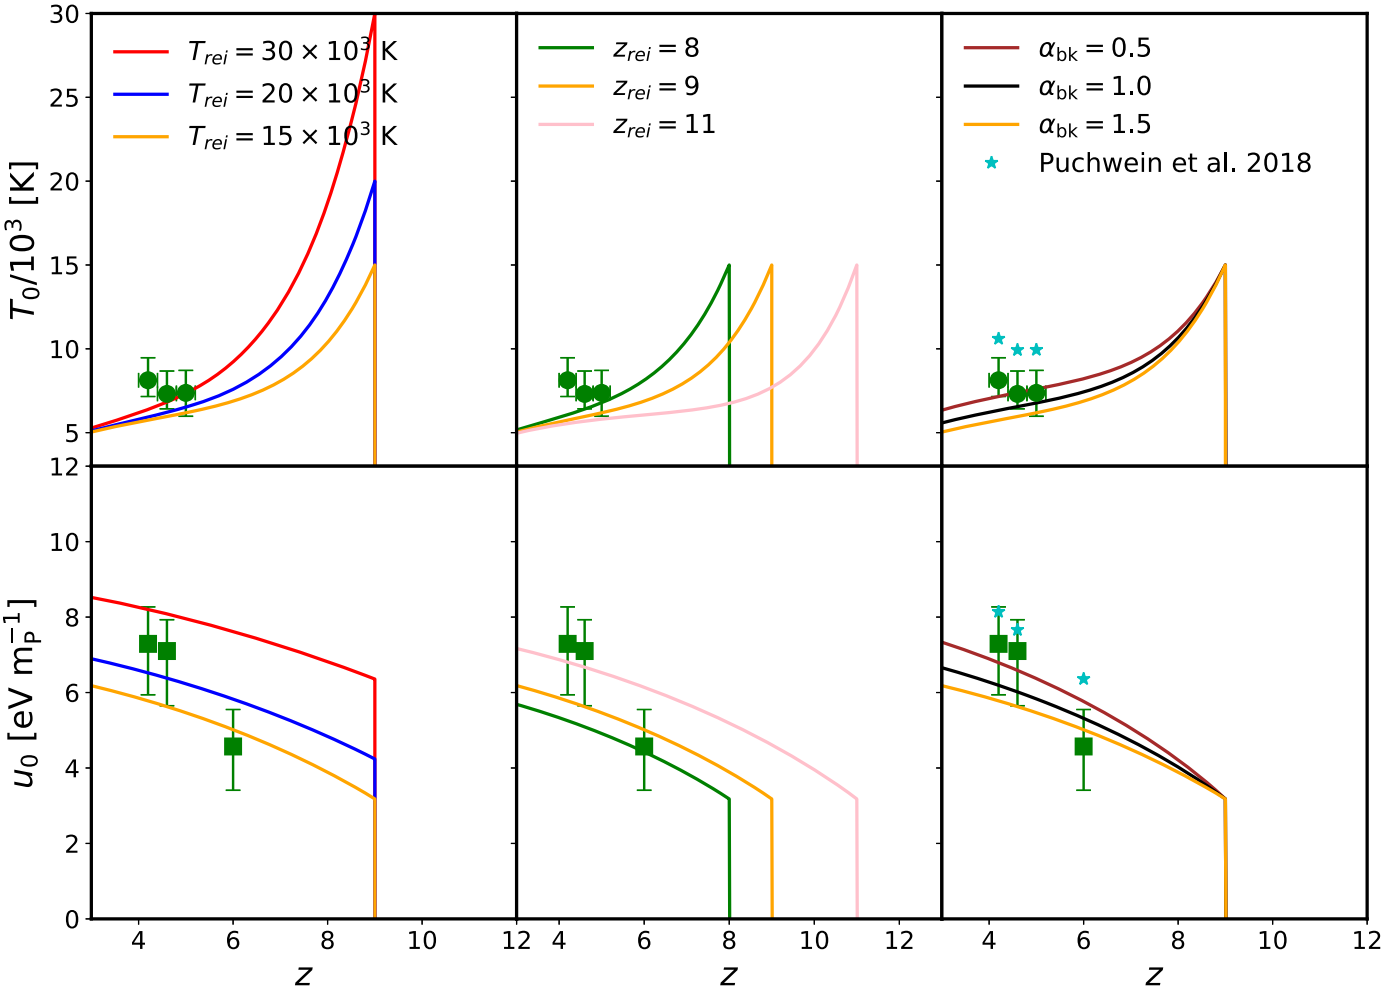

Figure 17. Evolution of the temperature of the IGM (top row) and the total energy injected per unit mass (bottom row) at the mean density, for toy models of instantaneous hydrogen reionization. A fiducial model with $T_{rei} = 15{,}000$ K, $z_{rei} = 9$, and $\alpha_{bk} = 1.5$ is plotted in orange in each panel. The left column shows the effect of varying the reionization temperature. The middle column shows how the evolution of T_0 and u_0 is influenced by the redshift of instantaneous reionization. The right column varies the spectral index of the post-reionization ionizing background. In all panels, the green points with 68% error bars are the observational constraints obtained in this work. For comparison, in the right column we also report the values of temperature and heating extracted from the empirically calibrated UVB model of Puchwein et al. 2019 (cyan stars).

using the thermal parameters in the lowest redshift bin ($z = 4.2$) only as upper limits because they may be affected by the extra heating due to the He II reionization. For each combination of parameters (α_{bk}, z_{rei}, T_{rei}) we obtain three likelihood values corresponding to the redshifts of the observational constraints: $L_{model}[<T_0^{z=4.2}, <u_0^{z=4.2}]$, $L_{model}[T_0^{z=4.6}, u_0^{z=4.6}]$, and $L_{model}[T_0^{z=5.0}, u_0^{z=6.0}]$. These probabilities describe how well a model can simultaneously fit the observed values of T_0 and u_0 at each redshift. A given L_{model} is obtained by associating the model's T_0 and u_0 values with the probability derived from the corresponding full posterior distribution in Figures 11, 12, or 13 (middle panels).

The final probability of each model is computed by multiplying the independent likelihood values obtained at each redshift:

$$L_{model} = L_{model}[<T_0^{z=4.2}, <u_0^{z=4.2}] \times L_{model}[T_0^{z=4.6}, u_0^{z=4.6}] \times L_{model}[T_0^{z=5.0}, u_0^{z=6.0}]. \quad (17)$$

8.4. Results

Before presenting the final results, we stress that, unlike previous attempts to constrain reionization using the instantaneous temperature alone (e.g., Theuns et al. 2002; Raskutti et al. 2012), the power of our approach relies on the simultaneous use of measurements of both T_0 and u_0. Figure 18 demonstrates how these measurements separately constrain the likelihood contours for our galaxy-driven reionization models. In the top panel the 68% and 95% probability contours are shown for the temperature constraints only, while in the middle panel they are given for the u_0 constraints only. The models that better fit the T_0 and u_0 data cover two different but intersecting regions in the T_{rei} versus z_{rei} parameter space. Applying both constraints simultaneously therefore reduces the allowed parameter space considerably (bottom panel of Figure 18).

Figure 19 shows the final 68% and 95% 2D probability contours for models of instantaneous reionization driven by softer ($\alpha_{bk} = 1.5$; green contours) and harder sources ($\alpha_{bk} = 0.5$; blue contours). For softer sources the favored

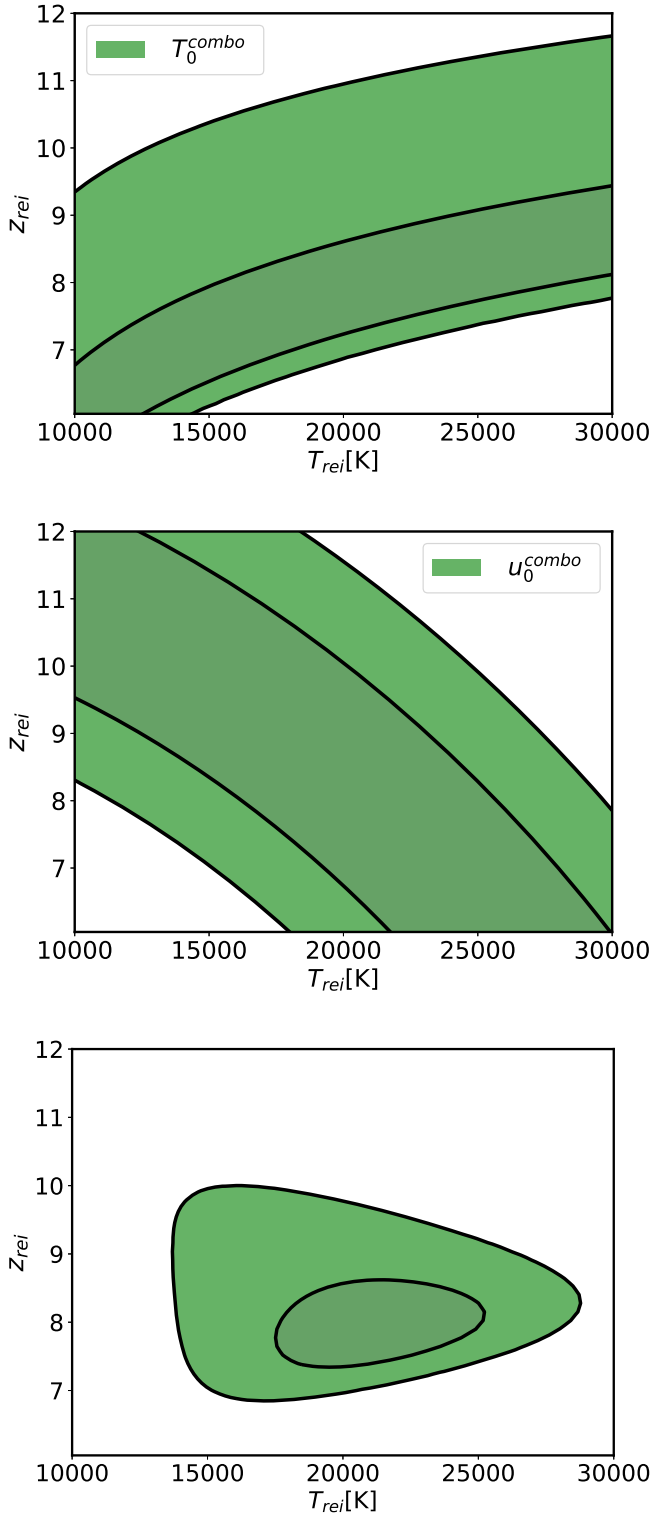

Figure 18. Constraints on instantaneous reionization parameters from our T_0 and u_0 measurements. The three panels show the 68% and 95% probability contours in the $T_{\rm rei}$ vs. $z_{\rm rei}$ parameter space for $\alpha_{\rm bk} = 1.5$, obtained when considering only the temperature results (top panel), only the u_0 measurements (middle panel), and both together (bottom panel).

models are the ones with $z_{\rm rei} \sim 8$ and reionization temperature of $20{,}000\,{\rm K} \lesssim T_{\rm rei} \lesssim 25{,}000\,{\rm K}$. These temperatures are consistent with the values predicted by radiative transfer models (e.g., Miralda-Escudé & Rees 1994; McQuinn 2012; D'Aloisio et al. 2018). For harder sources the thermal data prefer earlier reionizations and lower temperatures ($T_{\rm rei} \lesssim 20{,}000\,{\rm K}$). These results are driven by the fact that for harder post-reionization ionizing backgrounds the IGM temperature needs more time to cool in order to match the relatively low values observed at $z \lesssim 5$. We note that even lower $T_{\rm rei}$ would be needed to fit the observations if we included the contribution of He II photoheating, which has been conservatively excluded. Radiative transfer calculations may disfavor $T_{\rm rei} \lesssim 17{,}000\,{\rm K}$, as this would imply reionization front speeds unexpectedly low even for the early stages of reionization (D'Aloisio et al. 2018). Some of the parameter space in Figure 19 preferred by harder sources may therefore be disfavored on physical grounds.

In Figure 17, we also compare our constraints to the empirically calibrated UVB model recently presented by Puchwein et al. (2019). Their predictions for the instantaneous temperature are larger than our measurements and are inconsistent with our new constraints at around the 2σ level. Furthermore, the cumulative energy input into the IGM at mean density also exceeds our constraint at $z = 6$ and is again inconsistent at around 2σ. This suggests that there is slightly too much IGM heating at $z > 6$ in the fiducial Puchwein et al. (2019) model, under the assumption of a ΛCDM cosmology. This difference will be exacerbated further for models where AGN provide a substantial contribution to the photon budget for reionization.

Finally, we compare our results to constraints on the redshift of instantaneous reionization derived from the most recent Planck measurements of the Thompson scattering optical depth (Planck Collaboration et al. 2018). We marginalized over $T_{\rm rei}$ and $\alpha_{\rm bk}$ (where $\alpha_{\rm bk}$ was allowed to vary from 0.5 to 1.5) to obtain the 1D probability distribution on $z_{\rm rei}$ from the IGM thermal history. Figure 20 compares our marginalized constraints on $z_{\rm rei}$ (green solid line) to those derived for an instantaneous reionization from the Planck baseline optical depth constraint $\tau_{\rm e} = 0.0544 \pm 0.0073$ (based on Planck TT, TE, EE+lowE+lensing; Planck Collaboration et al. 2018) (red dotted–dashed line). The two distributions are broadly consistent and have comparable constraining power, with $z_{\rm rei} \simeq 8.5^{+1.1}_{-0.8}$ from the thermal history and $z_{\rm rei} \simeq 7.7^{+0.7}_{-0.7}$ from the Planck results. The combined probability distribution (blue dashed line) gives $z_{\rm rei} \simeq 8.1^{+0.5}_{-0.5}$. While we emphasize that the instantaneous reionization adopted in this section is simplistic, the proof of concept presented here demonstrates the potential of our observational constraints to inform reionization scenarios.

