



Inferring the Thermal History of the Intergalactic Medium from the Properties of the Hydrogen and Helium Ly α Forest

Bruno Villaseñor¹ , Brant Robertson¹ , Piero Madau¹ , and Evan Schneider² ¹ Department of Astronomy and Astrophysics, University of California, Santa Cruz, 1156 High Street, Santa Cruz, CA 95064 USA; brvillas@ucsc.edu² Department of Physics and Astronomy & Pittsburgh Particle Physics, Astrophysics, and Cosmology Center (PITT PACC), University of Pittsburgh, Pittsburgh, PA 15260, USA

Received 2021 October 29; revised 2022 May 17; accepted 2022 May 17; published 2022 July 4

Abstract

The filamentary network of intergalactic medium (IGM) gas that gives origin to the Ly α forest in the spectra of distant quasars encodes information on the physics of structure formation and the early thermodynamics of diffuse baryonic material. Here we use a massive suite of more than 400 high-resolution cosmological hydrodynamical simulations run with the Graphics Processing Unit–accelerated code Cholla to study the IGM at high spatial resolution maintained over the entire computational volume. The simulations capture a wide range of possible IGM thermal histories by varying the photoheating and photoionizing background produced by star-forming galaxies and active galactic nuclei. A statistical comparison of synthetic spectra with the observed 1D flux power spectra of hydrogen at redshifts $2.2 \leq z \leq 5.0$ and with the helium Ly α opacity at redshifts $2.4 < z < 2.9$ tightly constrains the photoionization and photoheating history of the IGM. By leveraging the constraining power of the available Ly α forest data to break model degeneracies, we find that the IGM experienced two main reheating events over 1.2 Gyr of cosmic time. For our best-fit model, hydrogen reionization completes by $z_R \approx 6.0$ with a first IGM temperature peak of $T_0 \approx 1.3 \times 10^4$ K and is followed by the reionization of He II that completes by $z_R \approx 3.0$ and yields a second temperature peak of $T_0 \approx 1.4 \times 10^4$ K. We discuss how our results can be used to obtain information on the timing and the sources of hydrogen and helium reionization.

Unified Astronomy Thesaurus concepts: [Hydrodynamical simulations \(767\)](#); [Large-scale structure of the universe \(902\)](#); [Lyman alpha forest \(980\)](#); [Computational methods \(1965\)](#)

1. Introduction

The neutral hydrogen and singly ionized helium components of gas near the cosmic mean density trace the distribution of matter in between galaxies and produce a “forest” of detectable Ly α absorption features in the spectra of distant quasars (e.g., Hernquist et al. 1996; Croft et al. 1998; Meiksin 2009; Slosar et al. 2011; McQuinn 2016; Worseck et al. 2019). The depth, shape, and location of absorption lines in the Ly α forest depend on the ionization degree and thermal state of this intergalactic medium (IGM), which are controlled by the uncertain UV radiation background produced by star-forming galaxies and active galactic nuclei (AGNs; e.g., Haardt & Madau 2012; Madau & Haardt 2015; Robertson et al. 2015; Faucher-Giguère 2020), and on its density and peculiar velocity fields shaped by gravity (Cen et al. 1994). Dark matter provides the backbone of large-scale structure in the universe, a web-like pattern present in embryonic form in the overdensity motif of the initial fluctuation field and sharpened by nonlinear gravitational dynamics (Bond et al. 1996). The Ly α forest traces this underlying “cosmic web” on scales and at redshifts that cannot be probed by any other observable. Because of its long cooling time, low-density gas at $z \sim 2\text{--}5$ that traces the underlying matter distribution retains some memory of when and how it was reheated and reionized at $z \gtrsim 6$ (Miralda-Escudé & Rees 1994). The physics that governs the properties of the IGM throughout these epochs remain similar, as the evolving cosmic UV emissivity and the transfer of that radiation through a medium made clumpy by gravity determine

both the details of the reionization process and the thermodynamics of the forest.

Understanding how reionization occurred, the nature of the early sources that drove it, the thermal history and fine-grained properties of hydrogen gas in the cosmic web, and how to extract crucial information on the cosmological model from observations of Ly α absorption are among the most important open questions in cosmology and key science drivers for numerous major new instruments and facilities. The promise of the Ly α forest for constraining cosmological physics, including the nature of dark matter and dark energy, has motivated in part the construction of the Dark Energy Spectroscopic Instrument (DESI Collaboration et al. 2016), which measures absorption-line spectra backlit by nearly a million quasars at $z > 2$, and the WEAVE survey (Pieri et al. 2016), which will observe more than 400,000 high-redshift quasars at $z > 2$. Interpreting such observations requires detailed cosmological hydrodynamical simulations that cover an extensive range of uncertain photoionization and photoheating histories and consistently maintain high resolution throughout a statistically representative sub-volume of the universe.

This paper extends research efforts directly focused on advancing the state of the art in modeling the IGM physical structure in cosmological simulations while still achieving high computational efficiency, thereby providing higher-fidelity physical models for interpreting Ly α forest data. In Villaseñor et al. (2021) we introduced the Cholla IGM Photoheating Simulations (CHIPS) to investigate how different photoheating histories and cosmological parameters impact the structure of the forest. Here we use a massive suite of more than 400 CHIPS simulations to study the IGM at a resolution of $49 h^{-1}$ ckpc maintained over $(50 h^{-1} \text{ cMpc})^3$ volumes. Performed with the GPU-native



Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](#). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

MPI-parallelized code Cholla (Schneider & Robertson 2015), these simulations span different amplitudes and peak redshifts of the H I and He II photoionization and photoheating rates.

To anticipate the results of our likelihood analysis constrained by the 1D flux power spectra $P(k)$ measured in eBOSS, Keck, and VLT data and the observed He II Ly α forest, we find that scenarios where hydrogen in the cosmic web was fully reionized by star-forming galaxies by redshift $z_R \approx 6.0$ and the double reionization of helium was completed by quasar sources about 1.2 billion years later are strongly favored by the data. Models that reionize hydrogen or helium at earlier or later cosmic times produce too much or too little cold gas and appear to be inconsistent with the observed $P(k)$ and He II Ly α opacity. Our approach differs from previous work in this field in the following aspects:

1. The simulation grid captures a wide range of possible thermal histories via a four-parameter scaling of the amplitude and timing of the (spatially uniform) metagalactic UV background (UVB) responsible for determining the ionization states and temperatures of the IGM (see Nasir et al. 2016; Oñorbe et al. 2017). We use the physically motivated model of Puchwein et al. (2019) as a template and vary the strength and redshift timing of their ionization and heating rates.
2. We do not modify, in post-processing, the mean transmitted flux $\langle F \rangle$ in the forest by recalibrating the Ly α optical depth, nor do we assume or rescale an instantaneous gas temperature–density relation (see Viel et al. 2013a; Iršič et al. 2017b; Boera et al. 2019; Walther et al. 2021). Indeed, we find from our simulations that the often assumed perfect power-law relationship between the temperature and density of the IGM does not provide a good approximation over the relevant density and redshift intervals.
3. Our likelihood analysis evaluates the performance of a given model in matching the observations over the complete self-consistently evolved reionization and thermal history of the IGM, i.e., over the full redshift range $2.2 \leq z \leq 5.0$ for the observed 1D flux power spectrum of hydrogen and over the redshift range $2.4 < z < 2.9$ for the Ly α opacity of He II. Since the properties of the gas at one redshift cannot be disentangled from its properties at previous epochs and the thermal and ionization structures of the IGM evolve with cosmic time along continuous trajectories, the marginalization over the parameter posterior distributions should not be performed independently at each redshift (see Bolton et al. 2014; Nasir et al. 2016; Hiss et al. 2018; Boera et al. 2019; Walther et al. 2019; Gaikwad et al. 2021).

This paper aims to find the optimal photoionization and photoheating rates that reproduce the observed properties of the hydrogen and helium Ly α forest. In Section 2 we describe the simulations used for this work, how we apply transformations to the UVB model from Puchwein et al. (2019) to generate our range of photoionization and photoheating rates, and the impact of the different UVB models on the statistics of the forest and the properties of the IGM. We follow by presenting the observational data and the methodology for the Bayesian

Markov Chain Monte Carlo (MCMC) inference used to constrain the model. Section 3 presents our result for the best-fit model and the comparison of the resulting properties of the forest and the thermal evolution of the IGM to the observational determinations and previous inferences. We summarize our results and conclusions in Section 4. In Appendix A we discuss resolution effects on the Ly α power spectrum $P(k)$ from our simulations. A quantitative study of the impact on $P(k)$ from rescaling the effective optical depth of the skewer sample is presented in Appendix B. In Appendix C we show the variation in the covariance matrix of the Ly α power spectrum from our simulations. We discuss in Appendix D how possible alterations to our model can modify the predicted temperature history of the IGM. Finally, Appendix E analyzes the accuracy of assuming a power-law relation for the density–temperature distribution of the gas in our simulations.

2. Methodology

For the study presented here, we compare the observed statistics of the Ly α forest to simulations that apply different models for the metagalactic UVB. In this section we briefly describe our simulation code and the method to extract Ly α spectra from the simulations. We then describe our simulation grid and the effects that the different UVB models have on the properties of the IGM. Finally, we present the observational measurements and the inference method used to constrain our model for the UVB photoionization and photoheating rates.

2.1. Simulations

The simulations used for this work were run with the cosmological hydrodynamics code Cholla (Schneider & Robertson 2015; Villasenor et al. 2021). Cholla evolves the equations of hydrodynamics on a uniform Cartesian grid using a finite-volume approach with a second-order Godunov scheme (Colella & Woodward 1984). The simulations track the ionization states of hydrogen and helium given by the photoionization from the UVB, recombination with free electrons, and collisional ionization. The nonequilibrium H + He chemical network is evolved simultaneously with the hydrodynamics using the GRACKLE library (Smith et al. 2017). We assume a spatially uniform, time-dependent UVB in the form of redshift-dependent photoionization rates per ion Γ and photoheating rates per ion \mathcal{H} for neutral hydrogen H I, neutral helium He I, and singly ionized helium He II. For a detailed description of the simulation code we refer the reader to the methodology section presented in Villasenor et al. (2021).

The initial conditions for our simulations were generated using the MUSIC code (Hahn & Abel 2011) for a flat Λ CDM cosmology with parameters $H_0 = 67.66 \text{ km s}^{-1}$, $\Omega_m = 0.3111$, $\Omega_\Lambda = 0.6889$, $\Omega_b = 0.0497$, $\sigma_8 = 0.8102$, and $n_s = 0.9665$, consistent with the constraints from Planck Collaboration et al. (2020). In future work, we plan to extend our analysis and include variation of the cosmological parameters (Bird et al. 2019; Ho et al. 2022). Unless otherwise stated, the volume and numerical size of our simulations correspond to $L = 50 \text{ h}^{-1} \text{ Mpc}$ and $N = 2 \times 1024^3$ cells and particles. The initial conditions for all runs were generated from identical random number seeds to preserve the same amplitude and phase for all initial Fourier modes across the simulation suite.

2.2. Synthetic Ly α Spectra

The Ly α forest sensitively probes the state of the baryons in the IGM, and absorption lines from the forest reflect the H I content and the temperature of the gas in the medium. To compare the properties of the IGM in our simulations directly to observations, we extract synthetic hydrogen Ly α forest spectra from the simulated boxes by measuring the H I density, temperature, and peculiar velocity of the gas along 12,228 skewers through the simulation volume, using 4096 skewers along each axis of the box. The optical depth τ as a function of velocity u along each skewer is computed by integrating the product of the Ly α scattering cross section and the number density of neutral hydrogen along the line of sight as described in Villasenor et al. (2021).

The transmitted flux F is computed from the optical depth τ along the skewers according to $F = \exp(-\tau)$. The power spectrum of the transmitted flux $P(k)$ is calculated as the average amplitude of the 1D Fourier transform of the flux fluctuations $\delta_F(u)$,

$$\delta_F(u) \equiv \frac{F(u) - \langle F \rangle}{\langle F \rangle}, \quad (1)$$

where $\langle F \rangle$ is the average transmitted flux over the skewer sample at a given redshift (see Section 5.4 from Villasenor et al. 2021 for a detailed description). Similarly, we extract the flux F_{HeII} transmitted through the He II Ly α forest from the simulations and compute the He II effective optical depth as $\tau_{\text{eff,HeII}} = -\ln \langle F_{\text{HeII}} \rangle$.

The top panel of Figure 1 shows the gas density distribution at redshift $z = 2$ from a section taken from one of our highest-resolution ($L = 50 h^{-1}$ Mpc, $N = 2 \times 2048^3$ cells and particles) simulations, where several skewers crossing the simulated box are shown as yellow lines. The bottom panels show the gas density surrounding a selected line of sight and the transmitted hydrogen Ly α flux along the skewer. The absorption lines in the forest probe the H I column density, the peculiar velocity, and the temperature of the gas along the line of sight.

