

On the rise and fall of galactic ionizing output at the end of reionization

Christopher Cain,^{1,2★} Anson D’Aloisio,² Garrett Lopez,² Nakul Gangolli² and Joshua T. Roth²

¹*School of Earth and Space Exploration, Arizona State University, Tempe, AZ 85281, USA*

²*Department of Physics and Astronomy, University of California, Riverside, CA 92521, USA*

Accepted 2024 April 30. Received 2024 April 26; in original form 2023 December 4

ABSTRACT

Quasar absorption spectra measurements suggest that reionization proceeded rapidly, ended late at $z \sim 5.5$, and was followed by a flat ionizing background evolution. Simulations that reproduce this behaviour often rely on a fine-tuned galaxy ionizing emissivity, which peaks at $z \sim 6$ –7 and drops a factor of 1.5–2.5 by $z \sim 5$. This is puzzling since the abundance of galaxies is observed to grow monotonically during this period. Explanations for this include effects such as dust obscuration of ionizing photon escape and feedback from photoheating of the IGM. We explore the possibility that this drop in emissivity is instead an artefact of one or more modelling deficiencies in reionization simulations. These include possibly incorrect assumptions about the ionizing spectrum and/or inaccurate modelling of IGM clumping. Our results suggest that the need for a drop could be alleviated if simulations are underestimating the IGM opacity from massive, star-forming haloes. Other potential modelling issues either have a small effect or require a steeper drop when remedied. We construct an illustrative model in which the emissivity is nearly flat at reionization’s end, evolving only ~ 0.05 dex at $5 < z < 7$. More realistic scenarios, however, require a ~ 0.1 –0.3 dex drop. We also study the evolution of the Ly α effective optical depth distribution and compare to recent measurements. We find that models that feature a hard ionizing spectrum and/or are driven by faint, low-bias sources most easily reproduce the mean transmission and optical depth distribution of the forest simultaneously.

Key words: radiative transfer – galaxies: high-redshift – galaxies: intergalactic medium – cosmology: dark ages, reionization, first stars.

1 INTRODUCTION

The past decade has seen great progress towards constraining the timing of cosmic reionization. Measurements of the Thomson scattering optical depth by Planck have localized the mid-point of reionization to $z \approx 7.5 \pm 0.8$ (Planck Collaboration VI 2020). QSO damping wing measurements have suggested that the IGM was significantly neutral at $z \sim 7$ –8 (Davies et al. 2018; Wang et al. 2020), while Ly α emitter surveys have hinted at large ionized bubbles at the same redshifts (Ouchi et al. 2018; Hu et al. 2019; Endsley & Stark 2022). The tail end of reionization has been probed by the Ly α forest of high-redshift QSOs (Fan et al. 2006; Becker et al. 2015; Bosman et al. 2018, 2022; Eilers, Davies & Hennawi 2018; Zhu et al. 2022) and measurements of the ionizing photon mean free path (MFP; Worseck et al. 2014; Becker et al. 2021; Zhu et al. 2023). These observations support a relatively late end to reionization at $z \approx 5$ –5.5 (Kulkarni et al. 2019; Keating et al. 2020a; Nasir & D’Aloisio 2020; Cain et al. 2021; Davies et al. 2021; Qin et al. 2021).

Despite our improving understanding of reionization’s timing, relatively little is known about the properties of the sources that drove it. Measurements of the high-redshift ultraviolet luminosity function (UVLF; Bouwens et al. 2015, 2021; Finkelstein et al. 2019) and simulations (e.g. Ocvirk et al. 2020; Kannan et al. 2022)

suggest that galaxies with physically reasonable ionizing properties could have completed reionization. Conversely, measurements of the quasar luminosity function (Willott et al. 2010; McGreer et al. 2013; Georgakakis et al. 2015; Matthee et al. 2024), simulations (Trebitsch et al. 2021; Kannan et al. 2022), and constraints on the temperature of the IGM (D’Aloisio et al. 2017) disfavour quasars as the dominant sources. However, there remains ongoing debate as to which galaxies were the dominant ionizing photon producers (Robertson et al. 2015; Finkelstein et al. 2019; Naidu et al. 2020), and how the ionizing properties of these galaxies evolved over time.

If galaxies drove reionization, the total ionizing photon emissivity is given by

$$\dot{N}_\gamma \equiv \langle f_{\text{esc}} \xi_{\text{ion}} \rangle \rho_{\text{UV}}, \quad (1)$$

where ρ_{UV} is the integrated galaxy UVLF, ξ_{ion} is the ionizing efficiency of galaxies, and f_{esc} is their ionizing photon escape fraction, and the average of $f_{\text{esc}} \xi_{\text{ion}}$ is over the galaxy population and weighted by the luminosity L_{UV} . The UVLF at $M_{\text{UV}} \lesssim -17$ has been measured out to $z \approx 10$ with *Hubble Space Telescope* (Finkelstein et al. 2019; Bouwens et al. 2021) and at higher redshifts with *JWST* (e.g. Adams et al. 2023), and most recently even fainter galaxies are being probed using strong lensing up to $z \sim 8$ (Atek et al. 2024). Assuming a constant $\log(\xi_{\text{ion}}) = 25.2$, Robertson et al. (2013) found that a constant $f_{\text{esc}} = 0.2$ is sufficient for galaxies to reionize the universe by $z \sim 6$ given reasonable assumptions about the faint end of the UVLF. The grey curve in Fig. 1 shows $\dot{N}_\gamma$ from Robertson et al. (2015).

* E-mail: clcain3@asu.edu




Figure 1. Ionizing emissivity versus redshift for several models and simulations in the literature. The faded grey curve shows the emissivity computed assuming $f_{\text{esc}} = 0.2$ and $\log(\xi_{\text{ion}}) = 25.2$, as in Robertson et al. (2015). The remaining curves come from simulations by Kulkarni et al. (2019), Keating et al. (2020a), Cain et al. (2021), Ocvirk et al. (2021), Yeh et al. (2022), and Gaikwad et al. (2023), which reproduce QSO observations at $5 < z < 6$. With the exception of THESAN-1 (Yeh et al. 2022), all the simulations find that a drop in $\dot{N}_\gamma$ starting at $z \sim 6.5$ is required to reproduce these observations. Our goal is to determine whether this drop is likely to be real, or an artefact of one or more modelling deficiencies in simulations.

It grows rapidly with cosmic time (decreasing redshift), tracing the growth of the UVLF and the underlying halo mass function.

However, QSO observations at the tail end of reionization may be telling a different story. Fig. 1 compares $\dot{N}_\gamma$ from Robertson et al. (2015) to that from several recent simulations in the literature. All these simulations reproduce reasonably well (or have been tuned to reproduce) the $5 < z < 6$ Ly α forest and/or MFP. All of these models, except THESAN-1 (Kannan et al. 2022; Yeh et al. 2022), display a *drop* in $\dot{N}_\gamma$, starting at $z \sim 6$ – 6.5 . The drop is necessary in these models to prevent the mean Ly α forest transmission from overshooting the measurements at $z < 6$. Its magnitude ranges from a factor of 1.5 to 2.5, and happens within 300 Myr or less. This behaviour is only explainable by a factor of several evolution in $\langle f_{\text{esc}} \xi_{\text{ion}} \rangle$ over this relatively short time period.

There are several reasons why f_{esc} and/or ξ_{ion} may decrease near the end of reionization. If f_{esc} and/or ξ_{ion} depends strongly on galaxy host halo mass (Finkelstein et al. 2019; Rosdahl et al. 2022; Yeh et al. 2022), the evolution of the halo mass function (HMF) could drive a decrease in $\langle f_{\text{esc}} \xi_{\text{ion}} \rangle$. The ionizing properties of galaxies (particularly faint ones) may be affected by feedback from IGM photoheating and/or supernovae (Shapiro, Giroux & Babul 1994; Kimm & Cen 2014; Wu et al. 2019a; Ocvirk et al. 2021). There may be evolution and/or variation in the intrinsic ionizing properties of high-redshift stellar populations (Maseda et al. 2020; Atek et al. 2022). Dust may reduce the ionizing photon output of the most massive galaxies (Kostyuk et al. 2023; Lewis et al. 2023).

For the most part, the studies referenced in Fig. 1 make little attempt to explain the physical mechanism(s) underlying the drop in $\dot{N}_\gamma$. The $\dot{N}_\gamma$ evolution in Kulkarni et al. (2019), Keating et al. (2020a), Cain et al. (2021), and Gaikwad et al. (2023) is not a prediction of any underlying galaxy model, but is instead tuned to reproduce Ly α forest and/or MFP observations. The fiducial simulation of Ocvirk et al. (2021), which has a drop, does include star formation. In their case,

the drop is driven by a combination of photoheating feedback, which suppresses gas accretion on to galaxies, and supernova feedback, which further disrupts star formation. However, they account for these effects by assuming a sharp temperature threshold for star formation, and it is unclear how much of the effect results from this choice. In THESAN-1, which uses the IllustrisTNG galaxy formation model (Vogelsberger et al. 2014; Weinberger et al. 2016; Pillepich et al. 2017), the global escape fraction declines by a factor of 2–3 during reionization (Yeh et al. 2022), in part due to feedback effects (Garaldi et al. 2022). However, this is not enough to produce a decrease in $\dot{N}_\gamma$, nor do they find that one is necessary to reproduce QSO observations. However, as we will see in Section 6.1, their agreement with QSO observations may be in part due to their use of the reduced speed of light.

Another possibility is that the drop in $\dot{N}_\gamma$ is an artefact of inaccurate or incomplete modelling of the IGM. Keating et al. (2020a) speculated that it may result from an inaccurate treatment of the IGM thermal history. Simulations may also be missing ionizing photon absorbers that would otherwise regulate the growth of the UV background (Cain et al. 2021). These absorbers could be missing due to a lack of spatial resolution, which is a problem for most reionization simulations in representative cosmological volumes (Emberson, Thomas & Alvarez 2013). Absorption in and around massive star-forming haloes may be sensitive to details of galaxy dynamics and evolution (see appendix C of Wu et al. 2019b). Simulations that model these details approximately (or not at all) may also be underestimating this source of IGM opacity. These, and several other potential IGM modelling considerations, including the spectrum of the ionizing radiation and behaviour of ionizing recombination radiation, affect the relationship between the Ly α forest and $\dot{N}_\gamma$. Our goal is to determine whether one or more of these issues, if resolved, might alleviate the need for a drop in $\dot{N}_\gamma$.

This work is organized as follows. In Section 2, we discuss several physical and numerical modelling effects that influence the relationship between $\dot{N}_\gamma$ and the Ly α forest. Section 3 describes our numerical methods. In Section 4, we study each modelling effect individually. We consider their combined effects, and implications for the properties of ionizing sources, in Section 5. In Section 6, we consider numerical (non-physical) modelling effects. We conclude in Section 7. Throughout this work, we assume the following cosmological parameters: $\Omega_m = 0.305$, $\Omega_\Lambda = 1 - \Omega_m$, $\Omega_b = 0.048$, $h = 0.68$, $n_s = 0.9667$, and $\sigma_8 = 0.82$, consistent with Planck Collaboration VI (2020) results. All distances are quoted in comoving units unless otherwise specified.

2 RELEVANT MODELLING CONSIDERATIONS

In this section, we will discuss several physical effects that affect the relationship between $\dot{N}_\gamma$ and the forest in simulations. We will also discuss how each of these might affect the need for a drop.

2.1 Spectrum of the ionizing radiation

The ionizing photon spectrum emitted by galaxies is often parametrized as a power law of the form

$$\frac{d\dot{N}_\gamma}{d\nu} \propto \frac{J_\nu}{h_p \nu} \propto \nu^{-\alpha-1}, \quad (2)$$

where J_ν is the galaxy’s spectral energy distribution above 1 Ryd, h_p is Planck’s constant, ν is frequency, and α is the spectral index. Typical values for α range from 0.5 to 2.5, depending on assumptions about the ionizing properties of stars (Bressan et al. 2012; Choi,

Conroy & Byler 2017) and spectral hardening by the interstellar and circumgalactic medium (ISM and CGM) (Madau 1995; Faucher-Giguère et al. 2009a; Haardt & Madau 2012). Ionizing radiation may be further hardened by the IGM after emerging from galaxies.

The Ly α forest is sensitive to α through the latter's effect on the photoionization rate Γ_{HI} and temperature T . In optically thin gas,

$$\Gamma_{\text{HI}} = \int_{\nu_0}^{4\nu_0} d\nu \frac{dN_\gamma}{d\nu} c \sigma_{\text{HI}}(\nu) = N_\gamma c \langle \sigma_{\text{HI}} \rangle, \quad (3)$$

where N_γ is the local ionizing photon number density, c is the speed of light, $\langle \sigma_{\text{HI}} \rangle$ is the spectrum-averaged H I cross-section, and $h_p \nu_0 \equiv 13.6 \text{ eV}$ is the ionization energy of H I.¹ For $\alpha = 2.5$ (0.5), $\langle \sigma_{\text{HI}} \rangle = 3.11 \times 10^{-18} \text{ cm}^2$ ($1.93 \times 10^{-18} \text{ cm}^2$), a factor of 1.6 difference. The temperature is sensitive to α through the photoheating rate,

$$\mathcal{H} = \int_{\nu_0}^{4\nu_0} d\nu \frac{dN_\gamma}{d\nu} c \sigma_{\text{HI}}(\nu) (h_p \nu - h_p \nu_0) \propto N_\gamma \langle E_{\text{ion, HI}} \rangle, \quad (4)$$

where $\langle E_{\text{ion, HI}} \rangle$ is the average energy injection per ionization per H I atom. Since harder photons carry more energy, smaller α will lead to larger $\langle E_{\text{ion, HI}} \rangle$ and a hotter IGM. For $\alpha = 2.5$ (0.5), $\langle E_{\text{ion, HI}} \rangle = 3.16$ (5.38) eV, a factor of 1.7 difference in heating rate.² The temperature the gas is heated to initially by reionization, T_{reion} , also depends on α , although D'Aloisio et al. (2019) found that sensitivity to be modest.

The opacity of the IGM to Ly α photons scales with Γ_{HI} and T as

$$\tau_{\text{Ly}\alpha} \propto n_{\text{HI}} \propto \frac{\alpha_A(T)}{\Gamma_{\text{HI}}} \propto T^{-0.7} \Gamma_{\text{HI}}^{-1}, \quad (5)$$

where α_A is the case A recombination coefficient of H II. Because Γ_{HI} and T scale oppositely with α , the effect of changing α will partially cancel in equation (5). The fact that $\tau_{\text{Ly}\alpha}$ scales more strongly with Γ_{HI} than with T suggests that it should increase for harder spectra (smaller α), resulting in less Ly α transmission at fixed N_γ . This would allow $\dot{N}_\gamma$ to be higher at fixed Ly α transmission, possibly alleviating the need for a drop in $\dot{N}_\gamma$.