9. Conclusions

In this work we have presented the first simultaneous constraints on the instantaneous temperature and integrated thermal history of the IGM at $z > 4$ and demonstrated how these results can be used to test different scenarios of hydrogen reionization. We have utilized a sample of 15 Keck/HIRES and VLT/UVES high-resolution and high-C/N spectra to obtain new measurements of the Lyα forest flux power spectrum over redshifts $4.0 \lesssim z \lesssim 5.2$, for the first time pushing the measurement down to the smallest scales currently accessible to high-resolution quasar spectra at these redshifts ($\log(k/{\rm km}^{-1}\,{\rm s}) \leqslant -0.7$). We fit the new flux power spectra to obtain robust constraints on the instantaneous IGM temperature, T_0, and integrated energy input per unit mass, u_0, marginalizing over the slope of the T–ρ relation and the effective optical depth, and assuming a ΛCDM cosmology.

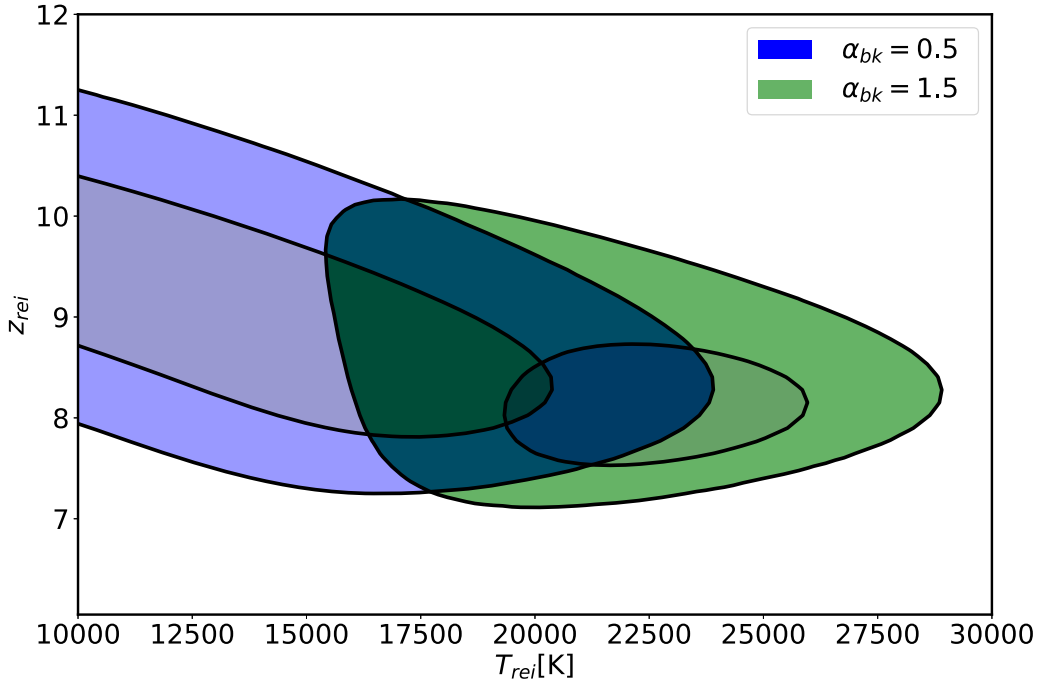

Figure 19. Constraints on instantaneous reionization models for two choices of the post-reionization ionizing background spectrum. The 2D 68% and 95% probability contours in the $T_{\rm rei}$ vs. $z_{\rm rei}$ parameter space are reported for softer ($\alpha_{\rm bk} = 1.5$; green contours) and harder ionizing ($\alpha_{\rm bk} = 0.5$; blue contours) backgrounds.

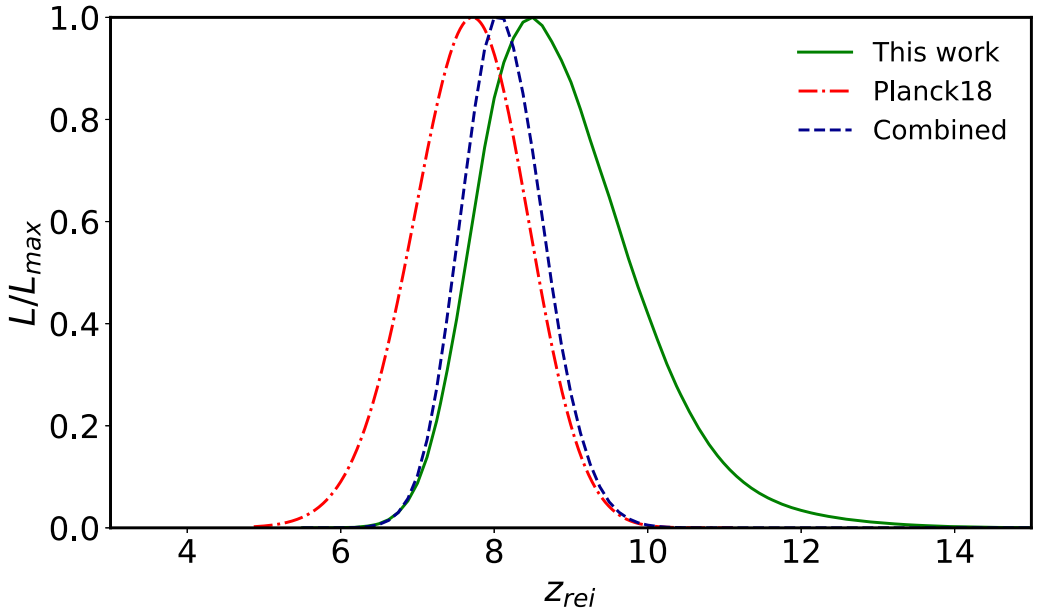

Figure 20. 1D probability distribution for $z_{\rm rei}$ obtained in this work after marginalizing over $T_{\rm rei}$ and $\alpha_{\rm bk}$ (green solid line) and from the Thompson optical depth constraints of Planck Collaboration et al. (2018) (red dot-dashed line), assuming an instantaneous reionization in both cases. The combined probability distribution distribution is plotted as a blue dashed line.

In agreement with previous results from the curvature method (Becker et al. 2011), we find temperatures of $T_0 \sim 7000$–8000 K and no strong temperature evolution over $4.2 \lesssim z \lesssim 5.0$. Our first constraints on u_0 show a significant increase from $u_0 \sim 4.5$ eV $m_{\rm p}^{-1}$ at $z > 6$ to 7.1 eV $m_{\rm p}^{-1}$ at $z > 4.2$–4.6. These results are consistent with a heating from reionization at $z > 6$ and residual photoionization heating over $z \sim 6$–4.

Our constraints on T_0 and u_0 can be used to test any reionization scenario for which the temperature and the energy injection into the IGM can be calculated. As a proof of concept we analyzed simplistic, semianalytical models of instantaneous reionization. These toy models depend on three parameters: the IGM temperature reached during reionization ($T_{\rm rei}$), the redshift of reionization ($z_{\rm rei}$), and the spectral index of the post-reionization UVB ($\alpha_{\rm bk}$), which is related to the sources driving the reionization process. We find that our measurements prefer instantaneous reionization redshifts near $z_{\rm rei} \sim 8$ with $T_{\rm rei} \sim 20{,}000$ K for a relatively soft UVB dominated by ionizing photons from star-forming galaxies. Our fully

marginalized constraints on the reionization redshift, $z_{\rm rei} \simeq 8.5^{+1.1}_{-0.8}$, are moreover comparable to those from recent Planck results.

While tests of more realistic scenarios of reionization are left for future work, the proof of concept presented here demonstrates the potential of the IGM thermal history at high redshift to impose tight constraints on the timing—and possibly the sources—of reionization.

We thank Simeon Bird and Anson D'Aloisio for helpful conversations. E.B. and G.D.B. were supported by the National Science Foundation through grant AST-1615814. J.S.B. acknowledges the support of a Royal Society University Research Fellowship. This work is based on observations made at the W. M. Keck Observatory, which is operated as a scientific partnership between the California Institute of Technology and the University of California; it was made possible by the generous support of the W. M. Keck Foundation. It also includes observations made with ESO Telescopes at the La Silla Paranal Observatory under program ID 092.A-0770. E.B. thanks Michael Murphy and the Centre of Astrophysics and Supercomputing at Swinburne for granting the access to the Swinburne supercomputer facility during the preparation of this work. The hydrodynamical simulations used in this work were performed with supercomputer time awarded by the Partnership for Advanced Computing in Europe (PRACE) 8th Call. We acknowledge PRACE for awarding us access to the Curie supercomputer, based in France at the Tres Grand Centre de Calcul (TGCC). This work also made use of the DiRAC High Performance Computing System (HPCS) at the University of Cambridge. These are operated on behalf of the STFC DiRAC HPC facility. This equipment is funded by BIS National E-infrastructure capital grant ST/J005673/1 and STFC grants ST/H008586/1 and ST/K00333X/1. We thank Volker Springel for making P-GADGET-3 available.

Appendix A
Systematic Effects

In this appendix we review some of the steps of our analysis to check and quantify possible systematic uncertainties arising from the specific strategies adopted.

A.1. Rolling Mean

As described in Section 5.1, we computed the flux contrast δ_F of Equation (3) using a rolling mean along the entire Lyα forest region. The application of this technique on both simulated and observational lines of sight guarantees a fair comparison between models and the real measurement when the continuum level is unknown, but it may introduce possible bias when comparing our power spectrum with previous works in which the power was computed from continuum-normalized spectra. Using the simulations, we tested different averaging functions and window sizes in order to minimize the impact of the rolling mean on the power spectrum at the relevant redshifts and to verify the ability of the rolling mean to capture continuum fluctuations. As demonstrated below, we found that a 40 h^{-1} cMpc boxcar rolling mean is able to recover the power at all relevant scales even in the presence of continuum fluctuations.

Figure 21 shows, for our three redshift bins, the comparison between the power spectrum computed from simulated data sets using the rolling mean technique (green dashed line) and using a fixed mean flux (black solid line). Both the synthetic samples of lines of sight used in this test have been created following the procedure described in Section 6. For the rolling mean model, we first imposed on each of the lines of sight a random continuum selected from the real continua fitted for the XQ-100 survey (López et al. 2016). We then run the 40 h^{-1} cMpc boxcar rolling mean directly on the total Lyα + continuum flux. Differences between the two power spectra are shown in the bottom panel of each plot and compared with the statistical error characterizing our observational sample (green shaded region; see Section 5.9). We note that, at all redshifts, the discrepancies between the two models always lie well within the statistical error, with systematic uncertainties $\sigma_{\rm roll}$ typically $\lesssim 0.20\sigma_{\rm stat}$. We therefore do not expect our results to be sensitive to this averaging choice.

A.2. Windowing Effects

As explained in Section 5.5, we compute the observational power spectrum in 20 h^{-1} cMpc sections of Lyα forest. Dividing the spectra into many small regions may introduce artificial excess power at intermediate and small scales owing to a windowing effect. This effect does not not appear in the simulations because of their periodicity.

In Figure 22 we present the effect on the flux power spectrum of dividing the spectra into smaller sections. At each redshift bin, the power computed from the largest 40 h^{-1} cMpc simulation box (baseline; solid black line) is compared with the power computed from the same simulation but dividing each of the native synthetic spectra into sections of 10 h^{-1} cMpc (blue dot-dashed line) and 20 h^{-1} cMpc (green dashed line). Changes in the power are reported as fractions of the baseline power spectrum values in the bottom panel of each plot. While cutting the spectra into 10 h^{-1} cMpc sections introduces an excess of power at the small scales ($\log(k/{\rm km}^{-1}\,{\rm s}) \gtrsim -1.1$) of $\sim$20%–25%, the windowing effect for the 20 h^{-1} cMpc sections is less significant, with variations in the power $\lesssim$8% at all scales. We therefore opted for the latter section size in our analysis.

A.3. Noise Subtraction

In this appendix we test the noise power subtraction method described in Section 5.6.2 using synthetic data sets generated from simulations. For each synthesized line of sight we use the error and the flux arrays of one of our observed spectra to fit a linear correlation between the signal and the noise level. Using these correlations, we then add noise to the simulated samples in a flux-dependent way. We construct synthetic samples of lines of sight following the procedure described in Section 6.2 using our 20 h^{-1} cMpc simulation box and compute the power spectrum with and without adding noise. In Figure 23 we show the noisy (red dot-dashed line) and noiseless (baseline; black solid line) power spectra obtained for each redshift bin.

We finally applied the noise power subtraction method to the noisy model and compare the corrected power spectrum (green dashed line in Figure 23) to the noiseless baseline. The noiseless power is recovered with errors of $\lesssim$2%. This suggests that this step of the data analysis is not introducing relevant systematic effects in the final results.

