

Reionization and galaxy inference from the high-redshift Ly α forest

Yuxiang Qin¹   ¹★, Andrei Mesinger¹ , Sarah E. I. Bosman^{2,3}  and Matteo Viel^{4,5,6,7}

¹*Scuola Normale Superiore, Piazza dei Cavalieri 7, I-56126 Pisa, Italy*

²*Department of Physics and Astronomy, University College London, Gower Street, London WC1E 6BT, UK*

³*Max-Planck-Institut für Astronomie, Königstuhl 17, D-69117 Heidelberg, Germany*

⁴*SISSA - International School for Advanced Studies, via Bonomea 265, I-34136 Trieste, Italy*

⁵*IFPU, Via Beirut 2, I-34014 Trieste, Italy*

⁶*INFN, Sezione di Trieste, via Valerio 2, 34127 Trieste, Italy*

⁷*INAF – Osservatorio Astronomico di Trieste, Via G.B. Tiepolo 11, I-34131 Trieste, Italy*

Accepted 2021 June 24. Received 2021 June 12; in original form 2021 January 22

ABSTRACT

The transmission of Lyman α (Ly α) in the spectra of distant quasars depends on the density, temperature, and ionization state of the intergalactic medium. Therefore, high-redshift ($z > 5$) Ly α forests could be invaluable in studying the late stages of the epoch of reionization (EoR), as well as properties of the sources that drive it. Indeed, high-quality quasar spectra have now firmly established the existence of large-scale opacity fluctuations at $z > 5$, whose physical origins are still debated. Here, we introduce a Bayesian framework capable of constraining the EoR and galaxy properties by forward-modelling the high- z Ly α forest. Using priors from galaxy and cosmic microwave background observations, we demonstrate that the final overlap stages of the EoR (when >95 per cent of the volume was ionized) should occur at $z < 5.6$, in order to reproduce the large-scale opacity fluctuations seen in forest spectra. However, it is the combination of patchy reionization and the inhomogeneous ultraviolet background that produces the longest Gunn–Peterson troughs. Ly α forest observations tighten existing constraints on the characteristic ionizing escape fraction of galaxies, with the combined observations suggesting $f_{\text{esc}} \approx 7_{-3}^{+4}$ per cent, and disfavoring a strong evolution with the galaxy’s halo (or stellar) mass.

Key words: galaxies: high-redshift–intergalactic medium–dark ages, reionization, first stars–diffuse radiation–early Universe–cosmology: theory.

1 INTRODUCTION

The presence of residual neutral hydrogen in the intergalactic medium (IGM) leads to a series of absorption lines in the spectra of distant galaxies and quasars, corresponding to photons redshifting into Lyman α (Ly α) resonance (Savaglio, Panagia & Padovani 2002; Faucher-Giguère et al. 2008a; Busca et al. 2013; Lee et al. 2013, 2014; Slosar et al. 2013; Iršič et al. 2017). This so-called Ly α forest provides invaluable insights into the structure and properties of the high-redshift IGM (e.g. Bolton et al. 2010, 2017; Lidz et al. 2010; Garzilli et al. 2012; Lee et al. 2015; Puchwein et al. 2015; Gaikwad et al. 2020), the cosmic radiation fields that regulate them (e.g. Bolton & Haehnelt 2007b; Faucher-Giguère et al. 2008b; Becker & Bolton 2013; D’Aloisio et al. 2018), and physical cosmology (e.g. Miralda-Escudé et al. 1996; Croft et al. 2002; Viel et al. 2005, 2013; Delubac et al. 2015; Bautista et al. 2017; Yèche et al. 2017).

Recent years especially have witnessed a large increase in the number of high-quality high-redshift ($z \gtrsim 5 - 6$) quasar spectra, probing well into the Epoch of Reionization (EoR; e.g. Fan et al. 2006; Becker, Rauch & Sargent 2007; Willott et al. 2010; Mortlock et al. 2011; Venemans et al. 2013; Becker et al. 2015; Wu et al. 2015;

Jiang et al. 2016; Bañados et al. 2018; Yang et al. 2020). These are being used to search for increasingly subtle EoR signatures.

Redward of the Ly α emission line, an incomplete EoR can be studied through absorption from the Ly α damping-wing profile (Bolton et al. 2011; Schroeder, Mesinger & Haiman 2013; Greig et al. 2017; Davies et al. 2018b; Greig, Mesinger & Bañados 2019; Wang et al. 2020), while on the blueward side, the additional ionizing contribution from the quasar itself facilitates Ly α transmission in the so-called near zone (e.g. Mesinger, Haiman & Cen 2004; Bolton & Haehnelt 2007a; Lidz et al. 2007; Maselli et al. 2007; Eilers et al. 2017, 2020; Davies, Hennawi & Eilers 2020). Blueward of the quasar near zone, the Ly α forest becomes more opaque and trace amounts of H I are sufficient to saturate transmission.

However, some transmission in the forest is seen even at the highest redshifts. In fact, the sightline-to-sightline scatter in the transmission has been suggested as a potential probe of the EoR (e.g. Becker et al. 2007; Gallerani et al. 2008; Bosman et al. 2018; Eilers, Davies & Hennawi 2018). The significant scatter recently observed on large scales (tens to hundred cMpc) is especially promising. For example, the $110 h^{-1}$ cMpc Gunn–Peterson trough observed in ULAS J0148+0600 cannot be explained by fluctuations in the gas density alone (e.g. Becker et al. 2015).

There have been several theoretical explanations for the sizeable sightline-to-sightline fluctuations, focusing on (i) gas temperature (e.g. D’Aloisio, McQuinn & Trac 2015; Keating, Puchwein &

* E-mail: Yuxiang.L.Qin@Gmail.com

Haehnelt 2018); (ii) rare sources (Chardin et al. 2015; Chardin, Puchwein & Haehnelt 2017; D’Aloisio et al. 2017; Meiksin 2020); (iii) the mean free path of ionizing UV photons (Davies & Furlanetto 2016; D’Aloisio et al. 2018); and (iv) late reionization (Kulkarni et al. 2019; Keating et al. 2020a; Nasir & D’Aloisio 2020).

Unfortunately, there are significant degeneracies between these models (see e.g. Nasir & D’Aloisio 2020), and most previous work has been fairly qualitative – showing a handful of models that agree with the observations to various degrees. Robust quantitative constraints require an exhaustive Bayesian inference framework. For instance, Choudhury, Paranjape & Bosman (2021) recently used seminumerical simulations to constrain the EoR history using forward-modelled cumulative distribution functions (CDFs) of the effective optical depth τ_{eff} . Their seminal work, appearing as this paper was nearing completion, used effective parameters to approximate inhomogeneous recombinations, in addition to assuming a constant mass-to-light ratio for galaxies.

Here, we showcase a fully Bayesian framework for interpreting the Ly α forest at high redshift using 21cm FAST¹ (Mesinger & Furlanetto 2007; Mesinger, Furlanetto & Cen 2011; Murray et al. 2020) and its Markov chain Monte Carlo (MCMC) driver, 21cMMC² (Greig & Mesinger 2015, 2017). Building on the model introduced by Park et al. (2019), we directly sample galaxy properties and forward-model the 3D light-cone of Ly α transmission. Our galaxy-driven model allows us to fold-in observations of high-redshift ultraviolet (UV) luminosity functions (LFs; Bouwens et al. 2015a, 2016; Oesch et al. 2016), in addition to EoR constraints from the Thomson scattering optical depth (τ_c ; Planck Collaboration VI 2020) and quasar dark fraction measurements (McGreer, Mesinger & D’Odorico 2015). All codes developed here are publicly available.

This paper is organized as follows. We present our model including galaxy properties, reionization, IGM temperature, and Ly α forests in Section 2. We then summarize the observed forest sample and free parameters used for inference in Sections 3 and 4. We present our inference results in Section 5 before concluding in Section 6. We assume a Lambda cold dark matter (Λ CDM) cosmology with parameters ($\Omega_m, \Omega_b, \Omega_\Lambda, h, \sigma_8, n_s = 0.31, 0.049, 0.69, 0.68, 0.81, 0.97$) chosen from the *TT, TE, EE + lowE + lensing + BAO* reconstruction in Planck Collaboration VI (2020).

2 MODELLING GALAXIES, THE IGM, AND LY α TRANSMISSION

We first summarize our parametrization of high-redshift galaxies, whose cosmic radiation fields govern the evolution of the high-redshift IGM. We then discuss the corresponding IGM properties, including inhomogeneous reionization, local recombination, and photoionization rates, before presenting the addition of two new outputs of 21 cm FAST developed in this work: post-reionization gas temperature and the Ly α optical depth.

2.1 High-redshift galaxies

We adopt the galaxy model of Park et al. (2019), which relates bulk galaxy properties to the halo mass function through power-law scalings. Such a parametrization can recover high-redshift galaxy UV LFs (see e.g. Moster, Naab & White 2013; Mutch et al. 2016; Sun & Furlanetto 2016; Tacchella et al. 2018; Behroozi et al. 2019; Yung

et al. 2019). Therefore, we can use galaxy observations, in addition to the Ly α forest, to constrain our model parameters. This improves on some previous forest studies (e.g. Mesinger & Furlanetto 2009; Choudhury et al. 2021) that assumed a constant mass-to-light ratio, and allows us to use physically meaningful priors when performing inference (e.g. ionizing escape and stellar mass fractions must be between 0 and 1).

Specifically, the number density of galaxies is described by the halo mass function, dn/dM_{vir} , with an additional factor of $\exp(-M_{\text{turn}}/M_{\text{vir}})$ accounting for inefficient star formation in low-mass haloes due to ineffective cooling, inhomogeneous feedback from reionization photoheating, and/or supernova feedback (Efstathiou 1992; Shapiro, Giroux & Babul 1994; Thoul & Weinberg 1996; Hui & Gnedin 1997; Wyithe & Loeb 2013; Hopkins et al. 2014, 2018; Sobacchi & Mesinger 2014; Mutch et al. 2016; Sun & Furlanetto 2016). The average stellar mass of a galaxy hosted by a halo of mass M_{vir} can be written as $M_* = M_{\text{vir}} \Omega_b / \Omega_m \times \min[1, f_{*,10}(M_{\text{vir}}/10^{10} M_\odot)^{\alpha_*}]$. Similarly, the UV ionizing escape fraction is taken to be $f_{\text{esc}} = \min[1, f_{\text{esc},10}(M_{\text{vir}}/10^{10} M_\odot)^{\alpha_{\text{esc}}}]$. Assuming the characteristic star formation time-scale is proportional to the halo dynamical time-scale, we take $M_*/\dot{M}_* = t_* H^{-1}$ where $H(z)$ represents the Hubble parameter at z . We can then estimate the non-ionizing UV luminosity through $L_{1500} = \dot{M}_* \times 8.7 \times 10^{27} \text{ erg s}^{-1} \text{ Hz}^{-1} M_\odot^{-1} \text{ yr}$ (Madau & Dickinson 2014) and compare against high-redshift observations (e.g. Bouwens et al. 2015a, 2016; Finkelstein et al. 2015; Livermore, Finkelstein & Lotz 2017; Atek et al. 2018; Ishigaki et al. 2018; Oesch et al. 2018; Bhatawdekar et al. 2019).

We thus have six free parameters to characterize the UV ionizing properties of high-redshift galaxies: $M_{\text{turn}}, f_{*,10}, \alpha_*, f_{\text{esc},10}, \alpha_{\text{esc}},$ and t_* . In the next subsection, we summarize how we calculate the IGM properties corresponding to a given galaxy model.

2.2 The IGM

We begin by generating a Gaussian realization of Λ CDM initial conditions in a periodic box with a side length of 500 cMpc and a cell resolution of ~ 0.39 cMpc (i.e. 500 cMpc/1280). Using second-order Lagrangian perturbation theory (Scoccimarro 1998) with high-resolution velocity and density fields,³ we evolve these fields towards lower redshifts and regrid them on to a lower resolution box (~ 1.95 cMpc; i.e. 500 cMpc/256) to calculate ionization fields.

2.2.1 Inhomogeneous reionization

Cosmological reionization by UV photons is effectively bimodal in terms of the neutral fraction: (almost) fully ionized regions surrounding the first groups of galaxies expand into (almost) fully neutral regions, eventually overlapping and completing reionization. Simulating the Ly α forest during the EoR thus requires tracking this inhomogeneous process, as well as estimating the residual H I inside the ionized component of the IGM. We summarize our procedures for these in turn.

³Rather than evolving the density field using lower resolution velocity fields (which was originally implemented in 21cmFAST to conserve RAM, given that velocity fields have much longer correlation lengths), here we use high resolution for all of the initial conditions (setting `PER-TURB.ON.HIGH.RES=TRUE`). This was included to guarantee the density fields are as accurate as possible for simulating the Ly α forest (Watkinson et al., in preparation).

¹<https://github.com/21cmfast/21cmFAST>

²<https://github.com/21cmfast/21cMMC>

To identify ionized cells, we use an excursion-set approach⁴ (Furlanetto, Zaldarriaga & Hernquist 2004). Centred on a cell at (\mathbf{r}, z) , we consider spherical volumes with decreasing radii, R , and corresponding overdensities, $\delta_{R|r,z} \equiv (\rho_b/\bar{\rho}_b - 1)_R$, where ρ_b and $\bar{\rho}_b$ are the baryon density and its cosmic mean. Using the corresponding conditional halo mass function (Barkana & Loeb 2005; Mesinger et al. 2011), we compute the cumulative number of ionizing photons per baryon in this spherical IGM patch by

$$\bar{n}_{\text{ion}} = \int dM_{\text{vir}} \frac{dn}{dM_{\text{vir}}} \exp\left(-\frac{M_{\text{turn}}}{M_{\text{vir}}}\right) M_* \rho_b^{-1} n_\gamma f_{\text{esc}}, \quad (1)$$

where $n_\gamma = 5000$ is the number of ionizing photons intrinsically emitted per stellar baryon (Barkana & Loeb 2005).

We follow Sobacchi & Mesinger (2014) to estimate the average number of recombinations per baryon, \bar{n}_{rec} . Using the probability distribution function (PDF; $dn/d\rho_{\text{sub}}$) of subgrid (unresolved by our simulation cell $\lesssim 1.95$ cMpc) densities (ρ_{sub}) from Miralda-Escudé, Haehnelt & Rees (2000), adjusted for the mean density in the cell, we calculate the recombination rate by

$$\dot{n}_{\text{rec}} = \int d\rho_{\text{sub}} \frac{dn}{d\rho_{\text{sub}}} \alpha_B f_{\text{H}} \rho_b^{-1} \rho_{\text{sub}}^2 (1 - x_{\text{HI}})^2, \quad (2)$$

where α_B , f_{H} , and x_{HI} are the case-B recombination coefficient,⁵ number fraction of hydrogen in the Universe, and the (residual) neutral hydrogen fraction of the subgrid gas element, respectively. The calculation of x_{HI} is presented in Section 2.2.2.

The cell is then considered to be ionized if the cumulative number of ionizing photons is larger than the number of baryons plus recombinations. Specifically, the cell is considered as ionized if at any radius, R ,

$$\bar{n}_{\text{ion}} - 1 \geq \bar{n}_{\text{rec}} = \left\langle \int_{z_{\text{ion}}}^z dz' \frac{dt}{dz'} \dot{n}_{\text{rec}} \right\rangle_R, \quad (3)$$

where z_{ion} is the reionization redshift of a cell and $\langle \rangle_R$ denotes averaging over all cells within the spherical H II region. We also approximate the local mean free path (R_{MFP}) using the largest radius at which this equation is valid. This is strictly true in the early stages of reionization that are not affected by IGM recombinations, but should also be a good approximation for the overlap stages as they likely evolve in a ‘photon-starved’ manner (Bolton & Haehnelt 2007b; Sobacchi & Mesinger 2014). In future work, we will generalize this derivation, which will allow us to extend our models to lower redshifts.

⁴Excursion-set algorithms generally do not conserve photons when H II bubbles overlap (e.g. Zahn et al. 2007; Paranjape & Choudhury 2014). In practice, this translates to a bias in the effective ionizing escape fraction (e.g. Hutter 2018). Using the updated, photon-conserving algorithm of 21 cmFAST v3, we quantify that this is a very minor effect for our model parametrization, resulting in a bias of ~ -0.2 for α_{esc} (Park et al., in preparation). In other words, the recovered posterior without ionizing photon conservation differs from the true posterior including ionizing photon conservation primarily through a translation in one of the parameters: $\alpha_{\text{esc}}^{\text{true}} \rightarrow \alpha_{\text{esc}}^{\text{recovered}} + 0.2$. However, as photon-conservation slows down our calculation by a factor of ~ 2 , we leave this option off in this proof-of-concept study, highlighting the resulting bias in the marginalized posterior of α_{esc} (see more in Section 5).