2.2 Ionizing recombination radiation

Ionizing recombination photons (IRPs) – produced by H I recombinations to the ground state – also affect the relationship between $\dot{N}_\gamma$ and the Ly α forest. These photons are produced mainly within dense gas clumps with high recombination rates, the same structures that act as sinks (see next section). Modelling IRPs self-consistently requires resolving the structures that create them, a challenging computational task. Thus, IRPs are often treated approximately assuming one of two limiting cases (although some codes include the full treatment, e.g. Rosdahl et al. 2013; Kannan et al. 2019). They are assumed to (1) have no effect on the local IGM or (2) get re-absorbed immediately by neighbouring H I after being emitted (the so-called on-the-spot approximation). These approximations are referred to as case A and B, respectively. Case A (B) is most accurate in low (high) density gas where the local MFP to IRPs is long (short).

Transmission in the $5 < z < 6$ Ly α forest is set by underdense gas, where it is safe to assume that the case A approximation is

¹The cut-off at 4 Ryd owes to absorption of He II ionizing photons within galaxies.

²To get these values, we first assume that the ionizing photon number spectrum is given by equation (2), with a sharp cut-off at 4 Ryd (the He II ionization threshold). To get the average energy injection per ionization over the whole spectrum, we evaluate $\langle E_{\text{ion, HI}} \rangle = \frac{\int_{\nu_0}^{4\nu_0} d\nu \frac{dN_\gamma}{d\nu} \sigma_{\text{HI}}^v (h\nu - 13.6 \text{ eV})}{\int_{\nu_0}^{4\nu_0} d\nu \frac{dN_\gamma}{d\nu} \sigma_{\text{HI}}^v}$, which assumes the optically thin limit for absorption.

valid locally. However, the fate of IRPs produced in dense gas clumps is less clear. The fraction of IRPs escaping these clumps, and their effect on the IGM opacity, are sensitive to the IGM column density distribution (Faucher-Giguère et al. 2009b; Altay et al. 2011; McQuinn, Oh & Faucher-Giguère 2011). If a significant fraction of IRPs escape, they may affect the properties of the sinks and the ionizing background. It is possible that they could affect the $\dot{N}_\gamma$ –forest relationship and the need for a drop in $\dot{N}_\gamma$.

2.3 Ionizing photon sinks

Ionizing photon sinks – dense gas clumps with high recombination rates – set the ionizing photon MFP and the reionization photon budget. They also regulate the ionizing background at the end of reionization. In Cain et al. (2021), we demonstrated that sinks could prevent the MFP from over-shooting the measurements at $z \sim 5$ without a drop in $\dot{N}_\gamma$.

It is possible that the smallest, most abundant sinks are missing from simulations. These occupy non-star-forming ‘mini-haloes’ with dark matter (DM) masses of 10^4 – $10^8 M_\odot$. They are photoevaporated over a ~ 100 Myr time-scale after the IGM surrounding them is reionized (Shapiro, Iliev & Raga 2004; Iliev, Shapiro & Raga 2005; Chan et al. 2024). Despite being short-lived, mini-haloes contribute significantly to the IGM opacity and significantly increase the photon budget (Park et al. 2016; D'Aloisio et al. 2020; Nasir et al. 2021; Chan et al. 2024). Mini-haloes are difficult to resolve in simulations because they can be as small as the pre-reionization baryon Jeans scale, which can be a kpc or less depending on the uncertain IGM thermal history. Their response to reionization is sensitive to the interplay between self-shielding and hydrodynamics, requiring fully coupled RT/hydro to model accurately. Another possibility is that simulations are underestimating the opacity from the most massive sinks, which occupy haloes massive enough ($> 10^9 M_\odot$) to self-shield long after reionization ends. The opacity after reionization is dominated by these objects (Prochaska, O'Meara & Worseck 2010), and they may be important during reionization (Muñoz et al. 2016). Since haloes this massive usually host star-forming galaxies (Okamoto, Gao & Theuns 2008; Finlator et al. 2017), their opacity may be sensitive to how galaxies form and evolve. For example, Wu et al. (2019a) found that how they modelled feedback in their galaxies significantly affected clumpiness of the surrounding IGM (their Appendix C). While several reionization simulations include galaxy formation models (e.g. Ocvirk et al. 2016; Rosdahl et al. 2018; Kannan et al. 2022), these differ in how they handle processes like AGN feedback and star formation. Many reionization simulations do not include galaxy physics (e.g. Keating et al. 2020a; Cain et al. 2021).

2.4 Clustering of the ionizing sources

The sources of reionization (galaxies) are clustered on spatial scales of 10s–100s of Mpc. If reionization was driven by the brightest galaxies (Naidu et al. 2020; Matthee et al. 2022), the sources would be more highly biased/clustered than if faint galaxies dominated (Finkelstein et al. 2019; Atek et al. 2024). Several recent studies, observational and theoretical, suggest fainter galaxies may have higher average f_{esc} and/or ξ_{ion} (Begley et al. 2022; Rosdahl et al. 2022; Saldana-Lopez et al. 2023; Atek et al. 2024). However, Matthee et al. (2022) and Naidu et al. (2022) recently argued that the brightest Lyman α emitters (LAEs) may have been the primary drivers. Their model may be supported by observations of bright LAEs at

$z \geq 7$, which hint at the presence of large, highly ionized regions surrounding them (Mason et al. 2018; Endsley & Stark 2022).

The clustering properties of the sources can affect the Ly α forest. The more biased the sources are, the more the ionizing background will fluctuate on large scales during reionization, resulting in regions with enhanced Ly α transmission near the brightest sources. This effect should be important as long as the MFP is less than or comparable to the clustering scale of the sources, which may be true even after reionization (Davies & Furlanetto 2016). The ionizing background surrounding the brightest sources may also affect the response of the sinks to reionization (Park et al. 2016; D’Aloisio et al. 2020; Chan et al. 2024).

3 NUMERICAL METHODOLOGY

In this section, we describe the numerical methods used in this work. Sections 3.1–3.3 review the methodology discussed in Cain et al. (2023), and the remaining sections introduce new features in our framework.

3.1 Large-scale radiative transfer

We ran radiative transfer (RT) simulations of reionization using the ray-tracing RT code first introduced in Cain et al. (2021) and further described in Cain et al. (2023), which we refer to as FlexRT (flexible radiative transfer) hereafter.

FlexRT solves the RT equation in post-processing on a time-series of cosmological density fields and with a uniform RT grid. Ionizing sources (see next section) are binned to their nearest RT cells and rays are cast from the centres of source cells. As rays travel, the optical depth through each intersected cell is computed, and photons are deposited accordingly. Rays can adaptively split and merge using the HealPix formalism (Gorski et al. 1999; Abel & Wandelt 2002; Trac & Cen 2007) to maintain a desired angular resolution. In this work we track 12 directions in the radiation field, with a maximum number of rays per cell of 28. Rays are deleted when their photon count is attenuated by a factor of 10^{10} .

Instead of solving for the H I number density ($n_{\text{H I}}$) in each cell, FlexRT adopts a more general approach to compute ionizing photon opacities. The optical depth at frequency ν encountered by ray j intersecting cell i is³

$$\tau_{ij}^{\nu} = \frac{x_{\text{ion}}^i \Delta s^{ij}}{\lambda_i^{\nu}}, \quad (6)$$

where Δs^{ij} is the distance travelled by ray j through cell i , λ_i^{ν} is the MFP at frequency ν , and x_{ion}^i is the ionized fraction. The photoionization rate is given by

$$\Gamma_{\text{H I}}^i = \frac{\langle \sigma_{\text{H I}} \rangle_v^i \langle \lambda \rangle_v^i}{x_{\text{ion}}^i V_{\text{cell}}^i \Delta t} \sum_{j=1}^{N_{\text{rays}}} \sum_{\nu=1}^{N_{\text{freq}}} N_{\gamma,0}^{ij,\nu} [1 - \exp(-\tau_{ij}^{\nu})], \quad (7)$$

where $N_{\gamma,0}^{ij,\nu}$ is the number of photons initially in frequency bin ν of ray j , $\langle \sigma_{\text{H I}} \rangle_v^i$ and $\langle \lambda \rangle_v^i$ are the H I cross-section and MFP averaged over the ionizing spectrum incident on cell i , and N_{freq} is the number

³To get this relationship between τ and x_{ion} , we assume that I-fronts travel along one cell axis and that rays travel along this direction and enter at the ionized side of the cell opposite the I-front. In this limit, the fraction of the cell traversed by the ray before reaching the I-front is x_{ion} . We note that our results are not very sensitive to this choice, since x_{ion} cancels out in equation (7) in the limit that $\tau \ll 1$.

of frequency bins. The inclusion of multifrequency RT is a new feature in FlexRT, and is further discussed in Section 3.4. The other quantities have the same meaning as in equation 1 of Cain et al. (2023). We will describe in Sections 3.3 and 3.4 how we calculate $\langle \lambda \rangle_{\nu}$. Appendix A gives the complete derivation of equation (7).⁴

We model the propagation of sub-resolved I-fronts in the ‘moving screen’ approximation, which assumes a sharp boundary between ionized and neutral gas. The I-front speed is given by

$$v_{\text{IF}} = \frac{F_{\gamma}^{\text{inc}}}{(1 + \chi)n_{\text{H}}}, \quad (8)$$

where F_{γ}^{inc} is the ionizing flux incident on the neutral part of the cell (after absorption by the intervening ionized part) and the factor of $1 + \chi = 1.082$ accounts for single ionization of He. The gas temperature behind the I-front, T_{reion} , is estimated using the flux-based method prescribed in D’Aloisio et al. (2019), and the temperature evolution thereafter is calculated using their equation (6).

3.2 Sources and density fields

Our sources are haloes taken from a DM-only simulations run with the N -body code used in Trac, Cen & Mansfield (2015) with a box size of $200 h^{-1} \text{ Mpc}$. We used $N_{\text{dm}} = 3600^3$ particles, resulting in a complete halo mass function (to within 10 per cent) down to a minimum halo mass of $3 \times 10^9 h^{-1} M_{\odot}$ (≈ 200 DM particles). Haloes were saved every 10 Myr from $z = 12$ to 4.8. We assigned UV luminosities to haloes by abundance-matching to the UVLF of Finkelstein et al. (2019). The global (integrated) ionizing emissivity, $\dot{N}_{\gamma}$, is a free function of redshift that we tune to match observations – we describe how we do this in Section 4. Ionizing photons are distributed between haloes according to

$$\dot{n}_{\gamma} \propto \begin{cases} L_{\text{UV}}^{\beta} & M_{\text{halo}} \geq M_{\text{min}} \\ 0 & M_{\text{halo}} < M_{\text{min}} \end{cases}, \quad (9)$$

where $\dot{n}_{\gamma}$ is the emissivity of a given halo, L_{UV} is its UV luminosity, M_{min} is the minimum mass of haloes that produce ionizing photons, and β parametrizes how the emissivity is distributed between bright and faint sources. We will use $M_{\text{min}} = 3 \times 10^9 h^{-1} M_{\odot}$ and $\beta = 1$ as our fiducial values, and will vary both in Section 5.2. Density fields are taken from an Eulerian hydro-dynamical simulation with $N = 1024^3$ gas cells run with the code of Trac & Pen (2004), with the same box size and large-scale initial conditions used in the N -body run. We re-binned the density fields to a coarse-grained resolution of $N_{\text{rt}} = 200^3$ for our FlexRT runs.

3.3 Sub-grid models for λ_{ν}

Our prescription for λ_{ν} in equation (7) follows the formalism outlined in (Cain et al. 2023), with the improvement that we now account for the frequency dependence of λ_{ν} (see next section). We use a suite of high-resolution hydro/RT simulations in $1 h^{-1} \text{ Mpc}$ volumes [like those in D’Aloisio et al. (2020)] run with a modified version of the code of Trac & Cen (2007). These simulations sample a range of box-scale overdensities (using DC modes, Gnedin, Kravtsov & Rudd 2011), $\Gamma_{\text{H I}}$, and reionization redshifts (z_{reion}) to calibrate a sub-grid model for λ_{ν} , which evolves in our simulations according to

$$\frac{d\lambda_{\nu}}{dt} = \left. \frac{\partial \lambda_{\nu}}{\partial t} \right|_{\Gamma_{\text{H I}}} + \left. \frac{\partial \lambda_{\nu}}{\partial \Gamma_{\text{H I}}} \right|_t \frac{d\Gamma_{\text{H I}}}{dt} - \frac{\lambda_{\nu} - \lambda_{\nu,0}}{t_{\text{relax}}}, \quad (10)$$

⁴Appendix A of Cain et al. (2023) derives only the monochromatic version of equation (7) (their equation 1).

The first term on the RHS gives the time evolution of λ_ν at fixed Γ_{HI} and the second gives the instantaneous dependence on Γ_{HI} . We interpolate the first term from our simulation suite and assume $\lambda_\nu \propto \Gamma_{\text{HI}}^{2/3}$ to calculate the second (motivated by theoretic expectations for an ionized IGM, e.g. Furlanetto & Oh 2005). The last term accounts for the evolution of λ_ν towards its constant- Γ_{HI} limit,⁵ $\lambda_{\nu,0}$, which we also interpolate from our simulations. We assume $t_{\text{relax}} = 100$ Myr for the timescale over which the IGM loses memory of its previous Γ_{HI} history. Further details can be found in the text and appendices of Cain et al. (2023).

In this work, we use two versions of our sub-grid model, both of which are described in Cain et al. (2023):

(i) **Full sinks:** This model uses the full, self-consistent evolution of λ_ν predicted by the our small-volume simulations suite combined with equation (7) and (10). It accounts for the pressure smoothing of the IGM after ionization and the effect of un-resolved self-shielded systems.

(ii) **Relaxed limit:** For this model, we extrapolate the low-redshift λ_ν in our $z_{\text{reion}} = 12$ sub-grid simulations to lower redshift assuming a power law. This treats the IGM as if it has been ionized for a long time, and is in a pressure smoothed equilibrium, at all times. In this model, the effects of small, short-lived sinks that are sensitive to pressure smoothing and photoevaporation (see Section 2.3) are neglected.