A.4. Instrumental Resolution

In this appendix we review the effect of uncertainties in the instrumental resolution when smoothing the synthetic spectra to

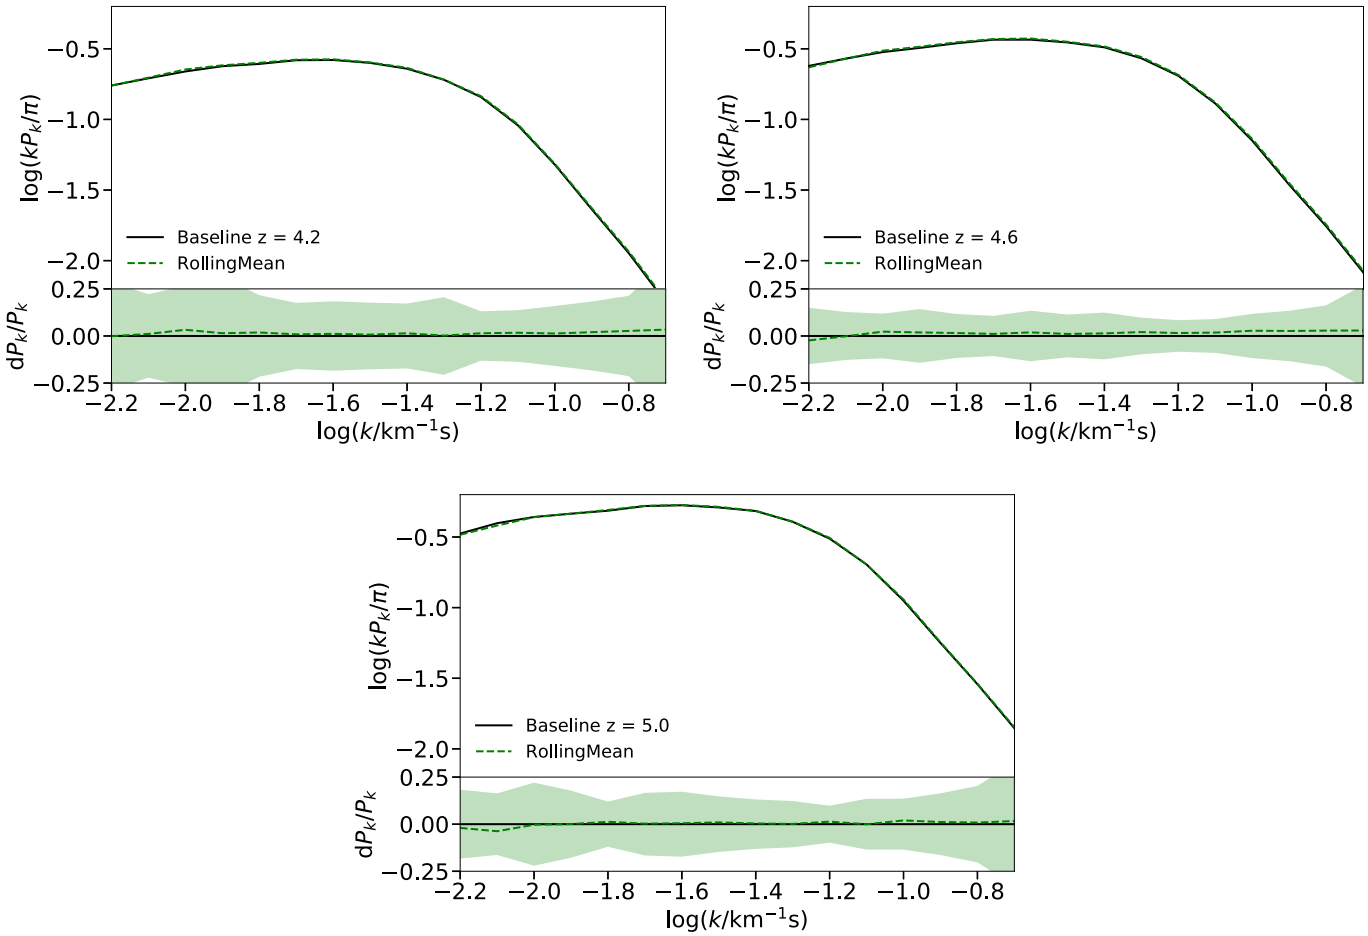

Figure 21. Effect of the rolling mean on the power spectrum measurements. At each of the three redshift bins we show the power spectrum computed using a 40 h^{-1} cMpc boxcar rolling mean (green dashed line) compared with the values obtained using a standard global flux average (black solid line). For the rolling mean calculation a random quasar continuum has been imposed on each synthetic line of sight. Differences between the two power spectra are reported in the bottom panel of each plot and are compared with the 68% statistical error for our observed sample (green shaded region).

match the resolution of the observed data. We calibrate the synthetic Lyα forest spectra using different instrumental resolutions, taking our nominal values of FWHM = 6 km s^{-1} as a baseline.

Figure 24 shows the variations in the power spectrum expected for a change of +10% (red dot-dashed line) and −10% (green dashed line) in the observed spectral resolution. We note that the power changes by $\lesssim$5% at all scales (see lower panel of each plot).

A.5. Masking Correction Function

In this appendix we analyze the systematic uncertainties arising from the choice of a particular simulation in computing the masking correction, $C_m(k)$, described in Section 5.6.1. For the fiducial masking correction we adopted the simulation S20_1z15 of Table 2. To demonstrate that this particular choice of model (with $T_0 \sim 7500$ K and $\gamma \sim 1.5$ at the redshifts of interest) does not relevantly affect the final power spectrum measurements, we compare our fiducial results with the measurements obtained when computing the masking correction from post-processed runs with different values of T_0 and γ. In particular, we tested two extreme cases: a colder model, $C_m^{\rm test1}(k)$, with $T_0 \sim 4000$ K and $\gamma \sim 1.5$, and an isothermal model, $C_m^{\rm test2}(k)$, with $T_0 \sim 7500$ K and $\gamma \sim 1.0$.

Figure 25 shows the power spectrum obtained using different masking correction functions for the redshift bins considered in our analysis. In all the cases, computing the masking correction using models with different thermal parameters affects the small scales only mildly, with variations in the power of $\lesssim$4% for scales $\log(k/{\rm km}^{-1}\,{\rm s}) \gtrsim -1.1$. We therefore do not expect that uncertainties in the masking correction function will relevantly affect our final constraints.

A.6. Metal Correction

The Lyα forest region is affected by narrow metal line contaminants that may increase the flux power at the small scales. To correct for the metal contribution, we subtract the metal power evaluated in Section 5.7 from the final power spectrum measurements. The effect of the metal subtraction on the power spectrum measurements is presented in Figure 26 for our three redshift bins. As expected, only the small scales ($\log(k/{\rm km}^{-1}\,{\rm s}) \gtrsim -1.0$) are significantly affected. The variations are relatively small (<10%) and well within the statistical errors of our final measurements.

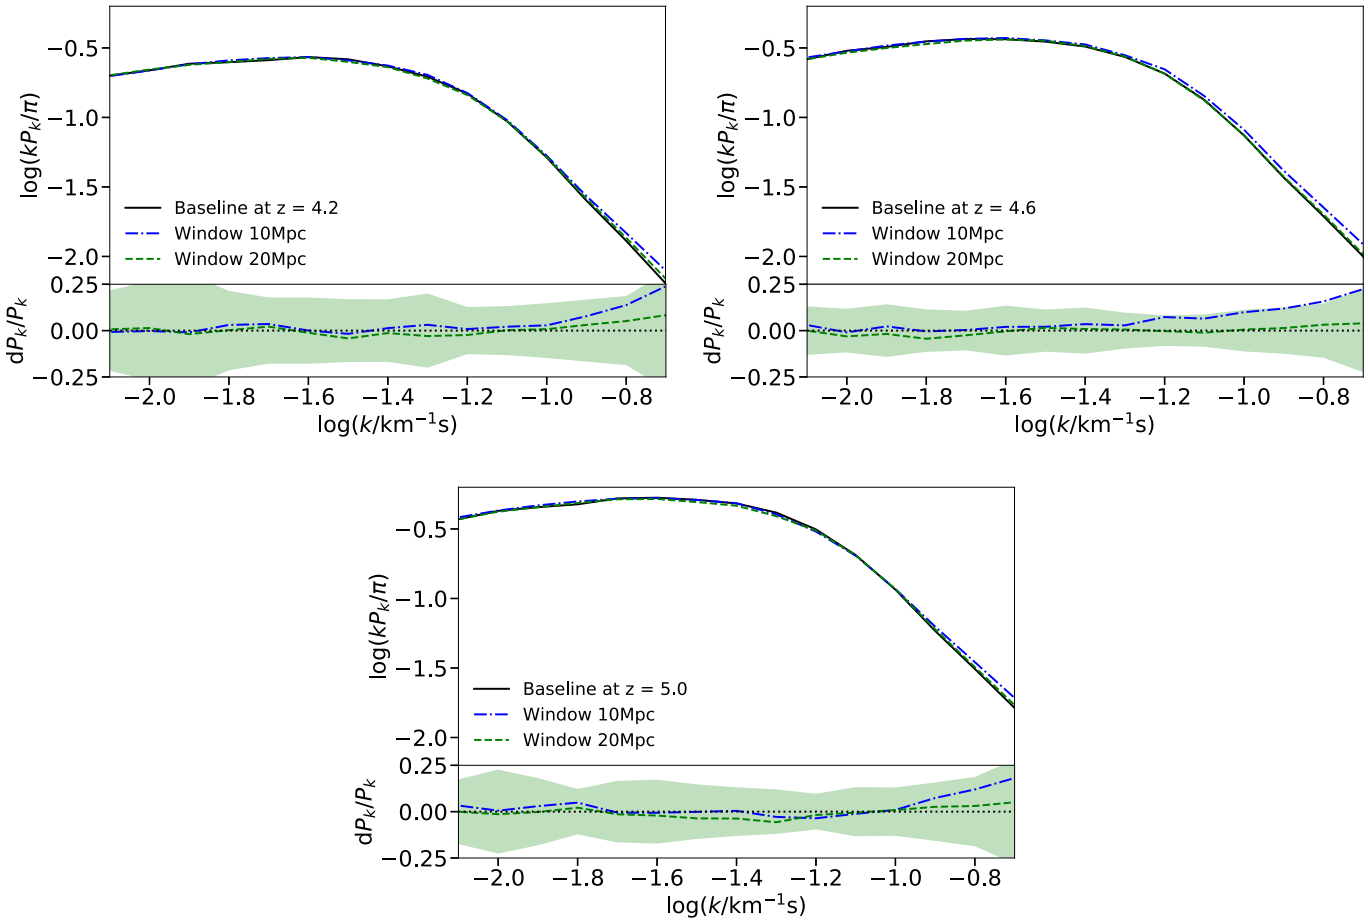

Figure 22. Windowing effect on the flux power spectrum. At each redshift, the baseline model (black solid line) has been computed from the largest simulation box (40 h^{-1} cMpc). We also show the power spectrum computed from the same simulation box but after dividing each of the native spectra into smaller sections of 10 h^{-1} cMpc (blue dot-dashed line) and 20 h^{-1} cMpc (green dashed line). Changes in the power as fractions of the baseline power spectrum are shown in the bottom panel of each plot. Given the small variations in the power for the 20 h^{-1} cMpc sections, we opted for this section size when computing the power spectrum from the observed data. For comparison, the green shaded region in the bottom panel shows the 68% statistical errors for our observational results.

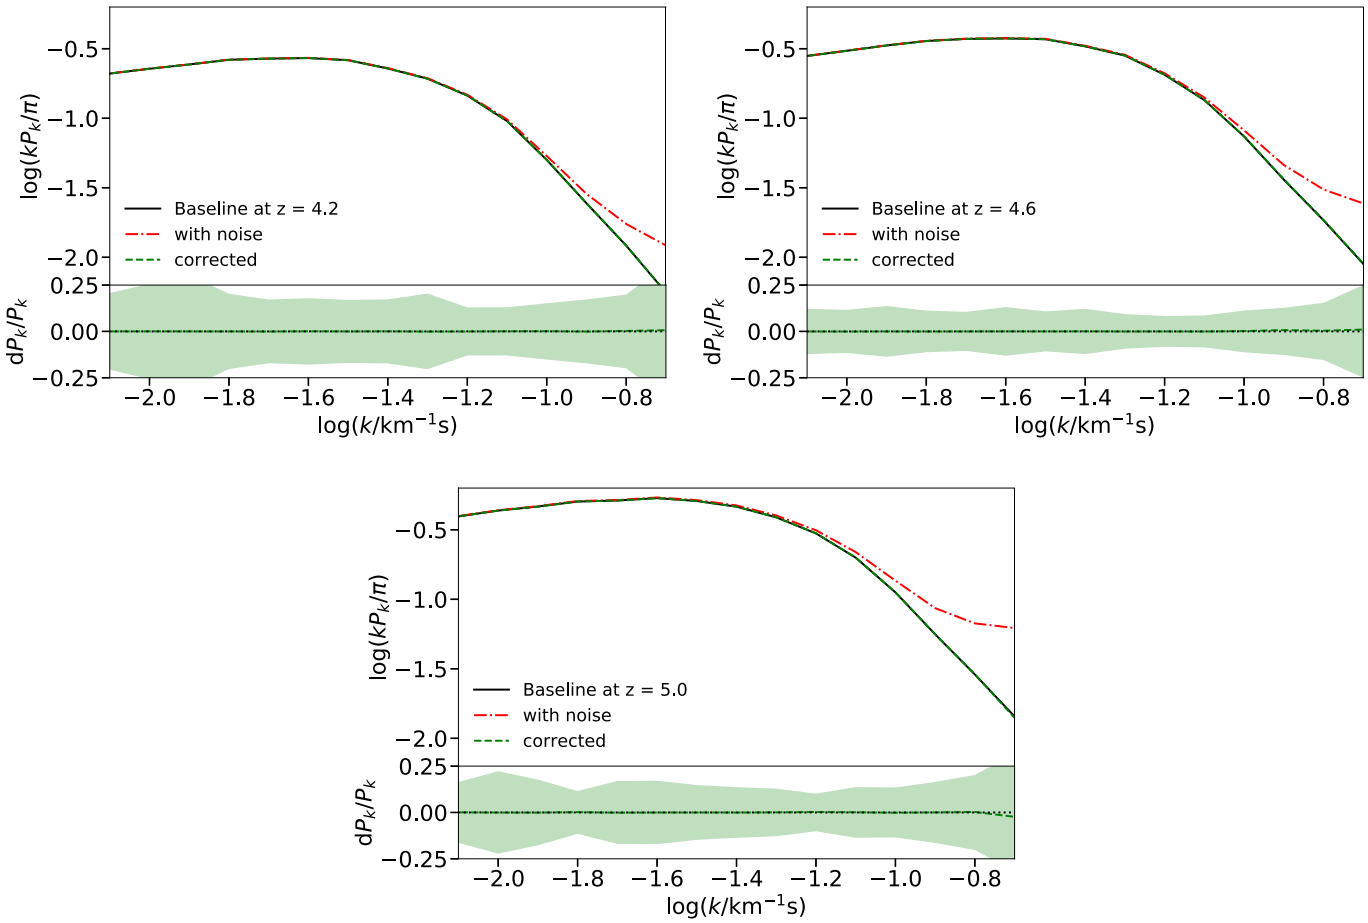

Figure 23. Noise power subtraction test. At each of the three redshift bins we show the power spectrum model for the 20 h^{-1} cMpc simulation box with (red dot-dashed line) and without (baseline; black solid line) the addition of realistic noise (see text for details). Applying the noise power subtraction method described in Section 5.6.2 to the noisy model, we obtain the corrected power spectrum (green dashed line), which is compared with the baseline in the lower panel of each plot. The noiseless power is recovered with high precision, with $\lesssim$2% changes from the baseline power at all scales. For comparison, the statistical 68% errors for the observational results are also shown (green shaded regions).