⁵The recombination coefficient depends on the gas temperature (T_g), $\alpha_B = 2.59 \times 10^{-10} (T_g/\text{K})^{-0.75} \text{ cm}^3 \text{ s}^{-1}$. For computational efficiency, we assume an average temperature of $T_g = 10^4 \text{ K}$ in equation (2) when computing the cumulative number of recombinations, used to identify if a region is ionized or not. We do account for local temperature fluctuations when computing the residual H I fraction inside the ionized IGM, as detailed in Section 2.2.2.

2.2.2 Residual neutral hydrogen inside the ionized IGM

Since a trace amount of neutral hydrogen can obscure all flux at the Ly α transition, it is important to determine the residual neutral hydrogen fraction within the cosmic H II regions. Assuming photoionization equilibrium in the reionized IGM, we evaluate x_{HI} by solving

$$x_{\text{HI}} f_{\text{ion,ss}} \Gamma_{\text{ion}} = \chi_{\text{He II}} n_{\text{H}} (1 - x_{\text{HI}})^2 \alpha_B, \quad (4)$$

where Γ_{ion} is the local photoionization background,⁶ $f_{\text{ion,ss}}$ is a self-shielding factor attenuating Γ_{ion} , $\chi_{\text{He II}} = 1.08$ accounts for singly ionized helium, and n_{H} is the hydrogen number density in the cell.

The photoionization rates before self-shielding are estimated assuming a stellar-driven UV background (UVB) with a spectral index⁷ of $\alpha_{\text{UVB}} \sim 5$ (Thoull & Weinberg 1996)

$$\Gamma_{\text{ion}} = (1 + z)^2 R_{\text{MFP}\sigma_{\text{H}}} \frac{\alpha_{\text{UVB}}}{\alpha_{\text{UVB}} + \beta_{\text{H}}} \frac{\bar{\rho}_b}{m_{\text{p}}} \bar{n}_{\text{ion}}, \quad (5)$$

where $\beta_{\text{H}} \sim 2.75$, m_{p} , and \bar{n}_{ion} are the spectral index of the H I photoionization cross-section, the proton mass, and the mean production rate of ionizing photons evaluated using equation (1) with M_* being replaced by \dot{M}_* . According to radiative transfer simulation results from Rahmati et al. (2013), the self-shielding factor depends on the local density (ρ_b), gas temperature (T_g) as well as the unattenuated photoionization rate (Γ_{ion}) and follows

$$f_{\text{ion,ss}} = 0.98 \left[1 + \left(\frac{\rho_b}{\rho_{\text{ss}}} \right)^{1.64} \right]^{-2.28} + 0.02 \left(1 + \frac{\rho_b}{\rho_{\text{ss}}} \right)^{-0.84}, \quad (6)$$

where

$$\rho_{\text{ss}} = 27 \bar{\rho}_b \left(\frac{T_g}{10^4 \text{ K}} \right)^{0.17} \left(\frac{\Gamma_{\text{ion}}}{10^{-12} \text{ s}^{-1}} \right)^{\frac{2}{3}} \left(\frac{1 + z}{10} \right)^{-3} \quad (7)$$

is the characteristic density for the onset of self-shielding (Schaye 2001).

2.2.3 Gas temperature of the ionized IGM

Inside the H II regions, we calculate the inhomogeneous gas temperature in each simulation cell following McQuinn & Upton Sanderbeck (2016):

$$T_g^\gamma = T_{\text{ion,I}}^\gamma \left[\left(\frac{\mathcal{L}}{\mathcal{L}_{\text{ion}}} \right)^3 \frac{\rho_b}{\rho_{\text{b,ion}}} \right]^{\frac{2\gamma}{3}} \frac{\exp(\mathcal{L}^{2.5})}{\exp(\mathcal{L}_{\text{ion}}^{2.5})} + T_{\text{lim}}^\gamma \frac{\rho_b}{\rho_b}. \quad (8)$$

Here, \mathcal{L} denotes $(1 + z)/7.1$, $\gamma = 1.7$ is the equation of state index, $T_{\text{ion,I}}$ represents the post-I-front temperature, here taken to be a constant ($T_{\text{ion,I}} = 2 \times 10^4 \text{ K}$),⁸ and $T_{\text{lim}} = 1.775 \mathcal{L} \times 10^4 \text{ K}$ refers to the final relaxation temperature. The subscript ‘ion’ indicates the

⁶Our model does not include small-scale fluctuations in Γ_{ion} , due to proximate galaxies. In most of the IGM, such Poisson fluctuations are negligible, and Γ_{ion} is determined by the combined radiation from many galaxies (e.g. Mesinger & Furlanetto 2009; Sadoun, Zheng & Miralda-Escudé 2017).

⁷The effective spectral index for a stellar-driven UVB could be somewhat harder than we assume (e.g. Becker & Bolton 2013; D’Aloisio et al. 2019); however, α_{UVB} is degenerate with the ionizing escape fraction in equation (5), and we treat the later as a free parameter in our analysis.

⁸The exact value depends on I-front speeds, and to a more minor extent, the spectral index of the ionizing background. However, the uncertainty and scatter in this value should be smaller than the scatter resulting from different reionization times (e.g. D’Aloisio et al. 2019; Davies et al. 2019; Zeng & Hirata 2021).

quantity is at the redshift of ionization, z_{ion} . The two terms on the RHS of equation (8) correspond to the initial and final temperatures of photoionized IGM.

The memory of the initial, post-reionization temperature fades within $\Delta z \sim 1-2$ (see the exponential term in equation 8), after which the IGM approaches an equilibrium temperature resulting from the balance between photoheating and various cooling processes (i.e. Hubble expansion, recombination, Compton scattering, and free-free radiation). Because reionization is ‘inside-out’ on large scales (see e.g. the review of Trac & Gnedin 2011), the underdense regions of the IGM reionize late. Thus, large-scale low-density regions at $z \sim 5-6$ should be *hotter* (due to the first term on the RHS of equation 8), resulting in a lower recombination rate coefficient in equation (4) and a corresponding *increase* in the Ly α transmission in the forest. On the other hand, the equilibrium temperature results in a temperature–density relation in which underdense regions are *colder* (see the second term on the RHS of equation 8; see also Hui & Gnedin 1997), and have a correspondingly *lower* Ly α transmission. However, when averaged over large scales, only the former effect remains. Therefore, one would expect voids to correspond to large-scale peaks in the Ly α forest transmission (e.g. D’Aloisio et al. 2015). This correlation between gas temperature and Ly α transmission is, however, opposite of what is inferred from ULAS J0148+0600 based on galaxy counts (Becker et al. 2018; Kashino et al. 2020), suggesting that opacity fluctuations are not dominated by the temperature field for this one sightline.

2.3 Building physical intuition from examples of forward-modelled data

Using the IGM properties described in the previous sections, we compute the corresponding Ly α optical depth in each simulation cell. As we are interested in large-scale effects and use relatively low resolution simulations, we adopt the fluctuating Gunn–Peterson approximation (FGPA;⁹ Gunn & Peterson 1965; Rauch 1998; Weinberg & et al. 1999) and estimate the optical depth by

$$\tau_\alpha = f_{\text{rescale}} \times \sqrt{\frac{3\pi\sigma_T}{8}} c_{f_\alpha} \lambda_\alpha H^{-1} n_{\text{H}} x_{\text{HI}}, \quad (9)$$

where σ_T , $f_\alpha = 0.416$, and $\lambda_\alpha = 1216\text{\AA}$ are the Thomson cross-section, oscillator strength, and Ly α rest-frame wavelength.

To better match observations, numerical simulations usually rescale the optical depth by a constant factor within a given redshift window (either implicitly or explicitly; e.g. Chardin et al. 2017; Bosman et al. 2018; D’Aloisio et al. 2018; Keating et al. 2018), such that the modelled mean Ly α transmission agrees with the measured values: $\tau_\alpha(\lambda_{\text{obs}}) \rightarrow \tau_\alpha(\lambda_{\text{obs}}) \times f_{\text{rescale}}$. This adds additional flexibility to the modelling (e.g. allowing for an arbitrary normalization and redshift evolution of the UV photoionization rate), ameliorates errors in the continuum subtraction, and/or compensates for modelling errors. However, such rescaling wastes the predictive power of the model’s emissivity and its redshift evolution. We further explain how we implement this rescaling factor in Section 4.

Below we showcase our procedure for generating mock Ly α forest data. As a specific illustrative example, we use the maximum a posteriori probability (MAP) model in the *forest* posterior (presented

in Section 5). This model is able to reproduce observations of the high-redshift galaxy UV LFs (Bouwens et al. 2015a; Oesch et al. 2018), the cosmic microwave background (CMB) optical depth (Planck Collaboration VI 2020), and the Ly α forest (Bosman et al. 2018). It corresponds to the following astrophysical model parameters: $(f_{*,10}, \alpha_*, f_{\text{esc},10}, \alpha_{\text{esc}}, M_{\text{turn}}, t_*) = (0.0448, 0.488, 0.0914, -0.298, 7.16 \times 10^8 M_\odot, 0.378)$ as well as $f_{\text{rescale}} = 0.9 + 0.2 \times (z - 5.7)$ for flux normalization (see more in Section 4).

2.3.1 Light-cones and cross-correlation coefficients

In Fig. 1, we show light-cone slices (with a thickness of 1.95 Mpc) through the following fields: (i) density ($1 + \delta$), (ii) large-scale photoionization rate (Γ_{ion}), (iii) gas temperature (T_g), (iv) the order unity fluctuations in the neutral hydrogen fraction from the EoR (x_{HI}), (v) the residual neutral hydrogen fraction inside the ionized IGM, and (vi) the Ly α transmission [$F \equiv \exp(-\tau_\alpha)$]. In addition, to quantify how these fields correlate with the underlying density on various scales, we calculate the cross-correlation coefficient (CCC) defined as the cross-correlation power between density and a field, $P_{\delta x}$, normalized by their autocorrelation power: $\text{CCC}_{\delta x}(k) = P_{\delta x}(k)/(\sqrt{P_{\delta\delta}(k)}\sqrt{P_{xx}(k)})$. Fig. 2 shows CCCs between F , T_g , Γ_{ion} , x_{HI} , and δ using a snapshot at $z \sim 5.8$ (see projections in Fig. 7). We summarize some key trends below.

(i) The underlying density (panel i) plays a key role in determining the UV ionizing background (panel ii) and the reionization morphology (panel iv). Regions near high-density peaks become ionized first with the H II bubbles spreading into the large-scale voids at later times (i.e. ‘inside-out’ reionization; Furlanetto et al. 2004; Iliiev et al. 2006; Lee et al. 2008; Choudhury, Haehnelt & Regan 2009; Friedrich et al. 2011; Bauer et al. 2015; Mesinger 2016; Hutter et al. 2017) and overlapping with each other to complete the EoR (at $z \sim 5.5$ in this model). This can also be seen in the CCCs: x_{HI} anticorrelates with δ and this anticorrelation strengthens towards smaller k (larger scales). A similar trend is also found for the photoionization rate that correlates with density. The characteristic scales of these correlations depend on the model (i.e. the luminosity-weighted galaxy bias) and stage of EoR (e.g. McQuinn et al. 2007; Zahn et al. 2011).

(ii) The IGM temperature (panel iii) shows both large-scale and small-scale structure. Large-scale hot regions are evident in the temperature maps down to $z \sim 5.2$, after the end of the EoR in this model. These trace the large-scale underdensities that were the last to reionize (panel iv). However, large-scale structure is not evident at lower redshifts, with the IGM cooling towards $T_{\text{lim}}(1 + \delta)^{1/\gamma}$ (see equation 8). Indeed, from Fig. 2 we confirm that $\text{CCC}_{\delta T_g}$ at $z = 5.8$ has a non-monotonic relation with k , which rises towards scales smaller than $k \sim 0.2 \text{ cMpc}^{-1}$. Smaller scales are dominated by the temperature–density relation, while larger scales are dominated by EoR morphology.

(iii) The residual neutral hydrogen fraction (panel v) has a complex, multiscale structure as it depends on the density, ionizing background, and temperature (see equation 4). On large scales ($\gtrsim 50 \text{ cMpc}$), the residual x_{HI} anticorrelates with the Γ_{ion} field, while on small scales it can be seen to correlate/anticorrelate with the density/temperature fields (see also D’Aloisio et al. 2015, 2018; Davies & Furlanetto 2016; Keating et al. 2018).

(iv) The Ly α optical depth (see the transmission in panel vi) generally decreases towards lower redshifts with fluctuations following the density field (panel i) and the inferred residual HI fraction (panel v; see equation 9). We can clearly see long patches with Ly α transmission lower than a thousandth on scales larger than 10 cMpc

⁹In Appendix A, we use high-resolution hydrodynamic simulations to quantify the error in modelling the large-scale Ly α transition caused by the FGPA. We include the resulting error covariance in our likelihood calculation, as detailed below.

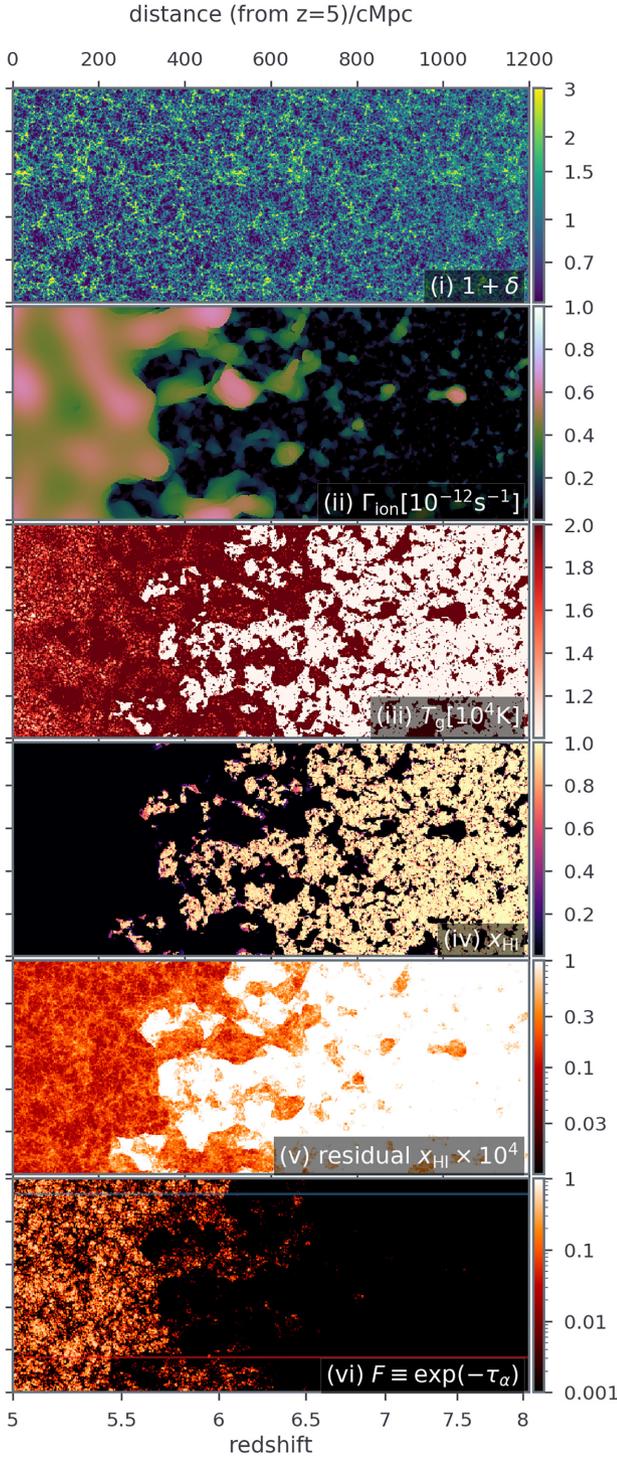


Figure 1. 2D slices through light-cones with a spatial scale of 500 cMpc on the vertical axis and a thickness of 1.95 Mpc, spanning a redshift range of $5 < z \leq 8$ (corresponding to a 1.2 cGpc sight length) for the reference model (i.e. the MAP model in the *forest* posterior in Fig. 5). From top to bottom, the panels correspond to: (i) overdensity ($1 + \delta$); (ii) locally averaged UVB (Γ_{ion} in units of 10^{-12}s^{-1}); (iii) temperature (T_{g}); (iv) neutral hydrogen fraction (x_{HI} on a linear scale between 0 and 1); (v) residual neutral hydrogen fraction (x_{HI} on a logarithmic scale between 10^{-6} and 10^{-4}); and (vi) Ly α transmission ($F \equiv e^{-\tau_{\alpha}}$). The red and blue lines in the bottom panel mark the two sample sightlines shown in Fig. 3.

