

PAPER

Cosmological constraints from the eBOSS Lyman- α forest using the PRIYA simulations

To cite this article: M.A. Fernandez *et al* JCAP07(2024)029

View the [article online](#) for updates and enhancements.

You may also like

- [Estimation of the Lyman- \$\alpha\$ signal of the EFILE diagnostic under static or radiofrequency electric field in vacuum](#)
Carlo POGGI, Théo GUILLAUME, Fabrice DOVEIL et al.

- [Quantum state preparation of hydrogen atoms by hyperfine quenching](#)
Amanda Alencar, I Prazeres, C R de Carvalho et al.

- [High-precision x-ray spectroscopy of highly charged ions with microcalorimeters](#)
S Kraft-Bermuth, V Andrianov, A Bleile et al.

RECEIVED: September 11, 2023

REVISED: May 29, 2024

ACCEPTED: June 27, 2024

PUBLISHED: July 17, 2024

Cosmological constraints from the eBOSS Lyman- α forest using the PRIYA simulations

M.A. Fernandez , Simeon Bird  and Ming-Feng Ho 

*Department of Physics and Astronomy, University of California Riverside,
900 University Ave, Riverside, CA 92521, U.S.A.*

E-mail: mfern027@ucr.edu, sbird@ucr.edu, mho026@ucr.edu

ABSTRACT: We present new cosmological parameter constraints from the eBOSS Lyman- α forest survey. We use a new theoretical model and likelihood based on the PRIYA simulation suite. PRIYA is the first suite to resolve the Lyman- α forest in a $(120 \text{ Mpc}/h)^3$ volume, using a multi-fidelity emulation technique. We use PRIYA to predict Lyman- α forest observables with $\lesssim 1\%$ interpolation error over an 11 dimensional (9 simulated, 2 in post-processing) parameter space. We identify an internal tension within the flux power spectrum data. Once the discrepant data is removed, we find the primeval scalar spectral index measured at a pivot scale of $k_0 = 0.78 \text{ Mpc}^{-1}$ to be $n_P = 1.009_{-0.018}^{+0.027}$ at 68% confidence. This measurement from the Lyman- α forest flux power spectrum alone is in reasonable agreement with Planck, and in tension with earlier eBOSS analyses. The amplitude of matter fluctuations is $\sigma_8 = 0.733_{-0.029}^{+0.026}$ at 68% confidence, in agreement with Dark Energy Survey weak lensing measurements and other small-scale structure probes and in tension with CMB measurements from Planck and ACT. The effective optical depth to Lyman- α photons from our pipeline is in good agreement with earlier high resolution measurements. We find a linear power at $z = 3$ and $k = 0.009 \text{ s/km}$ of $\Delta_L^2 = 0.302_{-0.027}^{+0.024}$ with a slope $n_{\text{eff}} = -2.264_{-0.018}^{+0.026}$. Our flux power spectrum only chains prefer a low level of heating during helium reionization. When we add IGM temperature data we find $n_P = 0.983 \pm 0.020$ and $\sigma_8 = 0.703_{-0.027}^{+0.023}$. Our chains prefer an early and long helium reionization event, as suggested by measurements from the helium Lyman- α forest. In the near future we will use our pipeline to infer cosmological parameters from the DESI Lyman- α data.

KEYWORDS: intergalactic media, Lyman alpha forest, Machine learning , cosmological simulations

ARXIV EPRINT: [2309.03943](https://arxiv.org/abs/2309.03943)

Contents

1	Introduction	1
2	Simulation suite and emulator	3
2.1	Cosmological & astrophysical parameters	4
2.2	Summary statistics: flux power and IGM temperature	5
2.3	Gaussian Process emulators	6
3	Inference scheme and likelihood function	7
3.1	Flux power spectrum data	8
3.2	IGM temperature at mean density data	8
3.3	Covariance matrix	9
3.4	Likelihood	11
3.5	Inference using simulation data	11
4	Results	14
4.1	Cosmological parameters	14
4.2	Reionization and other parameters	19
4.3	Parameter correlations	22
5	Discussion	23
5.1	The tension in the lowest redshift bins	23
5.2	Comparison of the posterior constraints to other Lyman- α analyses	25
5.3	Likelihood modifications	27
6	Conclusions	27
A	Leave-one-out versus emulator error	29
B	BOSS DR9 data	30
C	Posterior constraints from the IGM temperature alone	32
D	Full posteriors	33

1 Introduction

The Lyman- α forest [1–9] measures the distribution of neutral gas at relatively low densities. This gas traces the growth of cosmic structure, making the Lyman- α forest an exceptionally powerful cosmological probe, sensitive to the distribution of dark matter deep in the matter dominated era. Correlating absorption from different quasar sightlines has allowed detection of the baryon acoustic oscillations and constraints on the expansion of the universe [10–13]. The densities probed by the Lyman- α forest from redshift $z = 2\text{--}5$ are $\sim 1\text{--}100 \times$ the

cosmological mean density. For these redshifts and densities stellar winds and star formation effects are negligible, though feedback from black holes can be important [14, 15]. Thus the Lyman- α forest is able to measure the primordial fluctuations on some of the smallest scales available, $k \sim 1 \text{ h/Mpc}$ [16–23]. In addition, the Lyman- α forest is sensitive to the thermal and ionization history of the intergalactic medium (IGM) [24–31], and by constraining the smallest structures, the mass scale of thermal relic dark matter [22, 32–37].

The extended Baryon Oscillation Sky Survey (eBOSS), part of the Sloan Digital Sky Survey (SDSS) [38], has computed the 1D flux power spectrum along quasar sight lines for over 43,000 quasars, with a statistical error $\sim 1\%$ at some redshifts. This exceptional statistical error means that the error budget is dominated by systematic uncertainty, especially uncertainty in the resolution of the spectrograph on small scales [38]. The Dark Energy Spectroscopic Instrument (DESI) has improved the spectrograph resolution by a factor of two [39]. Thus, early data from DESI has measured the flux power spectrum at smaller scales ($k \gtrsim 0.035 \text{ km}^{-1} \text{ s}$) than SDSS [40, 41]. Future releases will measure higher redshifts ($z > 4.6$) and increase the number of Lyman- α forest quasar spectra by a factor of four over SDSS [42].

There are other high resolution, small sample datasets of quasar spectra, from which Lyman- α forest flux power measurements have been made [43–46]. Ref. [45] used spectra from multiple surveys (XQ-100, KODIAQ, and SQUAD) to measure the Lyman- α forest flux power at redshifts $z = 2\text{--}4.6$ and scales $k \approx 0.005\text{--}0.1 \text{ km}^{-1} \text{ s}$ (albeit with larger uncertainty than eBOSS), and ref. [46] presents recent cosmological constraints from these datasets.

Modeling the Lyman- α forest requires numerical simulations that are able to follow the distribution of gas on small scales. In this paper we present cosmological parameter inference using a new likelihood built on the PRIYA simulation suite [47]. The PRIYA simulations are in 120 Mpc/h boxes, and are comprised of 60 simulations with 2×1536^3 particles (mean inter-particle spacing of 78 kpc/h), as well as 3 simulations with 2×3072^3 particles (mean inter-particle spacing of 39 kpc/h). The higher of these two resolutions exceeds the resolution of state-of-the-art galaxy formation simulations such as Illustris-TNG [48]. PRIYA is run with the same highly scalable MP-Gadget code as the ASTRID simulation [49, 50]. PRIYA contains full hydrodynamic simulations with models of galaxy formation and black hole feedback to $z = 2.2$. PRIYA is thus the first cosmological simulation suite which achieves, in a single box, the required box size of 120 Mpc/h , capable of minimising sample variance in the Lyman- α forest [51], and a resolution high enough that it includes the pressure smoothing scale.¹ Importantly, this removes the need for the ‘splicing’ correction used in earlier work to combine different boxesizes into a single whole [22, 51].

Here, the PRIYA simulations are used to build multi-fidelity emulators [52–54] for the flux power spectrum and the IGM temperature at mean density. Each emulator is a surrogate model, able to reproduce the 1D flux power spectrum or mean IGM temperature for cosmological parameters (within the prior simulation volume) to $\sim 1\%$ accuracy. A multi-fidelity emulator combines two different resolution training samples. Many low fidelity samples are used to explore parameter space, and their output is corrected with a few high fidelity samples. A multi-fidelity emulator makes predictions for the highest resolution simulation at a fraction of the computational cost of a single fidelity emulator [53, 55]. Emulators

¹Our model boosts gas temperature during reionization, increasing the smoothing scale [47].

Simulation	Box Volume	N_{gas}	$M_{\text{gas}} (M_{\odot} \text{ h}^{-1})$
LF	$(120 \text{ Mpc h}^{-1})^3$	1536^3	$[5.29, 6.98] \times 10^6$
HF	$(120 \text{ Mpc h}^{-1})^3$	3072^3	$[6.73, 7.97] \times 10^5$

Table 1. Low-Fidelity (LF) and High-Fidelity (HF) simulation suite details. N_{gas} is the number of gas particles simulated, M_{gas} is the resulting mass resolution of those particles.

have been used to study various cosmological probes: the matter power spectrum [56–62], weak lensing shear [63, 64], the halo mass function [65–67], the 21-cm signal [68–71] and the Lyman- α forest [36, 52, 72–76].

Here, we present the first fully resolved multi-fidelity emulator based likelihood framework for the eBOSS Lyman- α forest and the first cosmological constraints derived from it. Our multi-fidelity emulator is similar to that described in ref. [54], but the simulation volume has been increased by a factor of 64, and the spatial resolution has been improved by a factor of 1.5. We also use mean IGM temperature data [30] to constrain the parameters of helium reionization, data which is ultimately derived from higher resolution quasar surveys [43–45].

In summary, our method is: (1) Construct an emulator for the 1D Lyman- α flux power spectrum and mean IGM temperature using the PRIYA simulations [47], section 2. (2) Augment observational errors with estimates of the residual theoretical uncertainty to build a covariance matrix, and correct the flux power spectra for metal contamination as described in section 3. (3) Use this emulator and likelihood to constrain cosmological parameters using Markov Chain Monte Carlo (MCMC), with results described in section 4. We discuss some caveats and compare to earlier work in section 5 and our conclusions are presented in section 6.

MCMC chains for all the results presented in this work along with files containing the training outputs used to construct the emulators,² as well as the code,³ which includes the emulator, likelihood, and integration with the Cobaya MCMC package, are available publicly.

2 Simulation suite and emulator

In this section, we briefly describe the properties of the simulations and emulator, and refer the reader to ref. [47] for the full details. The emulator allows predictions for the output of a simulation at an arbitrary set of cosmological parameters within our prior volume with an average interpolation error of 0.2% at low fidelity and 1% at high fidelity. Our multi-fidelity emulator combines simulations at different resolutions, following the scheme outlined in ref. [54]. The emulator combines low fidelity (LF) and high fidelity (HF) simulations. Box volume, number of gas particles, and gas particle mass resolution are reported in table 1. We ran a total of 60 low fidelity (LF) and 3 high fidelity (HF) simulations. For this work we have added 12 new LF simulations to those of ref. [47], which extend the simulated parameter range to better cover the posterior range allowed by SDSS DR14. Low fidelity simulations have 1536^3 particles, while high fidelity simulations have 3072^3 particles. Sampled parameters are chosen to maximise spread in parameter space, as described in ref. [47].

²<https://github.com/mafern/InferenceLyaData>.

³https://github.com/sbird/lya_emulator.

Our simulations still have limited resolution and a finite box. Ref. [47] showed that on these scales resolution is important at the 1% level. Compared to the literature, our resolution convergence is slightly better, due to the temperature boost we impart during H I reionization [77]. The finite box scatters the modes of the 1D flux power spectrum at the 2% level on large scales, due mostly to the limited number of helium reionization bubbles (30 Mpc/h across) that can fit into our volume. Section 3.3 describes how we attempt to marginalise out residual cosmic variance.

The range given for the gas mass resolution is due to the varying value of h in our simulation suite ($\Omega_b h^2$ is fixed at a value of 0.0224). We show in ref. [47] that this gas mass is sufficient for the scales and redshifts probed by the eBOSS flux power spectrum. Our simulations include a full galaxy physics model with star formation, stellar and AGN feedback and inhomogeneous reionization models. Simulations were performed using MP-Gadget,⁴ an N-body and smoothed particle hydrodynamics (SPH) code. MP-Gadget uses the gravitational timestepping algorithm from Gadget-4 [78], and various other algorithmic improvements [79]. Simulations are initialised at $z = 99$ and finish at $z = 2.2$. The galaxy formation model is similar to the ASTRID simulation [49, 50] and is described fully in ref. [47].

2.1 Cosmological & astrophysical parameters

Table 2 summarises the parameters that are varied across our suite of simulations, as well as their limits. We have expanded the limits from ref. [47] with 12 additional LF simulations covering the parameter ranges $\alpha_q = 2.5\text{--}3.0$ and $n_P = 1.0\text{--}1.05$. This was done so that the emulator covers the 2σ posterior range for n_P . Simulated parameters were chosen using a Latin hypercube. We model the primeval (that is, the pre-transfer function) power spectrum $P(k)$ using two parameters: a slope, n_P , and an amplitude, A_P :

$$P(k) = A_P \left(\frac{k}{0.78 \text{ Mpc}^{-1}} \right)^{n_P - 1}. \quad (2.1)$$

The simulation initial conditions are then generated using a set of transfer functions from CLASS. This parameterization is the same as that used by Planck, but with a different pivot scale, 0.78 Mpc^{-1} , rather than 0.05 Mpc^{-1} , reflecting the smaller scales probed by the forest.

We also vary the Hubble parameter h , and the total matter density through $\Omega_M h^2$, although we will see these are not strongly constrained by the Lyman- α forest. We add three parameters for the He II reionization model [80]: z_{Hei} and z_{Hef} are the redshifts for the start and end of He II reionization, and α_q is the quasar spectral index (which scales the peak temperature during He II reionization). z_{Hi} is the midpoint redshift of H I reionization. Finally, ϵ_{AGN} is the black hole feedback factor, to which the Lyman- α forest is insensitive.

There are two further parameters for the Lyman- α effective optical depth, varied by post-processing the artificial spectra.⁵ We parameterize the mean flux $\mathcal{F} = \exp(-\tau)$ by

⁴<https://github.com/MP-Gadget/MP-Gadget>.

⁵We have freedom to vary the Lyman- α mean flux as it is degenerate with the amplitude of the ultraviolet background (UVB).

Parameter	Minimum	Maximum	Description
n_P	0.8	1.05	Scalar spectral index
A_P	1.2×10^{-9}	2.6×10^{-9}	Power amplitude at $k = 0.78 \text{ Mpc}^{-1}$
h	0.65	0.75	Hubble parameter
$\Omega_0 h^2$	0.14	0.146	Total matter density
z_{Hei}	3.5	4.1	Start redshift of HeII reionization
z_{Hef}	2.6	3.2	End redshift of HeII reionization
α_q	1.3	3.0	Quasar spectral index during HeII reionization
z_{Hi}	6.5	8	Median redshift of HI reionization
ϵ_{AGN}	0.03	0.07	Thermal efficiency of black hole feedback
τ_0	0.75	1.25	Mean optical depth at $z = 3$ in eq. (2.3).
$d\tau_0$	-0.4	0.25	Mean optical depth redshift evolution in eq. (2.3).

Table 2. Summary of likelihood function parameters, together with the ranges covered by the emulator. We vary a total of 11 parameters: 4 for cosmology, 3 for helium reionization, 1 for hydrogen reionization, 1 for the strength of AGN feedback and 2 for the mean optical depth.

modifying the power law redshift evolution from ref. [81], as

$$\tau_{\text{H I}}^{\text{Kim}}(z) = 0.0023 \times (1 + z)^{3.65}, \quad (2.2)$$

$$\tau_{\text{H I}}^{\text{eff}}(z) = \tau_0 \left(\frac{\tau_{\text{H I}}^{\text{Kim}}(z)}{\tau_{\text{H I}}^{\text{Kim}}(3)} \right)^{d\tau_0} \tau_{\text{H I}}^{\text{Kim}}(z). \quad (2.3)$$

The parameters varied are τ_0 and $d\tau_0$, with $(1, 0)$ corresponding to the redshift evolution of ref. [81]. τ_0 is normalised at $z = 3$ so that $d\tau_0 > 0$ corresponds to a higher optical depth at $z > 3$ and a lower optical depth at $z < 3$. τ_0 changes the normalisation of the optical depth at $z = 3$. We choose prior ranges for τ_0 and $d\tau_0$ that comfortably cover the measurement error: $0.75 < \tau_0 < 1.25$ and $-0.4 < d\tau_0 < 0.25$. As the mean flux is chosen in post-processing, we can dramatically over-sample these parameters. We sample 10 linearly spaced values of the mean flux. Sampling values are independent of redshift and are chosen so that they include the value of $\tau_0(z)$ implied by our priors at the extreme ($z = 2.2$ and $z = 4.6$) redshifts included in eBOSS. For concreteness, we generate flux power spectra scaling the mean optical depth in each redshift bin by a factor between 0.66 and 1.3. This implies that $\tau_{\text{H I}}^{\text{eff}}(z = 3)$ varies between 0.238 and 0.47 and $\tau_{\text{H I}}^{\text{eff}}(z = 4)$ varies between 0.54 and 1.06. We thus produce a total of 600 LF and 30 HF simulated flux power spectra in each redshift bin, ten times the number of simulations.

2.2 Summary statistics: flux power and IGM temperature

Figure 1 shows an example of the gas density and temperature (colors) at $z = 4$ for both high and low resolution in our simulations, demonstrating how spectra connect to the matter density field. We generate a total of $3 \times 480^2 = 691,200$ spectra from each snapshot of each simulation, from $z = 4.6$ to $z = 2.2$ in increments of $\Delta z = 0.2$, with a pixel resolution of 10 km

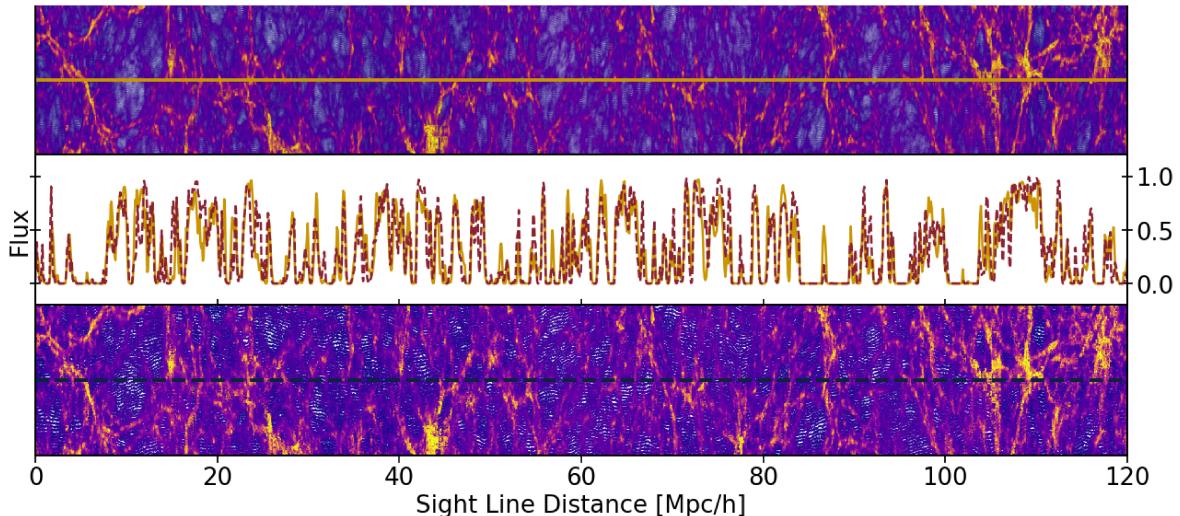


Figure 1. Example Lyman- α forest spectra and corresponding gas density and temperature (colors in top and bottom panels) from an LF and HF simulation at redshift $z = 4$. The top panel shows high resolution and the bottom panel shows low resolution. The middle panel shows the Lyman- α forest spectra for the skewers passing through the middle of the top panel (high resolution, yellow line) and bottom panel (low resolution, red dashed line).

s^{-1} . We generate Lyman- α forest absorption spectra using Fake Spectra Flux Extractor [82]⁶, described in ref. [83]. We compute the 1D flux power spectrum of the Lyman- α forest flux, averaged over all sightlines. The flux power is defined as

$$P_F(k) = |L^{-1}\tilde{\delta}_F^2(k)|. \quad (2.4)$$

$\tilde{\delta}_F^2(k)$ is the Fourier transform of the flux excess, $\delta_F(k) = F(k)/\langle F(k) \rangle - 1$, and L is the length of the sightline.