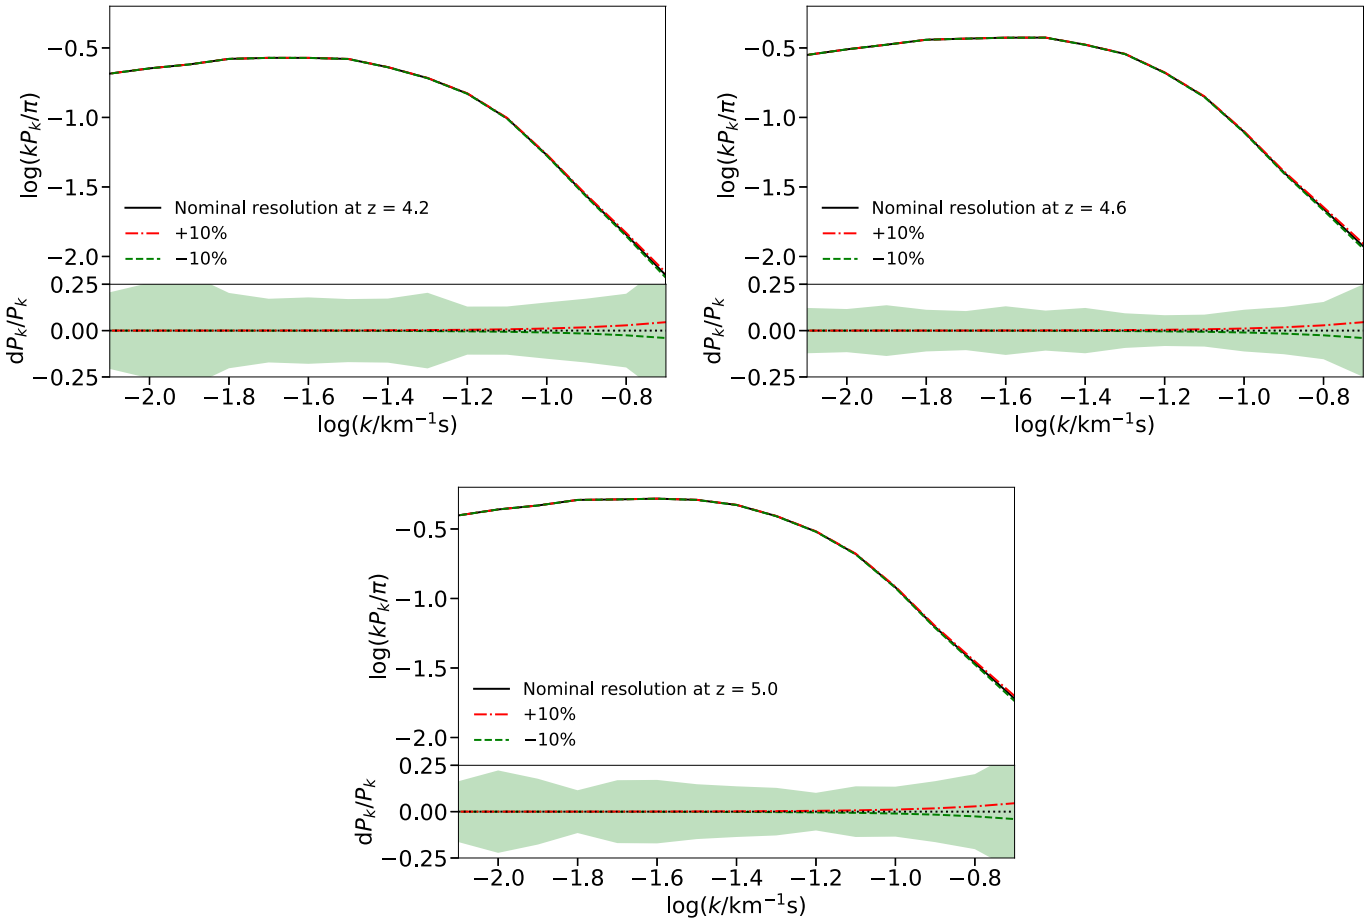

Figure 24. Effect of instrumental resolution on the Lyα flux power spectrum. For each of our redshift bins we show the power spectrum obtained from the 20 h^{-1} cMpc simulation box after smoothing the flux spectra by different instrumental resolutions. The baseline model (black solid line) has been calibrated assuming the correct nominal slit resolution, while a 10% lower (red dot-dashed line) and higher (green dashed line) resolution has been used for the smoothing of the other two models. Variations in the power spectrum at different scales are reported in the lower panels as fractions of the baseline power and are $\lesssim$5% at all scales. For comparison, the statistical 68% errors for the observational results are also shown (green shaded regions).

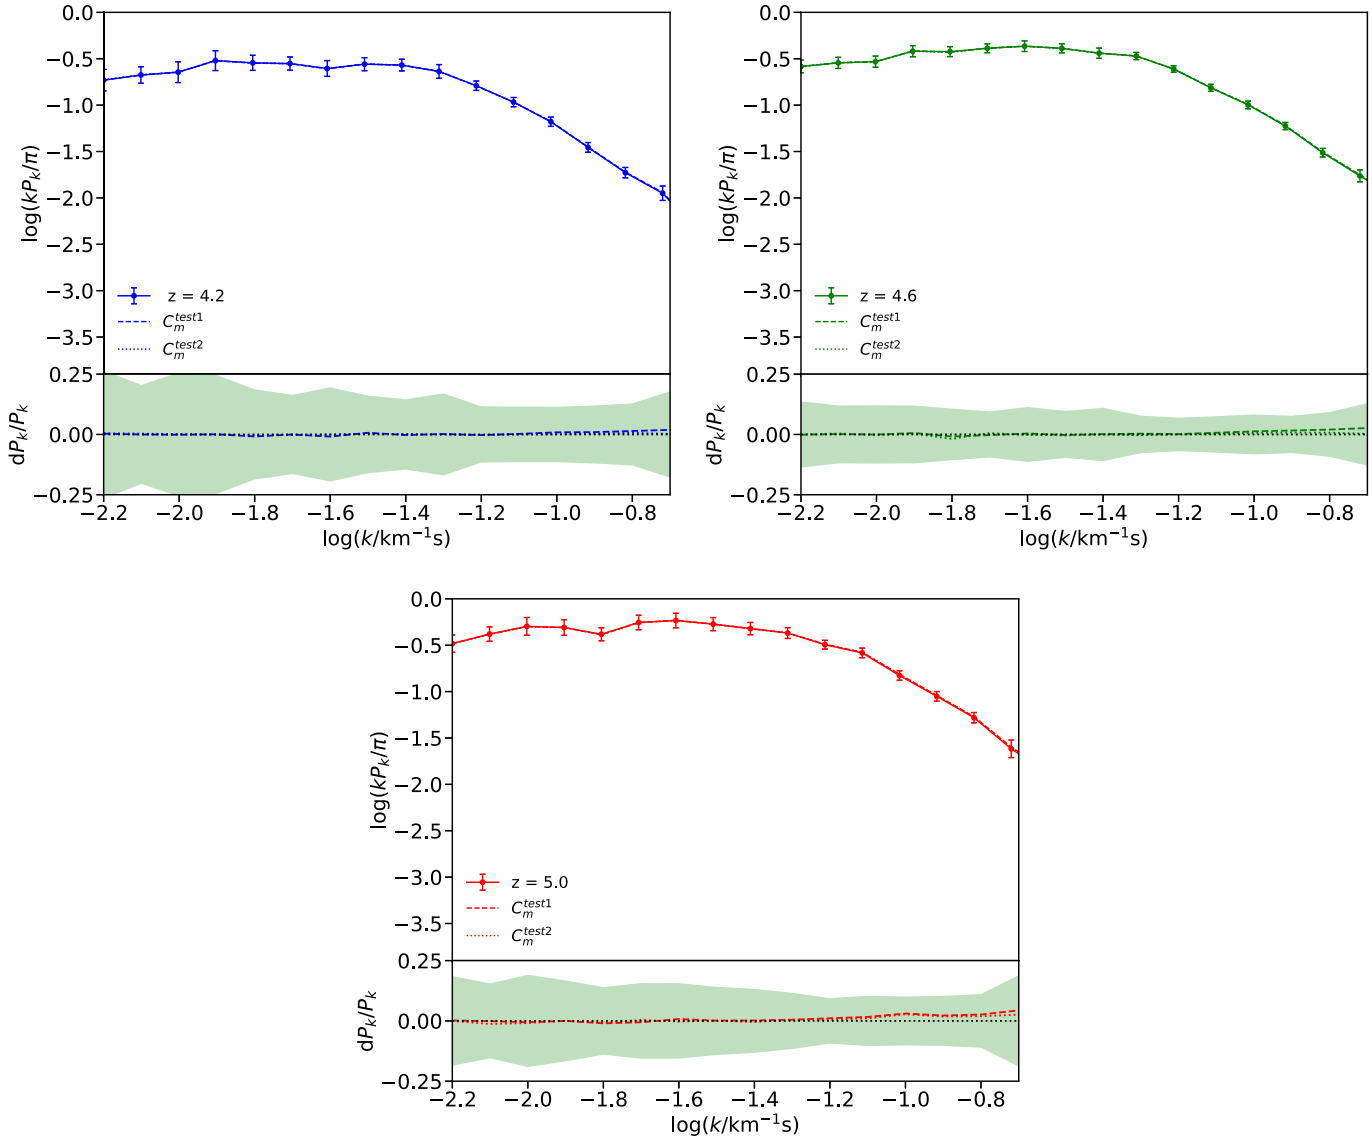

Figure 25. Effect of the choice of thermal model for calculating the masking correction function. For each redshift, the power spectrum corrected with the fiducial $C_m(k)$ function (solid colored lines and data points) is compared with the results obtained using masking correction function $C_m^{\rm test1}(k)$ computed from a post-processed model with $T_0 = 4000$ K and $\gamma = 1.5$ (dashed colored line) and $C_m^{\rm test2}(k)$ from a model with $\gamma = 1.0$ and a temperature $T_0 = 7500$ K (dotted colored line). Variations in the power spectrum as a fraction of the fiducial power are reported in the bottom section of each panel and compared with the statistical 68% uncertainties (green shaded regions). Uncertainties in the masking correction functions produce negligible variations ($\lesssim$4%) at the small scales of the final power spectrum measurement that are well within the corresponding statistical errors.

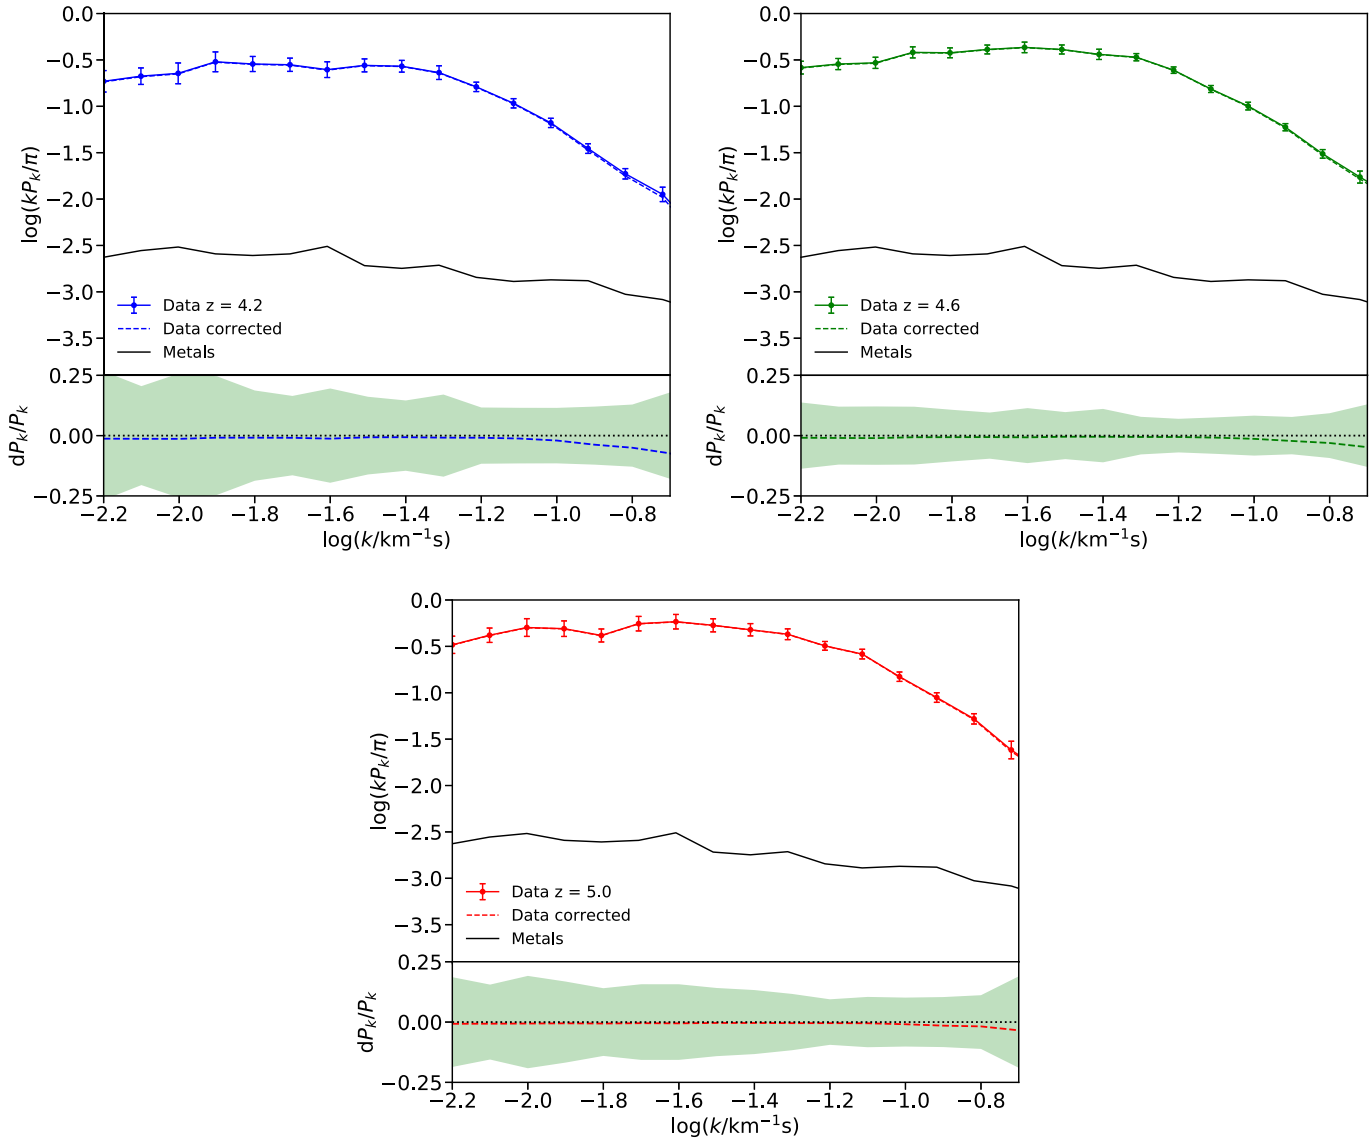

Figure 26. Effect of the metal correction on the Lyα power spectrum measurements for the three redshift bins considered in this work. In each panel the average metal power computed in Section 5.7 (black solid line) is subtracted from the total power (colored solid lines and data points) to obtain the corrected measurement (colored dashed lines). The effect of metals at small scales is more relevant in the lowest redshift bin because the amount of Lyα absorption is lower. Nevertheless, the changes in power due to the metal contribution, reported in the lower panels as a fraction of the total power, are always $<10\%$ and well within the statistical errors (green shaded regions).