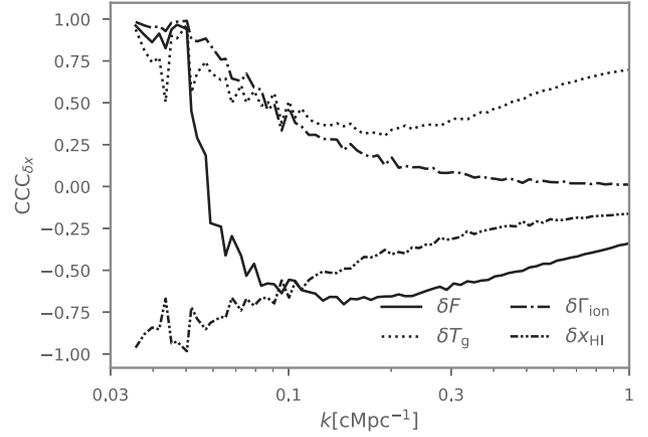


Figure 2. CCCs between the overdensity field (δ) and: (i) temperature (T_{g}), (ii) photonionization rate (Γ_{ion}), (iii) neutral hydrogen fraction (x_{HI}), (iv) Ly α transmission (F), for the reference model at $z = 5.8$ ($\bar{x}_{\text{HI}} \sim 0.14$).

post-reionization with some even surviving at $z \sim 5$. We see that the transmitted flux and density anticorrelate on scales smaller than $k \gtrsim 0.05 \text{ cMpc}^{-1}$ (i.e. $\lesssim 130 \text{ cMpc}$), with the strongest anticorrelation occurring around $k \sim 0.1 \text{ cMpc}^{-1}$ (i.e. $\sim 60 \text{ cMpc}$). On extremely large scales, however, Ly α transmission is tightly correlated with the underlying density through the ionizing background.

2.3.2 Spectra

In Fig. 3, we show two examples of 1D sightlines from this light-cone, denoted by the red (*Ly α Dark*) and blue lines (*Ly α Bright*) in panel (v) of Fig. 1. *Ly α Dark* was chosen as it exhibits a long, $110h^{-1} \text{ cMpc}$ GP trough over the redshift interval $5.52 \leq z \leq 5.88$, consistent with ULAS J0148+0600 (Becker et al. 2015; see the bottom panel of Fig. 3). *Ly α Dark* shows an average transmission of $< 10^{-4}$ over this redshift range, which is much lower than *Ly α Bright* (0.072). The latter corresponds to a more typical sightline at these redshifts.¹⁰

The panels in the figure correspond to the matter overdensity, ionizing background, temperature, neutral fraction, optical depth, and Ly α transmission, from top to bottom. We note that this MAP model resulted in a rescaling factor very close to unity ($f_{\text{rescale}} = 0.86\text{--}0.96$; Section 4). Thus, the theoretical spectra required only minor calibration using this hyperparameter, in order to be consistent with the observed data. We also list the mean values over the redshift interval $5.52 \leq z \leq 5.88$ in the bottom panel. We see that on average *Ly α Dark* has a lower ionization rate as well a lower temperature while the average densities along the two sightlines are both close to the cosmic mean.

We see that the different levels of Ly α transmission between these two sightlines are mostly driven by the neutral hydrogen and UVB fluctuations. In the region corresponding to the long GP trough, *Ly α Dark* pierces through a few remaining cosmic HI patches (with $\Gamma_{\text{ion}} \sim 0$). Although these HI patches combined span less than half of the GP through length, we see from the Γ_{ion} panel that the *ionized IGM between the HI patches is exposed to a below-average UVB* (see also Keating et al. 2020a; Nasir & D’Aloisio 2020). As discussed

¹⁰Among all $256^2 = 65\,536$ lines of sight in this MAP model, we find only 179 with a mean transmission lower than 0.001. This is broadly consistent with current observations (1 out of ~ 300 ; Bosman 2020).

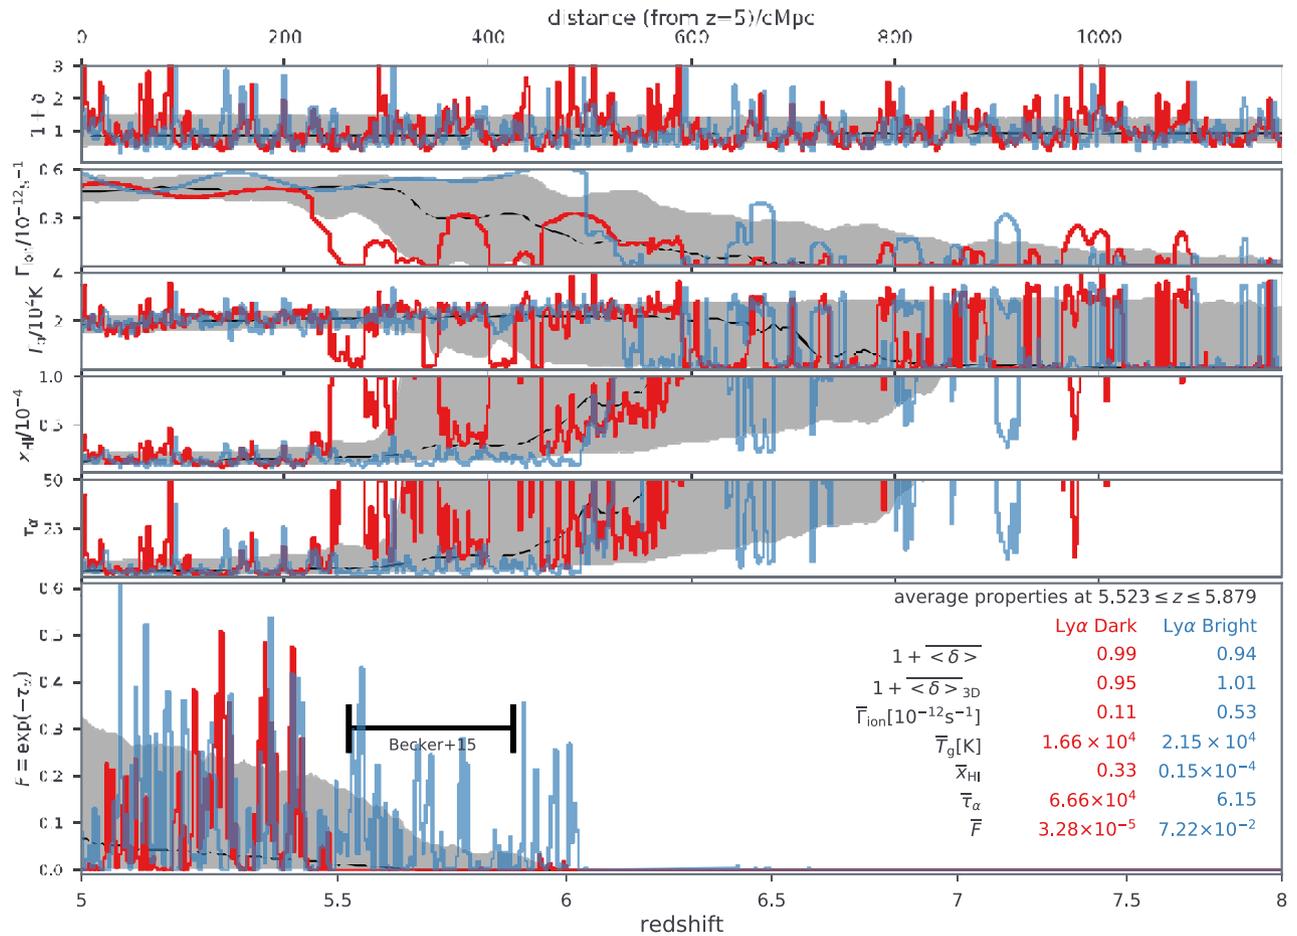


Figure 3. From top to bottom: overdensity ($1 + \delta$), locally averaged UVB (Γ_{ion}), gas temperature (T_g), residual neutral hydrogen fraction (x_{HI}), Ly α optical depth (τ_α), and transmission ($e^{-\tau_\alpha}$), along two particular lines of sight (i.e. *Ly α Dark* in red and *Ly α Bright* in blue; these two sightlines are also marked by the two horizontal lines in panel (v) of Fig. 1). The sightlines are shown with thick coloured curves while the median and [14,86] percentiles of the entire 3D light-cone are presented as the thin black lines and shaded regions. The average properties along the sightline at $5.523 \leq z \leq 5.879$ (the redshift range of the $110 h^{-1}$ cMpc Ly α GP trough from Becker et al. 2015) are listed in the bottom right corner. We also show the density averaged over a 3D volume corresponding to $50 \times 50 \times 110 h^{-3}$ cMpc 3 , highlighting that the large-scale environment of *Ly α Dark* is less dense than that of *Ly α Bright*; see text for more details.

above, both fields correlate with the underlying density on large scales: the regions last to ionize and those with a small UVB both correspond to large-scale underdensities with comparably few star-forming galaxies. This is quantitatively evident when we compute the volume-averaged density around the sightlines, showing that *Ly α Dark* lies in a large-scale underdensity, $1 + \bar{\delta}_{3D} = 0.95$ when averaged over a volume of $50 \times 50 \times 110 (h^{-1} \text{ Mpc})^3$. This picture is also consistent with follow-up observations of J0148+0600 that show a dearth of galaxies at the location of its GP trough (Becker et al. 2018; Kashino et al. 2020).

2.3.3 Fields that determine Ly α fluctuations

Using our MAP (reference) parameter set, we now further quantify which fields are most relevant for large-scale opacity fluctuations. In the left-hand panels of Fig. 4, we show the cumulative probability distributions (CDFs) of the effective optical depth defined as $\tau_{eff} \equiv -\ln[(\exp(-\tau_\alpha))_L]$, where we chose $L = 50 h^{-1}$ cMpc in order to facilitate comparison against published observations (shown as grey shaded regions and further discussed in Section 3). Thick, solid curves correspond to the mean CDFs from our MAP model,

for redshift bins spanning $z = 5-6$.¹¹ The other curves in these panels were constructed by removing fluctuations in a single-component field, and then recomputing the resulting optical depths (though keeping the same normalization factor, $f_{rescale}$). Specifically, we replace the density (*mean density*), photoionization rate (*mean Gamma*), and temperature (*mean Temp*) with their ionized volume-averaged values at each redshift. For *remove EoR*, we remove all cosmic HI patches by assigning to them the mean values of the photoionization rate and temperature in the ionized regions, and then recomputing τ_α in each cell.

In the right-hand panels, we also show the light-cone visualizations of the corresponding change in the optical depth with respect to the reference MAP model (see the transmission in Fig. 1). As expected, ignoring density fluctuations results in a large shift in the CDFs, but does not dramatically impact their shapes – there is only a minor decrease in the abundance of high τ_{eff} regions. Thus

¹¹The MAP model seems to be in (very) mild tension with the data at $z = 6$. We note, however, that preliminary, updated data from the XQR30 large VLT program (PI: V. D’Odorico) are in better agreement with this reference model at $z = 6$.

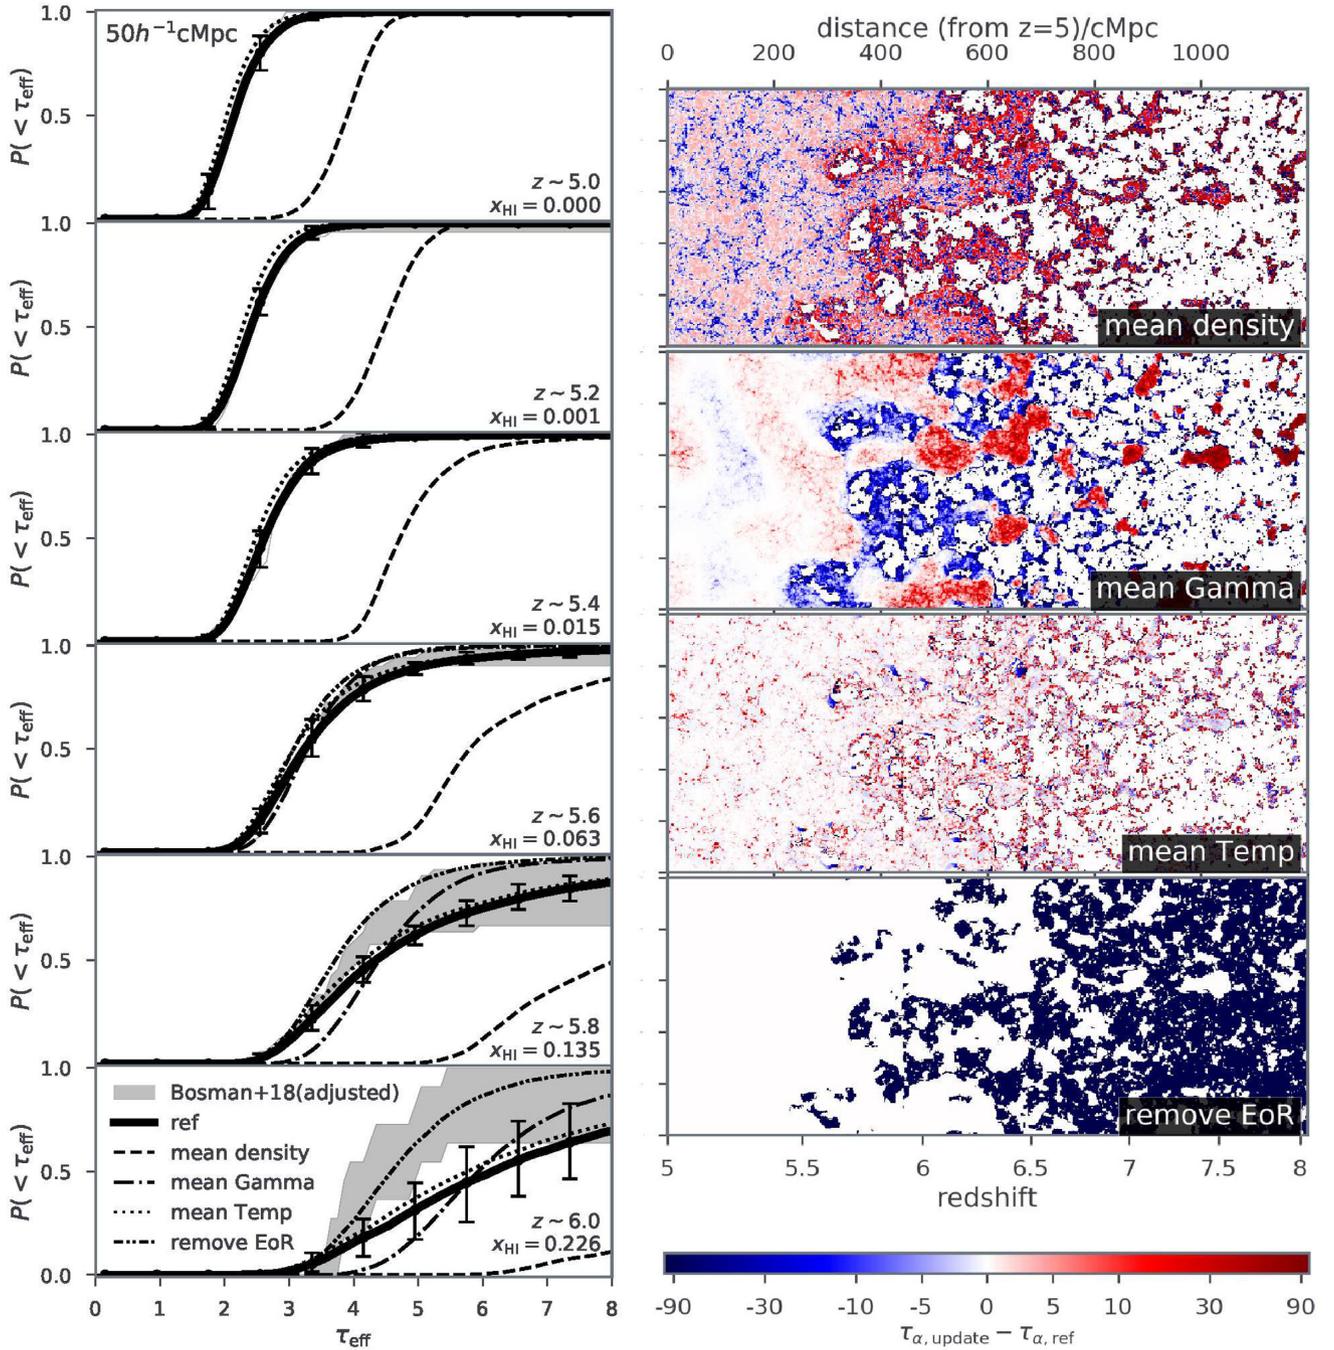


Figure 4. Left-hand panels: Mean CDFs of τ_{eff} averaged over $50 h^{-1} \text{cMpc}$ at $z = 5\text{--}6$ from our (ref) model are shown with thick black lines, with CV uncertainties ([14, 86] percentiles) indicated for a subset of the bins. We also show CDFs resulting from removing the spatial fluctuations in various component fields (see text for details). The grey shaded regions span the observational estimates (see Section 3). Right-hand panels: Light-cone slices illustrating the change in the Ly α optical depth relative to the reference model for each simulation that removes fluctuations in the indicated field. The colourbar is linear between 0 and ± 10 and logarithmic from ± 10 to ± 100 . These slices illustrate that large-scale fluctuations in this model are predominantly driven by a patchy EoR and a patchy UVB.

density fluctuations alone cannot efficiently generate long GP troughs (Becker et al. 2015). Ignoring the temperature fluctuations has a minor impact, mostly on small scales.