Our simulations contain a realistic population of DLAs, which we mask as in the observational pipeline. We extract the IGM temperatures at mean density directly from the simulation snapshots. First, the temperature and density for all the gas particles in the simulation are retrieved, then all particles that are within 5% of the critical density are retained. The median temperature of these retained particles is the IGM temperature at mean density. All of the Lyman- α forest flux power spectra, IGM temperatures, trained emulators, as well as select MCMC chains are available.⁷

2.3 Gaussian Process emulators

We use the Gaussian Process (GP) emulators described in refs. [47, 54] for the 1D flux power spectra and mean IGM temperature extracted from our simulations. The emulators interpolate over simulation outputs and make predictions for arbitrary parameter sets within the parameter limits shown in table 2. We use a multi-fidelity model, which allows simulations with different particle loads, and thus costs, to be combined together. Specifically, we combine simulations run at two different resolutions, high fidelity (HF) and low fidelity (LF) specified

⁶https://github.com/sbird/fake_spectra.

⁷<https://github.com/mafern/InferenceLyaData>.

in table 1. The multi-fidelity prediction for the HF outputs (shown here for the Lyman- α forest flux power spectrum, but equally valid for the IGM temperature) is given by a linear multi-fidelity model, at each redshift:

$$P_F^{\text{HF}}(k, \boldsymbol{\theta}|z) = \rho_z \cdot P_F^{\text{LF}}(k, \boldsymbol{\theta}|z) + \delta(k, \boldsymbol{\theta}|z), \quad (2.5)$$

where ρ_z is a constant parameter, and $\delta(k, \boldsymbol{\theta}|z)$ is a GP. We have simplified the model from ref. [54] by dropping the k dependence in ρ_z . All cosmology and scale dependence is thus in the additive GP $\delta(k, \boldsymbol{\theta}|z)$. We tested a GP emulator that included scale dependence in the ρ term, but found that it was harder to train and did not significantly improve the accuracy of the predictions. We train the GP emulators separately for each redshift. We implement our multi-fidelity models using Emukit [84].

Ref. [47] quantifies the accuracy of our emulator using a leave-one-out technique, in which one simulation is chosen as a ‘left out’ sample. A smaller emulator built excluding all samples from this simulation is used to predict the summary statistic for the left out sample. After computing leave-one-out interpolation errors for all potential test samples, we found on average 0.2% accuracy for the low-fidelity simulations and 1% for the high fidelity simulations. The last is likely a significant over-estimate of the actual error since the leave-one-out procedure in this case is missing 1/3 of the total training data. The interpolation errors from the mean IGM temperature emulator are significantly smaller than the 7% average uncertainty in mean IGM temperature measurements. Our likelihood function and emulator code is publicly available.⁸

3 Inference scheme and likelihood function

In this section, we describe the inference scheme and likelihood function by which our cosmological parameter constraints are derived from our emulator, the eBOSS flux power spectrum [38] and the mean IGM temperature. The overall inference scheme is:

1. Use the emulator to predict the flux power spectrum and IGM temperature at mean density for a set of input parameters (see table 2).
2. Calculate a likelihood comparing these predictions to their observational counterparts from eBOSS [38] and ref. [30].
3. Use Cobaya [85, 86] to run MCMC chains and compute posterior parameter constraints.

Section 3.1 discusses the flux power spectrum data, while section 3.2 discusses the IGM temperature data. We derive our covariance matrix in section 3.3. Details of the likelihood calculation used in the MCMC sampling are given in section 3.4. We validate our pipeline on simulated data in section 3.5.

⁸https://github.com/sbird/lya_emulator.

3.1 Flux power spectrum data

We use the observed Lyman- α forest flux power spectrum from [38], which is based on the Baryon Oscillation Spectroscopic Survey (BOSS) and extended-BOSS (eBOSS) quasar samples [87, 88]. In [38], the BOSS/eBOSS quasar samples are refined to remove spectra that have not been visually inspected, and to remove spectra with broad absorption lines. Sky lines and damped Lyman- α absorbers (DLAs) are masked. Our simulations include a realistic population of DLAs, which are masked in the same way.

The sample of Lyman- α forests from the set of remaining quasar spectra is then further refined based on cuts to the spectral resolution, signal-to-noise ratio, number of masked pixels, and forest length, with a final sample of about 43,000 spectra. The redshifts and scales covered by these observations set the redshift range and scales we use in our flux power spectrum emulator, namely $z = 2.2\text{--}4.6$ (redshift bin size of $\Delta z = 0.2$), and $k \approx 0.001\text{--}0.02 \text{ s/km}$ (over 35 linearly spaced bins, $\Delta k = 5.42 \times 10^{-4} \text{ s/km}$). Our emulator can easily be re-trained for the smaller scales probed by DESI. The uncertainty in the eBOSS 1D flux power varies with k and z , ranging from $> 10\%$ at $z > 4$ to $\sim 2\%$ at $z \leq 3$, and is often dominated by systematic uncertainty [38].

We apply correction terms to the Lyman- α forest flux power spectrum predicted by our emulator to model DLAs and metal contamination. We correct for DLAs using the template from ref. [89]. This allows us to account for differences in the DLA masking between our simulated pipeline and the observed pipeline. An example would be DLAs, or Lyman limit systems (LLS), which are not detected in the observational pipeline due to low spectral signal-to-noise. Note that our simulation includes a model that produces realistic populations of LLSs and DLAs, so the marginalised template allows for aspects in which the simulated model differs from the real Universe. In [89], there are four parameters, with sub-DLAs separate from LLSs, and DLAs divided into two categories. For each of the parameters, a redshift and scale dependent correction is applied, where a positive (negative) value for the parameter implies that our simulation has underestimated (overestimated) the number of absorbers in that category. We found that in practice our dataset was unable to measure separately all four of the column density bins. We thus simplify our likelihood by using only two additional free parameters, one parameter covering sub-DLAs and LLS, α_{lls} , and one parameter covering DLAs, α_{dla} . α_{lls} covers column densities between $1.6 \times 10^{17}\text{--}10^{20} \text{ cm}^{-2}$, and α_{dla} covers $10^{21}\text{--}10^{22.5} \text{ cm}^{-2}$.

We account for correlated Si III absorption within the Lyman- α forest following ref. [9]. Our likelihood includes an additional nuisance parameter, f_{SiIII} , which measures the amplitude of the metal contamination. An improved model could use the metal line distribution in a cosmological simulation.

3.2 IGM temperature at mean density data

We use the IGM temperatures at mean density from ref. [30], derived from simulation modeling of high resolution quasar spectra from the KODIAQ survey [90]. Ultimately the dataset is a relatively small, visually inspected set of high resolution quasar spectra. Importantly, these spectra are independent of the eBOSS quasar sample, justifying our choice of separate likelihood functions. We include IGM temperature data for $z = 2.2\text{--}3.8$, for consistency with

the available Lyman- α forest flux power data. The average uncertainty for this data set is $\approx 10\%$, whereas our IGM temperature emulator has an average uncertainty of $\sim 1\%$. Ref. [30] provides IGM temperatures derived from four different statistics: the Lyman- α forest flux power spectrum, curvature, wavelet decomposition, and Doppler width distribution. We use the Lyman- α forest flux power spectrum derived temperatures in the main body of this work, but show results using these other data sets in appendix C.

To derive temperatures from the observed quasar spectra, ref. [30] calculated several summary statistics and compared them to those derived from spectra drawn from simulations. The simulations they used were similar in resolution to our HF suite (gas mass resolution of $\sim 10^5 M_\odot$), though much smaller in volume (10 Mpc/h box side length). These observed IGM temperatures are themselves derived using a suite of simulations assuming fixed σ_8 and n_P taken from a Planck prior [30]. If the structure on Lyman- α forest scales differs significantly from the predictions of the Planck model, the derived IGM temperatures may be biased or inaccurate. In this case, combining the IGM temperatures with the flux power spectrum data could lead to inconsistent constraints. An improved analysis would directly model the observed small-scale 1D flux power spectrum in combination with the flux power spectrum from eBOSS. We will attempt this in future work.

We do not include data on the IGM temperature-density relation, γ . The constraints from ref. [30] are $\gamma(z=3) = 1.22 \pm 0.12$ at 68% confidence. A reasonable theoretical prior would be $1.1 \lesssim \gamma \lesssim 1.6$, so this measurement does not represent a strong constraint on our model. In addition, γ varies within our simulations, depending on the redshift at which a particular region of the box undergoes helium reionization [80].

3.3 Covariance matrix

In this section, we derive the covariance matrix, \mathbf{K} , that is used for our inference. We decompose \mathbf{K} as:

$$\mathbf{K} = \mathbf{K}_{\text{BOSS}} + \boldsymbol{\sigma}_{\text{GP}}(\mathbf{p}) \cdot \boldsymbol{\sigma}_{\text{GP}}^T(\mathbf{p}) + \boldsymbol{\sigma}_{\text{CV}} \cdot \boldsymbol{\sigma}_{\text{CV}}^T. \quad (3.1)$$

Here, \mathbf{K}_{BOSS} is the covariance matrix from the eBOSS pipeline [38], and is the largest term in the covariance matrix on most scales. We also add two extra terms which model theoretical error in our model. $\boldsymbol{\sigma}_{\text{GP}}(\mathbf{p})$ is the parameter dependent estimate of the interpolation error from the Gaussian process. Using mocks (see figure 3), we found that the GP error can sometimes unphysically drive the chain away from the edges of parameter space, where the expected interpolation error is large. We thus choose to omit it from the overall covariance matrix. Appendix A shows that its addition has a small effect on our final results.

The second theoretical error in our simulation suite (which dominates) is $\boldsymbol{\sigma}_{\text{CV}}$, which models residual sample variance from the finite box size, analogous to cosmic variance from the finite cosmological horizon.⁹ We include an estimate of sample variance using the leave-one-out errors discussed in ref. [47], a technique made possible by the inclusion of h in our simulation suite. The Hubble parameter does not directly affect the gravitational evolution in

⁹Ref. [91] reduced sample variance by interpolating the parameters of a higher order polynomial, rather than fitting the binned flux power spectrum directly. Our emulator is much less affected by sample variance as our simulated volume is 8 times larger.

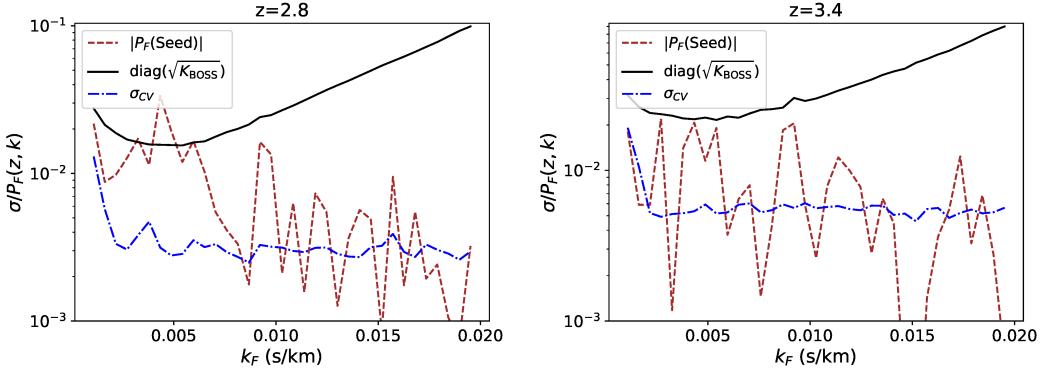


Figure 2. Elements of the covariance matrix, as a function of scale and for selected redshift bins $z = 2.8$ (left) and $z = 3.4$ (right). We have selected these redshifts as they are those where σ_{CV} is most important relative to K_{BOSS} . Shown are: 1) The ratio of the 1D flux power spectra between the two LF simulations with different seeds discussed in ref. [47] ($|P_F(\text{Seed})|$, brown, dashed). 2) the square root of the diagonal elements of the eBOSS covariance matrix K_{BOSS} (black solid) 3) the diagonal cosmic variance error estimated from leave-one-out errors, σ_{CV} (blue, dot-dashed). The y-axis label $\sigma/P_F(z, k)$ signifies that all errors are normalised by the 1D flux power spectrum, $P_F(k, z)$. Hence σ_{CV} is $\sigma_{\text{CV}}/P_F(z, k)$ and $|P_F(\text{Seed})|$ is $|P_F(\text{Seed}, z, k)/P_F(\text{DefaultSeed}, z, k) - 1|$.

our simulations due to Gadget’s use of Mpc/h units, and ref. [47] showed that the effect on the thermal history is small on the scales probed by eBOSS. However, in our parameterization, h also changes Ω_M ,¹⁰ and so the conversion of wavenumbers from h/Mpc to s/km. Individual Fourier modes thus move between bins depending on the value of h , mimicking the sample variance from different initial phases. We thus approximate σ_{CV}^2 with the averaged variance of the leave-one-out errors using the low fidelity simulations. Leave-one-out errors are found by building a reduced emulator, which is trained on all but a single sample, then evaluating the prediction accuracy for that left-out sample using the reduced emulator. This is then repeated, such that every sample takes a turn being left out (see figure 2 of [47]):

$$\sigma_{\text{CV}}^2 = \frac{1}{N_{\text{LF}}} \sum_i \left(P_F^{\text{Predict}}(k, z, p_i) - P_F^{\text{True}}(k, z, p_i) \right)^2. \quad (3.2)$$

Here the sum runs over all simulated low-fidelity parameter sets p_i and N_{LF} is the number of low-fidelity simulations.

Figure 2 shows the magnitude of the σ_{CV} term compared to the eBOSS errors. σ_{CV} is significant only on the largest scales, $k < 2.5 \times 10^{-3}$ s/km, as expected from an effect due to finite box size. In addition, there is a significant redshift dependence: σ_{CV} is important relative to the eBOSS errors only for $2.8 < z < 3.4$. For clarity, figure 2 shows only $z = 2.8$ and $z = 3.4$, which bracket the largest effect. For the redshift range $4.2 > z > 3.4$, σ_{CV} remains approximately constant, but becomes less significant as the eBOSS statistical errors increase. These details reveal the physical source of this large-scale variance. The relevant scale is close to the 20 Mpc/h size of the helium reionization bubbles, and the relevant redshift range is when our model performs helium reionization. Helium reionization bubbles are placed randomly around rare large halos, which creates sample variance in a finite box.

¹⁰ $\Omega_M h^2$ is a separate parameter and so kept fixed when varying h .

Figure 2 also shows the absolute value of the ratio of the flux power spectrum between two LF simulations with different initial structure seeds. If this ratio is consistently larger than σ_{CV} , it would indicate an extra source of cosmic variance not included by varying h . The average flux power spectrum ratio is similar to σ_{CV} on the largest scales measured, where it is most important. We have deliberately chosen to show the redshift bins near the ends of helium reionization where the effects of cosmic variance are most important. To evaluate the effect of a potential under-estimation of cosmic variance from these k-bins, we performed two extra chains where the last term in eq. (3.1) was $2 \times \sigma_{\text{CV}}$, and where $\sigma_{\text{CV}} = 0.02$. Our posterior constraints were almost unchanged in both cases, likely because K_{BOSS} is generally larger than 2%.

3.4 Likelihood

We use a log normal likelihood summed over all redshifts and, for the flux power, all scale bins:

$$\log \mathcal{L} = -\frac{1}{2} \sum_{z=2.2}^{z=4.6} \left(\left(\mathbf{P}_F^{\text{diff}} \right)^{\top} \cdot \mathbf{K}^{-1} \cdot \mathbf{P}_F^{\text{diff}} + \log(\det(\mathbf{K})) \right) \quad (3.3)$$

where $\mathbf{P}_F^{\text{diff}} = \mathbf{P}_F^{\text{sim}} - \mathbf{P}_F^{\text{obs}}$ is the vector difference between the simulation prediction and the observation. The covariance matrix, \mathbf{K} , is described in equation (3.1).

The likelihood for the IGM temperature is similar, but single valued per redshift. We compute the Lyman- α forest flux power and IGM temperature likelihoods separately and add the log likelihoods. We make use of the Cobaya package [85, 86, 92, 93] to run MCMC chains using this likelihood. The MCMC sampler uses the Metropolis method discussed in [92], and uses a Gaussian + exponential proposal distribution that dynamically learns the proposal covariance. Convergence is determined using the Gelman-Rubin statistic, R , also detailed in [92]. The chains presented here were run until a convergence of $R - 1 < 0.01$, with results plotted for those chains for samples at $R - 1 < 1$.

3.4.1 Priors

We use the parameter limits shown in table 2. As we showed in ref. [47], the AGN feedback parameter ϵ_{AGN} has minimal effect on the Lyman- α forest 1D flux power spectrum. Preliminary chains indicated that it is indeed poorly constrained by the data and has minimal correlations with other parameters. We use a strong Gaussian prior with $\mu = 0.05$ and $\sigma = 0.005$, which dominates over data constraints, and will omit constraints on ϵ_{AGN} from our results. We also place a weak Gaussian prior on the Hubble parameter, h , with $\mu = 0.70$ and $\sigma = 0.015$, as it is weakly constrained and this prior avoids the inference straying into areas near the edge of parameter volume where the emulation is less accurate. For all other parameters we use uniform priors within the parameter limits.

3.5 Inference using simulation data

In this section we test our inference framework with simulation outputs in place of the observational data, confirming that we recover the known input parameters. We first used the flux power spectrum from one of the three high fidelity simulations, and confirmed that

the maximum likelihood was indeed at the input parameter values for all parameters. All input parameters were recovered to better than one sigma.

We next ran chains using data from a low-fidelity simulation with a different random seed from the main PRIYA suite. For these runs only, we used an emulator built using the low fidelity suite. This test was designed to quantify whether the finite box size of our simulations can affect our parameter constraints. Figure 3 shows the results, with dashed black lines indicating the correct parameters. Note that figure 3 shows an estimate of the potential bias from cosmic variance on our results.

We have performed three runs. The first (Seed FPS) is our preferred error model, using the eBOSS covariance and a leave-one-out error term. The second adds an error term for the expected interpolation error from the Gaussian Process (Seed FPS + GPERR). Both of these runs use only information from the eBOSS flux power spectrum and thus do not provide strong constraints on the parameters of helium reionization. We therefore run another chain including constraints from the IGM temperature.

In all three runs, the optical depth parameters, τ_0 and $d\tau_0$, are tightly constrained around the true value in our pipeline, despite the effect of a different structure seed. The GPERR chain increases the uncertainty, especially on $d\tau_0$, but does not bias the measurement. The best estimate comes from the flux power spectrum data alone (Seed FPS). We also consistently recover the true values of the cosmological parameters n_P and A_P to better than $1 - \sigma$. Note that we deliberately constructed our test data, with a different structure seed, to be different from the training data. $\Omega_M h^2$ is poorly constrained in all chains, as expected given that our prior volume includes only a narrow range for $\Omega_M h^2$, motivated by Planck results. All parts of the prior range are within the $1 - \sigma$ posteriors.

The redshift of hydrogen reionization, z_{HI} , is estimated from the IGM temperature at $z > 3.6$ or from a large-scale increase in the flux power spectrum at $z > 4$ (see ref. [47]). The second effect is due to a scale-dependent bias arising from placement of the reionization bubbles [94]. Figure 3 indicates that this bias is sensitive to sample variance from the finite box, and so the hydrogen reionization redshift is not well measured by the flux power spectrum data alone. The three parameters which govern helium reionization, z_i^{HeII} , z_f^{HeII} and α_q , are well constrained by the IGM temperature data. The runs which do not include IGM temperature data have a slight preference for a larger α_q than the input value. As discussed above, the main effect of a different structure seed is through the placement of helium reionization bubbles. α_q is thus measured using a similar scale-dependent bias as z_{HI} , and so is slightly sensitive to the finite box size in the same way. However, the IGM temperature is sensitive to α_q through the peak temperature during helium reionization, and thus the chains including it correctly infer α_q .