2.3. Photoionization and Photoheating Rates

The ionization and thermal evolution of the IGM is primarily determined by the radiation emitted by star-forming galaxies and AGNs over cosmic history (McQuinn 2016; Upton Sanderbeck et al. 2016; Oñorbe et al. 2017). The photoionization and photoheating rates adopted in our simulations are computed from the intensity of the background radiation field, which is in turn determined by the emissivity of the radiating sources and the opacity of the IGM to ionizing photons. Recent models of the UVB (Khaire & Srianand 2019; Puchwein et al. 2019; Faucher-Giguère 2020), when applied to cosmological simulations, result in a hydrogen reionization era that completes by $z \sim 6-8$, in agreement with observational constraints (Davies et al. 2018; Planck Collaboration et al. 2020).

The updated model for the photoheating and photoionizing background presented in Puchwein et al. (2019, hereafter P19) adopts an improved treatment of the IGM opacity to ionizing radiation that consistently captures the transition from a neutral to an ionized IGM. To compute the intensity of the background radiation, the P19 model employs recent determinations of the ionizing emissivity due to stars and AGNs and of the H I absorber column density distribution and assumes an evolving

escape fraction of ionizing radiation from galaxies into the IGM that reaches 18%. When the P19 model is applied in cosmological simulations, hydrogen reionization completes at $z \sim 6$ consistently with recent measurements (Becker et al. 2001; Bosman et al. 2018; Becker et al. 2021; Qin et al. 2021). However, the subsequent evolution of the Ly α forest spectra measured in simulations that use the P19 model fails to reproduce the observed properties of the forest (Villasenor et al. 2021) and, in particular, does not agree with the observed power spectrum of the Ly α transmitted flux over the redshift range $2.2 \leq z \leq 5.0$. This work aims to present a new model photoionization and photoheating rates that result in an evolution of the IGM consistent with the observational measurements of the Ly α flux power spectrum and the He II effective optical depth.

2.4. Simulation Grid

To determine ionization and heating histories that result in properties of the IGM consistent with the observed Ly α flux power spectrum and He II effective opacity, we perform an unprecedented grid consisting of 400 cosmological simulations as a direct extension of the Cholla IGM Photoheating Simulations (CHIPS) suite originally presented in Villasenor et al. (2021). Each simulation in the CHIPS grid applies different photoionization and photoheating rates to model a variety of reionization and thermal histories and thereby produce different statistical properties for the Ly α forest. To generate different representations of the UVB, we modify the reference model from Puchwein et al. (2019) by rescaling the photoionization and photoheating rates (Γ and \mathcal{H} , respectively) by a constant factor β and shifting the redshift dependence of the rates by an offset Δz . The two transformations are expressed as

$$\begin{aligned} \Gamma(z) &\rightarrow \beta \Gamma^{\text{P19}}(z - \Delta z), \\ \mathcal{H}(z) &\rightarrow \beta \mathcal{H}^{\text{P19}}(z - \Delta z). \end{aligned} \quad (2)$$

Since the photoionization and photoheating rates for both H I and He I are dominated by the same sources, namely, star-forming galaxies at $z \gtrsim 5$ and AGNs at lower redshifts, and the radiation that ionizes both species is absorbed by intergalactic hydrogen, we modify the H I and He I photoionization and photoheating rates jointly by applying the transformations described by Equation (2), scaling and shifting by the parameters β_{H} and Δz_{H} , respectively. He II is reionized later in cosmic history primarily by the extreme-UV radiation emitted by AGNs, and we rescale and redshift-offset the photoionization and photoheating rates associated with He II by a second set of parameters β_{He} and Δz_{He} . Hence, each modified UVB model is characterized by the parameter vector $\theta = \{\beta_{\text{H}}, \Delta z_{\text{H}}, \beta_{\text{He}}, \Delta z_{\text{He}}\}$. The different photoionization and photoheating histories span all the combinations of the parameter values presented in Table 1.

The rescaling parameters β_{H} and β_{He} control the intensity of the background radiation, determine the efficiency with which H I and He II become ionized, and govern energy input into the IGM in the form of photoheating during the epochs of nonequilibrium reionization for hydrogen and helium. After reionization completes and the gas reaches photoionization equilibrium, the balance between ionizations from the background radiation and recombinations with free electrons

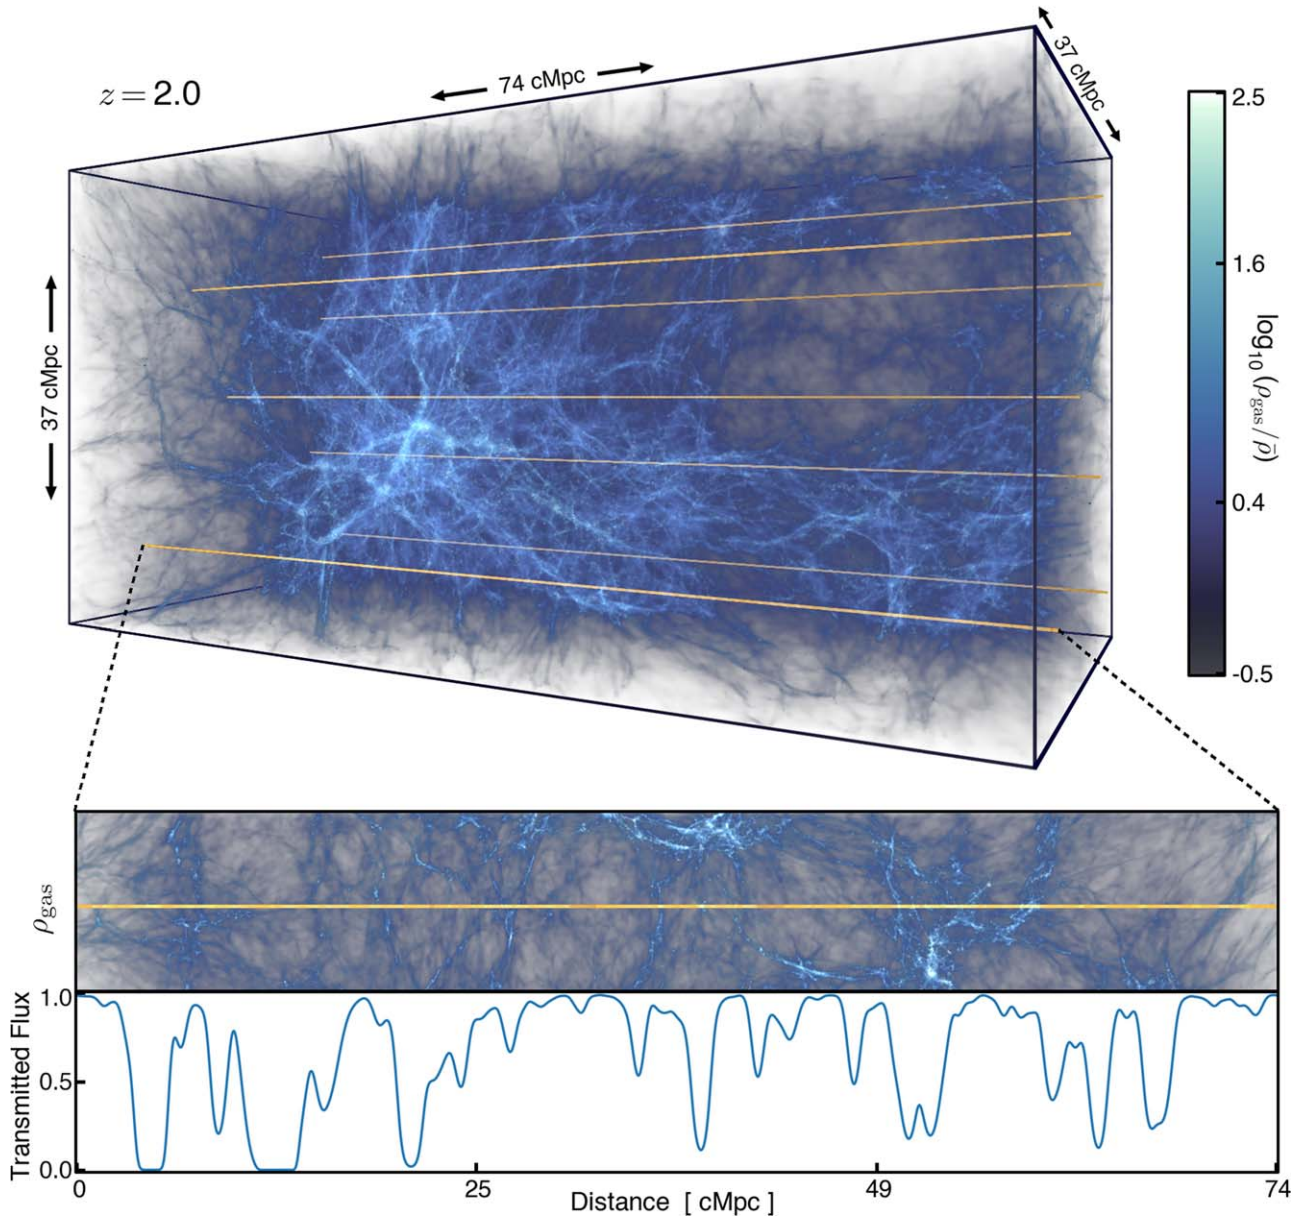


Figure 1. Large-scale distribution of gas density (top) from one of our highest-resolution cosmological simulations ($L = 50 h^{-1}$ Mpc, $N = 2 \times 2048^3$ cells and particles) at redshift $z = 2$ and a set of skewers crossing the simulated box (yellow lines). The bottom panels show the density of the gas surrounding a selected line of sight and the Ly α transmitted flux along the skewer. Absorption lines in the forest probe the H I column density, the peculiar velocity, and the temperature of the gas along the line of sight.

Table 1
CHIPS Simulation Grid

Parameter	Parameter Values
β_{H}	0.60, 0.73, 0.86, 1.00
Δz_{H}	-0.6, -0.4, -0.2, 0.0, 0.2
β_{He}	0.10, 0.30, 0.53, 0.76, 1.00
Δz_{He}	-0.1, 0.2, 0.5, 0.8

Note. The parameters β_{H} and Δz_{H} determine the amplitude and redshift offset of the H I and He I photoionization and photoheating rates, while β_{He} and Δz_{He} rescale and offset the He II rates.

determines the ionization state of H I and He II. At equilibrium, the ionized fraction of H I and He II is proportional to the photoionization rates Γ_{HI} and Γ_{HeII} , respectively, and inversely

proportional to the temperature-dependent radiative recombination rates $\alpha_{\text{HII}}(T)$ and $\alpha_{\text{HeIII}}(T)$. Therefore, by rescaling the photoionization and photoheating rates, we modify the evolution of the temperature and the ionization state of the gas in the IGM during and after H I and He II reionization.

The parameters Δz_{H} and Δz_{He} shift the redshift dependence of the photoionization and photoheating rates by a constant offset, affecting the timing of H I and He II reionization. In general, an offset of $\Delta z_{\text{H}} > 0$ or $\Delta z_{\text{He}} > 0$ moves H I or He II reionization to higher redshift and earlier cosmic time relative to the reference P19 model. Negative values of Δz_{H} or Δz_{He} shift reionization to lower redshift and later cosmic times. The offset in redshift of the models also affects the properties of the IGM after H I and He II reionization completes, as the photoheating and photoionization rates at a given redshift are generally modified when $\Delta z_{\text{H}} \neq 0$ or $\Delta z_{\text{He}} \neq 0$.