The most striking difference comes from the patchy EoR and UVB fields. We see that the fluctuations in both fields have a relatively large impact on the τ_{eff} CDFs during the final stages of the EoR ($z \gtrsim 5.5$), affecting the opacity on large scales. Removing

UVB fluctuations in the ionized IGM (i.e. assuming the ionized IGM component is exposed to a uniform UVB while keeping the neutral IGM component unchanged) makes the CDFs more narrow. Similarly, removing the EoR (i.e. artificially ionizing the neutral IGM component while keeping the ionized IGM component and its fluctuating UVB unchanged) also narrows the CDFs while at the same time shifts them towards low τ_{eff} values. In agreement with the

discussion in the previous section (see e.g. Keating et al. 2020a; Nasir & D’Aloisio 2020), their impact is additive – they both widen the CDFs, extending the high-value tails of τ_{eff} . Post-EoR as the mean free path increases, the fluctuations in the UVB become negligible when computing the τ_{eff} CDFs.¹²

3 OBSERVATIONAL DATA

We constrain our model using the most up-to-date, public Ly α forest sample¹³ from Bosman et al. (2018). Specifically, we use the ‘SILVER’ subsample of 51 quasars that have high-quality spectra with signal-to-noise ratio (S/N) ≥ 5.3 in the continuum. The spectra are binned over $50 h^{-1} \text{ cMpc}$, resulting in 217 flux measurements over the redshift interval $4.9 \leq z \leq 6.1$ (i.e. $\langle F \rangle_{50h^{-1} \text{ cMpc}}$ with uncertainties of $\delta_{\langle F \rangle_{50h^{-1} \text{ cMpc}}}$). We then have the following number of samples, $N_{\text{sample}} = 18, 47, 57, 51, 33,$ and 11 at $z = 5.0, 5.2, 5.4, 5.6, 5.8,$ and $6.0,$ respectively. The effective optical depth is then calculated as $\tau_{\text{eff}} = -\ln(\langle F \rangle_{50h^{-1} \text{ cMpc}})$ and, for non-detections, $\tau_{\text{eff}} \geq -\ln(2\delta_{\langle F \rangle_{50h^{-1} \text{ cMpc}}})$ as the lower limit.¹⁴

We then rescale these effective optical depths to account for the improved continuum reconstruction from Bosman et al. (2021). Specifically, we use the PCA-nominal method from Bosman et al. (2021), scaling the transmission by the ratio of the two continuum estimates, $\tau_{\text{eff}} = -\ln(\langle F \rangle_{50h^{-1} \text{ cMpc}} \mathcal{R})$, and the lower limits to $\tau_{\text{eff}} \geq -\ln(2\delta_{\langle F \rangle_{50h^{-1} \text{ cMpc}}} \mathcal{R})$ for non-detections. Here, $\mathcal{R} \equiv \langle F \rangle_{\text{obs}}^{\text{Bosman}+20} / \langle F \rangle_{\text{obs}}^{\text{Bosman}+18}$ represents the mean flux ratio between Bosman et al. (2018) and Bosman et al. (2021), which is 0.93, 0.87, 0.91, 0.87, 1.03, and 1.60 at $z = 5.0, 5.2, 5.4, 5.6, 5.8,$ and $6.0,$ respectively. We note that this only has a noticeable change to the CDF of τ_{eff} at $z = 6$ where the peak is at a slightly lower value.

The resulting observed CDFs of τ_{eff} are shown as the grey shaded regions in Fig. 4. When performing EoR inference from these observational estimates, we consider flat likelihoods between the two extremes (i.e. the grey shaded region; see more in Section 5.1).

4 SUMMARY OF MODEL PARAMETERS

Before proceeding to our MCMC results, we summarize the free parameters of our model and the associated prior ranges. For computational convenience in this proof-of-concept work, we restrict our parameter exploration to the most relevant astrophysical parameters (i.e. responsible for the largest variation in the forward-modelled data); for a complete list of parameters see Section 2 and Park et al. (2019). We also do not covary cosmological parameters, keeping the same underlying density field. In future work, we will relax this assumption and quantify joint constraints on astrophysical and cosmological parameters from the Ly α forest.

¹²We caution that our model does not include AGNs. AGNs are expected to dominate the UVB at lower redshifts, $z \lesssim 3-4$ (e.g. Madau & Haardt 2015; Mitra, Roy Choudhury & Ferrara 2015; Qin et al. 2017; Kulkarni et al. 2019); however, some works have evoked rare, bright AGNs to explain the observed large-scale τ_{eff} fluctuations at $z \sim 5-6$ (e.g. Chardin et al. 2017; Meiksin 2020). Regardless, our models are able to fit the data at $z > 5$ without AGNs (see e.g. Keating et al. 2020a; Nasir & D’Aloisio 2020).

¹³www.sarahbosman.co.uk/research

¹⁴Note that a few objects in Bosman et al. (2018) have different lower limits from $\tau_{\text{eff}} \geq -\ln(2\delta_{\langle F \rangle_{50h^{-1} \text{ cMpc}}})$ as they adopt the peak limit definition proposed by Becker et al. (2015).

Our model consists of six astrophysical parameters characterizing the UV emission of galaxies:

(i) $f_{*,10}$, the fraction of galactic baryons inside stars, defined for galaxies residing in haloes with $M_{\text{vir}} = 10^{10} M_{\odot}$. We vary $f_{*,10}$ between 10^{-3} and 1 with a flat prior in log space.

(ii) α_* , the power-law index of the stellar fraction to halo mass relation. We vary it between -0.5 and 1 with a flat prior, noting that high-redshift observations of galaxy UV LFs suggest $\alpha_* \sim 0.5$.

(iii) $f_{\text{esc},10}$, the UV ionizing escape fraction, defined for galaxies residing in haloes with $M_{\text{vir}} = 10^{10} M_{\odot}$. Similarly to $f_{*,10}$, we vary $f_{\text{esc},10}$ between 10^{-3} and 1 with a flat prior in log space.

(iv) α_{esc} , the power-law index of the ionizing escape fraction to halo mass relation. Although poorly known, some works suggest increasing or constant escape fractions towards lower mass galaxies (Ferrara & Loeb 2013; Kimm & Cen 2014; Paardekooper, Khochfar & Dalla Vecchia 2015; Xu et al. 2016). We thus have a wider prior on the negative range, varying α_{esc} between -1 and 0.5 (with a flat prior).

(v) M_{turn} , the turnover halo mass below which the number density of haloes hosting star-forming galaxies becomes exponentially suppressed. We vary M_{turn} between 10^8 and $10^{10} M_{\odot}$ with a flat prior in log space. The former corresponds to the atomic cooling threshold while the latter corresponds to current *HST* sensitivity limits beyond which we see a high occupancy fraction of star-forming galaxies.

(vi) t_* , the star formation time-scale as a fraction of the Hubble time. We vary t_* between 0 and 1 with a flat prior.

We also introduce two hyperparameters to describe the optical depth normalization factor, f_{rescale} from equation (9). Hyper (or nuisance) parameters can characterize systematics and errors, and are marginalized over to obtain the final constraints on the desired parameters. As mentioned in the associated discussion, the f_{rescale} normalization factor can account for errors in the continuum subtraction as well as modelling errors. However, it comes with the cost of losing the intrinsic predictive power of the model. In order to preserve some of the predictive power of our model, we assign these hyperparameters a prior centred around $f_{\text{rescale}} = 1$ (with unity corresponding to no calibration). The chosen widths are fairly arbitrary, with the fiducial Park et al. (2019) model included within 1σ ; however, below we also explore an extreme case that essentially allows infinitely wide priors. Specifically, the two hyperparameters of our fiducial MCMC are

(vii) $f_{\text{rescale}}(z = 5.7)$, the rescaling factor at $z = 5.7$, sampled in log space between $10^{-0.5}$ and $10^{0.5}$ with a Gaussian prior of a zero mean (i.e. $\langle \log_{10} f_{\text{rescale}}(z = 5.7) \rangle = 0$) and a width of $\sigma \log_{10} f_{\text{rescale}}(z = 5.7) = 1.5$.

(viii) df_{rescale}/dz , the slope of the rescaling factor as a function of redshift, sampled between -1.5 and 1.5 with a Gaussian prior of a zero mean (i.e. $\langle df_{\text{rescale}}/dz \rangle = 0$) and a width of $\sigma df_{\text{rescale}}/dz = 2.0$.

5 INFERRING GALAXY AND IGM PROPERTIES FROM THE LY α FOREST

In this section, we use the 21CMC sampler (Greig & Mesinger 2015, 2017, 2018) to perform three MCMC runs, quantifying EoR and galaxy constraints with and without the Ly α forest data. For computational efficiency, our MCMCs are done on smaller volumes (250 cMpc on a side) with the same cell size as the example shown in Section 2 (i.e. cell lengths of $250/128 = 1.95 \text{ cMpc}$).

5.1 Computing the likelihood

For a given combination of galaxy and hyperparameters, θ , we compute the associated likelihood from Ly α forest fluctuations, $\mathcal{L}_\alpha(\theta)$, according to the following:

- (i) We forward-model the corresponding 3D light-cone of the Ly α optical depth as described in Section 2.3.
- (ii) At a given redshift z , we compute the PDF¹⁵ of the effective optical depths from a sample of N_{sample} maximally separated (to minimize spatial coherence) sightlines, with N_{sample} corresponding to the sample size of the observational data at z as discussed in Section 3.
- (iii) We repeat step (ii) $N_{\text{realization}} = 150$ times, ensuring each sightline in the light-cone is not selected more than once. From these realizations, we compute the mean PDF (ϕ_{model}) and the cosmic variance (CV) error matrix (Σ_{CV} ; see more in Appendix B).
- (iv) We calculate the total error covariance matrix (ECM)¹⁶ as $\Sigma = \Sigma_{\text{CV}} + \Sigma_{\text{GP}}$, where Σ_{GP} corresponds to the error from the FGPA (see Appendix A).
- (v) We calculate the difference¹⁷ between the modelled mean PDF and the observed PDF, X .
- (vi) From (iv) and (v), we compute an χ^2 likelihood for this redshift, according to $\ln \mathcal{L}_z(\theta) = -0.5X^T \Sigma^{-1} X$.
- (vii) We repeat steps (ii) to (vi) for all redshift bins used in our analysis¹⁸ ($z = 5.4, 5.6, 5.8, \text{ and } 6.0$), summing up the log likelihoods: $\ln \mathcal{L}_\alpha(\theta) = \sum_z \ln \mathcal{L}_z(\theta)$.

Finally, we obtain the total likelihood with: $\mathcal{L}(\theta) = \mathcal{L}_\alpha \times \mathcal{L}_{\text{LFs}} \times \mathcal{L}_{\text{DF}} \times \mathcal{L}_{\tau_e}$, where the final three terms correspond to current, robust EoR constraints from: (i) the galaxy UV LFs at $z = 6-10$ from Bouwens et al. (2015a, 2016) and Oesch et al. (2018); (ii) the upper limit on the neutral hydrogen fraction at $z \sim 5.9$, $x_{\text{HI}} < 0.06 + 0.05(1\sigma)$, measured using the dark fraction of quasar spectra (McGreer et al. 2015); and (iii) the Thomson scattering optical depth of CMB photons reported by Planck Collaboration VI (2020), $\tau_e = 0.0561 \pm 0.0071(1\sigma)$, respectively. For more details about these observations and the functional forms of the corresponding likelihoods, see Park et al. (2019).

¹⁵As noted by Choudhury et al. (2021), we use PDFs instead of the more commonly presented CDFs when computing the likelihood, as CDFs have stronger covariances between bins.

¹⁶We do not include flux uncertainties in the total covariance matrix, as we verify they are far subdominant compared to the FGPA and CV errors for this observational sample. In future work, we will extend our framework and forward-model each individual sightline, including the corresponding flux errors.

¹⁷As mentioned in Section 3, we consider flat likelihoods between the two extreme cases shown in Fig. 4. For instance, assuming the observed number density in a particular τ_{eff} bin has the limits $\phi_{\text{obs,upper}}$ and $\phi_{\text{obs,lower}}$, we take $X = 0$ when $\phi_{\text{obs,lower}} \leq \phi_{\text{model}} \leq \phi_{\text{obs,upper}}$, and $X = \min(|\phi_{\text{model}} - \phi_{\text{obs,upper}}|, |\phi_{\text{model}} - \phi_{\text{obs,lower}}|)$ otherwise.

¹⁸We expect our models to become less accurate at lower redshifts since, (i) we do not account for the contribution of AGNs to the UVB; and (ii) we do not capture the spatial clustering of biased absorbers such as DLAs. As discussed above, we expect these approximations to be reasonable at $z > 5$; however, for this proof-of-concept work, we conservatively restrict ourselves to the highest redshift bins that are the most sensitive to the EoR. In future works, we will explore extending this range, including an additional population of bright AGNs.

5.2 Posteriors

We perform the following three MCMC runs:

- (i) *no_forest* – does not use the τ_{eff} PDFs, with the likelihood corresponding to $\mathcal{L}(\theta) = \mathcal{L}_{\text{LFs}} \times \mathcal{L}_{\text{DF}} \times \mathcal{L}_{\tau_e}$. This run roughly corresponds to our current state of knowledge, before accounting for Ly α opacity fluctuations. Here, we vary the astrophysical parameters (i)–(vi) from Section 4.
- (ii) *forest* – additionally includes the observed τ_{eff} PDFs discussed in Section 3; thus $\mathcal{L}(\theta) = \mathcal{L}_\alpha \times \mathcal{L}_{\text{LFs}} \times \mathcal{L}_{\text{DF}} \times \mathcal{L}_{\tau_e}$. Here, we vary the astrophysical and nuisance parameters (i)–(viii) from Section 4. *This corresponds to our fiducial MCMC run.*
- (iii) *forest_fluc* – unlike *forest*, does not sample the f_{rescale} hyperparameters from their priors. Instead, we normalize each PDF by ensuring the mean flux, $\langle \exp(-\tau_\alpha) \rangle$, matches the observed mean flux. This roughly corresponds to assuming infinitely wide priors for f_{rescale} . Therefore, this run only varies the astrophysical parameters (i)–(vi) from Section 4.

We monitor the progress of each MCMC, keeping only the samples after the posterior distributions stop evolving, and throwing out the preceding ‘burn-in’ samples. In the end, we discard 15k, 80k, and 80k simulations in the burn-in phase of *no_forest*, *forest* and *forest_fluc*, respectively, keeping 0.4M, 1.5M, and 1M simulations for the posteriors. These required approximately 0.3M CPUh to compute, corresponding to a month on a dedicated shared-memory server.

Fig. 5 presents the MCMC results including the marginalized posterior distributions of the model parameters (the corner plot in the lower left), the EoR histories (panels a), τ_{eff} distributions at $z = 5.4-6.0$ (i.e. CDF in panels b), galaxy UV LFs at $z = 6-15$ (panels c), the UVB evolution (panel d), and the PDF of the CMB Thomson scattering optical depth, τ_e (panel e). Below we discuss the posterior of each of these in turn.

5.2.1 no_forest

The *no_forest* posterior is shown with the orange, shaded regions. As noted by Park et al. (2019), current EoR observations already place constraints on some of our model parameters, even without making use of the τ_{eff} PDF data. In particular, galaxy LFs constrain the stellar to halo mass relation, parametrized here through $f_{*,10}/t_*$ and α_* , to within a factor of few (see e.g. Tacchella et al. 2018; Behroozi et al. 2019; Yung et al. 2019; Qin et al. 2020). This is also evident from the tight recovery of the UV LFs at the bright end ($M_{\text{UV}} \lesssim -15$). However, we do not detect a faint-end turnover in the LFs, resulting from inefficient star formation in galaxies hosted by haloes with masses less than M_{turn} ; current UV LFs only provide upper limits on this parameter ($\lesssim 5 \times 10^9 M_\odot$; e.g. Bouwens et al. 2015b). The ionizing escape fraction is only weakly constrained by the dark fraction and CMB limits on EoR timing: $f_{\text{esc},10} \sim 4-26$ per cent), while its scaling with the mass of the host halo remains uninformed (see e.g. Haardt & Madau 2012; Kuhlen & Faucher-Giguère 2012; Robertson et al. 2015; Price, Trac & Cen 2016).