The chain including Gaussian Process errors produces some incorrect parameter inferences, notably in α_q . This is because of the specific choice of simulated data, for which α_q is at the lower boundary of the emulator parameter range. The GP expects the emulator error to be larger near the boundary of the space, which penalises the fit in this region when the constraints from the data are weak. Notice that α_q is poorly measured by the flux power spectrum alone. In fact, our leave-one-out error calculation reveals that the flux power spectrum is reasonably accurately modelled even near the emulator boundaries, but, with

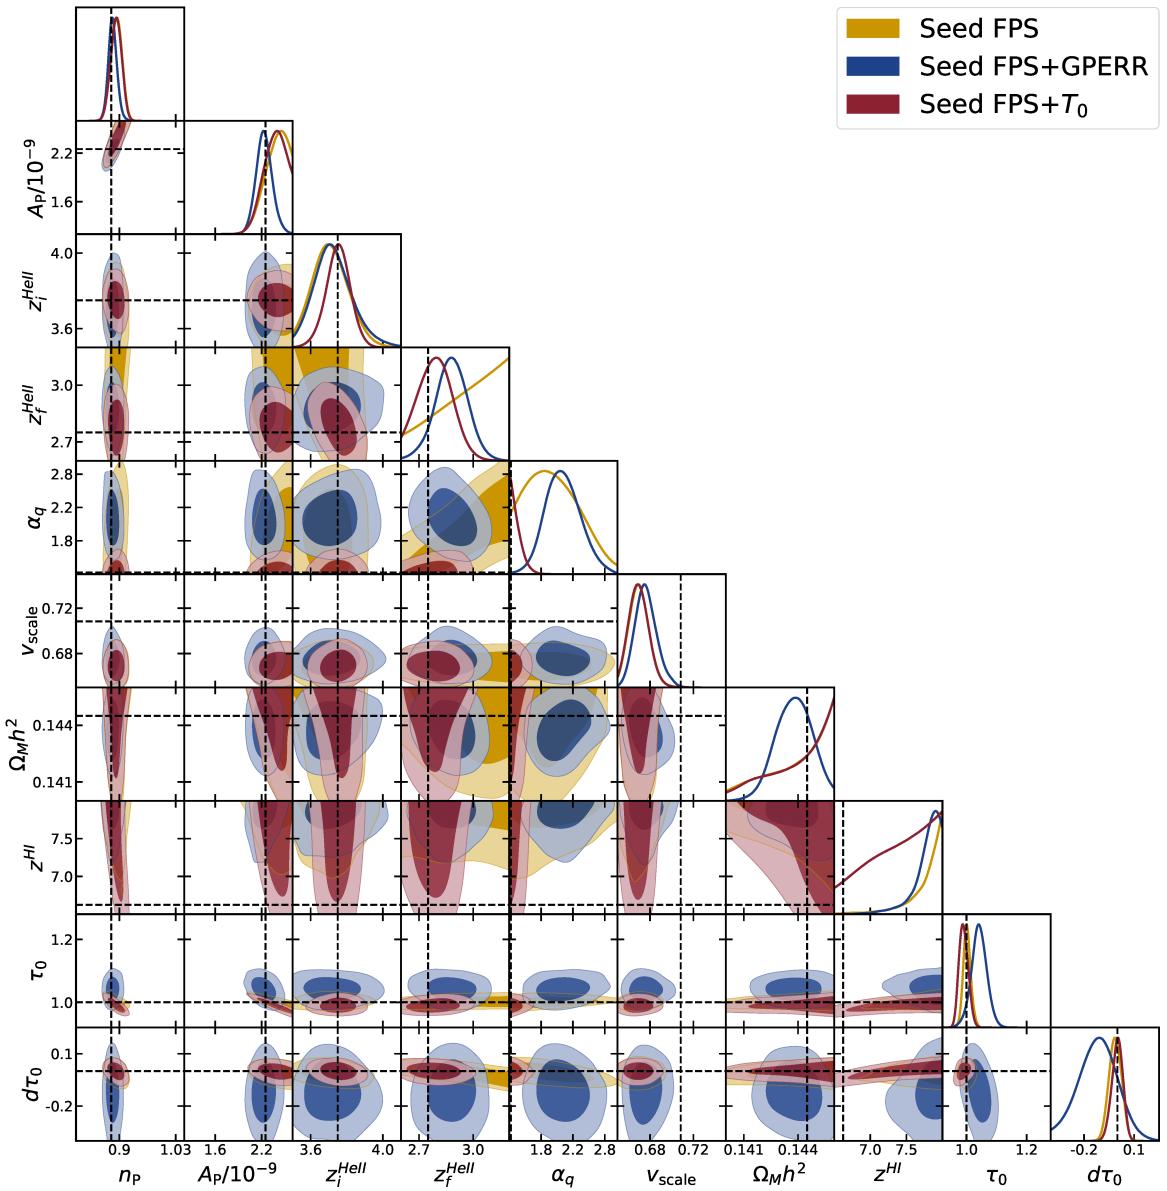


Figure 3. Posteriors using mock data, a simulation output with a different initial seed to the main PRIYA suite. The true parameter values for the input are indicated by the black dashed lines. Three chains are used: ‘Seed FPS’ uses the default error model, with the eBOSS covariance and leave-one-out errors. ‘Seed FPS + GPERR’ adds the covariance from the Gaussian Process. ‘Seed FPS + T_0 ’ uses the default error model but supplements the flux power spectrum data with information from the IGM temperature. v_{scale} is the Hubble parameter h , re-labelled for reasons explained in the text.

few simulations in this region, there is not enough information present for the GP to learn this and reduce the expected error. We show in appendix A that the posteriors from eBOSS data are not significantly changed by including GP error. Nevertheless, to be conservative our main results are reported with chains run omitting GP errors from the likelihood.

As discussed above, the Hubble parameter, h , does not affect the evolution of our simulations except through its effect on Ω_M (at fixed $\Omega_M h^2$) and thus the scaling between

velocity (km/s) and comoving (Mpc/h) units. Constraints on h (v_{scale}) are incorrect in the chains shown in figure 3, driven by sample variance in the finite box. We confirmed that gradients of the likelihood with respect to h are largest for the largest scales, particularly the first 4 k -bins. We computed the χ^2 per degree of freedom, which was ~ 0.9 for $h = 0.65$ and $\chi^2 \sim 1$ for the true input value, indicating over-fitting to the noise created by different structure seeds. We confirmed that fixing h to the known true value results in very small changes to the other parameters and their confidence intervals. There is thus zero cosmological information in our h constraints. In order to avoid unwarranted conclusions, we will henceforth relabel h as v_{scale} , emphasising that it merely controls the mapping between the native Fourier-space binning of the simulations and the observed velocity space of the spectra, and its inference is dominated by modelling error from sample variance.

The posterior v_{scale} from the simulation with a different seed does not match the input value. However, this is to be expected: v_{scale} is completely degenerate with the initial structure seed. The measured quantity is thus schematically a linear combination of h and a ‘cosmic variance parameter’ set by the initial distribution of structures. When the value of the cosmic variance parameter expected by the emulator differs from the true input, the emulator will measure a different posterior value for h , as shown in figure 3.

4 Results

In this section, we report posterior constraints on the parameters listed in table 2. Section 4.1 discusses the results for those parameters which are most strongly constrained by the flux power spectrum data. These are: the optical depth parameters, τ_0 and $d\tau_0$, the power spectrum parameters, n_P and A_P , and the matter density $\Omega_M h^2$. Section 4.2 then discusses the constraints on the other parameters, and shows the best fit to the IGM temperature at mean density data. These are the three parameters defining the He II reionization model, $z_i^{\text{He II}}$, $z_f^{\text{He II}}$, and α_q ; the parameter for the midpoint of H I reionization, $z^{\text{H I}}$, the strong absorber models, (α_{LLS} and α_{DLA}), the Silicon III correction (fSiIII) and the velocity to distance scale parameter v_{scale} . The same chains are used in all sections: we split parameters into two sections merely for readability. We show the full corner plot, containing all constrained parameters, in appendix D. Table 3 shows posterior parameter constraints. We also calculate the derived parameters A_s and σ_8 , Δ_L^2 and n_{eff} .

4.1 Cosmological parameters

Figure 4 shows the results of our chains for the cosmological parameters. We show three MCMC chains. Two chains are fit to the eBOSS flux power spectrum data only. The first fits to the full redshift range measured by eBOSS, $z = 2.2\text{--}4.6$, while the second fits a limited redshift range $z = 2.6\text{--}4.6$. The third chain uses the limited redshift range eBOSS dataset but adds the IGM temperature likelihood. The chain including the $z < 2.6$ data prefers lower n_P , lower A_P and higher τ_0 than the reduced redshift range. Figure 5 shows that the shift in posterior parameters is driven by the fit. The best-fit flux power spectrum to the data at $z \geq 2.6$ is a poor fit to the flux power spectra measured at $z = 2.2$ and $z = 2.4$. Since the lowest redshift bins have the smallest statistical error, when they are included

Parameter	FPS $z > 2.6$	FPS + T_0	FPS $z > 2.2$
	68% (95%)	68%, (95%)	68%, (95%)
$d\tau_0$	$-0.013_{-0.041}^{+0.047} \left(+0.083 \atop -0.089 \right)$	$0.009 \pm 0.045 \left(+0.089 \atop -0.087 \right)$	$-0.270 \pm 0.029 \left(+0.057 \atop -0.057 \right)$
τ_0	$1.082_{-0.026}^{+0.018} \left(+0.046 \atop -0.041 \right)$	$1.090 \pm 0.022 \left(+0.043 \atop -0.042 \right)$	$1.221_{-0.012}^{+0.021} (< 1.19)$
n_P	$1.009_{-0.018}^{+0.027} \left(+0.040 \atop -0.039 \right)$	$0.983 \pm 0.020 \left(+0.040 \atop -0.039 \right)$	$0.898_{-0.013}^{+0.012} \left(+0.025 \atop -0.023 \right)$
$A_P/10^{-9}$	$1.69_{-0.15}^{+0.14} \left(+0.30 \atop -0.28 \right)$	$1.46_{-0.13}^{+0.099} \left(+0.22 \atop -0.22 \right)$	$< 1.33 (< 1.44)$
$\Omega_M h^2$	$< 0.142 (< 0.144)$	$< 0.143 (--)$	$< 0.141 (< 0.142)$
z_i^{HeII}	—	$> 4.00 (> 3.87)$	$> 4.07 (> 4.01)$
z_f^{HeII}	$< 2.67 (< 2.80)$	$2.765_{-0.093}^{+0.080} \left(+0.14 \atop -0.16 \right)$	$< 2.70 (< 2.83)$
α_q	$> 2.86 (> 2.64)$	$1.74_{-0.21}^{+0.18} \left(+0.37 \atop -0.38 \right)$	$> 2.85 (> 2.65)$
z^{HI}	$< 7.01 (< 7.57)$	$7.24 \pm 0.38 (--)$	$7.28 \pm 0.37 (--)$
v_{scale}	$0.693 \pm 0.0085 \left(+0.018 \atop -0.017 \right)$	$0.688_{-0.0074}^{+0.013} \left(+0.018 \atop -0.025 \right)$	$0.695_{-0.0073}^{+0.0063} \left(+0.015 \atop -0.014 \right)$
α_{lss}	$0.193 \pm 0.035 \left(+0.069 \atop -0.069 \right)$	$0.196 \pm 0.033 \left(+0.063 \atop -0.065 \right)$	$0.042 \pm 0.018 (\pm 0.036)$
$\alpha_{\text{dla}}/10^{-2}$	$-0.28 \pm 0.64 \left(+1.2 \atop -1.3 \right)$	$-1.05 \pm 0.58 (\pm 1.1)$	$-0.87 \pm 0.34 \left(+0.67 \atop -0.66 \right)$
$f_{\text{SiIII}}/10^{-3}$	$9.63 \pm 0.51 \left(+9.8 \atop -1.0 \right)$	$9.63 \pm 0.51 \left(+0.97 \atop -1.0 \right)$	$8.83 \pm 0.39 (\pm 0.76)$
$A_s/10^{-9}$	$1.65_{-0.14}^{+0.12} \left(+0.27 \atop -0.25 \right)$	$1.52_{-0.13}^{+0.10} \left(+0.24 \atop -0.22 \right)$	$1.73_{-0.11}^{+0.076} \left(+0.19 \atop -0.17 \right)$
σ_8	$0.733_{-0.029}^{+0.026} \left(+0.057 \atop -0.053 \right)$	$0.703_{-0.027}^{+0.023} \left(+0.049 \atop -0.047 \right)$	$0.715_{-0.022}^{+0.015} \left(+0.038 \atop -0.034 \right)$
Δ_L^2	$0.302_{-0.027}^{+0.024} \left(+0.053 \atop -0.048 \right)$	$0.267_{-0.023}^{+0.018} \left(+0.042 \atop -0.039 \right)$	$0.2316_{-0.016}^{+0.0082} \left(+0.027 \atop -0.021 \right)$
n_{eff}	$-2.264_{-0.018}^{+0.026} \left(+0.038 \atop -0.042 \right)$	$-2.288 \pm 0.020 \left(+0.040 \atop -0.039 \right)$	$-2.376_{-0.013}^{+0.012} \left(+0.025 \atop -0.024 \right)$

Table 3. Posterior parameter constraints, including the derived parameters A_s and σ_8 , as well as the predicted linear theory power, Δ_L^2 , and slope, n_{eff} , evaluated at $k_P = 0.009 \text{ s/km}$, $z_P = 3$. Maximum posterior values, and 68% confidence limits are shown, with 95% confidence intervals in brackets. Each column shows a separate chain, from left to right: fits to the flux power spectrum alone from the reduced redshift range $z = 2.6\text{--}4.6$, fits to the flux power spectrum from the reduced redshift range $z = 2.6\text{--}4.6$ and the IGM temperature, and fits to the flux power spectrum alone from the full redshift range $z = 2.2\text{--}4.6$. Single sided limits are shown when one bound is larger than the prior volume of the emulator. ‘—’ denotes that both 68% and 95% constraints are wider than the prior volume of the emulator.

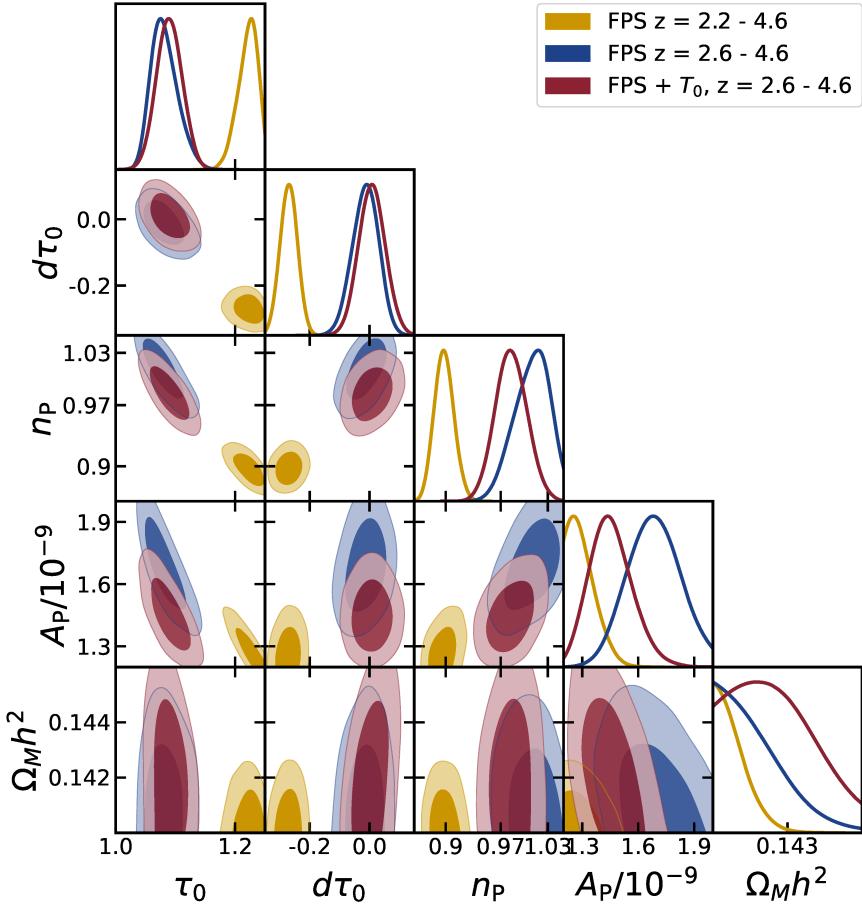


Figure 4. Posteriors for the optical depth and power spectrum parameters, τ_0 , $d\tau_0$, n_P , A_P , and $\Omega_M h^2$. Results are from three MCMC chains. ‘FPS $z = 2.2 - 4.6$ ’ (gold) uses the full redshift range eBOSS flux power spectrum dataset, ‘FPS $z = 2.6 - 4.6$ ’ (blue) uses a reduced redshift range eBOSS dataset flux power spectrum dataset, which removes the internal tension (section 5.1). The third chain, ‘FPS + T_0 , $z = 2.6 - 4.6$ ’ (red), uses the limited range eBOSS dataset but adds the IGM temperature constraints. Our preferred cosmological constraints are from ‘FPS $z = 2.6 - 4.6$ ’.

they drive the best-fit flux power spectrum to a region which is a poorer fit to the higher redshift data. We confirmed that a chain which included the $z = 2.4$ bin but not the $z = 2.2$ bin produced posterior constraints mid-way between the chain including $z = 2.2 - 4.6$ and the chain including $z = 2.6 - 4.6$. Table 4 shows this quantitatively. The chains which fit to $z > 2.6$ are a poor fit to the lowest two redshift bins. The chain fitting to $z = 2.2$ is a better fit to these bins, at the cost of an overall worse fit to redshifts $z > 2.6$. There is thus an internal tension in the eBOSS dataset, when compared to our model, driven by the lowest two redshift bins. We discuss possible reasons for this tension in section 5.1 and compare to the results of earlier analyses further in section 5.2. It is important not to over-interpret the posterior constraints from the $z = 2.2 - 4.6$ chain. When there is an internal tension in the data, the posteriors can be driven by noise in the dataset and may not be meaningful.

We also checked for other redshift bins where the fit is poor. Visually, figure 5 suggests that the fit is poor for $z = 4.0$ and $z = 4.2$. The χ^2 in table 4 is moderately higher other

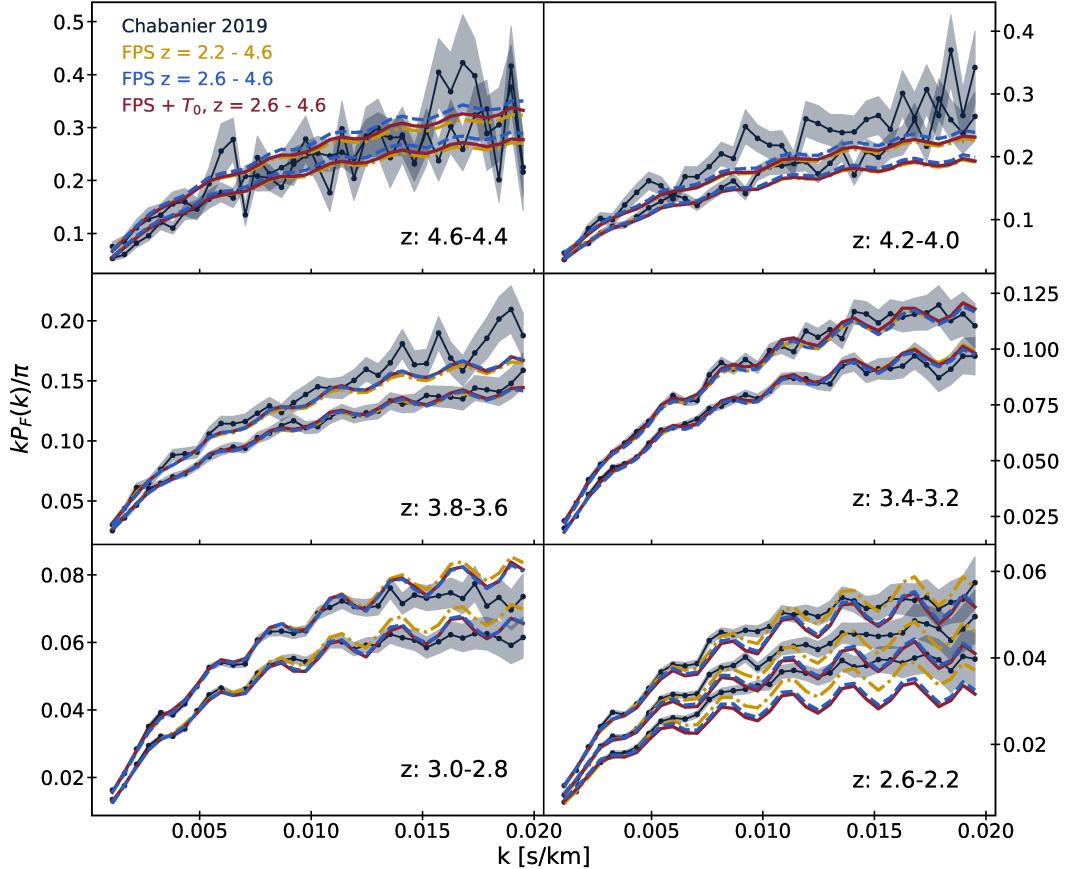


Figure 5. Observed Lyman- α forest flux power spectrum [38], from $z = 4.6$ to $z = 2.2$ (black lines and circles, with shading corresponding to one sigma uncertainty). Also shown are three predictions for the Lyman- α forest flux power spectrum from our multi-fidelity emulator corresponding to the maximum posterior input parameters compiled in table 3. The negative of the log-likelihood for these fits is compiled in table 4.

bins, but not significantly so, as the statistical errors are also large. It is possible that some element of the covariance matrix, theoretical or systematic, is moderately underestimated in these bins. Intriguingly, the fit to $z = 2.6$ is better when the $z < 2.6$ bins are included, which may suggest that the cause of the low redshift tension still has a small effect at $z = 2.6$.