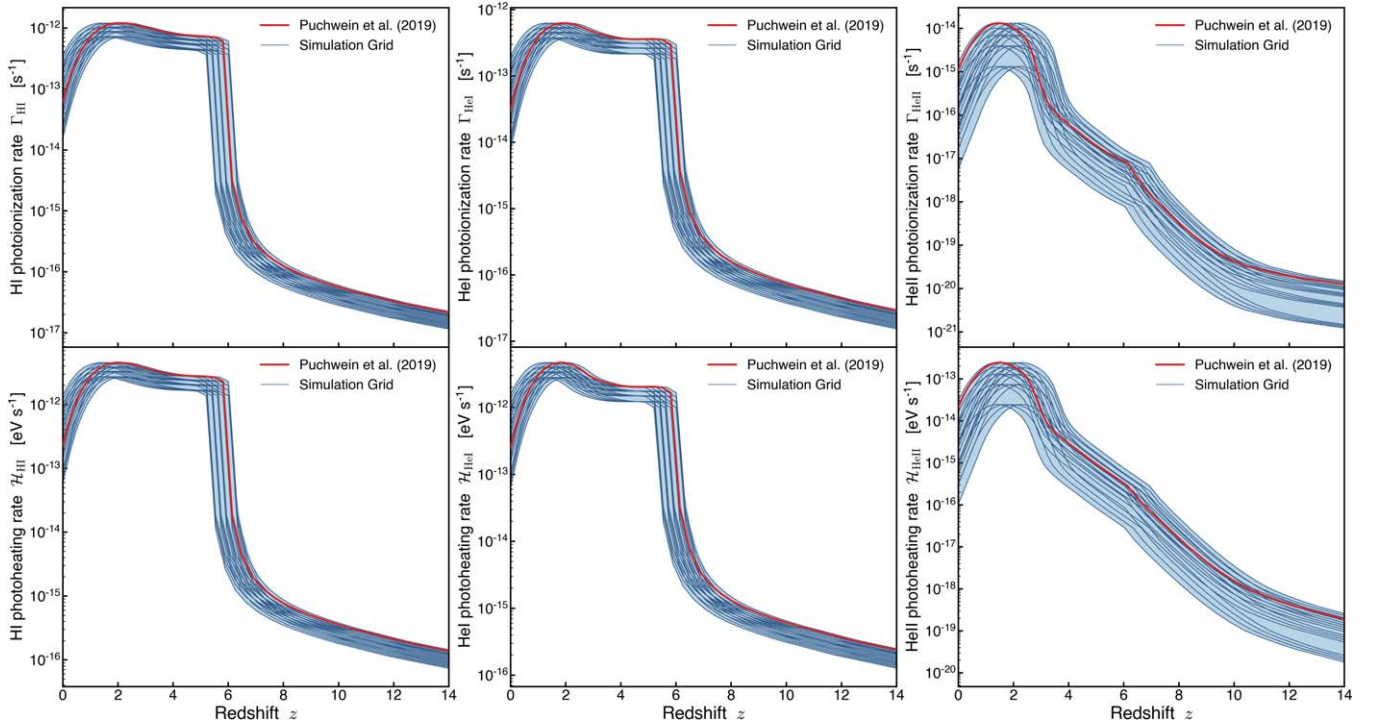


Figure 2. Photoionization (Γ , top) and photoheating (\mathcal{H} , bottom) rates for neutral hydrogen (H I; left), neutral helium (He I; middle), and singly ionized helium (He II; right) from the reference P19 model (Puchwein et al. 2019; red line), along with the photoionization and photoheating rates (blue lines) used for the 400 simulations of the CHIPS grid. The modified rates are generated by rescaling and shifting the reference P19 model as described in Section 2.4.

Figure 2 shows the photoionization and photoheating rates from the reference model by Puchwein et al. (2019), together with the modified rates adopted in the 400 simulations of the CHIPS grid. In Villasenor et al. (2021) we presented a comparison of the statistical properties of the Ly α forest and the thermal history of the IGM that result from a high-resolution simulation using the UVB model from Puchwein et al. (2019). We concluded that, in general, the gas in the simulation was too highly ionized after hydrogen reionization and possibly too hot during the epoch of helium reionization to be compatible with the observed statistics of the forest and other inferences of the thermal state of the IGM. We therefore do not include values of $\beta_{\text{H}} > 1$ or $\beta_{\text{He}} > 1$ in our grid, as such models would result in overall higher ionization fractions and temperatures of the IGM compared with the P19 case.

The simulations were run on the Summit system (Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory). Each simulation was performed on 128 GPUs and completed in less than two wall-clock hours. The cost of the entire grid of computations was only $\sim 16,000$ node hours. This work demonstrates that by taking advantage of an efficient code like Cholla and a capable system like Summit, future studies of the IGM using thousands of cosmological simulations are now possible.

2.5. Effects of UVB Models on the IGM Properties

The different photoionization and photoheating histories adopted in our simulations affect the ionization state of hydrogen and helium and the temperature of the IGM. Figure 3 shows the redshift evolution of the global properties of the IGM for each of the simulated histories. The top panels show the temperature of gas at mean density T_0 (left) and the index γ (right) of the power-law density–temperature relation

$T(\Delta) = T_0 \Delta^{\gamma-1}$, where $\Delta = \rho_{\text{gas}}/\bar{\rho}$. The bottom panels show the volume-weighted average fraction of neutral hydrogen x_{HI} (left) and singly ionized helium x_{HeII} (right).

As hydrogen becomes ionized at $z \gtrsim 5.5$, the gas in the IGM experiences a monotonic increase of T_0 while showing a close-to-isothermal distribution $\gamma \sim 1$. After hydrogen reionization ends at $z \sim 5.5$ – 6.5 , the gas cools primarily through the adiabatic expansion of the universe. During this period, the low-density gas cools faster and γ increases. This first epoch of cooling ends with the onset of helium (He II) reionization from the extreme-UV radiation emitted by AGNs at $z \lesssim 4$ – 5 , which reheats the IGM, increasing T_0 and decreasing γ . After the double reionization of helium completes ($z \sim 2.5$ – 3.5), the IGM cools monotonically by adiabatic expansion, increasing γ for a second time. Because of these two distinct photoheating epochs, the thermal state of the IGM in our simulations is more sensitive to variations in the hydrogen photoheating/photoionization parameters β_{H} and Δz_{H} at $z \gtrsim 5$ and more sensitive to the parameters β_{He} and Δz_{He} at $z \lesssim 5$ during the epoch of helium reionization.

For simulations with $\Delta z_{\text{H}} < 0$ the temperature peak from hydrogen reionization is shifted to later times (lower redshift) and the amplitude of the temperature peak depends on the value of β_{H} . Analogously, the parameters β_{He} and Δz_{He} determine the amplitude and timing of the second peak in T_0 caused by helium reionization. Positive values of Δz_{He} move helium reionization to higher redshifts compared with the reference P19 model, and higher values of β_{He} produce a higher peak in T_0 during the epoch $2.5 \lesssim z \lesssim 3.8$.

Variation in the timing of H and He reionization changes the cooling periods during which the power-law index γ increases. The different tracks of γ in our simulation grid then arise primarily from the different values of Δz_{H} and Δz_{He} adopted. In future work we plan to expand the flexibility of our

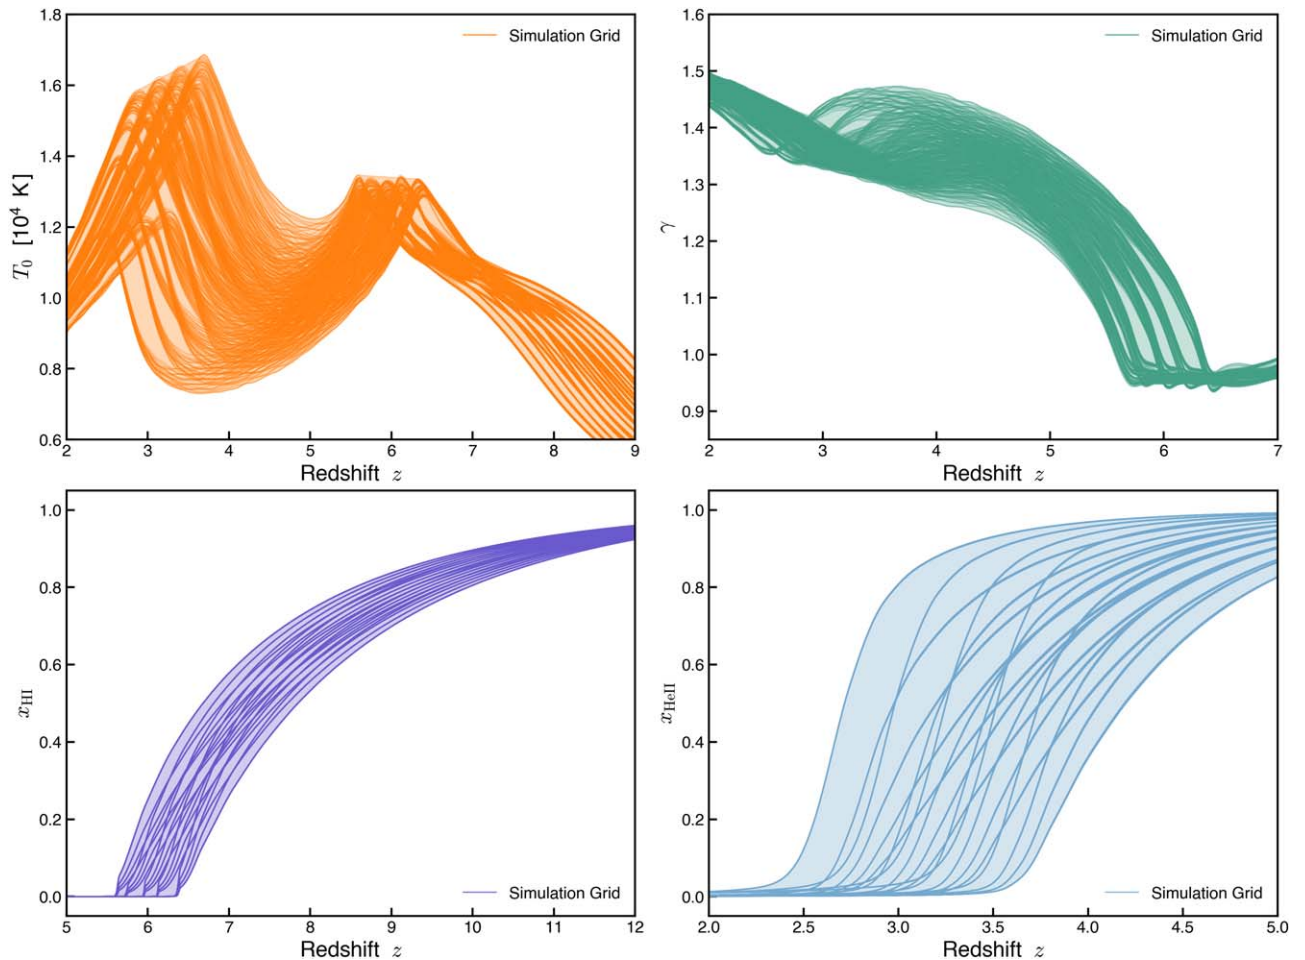


Figure 3. Evolution of global properties of the IGM computed from the 400 CHIPS simulations. The simulations evolve under different photoionization and photoheating rates, resulting in a large variety of ionization and thermal histories of the IGM. The top panels show the temperature, T_0 , of intergalactic gas at the mean density (left) and the index γ from the power-law density–temperature relation $T(\Delta) = T_0 \Delta^{\gamma-1}$ (right). The bottom panels show the volume-weighted average of the neutral hydrogen fraction x_{HI} (left) and the singly ionized helium fraction x_{HeII} (right). The amplitude and timing of the rates impact the thermal state of the IGM during H I and He II reionization. Simulations with higher values of β_{He} result in a higher temperature peak during He II reionization ($2.5 \lesssim z \lesssim 3.8$), and for simulations with $\Delta z_{\text{He}} > 0$ the epoch of He II reionization is shifted to earlier epochs. Analogously, negative values of Δz_{H} move the timing of H I reionization to later epochs, and simulations with different β_{H} show a different temperature peak during H I reionization at $z \sim 5.6$ – 6.3 .

simulations to sample the thermal state of the IGM by introducing density-dependent photoheating rates that will allow for larger variation in the evolution of γ .

The effects on the temperature of the IGM from the different helium reionization scenarios in our simulations are illustrated in Figure 4. The image displays the gas temperature of a slice through the IGM at $z = 3.6$ generated from a subset of 20 simulations that vary the parameters β_{He} and Δz_{He} controlling the He II photoionization and photoheating rates. Increases in β_{He} and Δz_{He} correspond to a larger extreme-UV background from AGNs and to a shift of the epoch of helium reionization to earlier cosmic times, respectively. Either effect causes the temperature of the IGM to increase at $z \sim 3.6$. Decreasing the He II photoheating rates or shifting helium reionization to later cosmic times (toward $z \sim 2.8$) decreases the temperature of IGM gas at $z \sim 3.6$.

2.6. Effects of UVB Models on the Ly α Forest Power Spectrum

The statistical properties of the Ly α forest provide insight into the state of the baryons in the IGM. The effective optical depth of the forest $\tau_{\text{eff,H}} = -\ln \langle F \rangle$ provides a global measurement of the overall H I content of the gas in the IGM, probes the hydrogen

ionization fraction, and allows for estimates of the intensity of the ionizing background radiation. The power spectrum $P(k)$ of the flux transmitted through the forest contains more information encoded across different spatial scales. On scales larger than a few Mpc the $P(k)$ is sensitive to the ionization fraction of hydrogen in the IGM and provides information similar to $\tau_{\text{eff,H}}$. This connection makes $P(k)$ and $\tau_{\text{eff,H}}$ a dependent pair of measurements, and Appendix B presents a detailed analysis about the effects that variations in $\tau_{\text{eff,H}}$ induce in $P(k)$. On scales smaller than a few comoving Mpc, structure in the forest is suppressed by pressure smoothing of the gas density fluctuations, as well as Doppler broadening of the absorption lines. These effects cause a cutoff in the dimensionless power spectrum $\Delta^2(k) = \pi^{-1} k P(k)$ for $k \gtrsim 0.02 \text{ s km}^{-1}$, making the flux power spectrum at intermediate and small scales a sensitive probe of the thermal state of IGM gas.