5.2.2 forest

Additionally including the τ_{eff} data has a dramatic impact on the posterior (shown with the red curves in Fig. 5). Most importantly, we note from panel (a1) that *the forest data requires late reionization*. The final overlap stages of the EoR, corresponding to when >95 percent of the volume was ionized, occur at $z <$

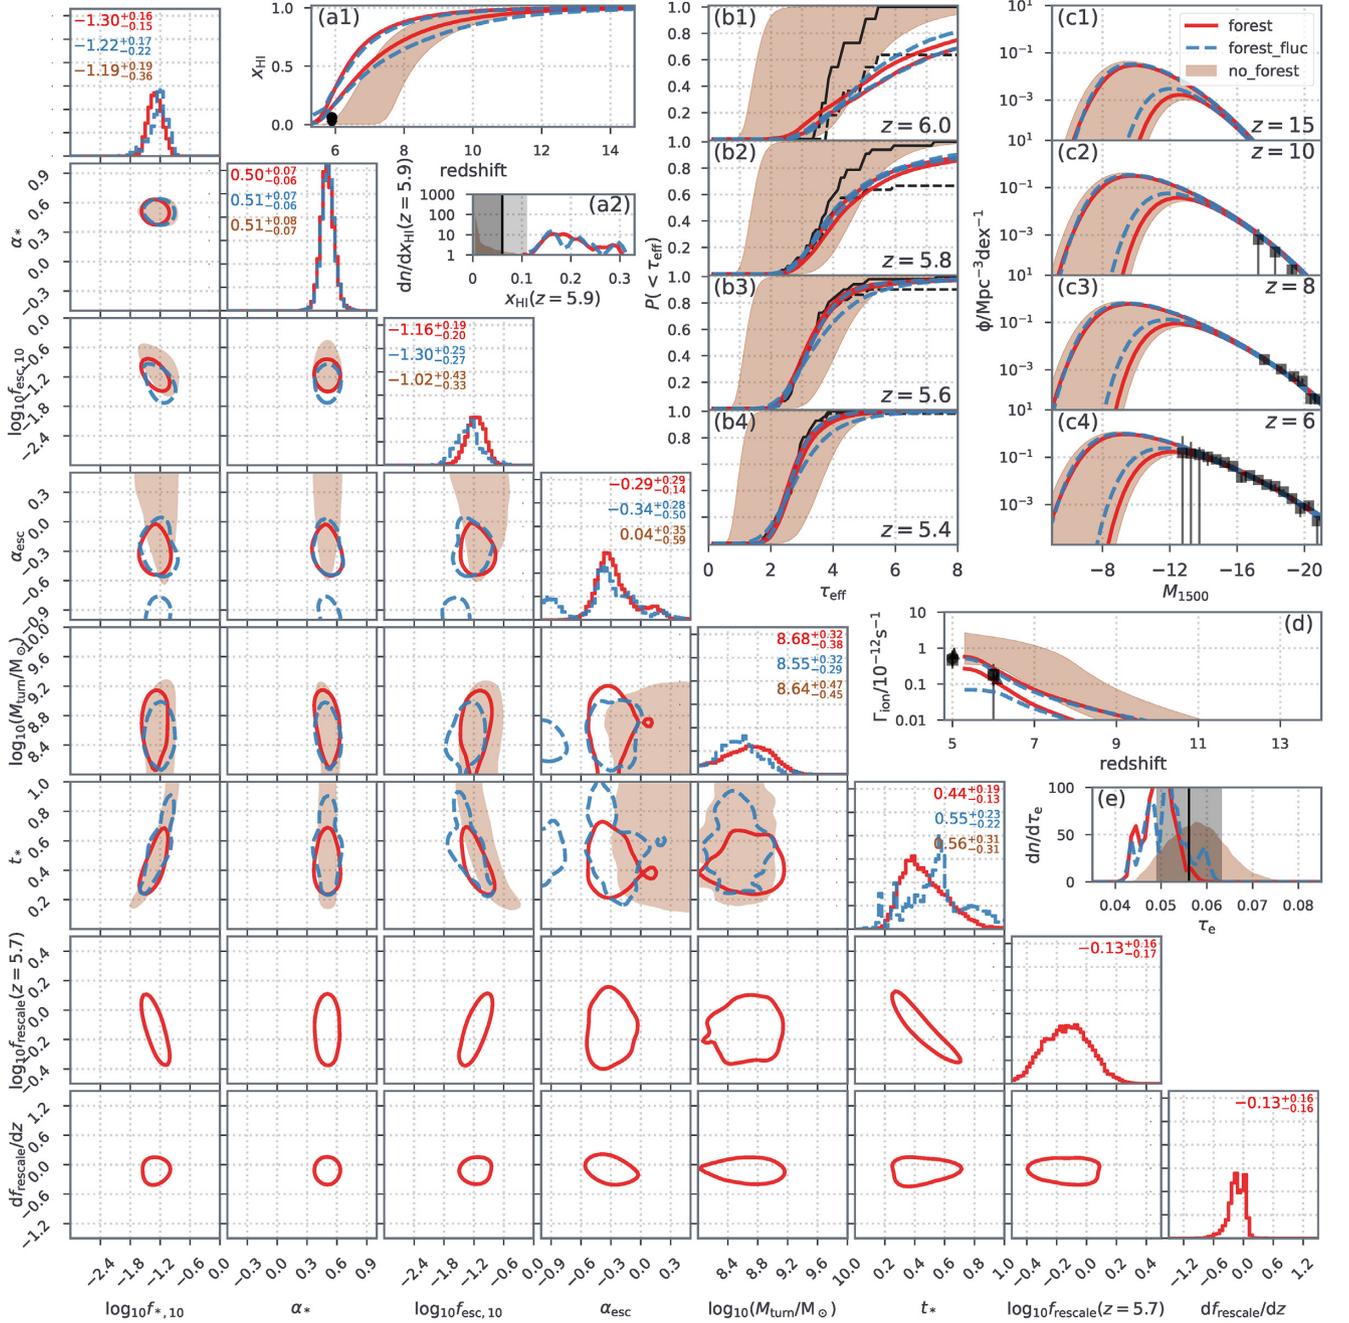


Figure 5. Marginalized posterior distributions of the model parameters with (forest; red solid lines and forest_fluc; blue dashed lines) and without (no_forest; brown shaded areas) the forest data (Bosman et al. 2018). All three results use the observed galaxy LFs at $z = 6-10$ (Bouwens et al. 2015a, 2017; Oesch et al. 2018), upper limits on the neutral fraction at $z \sim 5.9$ from quasar spectra (McGreer et al. 2015), and Thomson scattering optical depth of the CMB (Planck Collaboration VI 2020) as input constraints. The 2D distributions correspond to the 68th percentiles. The medians with [14, 86] percentiles for each parameter are presented in the upper corner of the 1D PDFs (from top to bottom: *forest*, *forest_fluc* and *no_forest*, respectively). Note that all 1D PDFs have been normalized to have area (or integral) under the histogram equal to 1. The upper right subpanels present the median and [14, 86] percentiles of the neutral hydrogen (x_{HI} ; panel a1); the CDFs of Ly α effective optical depths at $z = 5.4-6.0$ (panels b1–b4); the UV LFs at $z = 6-15$ (panels c1–c4); the evolution of the photoionization rate (Γ_{ion} ; panel d); and the PDFs of τ_e (panel e) and x_{HI} at $z = 5.9$ (panel a2) for the models presented in the posterior distributions. Observations including UV ionizing background measured by Bolton & Haehnelt (2007b), Calverley et al. (2011), and Wyithe & Bolton (2011) are indicated in black.

5.6. Our Bayesian framework provides statistical proof of previous suggestions that the EoR might have completed at $z < 6$ (Lidz et al. 2007; Mesinger 2010; Kulkarni et al. 2019; Keating et al. 2020a; Nasir & D’Aloisio 2020). Our constraints on overlap at $z < 5.6$ are also perfectly consistent with the recent, similar analysis

by Choudhury et al. (2021), using their indirect parametrization of ionizing sources and IGM recombinations.

Moreover, we find that the τ_{eff} PDFs and the dark fraction upper limits are in mild tension ($\sim 1.5\sigma$; see panel a2). The dark fraction tends to prefer earlier reionization, while the τ_{eff} PDFs require later

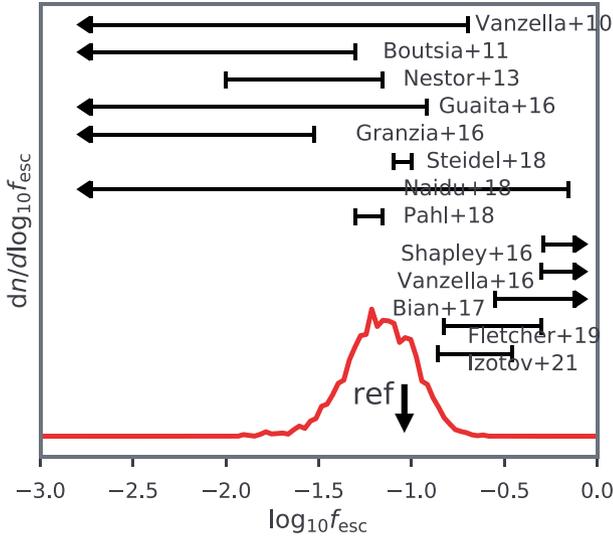


Figure 6. The posterior probability distribution function of $f_{\text{esc},10}$ from the *forest* result compared to observational estimates (Vanzella et al. 2010, 2016; Boutsia et al. 2011; Nestor et al. 2013; Grazian et al. 2016; Guaita et al. 2016; Shapley et al. 2016; Bian et al. 2017; Naidu et al. 2018; Steidel et al. 2018; Fletcher et al. 2019; Izotov et al. 2021; Pahl et al. 2021). The MAP (ref) model is indicated with an arrow.

reionization. Combining the two forces the marginalized posterior of EoR histories to have a very narrow tail below $z < 6$ (see panel a1). We will revisit this mild tension in future work, using updated estimates of both data sets from the XQR30 large observational program (PI: V. D’Odorico).

As it helps nail down the timing of the EoR, the forest data dramatically improve constraints on the ionizing escape fraction. To be consistent with all of the data, reionization would need to end late (preferred by the broad τ_{eff} PDFs) but not *too* late or with an extended tail towards low redshifts (preferred by the modest values of τ_e and the dark fraction). These limit the ionizing contribution of both early-forming, faint galaxies and their late-appearing, bright counterparts, resulting in α_{esc} that peaks around $-0.29^{+0.29}_{-0.14}$, a characteristic ionizing escape fraction of $f_{\text{esc},10} = 6.9^{+3.8}_{-2.6}$ per cent, as well as a weakly constrained ($\sim 1\sigma$) lower limit on M_{turn} (i.e. $\gtrsim 2 \times 10^8 M_{\odot}$).

As we mentioned previously, in this initial study we do not enforce UV photon conservation, in favour of simulation speed (setting PHOTON_CONS=FALSE in 21 cmFASTv3). This choice roughly leads to a +0.2 bias in the inferred α_{esc} parameter (Park et al., in preparation). Taking this into account, we predict that the true marginalized posterior of α_{esc} is consistent with zero and has a modest width of $\sigma_{\alpha_{\text{esc}}} \lesssim 0.3$. This suggests that *the forest data disfavour a strong evolution of the ionizing escape fraction with the galaxy’s halo (or stellar) mass* (i.e. large $|\alpha_{\text{esc}}|$), consistent with recent results from hydrodynamical simulations of a handful high- z galaxies in the relevant mass ranges (Kimm & Cen 2014; Ma et al. 2015, 2020; Xu et al. 2016).

In Fig. 6, we plot the *forest* posterior of $f_{\text{esc},10}$ together with some observational estimates.¹⁹ The fact that our data-constrained posterior is narrower than the spread in observational estimates

¹⁹Since our results imply only a weak mass evolution of the escape fraction, here we loosely treat $f_{\text{esc},10}$ as a characteristic f_{esc} when qualitatively comparing to observations.

demonstrates the power of our forward-modelling approach; however, we caution that as we do not match the redshifts, selection effects, and other systematics of these observations, this comparison is only intended to be illustrative. Stacking large samples of $z \sim 3$ –4 galaxies (e.g. Vanzella et al. 2010; Boutsia et al. 2011; Grazian et al. 2016; Guaita et al. 2016; Naidu et al. 2018; Steidel et al. 2018; Pahl et al. 2021) generally results in lower f_{esc} estimates, while larger values have been seen in some individual galaxies (e.g. Vanzella et al. 2016; Bian et al. 2017; Fletcher et al. 2019; Izotov et al. 2021). Since our model parametrizes population-averaged galaxy properties, it is interesting that it is consistent with the stacking results, and implies that the individual galaxies with high $f_{\text{esc}} \gtrsim 30$ per cent are rare objects prone to selection bias.

Finally, it is useful to point out a few sanity checks of our model. First, we see that the recovered hyperparameters, accounting for systematic errors in continuum subtraction and theoretical modelling, are consistent with $f_{\text{rescale}} = 1$. This implies that our intrinsic models of the forest can match the forest data, *without significant ‘tuning’*. Secondly, we note from panel (d) that our *forest* posterior matches perfectly with estimates of the mean ionizing background at $z = 5$ and 6 (denoted by black points with error bars; Bolton & Haehnelt 2007b; Calverley et al. 2011; Wyithe & Bolton 2011). It is worth noting that, although they are based on forest observations, *we do not use these estimates of the UVB in our likelihood*.

5.2.3 forest_flux

The *forest_flux* posterior is denoted with blue curves in Fig. 5. Recall that here we calibrate each forward-modelled τ_{eff} PDF such that the mean flux matches the data, for every redshift bin and every parameter sample. This corresponds to an extreme case in which we have no idea how to put priors on f_{rescale} . In effect, we are removing the ‘DC mode’ of our forest models, mostly comparing the shapes of the PDFs over the observable range.

Even when calibrating to the observed mean flux, we recover the result that EoR must finish at $z < 5.6$ to match the observed *widths* of the Ly α opacity distributions (panel a1). By comparing the blue curves to the red ones, we see that losing information on the mean flux expands the recovered posterior to include models with a very low ionizing background (see panel d). These are mostly sourced by small values of α_{esc} , corresponding to enhanced ionizing efficiencies in galaxies hosted by low-mass haloes. The corresponding EoR histories are also somewhat slower, since the fractional growth of less massive haloes is slower than that of more massive ones. As a result, the marginal posterior on α_{esc} is bimodal, with the smaller peak at low values being driven by a slightly better agreement with τ_e from *Planck* due to the more extended EoR.

5.3 Can we distinguish between different reionization morphologies?

In the previous section, we demonstrated that our models require reionization to be incomplete at $z < 6$ in order to match the observed τ_{eff} PDFs, while in Section 2.3.3 we saw that both the patchy UVB and patchy EoR are needed to reproduce the longest GP troughs (see also Keating et al. 2020a). In this section, we check if the forest opacity fluctuations can distinguish between different reionization morphologies, *at a fixed stage of the EoR*. Constraining the bias of the dominant galaxy population through the EoR morphology would be immensely powerful, allowing us to distinguish between different galaxy models that result in similar reionization histories

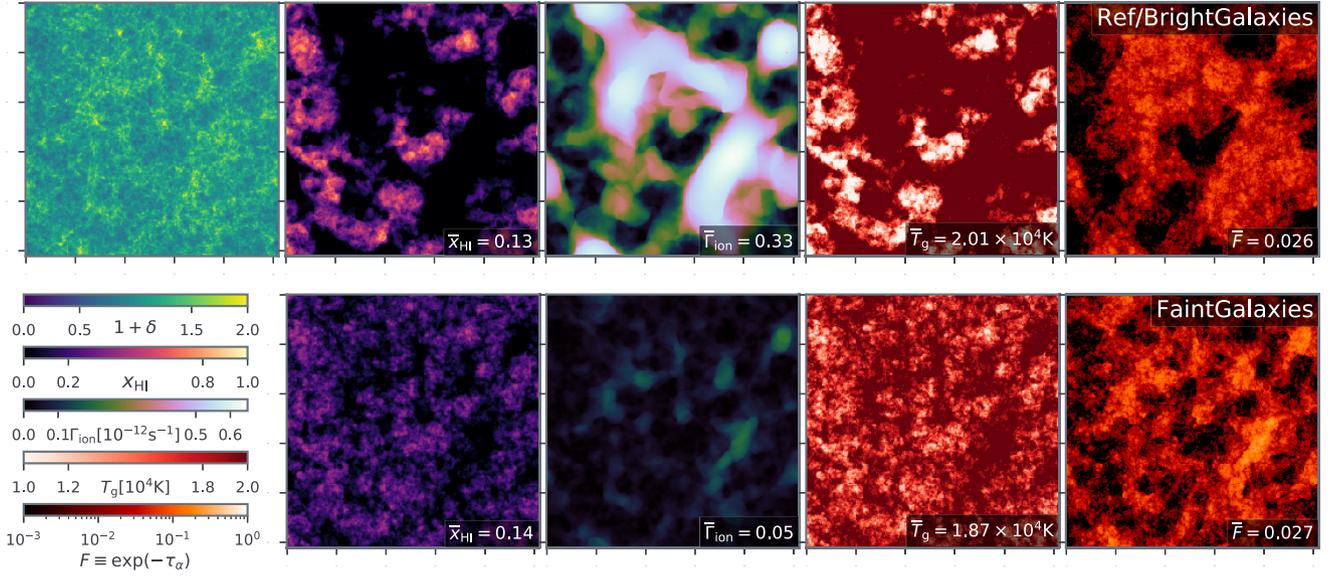


Figure 7. Slices (with a side length of $50 h^{-1}$ cMpc and depth of 2 cMpc) of the density ($1 + \delta$), neutral hydrogen fraction (x_{HI}), photonization rate (Γ_{ion}), gas temperature (T_{g}), and Ly α transmission flux from *Ref/BrightGalaxies* (upper subpanels) and *FaintGalaxies* (lower ones); see text for more details on these models. All slices correspond to $z = 5.8$ where the two models have the same filling factor of H I, and were rerun in larger volumes (500 cMpc) compared to the MCMC for better visualization.