Posterior constraints from the reduced redshift range flux power spectrum data show a mean optical depth in good agreement with other measurements. As a reminder, our τ_0 and $d\tau_0$ parameters measure deviations from the mean flux relation of ref. [81], so a value of $\tau_0 = 1$ and $d\tau_0 = 0$ corresponds to agreement with that model. The maximum posterior values are $\tau_0 = 1.1$, with a redshift variation $d\tau_0$ consistent with 0. This implies a mean optical depth at $z = 3$ of $\tau_{\text{H}_1}^{\text{eff}}(z = 3) = 0.398$, which is extremely close to the best-fit value of $\tau_{\text{H}_1}^{\text{eff}}(z = 3) = 0.4$ from ref. [95], and consistent within the error bars with $\tau_{\text{H}_1}^{\text{eff}}(z = 3) = 0.36 \pm 0.1$ from ref. [81].

The spectral index, n_P , is $n_P = 1.009_{-0.018}^{+0.027}$ (68% confidence) using only the flux power spectrum. On inclusion of the IGM temperature data, we find slightly tighter constraints of $n_P = 0.983 \pm 0.02$ at 68% confidence, lower by about $1 - \sigma$. Planck measures $n_s = 0.965 \pm 0.004$, consistent with the IGM temperature chain at $1 - \sigma$ and the flux power spectrum chain at

Redshift	4.6	4.4	4.2	4.0	3.8	3.6	3.4
FPS $z = 2.2\text{--}4.6$	23.8	22.1	28.3	33.9	18.7	8.7	16.1
FPS $z = 2.6\text{--}4.6$	27.8	19.1	23.0	26.5	15.6	7.4	16.5
FPS + T_0 $z = 2.6\text{--}4.6$	24.9	18.2	26.0	31.3	16.7	6.6	14.4
Redshift	3.2	3.0	2.8	2.6	2.4	2.2	Total
FPS $z = 2.2\text{--}4.6$	19.6	33.8	33.7	20.3	19.9	27.3	306
FPS $z = 2.6\text{--}4.6$	18.9	25.6	24.6	29.1	66.5	97.6	234
FPS+ T_0 $z = 2.6\text{--}4.6$	18.2	24.9	25.2	37.7	76.0	103	244

Table 4. Negative log-likelihood (χ^2) for the flux power spectrum for each redshift bin. Shown is the likelihood at the best fit parameters in each chain. We show chains fitting to the flux power spectrum only at $z = 2.2\text{--}4.6$, fitting to the flux power spectrum only at $z = 2.6\text{--}4.6$, and fitting to the flux power spectrum and IGM temperature at $z = 2.6\text{--}4.6$. The column labelled ‘Total’ excludes redshift bins not in the fit ($z = 2.2$ and $z = 2.4$ for the last two chains).

$2.5 - \sigma$ [96]. The matter density, $\Omega_M h^2$, is weakly constrained and not strongly affected by including IGM temperature data. Planck found $\Omega_M h^2 = 0.1424 \pm 0.001$, which is close to the posterior of our chains. The power spectrum amplitude is $A_P/10^{-9} = 1.69 \pm 0.14$ for the flux power spectrum reduced redshift result. The inclusion of the IGM temperature shrinks the constraints moderately and shifts the posterior value down by about $1.5 - \sigma$, driven by a correlation with α_q . Table 3 shows A_s , the power spectrum amplitude measured on large scales, which is related to A_P via:

$$A_s = (0.4/2\pi)^{n_P-1} A_P. \quad (4.1)$$

We find $A_s = (1.65^{+0.12}_{-0.14}) \times 10^{-9}$ for the flux power spectrum alone and $A_s = (1.53^{+0.10}_{-0.13}) \times 10^{-9}$ when including the IGM temperature. Planck [96] found a value of $A_s = (2.101^{+0.031}_{-0.034}) \times 10^{-9}$. We also derived the value of σ_8 implied by our parameters by using CLASS in post-processing [97]. For the flux power spectrum alone, we find $\sigma_8 = 0.733^{+0.026}_{-0.029}$, and when the IGM temperature is included, $\sigma_8 = 0.703^{+0.023}_{-0.027}$ (see table 3). The Planck result is $\sigma_8 = 0.811 \pm 0.006$ [96]. We thus measure a power spectrum amplitude around $3 - \sigma$ lower than Planck or ACT CMB lensing [98]. Interestingly, the dark energy survey year 3 results measure $\sigma_8 = 0.733^{+0.039}_{-0.049}$ [99], in good agreement with our results. Other small-scale structure probes vary [e.g. 100–102].

Figure 5 shows the Lyman- α forest flux power spectrum from [38], along with their estimated one sigma uncertainty (black). Also shown are predictions from our multi-fidelity emulator based on the maximum posterior input parameters from MCMC analysis with only the Lyman- α forest flux power emulator in the full and reduced redshift ranges, and MCMC analysis using both the IGM temperature and flux power emulators. The correlation between Lyman- α and Si III absorption is visible in the form of regular oscillations in the power spectrum (in section 3.4 we describe the correction we make for Si III). The best-fit flux power spectrum is not significantly affected by the inclusion of the T_0 data in the likelihood, although the fit is slightly worse at $z = 2.6$.

Table 3 shows the constraints our chains imply for the linear theory power at $z = 3$, $\Delta_L^2(k, z)$ and its slope, n_{eff} [17, 38, 91]. These are defined by:

$$\Delta_L^2(k_*, z_*) = k_*^3 P_L(k_*, z_*) / (2\pi^2) \quad (4.2)$$

$$n_{\text{eff}}(k_*, z_*) = \frac{d \ln P_L}{d \ln k}(k_*, z_*) . \quad (4.3)$$

Here $P_L(k, z)$ is the linear theory power spectrum, evaluated at $z_* = 3$ and $k_* = 0.0009 \text{ s/km}$. We find $\Delta_L^2 = 0.302_{-0.027}^{+0.024}$ and $n_{\text{eff}} = -2.264_{-0.018}^{+0.026}$ from our reduced redshift chains. The correlation coefficient between Δ_L and n_{eff} is 0.42 for the flux power spectrum only (FPS, $z > 2.6$) chain and 0.40 when the mean IGM temperature data is included (FPS + T_0).

4.2 Reionization and other parameters

Figure 6 shows the other parameters of our model from the same chains as figure 4. These are: three parameters of the helium reionization model ($z_i^{\text{He II}}$, $z_f^{\text{He II}}$, and α_q), the midpoint of H I reionization $z^{\text{H I}}$, the parameters of the strong absorber model (α_{LLS} , α_{DLA}), the strength of the metal contamination f_{SiIII} and the box velocity scale v_{scale} .

The combined data prefer an early start to helium reionization, $z_i^{\text{HeII}} > 3.87$ at 95% confidence.¹¹ Interestingly, this is in agreement with constraints from the helium Lyman- α forest, where regions of high transmission suggest that HeII reionization has already started at $z = 3.5$ [103, 104]. The end of helium reionization, z_f^{HeII} , is constrained by the flux power spectrum data alone to be $z_f^{\text{HeII}} < 2.8$. Adding the IGM temperature data reinforces this and constrains it to finish by $z = 2.6$, so that the IGM temperature can drop at lower redshifts. This is consistent with the He II Lyman- α forest, which suggests an end at $z \leq 2.7$ [105–107]. Our constraints on the timing of helium reionization hit our prior volume. Note this prior volume is set more by the redshift limits of our datasets than our simulation choices. The flux power spectrum at $z \geq 2.6$ alone, for example, cannot determine whether the IGM is cooling by $z = 2.2$. In addition, the upper prior limit on z_i^{HeII} is 4.1, which is larger than the highest redshift IGM temperature data.

The most significant effect of the IGM temperature data is on the spectral index during helium reionization, α_q . Smaller values of α_q correspond to a larger heating rate. The flux power spectrum data prefers a high value of $\alpha_q > 2.64$ at 95% confidence, which corresponds to minimal heating during helium reionization. However, as shown in figure 7, this high value of α_q produces an IGM temperature which is low and in disagreement with the data from ref. [30]. The posterior constraint on α_q when the IGM temperature data is included is lower than that with the flux power spectrum data alone by about $4 - \sigma$. Figure 5 shows that the flux power spectrum is not significantly different at the maximum posterior parameters of either chain. Appendix C shows the results of chains which include only the IGM temperature data. They are consistent with the combined chains, and the helium reionization parameters are constrained at similar values, although the maximum posterior α_q is slightly lower than in the joint chains.

There are two possible explanations for this discrepancy in α_q . The first is that the discrepancy is caused because the IGM temperature constraints assumed the Planck value of

¹¹The flux power only chain admits a solution where both helium and hydrogen reionization are late, but this is ruled out by the inclusion of IGM temperature data.

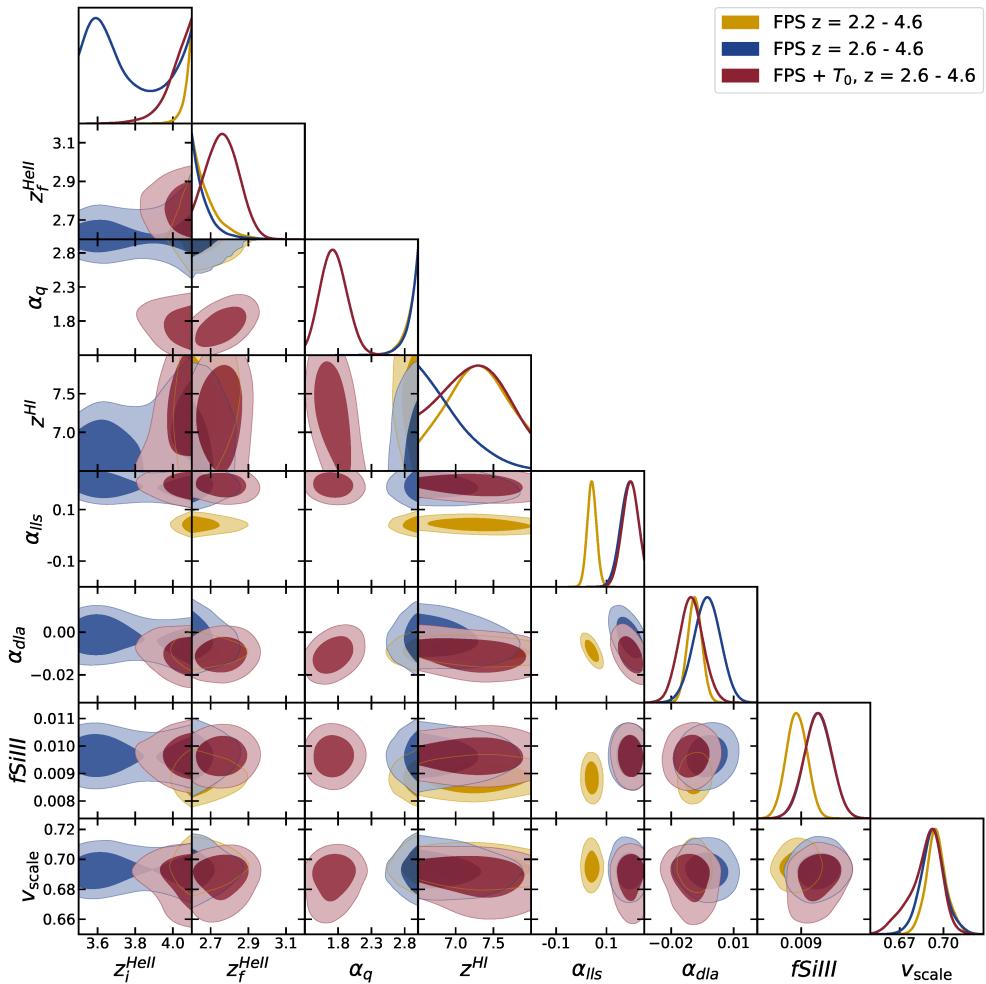


Figure 6. Posterior constraints for the parameters of the helium reionization model (z_i^{HeII} , z_f^{HeII} , α_q), the hydrogen reionization model (z^{HI}), the strong absorber models (α_{LLS} , α_{DLA}), the Silicon III correction (f_{SiIII}), and the velocity to distance scale parameter v_{scale} . Results are from the same three MCMC chains as figure 4. ‘FPS $z = 2.2\text{--}4.6$ ’ (gold) uses the full redshift range eBOSS flux power spectrum dataset, ‘FPS $z = 2.6\text{--}4.6$ ’ (blue) uses a reduced redshift range eBOSS dataset flux power spectrum dataset, which removes the internal tension (section 5.1). The third chain, ‘FPS $+T_0, z = 2.6\text{--}4.6$ ’ (red), uses the limited range eBOSS dataset but adds the IGM temperature constraints.

σ_8 , which is inconsistent with our results. Assuming a larger value for σ_8 will increase the flux power spectrum and require a greater degree of thermal heating to smooth the gas. The IGM temperature constraints from ref. [30] may thus be biased high. A consistent analysis of the flux power spectra from both eBOSS and XQ-100 varying A_P and α_q simultaneously would resolve this question, and we will perform it in future work. The second possibility is suggested by table 4, which reveals that the high α_q is driven by the fit to the $z = 2.6$ flux power spectrum bin. It is thus possible that the same effect which is responsible for the tension between the full redshift and reduced redshift chains is also present at $z = 2.6$ in the flux power spectrum, exploiting partial parameter degeneracies to move the posterior confidence intervals on α_q . We will in future work examine this possibility with the higher resolution DESI data.

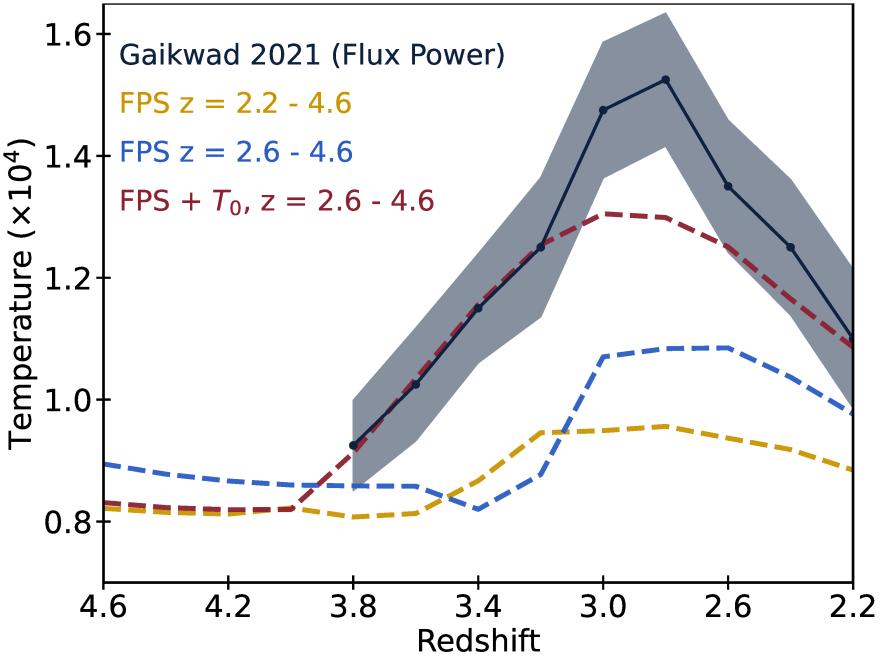


Figure 7. IGM temperatures at mean density from [30] (black lines and circles, with shading corresponding to one sigma uncertainty). Specifically their temperatures derived from the flux power spectrum calculated using high resolution Lyman- α forest spectra. Also shown are predictions for the IGM temperature from our multi-fidelity emulator corresponding to the same three maximum posterior input parameters used in figure 5.

The midpoint of hydrogen reionization is poorly constrained in all models. The redshift range explored here ($z = 2.2\text{--}4.6$) is well after the completion of hydrogen reionization, even in models where it ends late. All our chains are consistent with a midpoint of $z^{\text{H I}} \sim 7$. This is within the range allowed by other experiments: Planck suggests $z^{\text{H I}} = 7.7 \pm 0.7$, while an analysis of the Lyman- α emitter luminosity function finds $z^{\text{H I}} \sim 7.25$ [108].

We show results for the nuisance parameters associated with the strong absorbers and the Lyman- α -SiIII cross-correlation. As discussed in ref. [47], our simulation suite includes strong absorbers self-consistently using a galaxy formation model. The strong absorber parameters measure the difference between the strong absorber model in the simulation and the model in the observed spectra, so that $\alpha_{\text{LLS}} = 0$ means that our galaxy formation model is a good match to the circumgalactic gas in the observed Universe. DLAs are subtracted from both the observed and simulated spectra, and so α_{DLA} measures primarily the efficiency of the observational DLA finder. All our chains produce a posterior α_{DLA} tightly peaked and close to 0. The chains using the flux power alone are centered on $\alpha_{\text{DLA}} = 0$, while the chain including the IGM temperature prefers a slightly negative value, perhaps indicating that the eBOSS DLA finder includes some false positives.¹² Note that DESI includes improved DLA finding algorithms based on machine learning [109–111].

¹²As DLAs finders are sensitive to reduced flux transmission, regions with more Lyman- α forest absorption are more likely than average to be flagged as a DLA. The efficiency of the DLA finder may thus affect the flux power spectrum.

For the restricted redshift chains, we measure $\alpha_{\text{LLS}} \sim 0.19$, while for the full redshift range $\alpha_{\text{LLS}} \sim 0.04$. The preference for a non-zero α_{LLS} in the reduced redshift chains suggest that our simulations have fewer LLS than the real Universe. LLS are on the boundary of being optically thick and thus radiative transfer effects within the gas are important, making them the most difficult absorbers to model accurately. α_{LLS} can affect the flux power spectrum normalisation, and so it is reasonable to interpret the preference of the full redshift chains for a low α_{LLS} as an artifact of the fit. However, it is also possible that the low α_{LLS} points to the origin of the internal tension, a possibility we discuss further in section 5.1.

The SiIII cross-correlation is $f_{\text{SiIII}} = 0.0096 \pm 0.001$ (95% confidence) from our $z = 2.6\text{--}4.6$ chains. The full redshift range prefers a slightly lower value of 0.0085 ± 0.0008 , which is in good agreement with the measurement of 0.008 ± 0.001 from DR9 by ref. [112]. The effect of f_{SiIII} can be seen in the oscillations of the flux power spectrum in figure 5. The results for v_{scale} are dominated by the prior, as expected [22, 113]. Constraints are weaker than for the simulated data, likely because the simulated data did not include noise and so was over-fitting.