The different ionization and thermal histories produced by the range of photoionization and photoheating rates adopted in our simulations manifest as variations in the flux power spectrum of the Ly α forest. The effects on $P(k)$ from changing each of the four parameters β_{H} , Δz_{H} , β_{He} , or Δz_{He} independently is shown in Figure 5 for redshifts $z = 3$ (top) and $z = 4$ (bottom). The variation in $P(k)$ measured from our simulation grid over the




Figure 4. Gas temperature from a slice through the IGM at $z = 3.6$ in a subset of 20 simulations with different He II reionization scenarios. An increase in the parameters β_{He} and Δz_{He} corresponds to higher He II photoheating and a shift of the He II reionization epoch to earlier cosmic times (closer to $z \sim 3.6$), respectively. Either effect increases the temperature of the IGM at $z \sim 3.6$.

redshift range $2 \lesssim z \lesssim 5$ can be attributed mainly to three physical effects. First, since hydrogen is in photoionization equilibrium after H I reionization, changes to the photoionization rate Γ_{HI} from rescaling by β_{H} or applying a shift Δz_{H} alter the ionization fraction of hydrogen. This alteration globally affects the hydrogen effective optical depth $\tau_{\text{eff,H}}$, and, as a result, the overall normalization of $P(k)$ changes. Second, changes in the temperature of the IGM from the different hydrogen and helium reionization scenarios alter the recombination coefficient $\alpha_{\text{HII}}(T)$ in the IGM. In turn, changes to the recombination rate adjust the ionization fraction of hydrogen in the IGM and thereby the normalization of $P(k)$. Third, the different thermal histories of the IGM affect $P(k)$ on small scales through Doppler broadening of the absorption lines and the pressure smoothing of the density fluctuations. As shown in Figure 5, the parameters β_{H} and Δz_{H} mainly influence the normalization of $P(k)$ by changing the overall ionization fraction in the IGM, while the parameters β_{He} and Δz_{He} change the temperature of the IGM and thereby affect both the normalization and small-scale shape of $P(k)$.

2.7. Observational Data

For comparison with our simulations, we use the observational determinations of the flux power spectrum measured by the extended Baryon Oscillation Spectroscopy Survey (eBOSS; Chabanier et al. 2019) and separate measurements with the Keck Observatory and the Very Large Telescope (Iršič et al. 2017a;

Boera et al. 2019). The power spectrum estimates from Chabanier et al. (2019) probe mostly large scales ($0.001 \text{ s km}^{-1} \lesssim k \lesssim 0.02 \text{ s km}^{-1}$) in the redshift range $2.2 < z < 4.6$. The determinations from Iršič et al. (2017a) overlap with the eBOSS measurements on the large scales, albeit with lower precision, and extend to intermediate scales ($0.003 \text{ s km}^{-1} \lesssim k \lesssim 0.06 \text{ s km}^{-1}$) for redshifts $3.0 < z < 4.2$. The data from Boera et al. (2019) cover intermediate to small scales ($0.006 \text{ s km}^{-1} \lesssim k \lesssim 0.2 \text{ s km}^{-1}$) over the redshift range $4.2 < z < 5.0$. The combined data set spans a large redshift range from $z = 2.2$ to $z = 5.0$ and a wide range of scales, and it is shown along with our best-fit model $P(k)$ in Figure 6.

Figure 6 also shows the observational measurements of $P(k)$ presented by (Walther et al. 2018; purple open circles) for the redshift range $3.0 \leq z \leq 3.4$. We find that, in the overlapping range of scales ($0.003 \text{ s km}^{-1} \lesssim k \lesssim 0.02 \text{ s km}^{-1}$) and redshift ($2.2 \lesssim z \lesssim 3.4$), the estimates from Walther et al. (2018) show significant differences with those from eBOSS (Chabanier et al. 2019). The normalization and, in some cases, the shape of the large-scale $P(k)$ appear to be inconsistent between the two data sets. For several redshift bins (e.g., $z = 2.4$ and $z = 3.2$), a simple renormalization applied to the Walther et al. (2018) power spectrum would not be sufficient to match the large-scale measurements from eBOSS. Because of this discrepancy, we have not included the Walther et al. (2018) $P(k)$ determinations in our MCMC analysis, and we show them in

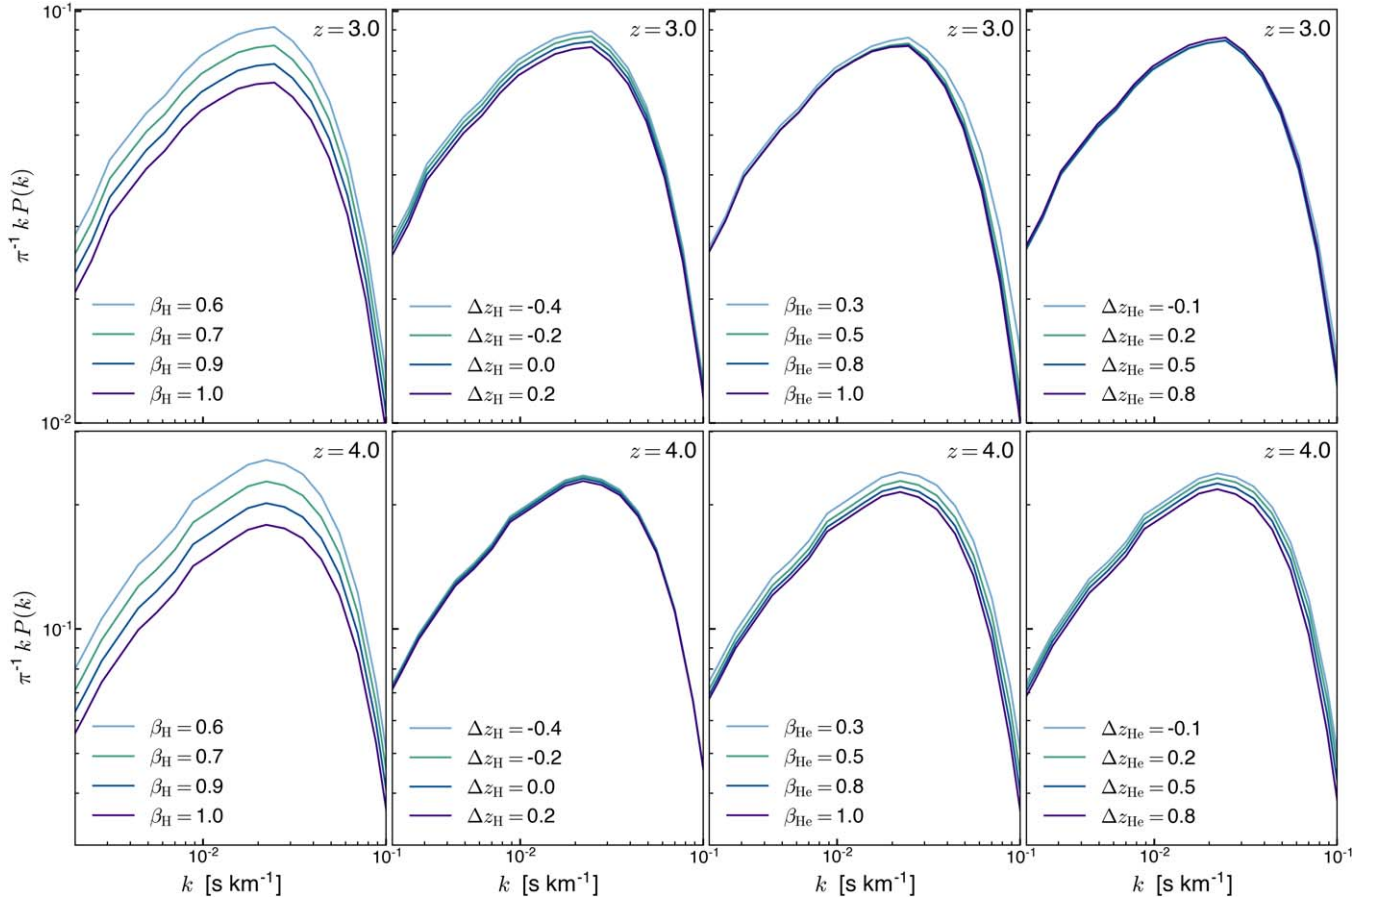


Figure 5. Sensitivity of the Ly α flux power spectrum $P(k)$ to independent variations of the parameters $\theta = \{\beta_{\text{H}}, \Delta z_{\text{H}}, \beta_{\text{He}}, \Delta z_{\text{He}}\}$ for redshifts $z = 3$ (top) and $z = 4$ (bottom). Independent changes of each parameter have different effects on the redshift-dependent $P(k)$. After hydrogen reionization completes, differences in the power spectrum at $z \lesssim 5.5$ arise from changes in the ionization state and temperature of the IGM. Variation of the parameters β_{H} and Δz_{H} mostly affect the ionization state of hydrogen and therefore the overall normalization of $P(k)$. Changes in the parameters β_{He} and Δz_{He} impact $P(k)$ through their effect on the temperature of the gas during and after helium reionization, as variations in the thermal state of the IGM control the ionization fraction of hydrogen by its effect on the recombination rate $\alpha_{\text{HI}}(T)$ and lead to the Doppler broadening of absorption lines and the smoothing of density fluctuations that suppress small-scale power ($k \gtrsim 0.02 \text{ s km}^{-1}$).

Figure 6 only for comparison with our modeling and other data sets.

To obtain a better determination of the He II photoionization and photoheating rates, we complement the power spectrum comparison with observational measurements of the helium effective optical depth $\tau_{\text{eff,HeII}}$ (Worseck et al. 2019) over the redshift range $2.4 \lesssim z \lesssim 2.9$ as additional constraints on our model. The data are shown in Figure 7, along with the corresponding evolution of $\tau_{\text{eff,HeII}}$ from our simulation grid and the best-fit model from our analysis. We do not include the observational lower limits at $z > 3$ as constraints in our MCMC analysis, but our best-fit model is consistent with those limits.

2.8. Systematic Uncertainties

When comparing models to observations, we include systematic uncertainties due to cosmological parameter variations and possible resolution limitations of the simulations. In Villasenor et al. (2021), we performed a study of the changes in the Ly α flux power spectrum $P(k)$ induced by small variations of the cosmological parameters within the constraints from Planck Collaboration et al. (2020). Our results suggested that uncertainties in the cosmological parameters could cause a fractional change of $\lesssim 5\%$ on the hydrogen effective optical depth in the redshift range $2 \lesssim z \lesssim 5$ and a similar $\lesssim 5\%$ effect in $P(k)$ for scales $0.002 \text{ s km}^{-1} \lesssim k \lesssim 0.2 \text{ s km}^{-1}$ and redshifts $2 \lesssim z \lesssim 5$.

For this reason, we include here an additional systematic uncertainty σ_{cosmo} of 5% to the observational determinations of the Ly α power spectrum. For the He II effective optical depth, we estimate similar variations of $\sim 5\%$ at $2 \lesssim z \lesssim 3$ from differences in the mean baryonic density of different cosmologies. We therefore include a $\sigma_{\text{cosmo}} = 5\%$ for the measurements of $\tau_{\text{eff,HeII}}$ as well.

In Appendix A we present a resolution convergence study where we compare the forest flux power spectrum from three simulations performed with the same cosmological parameters and photoionization and photoheating histories. The initial conditions used for the runs were generated to preserve the large-scale modes in common to each simulation, such that the properties of the simulations could be compared directly on shared spatial scales. The three simulations model a box of size $L = 50 h^{-1} \text{ Mpc}$ and $N = 512^3$, $N = 1024^3$, or $N = 2048^3$ cells and particles. The corresponding spatial resolutions are $\Delta x \simeq 98 h^{-1} \text{ Mpc}$, $\Delta x \simeq 49 h^{-1} \text{ Mpc}$, and $\Delta x \simeq 24 h^{-1} \text{ Mpc}$, respectively. In comparing the moderate-resolution ($\Delta x \simeq 49 h^{-1} \text{ Mpc}$) and high-resolution ($\Delta x \simeq 24 h^{-1} \text{ Mpc}$) simulations, we measure small fractional differences $\Delta P(z, k)/P(z, k)$ of $\lesssim 5\%$ for the large scales $k \lesssim 0.02 \text{ s km}^{-1}$. On small scales, $0.02 \text{ s km}^{-1} \lesssim k \lesssim 0.2 \text{ s km}^{-1}$, the fractional differences are slightly larger ($\lesssim 12\%$).