3.4 Multifrequency RT

We have extended our sub-grid formalism to allow for a multifrequency RT treatment. Unfortunately, our sub-grid simulations only saved the MFP averaged over the ionizing spectrum, and the Lyman Continuum (LyC, 912 Å) MFP. To approximate the full frequency dependence, we assume

$$\kappa_\nu = \kappa_{912} \left(\frac{\sigma_{\text{HI}}(\nu)}{\sigma_{\text{HI}}^{912}} \right)^{\beta_N - 1} \quad (11)$$

where $\kappa_\nu \equiv 1/\lambda_\nu$ is the absorption coefficient at frequency ν and κ_{912} is the LyC absorption coefficient. This form follows from equation (5) of Nasir et al. (2021) (see also Prochaska, Worseck & O'Meara 2009) assuming that the H I column density distribution is a power law of the form $f(N_{\text{HI}}) \propto N_{\text{HI}}^{-\beta_N}$. Our sub-grid simulations assume a power law ionizing spectrum of the form $J_\nu \propto \nu^{-\alpha}$ with $\alpha = 1.5$. We can then find β_N using

$$\frac{\langle \kappa_\nu \rangle_{\alpha=1.5}}{\kappa_{912}} = \frac{\alpha[1 - 4^{-\alpha-2.75(\beta_N-1)}]}{[\alpha + 2.75(\beta_N - 1)](1 - 4^{-\alpha})}, \quad (12)$$

where $\langle \kappa_\nu \rangle_{\alpha=1.5}$ is the frequency-averaged absorption coefficient in the sub-grid simulations. Then we calculate κ_ν in each frequency bin in FlexRT using equation (11). Note that $\langle \sigma_{\text{HI}} \rangle_\nu^i$ and $\langle \lambda \rangle_\nu^i$ in equation (7) are averaged over the spectrum incident on cell i in FlexRT, not the $\alpha = 1.5$ spectrum used in the sub-grid simulations. This assumes that λ_ν can be estimated for any incident spectrum using results from an $\alpha = 1.5$ simulation. In FlexRT, we use five frequency bins with energies 14.44, 16.64, 19.91, 25.47, and 37.6 eV, each containing the same number of photons.⁶ We test this procedure and give more details in Appendix B. We find that our approach is accurate to within 20 per cent even for negative values of α .

⁵That is, the limit in which Γ_{HI} has not changed for a long time.

⁶This particular binning is chosen to give approximately the correct average H I cross-section for an $\alpha = 1.5$ spectrum. The are also approximately the same frequency bins used in the high-resolution simulations.

3.5 Recombination radiation

We have added an approximate treatment of recombination radiation to FlexRT. We assume that a fraction $f_{\text{esc}}^{\text{rec}}$ of IRPs escape from the dense clumps that produce them, and the rest are absorbed locally. The effective recombination coefficient can be written as

$$\alpha(T) = f_{\text{esc}}^{\text{rec}} \alpha_A(T) + (1 - f_{\text{esc}}^{\text{rec}}) \alpha_B(T), \quad (13)$$

where T is temperature and α_A and α_B are the case A and B recombination coefficients, respectively. We assume equation (13) holds when modelling the ionizing photon opacity. For Ly α forest calculations (described in Section 3.7), we use α_A for the reasons explained in Section 2.2.

We assume that the ionized gas in each cell is in photoionization equilibrium.⁷ Then, we can write the recombination rate as a function of the ionizing absorption coefficient (Emberson et al. 2013), with $\kappa \propto \alpha(T)$ at fixed density and Γ_{HI} . Our small-volume simulations assume the case B recombination coefficient, so we can re-scale the absorption coefficient from those simulations like

$$\langle \kappa \rangle_\nu = \langle \kappa_B \rangle_\nu \frac{\alpha(T)}{\alpha_B(T)}, \quad (14)$$

where $\langle \kappa_B \rangle_\nu$ is the sub-grid model prediction.

Under these approximations, the emissivity of IRPs in cell i is

$$\dot{n}_\gamma^{\text{rec},i} = x_{\text{ion}}^i f_{\text{esc}}^{\text{rec}} \frac{\Gamma_{\text{HI}}}{\langle \sigma_{\text{HI}} \rangle_\nu^i} (\langle \kappa_A \rangle_\nu - \langle \kappa_B \rangle_\nu) (1 + \chi), \quad (15)$$

where $\langle \kappa_A \rangle_\nu$ is equation (14) evaluated for $f_{\text{esc}} = 1$, and the factor of $1 + \chi$ crudely accounts for ground state recombinations from He II.⁸ These photons are added to the emissivity from haloes in each cell. The energy of IRPs is

$$E_{\text{rec}} = 13.6 \text{ eV} + \frac{1}{2} m_e v^2, \quad (16)$$

where m_p is the electron mass and v is the relative velocity between the recombining electron and proton. The average kinetic energy is $\approx \frac{3}{2} kT$, where T is the gas temperature. For 10^4 K gas, this is 1.3 eV, yielding $E_{\text{rec}} = 14.9$ eV and an average H I cross-section of $\langle \sigma_{\text{HI}} \rangle_\nu = 4.93 \times 10^{-18} \text{ cm}^2$, a factor of 2 larger than the $\langle \sigma_{\text{HI}} \rangle_\nu = 2.55 \times 10^{-18} \text{ cm}^2$ for an $\alpha = 1.5$ spectrum. In our multifrequency FlexRT runs, we assign IRPs to our lowest energy bin (14.44 eV). We provide a complete derivation of equations (13)–(15) in Appendix C.

3.6 Modelling haloes as absorbers

Lastly, we added a simple prescription for the opacity from haloes hosting ionizing sources. This allows us to assess how missing opacity from massive haloes would affect the drop in $\dot{N}_\gamma$ (see Section 2.3). We treat haloes as spherically symmetric, optically thick absorbers with cross-section

$$\sigma_{\text{halo}}(M_{\text{halo}}, z, \Gamma_{\text{HI}}) = \pi R_{\text{halo}}^2(M_{\text{halo}}, z, \Gamma_{\text{HI}}), \quad (17)$$

where R_{halo} is the radius to which the halo is opaque. We write this as

$$R_{\text{halo}}(M_{\text{halo}}, z, \Gamma_{\text{HI}}) = f_{200}(M_{\text{halo}}, z, \Gamma_{\text{HI}}) R_{200}(M_{\text{halo}}, z), \quad (18)$$

⁷This is not true in our small volume simulations, which include non-equilibrium effects such as photoevaporation of self-shielded systems. However, it is likely a good approximation in the relaxed limit sub-grid model, since such processes are largely neglected there.

⁸This is appropriate because the recombination coefficients for H II and He II are very similar in the relevant temperature range.

where R_{200} is the radius within which the mean density is $200 \times$ the cosmic mean, and f_{200} parametrizes the opacity of the halo (including its sensitivity to $\Gamma_{\text{H I}}$). Halo masses in our N -body simulation are given by the mass enclosed within R_{200} , which is approximately the halo virial radius (Trac et al. 2015). We will describe momentarily how we estimate f_{200} .

The total opacity from haloes in cell i is given by

$$\kappa_{\text{halo}}^i = \frac{1}{V_{\text{cell}}^i} \sum_{j=1}^{N_{\text{halo}}} \sigma_{\text{halo}}(M_{ij}, z), \quad (19)$$

where the sum runs over all haloes occupying cell i . In what follows, we assume that all haloes resolved by our N -body simulations (with $M_{\text{halo}} \geq 3 \times 10^9 h^{-1} M_{\odot}$) contribute extra opacity. The total opacity in cell i is given by adding κ_{halo} to equation (6)⁹:

$$\tau_{ij}^v = x_{\text{ion}}^i \Delta s^{ij} \left(\frac{1}{\lambda_i^v} + \kappa_{\text{halo}}^i \right). \quad (20)$$

Finally, we need a prescription for f_{200} , for which we use an analytical model similar to that described in section 2 of Muñoz et al. (2016) (see also Theuns 2021; Theuns & Chan 2024). We assume all haloes have a spherically symmetric power-law density profile

$$\rho_{\text{gas}}(r) \propto r^{-\epsilon}. \quad (21)$$

The radius r_{abs} to which the halo is opaque to ionizing photons can be approximated by the condition

$$\tau_{\text{ion}} \approx n_{\text{H I}}(r_{\text{abs}}) \langle \sigma_{\text{H I}} \rangle r_{\text{abs}} = 1, \quad (22)$$

where τ_{ion} is approximately the optical depth encountered by a sightline intersecting the edge of the halo.¹⁰ We show in Appendix D that, assuming photoionization equilibrium, f_{200} scales like

$$f_{200} \propto \Gamma_{\text{H I}}^{-\frac{1}{2\epsilon-1}} M_{\text{halo}}^{\frac{1}{3(2\epsilon-1)}} (1+z)^{\frac{5}{2\epsilon-1}}, \quad (23)$$

where $\Gamma_{\text{H I}}$ is the photoionization rate in the IGM surrounding the halo. The positive scaling with redshift reflects the fact that haloes collapsing at higher z have larger physical densities at their virial radii. More details are given in Appendix D.

3.7 Modelling the Ly α forest

We compute Ly α forest statistics on a $N = 2048^3$ version of the same hydro run used to get our density fields. We traced 4000 sightlines through the box and mapped on the $\Gamma_{\text{H I}}$, T , and neutral fraction from FlexRT. In ionized RT cells (defined as those with $x_{\text{H I}} < 0.5$), we compute the residual neutral fraction assuming photoionization equilibrium. The spatial resolution of our hydro simulation is $97.7 h^{-1} \text{ kpc}$, which D'Aloisio et al. (2018) found to be un-converged in

⁹Note that the extra opacity from haloes does not contribute to $\langle \lambda \rangle_v$ in the numerator of equation (7). This is because the expression $\langle \lambda \rangle_v \langle \sigma_{\text{H I}} \rangle_v$ in the numerator of that equation is equal to the inverse of the residual H I number density in the *ionized* IGM, but the haloes are treated here as if they are completely self-shielding (neutral). Also, note that κ_{halo} is also multiplied by x_{ion} because not doing so would spuriously count opacity from haloes still within the neutral part of the cell.

¹⁰The analogous derivation in section 2 of Muñoz et al. (2016) identified the absorber radius as the distance at which the opacity approaching the centre of the halo from infinity reaches 1. This condition yields the same scaling relations found here, but gives an absorber radius about 70 per cent of that given by equation (22). This is an acceptable difference given the highly approximate nature of the model.

the mean Ly α transmission at the 50 per cent level. We correct for this using the approach outlined in Appendix A of D'Aloisio et al. (2018), and the correction factors from their table A1. With their correction, our effective spatial resolution for the forest calculation is $12.2 h^{-1} \text{ kpc}$, close to the convergence criteria found by Dougherty et al. (2023).

One problem with mapping the FlexRT fields directly to the hydro density grids is that correlations with density on scales smaller than the RT cell size ($1 h^{-1} \text{ Mpc}$) are missed. This is a problem for temperature, which correlates strongly with density on small scales. Hydro cells less dense than the coarse-grained RT cells they occupy contribute most of the Ly α transmission, and in general these have temperatures lower than that of their host RT cell. The IGM temperature is usually assumed to follow a power law in density of the form

$$T(\Delta) = T_0 \Delta^{\gamma-1} \quad (24)$$

where Δ is gas density in units of the mean, T_0 is the temperature at mean density, and the power law index γ is mainly a function of how recently the gas was ionized. Recently ionized gas ($z \approx z_{\text{reion}}$) is nearly isothermal ($\gamma = 1$), while for $z < z_{\text{reion}}$, $\gamma \rightarrow 5/3$. We estimate $\gamma(z, z_{\text{reion}})$ using the analytical solution for the IGM temperature given by McQuinn & Upton Sanderbeck (2016). Then we re-scale the temperatures in each hydro cell according to

$$T_{\text{hydro}}^i = T_{\text{RT}}^j \left(\frac{\Delta_{\text{hydro}}^i}{\Delta_{\text{RT}}^j} \right)^{\gamma(z, z_{\text{reion}}^j)-1}, \quad (25)$$

where T_{RT}^j is the temperature in RT cell j that contains hydro cell i , Δ_{hydro}^i is the density in hydro cell i , Δ_{RT}^j is the coarse-grained density in RT cell j , and z_{reion}^j is the reionization redshift of cell j . We find that this procedure results in a ~ 10 per cent correction to the forest mean flux. We test this procedure and give more details in Appendix E.

4 EFFECT OF INDIVIDUAL MODELLING CONSIDERATIONS

4.1 Reference model

In the following sections, we will study the IGM modelling effects described in Section 2. Starting from a reference model (described below), we will vary each modelling choice one at a time and assess how this changes the evolution of $\dot{N}_{\gamma}$ required by the forest. In *all* our FlexRT simulations, $\dot{N}_{\gamma}$ is a non-parametric function of redshift tuned such that the mean Ly α transmission, $\langle F_{\text{Ly } \alpha} \rangle$, agrees as closely as possible with recent measurements by Bosman et al. (2022) using the new XQR-30 QSO sample (D'Odorico et al. 2023). This is similar to the approach adopted in (Keating et al. 2020a, b). Our reference model has the following features:

- (i) A monochromatic spectrum with a $\sigma_{\text{H I}} = 2.55 \times 10^{-18} \text{ cm}^2$, as in Cain et al. (2021, 2023). This $\sigma_{\text{H I}}$ is the average value for a power law spectrum with $\alpha = 1.5$.
- (ii) The relaxed limit sub-grid model, described in Section 3.3. We chose this model so that our reference case would be representative of simulations that do not resolve the small, abundant sinks that are sensitive to photoevaporation and pressure smoothing. We do not include any extra opacity from haloes.
- (iii) We use the case B assumption for the ionizing opacity that is native to our sub-grid simulations (i.e. $f_{\text{esc}}^{\text{rec}} = 0$ in equation 13).