Figure 7 shows the IGM temperature from the flux power spectrum on small scales from ref. [30]. Also shown in figure 7 are predictions from our multi-fidelity emulator based on the maximum posterior input parameters from the same chains used in figure 5. Once the IGM temperature data is included in the fit, the chains are in good agreement. However, when it is not included the chains prefer a lower IGM temperature, as discussed above. The thermal history preferred by the full redshift range of the flux power spectrum is similar to that preferred by the restricted range flux power spectrum.

4.3 Parameter correlations

Figure 8 shows the correlations between our parameters, for the chain using the flux power spectrum from $z = 2.6\text{--}4.6$, as well as the IGM temperature. Most correlations are weak. We have deliberately chosen our pivot scale of 0.78 Mpc^{-1} to minimise the correlation between A_P and n_P , and the correlation matrix confirms it is weak, with a correlation coefficient $r = 0.2$. There is a correlation between τ_0 and $d\tau_0$ ($r = -0.54$), as the redshift bin which provides the strongest constraints on the optical depth is not exactly $z = 3$. The optical depth τ_0 is anti-correlated with both A_P ($r = -0.7$) and n_P ($r = -0.62$) as its main effect is to change the amplitude of the flux power spectrum.

There is a three-dimensional degeneracy between α_q , z_i^{HeII} and z_i^{HeIII} (see figure 8), which allows a wide range of α_q to fit the flux power spectrum data, and is only broken by information from the thermal history. Lower α_q corresponds to more heating from quasars during He II reionization. If He II reionization starts earlier or ends later, the IGM requires more heating from quasars to match the observations, while the opposite is true for late starting, or early ending He II reionization. Appendix C shows the results of chains which include only the IGM temperature data, which clearly shows this three-dimensional degeneracy: a slightly later start to helium reionization would require less total heating and thus a higher value of α_q . Several of these correlations could be broken by the inclusion of higher redshift thermal history data, or lower redshift flux power spectrum data.

Finally, the abundance of Lyman Limit Systems, α_{LLS} , exhibits several interesting correlations. α_{LLS} is anti-correlated with α_{DLA} ($r = -0.59$), as the flux power spectrum

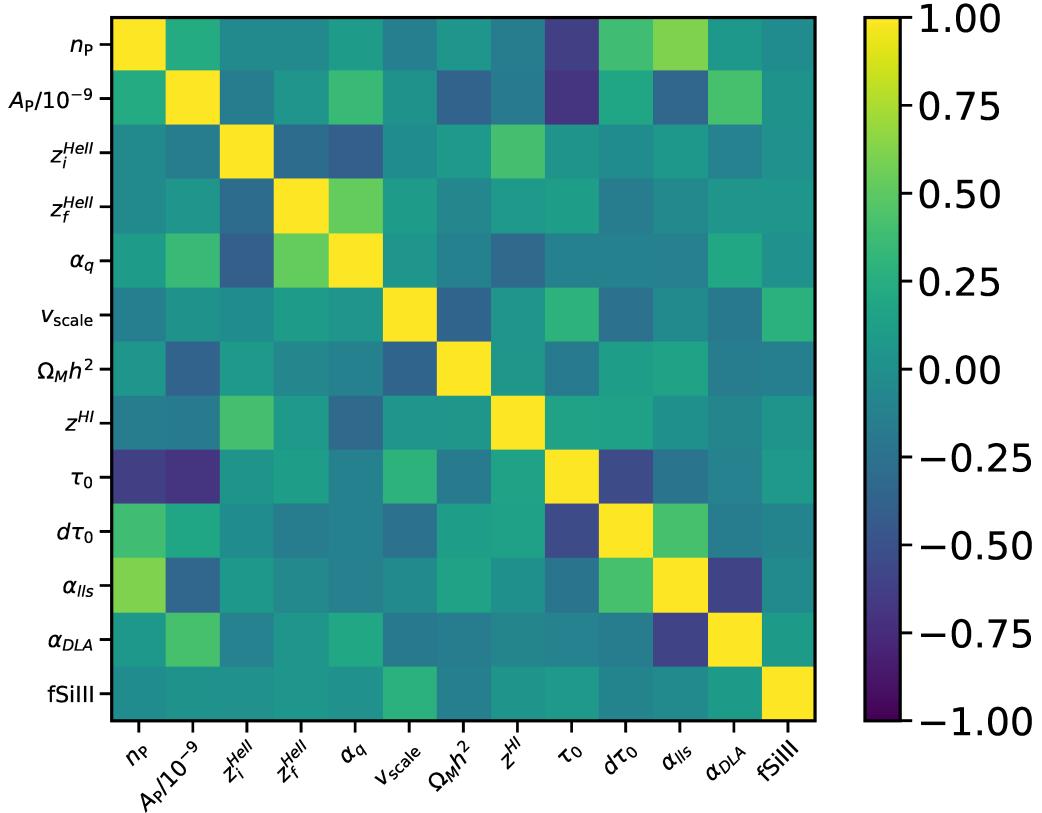


Figure 8. Correlation matrix between parameters for the chain using the flux power spectrum and IGM temperature for $z : 2.6\text{--}4.6$.

templates for strong absorbers have similar shapes in neighbouring column density bins. α_{LLS} is also correlated with n_P ($r = 0.61$) and $d\tau_0$ ($r = 0.42$), due to similarities in the shapes of their flux power spectrum templates. The combination of the three-way correlation between n_P , α_{LLS} and τ_0 is exploited by the chains to explain the inconsistent $z = 2.2, 2.4$ flux power spectrum bins and drives the discrepant constraints these chains show. This correlation may be reduced by the inclusion of extra small-scale data available in the DESI early data release.

5 Discussion

In this section we discuss the implications of the results in section 4. Section 5.1 discusses possible explanations for the internal tension in the data between $z = 2.2\text{--}2.4$ and $z \geq 2.6$. Section 5.2 compares our results to other datasets and earlier Lyman- α analyses. Section 5.3 discusses how our results are affected by modifications to our likelihood.

5.1 The tension in the lowest redshift bins

In this section, we discuss possible explanations for the internal tension between the flux power spectrum data at $z = 2.2\text{--}2.4$ and $z \geq 2.6$. There are two generic possibilities: either an important physical effect is missing from our simulation model, or there is a systematic in the dataset not captured by the systematic error budget.

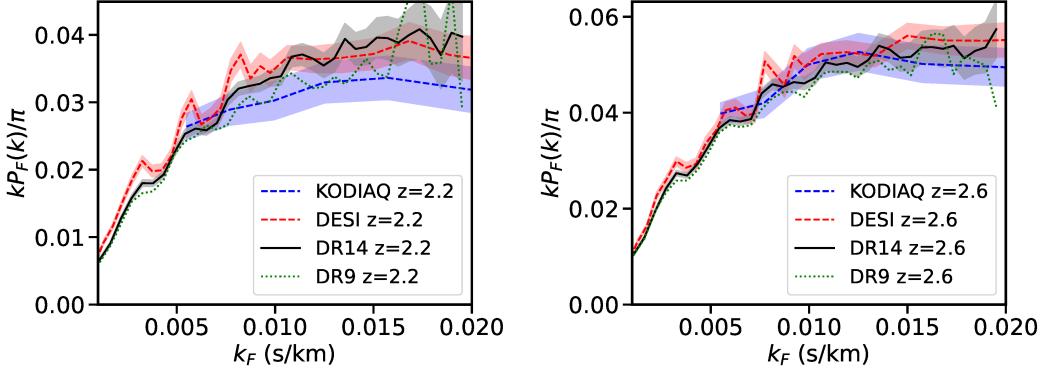


Figure 9. Observational 1D power spectrum data from SDSS DR14 (solid black) [38], SDSS DR9 (green dotted) [112] and KODIAQ/SQUAD [45]. Filled bands show the range covered by diagonal elements of the covariance matrix. (Left) At $z = 2.2$. (Right) At $z = 2.6$.

To evaluate the possibility of systematic error, we can look at independent measurements of the flux power spectrum on similar scales and at similar redshifts. Figure 9 shows different measurements of the 1D flux power, $P_F(k)$, at $z = 2.2$ and $z = 2.6$. We show the results from SDSS DR14 [38], SDSS DR9 [112], a recent analysis using high resolution spectra from KODIAQ/SQUAD [45], and from DESI Early Data Release data (DESI EDR) [40]. At $z = 2.6$ (and higher redshift bins) all analyses are in reasonably good agreement, given their respective statistical errors. However, this is not the case at $z = 2.2$, where there is a discrepancy between SDSS (DR14 and DR9), DESI and KODIAQ. SDSS DR14 and DR9 are in reasonable agreement. Appendix B shows the posterior parameter constraints for DR9.

The KODIAQ flux power spectrum is lower by around $1 - \sigma$ on the smallest scales measured by eBOSS, which could be due to the effect of continuum modelling in the KODIAQ data [45]. The DESI EDR data agrees well with eBOSS for $k > 0.01$ s/km, but is discrepant by $> 2\sigma$ for $k < 0.01$ s/km. This discrepancy is also present at $z = 2.4$, and is discussed in the DESI EDR papers, see ref. [41], appendix D. They ascribed 30% of the difference between eBOSS and DESI to continuum fitting, but the origin of the rest is currently unclear. Given the fairly large disagreements between different measurements at $z = 2.2$, systematic error is a highly plausible explanation for the internal tension we find. In future work we will combine our likelihood function with the DESI flux power spectra and investigate their cosmological implications.

We should also consider possible theoretical explanations. Explanations rooted in alternative early Universe models seem a priori unlikely as they would have to cause an effect only for $z < 2.6$, when the Universe is known to be matter dominated. However, there are a few possible astrophysical explanations. Feedback effects from AGN become increasingly important at low redshift. It is possible that a stronger AGN feedback prescription than we use, or than is implemented in current cosmological simulation suites, could efficiently disrupt gas at $z < 2.6$ on small scales and explain these results. Interestingly, such an AGN feedback model has recently been proposed as an explanation for the low value of $S_8 = \sigma_8(\Omega_M/0.3)^{0.5}$ preferred by some weak lensing surveys [114], which matches our results. Ref. [15] found that AGN feedback affects the forest at the 8% level at $z = 2$, but their AGN model has a different mean flux to

their non-AGN simulation, and once this is corrected for the effect is much smaller. Ref. [14] finds a 10% effect of AGN feedback, but their AGN model results in a star formation rate which is a factor of 10 too low at $z = 2$ [15]. Existing models of AGN feedback which also match the galaxy stellar mass function thus do not seem able to affect the Lyman- α forest at $z \sim 2$ enough to explain this tension. We note, however, that more aggressive models are possible.

DLAs are important at low redshift, and do affect the slope of the flux power spectrum. Our simulations include a population of DLAs in good agreement with observations [47], which are masked using the same procedure as the observational pipeline. We include a free parameter to model the residual power from any DLAs not detected by eBOSS, and the posterior value for this free parameter is consistent with zero. We considered separate constraints on the DLA efficiency at low and high redshift: $\alpha_{\text{DLA}}(z < 2.6)$ and $\alpha_{\text{DLA}}(z \geq 2.6)$. The high redshift DLA efficiency $\alpha_{\text{DLA}}(z \geq 2.6)$ remained consistent with zero. However, the low redshift parameter was negative; $\alpha_{\text{DLA}}(z < 2.6) = -0.0134 \pm 0.0074$, indicating an excess of large-scale power in eBOSS at $z < 2.6$. While most cosmological parameters were unchanged, n_P increased (by about $2 - \sigma$) to $n_P = 0.916 \pm 0.014$, reducing the internal tension by 1/3. Since this is only a partial resolution, the DLA parameter may be sensitive to a continuum fitting problem in the (relatively short) low redshift spectra.

At low redshift, the Lyman- α forest is increasingly contaminated by metal lines. We include a simple prescription for Si III and the inclusion of low redshift data does not drive the best-fit parameter for this model, preferring slightly less metal contamination. However, it is possible that a more sophisticated model could help reduce the tension.

One interesting but entirely speculative possibility is suggested by the LLS abundance, α_{LLS} , which is lower in the full redshift chains. Ref. [116] identified a systematic in the SDSS colour selection which causes quasar sightlines containing Lyman Limit Systems (LLS) to be preferentially selected for spectroscopic followup. Refs. [117, 118] showed that, due to the width of the u -band filter in SDSS, LLS are over-sampled for $z = 2.5\text{--}3.6$ for all quasars in the redshift range $z = 3.0\text{--}3.6$. It is thus possible that α_{LLS} depends on the quasar (not absorber) redshift. Note that the flux power spectrum we measure depends on the absorber redshift and so a simple redshift split would not detect this effect.¹³ A check for colour selection systematics would involve the flux power spectrum being computed from two different quasar redshift bins.

5.2 Comparison of the posterior constraints to other Lyman- α analyses

It is interesting to compare the results of our chains to those of ref. [22] (for DR14) and ref. [113] (for DR9). The most notable difference is that ref. [22] tested excluding the lowest two redshift bins and found minimal change in the posteriors of their cosmological parameter. This disagrees with our results. We believe this discrepancy can be ascribed to our different treatment of nuisance parameters. Ref. [22] employed correction functions for supernova feedback from the OWLs simulation suite [14] and for AGN feedback from the Horizon-AGN suite [15]. Each correction function is most significant at low redshift, and is included with

¹³We ran a chain with two α_{LLS} parameters for $z < 2.6$ and $z \geq 2.6$. The maximum likelihood for α_{LLS} at $z > 2.6$ was $\sim 2 - \sigma$ larger than in the full redshift chain, and n_P increased by $\sim 0.5\sigma$, but all other parameters were unchanged.

a free amplitude parameter which is marginalised over. In addition, earlier models were forced by computational limits to use the ‘splicing’ technique of ref. [51]. In this model multiple simulation boxes with over-lapping scale ranges are combined to model the scales probed by the Lyman- α forest. A single larger simulation with 2048^3 particles was used to generate a scale and redshift dependent correction function, and the amplitude of this correction was marginalised over with a Gaussian prior.

Our larger simulations can instead model all relevant scales in a single simulation and so do not need a free parameter for splicing. In addition, we self-consistently incorporate models for stellar and AGN feedback and star formation into our simulations. Thus the analysis of ref. [22] differs from ours in that it has three nuisance parameters, each of which affects the lowest redshift bins most strongly and each of which is marginalised over in the chains. Ref. [119] mentions that removing splicing reduces n_s significantly (although the posterior value of the splicing correction is not reported), which is what we would expect if the splicing correction were absorbing an internal tension. Thus we believe that the $z = 2.2$ and $z = 2.4$ redshift bins contribute only marginally to the cosmological constraints in ref. [22] and instead constrain splicing and AGN feedback. We should therefore compare the quantitative results of ref. [22] to our reduced redshift chains.

Ref. [22] found $n_s = 0.954 \pm 0.006$ for DR14 and $n_s = 0.938 \pm 0.010$ for DR9. Meanwhile the power spectrum amplitude as measured by σ_8 is 0.826 ± 0.02 in DR14 and similar in DR9. As shown in appendix B, we do not reproduce this shift in n_s , although we do observe a smaller shift in $d\tau_0$. Ref. [22] attribute this shift to the different catalogues for masking DLAs and BAL, and we marginalise out the DLA masking.

Ref. [113] find a mean optical depth of $\tau_{\text{eff}}(z = 3) = 0.0025 \pm 0.0001$ and $d\tau = 3.734 \pm 0.015$ in DR9. We define the mean optical depth relative to the power law $\tau_{\text{eff}} = 0.0023(1+z)^{3.65}$, so that in our parameterization these constraints correspond to $\tau_0 = 1.09 \pm 0.05$ and $d\tau_0 = 0.084 \pm 0.015$. Their optical depth measurements from $z = 2.2\text{--}4.4$ in DR9 are thus in good agreement with our measurements for the reduced redshift range of $z = 2.6\text{--}4.6$ from DR14. Appendix B shows that DR14 generally prefers a lower $d\tau_0$ than DR9. Ref. [22] do not report a value for $\tau_{\text{eff}}(z = 3)$ from DR14. A higher $d\tau_0$ implies a lower n_P . It is possible that the ref. [22] DR14 value of $d\tau_0$ is large and discrepant with other optical depth measurements, driving the change in n_P .

We find $n_P \sim 1.0$ and $\sigma_8 \sim 0.73$ from the reduced redshift flux power alone, a $3.5 - \sigma$ tension in σ_8 . Some of the differences between our constraints on A_P and σ_8 are due to a lever arm effect: for $n_s < 1$ the power spectrum amplitude will be increased on larger scales and measuring n_s will induce a correlation between n_s and σ_8 . A direct comparison can be made by comparing our constraints on Δ_L^2 and n_{eff} to those of ref. [38], who found $\Delta_L^2 = 0.31 \pm 0.02$ and $n_{\text{eff}} = -2.339 \pm 0.006$. Their measurement of n_{eff} is thus $\sim 4\sigma$ lower than the maximum posterior from our reduced redshift flux power spectrum chains, while their measurement of Δ_L^2 is in agreement. The discrepancy between our results thus lies in the measurement of the spectral slope. The ultimate source of this tension is not clear, but our simulation suite is substantially larger and more robust than that used by ref. [22]. A possibility is that the splicing correction they use does not fully account for the correlation of small and large scales (for example, the accuracy of the splice may be cosmology dependent).

Another possibility is that the emulator used is a polynomial expansion around a ‘best-fit’ simulation, with $n_s = 0.9624$. This is quite far from the posterior constraints, and so the accuracy of the polynomial emulator may be reduced. A third possibility is the lack of an explicit model for LLS. There is a correlation between n_P and α_{LLS} , as shown in appendix D, so that fixing $\alpha_{\text{LLS}} = 0$ could produce a low n_P .

Rather than use effective broken power laws for the IGM thermal history as a function of redshift, we have explicit physical models for hydrogen and helium reionization, which include the scale-dependent effects of patchy reionization. As shown in figure 7, the preferred IGM thermal history shows a temperature peak at $z = 2.8$ and thus $T_0(z)$ cannot be described by a power law broken at $z = 3$ as assumed in ref. [22] and many earlier works. Figure 5 shows that this does not affect the flux power spectrum on the scales measured by eBOSS. However, DESI data probes smaller scales, so it is not clear that this will be the case in future.

5.3 Likelihood modifications

We also considered modifications to the likelihood not shown here. First, we increased the observational uncertainty from eBOSS by a uniform factor of two. This increased the posterior uncertainties, but did not significantly resolve the internal tension at low redshift (which is many σ). We also considered removing the high redshift data, with $z > 3.8$, as is done in some earlier analyses [120]. We found that this made little difference as the statistical errors in the high redshift data are large and so they provide little information. We considered removing the largest and smallest scales with cuts in k . Several of the smallest scale bins are highly correlated, and so removing them either led to very poor constraints or had small effects, depending on the scale cut. Removing the largest bins on the largest scales increased the posterior uncertainty, but did not noticeably shift the posteriors. Thus none of these checks show any evidence for an internal tension between scales.

We tested whether a second mean flux rescaling slope would improve the fit to the observed Lyman- α forest flux power (figure 5), especially at lower redshifts and smaller scales. To do this, we added a second mean flux slope to the MCMC sampled parameters, and assigned each to a specific redshift range (we tested this using a redshift pivot of $z = 3$ and $z = 3.6$). Posterior constraints on the other parameters from a chain run using the second mean flux slope were unaffected and the fit was not improved.

6 Conclusions

We have developed a new likelihood and pipeline for the analysis of Lyman- α forest data and run MCMC chains using the Cobaya package [85, 86]. Our likelihood is built on a percent-level accurate emulator using the PRIYA simulations [47]. We use the multi-fidelity emulation technique [54] and a set of high resolution simulations to avoid the need to ‘splice’ together multiple simulations resolving different scales.

We model the Lyman- α forest 1D flux power spectrum from eBOSS [38], and include several simulated and post-processed parameters. Our main cosmological constraints are on the slope and amplitude of the primeval power spectrum on the scales probed by the Lyman- α forest (n_P and A_P). We augment our Lyman- α flux power spectrum likelihood

with information about the IGM thermal history [30]. With this information, we constrain the start, end, and heating rate for a patchy model of helium reionization ($z_i^{\text{He II}}$, $z_f^{\text{He II}}$, α_q), as well as the mean optical depth and its evolution with redshift (τ_0 and $d\tau_0$). Our likelihood includes corrections to the Lyman- α forest flux power spectrum from correlated Si III absorption and from the presence of Damped Lyman- α systems.