To approximate resolution effects from the grid of simulations used for our analysis ($N = 1024^3$, $\Delta x \simeq 49 h^{-1} \text{ Mpc}$), we




Figure 6. The transmitted flux power spectrum $P(k)$ from observations by eBOSS (Chabanier et al. 2019), Keck Observatory, and the Very Large Telescope (Iršič et al. 2017a; Boera et al. 2019) used to constrain models of the cosmic photoionization and photoheating history. The best-fit evolution of $P(k)$ marginalized over the posterior distribution of the parameters $\theta = \{\beta_{\text{H}}, \Delta z_{\text{H}}, \beta_{\text{He}}, \Delta z_{\text{He}}\}$ is shown with black curves, along with 95% confidence intervals (shaded bands). The fractional differences from the observations and the best-fit model are shown in the bottom part of each panel. Overall, the best-fit $P(k)$ is in good agreement with the large-scale power spectrum from eBOSS for $2.4 \leq z \leq 4.2$ and with the intermediate-scale data from Iršič et al. (2017a) at $3.0 \leq z \leq 4.2$. Our best-fit results also agree with the measurements from Boera et al. (2019) at $4.2 \leq z \leq 5.0$, showing 10%–30% differences mostly on the smallest scales ($0.1\text{--}0.2 \text{ s km}^{-1}$) and suggesting that the temperature of the IGM at this epoch could be slightly overestimated by the model. We also show the $P(k)$ determinations by Walther et al. (2018) for comparison. Due to discrepancies with the eBOSS results on large scales, we have not included the Walther et al. (2018) data points in our MCMC analysis.

add an additional systematic uncertainty σ_{res} to the observational determinations of the flux power spectrum and the He II effective optical depth. For $P(k)$, the additional uncertainty $\sigma_{\text{res}}(z, k) = \Delta P(z, k)$ is set equal to the difference between $P(k)$ from the $N = 1024^3$ and $N = 2048^3$ reference simulations used for our resolution study. For the He II effective optical depth the impact of resolution is a small increase of $\lesssim 3\%$ from the $N = 1024^3$ box to the $N = 2048^3$ run at $z < 3$; we then add an uncertainty of $\sigma_{\text{res}}(z) = 3\%$ to the estimate of $\tau_{\text{eff,HeII}}$. We note that the systematic errors added to $\tau_{\text{eff,HeII}}$ are significantly

smaller than the observational uncertainties $\sigma_{\text{obs}} \sim 12\% \text{--} 45\%$ of Worseck et al. (2019).

The total uncertainty applied to the observational determinations of $P(k)$ and $\tau_{\text{eff,HeII}}$ is finally given by the quadrature sum of the errors as

$$\sigma_{\text{total}} = \sqrt{\sigma_{\text{obs}}^2 + \sigma_{\text{cosmo}}^2 + \sigma_{\text{res}}^2}, \quad (3)$$

where σ_{obs} is the reported observational uncertainty in the flux power spectrum or helium opacity, respectively.

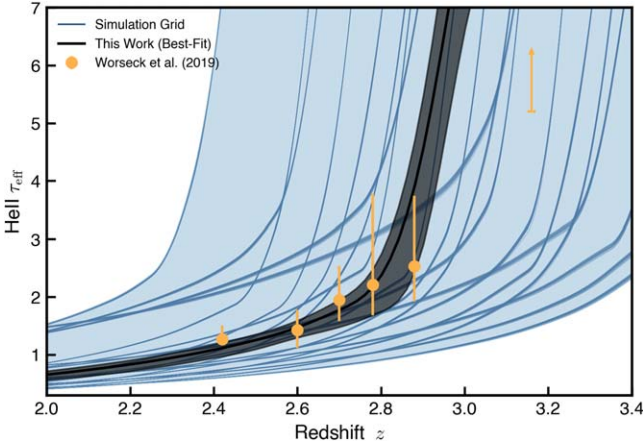


Figure 7. Evolution of the singly ionized helium (He II) effective optical depth $\tau_{\text{eff,HeII}}$ from our simulation grid (blue lines), along with the best-fit model (black line) and the 95% confidence interval (gray area) obtained from our MCMC marginalization. The orange circles show the observational measurements of $\tau_{\text{eff,HeII}}$ (Worseck et al. 2019). While only data in the redshift range $2.4 \lesssim z \lesssim 2.9$ were used as constraints for our statistical analysis, the observed lower limits at $z > 3$ are consistent with the model results.

In their study, Wolfson et al. (2021) showed the importance of using the covariance matrix when inferring the temperature of the IGM from measurements of the Ly α power spectrum and wavelet statistics. For our MCMC analysis we use the covariance matrices of $P(k)$ in the likelihood calculation (see Section 2.9). To reflect the increased uncertainty from Equation (3), we rescale the elements of the covariance matrices according to

$$\mathbf{C}[i, j] = \mathbf{C}_{\text{obs}}[i, j] \frac{\sigma_{\text{total},i} \sigma_{\text{total},j}}{\sigma_{\text{obs},i} \sigma_{\text{obs},j}}, \quad (4)$$

where \mathbf{C}_{obs} is the reported covariance matrix of $P(k)$ taken from the published observational data sets used for our analysis.

2.9. Inference of the UVB Model

To find the photoionization and photoheating rates that best reproduce the properties of the IGM encoded in the flux power spectrum of the Ly α forest $P(k)$ and the helium effective optical depth $\tau_{\text{eff,HeII}}$, we apply an MCMC sampler to compare the simulated $P(k)$ and $\tau_{\text{eff,HeII}}$ to the observational measurements over the redshift and frequency range where data are available. The likelihood function for the model given by the parameters $\theta = \{\beta_{\text{H}}, \Delta z_{\text{H}}, \beta_{\text{He}}, \Delta z_{\text{He}}\}$ is evaluated as

$$\ln \mathcal{L}(\theta) = -\frac{1}{2} \sum_z \left[\left(\frac{\tau_{\text{obs}}(z) - \tau(z|\theta)}{\sigma_{\tau}(z)} \right)^2 + \ln(2\pi\sigma_{\tau}(z)^2) \right] - \frac{1}{2} \sum_{\text{datasets}} \sum_z [\Delta^T \mathbf{C}^{-1} \Delta + \text{Indet}(\mathbf{C}) + N \ln 2\pi], \quad (5)$$

where the first term compares the He II effective optical depth measured from our simulations $\tau(z|\theta)$ for a given photoionization and photoheating model represented by the vector θ to the observational measurement $\tau_{\text{obs}}(z)$ from Worseck et al. (2019) with total (observational + systematic) uncertainty $\sigma_{\tau}(z)$. The second term compares the Ly α power spectrum, with Δ denoting the difference vector between the observations and the model $\Delta = P_{\text{obs}}(z, k) - P(z, k|\theta)$. Here \mathbf{C} corresponds to the

Table 2
Redshift Bin Contribution to the Likelihood

Type	z	$-\Delta \ln \mathcal{L}$	Type	z	$-\Delta \ln \mathcal{L}$
$P(k)$	2.2	330.6	$P(k)$	4.2	489.2
$P(k)$	2.4	363.3	$P(k)$	4.4	135.7
$P(k)$	2.6	229.2	$P(k)$	4.6	190.2
$P(k)$	2.8	297.0	$P(k)$	5.0	40.0
$P(k)$	3.0	215.1	$\tau_{\text{eff,HeII}}$	2.30	0.5
$P(k)$	3.2	134.4	$\tau_{\text{eff,HeII}}$	2.54	0.2
$P(k)$	3.4	113.8	$\tau_{\text{eff,HeII}}$	2.66	0.3
$P(k)$	3.6	84.1	$\tau_{\text{eff,HeII}}$	2.74	1.0
$P(k)$	3.8	137.1	$\tau_{\text{eff,HeII}}$	2.82	2.3
$P(k)$	4.0	180.3			

covariance matrix of size $N \times N$ associated with the observational determination, where N is the number of points of each measurement. To compute $P(z, k|\theta)$ and $\tau(z|\theta)$ for arbitrary values of the parameters θ not directly simulated by our grid, we perform a 4D linear interpolation of the 16 neighboring simulations in parameter space.

As described in Section 2.7, we employ the data sets from Chabanier et al. (2019; $2.2 \leq z \leq 4.6$), Iršič et al. (2017a; $3.0 \leq z \leq 4.2$), and Boera et al. (2019; $4.2 \leq z \leq 5.0$) for the observational measurements of the power spectrum used in our analysis. While there is some overlap in the measurements from the data sets, in general their determinations are consistent with each other. For this reason we include all the data points from each data set for the likelihood calculation. The only significant discrepancy is at $z = 4.6$, where $P(k)$ from Chabanier et al. (2019) is lower than the determination from Boera et al. (2019). We repeated our the analysis excluding the $z = 4.6$ measurement from Chabanier et al. (2019) and obtained similar posterior distributions. We conclude that this difference does not impact our result.

The contribution from each redshift bin to the total log likelihood $\ln \mathcal{L}$ (Equation (5)) from $P(k)$ and $\tau_{\text{eff,HeII}}$ in our analysis is presented in Table 2. The quantity $\Delta \ln \mathcal{L}$ is evaluated as the first and second terms of Equation (5) for $\tau_{\text{eff,HeII}}$ and $P(k)$, respectively, for each redshift bin. The power spectrum most strongly influences the log likelihood, with data from redshifts $z = 2.4$ and $z = 4.2$ inducing the largest fractional changes in the likelihood.

The covariance matrices of $P(k)$ are taken from the published observations. We note that Iršič et al. (2017a) provides the complete covariance of $P(k)$ across the seven redshift bins of their measurement. For this data set we employ the reported full covariance, and the residual vector Δ consists of the $P(k)$ difference from the model and observation concatenated over the seven redshift bins.

While our likelihood analysis uses the reported covariance matrices from the observations, in Appendix C we present the covariance of $P(k)$ measured from a subset of our simulations to quantify the differences induced by variation of our four model parameters. We show that the structure of the covariance is maintained across our simulations, and we measure relatively small variations between the different models.

We emphasize that our approach differs from previous studies of the thermal history of the IGM (e.g., Bolton et al. 2014; Nasir et al. 2016; Hiss et al. 2018; Boera et al. 2019; Walther et al. 2019; Gaikwad et al. 2021) in an important aspect. Typically, the method adopted to infer the thermal state of the IGM from observations of the Ly α forest involves

marginalizing over the thermal parameters T_0 and γ in the approximate power-law density–temperature relation (Hui & Gnedin 1997) $T(\Delta) = T_0 \Delta^{\gamma-1}$, where $\Delta = \rho_{\text{gas}}/\bar{\rho}$ is the gas overdensity. This marginalization is often performed independently for each redshift. Instead, our approach to find the optimal photoionization and photoheating rates that best reproduce the observational measurements is to compare the simulated $P(k)$ and $\tau_{\text{eff,HeII}}$ to the observations over the full redshift range where data are available, namely, $2.2 \leq z \leq 5.0$ for $P(k)$ and $2.2 < z < 3.0$ for $\tau_{\text{eff,HeII}}$.

In our approach, the performance for a given UVB model to match the observations is evaluated over the complete self-consistently evolved reionization and thermal history of the IGM that results from that model. Since the properties of the gas at one redshift cannot be disentangled from its properties at previous epochs, the thermal and ionization structure of the forest depends on the time-dependent photoheating and photoionization rate. Both T_0 and γ evolve along continuous trajectories with redshift, and we therefore marginalize over the full simulated histories of IGM properties.

Our simulations span a wide range of reionization histories for hydrogen in the IGM. Instead of following the common practice of rescaling the optical depth of the simulated skewers in post-processing to match the observed mean transmission of the forest, our method self-consistently follows the ionization evolution of hydrogen and the effective optical depth $\tau_{\text{eff,H}}$ encoded in the redshift-dependent power spectrum of the transmitted flux. Furthermore, during our inference procedure, we do not assume a power-law approximation for the density–temperature distribution of IGM gas or apply a post-processing procedure that artificially modifies the temperature of the gas in the simulations. Instead, our synthetic Ly α spectra reflect the real $\rho_{\text{gas}} - T$ distribution from the simulations. This improvement proves relevant, as we find that a single power law is not a good fit over the full range of gas densities responsible for the bulk of the Ly α absorption signal (see Appendix E).

The posterior distribution for our parameters $\theta = \{\beta_{\text{H}}, \Delta z_{\text{H}}, \beta_{\text{He}}, \Delta z_{\text{He}}\}$ resulting from the Bayesian inference procedure is shown in Figure 8. A clear global maximum of the posterior distribution is observed, and while the posterior shows other local maxima, their likelihoods are significantly lower than the global peak. The four model parameters are well constrained and show only small correlations that arise from the weak degeneracies in the resulting ionization and thermal histories produced by the different photoionization and photoheating rates. Our best-fit parameters and their 95% confidence limits are

$$\begin{aligned} \beta_{\text{H}} &= 0.81_{-0.03}^{+0.04} & \Delta z_{\text{H}} &= -0.09_{-0.24}^{+0.14} \\ \beta_{\text{He}} &= 0.47_{-0.09}^{+0.13} & \Delta z_{\text{He}} &= 0.25_{-0.07}^{+0.09}. \end{aligned} \quad (6)$$

To measure the properties of the IGM that result from our best-fit distribution, we sample $P(k)$, $\tau_{\text{eff,H}}$, and $\tau_{\text{eff,HeII}}$, together with the thermal parameters T_0 and γ , over the posterior distribution of the parameter vector θ , resulting in determinations of the highest-likelihood and 95% confidence interval for the forest statistics and thermal history. When necessary, we interpolate results for values of θ not directly simulated by our grid.