(e.g. McQuinn et al. 2007; Dixon et al. 2016; Mesinger, Greig & Sobacchi 2016; Ahn & Shapiro 2021). Indeed, the ability to measure the morphology of the EoR and epoch of heating is one of the main reasons the cosmic 21-cm signal will allow us to place \sim percent level constraints on the properties of the unseen first galaxies (see e.g. the recent review in Mesinger 2019).

In Fig. 7, we show slices through the fields of two models at $z = 5.8$, chosen to have the same neutral fraction ($\bar{x}_{\text{HI}} \sim 0.13$) but very different EoR morphologies. The top panels correspond to our reference (MAP) model, discussed already in Section 2.3. This model has $\alpha_{\text{esc}} = -0.298$ and $M_{\text{turn}} = 7.2 \times 10^8 M_{\odot}$, resulting in relatively massive galaxies driving the EoR; thus we label it *Ref/BrightGalaxies*. Below it we show a *FaintGalaxies* model with the following astrophysical parameters $(f_{*,10}, \alpha_*, f_{\text{esc},10}, \alpha_{\text{esc}}, M_{\text{turn}}, t_*) = (0.084, 0.483, 0.019, -0.834, 2.16 \times 10^8 M_{\odot}, 0.678)$. Note that we also include $f_{\text{rescale}} = 0.12$ in the *FaintGalaxies* model to rescale the Ly α transmission at $z = 5.8$ (see equation 9; compared to $f_{\text{rescale}} = 0.92$ in *BrightGalaxies*). Although both models have the same H II filling factor at $z = 5.8$, the EoR morphologies are noticeably different, as expected from the different biases of their corresponding dominant galaxy population. Specifically, *FaintGalaxies* is characterized with more numerous, smaller H II regions driven by its more abundant, yet fainter galaxies. As a result, the UVB is weaker, as the contribution of very distant sources is limited by the small sizes of the H II regions.

To quantify if the different EoR morphologies result in different τ_{eff} distributions, in Fig. 8 we compute the CDFs averaging over 10, 50 (the default value used in the MCMC), and $100 h^{-1}$ Mpc.²⁰

From the figure we see that the difference in the CDFs, when normalized to have the same Lyman alpha transmission, is modest

²⁰We note that the *FaintGalaxies* model has a different EoR evolution and is disfavoured by the total likelihood in *forest*. However, here we focus only on the CDFs at $z = 5.8$ where the two models have the same neutral fraction and comparable likelihoods (i.e. $\mathcal{L}_{z=5.8}$; see Section 5.1).

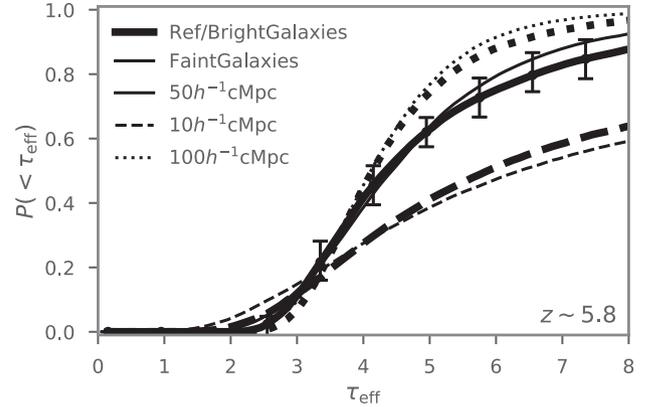


Figure 8. CDFs of τ_{eff} and $\ln \tau_{\text{eff}}$ at $z = 5.8$ averaged over 50, 10, and $100 h^{-1}$ cMpc from *Ref/BrightGalaxies* (thick curves) and *FaintGalaxies* (thin curves). Uncertainties ([14, 86] percentiles) for the $50 h^{-1}$ cMpc result in *Ref/BrightGalaxies* are indicated in a subset of the bins.

– smaller than the CV of the current sample (see the errorbar for the $50 h^{-1}$ cMpc result in *Ref/BrightGalaxies*). This is due to the fact that the dynamic range probed by the forest is too small to discriminate against different EoR morphologies at a fixed neutral fraction. Indeed, we confirm that the largest differences between the two models, when averaged over $10 h^{-1}$ cMpc, occurs at $\tau_{\text{eff}} \sim 90$. This is far beyond the observable range. In matching the same mean flux, the differences in τ_{eff} appear in the tail of the PDF, and are thus difficult to constrain. A larger quasar sample ($N_{\text{sample}} \sim 1000$) would narrow down the CV uncertainties within the accessible range; however, such a large sample might only become feasible with next-generation telescopes such as *Euclid*.

Therefore, we conclude that the current Ly α forest data are unlikely to be able to distinguish between different galaxy models

having the same global properties (such as the EoR history; see also Nasir & D’Aloisio 2020 for a similar conclusion). In future work, we will also forward-model the Lyman β forest, quantifying if its added dynamic range can further help constrain models (as has been implied by, e.g. Eilers et al. 2019; Keating et al. 2020b), despite the added CV from the overlapping, lower redshift Ly α forest. We will also investigate the constraining power of other summary statistics.

5.4 Comparison to previous works

Previous studies have made significant progress in understanding the large-scale Ly α opacity fluctuations at $z \gtrsim 5.5$. However, since forest simulations are generally computationally expensive, few studies do exhaustive parameter exploration. Instead, most focus on one or two of the following dominant sources of fluctuations – the UVB (either through rare sources or a short mean free path), the gas temperature, and late reionization. In this section, we compare our MAP model from the *forest* MCMC (labelled *ref*) to some of the recent works that managed to explain the large-scale opacity fluctuations. In Fig. 9, we show the average EoR history, photoionization rate, ionizing emissivity/recombination rate, and mean free path, from *ref*²¹ along with some other models and observational data.

As demonstrated quantitatively in Section 2.3.3, the patchy UVB is an important source of forest fluctuations. However, we find that UVB fluctuations alone are not sufficient, with late reionization required to fully explain the data (see e.g. Keating et al. 2020b; Choudhury et al. 2021). Some previous works have focused on the patchy UVB, sourced by either rare, bright active galactic nuclei (AGNs) or a small mean free path. We briefly discuss each in turn.

Although rare, bright sources could explain the τ_{eff} CDF (Chardin et al. 2015, 2017; D’Aloisio et al. 2017; Meiksin 2020), they somewhat struggle to reproduce the longest GP trough found in ULAS J0148+0600, and would likely be in tension with IGM temperature measurements (e.g. D’Aloisio et al. 2017). Furthermore, extrapolations of AGN LFs and semi-analytical models suggest they provide a negligible contribution to the UVB at $z > 5.5$ (e.g. Mitra et al. 2015; Manti et al. 2017; Parsa, Dunlop & McLure 2017; Qin et al. 2017; Garaldi, Compostella & Porciani 2019).

On the other hand, a small mean free path could modulate a galaxy-dominated UVB to the level required to explain the forest observations (e.g. Davies & Furlanetto 2016; D’Aloisio et al. 2018). However, these models generally require very short mean free paths (cf. the *early-reionization-short-mfp* model from Nasir & D’Aloisio 2020 and the *short mean free path* from D’Aloisio et al. 2018 denoted with the orange points in the bottom panel of Fig. 9). As discussed in D’Aloisio et al. (2018), observational estimates of the mean free path (e.g. Songaila & Cowie 2010; Worseck et al. 2014) may be biased high, due to contamination from the proximity zone in which the quasar flux dominates over that of the UVB. Indeed, the mean free path from our MAP model is in good agreement with the fiducial values computed in that work, but significantly above those required to explain the opacity fluctuations without evoking a late reionization. Unlike these works, our framework does not have a tuning knob for the mean free path, instead calculating this

²¹In principle, our Bayesian framework allows us to compare the posterior distributions of these quantities. However, we do not output all of these fields when running the MCMC. Therefore, here we only show the MAP model that was rerun to output the desired quantities.

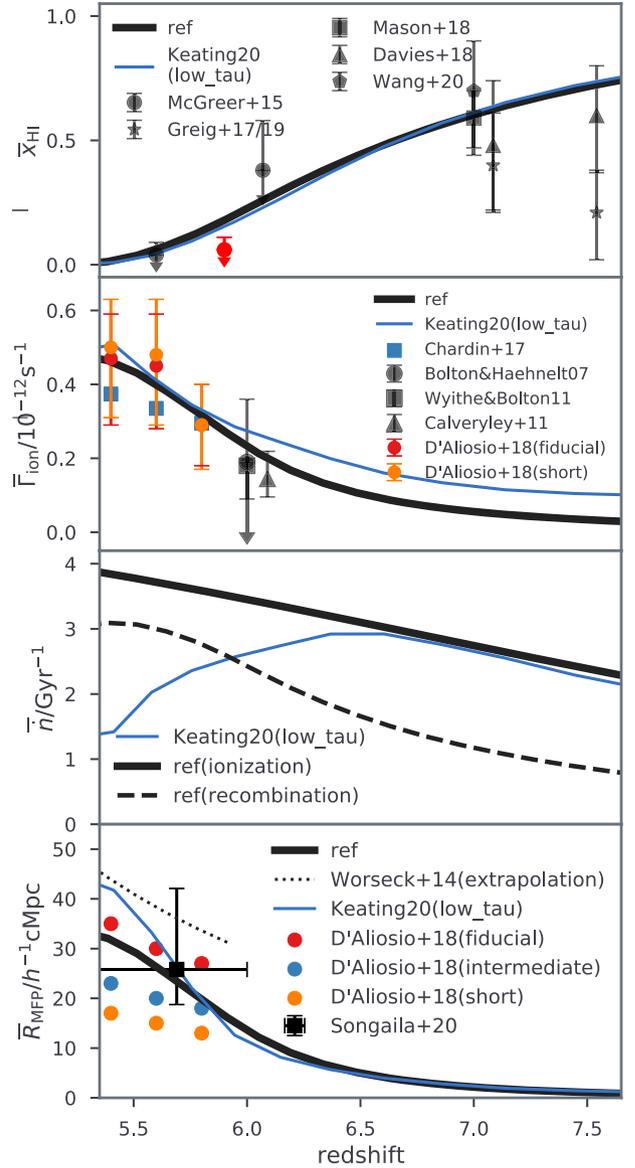


Figure 9. The average neutral hydrogen fraction, photoionization rate, ionizing emissivity/recombination rate, mean free path, and gas temperature of our MAP parameter combination (*ref*), along with some other models (Chardin et al. 2017; D’Aloisio et al. 2018; Keating et al. 2020b) and observational limits (Bolton & Haehnelt 2007b; Calverley et al. 2011; Wyithe & Bolton 2011; Worseck et al. 2014; McGreer et al. 2015; Greig et al. 2017, 2019; Mason et al. 2018; Davies et al. 2018b; Wang et al. 2020).

quantity directly from the source and sink distributions following Sobacchi & Mesinger (2014). Our galaxy models and associated parameter priors do not result in mean free paths of < 10 cMpc at $z < 6$.

In addition to a patchy UVB, large-scale temperature fluctuations have been evoked to explain the forest observations (D’Aloisio et al. 2015). However, as discussed in Section 2.3.3, we find temperature fluctuations have only a minor impact (see also Keating et al. 2018). This could be due to the fact that strong temperature fluctuations require fairly extended reionization histories, which are disfavoured in our posterior. Indeed, the fiducial model in D’Aloisio et al. (2015) is in mild tension with the latest *Planck* measurement (Planck

Collaboration VI 2020). Furthermore, these models require that troughs in the transmission come from large-scale overdensities, which host an overabundance of galaxies and thus reionize early. This is contrary to recent observations showing a dearth of galaxies around the long GP trough of ULAS J0148+0600 (Becker et al. 2018; Kashino et al. 2020).

Finally, late reionization (e.g. Mesinger 2010) has also been used to explain the τ_{eff} distributions (Kulkarni et al. 2019; Keating et al. 2020a, b; Nasir & D’Aloisio 2020). Our results are consistent with this claim (see also the recent work by Choudhury et al. 2021). In Fig. 9, we show the fiducial model from Keating et al. (2020b), together with our MAP model. Amazingly, the EoR histories are in perfect agreement, despite the fact that Keating et al. (2020b) did not perform Bayesian inference. However, their model requires a drop in the ionizing emissivity (\bar{n}_{ion}) at $z < 6.5$, which is difficult to justify physically and is in contradiction with the observed redshift evolution of the star formation rate density (e.g. Bouwens et al. 2015a). On the other hand, our model achieves the same EoR history without a non-monotonic evolution of the emissivity; the recombination rate increases, approaching the ionization rate, resulting in relatively slow, ‘photon starved’ end to reionization (e.g. Bolton & Haehnelt 2007b; Sobacchi & Mesinger 2014). One explanation for this difference might be that the simulations of Keating et al. (2020b), with a gas particle mass of $\sim 10^7 M_{\odot}$, could be underresolving small gas clumps and thus underestimating the impact of recombinations. Although the mean free path (being an instantaneous, volume-averaged quantity) does not directly show the cumulative impact of inhomogeneous recombinations, the rapid rise in their mean free path below $z < 6$ seen in the bottom panel supports this explanation.

6 CONCLUSIONS

In this work, we extend the Bayesian inference framework of 21cm FAST/21CMMC to forward-model the (low-resolution) Ly α forest. We run MCMCs by sampling empirical galaxy scaling relations and computing the corresponding 3D light-cones of the Ly α forest. With these, we quantify the additional constraining power provided by observations of large-scale opacity fluctuations (i.e. PDFs of τ_{eff} averaged over 50 cMpc/h; Bosman et al. 2018, 2021).

We find that, in order to be consistent with the observations, our models *require* late reionization. The final overlap stages of the EoR, corresponding to when >95 percent of the volume was ionized, occur at $z < 5.6$. Our Bayesian framework provides statistical proof of previous suggestions that the EoR might have completed reionization at $z < 6$ (Lidz et al. 2007; Mesinger 2010; Kulkarni et al. 2019; Keating et al. 2020a; Nasir & D’Aloisio 2020), and is perfectly consistent with the recent, similar analysis by Choudhury et al. (2021). Such late reionization is in mild tension ($\sim 1.5\sigma$) with the dark fraction upper limits from McGreer et al. (2015). In the future, we will revisit this mild tension using updated data from the XQR-30 large VLT program.

We also find that the forest data improve our current knowledge of galaxy UV ionizing properties. In particular, we find a weak ($\sim 1\sigma$) constraint on the turnover halo mass scale ($M_{\text{turn}} = 2 \times 10^8 - 10^9 M_{\odot}$), below which star formation stops being efficient. Moreover, we find that the late reionization preferred by the forest data tightens constraints on the ionizing escape fraction. Combined observations (i.e. galaxy UV LFs, CMB optical depth, dark fraction and forest) favour a characteristic ionizing escape fraction of $f_{\text{esc}} = 6.9^{+3.8}_{-2.6}$ per cent, and disfavour a strong evolution with the galaxy’s halo (or stellar) mass. Unfortunately, the τ_{eff} CDFs cannot distinguish among different

source/sink models that have different EoR morphologies but the same EoR history.

Using our MAP model, we demonstrate that large-scale opacity fluctuations are driven by a combination of both patchy reionization and spatial variations in the photoionizing background (with temperature inhomogeneities being subdominant). The cosmic H I patches and regions of weak UVB both corresponds to large-scale underdensities in the matter field. Thus, the longest GP troughs correlate with a relative dearth of galaxies, in agreement with observations (Becker et al. 2018; Kashino et al. 2020) and some previous models (e.g. Davies, Becker & Furlanetto 2018a).

Our inference framework can easily be extended to include different source models, such as AGNs and/or having a more complex parametrization of galaxy evolution. Using the Bayesian evidence, we can quantify if the data require the additional model complexity (e.g. Qin et al. 2021). We postpone such investigation to future work, applying them on upcoming, larger data sets.

ACKNOWLEDGEMENTS

We thank George Becker, Anson D’Aloisio, Frederick Davies, Laura Keating, and Girish Kulkarni for helpful comments on a draft version of this paper. This work was supported by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programmes (AIDA – #638809, First Light – #669253, and Cosmic Gas – #740246). The results presented here reflect the authors’ views; the ERC is not responsible for their use. We acknowledge computational resources of the HPC centre at SNS. MV is supported by grants INFN INDARK PD51 and ASI-INAF n.2017-14-H.0.

DATA AVAILABILITY

The data underlying this article will be shared on reasonable request to the corresponding author.