We found that the lowest redshift bins in the eBOSS flux power spectrum, at $z = 2.2$ and $z = 2.4$, produced results which were discrepant with those from higher redshifts. The flux power spectrum from the DESI early data release is also discrepant with eBOSS at these redshifts. It thus seems likely that this discrepancy is due to an as-yet unidentified systematic in the eBOSS pipeline at $z < 2.6$. We found that the discrepancy was reduced if the DLA finder efficiency was allowed to be redshift dependent. However, an unmodeled astrophysical effect, perhaps connected with AGN feedback, is still possible.

We added data on the IGM temperature at mean density to our chains, improving constraints on the thermal history. Our best-fit parameters from the flux power spectrum alone prefer a peak IGM temperature data much lower than the constraints of ref. [30], although the eBOSS flux power spectrum only weakly constrains helium reionization. It is possible that these constraints are affected by assuming the Planck value of σ_8 , which is inconsistent with our results. The low temperature is driven by improving the fit at $z = 2.6$, so it is also possible some residual effect connected with the tension between the low redshift bins is biasing the flux power only chains. Our least constraining and so most conservative constraints are those from the eBOSS flux power spectrum alone.

When removing the lowest redshift bins from the analysis, we find, for a primeval power spectrum with a pivot scale of 0.78 1/Mpc:

- A power spectrum slope of $n_P = 1.009_{-0.018}^{+0.027}$ from the eBOSS flux power alone and $n_P = 0.983 \pm 0.020$ when adding the IGM temperature data, both in reasonable ($2 - \sigma$) agreement with Planck.
- A power spectrum amplitude $A_P = 1.69_{-0.15}^{+0.14}$ from the eBOSS flux power spectrum, which translates to $A_s = (1.65_{-0.14}^{+0.12} \times 10^{-9})$ or $\sigma_8 = 0.733_{-0.029}^{+0.026}$, approximately $2-3\sigma$ lower than measurements from Planck, but in agreement with some other measurements from weak lensing or galaxy surveys.
- An early start and late finish to helium reionization, beginning at $z > 4.00$ and ending at $z = 2.765_{-0.093}^{+0.080}$ once IGM temperature data is included. Our data is consistent at 95% confidence with helium reionization still being underway at the lowest redshift eBOSS data we use, $z = 2.6$. When IGM temperature data is not included the chain prefers models where the IGM temperature does not increase substantially and thus cannot place strong constraints on the start and end of helium reionization.
- Weak constraints on hydrogen reionization and the matter density $\Omega_M h^2$.

In terms of the reduced likelihood of [17, 38, 91], we find a linear power of $\Delta_L^2 = 0.302_{-0.027}^{+0.024}$ and $n_{\text{eff}} = -2.264_{-0.018}^{+0.026}$ from our $z > 2.6$ flux power chains. The Δ_L^2 constraint is in agreement with earlier analyses, while n_{eff} is $4 - \sigma$ larger. When the IGM temperature

data is included, it reduces both the best-fit Δ_L^2 and n_{eff} to $0.267^{+0.018}_{-0.023}$ and -2.288 ± 0.020 , respectively. Ref. [121] noted that earlier analyses were in tension with the values implied by Planck of $\Delta_L^2 \sim 0.35$ and $n_{\text{eff}} \sim -2.305$. In our analysis the tension in n_{eff} is largely eliminated, while the Δ_L^2 tension remains.

In future work, we will combine our Lyman- α likelihood with other cosmological information, in particular the Planck CMB and Baryon Acoustic Oscillation measurements, with which we can constrain several extensions to the Λ CDM model. While we defer quantitative constraints to later papers, we are able to qualitatively discuss the constraints we expect. We will be able to constrain the running of the spectral index, $\alpha_s = \frac{dn_s}{d \ln k}$. Constraints on α_s come from the difference between the spectral index measured by the CMB on large scales and the spectral index measured by the Lyman- α forest on small scales. The sum of neutrino masses can be constrained via a comparison between the power spectrum amplitude on CMB and Lyman- α scales [122]. Our constraints are in strong tension with those of ref. [22], which will certainly affect neutrino mass constraints from the Lyman- α forest. Since our preferred power spectrum amplitude is lower than that of Planck, we will likely have a preference for a non-zero neutrino mass, although the strength of the preference and the value preferred is yet to be determined.

We will also incorporate new data sets into our likelihood. We will examine the posterior parameter constraints from the DESI EDR flux power spectrum. The statistical power of DESI EDR is currently weaker than that of eBOSS, but it is able to measure smaller scales ($k_F < 0.05$ s/km rather than $k_F < 0.02$ s/km for eBOSS). The higher resolution data may also improve the internal consistency of the dataset at $z < 2.6$. Finally, we can perform a joint analysis of eBOSS and the high resolution Lyman- α forest flux power spectra from ref. [45]. These smaller scales would directly measure the parameters of helium reionization, without the intermediate step of the IGM temperature, allowing an end-to-end validation of the consistency of our modelling of large and small scales.

Acknowledgments

MAF is supported by a National Science Foundation Graduate Research Fellowship under grant No. DGE-1326120. SB was supported by NSF grant AST-1817256 and by NASA-80NSSC21K1840. MQ was supported by NSF grant AST-2107821. MFH is supported by a NASA FINESST grant No. ASTRO20-0022. Computing resources were provided by Frontera LRAC AST21005. The authors acknowledge the Frontera computing project at the Texas Advanced Computing Center (TACC) for providing HPC and storage resources that have contributed to the research results reported within this paper. Frontera is made possible by National Science Foundation award OAC-1818253. URL: <http://www.tacc.utexas.edu>. Analysis computations were performed using the resources of the UCR HPCC, which were funded by grants from NSF (MRI-2215705, MRI-1429826) and NIH (1S10OD016290-01A1).

A Leave-one-out versus emulator error

In this appendix we evaluate the impact of including the Gaussian Process interpolation error, σ_{GP} , on the posterior parameters, as discussed in equation (3.1). Figure 10 compares the effect of including emulator errors, showing the training samples, GP emulator errors

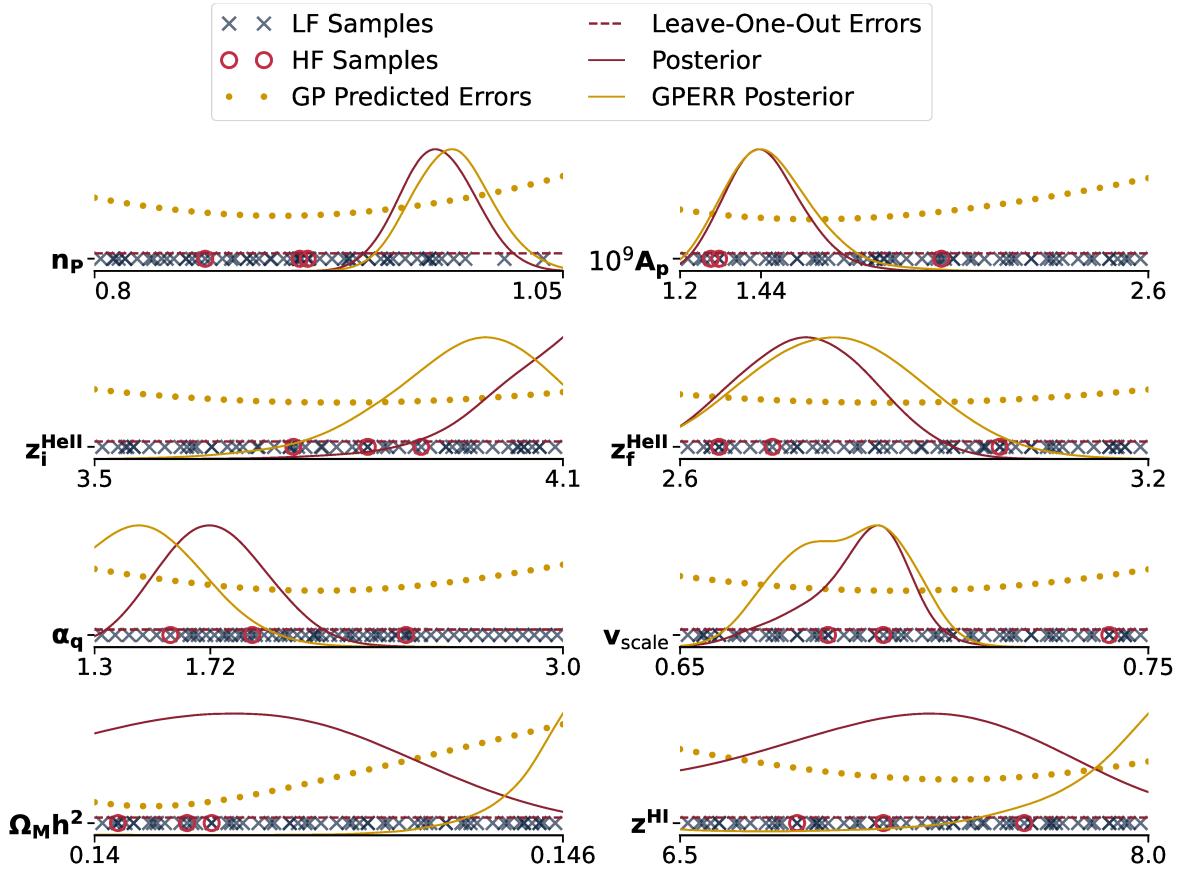


Figure 10. Emulator error and leave-one-out errors across parameter space. For eight of the input parameters, the training samples (grey crosses for LF, red circles for HF), GP emulator errors (yellow dots), and scale-averaged leave-one-out errors (red dashed) are shown. All chains contain the IGM temperature data, and the eBOSS flux power spectrum from $z = 2.6$ –4.6. Shown are 1D marginalised posteriors for the chains with the default likelihood (red) compared to chains run adding the GP emulator error to the likelihood (yellow).

and scale-averaged leave-one-out errors. The leave-one-out error is independent of position in parameter space, whereas the GP error is larger towards the edge of parameter space. The largest effect is on the matter density $\Omega_M h^2$ and the hydrogen reionization midpoint, which are very weakly measured by the data and so end up dominated by the GP error prior. There are smaller shifts in helium reionization, preferring a slightly later start and extra heating, albeit with shifts less than $1.5 - \sigma$. The cosmological parameters are unchanged.

B BOSS DR9 data

In this section we compare posteriors obtained using the flux power spectrum from ref. [112], which is based on BOSS DR9 quasar spectra. Figure 11 shows the posteriors for chains run with DR14, with both reduced and full redshift ranges, and chains with DR9, again with the reduced and full redshift ranges. We show chains using only the flux power spectrum likelihood, to emphasise any differences between the datasets.

The tension between $z < 2.6$ and $z \geq 2.6$ remains in DR9, and the parameter shifts by dropping the lowest redshift bins are similar. DR9 measures a similar A_P to DR14, and thus

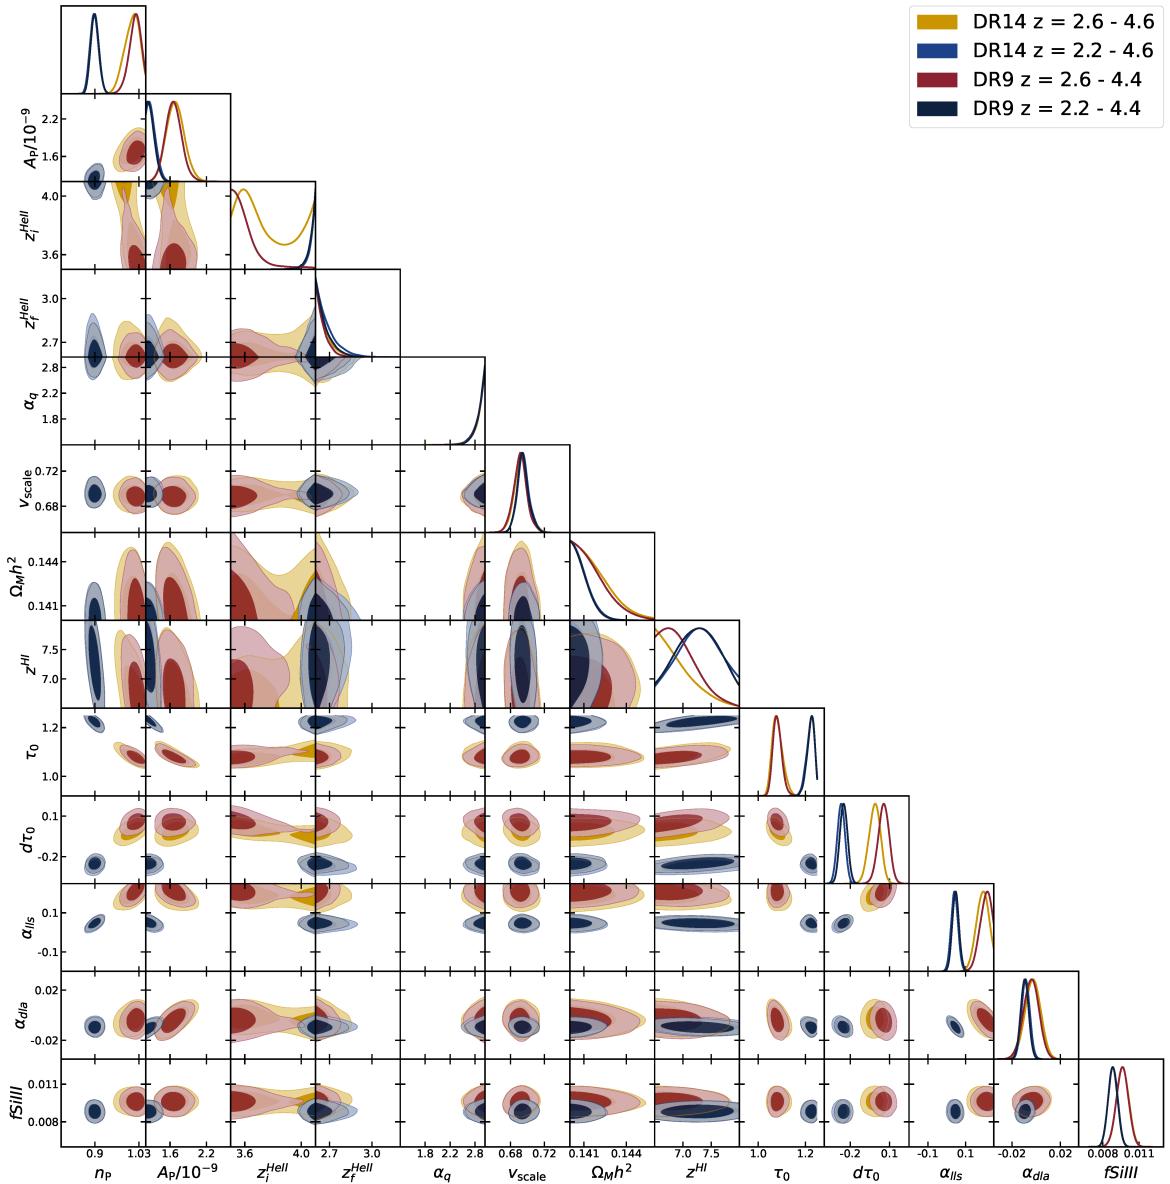


Figure 11. Posteriors for chains run using observations from the earlier SDSS data release, DR9, for the reduced redshift range (gold) and full redshift range (blue), compared to our main chains using DR14 with the reduced redshift range (black) and full redshift range (red).

also implies a low σ_8 . The τ_0 posteriors are also similar. The best-fit helium reionization model is similar in both datasets. Although DR14 includes many more quasars, the covariance matrix is often dominated by systematic error. At $z = 2.6$ DR14 and DR9 have very similar errors, while DR14 has smaller measurement uncertainty at $z \geq 3$. The posterior uncertainties for DR9 are for most parameters similar to those for DR14. The largest parameter changes are in the posterior values of $d\tau_0$ and z_i^{HeII} , although both are less than $1 - \sigma$. Particularly, DR9 prefers a low value of z_i^{HeII} , plausibly because of the inclusion of high redshift data in DR14. The good consistency between our posterior constraints in DR9 and DR14 is a validation of the robustness of our analysis.

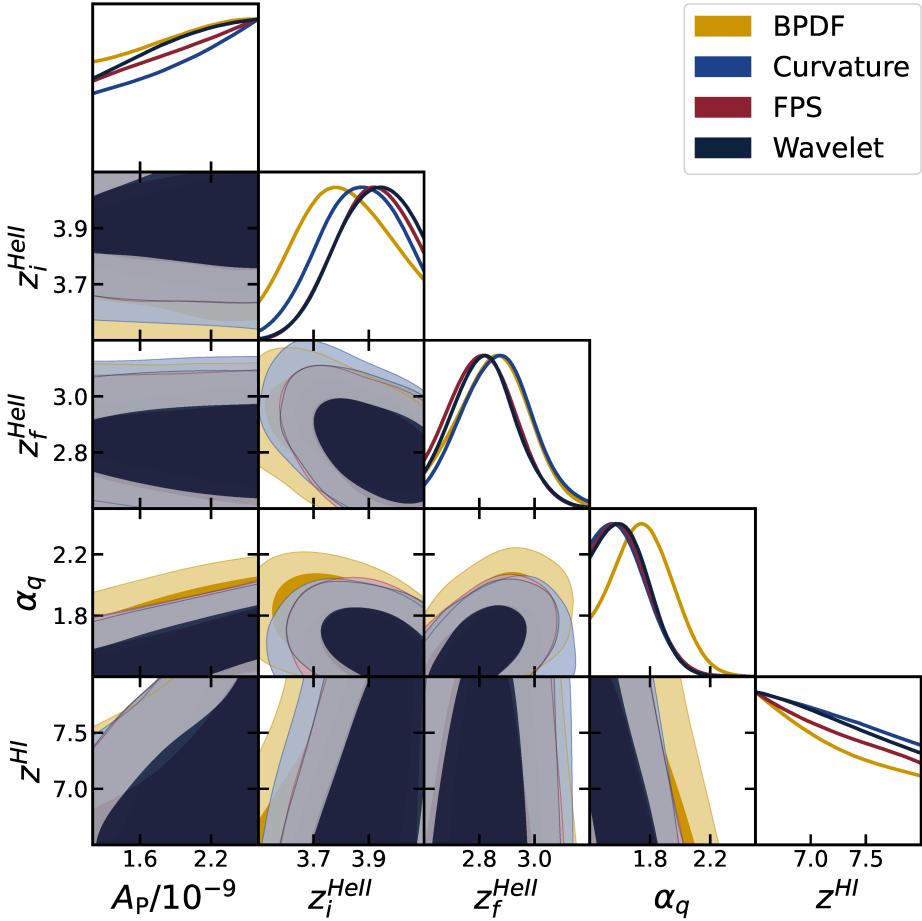


Figure 12. Posteriors for chains run with only the IGM temperature likelihood. Shown are chains using each of the four observational measurements of the IGM temperature: using the flux power spectrum (blue), using the Doppler width distribution (BPDF, yellow), using the curvature statistic (red), and using a wavelet decomposition (black). The main results of this work use the flux power derived IGM temperatures.

C Posterior constraints from the IGM temperature alone

In this section, we present results from chains run using only the IGM temperature likelihood. Shown in figure 12 are four chains, each using one of the observational measurements of the IGM temperature at mean density, derived from different Lyman- α forest summary statistics: the flux power spectrum, the Doppler width distribution, the curvature statistic, and a wavelet decomposition. Most of the cosmology parameters are entirely unconstrained by the IGM temperature and are omitted from the figure. We show A_P for reference. The three He II reionization parameters are well constrained by the IGM temperature history. There is very little difference between the different IGM temperature observations, with the BPDF derived temperature differing the most, specifically preferring a later start to He II reionization, and less heating, corresponding to a lower IGM temperature. However, differences are well within 1σ .