3. Results and Discussion

By comparing the flux power spectrum and the He II effective opacity in our CHIPS simulation grid to observational determinations, we can infer a set of photoionization and

photoheating histories that, when input in cosmological hydrodynamical simulations, result in statistical properties of the Ly α forest that are consistent with observations. In this section, we present the best-fit rates obtained from our inference procedure, as well as the Ly α forest statistics and thermal evolution of the IGM produced by our best-fit UVB model. We compare our results to previous work and finalize our discussion by describing the limitations of our method.

3.1. Best-Fit Photoionization and Photoheating Rates

Figure 9 shows our best-fit model for the photoionization and photoheating rates, along with the corresponding 95% confidence interval that results from our MCMC marginalization of the UVB rates over the posterior distribution of the model parameters obtained from our MCMC analysis. We note that the transformations applied in this work to generate new photoionization and photoheating rates from the reference model (Puchwein et al. 2019) are relatively simple and preserve the functional form of the P19 model. While we allow for orders-of-magnitude variations in the rates, the flexibility of the ionization and thermal histories sampled here is limited by the fixed shape of the UVB model employed in our simulation grid. A study that allows for more flexibility in the photoionization and photoheating rates of hydrogen and helium will be the scope of future work.

3.2. P(k) Model Comparison with the Data

Figure 6 shows the evolution of the best-fit flux power spectrum and 95% confidence intervals over the redshift range $2.2 \leq z \leq 5.0$ that result from marginalizing $P(k)$ over the posterior distribution of model parameters $\theta = \{\beta_{\text{H}}, \Delta z_{\text{H}}, \beta_{\text{He}}, \Delta z_{\text{He}}\}$. Our best-fit synthetic power spectrum shows good agreement with the large-scale $P(k)$ measured by the eBOSS experiment (Chabanier et al. 2019) in the range $2.4 \lesssim z \lesssim 4.2$, suggesting that the mean transmission $\langle F \rangle$ of the forest inferred by our analysis is consistent with the measurements by Chabanier et al. (2019). Only for $z = 2.2$ and $z = 4.4$ – 4.6 do our results show significant differences from the eBOSS data set. At $z = 2.2$, the $P(k)$ from eBOSS is higher than our results by $\sim 8\%$ – 20% on scales $0.008 \text{ s km}^{-1} \lesssim k \lesssim 0.02 \text{ s km}^{-1}$. This modest tension may suggest that the hydrogen opacity $\tau_{\text{eff,H}}$ is underestimated by $\sim 10\%$ in our modeling relative to eBOSS. At $z = 4.4$ and $z = 4.6$ the opposite is true, and our best-fit $P(k)$ on large scales is $\sim 15\%$ and $\sim 25\%$ higher than the eBOSS measurements, respectively. These small discrepancies could be alleviated, e.g., by a small 15% decrease of the H I photoionization rate at $z = 2.2$ and by a comparable small increase in the same quantity at $z = 4.4$ – 4.6 by $\sim 10\%$ – 20% .

Our results also agree on large and intermediate scales ($0.003 \text{ s km}^{-1} \lesssim k \lesssim 0.06 \text{ s km}^{-1}$) with the estimates of Iršič et al. (2017a). The best-fit model reproduces the turnover in the observed dimensionless power spectrum $\Delta^2(k) = \pi^{-1} k P(k)$ at $k \sim 0.02$ – 0.03 s km^{-1} and generally lies within the observational uncertainties at intermediate scales $0.01 \text{ s km}^{-1} \lesssim k \lesssim 0.06 \text{ s km}^{-1}$. Only at redshifts $z = 3.4$ and $z = 3.8$ do the $P(k)$ measurements show some differences relative to the model. At $z = 3.4$ the data are higher than the model by $\sim 5\%$ – 20% . A similar discrepancy is observed when comparing Iršič et al. (2017a) with the determinations by eBOSS at the same redshift, suggestive of a slightly higher H I opacity $\tau_{\text{eff,H}}$ in the former sample. Differences with the model are more significant at $z = 3.8$, where on intermediate scales

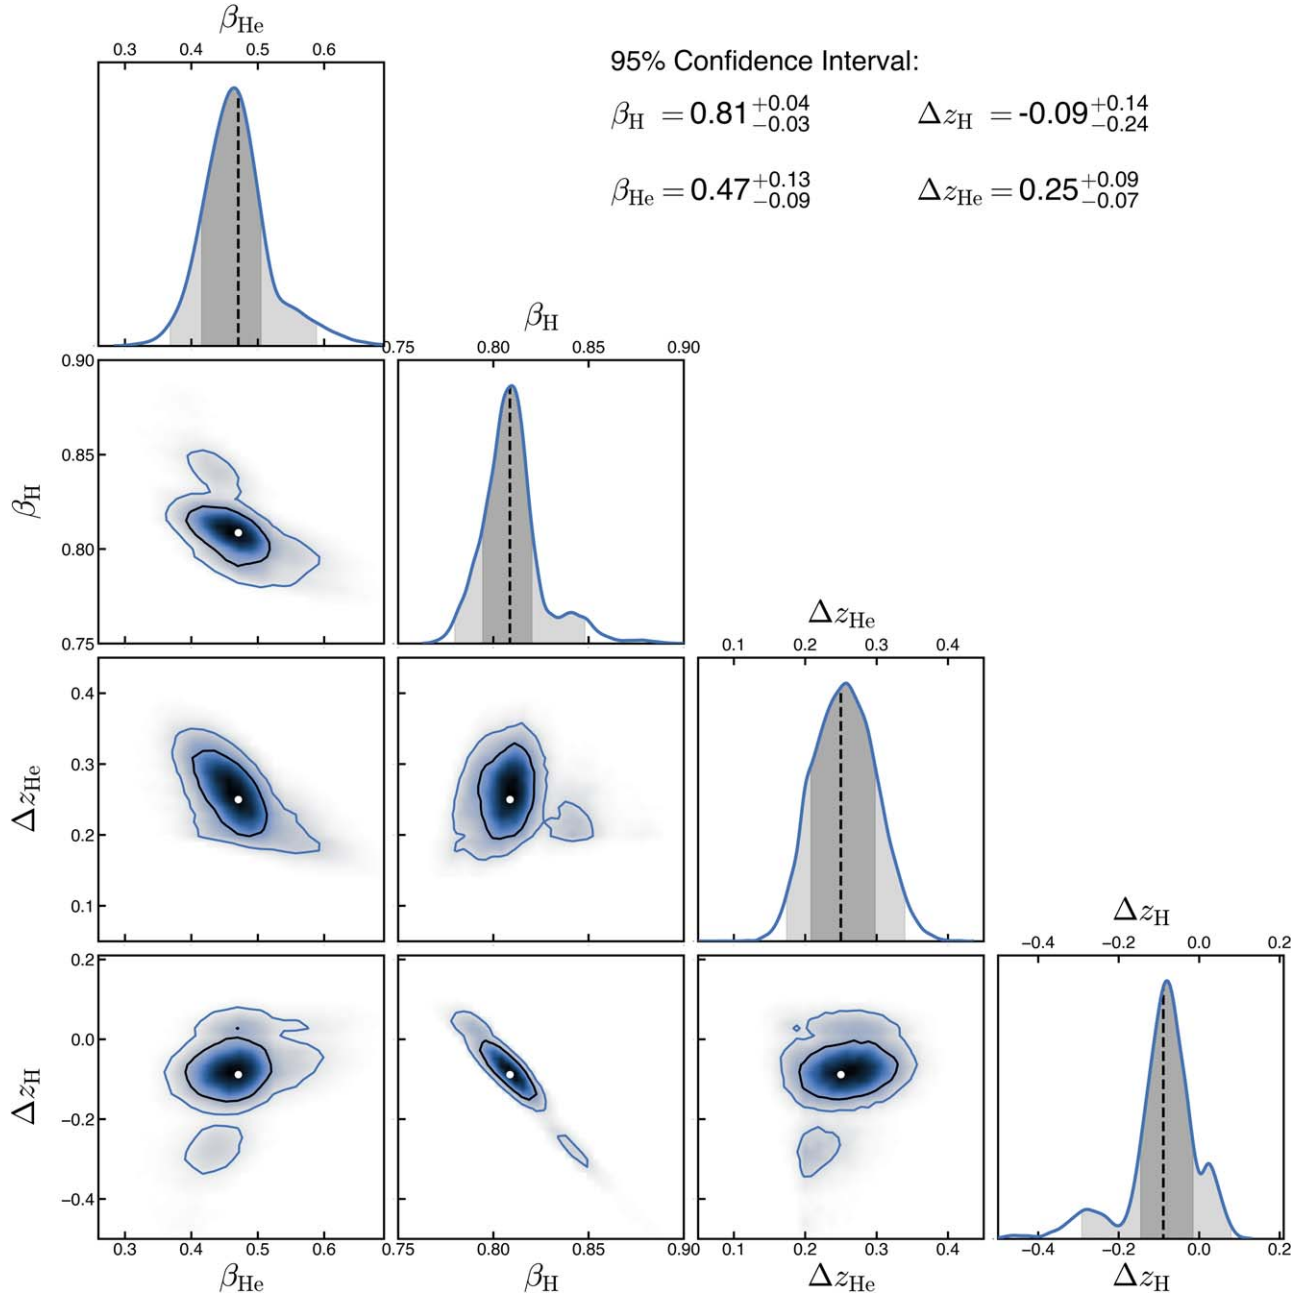


Figure 8. Results from the Bayesian inference procedure, showing 1D and 2D projections of the posterior distributions for the parameters $\theta = \{\beta_{\text{H}}, \Delta z_{\text{H}}, \beta_{\text{He}}, \Delta z_{\text{He}}\}$. The parameter constraints were obtained by fitting the observed flux power spectrum of the Ly α forest and the He II effective optical depth (Iršič et al. 2017a; Boera et al. 2019; Chabanier et al. 2019; Worseck et al. 2019) with a grid of CHIPS simulations. The posterior distribution shows a clear global maximum, and while other local maxima are present, their peak likelihoods are significantly lower than the global maximum. The resulting best-fit parameters and their 95% confidence intervals are shown in the upper right corner.

($k \gtrsim 0.2 \text{ s km}^{-1}$) the measurements of Iršič et al. (2017a) are lower than the model by $\sim 10\%$ – 20% , while on large scales ($k \lesssim 0.2 \text{ s km}^{-1}$) their estimates are higher than both the model and the determinations by eBOSS by $\sim 5\%$ – 30% .

Our model is in good agreement with the high-redshift measurements of $P(k)$ by Boera et al. (2019), with minor differences that could be addressed by small modifications to the early photoheating history. At $z = 4.2$, $z = 4.6$, and $z = 5.0$, our best-fit $P(k)$ is consistent with their data points on large scales $k \lesssim 0.02 \text{ s km}^{-1}$, suggesting that our inferred IGM HI opacity matches that measured by Boera et al. (2019). The model also reproduces the cutoff in $\Delta^2(k)$ at $k \sim 0.02$ – 0.03 s km^{-1} , and the

consistency with the observations extends to small scales $k \lesssim 0.1 \text{ s km}^{-1}$. Discrepancies appear only on the smallest scales $0.1 \text{ s km}^{-1} \lesssim k \lesssim 0.2 \text{ s km}^{-1}$, where the model has less power ($\sim 10\%$ – 30%) than that of Boera et al. (2019). This may suggest that the temperature of the IGM has been overestimated by the model in the redshift range $4 \lesssim z \lesssim 5$ (see Section 3.3 for a discussion of this issue).