Appendix B
Covariance Matrix Uncertainties

In this appendix we test how strongly the choice of simulation model for the covariance matrix regularization affects the final constraints on T_0 and u_0. In principle, off-diagonal coefficients of the covariance matrix will mildly depend on the shape of the power spectrum and therefore on the thermal parameters characterizing the models.

Figure 27 shows the results for T_0 and u_0 obtained when using a covariance matrix derived from the fiducial model S40_1z15 (Covariance matrix_1) and from the S40_1z9 model (Covariance matrix_2). Both the nominal results and error estimates shows only modest ($\lesssim$3%) changes.

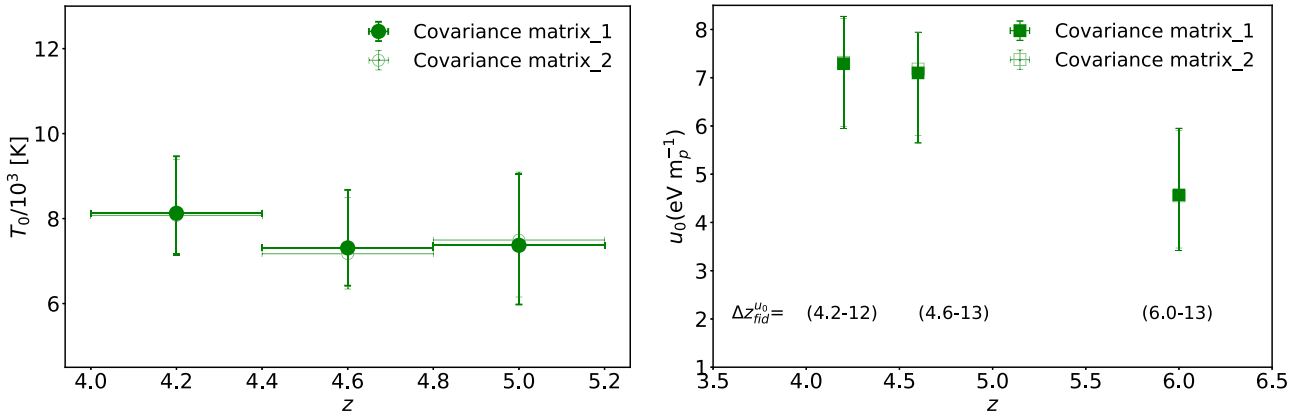

Figure 27. Effect of varying the simulation used for the covariance matrix regularization. The measurements of T_0 (left panel) and u_0 (right panel) are reported for two different choices of simulations for the covariance matrix regularization: model S40_1z15 (filled markers) and model S40_1z9 (open markers).

Appendix C
Comparison with Viel et al. (2013)

In this appendix we compare our power spectrum measurements to those from Viel et al. (2013a). While Viel et al. (2013a) included somewhat lower resolution spectra from Magellan/MIKE, for consistency we limit our comparison to their measurements obtained with the HIRES spectrograph.

We consider two main differences between our new estimates and the older Viel et al. (2013a) power spectra: flux contrast estimators and cosmic variance. First, Viel et al. (2013a) normalized the spectra using a spline continuum estimate and then computed the mean flux in large sections of data. To test whether this affected the results, we recomputed the Viel et al. (2013a) power spectra by applying the procedure described in Section 5 to the Viel et al. data, using the same sections of spectra. Note that, given the somewhat larger scales probed by Viel et al., we do not expect any relevant effect due to noise or resolution. We verified that we were able to reproduce the previous results with good precision and no significant bias was introduced by the different flux contrast estimators.

We next consider whether sample size may be playing a role. In each of our redshift bins the number of independent lines of sight contributing to our measurement is always more than double the number in Viel et al. (2013a). In particular, for the $z = 4.2$ bin we used 12 lines of sight versus 4 in Viel et al. (2013a), for the $z = 5.0$ bin we used 12 versus 5, while at $z = 4.6$, where the largest differences between the power spectra are seen, we used 15 quasars versus 5 in the previous work.

To determine whether cosmic variance can explain the discrepancy between our results and those of Viel et al. (2013a), we computed the power spectrum from subsamples of our data. Figure 28 shows the comparison between the Viel et al. (2013a) power spectra (black dashed lines) and the 68% (darker shaded regions) and 95% (lighter shaded regions) contours of the distribution of power spectrum realizations obtained from a Monte Carlo sampling of our lines of sight. For each redshift bin we randomly select from our sample the same number of lines of sight contributing to the Viel et al. (2013a) measurement and use them to compute the power spectrum. Repeating the process for many (∼200) realizations, we verify that the Viel et al. (2013a) results fall within this distribution. At $z = 4.2$ and $z = 5.0$ most of the Viel et al. (2013a) values

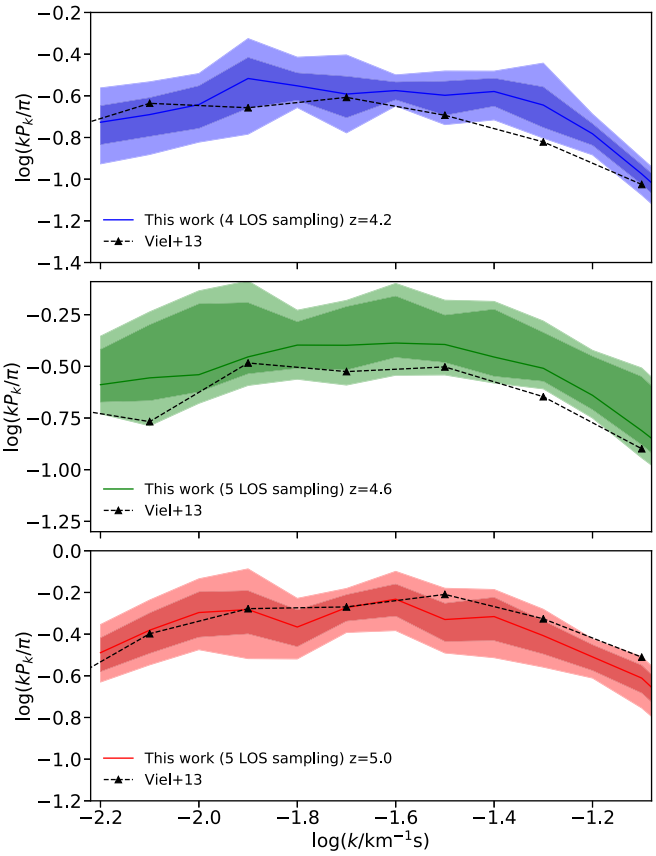

Figure 28. Comparison with previous HIRES results. For each redshift bin, the power spectrum measurements obtained by Viel et al. (2013a) (black dashed lines) are compared with the 68% (darker shaded regions) and 95% (lighter shaded regions) contours of the power spectrum realizations obtained with a Monte Carlo sampling of our lines of sight (LOS), when using the sample size of Viel et al. (see text for details). At all redshifts there is a broad consistency with the Viel et al. (2013a) values.

fall within the 68% contours. For the $z = 4.6$ bin the agreement is slightly worse but still largely within the 95% region. We note that the errors in different k bins are correlated, as discussed in Section 5.9. We further note that our sample size is still relatively modest, and our Monte Carlo technique is likely to underestimate the cosmic variance at the 95% level.

Appendix D
Comparison with Iršič et al. (2017)

In this appendix we compare our power spectrum measurement at $z = 4.2$ to the one from Iršič et al. (2017a). These authors use spectra from the XQ-100 Legacy Survey (López et al. 2016), collected using the lower-resolution VLT/X-Shooter spectrograph ($R \sim 11\,\mathrm{km\,s^{-1}}$ corresponding to an FWHM $= 26\,\mathrm{km\,s^{-1}}$ for the VIS arm). We therefore limited our comparison to the power spectrum scales where the resolution correction for the Iršič et al. (2017a) measurements is $\lesssim$20%, corresponding to $\log(k/\mathrm{km^{-1}\,s}) < -1.4$.

Figure 29 presents the comparison between the Iršič et al. (2017a) power spectrum (black points) and our measurements (blue points) at $z = 4.2$. For both data sets we also include the final 68% uncertainties (colored shaded regions). Even if the X-Shooter power spectrum shows tendentially lower values than ours, the two measurements are consistent within the 68% uncertainties at all scales but one point (at the scale $\log(k/\mathrm{km^{-1}\,s}) = -1.92$). A possible explanation for the offset in the power at scales $\log(k/\mathrm{km^{-1}\,s}) \gtrsim -1.9$ may be differences in the spectra samples and in the redshift coverage of the $z = 4.2$ bin: while we include spectra falling within a broad $\Delta z = 0.4$ redshift bin, Iršič et al. (2017a) adopt a narrower $\Delta z = 0.2$ bin. Nevertheless, as for the Viel et al. (2013a) power spectrum, we verified that, using the same X-Shooter spectra, we were able to reproduce the Iršič et al. (2017a) results at the considered scales without introducing any significant bias due to possible differences in the analysis.

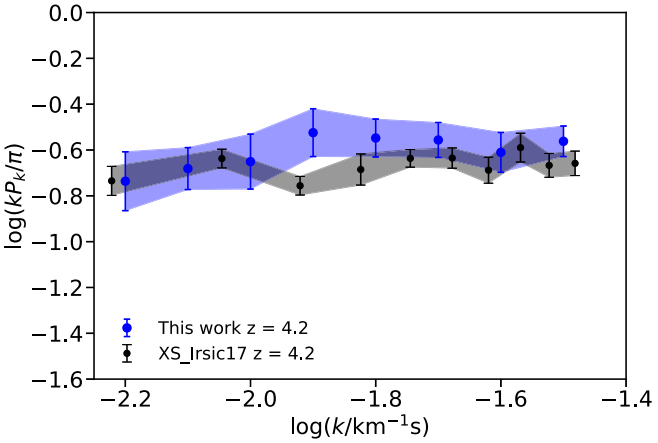

Figure 29. Comparison with previous X-Shooter results at $z = 4.2$, limited to the scales where the resolution correction for the X-Shooter data is $\lesssim$20%. The power spectrum obtained by Iršič et al. (2017a) (black points) is compared with our results for the lowest redshift bin (blue points). The two measurements are generally consistent within the 68% uncertainties (colored shaded regions).

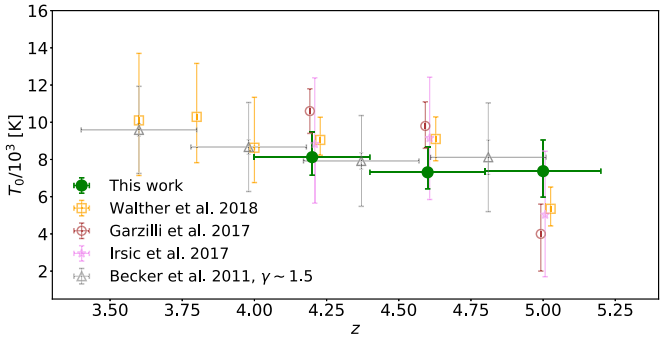

Figure 30. Comparison with previous temperature measurements over the redshift range covered by our analysis. Our results (green points) are compared with the temperature values obtained from the curvature analysis of Becker et al. 2011 (gray triangles) and with the thermal estimates from the power spectrum statistic of Iršič et al. (2017b) (pink stars), Garzilli et al. (2017) (brown circles), and Walther et al. (2018) (orange squares). Note that the temperature values from Iršič et al. and Garzilli et al. and the $z \gtrsim 4.2$ results from Walther et al. are all obtained from the flux power spectrum measurements of Viel et al. (2013), although calibrated with different sets of simulations; therefore, they are not fully independent. Vertical error bars are 68% statistical uncertainties for all the data. For Becker et al. the nominal errors have been increased to include the Jeans smoothing uncertainty estimated by those authors.

Appendix E
Comparison to Previous Temperature Measurements at High z

In Figure 30 we compare our IGM temperature measurements to previous constraints from the literature over the redshift range covered by our analysis. We note that the temperature values from Iršič et al. (2017b) (pink stars) and Garzilli et al. (2017) (brown circles) and the $z \gtrsim 4.2$ results from Walther et al. (2018) (orange squares) are all obtained from the flux power spectrum measurements of Viel et al. (2013a), although calibrated with different sets of simulations; therefore, they are not fully independent. Among these constraints, the larger error bars reported for Iršič et al. (2017b) reflect a 1.5 correction factor applied by these authors to the nominal errors associated with the Viel et al. (2013a) measurements.