Figure 2. Top two rows, in clockwise order from top left: ionized fraction x_{ion} , LyC MFP $\lambda_{912}^{\text{mfp}}$, ionizing emissivity $\dot{N}_\gamma$, IGM temperature at mean density T_0 , mean Ly α forest transmission $\langle F_{\text{Ly}\alpha} \rangle$, and electron scattering optical depth to the CMB (τ_{es}). We compare our reference model (black solid curve) to observations from (Becker et al. 2011; Becker & Bolton 2013; Worseck et al. 2014; McGreer et al. 2015; Greig et al. 2016; Bosman et al. 2018; Davies et al. 2018; Eilers et al. 2018; Mason et al. 2018, 2019; Boera et al. 2019; Greig, Mesinger & Bañados 2019; Walther et al. 2019; Gaikwad et al. 2020; Planck Collaboration VI 2020; Wang et al. 2020; Yang et al. 2020a, b; Becker et al. 2021; Bosman et al. 2022; Zhu et al. 2022; Gaikwad et al. 2023; Jin et al. 2023; Zhu et al. 2023). The upper right panel displays the same models shown in Fig. 1 for reference (faded curves). All the simulations in this work have their emissivity histories calibrated to match the $\langle F_{\text{Ly}\alpha} \rangle$ measurements from Bosman et al. (2022) as closely as possible over $4.8 < z < 6$. The bottom two rows show the CDF of τ_{eff} , computed over $50 h^{-1}\text{Mpc}$ segments of the Ly α forest, compared to the measurements of Bosman et al. (2018, 2022) and Eilers et al. (2018). See text for details.

Fig. 2 compares our reference model with a wide range of observations and other simulations. The top two rows show, clockwise from the top left, the ionized fraction x_{ion} , LyC MFP $\lambda_{912}^{\text{mfp}}$, $\dot{N}_\gamma$, IGM temperature at mean density T_0 , mean Ly α forest transmission $\langle F_{\text{Ly}\alpha} \rangle$, and electron scattering optical depth to the CMB, τ_{es} . We estimate $\lambda_{912}^{\text{mfp}}$ by stacking mock spectra calculated from random starting points and fitting to the model of Prochaska et al. (2009), as in Cain et al. (2021).¹¹ The upper right panel shows $\dot{N}_\gamma$ for the reference model (black solid) and the same curves from Fig. 1. The other panels show a collection of recent observational measurements from the literature, referenced in the figure caption. The red points in the second-from-top middle panel show the most recent measurements of $\langle F_{\text{Ly}\alpha} \rangle$ by Bosman et al. (2022) at $4.8 < z < 6$, against which we calibrate our models.

The bottom two rows show the cumulative distribution function (CDF) of effective Ly α optical depth, τ_{eff} , measured over $50 h^{-1}$ Mpc segments of the forest. Each panel shows a different redshift between $z = 5$ and 6 . We compare to measurements by Bosman et al. (2022) (red), Bosman et al. (2018) (green), and Eilers et al. (2018) (blue). For observations, shaded regions denote limiting assumptions about the value of τ_{eff} for null detections.¹² In each redshift bin, we randomly draw a number of $50 h^{-1}$ Mpc segments from our simulation sightlines equal to the corresponding number in the Bosman et al. (2022) data set, and compute the CDF 500 times. The shaded grey curves show the 10–90 per cent range from this procedure.

Our reference model is in broad agreement with the observations in Fig. 2. Our reionization history ends slightly too late for the dark pixel constraints from McGreer, Mesinger & D’Odorico (2015), but is in agreement with updated limits from Jin et al. (2023) and recent dark gap constraints from Zhu et al. (2022). The MFP is in reasonable agreement with measurements of Worseck et al. (2014) and Gaikwad et al. (2023), but is slightly too high for those of Becker et al. (2021) and Zhu et al. (2023). This could owe to the relaxed limit model having too few sinks, and/or to our use of the case B approximation for recombination radiation. Our $\dot{N}_\gamma$ is qualitatively similar to those discussed in Section 1, with a factor of ~ 2 decline between $z = 6$ and the end of the simulation at $z = 4.8$. Our thermal history is in agreement with the recent measurements of Gaikwad et al. (2020), but is slightly too high for measurements at $z \leq 5$.

Our τ_{eff} distributions are also in broad agreement with measurements. Note that our simulations are not explicitly calibrated to match the τ_{eff} distribution of Bosman et al. (2022), only $\langle F_{\text{Ly}\alpha} \rangle$. So, to some extent, our models can be judged by their agreement with τ_{eff} measurements. We see that at $5.2 < z < 5.8$, our τ_{eff} distributions are slightly wider than the Bosman et al. (2022) measurements. Bosman et al. (2022) found that their τ_{eff} distributions are in good agreement with simulations in which reionization is completely over by $z = 5.3$. The fact that our reference model ends reionization slightly later than this (~ 5.1) is likely the reason for this disagreement.

4.2 Ionizing spectrum

We start by varying the properties of the ionizing spectrum. Our first model, which we label ‘fiducial (Fid.) multifrequency’, adopts the

¹¹Roth et al. (2023) find that this estimator for the MFP agrees well with the measurements using QSOs, even at $z = 6$ when reionization is ongoing (see also Satyavolu et al. 2023).

¹²The lower bound assumes $\tau_{\text{eff}} = \infty$, while the upper bound assumes $\tau_{\text{eff}} = -\ln(2\sigma)$, where σ is the uncertainty of the mean transmission.




Figure 3. Comparison of our Fid. Multifrequency (orange dashed) and hard multifrequency (blue dotted) models to our reference model (solid black). From top to bottom, the panels show $\dot{N}_\gamma$, $\langle F_{\text{Ly}\alpha} \rangle$, and T_0 . All models are calibrated to match the measurements of $\langle F_{\text{Ly}\alpha} \rangle$ by Bosman et al. (2022). Harder spectra work in the direction of making the drop in $\dot{N}_\gamma$ shallower, but the effect is small compared to the size of the drop itself. The hard multifrequency model is also in significant tension with most recent IGM temperature measurements.

multifrequency treatment described in Section 3.4, still assuming $\alpha = 1.5$ for sources. Comparing this model to the reference model (which employs the monochromatic approximation) isolates the effects of IGM filtering.¹³ Our second model, labelled ‘hard multifrequency’, also uses multifrequency RT but assumes $\alpha = 0.5$, on the low end for reasonable stellar population models (see Section 2.1).

Fig. 3 shows $\dot{N}_\gamma$ (top), $\langle F_{\text{Ly}\alpha} \rangle$ (middle), and T_0 (bottom) for these models (orange dashed and blue dotted, respectively) compared to the reference model (black solid). By construction, all three models have nearly identical $\langle F_{\text{Ly}\alpha} \rangle$, in agreement with Bosman et al. (2022). All three have fairly similar $\dot{N}_\gamma$ at $z < 6$, with the Fid. Multifrequency and hard multifrequency models having a slightly shallower drop than the reference case. The effects of multifrequency RT, including a harder spectrum, on $\dot{N}_\gamma$ are about 0.05 dex, much smaller than the ~ 0.3 dex drop in the reference model. The smallness of the effect likely owes to the partial cancellation of $\Gamma_{\text{H I}}$ and T effects in the $\langle F_{\text{Ly}\alpha} \rangle$, discussed in Section 2.1. Recall that a harder spectrum leads to lower $\Gamma_{\text{H I}}$ and higher temperatures, which have opposite effects on the forest transmission.

¹³To maximize the effect of filtering, we assume $\beta_{\text{N}} = 2$ in equation (11) instead of calculating it self-consistently using equation (12) in this section.




Figure 4. Effect of IRPs on the drop in $\dot{N}_\gamma$. Top: $\dot{N}_\gamma$ for our recombination model (orange dashed) and reference model (black solid). Bottom: $\langle F_{\text{Ly}\alpha} \rangle$ compared to Bosman et al. (2022) measurements. Ionizing recombination radiation has a small effect on the drop, even for parameters that maximize its effect. See text for details.

We see from the bottom panel that these models also have significantly higher IGM temperatures. T_0 is ~ 10 per cent (20 per cent) higher in the Fid. Multifrequency (hard multifrequency) model than in the reference case. The hard multifrequency case is in $\sim 2\sigma$ tension with the Gaikwad et al. (2023) measurements at $z = 5.4$ and 5.6 , and in $\sim 3\sigma$ tension with the Boera et al. (2019) measurement at $z = 5$. Even harder spectra, which may allow for a shallower drop, would only worsen this disagreement. We conclude that although a harder ionizing spectrum works in the direction of making the drop shallower, on its own it is an unlikely explanation.

4.3 Ionizing recombination radiation

Next, we study the effect of ionizing recombination radiation. Our ‘recombination radiation’ model assumes $f_{\text{esc}}^{\text{rec}} = 1$, which corresponds to the case A limit and maximally contrasts the reference model. To facilitate a fair comparison with our monochromatic reference model, here we only include two frequency bins: $h_p\nu = 19.0$ eV for sources (which gives the same $\sigma_{\text{H I}}$ assumed in the reference model) and 13.6 eV for IRPs. The choice of 13.6 eV maximizes the ‘softening’ effect of IRPs on the ionizing spectrum.¹⁴ Fig. 4 shows $\dot{N}_\gamma$ (top) and $\langle F_{\text{Ly}\alpha} \rangle$ (bottom) for our Recombination Radiation model (orange dashed) and our reference model (black solid). As in Fig. 3, the differences in $\dot{N}_\gamma$ between models are modest. The recombination radiation model has a slightly flatter $\dot{N}_\gamma$ at $z = 6$, with a slightly steeper drop between $z = 5.5$ and 4.8 . Given that the model in Fig. 4 likely overestimates the effect of recombination radiation, we conclude that IRPs probably have a small effect on $\dot{N}_\gamma$ at fixed $\langle F_{\text{Ly}\alpha} \rangle$.

However, IRPs do have an appreciable effect on the reionization history and the τ_{eff} distribution. Fig. 5 shows the $z < 6$ reionization

history (left) and τ_{eff} CDFs at $z = 5.2$ and 5.4 (middle and right, respectively). We also show our hard multifrequency model (blue dotted curve and shaded regions) from the previous section. The recombination radiation (hard multifrequency) model ends reionization later (earlier) than the reference case, and displays worse (better) agreement with τ_{eff} measurements. The former is in mild tension with the $z = 5.5$ neutral fraction constraint of Zhu et al. (2022). In the hard multifrequency model, at fixed $\langle F_{\text{Ly}\alpha} \rangle$, $\dot{N}_\gamma$ is higher and reionization ends earlier, while the opposite is true in the recombination radiation model.

The main reason for later reionization history in the recombination radiation model is the soft ionizing spectrum of the IRPs. Since the spectrum is softer in that model than the reference case, $\Gamma_{\text{H I}}$ and $\langle F_{\text{Ly}\alpha} \rangle$ are higher at fixed $\dot{N}_\gamma$, and so reionization ends later at fixed $\langle F_{\text{Ly}\alpha} \rangle$. We have checked that if the IRPs are assigned the same frequency as the sources, the differences with the reference model largely disappear. A key takeaway from Fig. 5 is that harder (softer) ionizing spectra can lead to earlier (later) reionization and better (worse) agreement with measurements of the τ_{eff} distribution at fixed $\langle F_{\text{Ly}\alpha} \rangle$.

4.4 Missing sinks

Ionizing photon sinks regulate the ionizing background, making them promising candidates to explain the drop in $\dot{N}_\gamma$. In this section, we consider whether a population of sinks that are missing from simulations could explain the drop.

4.4.1 Missing unresolved sinks

One reason for sinks to be missing in a simulation is that they are unresolved. A natural way for us to assess the effect of the smallest sinks is to compare our relaxed limit and full sinks sub-grid models, described in Section 3.3. The latter includes the effects of these objects on the IGM opacity, while the former is designed to neglect them. The orange dashed curve in Fig. 6 shows x_{ion} (left), $\dot{N}_\gamma$ (middle), and $\langle F_{\text{Ly}\alpha} \rangle$ for a simulation that uses the full sinks treatment, but is otherwise the same as our reference case (black solid curve). We emphasize again that this comparison is made at fixed $\langle F_{\text{Ly}\alpha} \rangle$, as the right panel shows.

A surprising result is that including the smallest sinks does not alleviate the need for a drop – in fact, the drop is ~ 0.05 dex steeper than in the reference case. This is because the smallest sinks are destroyed by photoevaporation and pressures smoothing effects within a few hundred Myr, and only patches of the IGM that ionized recently still have them. So, although they increase the photon budget (reflected by a later end to reionization and higher $\dot{N}_\gamma$), they do not survive long enough to help regulate $\Gamma_{\text{H I}}$ after reionization ends. The result is that an even steeper drop in $\dot{N}_\gamma$ is required. For this reason, attempts to correct the LyC opacity in simulations for missing small-scale power are unlikely to eliminate the need for a drop.

4.4.2 Missing massive sinks

It is also possible that simulations are underpredicting the opacity from massive haloes. We have run a simulation that uses the halo absorber model described in Section 3.6, with all else the same as in the reference case. We assume fairly shallow density profiles for our haloes, with $\epsilon = 1.5$ (compared to 2 for isothermal profiles). This choice is motivated by the possibility that feedback drives significant outflows (e.g. Weldon et al. 2022), possibly resulting in significant

¹⁴IRPs do not have a net heating or cooling effect, so this choice does not affect the IGM temperature.




Figure 5. Sensitivity of the reionization history and τ_{eff} distribution to the ionizing spectrum and recombination radiation. From left to right, the panels show the reionization history and τ_{eff} distribution at $z = 5.2$ and 5.4 . The main effect of recombination radiation is to soften the ionizing spectrum, so we can think of this model as a ‘soft spectrum’ scenario. Models with harder spectra end reionization earlier when calibrated to the same $\langle F_{\text{Ly}\alpha} \rangle$ evolution, in better agreement with constraints on the neutral fraction at $z < 6$. For this reason, they have narrower τ_{eff} distributions, in better agreement with recent measurements.




Figure 6. Effect of sinks on the drop in $\dot{N}_\gamma$. We show x_{ion} , $\dot{N}_\gamma$, and $\langle F_{\text{Ly}\alpha} \rangle$ (left to right). We compare our reference model (black solid) to a model that includes the smallest, hardest-to-resolve sinks (orange dashed) and a model with extra opacity from massive, star-forming haloes (blue dotted). The former requires a higher photon budget and ends reionization later, but has an ~ 0.05 dex steeper drop than the reference case. This is because the smallest sinks do not survive very long and cannot help regulate the ionizing background at reionization’s end. Sinks in massive haloes are not destroyed by photoevaporation and pressure smoothing, and rapidly grow in abundance at lower redshifts. These sinks effectively regulate Γ_{HI} , reducing the drop to only ~ 0.05 dex.

opacity outside the halo virial radius ($f_{200} > 1$). The blue dotted curve in Fig. 6 shows results for our ‘massive sinks’ model. Unlike our full sinks model, our halo absorber model does make the drop in $\dot{N}_\gamma$ shallower. Moreover, $\dot{N}_\gamma$ steadily rises between $z = 6$ and 8 instead of staying flat. This is because the abundance of $M \geq 3 \times 10^9 h^{-1} M_\odot$ haloes grows rapidly throughout reionization, increasing by a factor of 2.6 (3.7) from $z = 8$ to 6 (5).¹⁵ Because these sinks become more abundant rapidly, they can effectively regulate the growth of Γ_{HI} and avoid the need for a drop in $\dot{N}_\gamma$. Moreover, these sinks are not evaporated by reionization, so they survive long enough to regulate the post-reionization ionizing background.