3.3. Evolution of the IGM Temperature

The flux power spectrum and helium opacity tightly constrain the time-dependent photoionization and photoheating rates, which in turn determine the IGM ionization and thermal

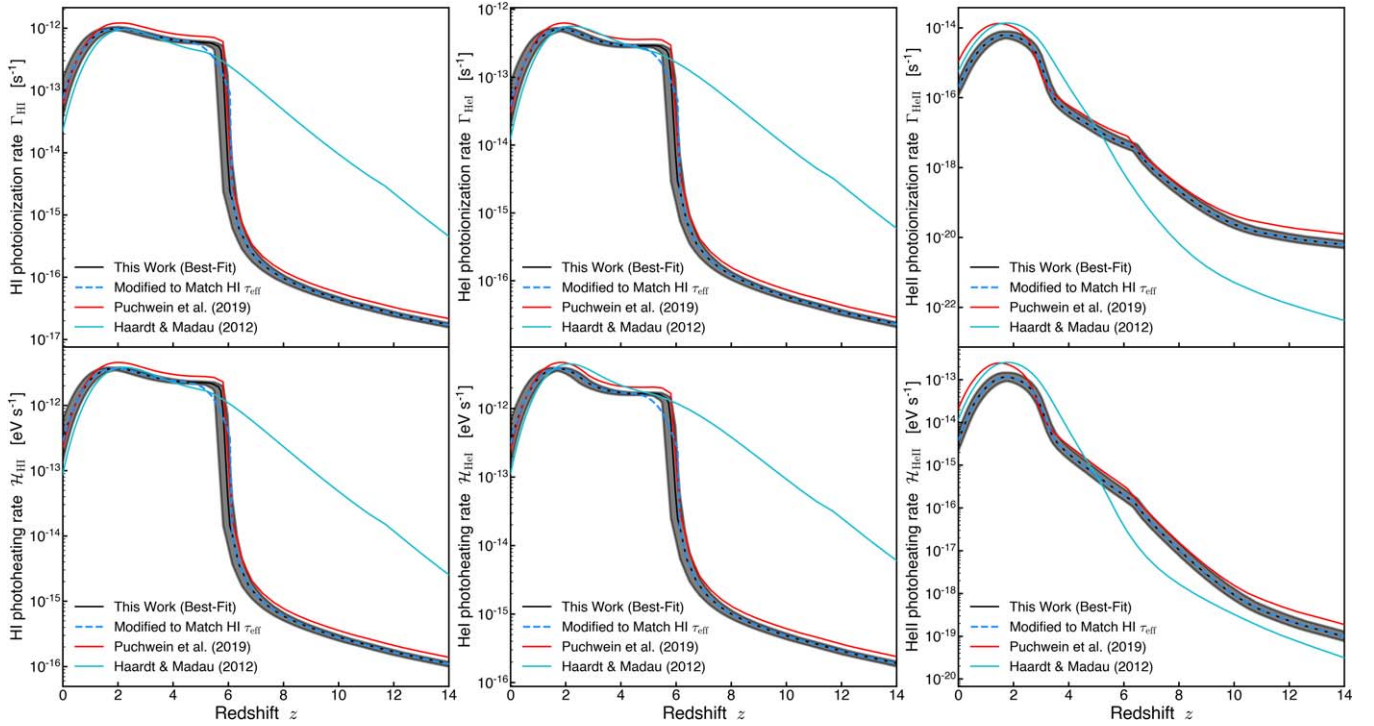


Figure 9. Best-fit (black lines) and 95% confidence intervals (gray bands) for the photoionization (Γ ; top) and photoheating (\mathcal{H} ; bottom) rates for neutral hydrogen (H I; left), neutral helium (He I; middle), and singly ionized helium (He II; right) obtained from our MCMC analysis. The *modified* H I and He I photoionization and photoheating rates (dashed blue lines) are identical to the reference best-fit model except for the redshift range $4.8 \leq z \leq 6.1$, where they have been modified to produce an evolution of the hydrogen effective optical depth consistent with the observational determinations of Bosman et al. (2018) for $z > 5$ (see Sections 3.4 and 3.5 for details). For reference, we also show the models from Puchwein et al. (2019; red) and Haardt & Madau (2012; cyan).




Figure 10. Redshift evolution of the gas temperature from a high-resolution simulation ($L = 50 h^{-1}$ Mpc, $N = 2048^3$ cells and particles) that employed our best-fit model for the photoheating and photoionization rates. The image displays the monotonic increase in the temperature of the IGM due to hydrogen reionization for $z \gtrsim 6.0$, followed by an epoch of cooling of the IGM due to cosmic expansion. The onset of helium reionization ($z \sim 4.5$) initiates a second epoch of heating of the IGM that ends at $z \sim 3$ when He II reionization completes. A second epoch of cooling due to cosmic expansion then follows. The temperature increase of gas collapsing into the filamentary cosmic web as large-scale structure develops is also visible in the image.

history. The redshift evolution of the gas temperature is illustrated in Figure 10, which is generated from a slice through a high-resolution simulation ($L = 50 h^{-1}$ Mpc, $N = 2048^3$ cells and particles) using our best-fit photoionization and photoheating rates. The figure shows the monotonic increase in the temperature of the IGM during hydrogen reionization at $z \gtrsim 6$. After hydrogen reionization completes by $z \sim 6$, the input of energy into the IGM falls dramatically, and the gas then cools primarily through adiabatic expansion. This first epoch of IGM cooling lasts until the onset of helium reionization ($z \sim 4.5$) when extreme-UV radiation from AGNs ionizes He II atoms

and drives a second epoch of reheating that completes by $z \sim 3$ and is followed by a second epoch of adiabatic cooling.

The thermal state of diffuse IGM gas is often modeled with the power-law relation (Hui & Gnedin 1997; Puchwein et al. 2015; McQuinn 2016)

$$T(\Delta) = T_0 \Delta^{\gamma-1}. \quad (7)$$

We fit the power-law relation to the gas density–temperature distribution in each of the simulations from the CHIPS grid and at multiple epochs, $2 \leq z \leq 9$, following the procedure presented in Villasenor et al. (2021). We restrict the fit to the

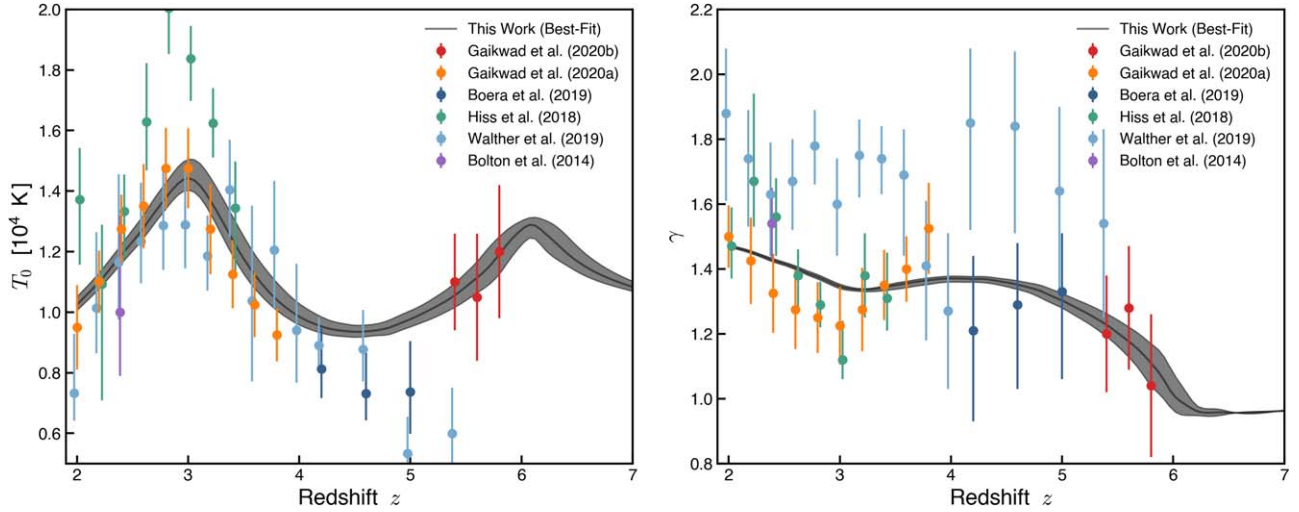


Figure 11. Redshift evolution of the parameters T_0 and γ (Equation (7)) from the best-fit model (black lines) and 95% confidence interval (gray band) obtained from our MCMC analysis. The data points show the values of T_0 and γ inferred from observations of the Ly α forest by Bolton et al. (2014), Hiss et al. (2018), Boera et al. (2019), Walther et al. (2019), and Gaikwad et al. (2020, 2021). Our results reveal two peaks in the evolution of T_0 due to hydrogen reionization at $z \sim 6$ and helium reionization at $z \sim 3$ and are consistent with previous measurements from Gaikwad et al. (2020, 2021).

overdensity range $0 \leq \log_{10} \Delta \leq 1$, as we find that in our simulations a single power law does not accurately describe the wider range $-1 \leq \log_{10} \Delta \leq 1$ (see Appendix E).

Figure 11 shows the redshift evolution of the parameters T_0 and γ from our best-fit model and the 95% confidence interval that results from our MCMC marginalization over the posterior distribution of the photoionization and photoheating rates. For comparison, we also depict the data points for these parameters inferred from the properties of the Ly α forest by Bolton et al. (2014), Hiss et al. (2018), Boera et al. (2019), Walther et al. (2019), Gaikwad et al. (2020), and Gaikwad et al. (2021).

The inferences from Boera et al. (2019) and Walther et al. (2019) follow similar methodologies. They generate flux power spectra from simulations run with different thermal histories, resulting in multiple trajectories for the evolution of T_0 and γ . For each redshift bin they determine the best-fit T_0 , γ , and mean transmitted flux $\langle F \rangle$ by performing Bayesian inference and comparing the simulated flux power spectra to observations of the Ly α forest $P(k)$. Bolton et al. (2014) and Hiss et al. (2018) measure a set of values for the Doppler parameter b and HI column density N_{HI} directly from the forest by decomposing the absorption spectra into a collection of Voigt profiles. They infer the parameters T_0 and γ by comparing simulations with different $b - N_{\text{HI}}$ distributions to the observed one. Gaikwad et al. (2020) follow a similar approach by comparing simulated Ly α forest spectra to Voigt profiles fitted to the observed transmission spikes in the inverse transmitted flux $1 - F$ at $z > 5$. Gaikwad et al. (2021) report more precise determinations by inferring T_0 and γ from the combined constraints obtained through a comparison of simulated Ly α forest absorption with the observed flux power spectra, $b - N_{\text{HI}}$ distributions, wavelet statistics, and curvature statistics.

As shown in Figure 11, the temperature evolution from our best-fit model presents a first peak ($T_0 \simeq 1.3 \times 10^4$ K) at the end of hydrogen reionization ($z \sim 6.0$), followed by an epoch of adiabatic cooling from cosmic expansion. Our results agree well with the high-redshift measurements of T_0 and γ at $5.4 \leq z \leq 5.8$ from Gaikwad et al. (2020). We note that their estimates also suggest a period of cooling at these epochs, and

from their result it is possible to infer a peak in T_0 from H reionization sometime at redshift $z \gtrsim 5.8$.

In our model, the IGM continues to cool until the onset of helium reionization, and the temperature reaches a local minimum of $T_0(z \sim 4.5) \simeq 9.5 \times 10^3$ K. Evidence of this transition can also be seen in the measurements from Boera et al. (2019), where T_0 shows little evolution from $z = 5.0$ to $z = 4.6$ and then a slight increase to $z = 4.2$. Nevertheless, there are significant differences between T_0 from the model at $4 \lesssim z \lesssim 5$ and the measurements from Boera et al. (2019), as the temperature predicted by our model is higher than their inferred values of $T_0 \sim 7.4 \times 10^3$ K and $T_0 \sim 8.1 \times 10^3$ K at $z = 4.6 - 5$ and $z = 4.2$, respectively. The higher temperatures in our model reflect a suppressed power spectrum of the Ly α flux on small scales ($0.1 \text{ s km}^{-1} \lesssim k \lesssim 0.2 \text{ s km}^{-1}$) compared to the $P(k)$ measurement from Boera et al. (2019) at $4.2 \leq z \leq 5.0$ (see Figure 6). Decreasing the photoheating from the UVB during $z \gtrsim 4$ would decrease the temperature of the IGM at this epoch and potentially alleviate this discrepancy.

In Appendix D we present scenarios where the intermediate-redshift IGM is set to be colder compared to our model by decreasing the best-fit HI and HeI photoheating rates at $4.2 \leq z \leq 6.2$. We find that reducing \mathcal{H}_{HI} and \mathcal{H}_{HeI} by $\sim 80\%$ at $z \sim 6$ decreases the IGM temperature T_0 by $\sim 20\%$, making it consistent with the estimates from Boera et al. (2019) at $4.2 \leq z \leq 5.0$ with minimal impact in T_0 at $z \lesssim 3.5$ (see Figure 21). Nevertheless, we find that such colder evolution of T_0 is in conflict with the $z \sim 5.4$ estimate from Gaikwad et al. (2020) (see Figure 21). This conflict indicates some degree of tension between the higher $T_0 = (1.10 \pm 0.16) \times 10^4$ K at $z \sim 5.4$ from Gaikwad et al. (2020) and the low $T_0 = 7.37_{-1.39}^{+1.13} \times 10^3$ K at $z \sim 5.0$ from Boera et al. (2019).