The midpoint of H I reionization has a marginal preference for a late midpoint, but is very weakly constrained. The IGM temperature provides information on z_{HI} as the IGM cools from the completion of H I reionization, setting the temperature before the onset of

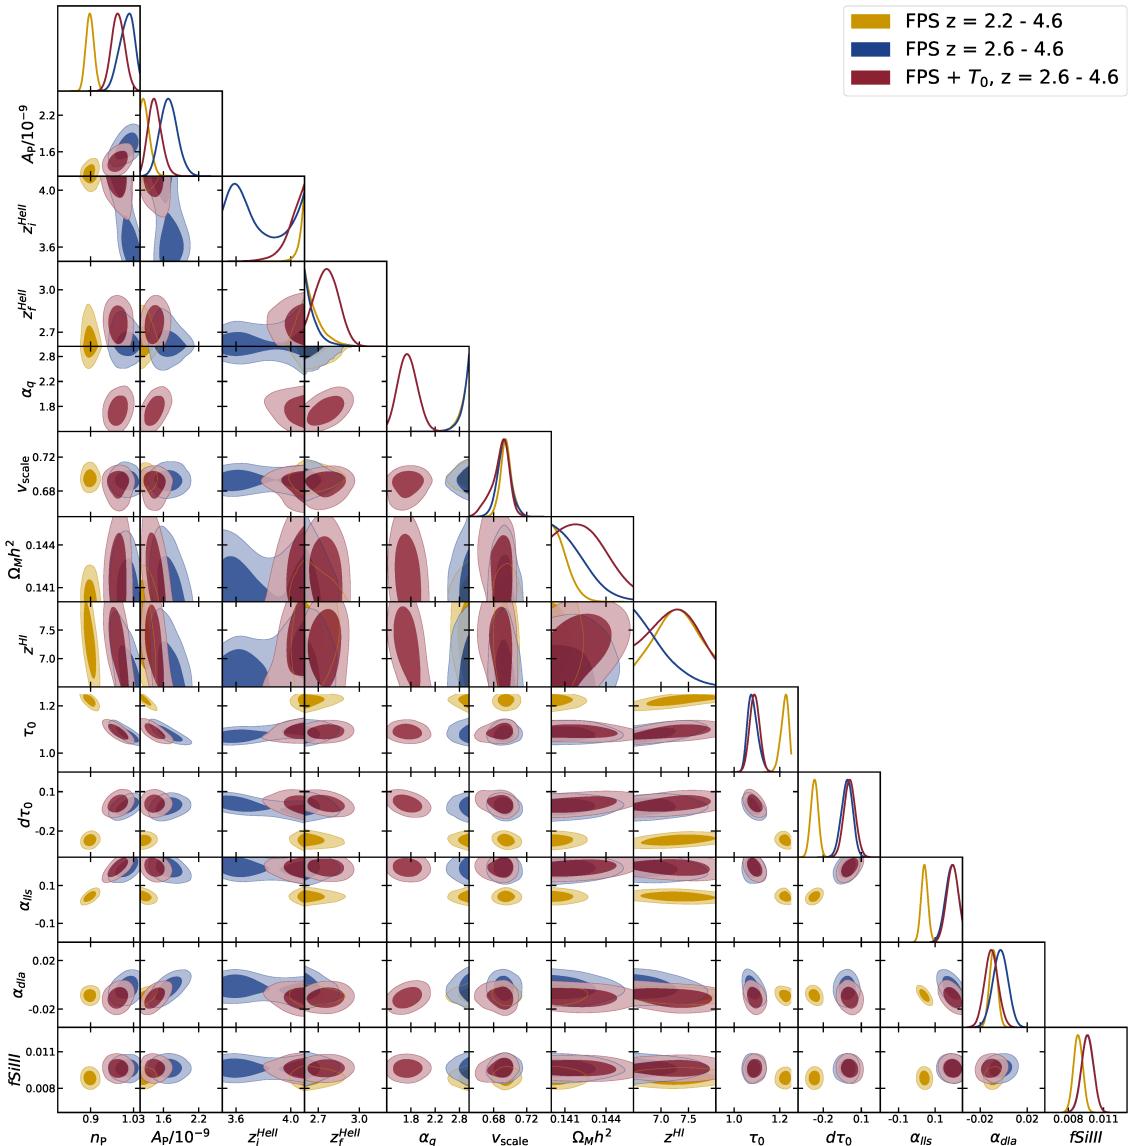


Figure 13. Posteriors for the full set of simulation parameters, both cosmological and astrophysical. Shown are the same chains from the figures in section 4, thus the correlations between the parameters in figure 4 and those in figure 6 are the only new information here. ‘FPS $z = 2.2\text{--}4.6$ ’ (gold) uses the full redshift range eBOSS flux power spectrum dataset, ‘FPS $z = 2.6\text{--}4.6$ ’ (blue) uses a reduced redshift range eBOSS dataset flux power spectrum dataset, which removes the internal tension (section 5.1). The third chain, ‘FPS $+T_0, z = 2.6\text{--}4.6$ ’ (red), uses the limited range eBOSS dataset but adds the IGM temperature constraints.

helium reionization. Incorporating measurements of the IGM temperature at $z > 3.8$ [e.g. 123] could substantially improve these constraints and we may do so in future work.

D Full posteriors

Figure 13 presents the full posteriors, including the correlations between the cosmology and astrophysics parameter sets, using the same chains discussed extensively in section 4. Correlations are discussed in section 4.3.

References

- [1] J.E. Gunn and B.A. Peterson, *On the Density of Neutral Hydrogen in Intergalactic Space*, *Astrophys. J.* **142** (1965) 1633 [[INSPIRE](#)].
- [2] R.A.C. Croft, D.H. Weinberg, N. Katz and L. Hernquist, *Recovery of the Power Spectrum of Mass Fluctuations from Observations of the Ly α Forest*, *Astrophys. J.* **495** (1998) 44 [[astro-ph/9708018](#)] [[INSPIRE](#)].
- [3] T. Theuns et al., *P^3M -SPH simulations of the Ly α forest*, *Mon. Not. Roy. Astron. Soc.* **301** (1998) 478 [[astro-ph/9805119](#)] [[INSPIRE](#)].
- [4] P. McDonald et al., *The Observed Probability Distribution Function, Power Spectrum, and Correlation Function of the Transmitted Flux in the Ly α Forest*, *Astrophys. J.* **543** (2000) 1 [[astro-ph/9911196](#)] [[INSPIRE](#)].
- [5] L. Hui et al., *On Estimating the QSO Transmission Power Spectrum*, *Astrophys. J.* **552** (2001) 15 [[astro-ph/0005049](#)] [[INSPIRE](#)].
- [6] M. Viel et al., *Probing the intergalactic medium with the Ly α forest along multiple lines of sight to distant QSOs*, *Mon. Not. Roy. Astron. Soc.* **329** (2002) 848 [[astro-ph/0105233](#)] [[INSPIRE](#)].
- [7] X.-H. Fan et al., *Constraining the Evolution of the Ionizing Background and the Epoch of Reionization with $z \sim 6$ Quasars. II. A Sample of 19 Quasars*, *Astron. J.* **132** (2006) 117 [[astro-ph/0512082](#)] [[INSPIRE](#)].
- [8] M. Viel and M.G. Haehnelt, *Cosmological and astrophysical parameters from the Sloan Digital Sky Survey flux power spectrum and hydrodynamical simulations of the Lyman α forest*, *Mon. Not. Roy. Astron. Soc.* **365** (2006) 231 [[astro-ph/0508177](#)] [[INSPIRE](#)].
- [9] SDSS collaboration, *The Ly α Forest Power Spectrum from the Sloan Digital Sky Survey*, *Astrophys. J. Suppl.* **163** (2006) 80 [[astro-ph/0405013](#)] [[INSPIRE](#)].
- [10] BOSS collaboration, *The Lyman- α forest in three dimensions: measurements of large scale flux correlations from BOSS 1st-year data*, *JCAP* **09** (2011) 001 [[arXiv:1104.5244](#)] [[INSPIRE](#)].
- [11] BOSS collaboration, *Measurement of baryon acoustic oscillations in the Lyman- α forest fluctuations in BOSS data release 9*, *JCAP* **04** (2013) 026 [[arXiv:1301.3459](#)] [[INSPIRE](#)].
- [12] EBOSS collaboration, *The Completed SDSS-IV Extended Baryon Oscillation Spectroscopic Survey: Baryon Acoustic Oscillations with Ly α Forests*, *Astrophys. J.* **901** (2020) 153 [[arXiv:2007.08995](#)] [[INSPIRE](#)].
- [13] A. Cuceu et al., *Constraints on the Cosmic Expansion Rate at Redshift 2.3 from the Lyman- α Forest*, *Phys. Rev. Lett.* **130** (2023) 191003 [[arXiv:2209.13942](#)] [[INSPIRE](#)].
- [14] M. Viel, J. Schaye and C.M. Booth, *The impact of feedback from galaxy formation on the Lyman α transmitted flux*, *Mon. Not. Roy. Astron. Soc.* **429** (2013) 1734 [[arXiv:1207.6567](#)] [[INSPIRE](#)].
- [15] S. Chabanier et al., *The impact of AGN feedback on the 1D power spectra from the Ly α forest using the Horizon-AGN suite of simulations*, *Mon. Not. Roy. Astron. Soc.* **495** (2020) 1825 [[arXiv:2002.02822](#)] [[INSPIRE](#)].
- [16] M. Viel, M.G. Haehnelt and V. Springel, *Inferring the dark matter power spectrum from the Lyman α forest in high-resolution QSO absorption spectra*, *Mon. Not. Roy. Astron. Soc.* **354** (2004) 684 [[astro-ph/0404600](#)] [[INSPIRE](#)].
- [17] SDSS collaboration, *The Linear Theory Power Spectrum from the Ly α Forest in the Sloan Digital Sky Survey*, *Astrophys. J.* **635** (2005) 761 [[astro-ph/0407377](#)] [[INSPIRE](#)].
- [18] M. Viel, M.G. Haehnelt and A. Lewis, *The Lyman α forest and WMAP year three*, *Mon. Not. Roy. Astron. Soc.* **370** (2006) L51 [[astro-ph/0604310](#)] [[INSPIRE](#)].

- [19] SDSS collaboration, *Cosmological parameter analysis including SDSS Ly α forest and galaxy bias: Constraints on the primordial spectrum of fluctuations, neutrino mass, and dark energy*, *Phys. Rev. D* **71** (2005) 103515 [[astro-ph/0407372](#)] [[INSPIRE](#)].
- [20] U. Seljak, A. Slosar and P. McDonald, *Cosmological parameters from combining the Lyman- α forest with CMB, galaxy clustering and SN constraints*, *JCAP* **10** (2006) 014 [[astro-ph/0604335](#)] [[INSPIRE](#)].
- [21] C. Yèche, N. Palanque-Delabrouille, J. Baur and H. du Mas des Bourboux, *Constraints on neutrino masses from Lyman-alpha forest power spectrum with BOSS and XQ-100*, *JCAP* **06** (2017) 047 [[arXiv:1702.03314](#)] [[INSPIRE](#)].
- [22] N. Palanque-Delabrouille et al., *Hints, neutrino bounds, and WDM constraints from SDSS DR14 Lyman- α and Planck full-survey data*, *JCAP* **04** (2020) 038 [[arXiv:1911.09073](#)] [[INSPIRE](#)].
- [23] M. Garny, T. Konstandin, L. Sagunski and M. Viel, *Neutrino mass bounds from confronting an effective model with BOSS Lyman- α data*, *JCAP* **03** (2021) 049 [[arXiv:2011.03050](#)] [[INSPIRE](#)].
- [24] J.S. Bolton et al., *Possible evidence for an inverted temperature-density relation in the intergalactic medium from the flux distribution of the Ly α forest*, *Mon. Not. Roy. Astron. Soc.* **386** (2008) 1131 [[arXiv:0711.2064](#)] [[INSPIRE](#)].
- [25] J.S. Bolton, G.D. Becker, M.G. Haehnelt and M. Viel, *A consistent determination of the temperature of the intergalactic medium at redshift $\langle z \rangle = 2.4$* , *Mon. Not. Roy. Astron. Soc.* **438** (2014) 2499 [[arXiv:1308.4411](#)] [[INSPIRE](#)].
- [26] F. Nasir, J.S. Bolton and G.D. Becker, *Inferring the IGM thermal history during reionization with the Lyman α forest power spectrum at redshift $z \simeq 5$* , *Mon. Not. Roy. Astron. Soc.* **463** (2016) 2335 [[arXiv:1605.04155](#)] [[INSPIRE](#)].
- [27] M. Walther, J. Oñorbe, J.F. Hennawi and Z. Lukić, *New Constraints on IGM Thermal Evolution from the Ly α Forest Power Spectrum*, *Astrophys. J.* **872** (2019) 13 [[arXiv:1808.04367](#)] [[INSPIRE](#)].
- [28] E. Boera, G.D. Becker, J.S. Bolton and F. Nasir, *Revealing Reionization with the Thermal History of the Intergalactic Medium: New Constraints from the Ly α Flux Power Spectrum*, *Astrophys. J.* **872** (2019) 101 [[arXiv:1809.06980](#)] [[INSPIRE](#)].
- [29] X. Wu et al., *Imprints of temperature fluctuations on the $z \sim 5$ Lyman- α forest: a view from radiation-hydrodynamic simulations of reionization*, *Mon. Not. Roy. Astron. Soc.* **490** (2019) 3177 [[arXiv:1907.04860](#)] [[INSPIRE](#)].
- [30] P. Gaikwad, R. Srianand, M.G. Haehnelt and T.R. Choudhury, *A consistent and robust measurement of the thermal state of the IGM at $2 \leq z \leq 4$ from a large sample of Ly α forest spectra: evidence for late and rapid He II reionization*, *Mon. Not. Roy. Astron. Soc.* **506** (2021) 4389 [[arXiv:2009.00016](#)] [[INSPIRE](#)].
- [31] B. Villasenor, B. Robertson, P. Madau and E. Schneider, *Inferring the Thermal History of the Intergalactic Medium from the Properties of the Hydrogen and Helium Ly α Forest*, *Astrophys. J.* **933** (2022) 59 [[arXiv:2111.00019](#)] [[INSPIRE](#)].
- [32] M. Viel et al., *Constraining warm dark matter candidates including sterile neutrinos and light gravitinos with WMAP and the Lyman- α forest*, *Phys. Rev. D* **71** (2005) 063534 [[astro-ph/0501562](#)] [[INSPIRE](#)].
- [33] M. Viel, G.D. Becker, J.S. Bolton and M.G. Haehnelt, *Warm dark matter as a solution to the small scale crisis: New constraints from high redshift Lyman- α forest data*, *Phys. Rev. D* **88** (2013) 043502 [[arXiv:1306.2314](#)] [[INSPIRE](#)].

- [34] V. Iršič et al., *New constraints on the free-streaming of warm dark matter from intermediate and small scale Lyman- α forest data*, *Phys. Rev. D* **96** (2017) 023522 [[arXiv:1702.01764](https://arxiv.org/abs/1702.01764)] [[INSPIRE](#)].
- [35] A. Garzilli, A. Magalich, O. Ruchayskiy and A. Boyarsky, *How to constrain warm dark matter with the Lyman- α forest*, *Mon. Not. Roy. Astron. Soc.* **502** (2021) 2356 [[arXiv:1912.09397](https://arxiv.org/abs/1912.09397)] [[INSPIRE](#)].
- [36] K.K. Rogers and H.V. Peiris, *Strong Bound on Canonical Ultralight Axion Dark Matter from the Lyman-Alpha Forest*, *Phys. Rev. Lett.* **126** (2021) 071302 [[arXiv:2007.12705](https://arxiv.org/abs/2007.12705)] [[INSPIRE](#)].
- [37] B. Villasenor, B. Robertson, P. Madau and E. Schneider, *New constraints on warm dark matter from the Lyman- α forest power spectrum*, *Phys. Rev. D* **108** (2023) 023502 [[arXiv:2209.14220](https://arxiv.org/abs/2209.14220)] [[INSPIRE](#)].
- [38] EBOSS collaboration, *The one-dimensional power spectrum from the SDSS DR14 Ly α forests*, *JCAP* **07** (2019) 017 [[arXiv:1812.03554](https://arxiv.org/abs/1812.03554)] [[INSPIRE](#)].
- [39] DESI collaboration, *Overview of the Instrumentation for the Dark Energy Spectroscopic Instrument*, *Astron. J.* **164** (2022) 207 [[arXiv:2205.10939](https://arxiv.org/abs/2205.10939)] [[INSPIRE](#)].
- [40] N.G. Karaçaylı et al., *Optimal 1D Ly α Forest Power Spectrum Estimation — III. DESI early data*, *Mon. Not. Roy. Astron. Soc.* **528** (2024) 3941 [[arXiv:2306.06316](https://arxiv.org/abs/2306.06316)] [[INSPIRE](#)].
- [41] DESI collaboration, *The Dark Energy Spectroscopic Instrument: One-dimensional power spectrum from first Lyman- α forest samples with Fast Fourier Transform*, *Mon. Not. Roy. Astron. Soc.* **526** (2023) 5118 [[arXiv:2306.06311](https://arxiv.org/abs/2306.06311)] [[INSPIRE](#)].
- [42] DESI collaboration, *The DESI Experiment Part I: Science, Targeting, and Survey Design*, [arXiv:1611.00036](https://arxiv.org/abs/1611.00036) [[INSPIRE](#)].
- [43] V. Iršič et al., *The Lyman α forest power spectrum from the XQ-100 Legacy Survey*, *Mon. Not. Roy. Astron. Soc.* **466** (2017) 4332 [[arXiv:1702.01761](https://arxiv.org/abs/1702.01761)] [[INSPIRE](#)].
- [44] A. Day, D. Tytler and B. Kambalur, *Power spectrum of the flux in the Lyman-alpha forest from high-resolution spectra of 87 QSOs*, *Mon. Not. Roy. Astron. Soc.* **489** (2019) 2536 [[INSPIRE](#)].
- [45] N.G. Karaçaylı et al., *Optimal 1D Ly α forest power spectrum estimation — II. KODIAQ, SQUAD, and XQ-100*, *Mon. Not. Roy. Astron. Soc.* **509** (2022) 2842 [[arXiv:2108.10870](https://arxiv.org/abs/2108.10870)] [[INSPIRE](#)].
- [46] M. Esposito et al., *Weighing cosmic structures with clusters of galaxies and the intergalactic medium*, *Mon. Not. Roy. Astron. Soc.* **515** (2022) 857 [[arXiv:2202.00974](https://arxiv.org/abs/2202.00974)] [[INSPIRE](#)].
- [47] S. Bird et al., *PRIYA: a new suite of Lyman- α forest simulations for cosmology*, *JCAP* **10** (2023) 037 [[arXiv:2306.05471](https://arxiv.org/abs/2306.05471)] [[INSPIRE](#)].
- [48] V. Springel et al., *First results from the IllustrisTNG simulations: matter and galaxy clustering*, *Mon. Not. Roy. Astron. Soc.* **475** (2018) 676 [[arXiv:1707.03397](https://arxiv.org/abs/1707.03397)] [[INSPIRE](#)].
- [49] S. Bird et al., *The ASTRID simulation: galaxy formation and reionization*, *Mon. Not. Roy. Astron. Soc.* **512** (2022) 3703 [[arXiv:2111.01160](https://arxiv.org/abs/2111.01160)] [[INSPIRE](#)].
- [50] Y. Ni et al., *The ASTRID simulation: the evolution of supermassive black holes*, *Mon. Not. Roy. Astron. Soc.* **513** (2022) 670 [[arXiv:2110.14154](https://arxiv.org/abs/2110.14154)] [[INSPIRE](#)].
- [51] A. Borde et al., *New approach for precise computation of Lyman- α forest power spectrum with hydrodynamical simulations*, *JCAP* **07** (2014) 005 [[arXiv:1401.6472](https://arxiv.org/abs/1401.6472)] [[INSPIRE](#)].
- [52] S. Bird et al., *An emulator for the Lyman- α forest*, *JCAP* **02** (2019) 050 [[arXiv:1812.04654](https://arxiv.org/abs/1812.04654)] [[INSPIRE](#)].

[53] M.-F. Ho, S. Bird and C.R. Shelton, *Multifidelity emulation for the matter power spectrum using Gaussian processes*, *Mon. Not. Roy. Astron. Soc.* **509** (2021) 2551 [[arXiv:2105.01081](https://arxiv.org/abs/2105.01081)] [[INSPIRE](#)].

[54] M.A. Fernandez, M.-F. Ho and S. Bird, *A multifidelity emulator for the Lyman- α forest flux power spectrum*, *Mon. Not. Roy. Astron. Soc.* **517** (2022) 3200 [[arXiv:2207.06445](https://arxiv.org/abs/2207.06445)] [[INSPIRE](#)].