REFERENCES

- Ahn K., Shapiro P. R., 2021, *ApJ*, 914, 44
 Atek H., Richard J., Kneib J.-P., Schaerer D., 2018, *MNRAS*, 479, 5184
 Bañados E. et al., 2018, *Nature*, 553, 473
 Barkana R., Loeb A., 2005, *ApJ*, 626, 1
 Bauer A., Springel V., Vogelsberger M., Genel S., Torrey P., Sijacki D., Nelson D., Hernquist L., 2015, *MNRAS*, 453, 3593
 Bautista J. E. et al., 2017, *A&A*, 603, A12
 Becker G. D., Bolton J. S., 2013, *MNRAS*, 436, 1023
 Becker G. D., Rauch M., Sargent W. L. W., 2007, *ApJ*, 662, 72
 Becker G. D., Bolton J. S., Madau P., Pettini M., Ryan-Weber E. V., Venemans B. P., 2015, *MNRAS*, 447, 3402
 Becker G. D., Davies F. B., Furlanetto S. R., Malkan M. A., Boera E., Douglass C., 2018, *ApJ*, 863, 92
 Behroozi P., Wechsler R. H., Hearin A. P., Conroy C., 2019, *MNRAS*, 488, 3143
 Bhatwadekar R., Conselice C. J., Margalef-Bentabol B., Duncan K., 2019, *MNRAS*, 486, 3805
 Bian F., Fan X., McGreer I., Cai Z., Jiang L., 2017, *ApJ*, 837, L12
 Bolton J. S., Haehnelt M. G., 2007a, *MNRAS*, 374, 493
 Bolton J. S., Haehnelt M. G., 2007b, *MNRAS*, 382, 325
 Bolton J. S., Becker G. D., Wyithe J. S. B., Haehnelt M. G., Sargent W. L. W., 2010, *MNRAS*, 406, 612
 Bolton J. S., Haehnelt M. G., Warren S. J., Hewett P. C., Mortlock D. J., Venemans B. P., McMahon R. G., Simpson C., 2011, *MNRAS*, 416, L70

- Bolton J. S., Puchwein E., Sijacki D., Haehnelt M. G., Kim T.-S., Meiksin A., Regan J. A., Viel M., 2017, *MNRAS*, 464, 897
- Bosman S., 2020, Zenodo Dataset, Vol. 36, Zenodo, p. 34964
- Bosman S. E. I., Fan X., Jiang L., Reed S., Matsuoka Y., Becker G., Haehnelt M., 2018, *MNRAS*, 479, 1055
- Bosman S. E. I., Durovčková D., Davies F. B., Eilers A.-C., 2021, *MNRAS*, 503, 2077
- Boutsia K. et al., 2011, *ApJ*, 736, 41
- Bouwens R. J. et al., 2015a, *ApJ*, 803, 34
- Bouwens R. J., Illingworth G. D., Oesch P. A., Caruana J., Holwerda B., Smit R., Wilkins S., 2015b, *ApJ*, 811, 140
- Bouwens R. J. et al., 2016, *ApJ*, 830, 67
- Bouwens R. J., Oesch P. A., Illingworth G. D., Ellis R. S., Stefanon M., 2017, *ApJ*, 843, 129
- Busca N. G. et al., 2013, *A&A*, 552, A96
- Calverley A. P., Becker G. D., Haehnelt M. G., Bolton J. S., 2011, *MNRAS*, 412, 2543
- Chardin J., Haehnelt M. G., Aubert D., Puchwein E., 2015, *MNRAS*, 453, 2943
- Chardin J., Puchwein E., Haehnelt M. G., 2017, *MNRAS*, 465, 3429
- Choudhury T. R., Haehnelt M. G., Regan J., 2009, *MNRAS*, 394, 960
- Choudhury T. R., Paranjape A., Bosman S. E. I., 2021, *MNRAS*, 501, 5782
- Croft R. A. C., Weinberg D. H., Bolte M., Burles S., Hernquist L., Katz N., Kirkman D., Tytler D., 2002, *ApJ*, 581, 20
- D'Aloisio A., McQuinn M., Trac H., 2015, *ApJ*, 813, L38
- D'Aloisio A., Upton Sanderbeck P. R., McQuinn M., Trac H., Shapiro P. R., 2017, *MNRAS*, 468, 4691
- D'Aloisio A., McQuinn M., Davies F. B., Furlanetto S. R., 2018, *MNRAS*, 473, 560
- D'Aloisio A., McQuinn M., Maupin O., Davies F. B., Trac H., Fuller S., Upton Sanderbeck P. R., 2019, *ApJ*, 874, 154
- Davies F. B., Furlanetto S. R., 2016, *MNRAS*, 460, 1328
- Davies F. B., Becker G. D., Furlanetto S. R., 2018a, *ApJ*, 860, 155
- Davies F. B. et al., 2018b, *ApJ*, 864, 142
- Davies J. E., Mutch S. J., Qin Y., Mesinger A., Poole G. B., Wyithe J. S. B., 2019, *MNRAS*, 489, 977
- Davies F. B., Hennawi J. F., Eilers A.-C., 2020, *MNRAS*, 493, 1330
- Delubac T. et al., 2015, *A&A*, 574, A59
- Dixon K. L., Iliiev I. T., Mellema G., Ahn K., Shapiro P. R., 2016, *MNRAS*, 456, 3011
- Efstathiou G., 1992, *MNRAS*, 256, 43P
- Eilers A.-C., Davies F. B., Hennawi J. F., Prochaska J. X., Lukić Z., Mazzucchelli C., 2017, *ApJ*, 840, 24
- Eilers A.-C., Davies F. B., Hennawi J. F., 2018, *ApJ*, 864, 53
- Eilers A.-C., Hennawi J. F., Davies F. B., Oñorbe J., 2019, *ApJ*, 881, 23
- Eilers A.-C. et al., 2020, *ApJ*, 900, 37
- Fan X. et al., 2006, *AJ*, 132, 117
- Faucher-Giguère C.-A., Prochaska J. X., Lidz A., Hernquist L., Zaldarriaga M., 2008a, *ApJ*, 681, 831
- Faucher-Giguère C.-A., Lidz A., Hernquist L., Zaldarriaga M., 2008b, *ApJ*, 682, L9
- Ferrara A., Loeb A., 2013, *MNRAS*, 431, 2826
- Finkelstein S. L. et al., 2015, *ApJ*, 810, 71
- Fletcher T. J., Tang M., Robertson B. E., Nakajima K., Ellis R. S., Stark D. P., Inoue A., 2019, *ApJ*, 878, 87
- Friedrich M. M., Mellema G., Alvarez M. A., Shapiro P. R., Iliiev I. T., 2011, *MNRAS*, 413, 1353
- Furlanetto S. R., Zaldarriaga M., Hernquist L., 2004, *ApJ*, 613, 1
- Gaikwad P. et al., 2020, *MNRAS*, 494, 5091
- Gallerani S., Ferrara A., Fan X., Choudhury T. R., 2008, *MNRAS*, 386, 359
- Garaldi E., Compostella M., Porciani C., 2019, *MNRAS*, 483, 5301
- Garzilli A., Bolton J. S., Kim T.-S., Leach S., Viel M., 2012, *MNRAS*, 424, 1723
- Grazian A. et al., 2016, *A&A*, 585, A48
- Greig B., Mesinger A., 2015, *MNRAS*, 449, 4246
- Greig B., Mesinger A., 2017, *MNRAS*, 472, 2651
- Greig B., Mesinger A., 2018, *MNRAS*, 477, 3217
- Greig B., Mesinger A., Haiman Z., Simcoe R. A., 2017, *MNRAS*, 466, 4239
- Greig B., Mesinger A., Bañados E., 2019, *MNRAS*, 484, 5094
- Guaity L. et al., 2016, *A&A*, 587, A133
- Gunn J. E., Peterson B. A., 1965, *ApJ*, 142, 1633
- Haardt F., Madau P., 2012, *ApJ*, 746, 125
- Hopkins P. F., Kereš D., Oñorbe J., Faucher-Giguère C.-A., Quataert E., Murray N., Bullock J. S., 2014, *MNRAS*, 445, 581
- Hopkins P. F. et al., 2018, *MNRAS*, 480, 800
- Hui L., Gnedin N. Y., 1997, *MNRAS*, 292, 27
- Hutter A., 2018, *MNRAS*, 477, 1549
- Hutter A., Dayal P., Müller V., Trott C. M., 2017, *ApJ*, 836, 176
- Iliiev I. T., Mellema G., Pen U.-L., Merz H., Shapiro P. R., Alvarez M. A., 2006, *MNRAS*, 369, 1625
- Iršič V. et al., 2017, *MNRAS*, 466, 4332
- Ishigaki M., Kawamata R., Ouchi M., Oguri M., Shimasaku K., Ono Y., 2018, *ApJ*, 854, 73
- Izotov Y. I., Worseck G., Schaerer D., Guseva N. G., Chisholm J., Thuan T. X., Fricke K. J., Verhamme A., 2021, *MNRAS*, 503, 1734
- Jiang L. et al., 2016, *ApJ*, 833, 222
- Kashino D., Lilly S. J., Shibuya T., Ouchi M., Kashikawa N., 2020, *ApJ*, 888, 6
- Keating L. C., Puchwein E., Haehnelt M. G., 2018, *MNRAS*, 477, 5501
- Keating L. C., Weinberger L. H., Kulkarni G., Haehnelt M. G., Chardin J., Aubert D., 2020a, *MNRAS*, 491, 1736
- Keating L. C., Kulkarni G., Haehnelt M. G., Chardin J., Aubert D., 2020b, *MNRAS*, 497, 906
- Kimm T., Cen R., 2014, *ApJ*, 788, 121
- Kuhlen M., Faucher-Giguère C.-A., 2012, *MNRAS*, 423, 862
- Kulkarni G., Keating L. C., Haehnelt M. G., Bosman S. E. I., Puchwein E., Chardin J., Aubert D., 2019, *MNRAS*, 485, L24
- Lee K.-G., Cen R., Gott J. Richard III, Trac H., 2008, *ApJ*, 675, 8
- Lee K.-G. et al., 2013, *AJ*, 145, 69
- Lee K.-G. et al., 2014, *ApJ*, 795, L12
- Lee K.-G. et al., 2015, *ApJ*, 799, 196
- Lidz A., McQuinn M., Zaldarriaga M., Hernquist L., Dutta S., 2007, *ApJ*, 670, 39
- Lidz A., Faucher-Giguère C.-A., Dall'Aglio A., McQuinn M., Fechner C., Zaldarriaga M., Hernquist L., Dutta S., 2010, *ApJ*, 718, 199
- Livermore R. C., Finkelstein S. L., Lotz J. M., 2017, *ApJ*, 835, 113
- Ma X., Kasen D., Hopkins P. F., Faucher-Giguère C.-A., Quataert E., Kereš D., Murray N., 2015, *MNRAS*, 453, 960
- Ma X., Quataert E., Wetzel A., Hopkins P. F., Faucher-Giguère C.-A., Kereš D., 2020, *MNRAS*, 498, 2001
- McGreer I. D., Mesinger A., D'Odorico V., 2015, *MNRAS*, 447, 499
- McQuinn M., Upton Sanderbeck P. R., 2016, *MNRAS*, 456, 47
- McQuinn M., Lidz A., Zahn O., Dutta S., Hernquist L., Zaldarriaga M., 2007, *MNRAS*, 377, 1043
- Madau P., Dickinson M., 2014, *ARA&A*, 52, 415
- Madau P., Haardt F., 2015, *ApJ*, 813, L8
- Manti S., Gallerani S., Ferrara A., Greig B., Feruglio C., 2017, *MNRAS*, 466, 1160
- Maselli A., Gallerani S., Ferrara A., Choudhury T. R., 2007, *MNRAS*, 376, L34
- Mason C. A., Treu T., Dijkstra M., Mesinger A., Trenti M., Pentericci L., de Barros S., Vanzella E., 2018, *ApJ*, 856, 2
- Meiksin A., 2020, *MNRAS*, 491, 4884
- Mesinger A., 2010, *MNRAS*, 407, 1328
- Mesinger A., ed., 2016, Understanding the Epoch of Cosmic Reionization, Vol. 423, Springer International Publishing, Switzerland
- Mesinger A., ed., 2019, The Cosmic 21-cm Revolution. IOP Publishing, Bristol, UK, p. 2514
- Mesinger A., Furlanetto S., 2007, *ApJ*, 669, 663
- Mesinger A., Furlanetto S., 2009, *MNRAS*, 400, 1461
- Mesinger A., Haiman Z., Cen R., 2004, *ApJ*, 613, 23
- Mesinger A., Furlanetto S., Cen R., 2011, *MNRAS*, 411, 955
- Mesinger A., Greig B., Sobacchi E., 2016, *MNRAS*, 459, 2342
- Miralda-Escudé J., Cen R., Ostriker J. P., Rauch M., 1996, *ApJ*, 471, 582
- Miralda-Escudé J., Haehnelt M., Rees M. J., 2000, *ApJ*, 530, 1

- Mitra S., Roy Choudhury T., Ferrara A., 2015, *MNRAS*, 454, L76
- Mortlock D. J. et al., 2011, *Nature*, 474, 616
- Moster B. P., Naab T., White S. D. M., 2013, *MNRAS*, 428, 3121
- Murray S., Greig B., Mesinger A., Muñoz J., Qin Y., Park J., Watkinson C., 2020, *J. Open Source Softw.*, 5, 2582
- Mutch S. J., Geil P. M., Poole G. B., Angel P. W., Duffy A. R., Mesinger A., Wyithe J. S. B., 2016, *MNRAS*, 462, 250
- Naidu R. P., Forrest B., Oesch P. A., Tran K.-V. H., Holden B. P., 2018, *MNRAS*, 478, 791
- Nasir F., D'Aloisio A., 2020, *MNRAS*, 494, 3080
- Nestor D. B., Shapley A. E., Kornei K. A., Steidel C. C., Siana B., 2013, *ApJ*, 765, 47
- Oesch P. A. et al., 2016, *ApJ*, 819, 129
- Oesch P. A., Bouwens R. J., Illingworth G. D., Labbé I., Stefanon M., 2018, *ApJ*, 855, 105
- Paardekooper J.-P., Khochfar S., Dalla Vecchia C., 2015, *MNRAS*, 451, 2544
- Pahl A. J., Shapley A., Steidel C. C., Chen Y., Reddy N. A., 2021, preprint ([arXiv:2104.02081](https://arxiv.org/abs/2104.02081))
- Paranjape A., Choudhury T. R., 2014, *MNRAS*, 442, 1470
- Park J., Mesinger A., Greig B., Gillet N., 2019, *MNRAS*, 484, 933
- Parsa S., Dunlop J. S., McLure R. J., 2017, *MNRAS*, 474, 2904
- Planck Collaboration XIII, 2016, *A&A*, 594, A13
- Planck Collaboration VI, 2020, *A&A*, 641, A6
- Price L. C., Trac H., Cen R., 2016, preprint ([arXiv:1605.03970](https://arxiv.org/abs/1605.03970))
- Puchwein E., Bolton J. S., Haehnelt M. G., Madau P., Becker G. D., Haardt F., 2015, *MNRAS*, 450, 4081
- Qin Y. et al., 2017, *MNRAS*, 472, 2009
- Qin Y., Poulin V., Mesinger A., Greig B., Murray S., Park J., 2020, *MNRAS*, 499, 550
- Qin Y., Mesinger A., Greig B., Park J., 2021, *MNRAS*, 501, 4748
- Rahmati A., Pawlik A. H., Raičević M., Schaye J., 2013, *MNRAS*, 430, 2427
- Rauch M., 1998, *ARA&A*, 36, 267
- Robertson B. E., Ellis R. S., Furlanetto S. R., Dunlop J. S., 2015, *ApJ*, 802, L19
- Sadoun R., Zheng Z., Miralda-Escudé J., 2017, *ApJ*, 839, 44
- Savaglio S., Panagia N., Padovani P., 2002, *ApJ*, 567, 702
- Schaye J., 2001, *ApJ*, 559, 507
- Schroeder J., Mesinger A., Haiman Z., 2013, *MNRAS*, 428, 3058
- Scoccimarro R., 1998, *MNRAS*, 299, 1097
- Shapiro P. R., Giroux M. L., Babul A., 1994, *ApJ*, 427, 25
- Shapley A. E., Steidel C. C., Strom A. L., Bogosavljević M., Reddy N. A., Siana B., Mostardi R. E., Rudie G. C., 2016, *ApJ*, 826, L24
- Slosar A. et al., 2013, *J. Cosmol. Astropart. Phys.*, 2013, 026
- Sobacchi E., Mesinger A., 2014, *MNRAS*, 440, 1662
- Songaila A., Cowie L. L., 2010, *ApJ*, 721, 1448
- Springel V. et al., 2005, *Nature*, 435, 629
- Steidel C. C., Bogosavljević M., Shapley A. E., Reddy N. A., Rudie G. C., Pettini M., Trainor R. F., Strom A. L., 2018, *ApJ*, 869, 123
- Sun G., Furlanetto S. R., 2016, *MNRAS*, 460, 417
- Tacchella S., Bose S., Conroy C., Eisenstein D. J., Johnson B. D., 2018, *ApJ*, 868, 92
- Thoul A. A., Weinberg D. H., 1996, *ApJ*, 465, 608
- Trac H. Y., Gnedin N. Y., 2011, *Adv. Sci. Lett.*, 4, 228
- Vanzella E. et al., 2010, *ApJ*, 725, 1011
- Vanzella E. et al., 2016, *ApJ*, 825, 41
- Venemans B. P. et al., 2013, *ApJ*, 779, 24
- Viel M., Haehnelt M. G., 2006, *MNRAS*, 365, 231
- Viel M., Lesgourgues J., Haehnelt M. G., Matarrese S., Riotto A., 2005, *Phys. Rev. D*, 71, 063534
- Viel M., Becker G. D., Bolton J. S., Haehnelt M. G., 2013, *Phys. Rev. D*, 88, 043502
- Wang F. et al., 2020, *ApJ*, 896, 23
- Weinberg D., et al., 1999, in Banday A. J., Sheth R. K., da Costa L. N., eds, *Evolution of Large Scale Structure : From Recombination to Garching*. p. proceeding of the MPA-ESO cosmology conference, Garching, Germany, p. 346
- Willott C. J. et al., 2010, *AJ*, 139, 906
- Worseck G. et al., 2014, *MNRAS*, 445, 1745
- Wu X.-B. et al., 2015, *Nature*, 518, 512
- Wyithe J. S. B., Bolton J. S., 2011, *MNRAS*, 412, 1926
- Wyithe J. S. B., Loeb A., 2013, *MNRAS*, 428, 2741
- Xu H., Wise J. H., Norman M. L., Ahn K., O'Shea B. W., 2016, *ApJ*, 833, 84
- Yang J. et al., 2020, *ApJ*, 897, L14
- Yèche C., Palanque-Delabrouille N., Baur J., du Mas des Bourboux H., 2017, *J. Cosmol. Astropart. Phys.*, 2017, 047
- Yung L. Y. A., Somerville R. S., Popping G., Finkelstein S. L., Ferguson H. C., Davé R., 2019, *MNRAS*, 490, 2855
- Zahn O., Lidz A., McQuinn M., Dutta S., Hernquist L., Zaldarriaga M., Furlanetto S. R., 2007, *ApJ*, 654, 12
- Zahn O., Mesinger A., McQuinn M., Trac H., Cen R., Hernquist L. E., 2011, *MNRAS*, 414, 727
- Zeng C., Hirata C. M., 2021, *ApJ*, 906, 124