There are a few caveats worth mentioning. First, our treatment of haloes as spherically symmetric, ‘billiard-ball’ absorbers with a single density profile is overly simplistic. Real galaxies have much more complex geometries and dynamics, and as such our model may not capture their opacity very well. Secondly, the choice of including extra opacity for $M_{\text{halo}} > 3 \times 10^9 h^{-1} M_\odot$ is arbitrary, being set by the completeness limit of our N -body simulation. Lower mass cutoffs

would result in more opacity from haloes, but would also cause κ_{halo} to grow less quickly (or perhaps decline) with cosmic time (see Fig. D1 and surrounding discussion). This would also increase the budget of ionizing photons required to complete reionization. Finally, our choice of $\epsilon = 1.5$ may not be realistic – a smaller value for ϵ results in less opacity from haloes, all else being the same. A key ‘performance test’ for models like this will be to see if their MFP is consistent with measurements at $5 < z < 6$ (Gaikwad et al. 2023; Zhu et al. 2023; Davies et al. 2024). We provide such a test in Section 5.1 and will further address this point in a forthcoming paper.

5 IMPLICATIONS FOR THE PROPERTIES OF SOURCES

5.1 Realistic and optimistic scenarios

In this section, we combine the effects studied in Section 4. Table 1 lists the physical effects we studied and describes how each changes the drop in $\dot{N}_\gamma$ and the reionization history when $\langle F_{\text{Ly}\alpha} \rangle$ is held fixed. Green text denotes a shallower drop in $\dot{N}_\gamma$ or an earlier reionization history, and red the opposite. We consider two models in this section: a ‘realistic’ model and an ‘optimistic’, shown in Fig. 7 by the orange

¹⁵The mass contained in these haloes grows by a factor of 4.1 (7.8) between $z = 8$ and 6 (5).

Table 1. Summary of the modelling considerations studied in Section 4 and their effects on $\dot{N}_\gamma$ and the reionization history. The second and third columns describe their effect on the drop in $\dot{N}_\gamma$ and the reionization history.

Physical effect	Emissivity drop	End of reionization
Harder ionizing spectrum	< 10 per cent difference	Earlier
Recombination radiation	< 10 per cent difference	Later
Missing small sinks	Steeper	Later
Missing massive sinks	Shallower	$\Delta z < 0.1$

dashed and blue dotted curves, respectively. The top row has the same format as Fig. 6, while the bottom row shows the τ_{eff} CDF at $z = 5.2, 5.4$, and 5.6 .

The optimistic model is a combination of our massive sinks and hard multifrequency model, and ignores the effect of un-resolved sinks and recombination radiation. This model has almost no drop and an early end to reionization, the latter of which results in good agreement with the measured τ_{eff} distribution. However, even this model would require some decline in the ionizing output of galaxies during reionization. The Robertson et al. (2015) emissivity, which assumes constant f_{esc} and ξ_{ion} , increases by ~ 0.4 dex between $z = 8$ and 5 (Fig. 1). So, $\langle f_{\text{esc}} \xi_{\text{ion}} \rangle$ would have to fall by a similar factor over this redshift range to produce the Optimistic $\dot{N}_\gamma$. We note that the Optimistic model has a higher ionizing budget – 3.2 photons per H atom – than the reference model (2.2 photons/H atom).

Our realistic model includes reasonable prescriptions for all the effects in 1. We include multifrequency RT as described in Section 3.4, but keep our fiducial $\alpha = 1.5$ spectrum. We assume $f_{\text{esc}}^{\text{rec}} = 0.5$ for recombination radiation, in between case A and B. We use our full sinks sub-grid model to account for the smallest sinks, and include our halo absorber model, but with $\epsilon = 2$ (isothermal density profiles). The last choice reduces the extra opacity from haloes (see Fig. D1 of Appendix D). This model has an $\dot{N}_\gamma$ between those of the reference and optimistic models – flat until $z = 5.5$, then dropping ~ 0.15 dex by $z = 4.8$, suggesting that realistic IGM models will likely require a drop. However, the details of the drop are sensitive to IGM modelling effects. This model also ends reionization later than the reference model and displays worse agreement with the τ_{eff} CDF (similar to that of the recombination radiation model in Section 4.3). This owes to the inclusion of recombination radiation and small sinks (see Figs 4 and 6). This may hint that recombination radiation and/or small sinks play a smaller role in reality than they do in our Realistic model – although see the next section for an alternative solution.

As mentioned in Section 4.4.2, an important check for models that resolve the drop with extra sinks is that they agree with measurements of the MFP. Fig. 8 shows the MFP for our reference, optimistic, and realistic models. We see that the optimistic model has a lower MFP than the reference case, but that it actually agrees better with $z < 5.5$ measurements. This shows that it is possible to alleviate the emissivity drop with massive sinks without undershooting MFP measurements. However, the realistic model undershoots measurements at $z \leq 5.1$, due to its inclusion of recombination radiation and the smallest sinks.¹⁶ This result highlights the fact that the relationship between the MFP and the forest is sensitive to the effects studied in this work.

An important caveat is that we have not exhaustively treated everything that could affect the $\dot{N}_\gamma$ –forest relationship. For example,

¹⁶The former increases the opacity of the gas by a factor of $\sim f_{\text{esc}}^{\text{rec}} \frac{\alpha_A - \alpha_B}{\alpha_B} + 1$ (equation 14), but does so without affecting the ionizing photon budget, since IRPs compensate for the extra absorption (equation 15)

the slope of the H I column density distribution, β_N (equation 12), is not a free parameter in our model, being inferred from our sub-grid simulations. Treating this as a free parameter would widen the range of possible sinks models. We also have not fully explored the parameter space of our IGM prescriptions. For example, our halo density profile parameter ϵ could have time-dependence, and so could $f_{\text{esc}}^{\text{rec}}$ or the spectral index α . We will study these points in more detail in a forthcoming paper.

5.2 Clustering of the ionizing sources

The last physical effect we will study is that of the clustering of ionizing sources (see Section 2.4). We will consider three sources models (all assuming our realistic IGM model):

(i) Fiducial sources: This model, which we have assumed so far, has $\beta = 1$ (that is, $\dot{n}_\gamma \propto L_{\text{UV}}$), and includes haloes down to the completeness limit of our N -body simulation, $M_{\text{halo}} > 3 \times 10^9 h^{-1} M_\odot$.

(ii) Democratic sources: We assume $\beta = 0$, which assigns all haloes the same ionizing emissivity, independent of L_{UV} . We also include all the haloes in the N -body simulation even below the completeness limit.¹⁷ This model maximizes the contribution of the faintest, least clustered ionizing sources to the ionizing budget.

(iii) Oligarchic sources: Our last model assumes $\beta = 1$ and that only $M_{\text{halo}} > 2 \times 10^{10} h^{-1} M_\odot$ haloes host ionizing sources, such that only the rarest and most clustered sources contribute to the ionizing budget. This model is motivated by the bright-galaxy driven models proposed by Naidu et al. (2020) and Matthee et al. (2022) and contrasts the democratic model.

Fig. 9 shows our source models in the same format as Fig. 11. These models have significantly different reionization histories, with the democratic (oligarchic) sources model ending reionization $\Delta z \sim 0.2$ earlier ($\Delta z \sim 0.15$ later) than the fiducial case. The democratic sources model agrees much better with the τ_{eff} CDF than the fiducial or oligarchic models. This is because during reionization, this model is less transmissive in biased regions close to sources than the other models at fixed neutral fraction. So, reionization is allowed to end sooner without overshooting measurements. However, this model has a larger drop in $\dot{N}_\gamma$ than the fiducial sources case (~ 0.25 versus ~ 0.15 dex). The earlier end to reionization in this case leads to a faster buildup of photons in the IGM by $z \sim 5$, and compensating for this requires a larger drop in $\dot{N}_\gamma$.

The oligarchic model is in significant tension with the τ_{eff} CDF and some of the neutral fraction constraints from dark gaps/pixels. Up to $z = 6.5$, $\dot{N}_\gamma$ is flat, then increases by 0.1 dex between $z = 6.5$ and 5.5 and then rapidly decreases by 0.2 dex after 5.5 . This odd behaviour is necessary to match $\langle F_a \rangle$ measurements. Taken together with the very high neutral fraction, this ‘artificial’ behaviour of $\dot{N}_\gamma$ seems to suggest that the forest disfavors the oligarchic sources scenario. However, we caution there is at least partial degeneracy between the effects of the source model and the assumed IGM parameters.

¹⁷The smallest halo identified in our N -body simulation has $M_{\text{halo}} = 7 \times 10^8 h^{-1} M_\odot$. Haloes below the completeness limit sometimes appear and disappear in subsequent snapshots. This behaviour is unphysical, though it may mimic to some extent the bursty nature of star formation and associated ionizing photon production in low-mass haloes (Emami et al. 2019; Dome et al. 2024). Since our goal is to qualitatively assess the effect of including fainter sources, so this is not a significant concern.




Figure 7. Comparison of optimistic (blue dotted) and realistic (orange dashed) models. The top row has the same format as Fig. 6, while the row shows the τ_{eff} distribution at $z = 5.2, 5.4$, and 5.6 . The former includes only effects that make the drop in $\dot{N}_\gamma$ smaller and/or the reionization history end earlier. The realistic model includes ‘intermediate’ treatments of the effects discussed in Section 4.2–4.4. The Optimistic model has almost no drop in $\dot{N}_\gamma$ and ends reionization early enough to agree well with the measured τ_{eff} CDF. However, even a flat $\dot{N}_\gamma$ implies significant evolution of source properties. Our realistic model requires some drop and ends reionization later than suggested by the τ_{eff} CDF.




Figure 8. Ionizing photon MFP for the reference, optimistic, and realistic models. The optimistic model has a shorter MFP than the reference case, and is actually in slightly better agreement with measurements at $z \lesssim 5.5$. The MFP in the realistic model is too short for measurements, mainly due to recombination radiation and the inclusion of the smallest sinks.

6 NUMERICAL EFFECTS

In this section, we will briefly study two numerical effects that affect the $\dot{N}_\gamma$ –forest relationship – the reduced speed of light approximation (RSLA) and the numerical convergence of the Ly α forest itself.

6.1 Reduced speed of light approximation

A self-consistent numerical treatment of reionization requires solving the RT equation. The maximum time-step required to resolve the transport of ionizing radiation in the Eulerian frame is the light-crossing time of the spatial resolution element, $\Delta x_{\text{cell}}/c$, where Δx_{cell} is the RT cell size and c is the speed of light. This time-scale is usually much shorter than any other time-scales in the problem. Because of this, a commonly used time-saving approach is to decrease the speed of light in the simulation, which increases the time-step without compromising the accuracy. This is called the RSLA, and it has been used in a number of reionization studies (e.g. Katz et al. 2017; D’Aloisio et al. 2020; Kannan et al. 2022). However, the RSLA can cause inaccuracies in large-scale reionization simulations. Using moment-based M1 RT,¹⁸ Ocvirk et al. (2019) demonstrated that the RSLA can lead to a significant underestimate of the ionizing background near the end of reionization, even if the reionization history is converged¹⁹

Fig. 10 compares two models to the reference case (in the same format as Fig. 4). The first (the ‘RSLA’ model, orange dashed) uses $\tilde{c} = 0.2$ (as in the recent THESAN simulations; Kannan et al. 2022) and has $\dot{N}_\gamma$ re-calibrated to match $\langle F_{\text{Ly}\alpha} \rangle$ measurements. The RSLA

¹⁸Wu, McQuinn & Eisenstein (2021) showed that moment-based RT may itself lead to inaccuracies in the post-reionization ionizing background. However, based on the arguments in this section, we believe that their result should hold for all RT methods.

¹⁹A related effect is that the RSLA leads to underestimates of the speed of I-fronts. This is studied in Deparis et al. (2019).




Figure 9. Effect of source clustering on the $\dot{N}_\gamma$ -forest relationship, in the same format as Fig. 11. The orange dashed and blue dotted lines compare our democratic and oligarchic sources models to our fiducial sources model (green solid). All three assume the realistic IGM model from the previous section. The democratic (oligarchic) models end reionization earlier (later) than the fiducial case and agree better (worse) with τ_{eff} measurements. The latter is in some tension with recent neutral fraction constraints from dark gaps and dark pixels. The democratic model requires a greater drop in $\dot{N}_\gamma$, while $\dot{N}_\gamma$ for the oligarchic model at $5 < z < 6$ is non-monotonic and seems ‘artificial’.




Figure 10. Effect of the RSLA on the $\dot{N}_\gamma$ - $\langle F_{\text{Ly}\alpha} \rangle$ relationship. The orange dashed curve shows a model calibrated to match the forest, but with $\tilde{c} = 0.2c$. This model requires no drop in $\dot{N}_\gamma$. The blue dotted curve assumes the same $\dot{N}_\gamma$ but sets $\tilde{c} = 1$, causing $\langle F_{\text{Ly}\alpha} \rangle$ to overshoot measurements considerably.

model has a similar reionization history, thermal history, and MFP as the reference case. Strikingly, it has no drop in $\dot{N}_\gamma$. The blue dotted curve shows what happens when we use the $\dot{N}_\gamma$ from the RSLA model, but set $\tilde{c} = 1$. This case recovers the expected result – an overshoot of $\langle F_{\text{Ly}\alpha} \rangle$ by a factor of 2.5–6. We see that using the RSLA can spuriously lead to the conclusion that no drop is needed.

The reason for this is as follows. Early in reionization, $\Gamma_{\text{H I}}$ in ionized bubbles is independent of $\tilde{c}$. When the IGM is still a collection of isolated ionized bubbles, the photon density N_γ in these bubbles is overestimated by a factor of $c/\tilde{c}$. Since $\Gamma_{\text{H I}} \propto N_\gamma c$ in ionized gas, $\Gamma_{\text{H I}}$ is independent of $\tilde{c}$. Later on, as ionized bubbles overlap, N_γ becomes similar for simulations with different $\tilde{c}$, and $\Gamma_{\text{H I}}$ is underestimated by a factor of $\approx \tilde{c}/c$. This behaviour is illustrated in fig. 3 of Ocvirk et al. (2019). The transition between these regimes happens in the last half of reionization, when $\Gamma_{\text{H I}}$ is increasing, such that the RSLA artificially slows down the growth of $\Gamma_{\text{H I}}$.