[55] M. Kennedy, *Predicting the output from a complex computer code when fast approximations are available*, *Biometrika* **87** (2000) 1.

[56] K. Heitmann et al., *The Coyote Universe. II. Cosmological Models and Precision Emulation of the Nonlinear Matter Power Spectrum*, *Astrophys. J.* **705** (2009) 156 [[arXiv:0902.0429](https://arxiv.org/abs/0902.0429)] [[INSPIRE](#)].

[57] K. Heitmann et al., *The Coyote Universe Extended: Precision Emulation of the Matter Power Spectrum*, *Astrophys. J.* **780** (2014) 111 [[arXiv:1304.7849](https://arxiv.org/abs/1304.7849)] [[INSPIRE](#)].

[58] E. Lawrence et al., *The Mira-Titan Universe. II. Matter Power Spectrum Emulation*, *Astrophys. J.* **847** (2017) 50 [[arXiv:1705.03388](https://arxiv.org/abs/1705.03388)] [[INSPIRE](#)].

[59] B. Giblin, M. Cataneo, B. Moews and C. Heymans, *On the road to per cent accuracy — II. Calibration of the non-linear matter power spectrum for arbitrary cosmologies*, *Mon. Not. Roy. Astron. Soc.* **490** (2019) 4826 [[arXiv:1906.02742](https://arxiv.org/abs/1906.02742)] [[INSPIRE](#)].

[60] Euclid collaboration, *Euclid preparation: IX. EuclidEmulator2 — power spectrum emulation with massive neutrinos and self-consistent dark energy perturbations*, *Mon. Not. Roy. Astron. Soc.* **505** (2021) 2840 [[arXiv:2010.11288](https://arxiv.org/abs/2010.11288)] [[INSPIRE](#)].

[61] G. Aricò et al., *The BACCO simulation project: a baryonification emulator with neural networks*, *Mon. Not. Roy. Astron. Soc.* **506** (2021) 4070 [[arXiv:2011.15018](https://arxiv.org/abs/2011.15018)] [[INSPIRE](#)].

[62] S.K. Giri and A. Schneider, *Emulation of baryonic effects on the matter power spectrum and constraints from galaxy cluster data*, *JCAP* **12** (2021) 046 [[arXiv:2108.08863](https://arxiv.org/abs/2108.08863)] [[INSPIRE](#)].

[63] J. Harnois-Déraps, B. Giblin and B. Joachimi, *Cosmic Shear Covariance Matrix in w CDM: Cosmology Matters*, *Astron. Astrophys.* **631** (2019) A160 [[arXiv:1905.06454](https://arxiv.org/abs/1905.06454)] [[INSPIRE](#)].

[64] C.T. Davies et al., *Constraining cosmology with weak lensing voids*, *Mon. Not. Roy. Astron. Soc.* **507** (2021) 2267 [[arXiv:2010.11954](https://arxiv.org/abs/2010.11954)] [[INSPIRE](#)].

[65] T. McClintock et al., *The Aemulus Project. II. Emulating the Halo Mass Function*, *Astrophys. J.* **872** (2019) 53 [[arXiv:1804.05866](https://arxiv.org/abs/1804.05866)] [[INSPIRE](#)].

[66] T. Nishimichi et al., *Dark Quest. I. Fast and Accurate Emulation of Halo Clustering Statistics and Its Application to Galaxy Clustering*, *Astrophys. J.* **884** (2019) 29 [[arXiv:1811.09504](https://arxiv.org/abs/1811.09504)] [[INSPIRE](#)].

[67] S. Bocquet et al., *The Mira-Titan Universe. III. Emulation of the Halo Mass Function*, *Astrophys. J.* **901** (2020) 5 [[arXiv:2003.12116](https://arxiv.org/abs/2003.12116)] [[INSPIRE](#)].

[68] N.S. Kern et al., *Emulating Simulations of Cosmic Dawn for 21 cm Power Spectrum Constraints on Cosmology, Reionization, and X-Ray Heating*, *Astrophys. J.* **848** (2017) 23 [[arXiv:1705.04688](https://arxiv.org/abs/1705.04688)] [[INSPIRE](#)].

[69] A. Cohen, A. Fialkov, R. Barkana and R. Monsalve, *Emulating the Global 21-cm Signal from Cosmic Dawn and Reionization*, [arXiv:1910.06274](https://arxiv.org/abs/1910.06274) [[DOI:10.1093/mnras/staa1530](https://doi.org/10.1093/mnras/staa1530)] [[INSPIRE](#)].

[70] H.T.J. Bevins et al., *GLOBALEMU: a novel and robust approach for emulating the sky-averaged 21-cm signal from the cosmic dawn and epoch of reionization*, *Mon. Not. Roy. Astron. Soc.* **508** (2021) 2923 [[arXiv:2104.04336](https://arxiv.org/abs/2104.04336)] [[INSPIRE](#)].

- [71] C.H. Bye, S.K.N. Portillo and A. Fialkov, *21cmVAE: A Very Accurate Emulator of the 21 cm Global Signal*, *Astrophys. J.* **930** (2022) 79 [[arXiv:2107.05581](https://arxiv.org/abs/2107.05581)] [[INSPIRE](#)].
- [72] K.K. Rogers et al., *Bayesian emulator optimisation for cosmology: application to the Lyman-alpha forest*, *JCAP* **02** (2019) 031 [[arXiv:1812.04631](https://arxiv.org/abs/1812.04631)] [[INSPIRE](#)].
- [73] C. Pedersen et al., *An emulator for the Lyman- α forest in beyond- Λ CDM cosmologies*, *JCAP* **05** (2021) 033 [[arXiv:2011.15127](https://arxiv.org/abs/2011.15127)] [[INSPIRE](#)].
- [74] M. Walther et al., *Simulating intergalactic gas for DESI-like small scale Lyman α forest observations*, *JCAP* **04** (2021) 059 [[arXiv:2012.04008](https://arxiv.org/abs/2012.04008)] [[INSPIRE](#)].
- [75] K.K. Rogers and H.V. Peiris, *General framework for cosmological dark matter bounds using N-body simulations*, *Phys. Rev. D* **103** (2021) 043526 [[arXiv:2007.13751](https://arxiv.org/abs/2007.13751)] [[INSPIRE](#)].
- [76] L. Cabayol-Garcia, J. Chaves-Montero, A. Font-Ribera and C. Pedersen, *A neural network emulator for the Lyman- α forest 1D flux power spectrum*, *Mon. Not. Roy. Astron. Soc.* **525** (2023) 3499 [[arXiv:2305.19064](https://arxiv.org/abs/2305.19064)] [[INSPIRE](#)].
- [77] A. D'Aloisio et al., *Heating of the Intergalactic Medium by Hydrogen Reionization*, *Astrophys. J.* **874** (2019) 154 [[arXiv:1807.09282](https://arxiv.org/abs/1807.09282)] [[INSPIRE](#)].
- [78] V. Springel, R. Pakmor, O. Zier and M. Reinecke, *Simulating cosmic structure formation with the GADGET-4 code*, *Mon. Not. Roy. Astron. Soc.* **506** (2021) 2871 [[arXiv:2010.03567](https://arxiv.org/abs/2010.03567)] [[INSPIRE](#)].
- [79] S. Bird, Y. Feng, C. Pedersen and A. Font-Ribera, *More accurate simulations with separate initial conditions for baryons and dark matter*, *JCAP* **06** (2020) 002 [[arXiv:2002.00015](https://arxiv.org/abs/2002.00015)] [[INSPIRE](#)].
- [80] P. Upton Sanderbeck and S. Bird, *Inhomogeneous He II reionization in hydrodynamic simulations*, *Mon. Not. Roy. Astron. Soc.* **496** (2020) 4372 [[arXiv:2002.05733](https://arxiv.org/abs/2002.05733)] [[INSPIRE](#)].
- [81] T.-S. Kim et al., *An improved measurement of the flux distribution of the Ly α forest in QSO absorption spectra: the effect of continuum fitting, metal contamination and noise properties*, *Mon. Not. Roy. Astron. Soc.* **382** (2007) 1657 [[arXiv:0711.1862](https://arxiv.org/abs/0711.1862)] [[INSPIRE](#)].
- [82] S. Bird, *FSFE: Fake Spectra Flux Extractor*, Astrophysics Source Code Library, [ascl:1710.012](https://ascl.net/1710.012) (2017).
- [83] S. Bird et al., *Reproducing the kinematics of damped Lyman α systems*, *Mon. Not. Roy. Astron. Soc.* **447** (2015) 1834 [[arXiv:1407.7858](https://arxiv.org/abs/1407.7858)] [[INSPIRE](#)].
- [84] A. Paleyes et al., *Emulation of physical processes with Emukit*, [arXiv:2110.13293](https://arxiv.org/abs/2110.13293).
- [85] J. Torrado and A. Lewis, *Cobaya: Code for Bayesian Analysis of hierarchical physical models*, *JCAP* **05** (2021) 057 [[arXiv:2005.05290](https://arxiv.org/abs/2005.05290)] [[INSPIRE](#)].
- [86] J. Torrado and A. Lewis, *Cobaya: Bayesian analysis in cosmology*. Astrophysics Source Code Library, [ascl:1910.019](https://ascl.net/1910.019) (2019).
- [87] BOSS collaboration, *The Baryon Oscillation Spectroscopic Survey of SDSS-III*, *Astron. J.* **145** (2013) 10 [[arXiv:1208.0022](https://arxiv.org/abs/1208.0022)] [[INSPIRE](#)].
- [88] EBOSS collaboration, *The SDSS-IV extended Baryon Oscillation Spectroscopic Survey: Overview and Early Data*, *Astron. J.* **151** (2016) 44 [[arXiv:1508.04473](https://arxiv.org/abs/1508.04473)] [[INSPIRE](#)].
- [89] K.K. Rogers et al., *Simulating the effect of high column density absorbers on the one-dimensional Lyman α forest flux power spectrum*, *Mon. Not. Roy. Astron. Soc.* **474** (2018) 3032 [[arXiv:1706.08532](https://arxiv.org/abs/1706.08532)] [[INSPIRE](#)].
- [90] J.M. O'Meara et al., *The Second Data Release of the KODIAQ Survey*, *Astron. J.* **154** (2017) 114 [[arXiv:1707.07905](https://arxiv.org/abs/1707.07905)].

- [91] C. Pedersen, A. Font-Ribera and N.Y. Gnedin, *Compressing the Cosmological Information in One-dimensional Correlations of the Lyman- α Forest*, *Astrophys. J.* **944** (2023) 223 [[arXiv:2209.09895](https://arxiv.org/abs/2209.09895)] [[INSPIRE](#)].
- [92] A. Lewis, *Efficient sampling of fast and slow cosmological parameters*, *Phys. Rev. D* **87** (2013) 103529 [[arXiv:1304.4473](https://arxiv.org/abs/1304.4473)] [[INSPIRE](#)].
- [93] A. Lewis and S. Bridle, *Cosmological parameters from CMB and other data: A Monte Carlo approach*, *Phys. Rev. D* **66** (2002) 103511 [[astro-ph/0205436](https://arxiv.org/abs/astro-ph/0205436)] [[INSPIRE](#)].
- [94] P. Montero-Camacho, C.M. Hirata, P. Martini and K. Honscheid, *Impact of inhomogeneous reionization on the Lyman- α forest*, *Mon. Not. Roy. Astron. Soc.* **487** (2019) 1047 [[arXiv:1902.02892](https://arxiv.org/abs/1902.02892)] [[INSPIRE](#)].
- [95] G.D. Becker, P.C. Hewett, G. Worseck and J.X. Prochaska, *A Refined Measurement of the Mean Transmitted Flux in the Ly α Forest over $2 < z < 5$ Using Composite Quasar Spectra*, *Mon. Not. Roy. Astron. Soc.* **430** (2013) 2067 [[arXiv:1208.2584](https://arxiv.org/abs/1208.2584)] [[INSPIRE](#)].
- [96] PLANCK collaboration, *Planck 2018 results. VI. Cosmological parameters*, *Astron. Astrophys.* **641** (2020) A6 [*Erratum ibid.* **652** (2021) C4] [[arXiv:1807.06209](https://arxiv.org/abs/1807.06209)] [[INSPIRE](#)].
- [97] J. Lesgourgues, *The Cosmic Linear Anisotropy Solving System (CLASS) I: Overview*, [arXiv:1104.2932](https://arxiv.org/abs/1104.2932) [[INSPIRE](#)].
- [98] ACT collaboration, *The Atacama Cosmology Telescope: A Measurement of the DR6 CMB Lensing Power Spectrum and Its Implications for Structure Growth*, *Astrophys. J.* **962** (2024) 112 [[arXiv:2304.05202](https://arxiv.org/abs/2304.05202)] [[INSPIRE](#)].
- [99] DES collaboration, *Dark Energy Survey Year 3 results: Cosmological constraints from galaxy clustering and weak lensing*, *Phys. Rev. D* **105** (2022) 023520 [[arXiv:2105.13549](https://arxiv.org/abs/2105.13549)] [[INSPIRE](#)].
- [100] M.M. Ivanov, M. Simonović and M. Zaldarriaga, *Cosmological Parameters from the BOSS Galaxy Power Spectrum*, *JCAP* **05** (2020) 042 [[arXiv:1909.05277](https://arxiv.org/abs/1909.05277)] [[INSPIRE](#)].
- [101] E. Abdalla et al., *Cosmology intertwined: A review of the particle physics, astrophysics, and cosmology associated with the cosmological tensions and anomalies*, *JHEAp* **34** (2022) 49 [[arXiv:2203.06142](https://arxiv.org/abs/2203.06142)] [[INSPIRE](#)].
- [102] B. Yu, U. Seljak, Y. Li and S. Singh, *RSD measurements from BOSS galaxy power spectrum using the halo perturbation theory model*, *JCAP* **04** (2023) 057 [[arXiv:2211.16794](https://arxiv.org/abs/2211.16794)] [[INSPIRE](#)].
- [103] G. Worseck, J.X. Prochaska, J.F. Hennawi and M. McQuinn, *Early and Extended Helium Reionization Over More Than 600 Million Years of Cosmic Time*, *Astrophys. J.* **825** (2016) 144 [[arXiv:1405.7405](https://arxiv.org/abs/1405.7405)] [[INSPIRE](#)].
- [104] K. Makan et al., *New Evidence for Extended He II Reionization at $z \gtrsim 3.5$ from He II Lyman Alpha and Beta Transmission Spikes*, *Astrophys. J.* **912** (2021) 38 [[arXiv:2012.07876](https://arxiv.org/abs/2012.07876)] [[INSPIRE](#)].
- [105] M. McQuinn, *The Implications of Gunn-Peterson Troughs in the He II Ly α Forest*, *Astrophys. J. Lett.* **704** (2009) L89 [[arXiv:0905.0481](https://arxiv.org/abs/0905.0481)] [[INSPIRE](#)].
- [106] G. Worseck et al., *The End of Helium Reionization at $z \sim 2.7$ Inferred from Cosmic Variance in HST/COS HeII Ly α Absorption Spectra*, *Astrophys. J. Lett.* **733** (2011) L24 [[arXiv:1103.5752](https://arxiv.org/abs/1103.5752)] [[INSPIRE](#)].
- [107] G. Worseck, F.B. Davies, J.F. Hennawi and J.X. Prochaska, *The Evolution of the HeII-Ionizing Background at Redshifts $2.3 < z < 3.8$ Inferred from a Statistical Sample of 24 HST/COS HeII Ly α Absorption Spectra*, *Astrophys. J.* **875** (2019) 111 [[arXiv:1808.05247](https://arxiv.org/abs/1808.05247)] [[INSPIRE](#)].
- [108] A.M. Morales et al., *The Evolution of the Lyman-alpha Luminosity Function during Reionization*, *Astrophys. J.* **919** (2021) 120 [[arXiv:2101.01205](https://arxiv.org/abs/2101.01205)] [[INSPIRE](#)].

[109] M.-F. Ho, S. Bird and R. Garnett, *Detecting Multiple DLAs per Spectrum in SDSS DR12 with Gaussian Processes*, *Mon. Not. Roy. Astron. Soc.* **496** (2020) 5436 [[arXiv:2003.11036](https://arxiv.org/abs/2003.11036)] [[INSPIRE](#)].

[110] D. Parks, J.X. Prochaska, S. Dong and Z. Cai, *Deep learning of quasar spectra to discover and characterize damped Ly α systems*, *Mon. Not. Roy. Astron. Soc.* **476** (2018) 1151 [[arXiv:1709.04962](https://arxiv.org/abs/1709.04962)].

[111] B. Wang et al., *Deep Learning of Dark Energy Spectroscopic Instrument Mock Spectra to Find Damped Ly α Systems*, *Astrophys. J. Supp.* **259** (2022) 28 [[arXiv:2201.00827](https://arxiv.org/abs/2201.00827)] [[INSPIRE](#)].

[112] BOSS collaboration, *The one-dimensional Ly α forest power spectrum from BOSS*, *Astron. Astrophys.* **559** (2013) A85 [[arXiv:1306.5896](https://arxiv.org/abs/1306.5896)] [[INSPIRE](#)].

[113] N. Palanque-Delabrouille et al., *Neutrino masses and cosmology with Lyman-alpha forest power spectrum*, *JCAP* **11** (2015) 011 [[arXiv:1506.05976](https://arxiv.org/abs/1506.05976)] [[INSPIRE](#)].

[114] A. Amon and G. Efstathiou, *A non-linear solution to the S_8 tension?*, *Mon. Not. Roy. Astron. Soc.* **516** (2022) 5355 [[arXiv:2206.11794](https://arxiv.org/abs/2206.11794)] [[INSPIRE](#)].

[115] J. Schaye et al., *The physics driving the cosmic star formation history*, *Mon. Not. Roy. Astron. Soc.* **402** (2010) 1536 [[arXiv:0909.5196](https://arxiv.org/abs/0909.5196)] [[INSPIRE](#)].

[116] J.X. Prochaska, G. Worseck and J.M. O'Meara, *A Direct Measurement of the Intergalactic Medium Opacity to H I Ionizing Photons*, *Astrophys. J. Lett.* **705** (2009) L113 [[arXiv:0910.0009](https://arxiv.org/abs/0910.0009)] [[INSPIRE](#)].

[117] G. Worseck and J.X. Prochaska, *GALEX Far-ultraviolet Color Selection of UV-bright High-redshift Quasars*, *Astrophys. J.* **728** (2011) 23 [[arXiv:1004.3347](https://arxiv.org/abs/1004.3347)] [[INSPIRE](#)].

[118] M. Fumagalli, J.M. O'Meara, J.X. Prochaska and G. Worseck, *Dissecting the properties of optically-thick hydrogen at the peak of cosmic star formation history*, *Astrophys. J.* **775** (2013) 78 [[arXiv:1308.1101](https://arxiv.org/abs/1308.1101)] [[INSPIRE](#)].

[119] N. Palanque-Delabrouille et al., *Constraint on neutrino masses from SDSS-III/BOSS Ly α forest and other cosmological probes*, *JCAP* **02** (2015) 045 [[arXiv:1410.7244](https://arxiv.org/abs/1410.7244)] [[INSPIRE](#)].

[120] S. Bird, H.V. Peiris, M. Viel and L. Verde, *Minimally parametric power spectrum reconstruction from the Lyman α forest*, *Mon. Not. Roy. Astron. Soc.* **413** (2011) 1717 [[arXiv:1010.1519](https://arxiv.org/abs/1010.1519)] [[INSPIRE](#)].

[121] K.K. Rogers and V. Poulin, *5 σ tension between Planck cosmic microwave background and eBOSS Lyman-alpha forest and constraints on physics beyond Λ CDM*, [arXiv:2311.16377](https://arxiv.org/abs/2311.16377) [[INSPIRE](#)].

[122] C. Pedersen et al., *Massive neutrinos and degeneracies in Lyman-alpha forest simulations*, *JCAP* **04** (2020) 025 [[arXiv:1911.09596](https://arxiv.org/abs/1911.09596)] [[INSPIRE](#)].

[123] P. Gaikwad et al., *Measuring the photoionization rate, neutral fraction, and mean free path of H I ionizing photons at $4.9 \leq z \leq 6.0$ from a large sample of XShooter and ESI spectra*, *Mon. Not. Roy. Astron. Soc.* **525** (2023) 4093 [[arXiv:2304.02038](https://arxiv.org/abs/2304.02038)] [[INSPIRE](#)].