While generally within the 68% uncertainties of our new results, previous estimates from the flux power spectrum statistic suggest a significant decrease in temperature at $z \gtrsim 4.6$. Our thermal constraints, however, obtained from a larger and higher-quality sample of high-resolution spectra, do not show any strong evolution in the temperature at these redshifts (see also Appendix C for a comparison between our recent flux power spectrum measurements and the ones from Viel et al. 2013a).

Appendix F
Numerical Convergence

In this appendix we examine the convergence of the Lyα flux power spectrum in the simulations used in this work. We used multiple simulations with the thermal history model 1_z15 in Table 2. All of the synthetic Lyα forest lines of sight were produced using the procedure described in Section 6.1. The flux power spectrum was then computed as in Section 6.2.

The tests are shown in Figure 31, where the convergence with box size for a fixed mass resolution ($M_{\rm gas} = 9.97 \times 10^4\ h^{-1} M_{\odot}$) is reported in the left column and the convergence

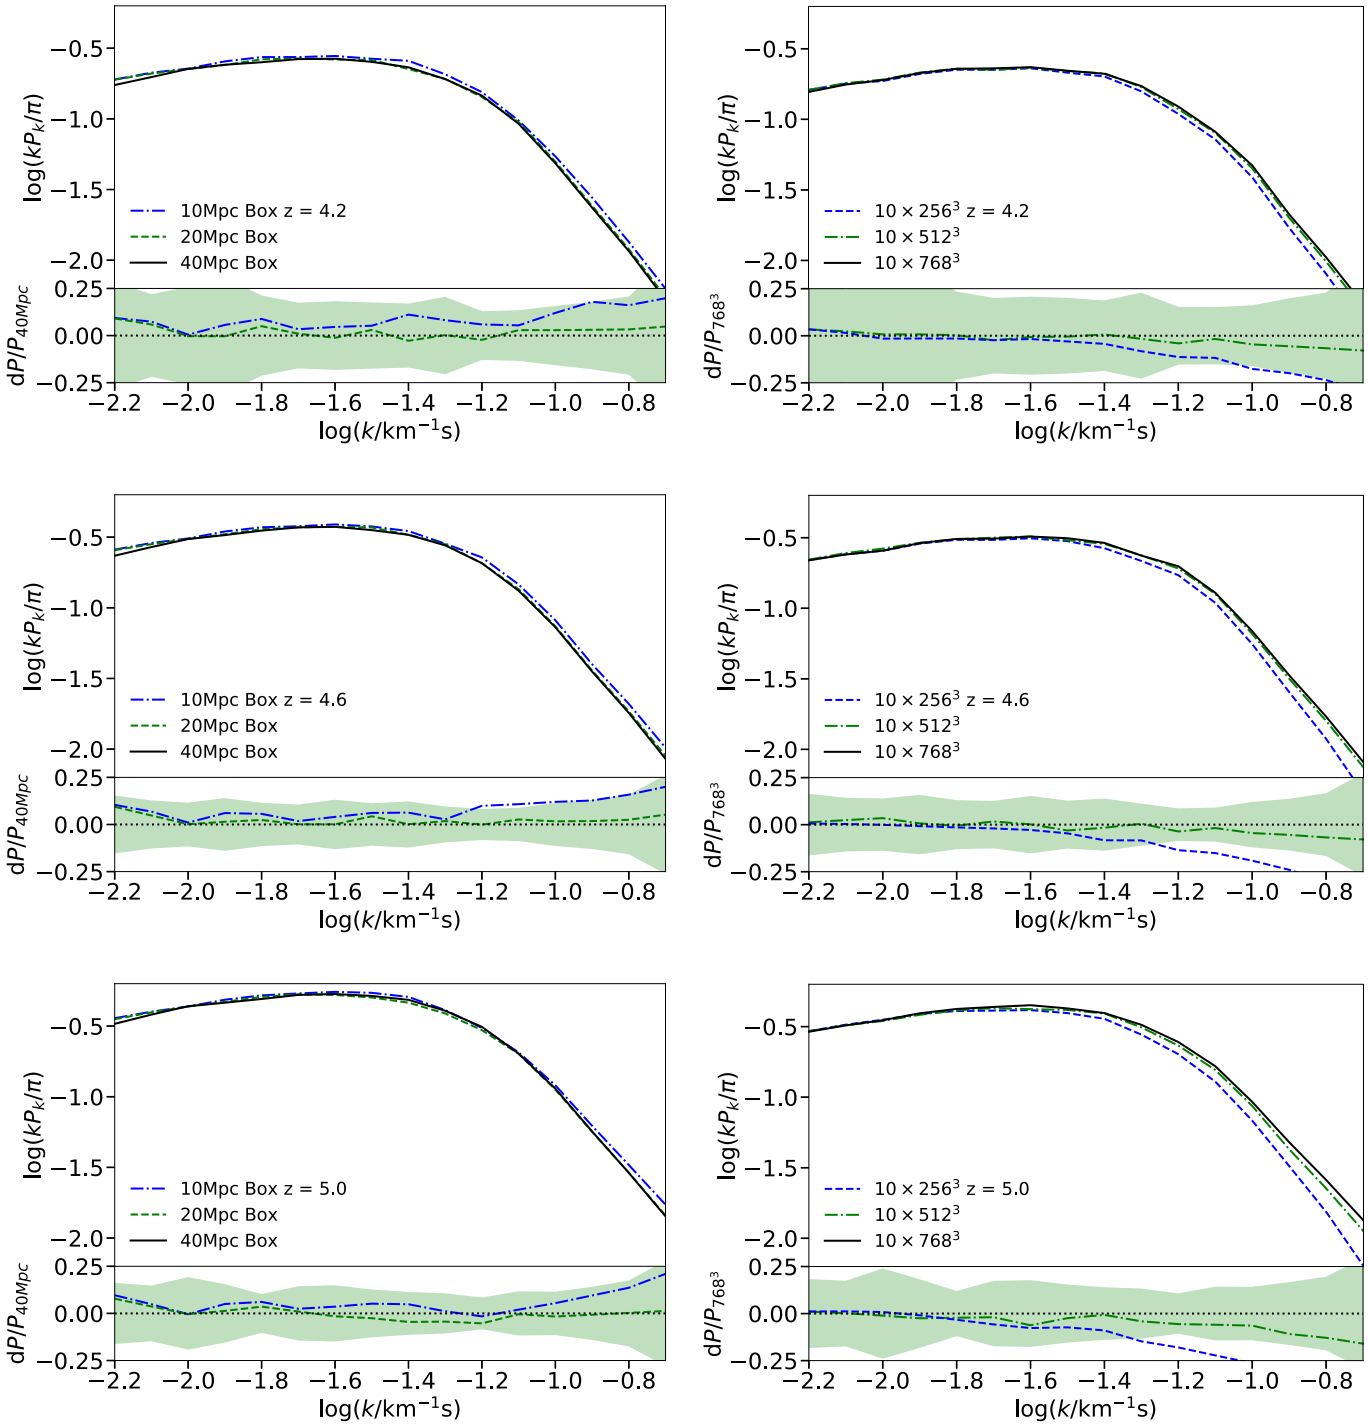

Figure 31. Convergence of the flux power spectrum with box size and mass resolution for the redshifts relevant in this work. Power spectra for this test were computed from noise-free mock lines of sight using the procedure described in Section 6.2. The left column shows the convergence with box size at a fixed mass resolution ($M_{\rm gas} = 9.97 \times 10^4\ h^{-1}\ M_{\odot}$), while the right column displays the convergence with respect to the highest mass resolution model (S10_1z15_768 in Table 2) for a fixed box size ($L = 10\ h^{-1}$ cMpc). Variations in the power spectrum as a fraction of the reference power are plotted in the bottom section of each panel and compared with the statistical 68% uncertainties (green shaded regions). Both resolution and box size corrections have been applied to our fitting models.

with mass resolution for a fixed box size ($L = 10\ h^{-1}$ cMpc) is displayed on the right. The results show that a small correction for both box size and mass resolution needs to be applied to the power spectra derived from our nominal 10 h^{-1} cMpc simulations. When increasing the box size, the power decrease up to $\sim$15%, particularly at very small scales ($\log(k/\mathrm{km}^{-1}\,\mathrm{s}) > -1$). In contrast, when increasing the mass resolution, the power toward small scales ($\log(k/\mathrm{km}^{-1}\,\mathrm{s}) \gtrsim -1.4$) increases progressively, reaching a correction of $\sim$15% at $z = 5.0$ for our nominal mass resolution. Because the corrections are in opposite directions, the final scale factor for box and mass resolution convergence is $\lesssim$5% at all scales. However, we note that in principle the mass resolution convergence may depend on the underlying IGM density structure and, consequently, on the choice of a particular thermal history model. We therefore verified that the entity of this possible systematic effect was negligible when compared to the 68% statistical uncertainties characterizing the observational data and that the final thermal constraints were not affected by the numerical corrections.

We further verified, using lower mass resolution simulations, that increasing the box size up to $L = 160\ h^{-1}$ cMpc for a fixed mass resolution does not introduce additional power at the scales considered in this work.

Appendix G
Effective Optical Depth Evolution

As explained in Section 6.1, when constructing mock samples, we account for the mild redshift evolution of the mean flux along the line of sight by initially rescaling the effective Lyα optical depth using Equation (10). In this appendix we show how the choice of this relation for the τ_{eff} evolution, while somewhat arbitrary, represents a reasonable transition between the measurements of Becker et al. (2013) and the newer results of Bosman et al. (2018). Figure 32 presents the comparison among the different τ_{eff} evolutions. The results from the analysis of Becker et al. (2013) (gray triangles) at $z < 4$ are smoothly connected to the most recent measurements of Bosman et al. (2018) (red squares) at $z > 5$ by the fiducial fit adopted in this work (green dashed line). For comparison, we also show the constraints on the optical depth obtained from our MCMC chains and reported in Table 5 (green points). We recover values broadly consistent with Equation (10) for the three redshift bins.

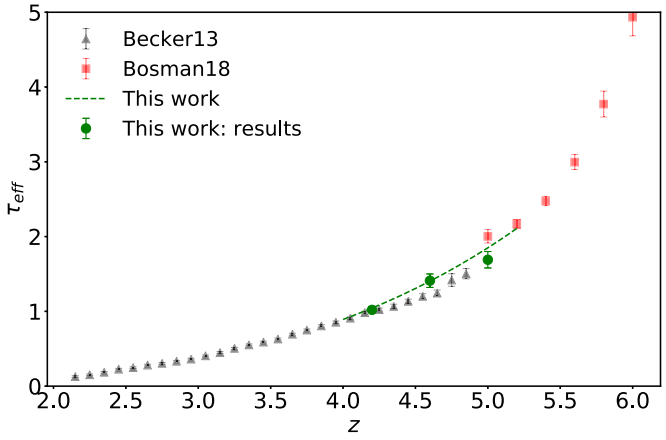

Figure 32. Evolution of effective optical depth with redshift. Measurements from Becker et al. (2013) (gray triangles) and Bosman et al. (2018) (red squares) are shown along with the fiducial evolution of Equation (10) (green dashed line). Constraints from our MCMC analysis are also reported for comparison (green data points). Error bars represent 68% uncertainties for all data points.

Appendix H
Interpolation Uncertainties

In this appendix we describe the test performed to verify the interpolation scheme implemented in the MCMC analysis. For this we remove one model from the interpolation grid (in the example below we excluded the model S10_0.55z9 of Table 2) and test how well the thermal parameters for this model are recovered when it is used to generate artificial data.

Figure 33 displays the correct values (red squares) overlaid on the parameter constraints obtained from the MCMC analysis at $z = 5.0$. The thermal parameters T_0 and u_0 are recovered accurately by the analysis, with discrepancies $\lesssim$5%. The

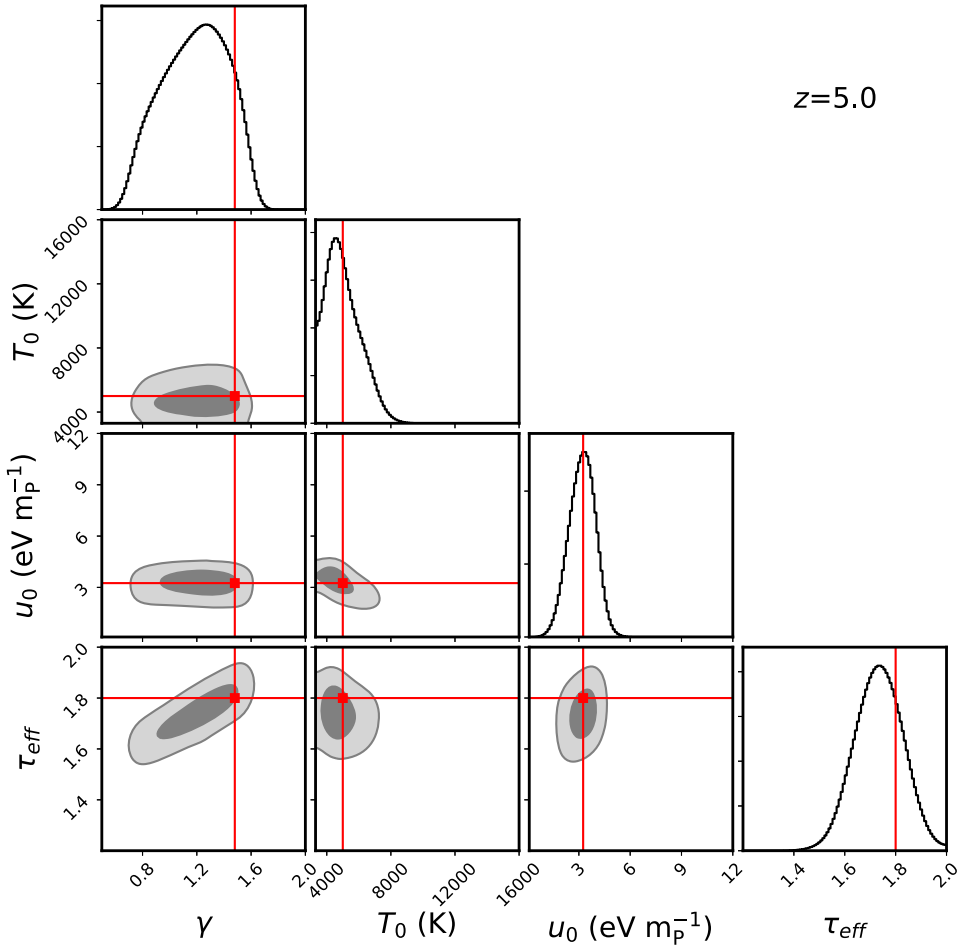

Figure 33. Probability distributions for the parameters T_0, u_0, γ, and $\tau_{\rm eff}$ at $z = 5$, recovered when fitting the synthetic power spectrum generated from the model S10_0.55z9. For this test the model has been removed from the set of comparison simulations. Contours plots show the 68% and 95% 2D probability distributions, while the black histograms display the 1D marginalized posterior distributions for each parameter. The input model parameters are presented for comparison (red squares).

correct values of γ and $\tau_{\rm eff}$ fall within the 1σ probability distribution, although the peaks of their posterior distributions are somewhat biased toward lower values. As expected, the power spectrum at these redshifts is not sensitive enough to break the degeneracy between γ and $\tau_{\rm eff}$. The poor constraints on γ, however, do not affect our constraints on T_0 and u_0.