In our RSLA model, the slower growth of $\Gamma_{\text{H I}}$ demands a larger $\dot{N}_\gamma$ to reproduce the same Ly α forest properties as the reference case. As a result, instead of declining at $z < 6.5$, $\dot{N}_\gamma$ continues to increase. This is a possible explanation for why the THESAN simulations were able to reproduce forest observations reasonably well without any drop in $\dot{N}_\gamma$. However, it is also possible that their treatment of galaxies played a role in reducing the drop, as illustrated in Sections 4.4.2 and 5.1. They also use a different RT method than ours (moment-based versus ray tracing). It is therefore difficult to assess exactly how large this effect is in THESAN, although Fig. 10 suggests that it is likely significant.




Figure 11. Effect of an under-resolved $\text{Ly}\alpha$ forest. The orange dashed curve shows a model where we remove the spatial resolution correction applied to our forest calculations and re-calibrate to match measurements. The top row shares the same format as Fig. 6, and the bottom row shows the τ_{eff} at $z = 5.2, 5.4$, and 5.6 (left to right). This model ends reionization earlier, has a slightly steeper drop in $\dot{N}_\gamma$, and agrees better with the measured τ_{eff} CDF than does the reference model. This shows that in self-consistent simulations, an un-resolved $\text{Ly}\alpha$ forest can bias conclusions about reionization.

6.2 Numerical convergence of the $\text{Ly}\alpha$ forest

Convergence of the forest itself could also affect the $\dot{N}_\gamma$ –forest relationship. At $5 < z < 6$, even the mean density IGM is mostly opaque to $\text{Ly}\alpha$ photons, so the mean transmission is set underdense gas. Capturing the distribution of underdensities to which the forest is sensitive requires high spatial resolution. It also requires a fairly large volume to include a representative distribution of large-scale voids, where underdense gas is more common. Simulations that fail to meet one or both of these requirements underestimate the mean transmission of the forest. Recently, Doughty et al. (2023) showed that box sizes of $\geq 20 h^{-1} \text{ Mpc}$ and spatial resolution of $\leq 10 h^{-1} \text{ kpc}$ are necessary to converge on the mean transmission of the forest in uniform grid simulations. The requirements are likely more stringent for smoothed particle hydrodynamics/adaptive mesh simulations, which have worse resolution in underdensities (Bolton & Becker 2009).

To assess the effect of forest convergence, we calibrated a model like our reference case, but without the spatial resolution correction for the $\text{Ly}\alpha$ forest described in Section 3.7. This reduced our effective cell size in the forest calculation to the hydro resolution, $\Delta x = 97.6 h^{-1} \text{ kpc}$. Fig. 11 compares this ‘unconverged forest’ model to our reference case, in the same format as Fig. 7. At fixed $\langle F_{\text{Ly}\alpha} \rangle$, the unconverged model ends reionization $\Delta z \sim 0.2$ earlier than the reference case, displays better overall agreement with τ_{eff} measurements, and has a slightly steeper drop in $\dot{N}_\gamma$ (by ~ 0.05 dex). This is because $\langle F_{\text{Ly}\alpha} \rangle$ is underestimated in this model, allowing for a larger $\Gamma_{\text{H I}}$ at fixed $\langle F_{\text{Ly}\alpha} \rangle$. This in turn allows $\dot{N}_\gamma$ to be higher and reionization to end sooner.

This result, like the previous one, is a cautionary tale for modelling the $\text{Ly}\alpha$ forest. Simulations that are calibrated to match measurements of the mean forest transmission may come to incorrect conclusions if they are under-resolved in the forest. In our case, this caused accidentally good agreement with τ_{eff} measurements because the reionization history was allowed to end earlier than it did when the forest was converged. We caution that this is only an issue for simulations that self-consistently model the interplay between $\dot{N}_\gamma$, the reionization history, and $\langle F_{\text{Ly}\alpha} \rangle$. We emphasize that these differences arise primarily because the neutral fractions in the models are different. At fixed neutral fraction and $\langle F_{\text{Ly}\alpha} \rangle$, we find little difference in the shape of the τ_{eff} distribution between models.

7 CONCLUSIONS

In this work, we have studied the implications of $\text{Ly}\alpha$ forest measurements at $5 < z < 6$ for the evolution of the ionizing emissivity at the end of reionization. Several recent works have found that matching these measurements requires a drop in emissivity in this redshift range, requiring strong evolution in the ionizing properties of galaxies. This work investigated the possibility that this drop is an artefact of inadequate modelling of the IGM in reionization simulations. We have done this by comparing reionization simulations run with our radiative transfer code, FlexRT (Cain et al. 2021, 2023), with different IGM modelling assumptions. Our main conclusions can be summarized as follows:

- (i) At fixed forest transmission, $\dot{N}_\gamma$ is fairly insensitive to the hardness of the ionizing spectrum. Models with hard ionizing spectra ($\alpha \leq 0.5$) result in slightly shallower (by ~ 0.05 dex)

drops and earlier reionization histories, the latter yielding better agreement with the distribution of Ly α opacities (τ_{eff}). However, such models also yield higher IGM temperatures than suggested by measurements.

(ii) Accounting for ionizing recombination radiation also has a small effect on the behaviour of $\dot{N}_\gamma$. However, the ‘spectral softening’ of the radiation field by recombination photons requires a later reionization history to match the forest transmission, resulting in worse agreement with the τ_{eff} distribution.

(iii) Accounting for the effect of the small, hard-to-resolve ionizing photon sinks on the IGM opacity requires a slightly steeper drop in $\dot{N}_\gamma$ and pushes the end of reionization later. These sinks increase the reionization photon budget significantly, but they are destroyed on a time-scale of a few hundred Myr after ionization. Because these sinks do not survive long after reionization, they do not alleviate the need for a drop in $\dot{N}_\gamma$.

(iv) Unlike the small sinks, missing opacity from massive sinks may be able to explain the drop. These structures are large enough to self-shield against the ionizing background, allowing them to survive past the end of reionization. They also become much more numerous at lower redshifts. So, they are able to regulate the ionizing background and in some scenarios eliminate the need for a drop.

(v) Under optimistic assumptions about the properties of the IGM, agreement with the forest transmission can be achieved without a drop in $\dot{N}_\gamma$. In this scenario, massive sinks contribute significantly to the IGM opacity, the ionizing spectrum is very hard, and the smallest sinks play a minimal role. However, even models like this require $\dot{N}_\gamma$ to grow less steeply than the UVLF, suggesting some evolution in underlying source properties. Models that include both small and massive sinks and a more reasonable ionizing spectrum still require some drop.

(vi) Scenarios where reionization is driven by lower-mass, less clustered ionizing sources yield earlier reionization histories (in better agreement with τ_{eff}) but also a steeper drop in $\dot{N}_\gamma$. In models driven by bright sources, reionization may have to end later than recent measurements suggest to match forest transmission measurements.

(vii) The RSLA artificially blunts the growth of the ionizing background at the end of reionization. This can erroneously lead to the conclusion that a drop in $\dot{N}_\gamma$ is not required to match forest observations.

(viii) Insufficient spatial resolution in numerical simulations can lead to an underestimation of the Ly α forest mean transmission. This can result in reionization ending earlier than it would otherwise. In some cases, this can cause spuriously good agreement with the measured τ_{eff} distribution.

We are led to conclude that either (1) the forest demands at least some evolution in the ionizing properties of sources near reionization’s end or (2) there are other effects setting the relationship between $\dot{N}_\gamma$ and the forest transmission that we have not considered here. The first possibility would have important implications for models of the high- z galaxy population. *JWST* has already begun to discover galaxies at $z \geq 6$ that may have exceptionally high ionizing efficiency (Atek et al. 2024; Cameron et al. 2023), hinting at such evolution. Improved modelling of the IGM will be key to helping to understand the implications of observations like these over the next few years. Further work should also explore in more detail the possible mechanisms that may drive an evolution in source properties, which we do not address directly in this work.

ACKNOWLEDGEMENTS

CC acknowledges helpful conversations with Yongda Zhu, George Becker, Sarah Bosman, Bayu Wilson, and Pierre Ocvirk. CC also thanks Pierre Ocvirk for inspiring the title. AD’s group was supported by NASA 19-ATP19-0191, NSF AST-2045600, and JWST-AR-02608.001-A. CC was supported by a combination of the above and the Beus Center for Cosmic Foundations while this work was ongoing. All computations were made possible by NSF XSEDE allocations TG-PHY210041 and TG-PHY230063, and by the NASA HEC Program through the NAS Division at Ames Research Center.

DATA AVAILABILITY

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

REFERENCES

- Abel T., Wandelt B. D., 2002, *MNRAS*, 330, L53
 Adams N. J. et al., 2024, *ApJ*, 965, 169
 Altay G., Theuns T., Schaye J., Crighton N. H. M., Dalla Vecchia C., 2011, *ApJ*, 737, L37
 Atek H., Furtak L. J., Oesch P., van Dokkum P., Reddy N., Contini T., Illingworth G., Wilkins S., 2022, *MNRAS*, 511, 4464
 Atek H. et al., 2024, *Nature*, 626, 975
 Becker G. D., Bolton J. S., 2013, *MNRAS*, 436, 1023
 Becker G. D., Bolton J. S., Haehnelt M. G., Sargent W. L. W., 2011, *MNRAS*, 410, 1096
 Becker G. D., Bolton J. S., Madau P., Pettini M., Ryan-Weber E. V., Venemans B. P., 2015, *MNRAS*, 447, 3402
 Becker G. D., D’Aloisio A., Christenson H. M., Zhu Y., Worseck G., Bolton J. S., 2021, *MNRAS*, 508, 1853
 Begley R. et al., 2022, *MNRAS*, 513, 3510
 Boera E., Becker G. D., Bolton J. S., Nasir F., 2019, *ApJ*, 872, 101
 Bolton J. S., Becker G. D., 2009, *MNRAS*, 398, L26
 Bosman S. E. I., Fan X., Jiang L., Reed S., Matsuoka Y., Becker G., Haehnelt M., 2018, *MNRAS*, 479, 1055
 Bosman S. E. I. et al., 2022, *MNRAS*, 514, 55
 Bouwens R. J., Illingworth G. D., Oesch P. A., Caruana J., Holwerda B., Smit R., Wilkins S., 2015, *ApJ*, 811, 140
 Bouwens R. J. et al., 2021, *AJ*, 162, 47
 Bressan A., Marigo P., Girardi L., Salasnich B., Dal Cero C., Rubele S., Nanni A., 2012, *MNRAS*, 427, 127
 Cain C., D’Aloisio A., Gangolli N., Becker G. D., 2021, *ApJ*, 917, L37
 Cain C., D’Aloisio A., Gangolli N., McQuinn M., 2023, *MNRAS*, 522, 2064, <https://ui.adsabs.harvard.edu/abs/2023MNRAS.522.2047C>
 Cameron A. J., Katz H., Witten C., Saxena A., Laporte N., Bunker A. J., 2023, preprint (arXiv:2311.02051)
 Chan T. K., Benitez-Llambay A., Theuns T., Frenk C., Bower R., 2024, *MNRAS*, 528, 1296, <https://ui.adsabs.harvard.edu/abs/2024MNRAS.528.1296C>
 Choi J., Conroy C., Byler N., 2017, *ApJ*, 838, 159
 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
 D’Aloisio A., McQuinn M., Trac H., Cain C., Mesinger A., 2020, *ApJ*, 898, 149
 D’Odorico V. et al., 2023, *MNRAS*, 523, 1399
 Davies F. B., Furlanetto S. R., 2016, *MNRAS*, 460, 1328
 Davies F. B. et al., 2018, *ApJ*, 864, 142
 Davies F. B., Bosman S. E. I., Furlanetto S. R., Becker G. D., D’Aloisio A., 2021, *ApJ*, 918, L35