APPENDIX A: TESTING THE FLUCTUATING GUNN-PETERSON APPROXIMATION

By computational necessity, our forward-models of the Lyman alpha forest are low resolution (~ 2 cMpc cells), and use the FGPA. The lack of small-scale structure in our models could also impact the large-scale opacity fluctuations we use as our summary statistic (see e.g. Viel & Haehnelt 2006 for an example of the impact on the flux power spectrum at lower redshifts where the transmission is higher). In order to account for this source of inaccuracy, here we compute an ECM, using a high-resolution hydro simulation from the Sherwood suite (Bolton et al. 2017).

The simulation used in this section was run with an updated version of GADGET-2 (Springel et al. 2005) and assumes a Λ CDM cosmology with parameters (Ω_m , Ω_b , Ω_Λ , h , σ_8 , n_s = 0.31, 0.048, 0.69, 0.68, 0.83, 0.96) from Planck Collaboration XIII (2016). It includes 512^3 baryonic and 512^3 dark matter particles within a cube of $80 h^{-1}$ cMpc on a side. The forest is calculated on a 2048^3 Eulerian grid, with a corresponding resolution of 0.057 cMpc (i.e. $80 h^{-1}$ cMpc/2048).

We compute Ly α spectra and τ_{eff} PDFs (i) assuming the FGPA on fields smoothed down to ~ 2 cMpc resolution; and (ii) using the native high-resolution fields (0.057 cMpc), including peculiar velocities, and a Voigt profile for the Ly α absorption. The former corresponds to approximations used in our forward-models, while the latter we take as the ‘true’ spectra.

We present the results in Fig. A1, including the hydrogen density, peculiar velocity, gas temperature, residual neutral hydrogen fraction, and the inferred Ly α optical depth as well as its transmission after rescaling the mean flux to the observed one (i.e. 0.0765; Bosman et al. 2021). Since the hydro simulation does not include patchy reionization, we focus on the $z = 5.4$ snapshot (this is at the lowest end of the redshift range of interest, where patchy reionization should have the smallest impact; see Fig. 4).

As expected, having a lower resolution reduces the small-scale structure in the forest. Compared to the high-resolution forest, the FGPA has a smaller variance on small scales and a ~ 30 per cent larger f_{rescale} is required to match the same mean flux.

However, the differences are much smaller in the τ_{eff} CDFs, averaged over $50 h^{-1}$ cMpc. These are shown in the top panel of Fig. A2 for both the high-resolution spectra and the FGPA. We include also the observational data in grey. As expected, the simulation cannot match the observed distributions, owing to its small size and uniform UVB. However, the fact that the FGPA and the high-resolution spectra can produce comparable large-scale opacity fluctuations is highly encouraging of our approach.

In Appendix B, we show the corresponding ECM. This error is added to the CV error used when computing the forest likelihood in our MCMCs.

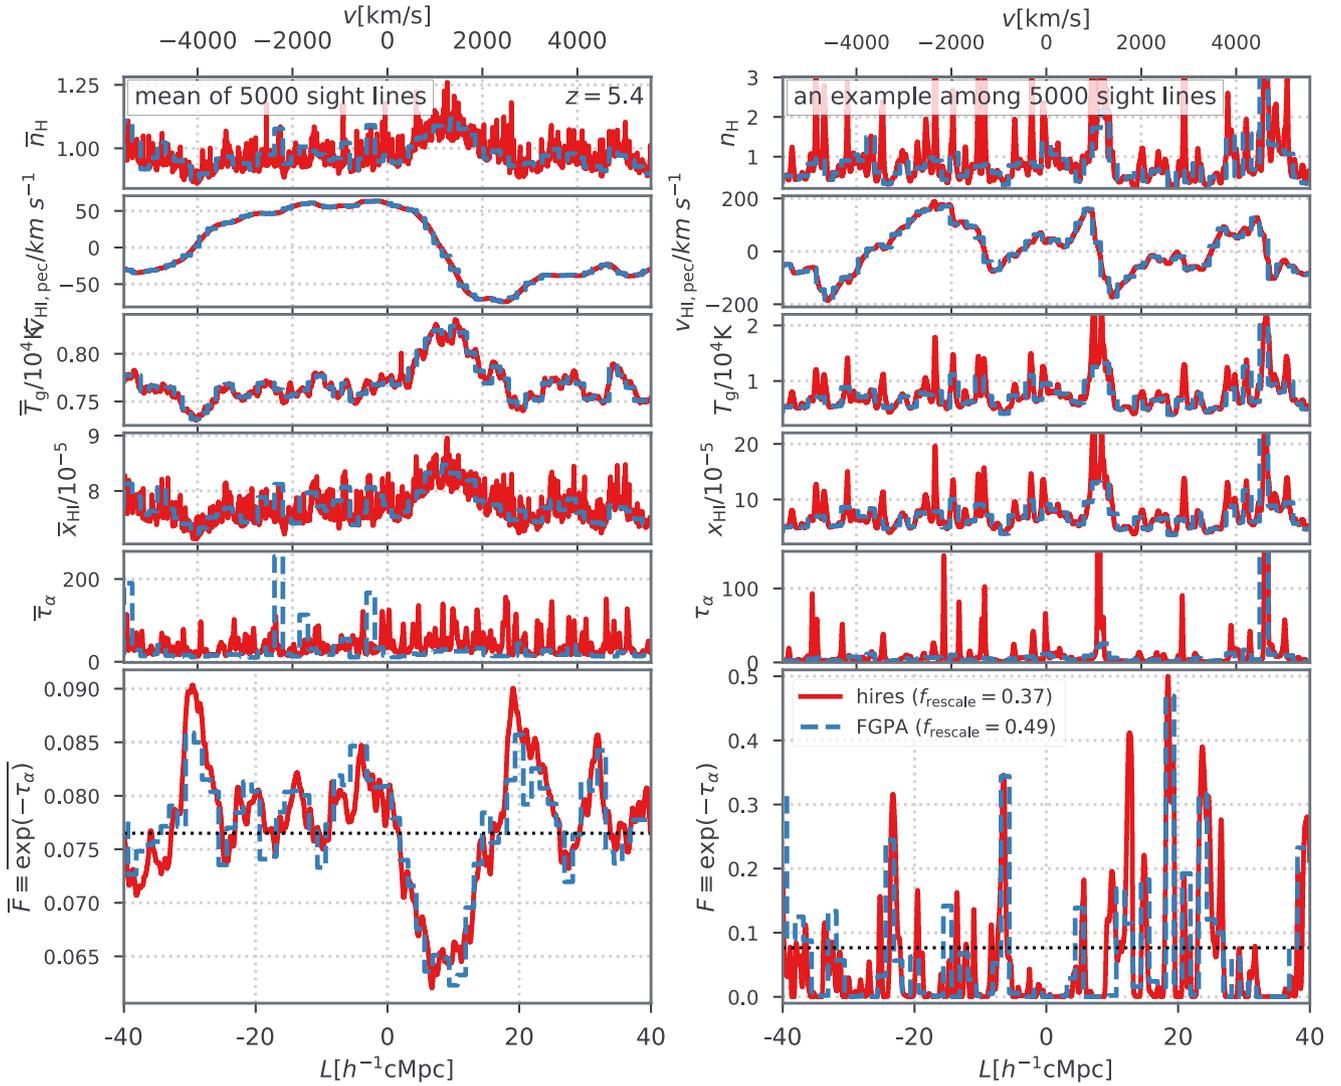


Figure A1. Properties including, from top to bottom, gas density, peculiar velocity, gas temperature, residual neutral hydrogen fraction, Ly α optical depth and its transmission after rescaling the mean flux to the observed one (indicated by the horizontal dotted lines; Bosman et al. 2021), along 5000 lines of sight across the entire Sherwood simulation box ($80 h^{-1} \text{ cMpc}$) at $z = 5.4$. Note that the mean of all sightlines are presented on the left while the right ones show one example sightline. The red solid curves indicate results from the reference model having high spectral resolution and integrating over the full Ly α profile when evaluating its optical depth (*hires*). The blue dashed lines correspond to results assuming FGPA and a lower resolution similar to what is used in the main context (FGPA).

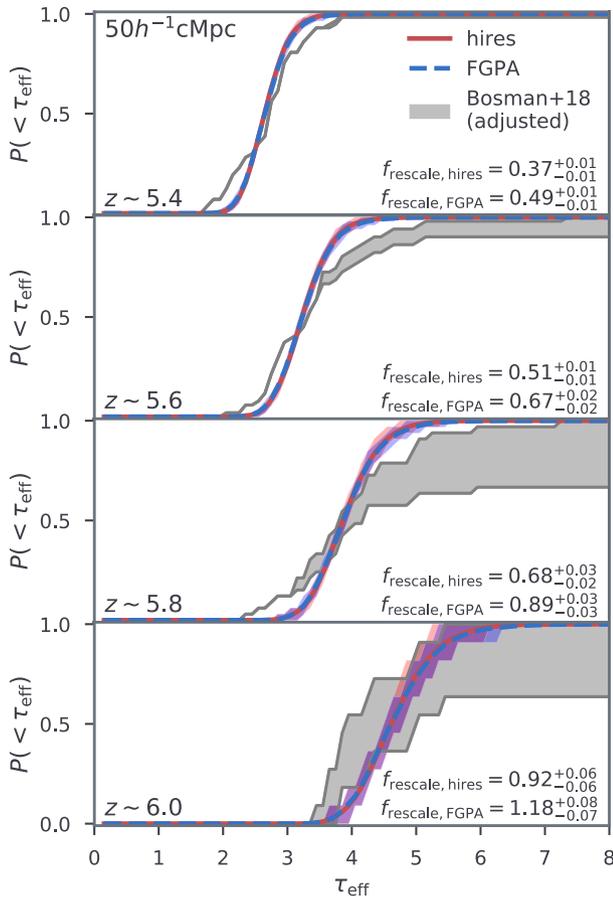


Figure A2. The CDF of τ_{eff} (averaged over $50 h^{-1} \text{cMpc}$) at $z = 5.4-6$ from the high-resolution Sherwood hydrodynamical simulation (*hires*; red) and assuming FGPA with low resolution (*FGPA*; blue). The coloured lines with shaded region indicate the mean and [14,86] percentiles drawn from 500 realizations. The observed CDFs are indicated in grey (see more in Section 3).

APPENDIX B: ERROR COVARIANCE MATRICES

We present the ECM²² of the τ_{eff} PDF introduced by making the FGPA or from the CV at $z = 5.4-6$ in Fig. B1. Since ECMs are symmetric, we only show half of the FGPA and CV matrices, and present them together (i.e. top is FGPA and bottom is CV, respectively). For this plot, the CV is estimated using the 500 realizations²³ drawn from our MAP model (see more in Section 2). We note that the CV ECM is recomputed on the fly for each sample in our MCMC; unfortunately, it is computationally impractical to do this also for the FGPA ECM.²⁴ The latter was computed as described

²²An identity matrix with a normalization of 10^{-5} is imposed on the total ECM as a precautionary measure.

²³For computing efficiency, only 150 realizations are generated during the MCMC.

²⁴As we do not know a priori what are the true values of these fields in the real Universe, the FGPA ECM should be recomputed for each forward-modelled universe. In other words, we use $\Sigma_{\text{GP}}(\theta = \theta_{\text{Sherwood}})$ in our MCMC, where θ_{Sherwood} encapsulates all of the choices and approximations made to generate the Sherwood simulation. However, we should instead know the general error covariance, $\Sigma_{\text{GP}}(\theta)$, evaluating it on-the-fly for any astrophysical parameter combination θ . Unfortunately, this is computationally impractical. Using different hydro simulations, in the future we will explore how the covariance

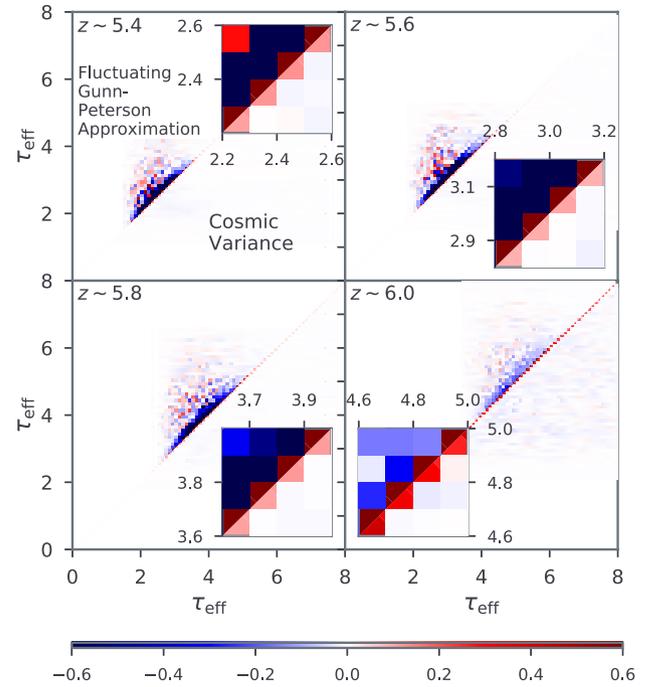


Figure B1. The ECMs of the τ_{eff} PDF introduced by the FGPA and CV at $z = 5-6$. The two symmetric matrices are presented together in each subpanel separated by the diagonal. The coefficients of the ECMs are represented by varying colours shown in the colourbar. A zoom-in subpanel is presented to show more details around the maximum diagonal element of either matrix.

in the previous Appendix at $z = 5.4$, and adjusted for higher redshifts by shifting the PDF such that the mean flux matches the observations (see the corresponding CDFs in Fig. A2).

There are both positive (correlation; red) and negative values (anticorrelation; blue) between pairs of different τ_{eff} bins, with more showing anticorrelation when being closer to the diagonal (i.e. nearby bins). This is expected as increasing one histogram bin can be roughly compensated by a decrease in a nearby bin. The dominant component in the total ECM is caused by the diagonal coefficients of the FGPA, though its coefficients along the diagonal (i.e. uncertainties in each τ_{eff} bin of the PDF) decreases towards lower and higher redshifts. This is because when the sample size decreases, the PDF, with a lower value, also possesses a smaller absolute difference between the FGPA and the reference results (see e.g. the corresponding CDF in the lower right-hand panel of Fig. A1). On the other hand, the CV ECM shows the opposite trend, which is also caused by the sample size. Since different lines of sight are randomly selected when estimating the CV, a smaller sample size leads to larger differences between different realizations and therefore larger CV.

matrix changes for a few different values of θ , and estimate the corresponding impact on the posteriors.

This paper has been typeset from a \LaTeX file prepared by the author.