We have also tested how well our interpolation scheme was able to recover the thermal parameters of a completely independent model. We fit the power spectrum extracted from a simulation using the UVB model of Puchwein et al. (2019) assuming nonequilibrium ionization, and we verified that the values of T_0 and u_0 (see Figure 17) were recovered within the 68% uncertainties given by our MCMC method. This gives some reassurance that, as intended, our results do not significantly depend on the specific thermal histories adopted for the modeling in this work.

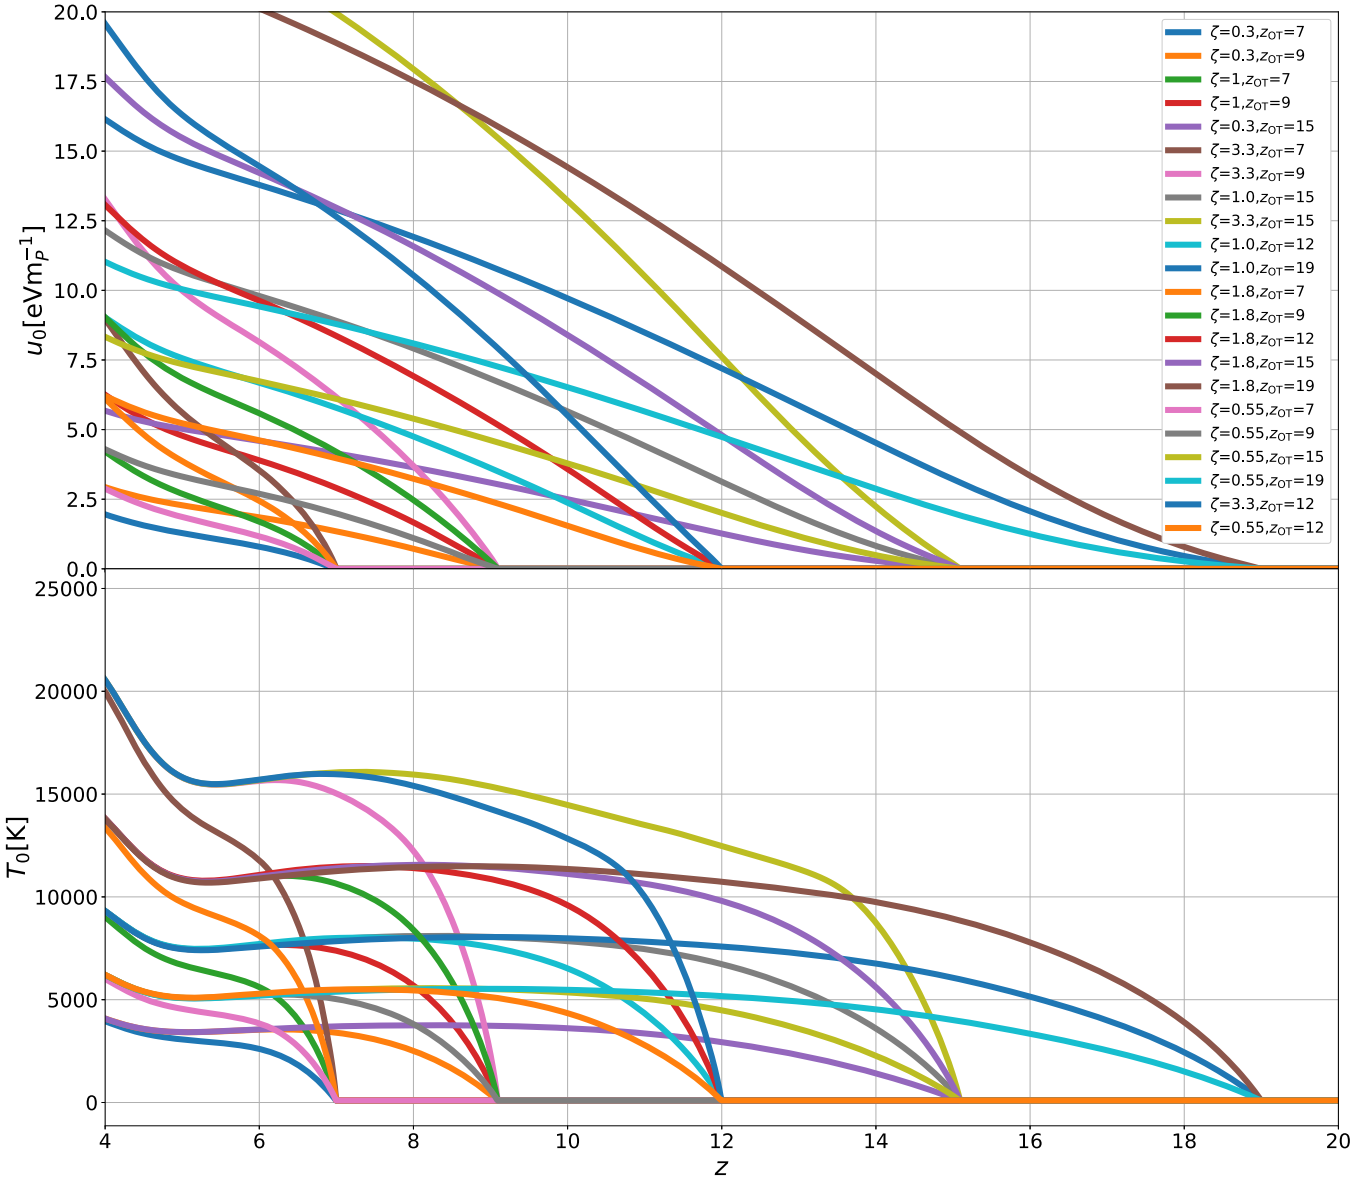

Figure 34. Evolution of the parameters governing the thermal state of the IGM in all of the simulations used for this analysis. Top panel: evolution as a function of redshift of the cumulative energy per unit mass, u_0, for models with a wide range of optically thin redshifts (z_{OT}) and photoheating rates ($\propto\zeta$). Bottom panel: corresponding evolution of the gas temperature at the mean density, T_0.

Appendix I
Thermal History Overview

For illustrative purposes we show in Figure 34 the evolution of the thermal parameters u_0 and T_0 for all the simulations listed in Table 2. While these models are not meant to represent realistic reionization scenarios, they provide a wide range of thermal histories and can be used to explore the thermal state of the IGM in a relatively model-independent way.

Appendix J
Integrated Heating versus Real-space Flux Cutoff Scale

In this appendix we show the relationship between u_0 and the characteristic real-space flux power cutoff scale, λ_P, as defined by Kulkarni et al. (2015). At each redshift we compute λ_P for all the models of Table 2 following the method described in Kulkarni et al. We then fit a relationship between the corresponding u_0 computed over the fiducial redshift range. Figure 35 shows the best-fitting relationship between u_0 and λ_P for the redshifts considered in this work. While a certain level of scatter about the fit is present at all redshifts, there is clearly a positive correlation between the two variables. Using the current constraints on u_0, we can then attempt to obtain a rough estimate of λ_P (green square with error bars).

We note that these estimates for λ_P are smaller than the recent constraints at $2 < z < 4$ from quasar pairs presented by Rorai et al. (2017a), suggesting that the pressure smoothing scale increases toward lower redshift as the IGM is photoheated

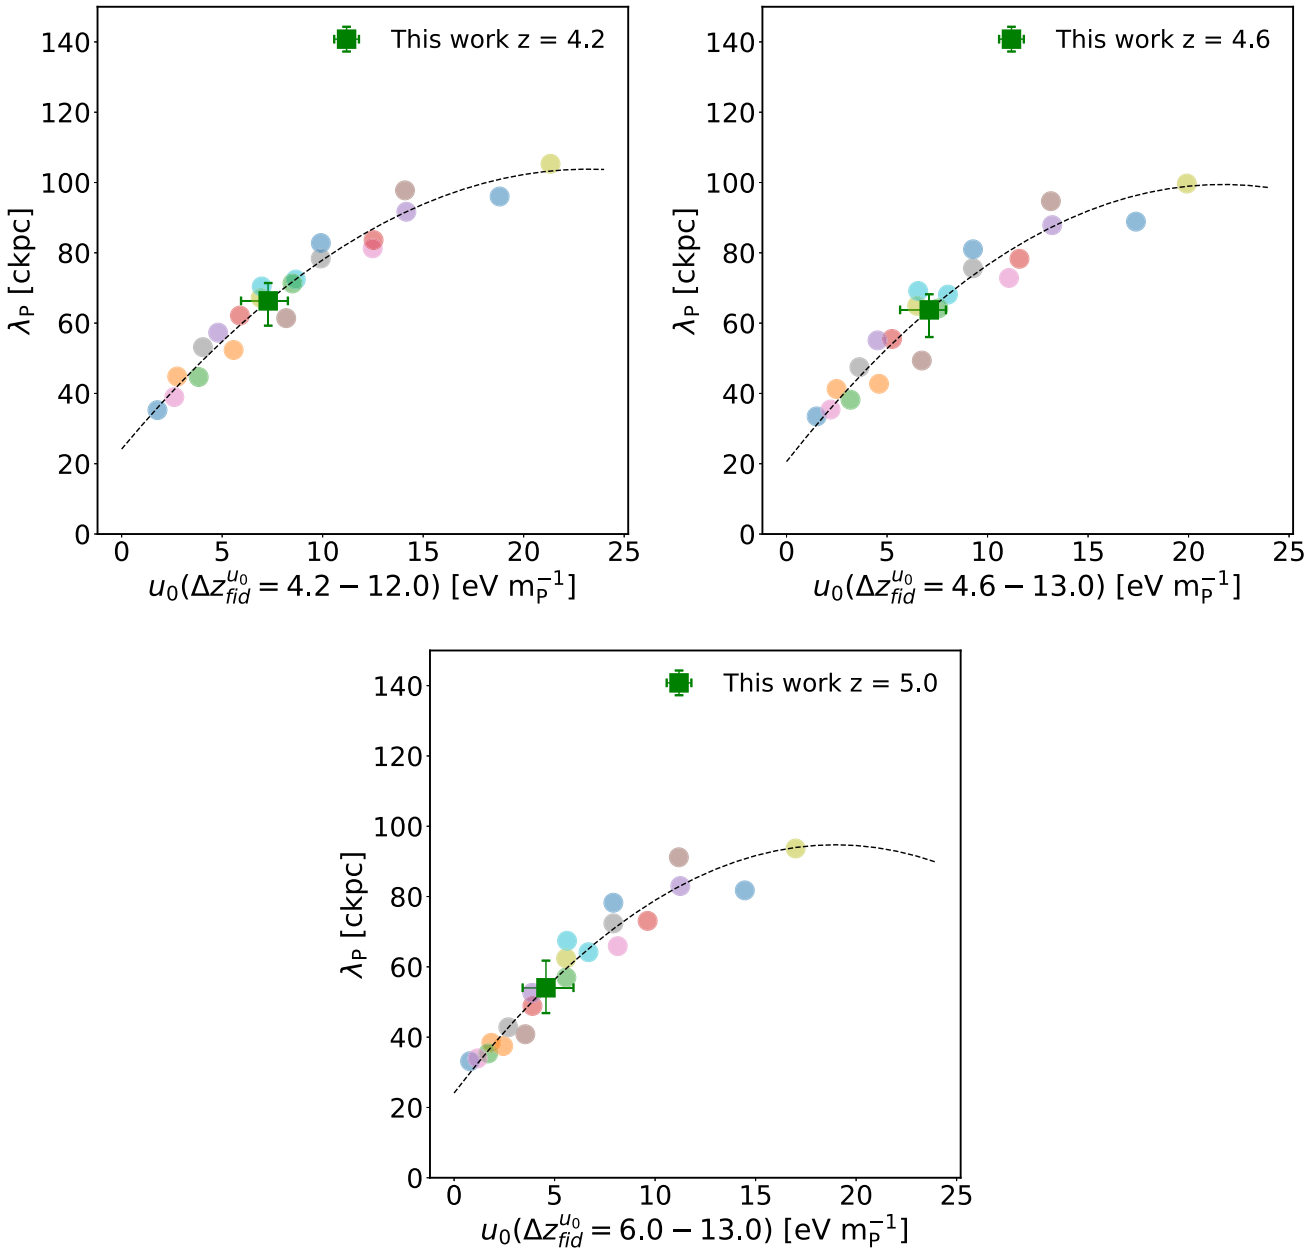

Figure 35. Relationship between the integrated heating per unit mass u_0 and the real-space flux power cutoff scale λ_P of Kulkarni et al. (2015). Colored points correspond to the different simulations of Table 2. While some scatter about the fit (black dashed line) is always present, there is a significant positive correlation between the two variables. For reference, along the fit at each redshift we plot our value of u_0 with the corresponding value of λ_P (green squares with 68% errors).

further (see the instantaneous temperature measurements presented by Becker et al. 2011)

Appendix K
Lyα Flux Power Spectrum Measurements

In Tables 6–8 we report the power spectrum measurements obtained in this work for the three redshift bins centered at $z = 4.2$, 4.6, and 5.0. In each table the values of the power spectrum obtained with (Column (3)) and without (Column (2)) instrumental resolution and pixel size correction (R.C.) are reported for each scale (Column (1)). The corresponding 68% uncertainties are shown in Column (4). The covariance matrices for the power spectrum measurements may be found in the online version of this article.