- Deparis N., Aubert D., Ocvirk P., Chardin J., Lewis J., 2019, *A&A*, 622, A142
- Dome T., Tacchella S., Fialkov A., Dekel A., Ginzburg O., Lapiner S., Looser T. J., 2024, *MNRAS*, 527, 2139, <https://ui.adsabs.harvard.edu/abs/2024MNRAS.527.2139D>
- Doughty C. C., Hennawi J. F., Davies F. B., Lukić Z., Oñorbe J., 2023, *MNRAS*, 525, 3790, <https://ui.adsabs.harvard.edu/abs/2023MNRAS.525.3790D>
- Eilers A.-C., Davies F. B., Hennawi J. F., 2018, *ApJ*, 864, 33
- Emami N., Siana B., Weisz D. R., Johnson B. D., Ma X., El-Badry K., 2019, *ApJ*, 881, 71
- Emberson J. D., Thomas R. M., Alvarez M. A., 2013, *ApJ*, 763, 146
- Endsley R., Stark D. P., 2022, *MNRAS*, 511, 6042
- Fan X. et al., 2006, *AJ*, 132, 117
- Faucher-Giguère C.-A., Lidz A., Zaldarriaga M., Hernquist L., 2009a, *ApJ*, 703, 1416
- Faucher-Giguère C.-A., Lidz A., Zaldarriaga M., Hernquist L., 2009b, *ApJ*, 703, 1416
- Finkelstein S. L. et al., 2019, *ApJ*, 879, 36
- Finlator K. et al., 2017, *MNRAS*, 464, 1633
- Furlanetto S. R., Oh S. P., 2005, *MNRAS*, 363, 1031
- Gaikwad P. et al., 2020, *MNRAS*, 494, 5091
- Gaikwad P. et al., 2023, *MNRAS*, 525, 4093, <https://ui.adsabs.harvard.edu/abs/2023MNRAS.525.4093G>
- Garaldi E., Kannan R., Smith A., Springel V., Pakmor R., Vogelsberger M., Hernquist L., 2022, *MNRAS*, 512, 4909, <https://ui.adsabs.harvard.edu/abs/2022MNRAS.512.4909G>
- Georgakakis A. et al., 2015, *MNRAS*, 453, 1946
- Gnedin N. Y., Kravtsov A. V., Rudd D. H., 2011, *ApJS*, 194, 46
- Gorski K. M., Wandelt B. D., Hansen F. K., Hivon E., Banday A. J., 1999, preprint (astro-ph/9905275)
- Greig B., Mesinger A., Haiman Z., Simcoe R. A., 2016, *MNRAS*, 466, 4239
- Greig B., Mesinger A., Bañados E., 2019, *MNRAS*, 484, 5094
- Haardt F., Madau P., 2012, *ApJ*, 746, 125
- Hu W. et al., 2019, *ApJ*, 886, 90
- Iliev I. T., Shapiro P. R., Raga A. C., 2005, *MNRAS*, 361, 405
- Jin X. et al., 2023, *ApJ*, 942, 59
- Kannan R., Vogelsberger M., Marinacci F., McKinnon R., Pakmor R., Springel V., 2019, *MNRAS*, 485, 117
- Kannan R., Garaldi E., Smith A., Pakmor R., Springel V., Vogelsberger M., Hernquist L., 2022, *MNRAS*, 511, 4005
- Katz H., Kimm T., Sijacki D., Haehnelt M. G., 2017, *MNRAS*, 468, 4831
- 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
- Kostyuk I., Nelson D., Ciardi B., Glatzle M., Pillepich A., 2023, *MNRAS*, 521, 3077
- Kulkarni G., Keating L. C., Haehnelt M. G., Bosman S. E. I., Puchwein E., Chardin J., Aubert D., 2019, *MNRAS*, 485, L24
- Lewis J. S. W., Ocvirk P., Dubois Y., Aubert D., Chardin J., Gillet N., Thélie É., 2023, *MNRAS*, 519, 5987
- Madau P., 1995, *ApJ*, 441, 18
- Maseda M. V. et al., 2020, *MNRAS*, 493, 5120
- Mason C. A. et al., 2018, *ApJ*, 857, L11
- Mason C. A. et al., 2019, *MNRAS*, 485, 3947
- Matthee J. et al., 2022, *MNRAS*
- Matthee J. et al., 2024, *ApJ*, 963, 129, <https://ui.adsabs.harvard.edu/abs/2024ApJ...963..129M>
- McGreer I. D. et al., 2013, *ApJ*, 768, 105
- McGreer I. D., Mesinger A., D’Odorico V., 2015, *MNRAS*, 447, 499
- McQuinn M., Upton Sanderbeck P. R., 2016, *MNRAS*, 456, 47
- McQuinn M., Oh S. P., Faucher-Giguère C.-A., 2011, *ApJ*, 743, 82
- Muñoz J. A., Oh S. P., Davies F. B., Furlanetto S. R., 2016, *MNRAS*, 455, 1385
- Naidu R. P., Tacchella S., Mason C. A., Bose S., Oesch P. A., Conroy C., 2020, *ApJ*, 892, 109
- Naidu R. P. et al., 2022, *MNRAS*, 510, 4582
- Nasir F., D’Aloisio A., 2020, *MNRAS*, 494, 3080
- Nasir F., Cain C., D’Aloisio A., Gangolli N., McQuinn M., 2021, *ApJ*, 923, 161
- Ocvirk P. et al., 2016, *MNRAS*, 463, 1462
- Ocvirk P., Aubert D., Chardin J., Deparis N., Lewis J., 2019, *A&A*, 626, A77
- Ocvirk P. et al., 2020, *MNRAS*, 496, 4087
- Ocvirk P., Lewis J. S. W., Gillet N., Chardin J., Aubert D., Deparis N., Thélie É., 2021, *MNRAS*, 507, 6108
- Okamoto T., Gao L., Theuns T., 2008, *MNRAS*, 390, 920
- Ouchi M. et al., 2018, *PASJ*, 70, S13
- Park H., Shapiro P. R., Choi J.-h., Yoshida N., Hirano S., Ahn K., 2016, *ApJ*, 831, 86
- Pillepich A. et al., 2017, *MNRAS*, 473, 4077
- Planck Collaboration VI, 2020, *A&A*, 641, A6
- Prochaska J. X., Worseck G., O’Meara J. M., 2009, *ApJ*, 705, L113
- Prochaska J. X., O’Meara J. M., Worseck G., 2010, *ApJ*, 718, 392
- Qin Y., Mesinger A., Bosman S. E. I., Viel M., 2021, *MNRAS*, 506, 2390
- Robertson B. E. et al., 2013, *ApJ*, 768, 71
- Robertson B. E., Ellis R. S., Furlanetto S. R., Dunlop J. S., 2015, *ApJ*, 802, L19
- Rosdahl J., Blaizot J., Aubert D., Stranex T., Teyssier R., 2013, *MNRAS*, 436, 2188
- Rosdahl J. et al., 2018, *MNRAS*, 479, 994
- Rosdahl J. et al., 2022, *MNRAS*, 515, 2386
- Roth J. T., D’Aloisio A., Cain C., Wilson B., Zhu Y., Becker G. D., 2023, *MNRAS*, 530, 5209, <https://ui.adsabs.harvard.edu/abs/2024MNRAS.530.5209R>
- Saldana-Lopez A. et al., 2023, *MNRAS*, 522, 6295
- Satyavolu S., Kulkarni G., Keating L. C., Haehnelt M. G., 2023, preprint (arXiv:2311.06344)
- Shapiro P. R., Giroux M. L., Babul A., 1994, *ApJ*, 427, 25
- Shapiro P. R., Iliev I. T., Raga A. C., 2004, *MNRAS*, 348, 753
- Theuns T., 2021, *MNRAS*, 500, 2741
- Theuns T., Chan T. K., 2024, *MNRAS*, 527, 689
- Trac H., Cen R., 2007, *ApJ*, 671, 1
- Trac H., Pen U.-L., 2004, *New Astron.*, 9, 443
- Trac H., Cen R., Mansfield P., 2015, *ApJ*, 813, 54
- Trebtsch M. et al., 2021, *A&A*, 653, A154
- Vogelsberger M. et al., 2014, *MNRAS*, 444, 1518
- Walther M., Oñorbe J., Hennawi J. F., Lukić Z., 2019, *ApJ*, 872, 13
- Wang F. et al., 2020, *ApJ*, 896, 23
- Weinberger R. et al., 2016, *MNRAS*, 465, 3291
- Weldon A. et al., 2022, *MNRAS*, 515, 841
- Willott C. J. et al., 2010, *AJ*, 139, 906
- Worseck G. et al., 2014, *MNRAS*, 445, 1745
- Wu X., Kannan R., Marinacci F., Vogelsberger M., Hernquist L., 2019a, *MNRAS*, 488, 419
- Wu X., McQuinn M., Kannan R., D’Aloisio A., Bird S., Marinacci F., Davé R., Hernquist L., 2019b, *MNRAS*, 490, 3177
- Wu X., McQuinn M., Eisenstein D., 2021, *J. Cosmol. Astropart. Phys.*, 2021, 042
- Yang J. et al., 2020a, *ApJ*, 897, L14
- Yang J. et al., 2020b, *ApJ*, 904, 26
- Yeh J. Y. C. et al., 2022, *MNRAS*, 520, 2757, <https://ui.adsabs.harvard.edu/abs/2023MNRAS.520.2757Y>
- Zhu Y. et al., 2022, *ApJ*, 932, 76
- Zhu Y. et al., 2023, *ApJ*, 955, 115, <https://ui.adsabs.harvard.edu/abs/2023ApJ...955..115Z>
- Davies F. B. et al., 2024, *ApJ*, 965, 134, <https://ui.adsabs.harvard.edu/abs/2024ApJ...965..134D>

APPENDIX A: MULTIFREQUENCY GENERALIZATION FOR Γ_{H1} (EQUATION 7)

Consider an infinitely sharp I-front traveling along one axis of cell i . Then ray j intersecting cell i will travel a distance $x_{ion}^{ij} \Delta s^{ij}$ (recall

Δs^{ij} is the total path length of ray j through cell i) before reaching neutral gas. The number of photons absorbed over this distance is

$$N_{\text{abs}}^i = \sum_{j=1}^{N_{\text{rays}}} \sum_v N_{0,v}^{ij} \left(1 - \exp \left[\frac{-x_{\text{ion}}^i \Delta s^{ij}}{\lambda_v^i} \right] \right), \quad (\text{A1})$$

where the outer sum runs over all rays j intersecting cell i and the inner sum runs over all frequency bins of ray j . Here, $N_{0,v}^{ij}$ is the number of photons in ray j entering cell i at frequency v and λ_v^i is the frequency-dependent MFP. During a time step Δt , $\Gamma_{\text{H I}}$ in ionized gas behind the I-front is

$$\Gamma_{\text{H I}}^i = \frac{\text{\# of photons absorbed per time}}{\text{\# of H I atoms in ionized gas}} = \frac{N_{\text{abs}}^i / \Delta t}{n_{\text{H I}}^i x_{\text{ion}}^i V_{\text{cell}}}, \quad (\text{A2})$$

where $x_{\text{ion}}^i V_{\text{cell}}$ is the ionized volume of cell i and

$$n_{\text{H I}}^i \equiv \frac{\langle \Gamma_{\text{H I}} n_{\text{H I}} \rangle_V}{\langle \Gamma_{\text{H I}} \rangle_V} \quad (\text{A3})$$

is the $\Gamma_{\text{H I}}$ -weighted H I number density (the V subscript denotes a volume average). In appendix C of Cain et al. (2023), we showed that the frequency-averaged MFP in our sub-grid simulations is given by

$$\langle \lambda_v^{-1} \rangle_v^i = \frac{\langle n_{\text{H I}} \Gamma_{\text{H I}} \rangle_V}{F_\gamma} = \frac{1}{F_\gamma} \int_{\nu_{\text{H I}}}^{4\nu_{\text{H I}}} d\nu \frac{I_\nu}{h\nu} \lambda_v^{-1}, \quad (\text{A4})$$

where I_ν and F_γ are the specific intensity and ionizing photon flux at the source planes in the sub-grid simulations, respectively. Using this result, we can write

$$\frac{n_{\text{H I}}^i}{F_\gamma} = \frac{1}{\langle \Gamma_{\text{H I}} \rangle_V} \int_{\nu_{\text{H I}}}^{4\nu_{\text{H I}}} d\nu \frac{1}{F_\gamma} \frac{I_\nu}{h\nu} \lambda_v^{-1} \approx \frac{\langle \lambda_v^{-1} \rangle_v^i}{\langle \Gamma_{\text{H I}} \rangle_V}, \quad (\text{A5})$$

where $\langle \lambda_v^{-1} \rangle_v^i$ is the opacity averaged over the spectrum incident on cell i . Equation (A5) ignores IGM filtering over distances smaller than the cell size, but this approximation holds as long as $\lambda_{912}^{\text{mfp}} \gg \Delta x_{\text{cell}}$. Combining equation (A1)–(A5) yields

$$\Gamma_{\text{H I}}^i = \frac{\sum_{j=1}^{N_{\text{rays}}} \sum_v N_{0,v}^{ij} \left(1 - \exp \left[\frac{-x_{\text{ion}}^i \Delta s^{ij}}{\lambda_v^i} \right] \right)}{(\langle \lambda_v^{-1} \rangle_v F_\gamma / \langle \Gamma_{\text{H I}} \rangle_V) x_{\text{ion}}^i V_{\text{cell}} \Delta t}. \quad (\text{A6})$$

The ratio $F_\gamma / \langle \Gamma_{\text{H I}} \rangle_V$ can be simplified as long as $F_\gamma \langle \sigma_{\text{H I}} \rangle_v^i \approx \langle \Gamma_{\text{H I}} \rangle_V$, where $\langle \sigma_{\text{H I}} \rangle_v^i$ is the H I cross-section averaged over the spectrum incident on cell i . Under this approximation,

$$\Gamma_{\text{H I}}^i = \frac{\langle \sigma_{\text{H I}} \rangle_v^i \sum_{j=1}^{N_{\text{rays}}} \sum_v N_{0,v}^{ij} \left(1 - \exp \left[\frac{-x_{\text{ion}}^i \Delta s^{ij}}{\lambda_v^i} \right] \right)}{\langle \lambda_v^{-1} \rangle_v^i x_{\text{ion}}^i V_{\text{cell}} \Delta t}. \quad (\text{A7})$$

If we define $\langle \lambda_v \rangle_v \equiv \langle \lambda_v^{-1} \rangle_v^{-1}$, then equation (A8) can be written in a form similar to equation (1) of Cain et al. (2021),

$$\Gamma_{\text{H I}}^i = \frac{\langle \lambda_v \rangle_v^i \langle \sigma_{\text{H I}} \rangle_v^i \sum_{j=1}^{N_{\text{rays}}} \sum_v N_{0,v}^{ij} \left(1 - \exp \left[\frac{-x_{\text{ion}}^i \Delta s^{ij}}{\lambda_v^i} \right] \right)}{x_{\text{ion}}^i V_{\text{cell}} \Delta t}. \quad (\text{A8})$$

APPENDIX B: TESTING THE MULTIFREQUENCY MODEL AND EFFECT IGM FILTERING

In this appendix, we validate our multifrequency RT treatment described in Section 3.4. We tested this procedure using a set of small-volume hydro/RT simulations similar to the ones used to calibrate our sub-grid model. We have run three simulations in $0.512 h^{-1} \text{ Mpc}$ boxes with $N = 512^3$ gas/RT cells, each with $\Gamma_{\text{H I}} = 3 \times 10^{-13}$, $z_{\text{reion}} = 8$, and with the box-scale mean density equal to the cosmic mean. Our simulations have different power law spectra with indices $\alpha =$




Figure B1. Direct test of our multifrequency procedure described in Section 3.4. We have run small-scale simulations with $\alpha = 1.5, 0.5$, and -0.5 , and used the procedure in that section to compute $\langle \lambda_v \rangle_{\text{spec}}$, which is the frequency-averaged MFP estimated using only information from the $\alpha = 1.5$ simulation. Here, we show the ratio between $\langle \lambda_v \rangle_{\text{spec}}$ and the true value extracted from each simulation $\langle \lambda_v \rangle_{\text{sim}}$. The agreement is within 1 per cent for $\alpha = 1.5$ and always within 20 per cent for the others.

1.5 (as used in our sub-grid simulations), 0.5, and -0.5 , each with five frequency bins spanning 1–4 Ryd.

For each α , we use the result of the $\alpha = 1.5$ simulation to estimate the frequency-averaged MFP, which we call $\langle \lambda_v \rangle_{\text{spec}}$, following the procedure described in Section 3.4. Next, we calculate the actual average MFP in each simulation, $\langle \lambda_v \rangle_{\text{sim}}$. Fig. B1 shows the ratio of these quantities for each value of α . As expected, we find very close agreement for $\alpha = 1.5$. For $\alpha = 0.5$ (-0.5), the re-constructed $\langle \lambda_v \rangle_{\text{spec}}$ differs from the truth by at most 10 per cent (20 per cent), with the agreement improving to better than 5 per cent $\Delta z = 1$ after the gas is ionized. This disagreement likely owes to differences in the self-shielding properties of the gas in simulations with different α , which are unable to account for with our current approach. Given that $\langle \lambda_v \rangle_{\text{sim}}$ for the $\alpha = 1.5$ and -0.5 simulations are different by almost a factor of 3, this level of agreement is acceptable for our purposes.

Next, we illustrate the importance of including multifrequency RT in our FlexRT simulations for the Ly α forest. To show this, we have run an illustrative simulation with the same box size and fiducial source prescription as the simulations presented in this paper. The simulation includes multifrequency RT with $\alpha = 1.5$ and uses our full sinks sub-grid model for the IGM. In the left panel of In Fig. B2, we show a slice through the column density distribution slope β_N at $z = 6$, computed using equation (12). We find $\beta_N \approx 1.6$ – 1.9 in highly ionized gas, with the lowest values in the most overdense cells, consistent with the findings of McQuinn et al. (2011).²⁰ We find lower values ($\beta_N \approx 1.4$ – 1.5) in the most recently ionized gas close to I-fronts. This is consistent with the results of Nasir et al. (2021) – they found that recently ionized gas has a shallower column density slope owing to abundant tiny sinks that have not yet photoevaporated. These values are somewhat higher than those typically assumed for

²⁰They found $\beta_N \approx 1.8$ for optically thin systems.




Figure B2. Visualization of the effect of IGM filtering at $z = 6$ in our multifrequency treatment described in Section 3.4. This simulation assumes $\alpha = 1.5$ for sources and uses our full sinks sub-grid model. Left: β_N estimated using equation (12). We find $\beta_N \approx 1.6$ – 1.9 in most ionized gas, with smaller values (≈ 1.4 – 1.5) close to I-fronts. Middle: ratio of the IGM heating rate with the value expected for an $\alpha = 1.5$ spectrum. IGM heating rates are enhanced by as much as a factor of 1.5 close to ionization fronts, where the effect of IGM filtering is greatest. Right: spectral index (α_{eff}) that would reproduce the heating rates shown in the middle panel. We see that α_{eff} can be as small as 0 close to ionization fronts.

the IGM at $z \sim 6$ (e.g. Becker et al. (2021) assumes $\beta_N = 1.3$). Larger values of β_N result in a stronger frequency dependence of the ionizing opacity (equation 11) and more IGM filtering.

The middle panel of Fig. B2 shows the ratio of the photoheating rate $\mathcal{H}$ in our simulation with that expected for an $\alpha = 1.5$ spectrum (equation 4). We see that near the centres of ionized bubbles where the sources are clustered, this ratio is close to 1. This is because the radiation spectrum in those regions is on average close to that emitted by the sources, since most of the radiation has not yet been filtered by the IGM. Near the edges of the ionized bubbles, where the radiation has traveled several MFPs on average, this ratio gets as large as 1.5, demonstrating that it has been hardened significantly by the IGM. The right panel converts this enhanced heating rate into an effective spectral index α_{eff} , defined such that $\mathcal{H}(\alpha_{\text{eff}}) \equiv \mathcal{H}$ for each cell. We see that α_{eff} declines from ≈ 1.5 at the centre of ionized bubbles to ≈ 0 near their edges. These illustrations highlight the importance of IGM filtering effects for both the averages and spatial fluctuations of the heating rate and photoionization rate in the IGM, which are crucial for accurately modelling the Ly α forest.

APPENDIX C: DERIVATION AND TESTS OF THE RECOMBINATION RADIATION MODEL (EQUATIONS 13–15)

Here, we will derive and test our model for recombination radiation (equations 13–15). We start with equation (13), which describes the effective recombination coefficient if a fraction $f_{\text{esc}}^{\text{rec}}$ of IRPs escape the dense clumps where they are produced. Assuming photoionization equilibrium,

$$\langle \Gamma_{\text{HI}} n_{\text{HI}} \rangle_V = \langle \alpha(T) n_e n_{\text{HI}} \rangle_V, \quad (\text{C1})$$

where the average is over volume. In appendix B of Cain et al. (2022), we showed that the left hand side of this equation can be expressed in terms of the frequency-averaged MFP $\langle \lambda_v^{-1} \rangle_v^{-1}$,

$$\langle \Gamma_{\text{HI}} n_{\text{HI}} \rangle_V = \frac{\langle \Gamma_{\text{HI}} \rangle_V}{\langle \lambda_v^{-1} \rangle_v^{-1} \langle \sigma_{\text{HI}} \rangle_V}. \quad (\text{C2})$$

Combining these equations yields

$$\langle \lambda_v^{-1} \rangle_v^{-1} = \frac{\Gamma_{\text{HI}}}{\langle \sigma_{\text{HI}} \rangle_V \langle \alpha(T) n_e n_{\text{HI}} \rangle_V}. \quad (\text{C3})$$

Let the MFP extracted from our sub-grid simulations, which assume the case B recombination approximation, be denoted $\langle \lambda_{v,B}^{-1} \rangle_v^{-1} \equiv \langle \kappa_B \rangle_v$. Since equation (C3) also holds in this case, we have

$$\langle \lambda_v^{-1} \rangle_v^{-1} \equiv \langle \kappa \rangle_v = \langle \kappa_B \rangle_v \frac{\langle \alpha(T) n_e n_{\text{HI}} \rangle_V}{\langle \alpha_B(T) n_e n_{\text{HI}} \rangle_V} \approx \langle \kappa_B \rangle_v \frac{\alpha(T)}{\alpha_B(T)}, \quad (\text{C4})$$

which is equation (14). In the final step, we have assumed that the ratio of the recombination rates in the sub-grid simulations can be approximated by the ratio of the recombination coefficients evaluated at the temperature computed in FlexRT.

Lastly, we derive equation (15). The production rate of IRPs that are not immediately re-absorbed is given by

$$\dot{n}_\gamma^{\text{rec}} \approx x_{\text{ion}} (\langle \alpha(T) n_e n_{\text{HI}} \rangle_V - \langle \alpha_B(T) n_e n_{\text{HI}} \rangle_V (1 + \chi)), \quad (\text{C5})$$

where the factor of $1 + \chi$ approximately accounts for ground state recombinations from He II. Applying equation (C3) on both terms inside parentheses yields

$$\begin{aligned} \dot{n}_\gamma^{\text{rec}} &= x_{\text{ion}} \left(\frac{\Gamma_{\text{HI}}}{\langle \sigma_{\text{HI}} \rangle_V \langle \lambda_v^{-1} \rangle_v^{-1}} - \frac{\Gamma_{\text{HI}}}{\langle \sigma_{\text{HI}} \rangle_V \langle \lambda_{v,B}^{-1} \rangle_v^{-1}} \right) (1 + \chi) \\ &= x_{\text{ion}} \frac{\Gamma_{\text{HI}}}{\langle \sigma_{\text{HI}} \rangle_V} (\langle \kappa \rangle_v - \langle \kappa_B \rangle_v (1 + \chi)). \end{aligned} \quad (\text{C6})$$

Combining with equation (13) and (14) yields

$$\dot{n}_\gamma^{\text{rec}} = x_{\text{ion}} \frac{\Gamma_{\text{HI}}}{\langle \sigma_{\text{HI}} \rangle_V} (f_{\text{esc}}^{\text{rec}} \langle \kappa_A \rangle_v + (1 - f_{\text{esc}}^{\text{rec}}) \langle \kappa_B \rangle_v - \langle \kappa_B \rangle_v (1 + \chi)) \quad (\text{C7})$$

which simplifies to equation (15).

APPENDIX D: DERIVATION OF THE HALO ABSORBER MODEL

In this appendix, we provide a more complete derivation of our halo absorber model described in Section 3.6, and explore some of its key features. First, we can approximately write M_{200} as a function of the gas density profile of the halo,

$$M_{200} \approx \frac{\Omega_m}{\Omega_b} \int_0^{R_{200}} dr (4\pi r^2) \rho_{\text{gas}}(r) = \frac{\Omega_m}{\Omega_b} \int_0^{R_{200}} dr (4\pi r^2) \frac{n_{\text{H}}(r) m_{\text{p}}}{X_{\text{Hy}}}, \quad (\text{D1})$$

where we have assumed that the gas in the halo traces the dark matter




Figure D1. Features of the halo absorber model described in Section 3.6. The black solid curve in the top panel shows f_{200} for a halo with $\Gamma_{-12} = 0.3$, $M_{\text{halo}} = 10^{10} h^{-1} M_{\odot}$, and $\epsilon = 2$. The other curves show the effects of varying these parameters one at a time. The bottom panel shows the total opacity from haloes, integrated over the HMF, for $\epsilon = 2$ and 1.5 (and $\Gamma_{-12} = 0.3$). See text for details.

with fraction Ω_b/Ω_m , and that the halo is spherically symmetric. In the second equality, m_p is the proton mass and X_{Hy} is the hydrogen fraction. Combining equation (D1) with equation (21) and the definition of R_{200} and integrating yields a solution for $n_{\text{H}}(r)$:

$$n_{\text{H}}(r) = \frac{\Omega_b}{\Omega_m} \frac{X_{\text{Hy}}}{m_p} (3 - \epsilon) \frac{200 \rho_{\text{crit}}(z)}{3} \left(\frac{r}{R_{200}} \right)^{-\epsilon}. \quad (\text{D2})$$

Putting this result into equation (22) and assuming photoionizational equilibrium gives

$$r_{\text{abs}} = \left[\frac{(\sigma_{\text{HI}}) \alpha_{\text{B}}(T_0)(1 + \chi)}{\Gamma_{\text{HI}}} \right]^{\frac{1}{2\epsilon-1}} \times \left[\frac{200 \Omega_b X_{\text{Hy}} (3 - \epsilon) \rho_{\text{crit}}(z)}{3 \Omega_m m_p} \right]^{\frac{2}{2\epsilon-1}} R_{200}^{\frac{2\epsilon}{2\epsilon-1}}, \quad (\text{D3})$$

where f_{200} is the right-hand side divided by R_{200} (equation 18). We assume $T_0 = 10^4$ K for ionized gas in and around haloes. The scaling relations in equation (23) follow from equation (D3).

The top panel of Fig. D1 shows how f_{200} responds to each of the parameters in our model as a function of redshift. The black solid curve shows a halo with $\Gamma_{-12} = 0.3$, $M_{\text{halo}} = 10^{10} h^{-1} M_{\odot}$, and $\epsilon = 2$ (isothermal density profile). The blue dotted, red dashed, and magenta dot-dashed curves show the effect of changing Γ_{-12} , ϵ , and M_{halo} , respectively, one at a time (see legend). Higher Γ_{-12} (M_{halo}) results in lower (higher) f_{200} , but does not affect the shape of




Figure D2. Power-law index of the TDR, $\gamma(z, z_{\text{re}})$, used to correct TDR when mapping our coarse-grained RT temperatures onto the high-resolution density grid for our Ly α forest calculations. γ approaches its limiting value of $5/3$ when $z < z_{\text{reion}}$, while when $z = z_{\text{reion}}$, $\gamma = 1$ (for isothermal gas). These values are calculated by fitting the solution of the McQuinn & Upton Sanderbeck (2016) IGM temperature model to a power law at $\Delta \leq 1$.

its redshift evolution. Smaller ϵ leads to both larger f_{200} and much steeper redshift evolution.

The bottom panel shows the total Lyman limit opacity from haloes, integrated over the HMF, for the black and red curves in the top panel. We see that the model with $\epsilon = 1.5$ (our Optimistic model) results in $2.5\text{--}5 \times$ more opacity from haloes than $\epsilon = 2$ (our realistic model). However, the opacity from haloes begins to turn over at $z = 7$ in the former, as effect of f_{200} decreasing outpaces that of the rapidly growing HMF. This turnover is steeper in the actual simulations because Γ_{HI} increases rapidly at $z < 6$. This decline in halo opacity at the end of reionization may be one reason why our optimistic models allows for only a flat ionizing emissivity rather than the rapidly rising one predicted by Robertson et al. (2015). In a forthcoming paper, we will explore the parameter space of this halo absorber model in more detail, and determine whether models that do allow for a rapidly rising emissivity are physically reasonable.

APPENDIX E: CORRECTING THE IGM TEMPERATURE-DENSITY RELATION

In this appendix, we describe in more detail our procedure for correction the IGM temperature–density relation (TDR) in our Ly α forest calculation, described in Section 3.7. We describe the expected TDR in each FlexRT cell using the power-law index $\gamma(z, z_{\text{reion}})$ (see equation 24). We estimate $\gamma(z, z_{\text{reion}})$ by fitting the solution of the McQuinn & Upton Sanderbeck (2016) analytic temperature model to a power law at $\Delta \leq 1$ (the density range that sets high- z forest transmission). We then assign γ values to each cell using the FlexRT estimate of z_{reion} , defined as the redshift at which the neutral fraction crosses 0.5. Fig. D2 shows γ versus z versus z_{reion} for $5 < z, z_{\text{reion}} < 12$. Freshly ionized gas ($z = z_{\text{reion}}$) is isothermal ($\gamma = 1$), while for $z < z_{\text{reion}}$, $\gamma \rightarrow 5/3$.

In Fig. D3, we show the effect of applying this correction to the TDR in our high-resolution simulation used to calculate forest statistics. The left panel shows the TDR in FlexRT, which shows a clear power law with a population of lower low density cells at higher temperatures. These hotter underdense cells are those most recently reionized. Mapping these temperatures onto the high-




Figure D3. The TDR in our $\text{Ly}\alpha$ forest calculation before and after the correction described in this section. Left: the TDR in FlexRT at $z = 5.6$ in one of our simulations. Middle: TDR on the high-resolution density grid using a straight mapping of the FlexRT temperatures. Right: TDR on the high-resolution density grid after applying our correction. We see that this procedure approximately restores the qualitative features seen in the TDR in FlexRT: a tight power law with a population of recently ionized, underdense cells at much higher temperatures.

resolution hydro grid used for the forest almost completely erases these features, as the middle panel shows. The right panel shows the TDR on the hydro grid after applying our procedure, which recovers the qualitative behaviour seen in the left panel. The correction also reduces the total forest transmission by 10–20 per cent at $5 < z < 6$,

since the lowest-density cells with the highest transmission end up colder than they would be otherwise.

This paper has been typeset from a $\text{\LaTeX}$ file prepared by the author.