Table 6
Power Spectrum Measurement for the Redshift Bin Centered at $z = 4.2$

$\log(k/\mathrm{km}^{-1}\ \mathrm{s})$	$P_k(\mathrm{km\ s}^{-1})$	$P_k(\mathrm{km\ s}^{-1})$(R.C.)	σ
−2.2	91.4065	91.4324	27.0528
−2.1	82.4448	82.4819	17.3864
−2.0	70.2289	70.2789	19.4023
−1.9	74.6290	74.7128	17.9354
−1.8	56.1625	56.2625	10.6905
−1.7	43.7497	43.8733	7.6836
−1.6	30.6775	30.8155	6.1441
−1.5	27.2371	27.4304	4.1755
−1.4	21.1838	21.4225	2.9169
−1.3	14.3394	14.5968	2.3634
−1.2	7.9927	8.2213	0.9394
−1.1	4.2090	4.4020	0.4863
−1.0	2.0377	2.1891	0.2347
−0.9	0.8415	0.9444	0.1039
−0.8	0.3525	0.4241	0.0493
−0.7	0.1638	0.2208	0.0281

Note. Values in Column (3) have been corrected for instrumental resolution. The reported values have been obtained from the analysis of 51 sections of 20 h^{-1} cMpc of Lyα forest with $z\ \epsilon$ [4.0, 4.4), extracted from a total of 12 quasar lines of sight. The mean redshift for this bin is $\bar{z} = 4.24$.

Table 7
Power Spectrum Measurement for the Redshift Bin Centered at $z = 4.6$

$\log(k/\mathrm{km}^{-1}\ \mathrm{s})$	$P_k(\mathrm{km\ s}^{-1})$	$P_k(\mathrm{km\ s}^{-1})$(R.C.)	σ
−2.2	128.8440	128.8804	15.5333
−2.1	111.7963	111.8463	13.5236
−2.0	91.4603	91.5253	10.6477
−1.9	94.5054	94.6114	11.2976
−1.8	74.2880	74.4201	7.6046
−1.7	64.1286	64.3093	6.5800
−1.6	53.5776	53.8172	5.3659
−1.5	40.5199	40.8068	3.6843
−1.4	28.5061	28.8268	3.0231
−1.3	21.0848	21.4623	1.6939
−1.2	12.1394	12.4857	0.8407
−1.1	6.0252	6.3006	0.4470
−1.0	3.1159	3.3457	0.2371
−0.9	1.4523	1.6273	0.1078
−0.8	0.5897	0.7071	0.0584
−0.7	0.2588	0.3465	0.0342

Note. Values in Column (3) have been corrected for instrumental resolution. The reported values have been obtained from the analysis of 114 sections of 20 h^{-1} cMpc of Lyα forest with $z\ \epsilon$ [4.4, 4.8), extracted from a total of 15 quasar lines of sight. The mean redshift for this bin is $\bar{z} = 4.58$.

Table 8
Power Spectrum Measurement for the Redshift Bin Centered at $z = 5.0$

$\log(k/\mathrm{km}^{-1}\ \mathrm{s})$	$P_k(\mathrm{km\ s}^{-1})$	$P_k(\mathrm{km\ s}^{-1})$(R.C.)	σ
−2.2	162.4708	162.5166	30.6076
−2.1	163.8056	163.8787	28.1692
−2.0	157.6143	157.7257	25.4039
−1.9	121.8037	121.9401	23.5215
−1.8	81.6827	81.8278	11.8521
−1.7	87.0669	87.3118	12.4211
−1.6	72.5758	72.8998	10.3949
−1.5	52.7895	53.1629	7.4238
−1.4	37.5451	37.9670	4.5890
−1.3	26.6637	27.1404	2.6700
−1.2	15.8800	16.3323	1.4250
−1.1	10.2757	10.7436	1.0259
−1.0	4.6382	4.9787	0.5013
−0.9	2.1834	2.4443	0.2166
−0.8	1.0184	1.2185	0.1182
−0.7	0.3671	0.4897	0.0613

Note. Values in Column (3) have been corrected for instrumental resolution. The reported values have been obtained from the analysis of 44 sections of 20 h^{-1} cMpc of Lyα forest with $z\ \epsilon$ [4.8, 5.2), extracted from a total of 12 quasar lines of sight. The mean redshift for this bin is $\bar{z} = 4.95$.

References

Abel, T., & Haehnelt, M. G. 1999, ApJL, 520, L13
Becker, G. D., Bolton, J. S., Haehnelt, M. G., & Sargent, W. L. W. 2011, MNRAS, 410, 1096
Becker, G. D., Bolton, J. S., & Lidz, A. 2015a, PASA, 32, e045
Becker, G. D., Bolton, J. S., Madau, P., et al. 2015b, MNRAS, 447, 3402
Becker, G. D., Hewett, P. C., Worseck, G., & Prochaska, J. X. 2013, MNRAS, 430, 2067
Boera, E., Murphy, M. T., Becker, G. D., & Bolton, J. S. 2014, MNRAS, 441, 1916
Boera, E., Murphy, M. T., Becker, G. D., & Bolton, J. S. 2016, MNRAS, 456, L79
Bolton, J. S., & Becker, G. D. 2009, MNRAS, 398, L26
Bolton, J. S., Becker, G. D., Haehnelt, M. G., & Viel, M. 2014, MNRAS, 438, 2499
Bolton, J. S., & Haehnelt, M. G. 2007, MNRAS, 382, 325
Bolton, J. S., Puchwein, E., Sijacki, D., et al. 2017, MNRAS, 464, 897
Bolton, J. S., Viel, M., Kim, T.-S., Haehnelt, M. G., & Carswell, R. F. 2008, MNRAS, 386, 1131
Bosman, S. E. I., Fan, X., Jiang, L., et al. 2018, MNRAS, 479, 1055
Bouwens, R. J., Illingworth, G. D., Oesch, P. A., et al. 2015, ApJ, 811, 140
Bouwens, R. J., Smit, R., Labbé, I., et al. 2016, ApJ, 831, 176
Caruana, J., Bunker, A. J., Wilkins, S. M., et al. 2014, MNRAS, 443, 2831
D'Aloisio, A., McQuinn, M., Maupin, O., et al. 2018, ApJ, submitted (arXiv:1807.09282)
D'Aloisio, A., Upton Sanderbeck, P. R., McQuinn, M., Trac, H., & Shapiro, P. R. 2017, MNRAS, 468, 4691

Davies, F. B., Hennawi, J. F., Bañados, E., et al. 2018, ApJ, 864, 142
Dekker, H., D'Odorico, S., Kaufer, A., Delabre, B., & Kotzlowski, H. 2000, Proc. SPIE, 4008, 534
Eilers, A.-C., Davies, F. B., & Hennawi, J. F. 2018, ApJ, 864, 53
Fan, Z., Ma, J., de Grijs, R., Yang, Y., & Zhou, X. 2006, MNRAS, 371, 1648
Finkelstein, S. L. 2016, PASA, 33, e037
Garzilli, A., Boyarsky, A., & Ruchayskiy, O. 2017, PhLB, 773, 258
Giallongo, E., Grazian, A., Fiore, F., et al. 2015, A&A, 578, A83
Gnedin, N. Y., & Hui, L. 1998, MNRAS, 296, 44
Greig, B., Mesinger, A., Haiman, Z., & Simcoe, R. A. 2017, MNRAS, 466, 4239
Haardt, F., & Madau, P. 2012, ApJ, 746, 125
Hiss, H., Walther, M., Hennawi, J. F., et al. 2018, ApJ, 865, 42
Horne, K. 1986, PASP, 98, 609
Hui, L., & Gnedin, N. Y. 1997, MNRAS, 292, 27
Iršič, V., Slosar, A., Bailey, S., et al. 2013, JCAP, 9, 016
Iršič, V., Viel, M., Berg, T. A. M., et al. 2017a, MNRAS, 466, 4332
Iršič, V., Viel, M., Haehnelt, M. G., et al. 2017b, PhRvD, 96, 023522
Jones, A., Noll, S., Kausch, W., Szyszka, C., & Kimeswenger, S. 2013, A&A, 560, A91
Kelson, D. D. 2003, PASP, 115, 688
Kim, T.-S., Viel, M., Haehnelt, M. G., Carswell, R. F., & Cristiani, S. 2004, MNRAS, 347, 355
Kuhlen, M., & Faucher-Giguère, C.-A. 2012, MNRAS, 423, 862
Kulkarni, G., Hennawi, J. F., Oñorbe, J., Rorai, A., & Springel, V. 2015, ApJ, 812, 30
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
Lidz, A., Heitmann, K., Hui, L., et al. 2006, ApJ, 638, 27
Lomb, N. R. 1976, Ap&SS, 39, 447
López, S., D'Odorico, V., Ellison, S. L., et al. 2016, A&A, 594, A91
Madau, P., & Haardt, F. 2015, ApJL, 813, L8
Mason, C. A., Treu, T., Dijkstra, M., et al. 2018, ApJ, 856, 2
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., et al. 2005, ApJ, 635, 761
McGreer, I. D., Mesinger, A., & D'Odorico, V. 2015, MNRAS, 447, 499
McQuinn, M. 2012, MNRAS, 426, 1349
McQuinn, M., Lidz, A., Zaldarriaga, M., et al. 2009, ApJ, 694, 842
McQuinn, M., & Upton Sanderbeck, P. R. 2016, MNRAS, 456, 47
Miralda-Escudé, J., & Rees, M. J. 1994, MNRAS, 266, 343
Mortlock, D. J., Warren, S. J., Venemans, B. P., et al. 2011, Natur, 474, 616
Murphy, M. T., Kacprzak, G. G., Savorgnan, G. A. D., & Carswell, R. F. 2019, MNRAS, 482, 3458
Nasir, F., Bolton, J. S., & Becker, G. D. 2016, MNRAS, 463, 2335
Noll, S., Kausch, W., Barden, M., et al. 2012, A&A, 543, A92
Oñorbe, J., Hennawi, J. F., Lukić, Z., & Walther, M. 2017, ApJ, 847, 63
Palanque-Delabrouille, N., Yèche, C., Baur, J., et al. 2015, JCAP, 11, 011
Palanque-Delabrouille, N., Yèche, C., Borde, A., et al. 2013, A&A, 559, A85
Parsa, S., Dunlop, J. S., & McLure, R. J. 2018, MNRAS, 474, 2904
Planck Collaboration et al. 2014, A&A, 571, A16
Planck Collaboration et al. 2018, arXiv:1807.06209
Puchwein, E., Bolton, J. S., Haehnelt, M. G., et al. 2015, MNRAS, 450, 4081
Puchwein, E., Haardt, F., Haehnelt, M. G., & Madau, P. 2019, MNRAS, in press (doi:10.1093/mnras/stz222)
Raskutti, S., Bolton, J. S., Wyithe, J. S. B., & Becker, G. D. 2012, MNRAS, 421, 1969
Ricotti, M., Gnedin, N. Y., & Shull, J. M. 2000, ApJ, 534, 41
Rorai, A., Becker, G. D., Haehnelt, M. G., et al. 2017a, MNRAS, 466, 2690
Rorai, A., Carswell, R. F., Haehnelt, M. G., et al. 2018, MNRAS, 474, 2871
Rorai, A., Hennawi, J. F., Oñorbe, J. F., et al. 2017b, Sci, 356, 418
Rorai, A., Hennawi, J. F., & White, M. 2013, ApJ, 775, 81
Rudie, G. C., Steidel, C. C., & Pettini, M. 2012, ApJL, 757, L30
Sadoun, R., Zheng, Z., & Miralda-Escudé, J. 2017, ApJ, 839, 44
Scargle, J. D. 1982, ApJ, 263, 835
Schaye, J., Theuns, T., Rauch, M., Efstathiou, G., & Sargent, W. L. W. 2000, MNRAS, 318, 817
Schmidt, K. B., Treu, T., Bradač, M., et al. 2016, ApJ, 818, 38
Scott, J., Bechtold, J., Dobrzycki, A., & Kulkarni, V. P. 2000, ApJS, 130, 67
Shull, J. M., Stevans, M., & Danforth, C. W. 2012, ApJ, 752, 162
Simcoe, R. A., Sullivan, P. W., Cooksey, K. L., et al. 2012, Natur, 492, 79
Songaila, A., & Cowie, L. L. 2010, ApJ, 721, 1448
Springel, V. 2005, MNRAS, 364, 1105
Syphers, D., & Shull, J. M. 2014, ApJ, 784, 42
Telfer, R. C., Zheng, W., Kriss, G. A., & Davidsen, A. F. 2002, ApJ, 565, 773
Theuns, T., Zaroubi, S., Kim, T.-S., Tzanavaris, P., & Carswell, R. F. 2002, MNRAS, 332, 367
Upton Sanderbeck, P. R., D'Aloisio, A., & McQuinn, M. J. 2016, MNRAS, 460, 1885
Viel, M., Becker, G. D., Bolton, J. S., & Haehnelt, M. G. 2013a, PhRvD, 88, 043502
Viel, M., Haehnelt, M. G., & Springel, V. 2004, MNRAS, 354, 684
Viel, M., Schaye, J., & Booth, C. M. 2013b, MNRAS, 429, 1734
Vogt, S. S., Allen, S. L., Bigelow, B. C., et al. 1994, Proc. SPIE, 2198, 362
Walther, M., Hennawi, J. F., Hiss, H., et al. 2018, ApJ, 852, 22
Walther, M., Oñorbe, J., Hennawi, J. F., & Lukić, Z. 2019, ApJ, in press (arXiv:1808.04367)
Worseck, G., Prochaska, J. X., McQuinn, M., et al. 2011, ApJL, 733, L24
Worseck, G., & Wisotzki, L. 2006, A&A, 450, 495
Yèche, C., Palanque-Delabrouille, N., Baur, J., & du Mas des Bourboux, H. 2017, JCAP, 6, 047