After $z \sim 4.5$, radiation from AGN ionizes HeII atoms in the universe and heats the IGM for a second time. Our model predicts that T_0 increases monotonically until HeII reionization completes at $z \sim 3$, resulting in a second peak in the temperature ($T_0 \simeq 1.4 \times 10^4$ K) followed by a second epoch of cooling due to cosmic expansion. Our results for the evolution of T_0 during $z \lesssim 4.5$ are consistent with the determinations from Gaikwad et al. (2021) and Walther et al. (2019) that show a similar T_0 history

within the uncertainties during and after He II reionization, as both show a peak in T_0 at $z \sim 2.8\text{--}3.0$. Our $T_0(z)$ results are higher yet consistent within the uncertainties from the measurement by Bolton et al. (2014) at $z = 2.4$. The results presented by Hiss et al. (2018) also show the effects of He II reionization on the temperature of the IGM in the form of a peak in the temperature at $z \sim 2.8$, but their peak value of $T_0 \sim 2 \times 10^4$ K is significantly higher than our result and the measurements from Gaikwad et al. (2020) and Walther et al. (2019).

The right panel of Figure 11 shows our result for the evolution of the density–temperature power-law index γ (black line and shaded 95% confidence interval). At the end of hydrogen reionization, the gas in the IGM is mostly isothermal ($\gamma \sim 1$). As the IGM cools and the low-density gas cools more efficiently, the index γ increases in the interval $4.5 \lesssim z \lesssim 6$. During the reheating of the IGM from He II reionization, low-density gas heats faster and γ decreases until helium reionization completes. After helium reionization, cooling from cosmic expansion causes an increase on γ for a second time.

The evolution of the power-law index in our model is consistent with measurements from Hiss et al. (2018), Boera et al. (2019), Gaikwad et al. (2020), and Gaikwad et al. (2021) and shows deviations only for a few redshift bins after He II reionization completes. The transition in γ after He II reionization in our model is not as pronounced as the determinations from Gaikwad et al. (2021) and Hiss et al. (2018).

The results from Walther et al. (2019) show significantly higher values of γ compared to all the other measurements. We have evaluated the plausibility of a steep density–temperature relation ($\gamma > 1.6$) by simulating the extreme case in which all photoheating and photoionization from the UVB stop after hydrogen reionization completes, i.e., $\Gamma = 0$ and $\mathcal{H} = 0$ for $z > 6$. We find that in the absence of external heating, as the IGM cools by adiabatic expansion, the overdensities cool down at a slower rate from compression by gravitational collapse. Here γ tends to increase with decreasing redshift at a roughly constant rate of $\Delta\gamma/|\Delta z| \sim 0.18$. Starting from an isothermal distribution of the gas in the IGM when H reionization finishes ($\gamma = 1$), it takes a change in redshift $|\Delta z| \sim 3\text{--}3.5$ for the gas distribution to steepen to $\gamma \sim 1.6$. Hence, we can reproduce values of $\gamma > 1.6$ at $z \sim 5$ only if hydrogen reionization completes very early at $z > 8$.

3.4. Evolution of the Hydrogen Effective Optical Depth

The H I effective optical depth $\tau_{\text{eff,H}} = -\ln(F)$ measured from the Ly α forest reflects the overall H I content of the gas in the IGM. Hence, $\tau_{\text{eff,H}}$ probes the ionization state of hydrogen in the medium and can be used to constrain the intensity of the ionizing UVB. In our work, constraints obtained for the H I photoionization rate Γ_{HI} derive from the power spectrum of the Ly α transmitted flux itself, as we do not include the observational determinations of $\tau_{\text{eff,H}}$ as constraints in our inference procedure.

The power spectrum $P(k)$ of the flux fluctuations (Equation (1)) is itself sensitive to the hydrogen effective optical depth. Because of the nonlinear relation $F = \exp(-\tau)$, the normalization of $P(k)$ on most scales relevant to this work ($0.002 \text{ s km}^{-1} \lesssim k \lesssim 0.1 \text{ s km}^{-1}$) is affected by the value of $\tau_{\text{eff,H}}$ obtained from the skewer sample used for the measurement. Thus, including the effective optical depth of the forest does not provide additional independent information for




Figure 12. Redshift evolution of the hydrogen effective optical depth $\tau_{\text{eff,H}}$ from our best-fit determination of the photoheating and photoionization rates (black line) and the corresponding 95% confidence interval. Data points show the observational measurements of $\tau_{\text{eff,H}}$ from Fan et al. (2006), Becker et al. (2013), Bosman et al. (2018), Eilers et al. (2018), Boera et al. (2019), and Yang et al. (2020a). The model results show consistency with the measurement from Becker et al. (2013; yellow) for $2.5 \lesssim z \lesssim 4.2$ and are in good agreement with the determination from Boera et al. (2019; green) for $4.2 \lesssim z \lesssim 5.0$. At high redshift ($z > 5$) the results from Yang et al. (2020a) lie significantly higher than those from Eilers et al. (2018) and Bosman et al. (2018) by $\sim 10\%$ – 30% . In the redshift range $5 \lesssim z \lesssim 5.8$, the model shows lower $\tau_{\text{eff,H}}$ compared with the observations. By modifying the best-fit H I photoionization rate Γ_{HI} as shown in Section 3.7, we can obtain a high- z evolution of $\tau_{\text{eff,H}}$ (dashed blue) consistent with the measurement from Bosman et al. (2018) and Fan et al. (2006).

constraining the model. See Appendix B for a discussion on the impact that H I $\tau_{\text{eff,H}}$ has on the Ly α flux power spectrum.

Figure 12 shows the redshift dependence of $\tau_{\text{eff,H}}$ from our best-fit determination of the photoheating and photoionization rates (black line) and the corresponding 95% confidence interval. Data points in the figure show the observational measurements of $\tau_{\text{eff,H}}$ reported by Fan et al. (2006), Becker et al. (2013), Bosman et al. (2018), Eilers et al. (2018), Boera et al. (2019), and Yang et al. (2020a). Our results are consistent with the evolution of H I $\tau_{\text{eff,H}}$ measured by Becker et al. (2013; yellow circles) for the redshift range $2.5 \lesssim z \lesssim 4.2$. Our model results in a more opaque IGM compared to their measurements at lower redshifts $2.2 \lesssim z \lesssim 2.5$ and higher redshifts $4.2 \lesssim z \lesssim 4.8$. Our model agrees well with the determination from Boera et al. (2019; green circles) during the redshift range $4.2 \lesssim z \lesssim 5.0$.

At high redshift ($z > 5$), the measurements of the H I effective optical depth from Bosman et al. (2018; red circles), Eilers et al. (2018; cyan circles), and Fan et al. (2006; orange circles) are similar, with only small differences ($< 12\%$) toward higher $\tau_{\text{eff,H}}$ from Eilers et al. (2018) compared with Bosman et al. (2018). The measurements by Yang et al. (2020a; purple circles) suggest a more opaque IGM with a $\tau_{\text{eff,H}}$ that is significantly higher ($\sim 20\%$ – 30%) compared to the measurements by Bosman et al. (2018).

Shortly after hydrogen reionization completes ($5 \lesssim z \lesssim 5.8$), our best-fit UVB model significantly underestimates $\tau_{\text{eff,H}}$ compared with the observational measurements, suggesting that the hydrogen in the IGM is overly ionized in our model at these redshifts. To address this possible discrepancy, we can modify our best-fit result for the H I photoionization rate such that Γ_{HI} is reduced only in the redshift range $4.8 < z < 5.8$ and

increased for $5.8 < z < 6.1$ (see Sections 3.5 and 3.7). As shown in Figure 12, the high-redshift evolution ($z > 5$) of $\tau_{\text{eff,H}}$ from the modified model (dashed blue line) is consistent with the measurements from Bosman et al. (2018). The subsequent evolution at redshifts $z < 4.8$ remains virtually unchanged from the best-fit model, as hydrogen is in photoionization equilibrium at these times and the ionization fraction is therefore determined by the instantaneous amplitude of the HI photoionization rate Γ_{HI} . We refer the reader to Section 3.7 for a discussion on the effect that the modified UVB model has on the properties of the gas in the IGM.

By providing a simple modification to our best-fit UVB model that allows us to change the high-redshift evolution of the hydrogen effective optical depth to achieve consistency with the observation and with minimal impact on the subsequent evolution of the properties of the IGM for $z \lesssim 5.0$, we show that the high- z discrepancy of the observed $\tau_{\text{eff,H}}$ and the model is not a significant challenge to our results and the conclusions of this work.

3.5. Hydrogen Photoionization Rate

Our best-fit model results for the hydrogen photoionization rate Γ_{HI} provide several opportunities for comparisons with observations, even though observationally inferred Γ_{HI} measurements are not used to constrain our model. There are observational determinations of Γ_{HI} informed by simulations where the photoionization rate is rescaled to match the observational (F) (Becker & Bolton 2013; D’Aloisio et al. 2018). Our results can also be compared to estimates of Γ_{HI} from the quasar proximity effect and the size of the near zone of high Ly α transmission around quasars (Calverley et al. 2011; Wyithe & Bolton 2011). Observations have measured Γ_{HI} by detecting the fluorescent Ly α emission produced by the Lyman limit systems (LLSs) illuminated by background radiation (Gallego et al. 2021). Finally, there are Γ_{HI} determinations from combining the probability density function and power spectrum of the Ly α transmitted flux from observations with simulations that apply different photoionization rates Γ_{HI} (Gaikwad et al. 2017).

Figure 13 shows our result for the HI photoionization rate with the corresponding 95% confidence limits (black line and shaded band), along with the observational inferences of Γ_{HI} mentioned above. Our result is consistent with the previous observational determinations that show a rapid evolution in Γ_{HI} for $z \gtrsim 5.6$, followed by a gradual increase during $2 \lesssim z \lesssim 5.6$ and a rapid decrease at $z < 2$. The only visible differences with Becker & Bolton (2013) occur in the redshift range $4 \lesssim z \lesssim 4.8$. Their measurement was obtained by tuning the photoionization rate Γ_{HI} in simulations such that the Ly α effective optical depth $\tau_{\text{eff,H}}$ was consistent with the observational measurement from Becker et al. (2013). The higher estimate of Γ_{HI} from their result reflects the lower $\tau_{\text{eff,H}}$ from Becker et al. (2013) compared with the evolution of $\tau_{\text{eff,H}}$ from our model for the redshift range $4.2 \lesssim z \lesssim 4.8$, as shown in Figure 12.

As described in Section 3.4, shortly after hydrogen reionization completes, our best-fit model significantly underestimates the Ly α effective optical depth $\tau_{\text{eff,H}}$ compared with the observations in the redshift range $5 \lesssim z \lesssim 5.8$. To address this discrepancy, we presented an alternative model where the sharp transition in Γ_{HI} at $z \sim 5.6$ from the original best-fit model is replaced by a softer increase that extends over the redshift range $4.8 < z < 5.8$ (dashed blue line in Figure 13). Decreasing Γ_{HI} during this epoch increases the neutral fraction

of hydrogen in the IGM in photoionization equilibrium, thereby increasing $\tau_{\text{eff,H}}$. Our modified model for Γ_{HI} was chosen such that the resulting evolution of $\tau_{\text{eff,H}}$ is consistent with the observational measurement presented by Bosman et al. (2018; dashed blue line in Figure 12), and the altered transition of Γ_{HI} from our modified model is still within the uncertainties of the observational inference by D’Aloisio et al. (2018) in the redshift interval $4.8 \lesssim z \lesssim 5.8$.

3.6. Ionization History

We present the redshift evolution of the volume-weighted neutral fraction of hydrogen x_{HI} resulting from our best-fit determination of the UVB model and the corresponding 95% confidence limits (black line and shaded band) in Figure 14. For comparison we show several observational estimates. We show constraints from the optical depth of the Ly α , Ly β , and Ly γ transitions in the forest (Fan et al. 2006). We also show constraints on the IGM neutrality from properties of Ly α emission from galaxies at high redshift (Mason et al. 2018; Hoag et al. 2019; Mason et al. 2019) and the damping wing absorption in the spectra of $z \gtrsim 7$ quasars (Greig et al. 2017, 2019; Yang et al. 2020b; Jung et al. 2020; Wang et al. 2020). Finally, we show constraints from the covering fraction of dark pixels in the Ly α /Ly β forest of high- z quasars (McGreer et al. 2011, 2015).

Our model results in a prolonged hydrogen reionization history, extending from $x_{\text{HI}} \sim 0.9$ at $z \sim 11$ to $x_{\text{HI}} \sim 0.1$ at $z \sim 6.5$. The duration results in part from the gradually increasing ionization rate $\Gamma_{\text{HI}} < 1 \times 10^{-15} \text{ s}^{-1}$ at $z > 6.5$ associated with radiation emitted by early star-forming galaxies.

For $7 \lesssim z \lesssim 8$, the observational estimates display a wide range of x_{HI} , from a highly ionized ($x_{\text{HI}} \sim 0.8$) to a mostly neutral ($x_{\text{HI}} \sim 0.2$) IGM. Our model lies within this range, and at $z = 7$ our result is in agreement with the $x_{\text{HI}} \sim 0.4$ estimates from Greig et al. (2017) and Yang et al. (2020b), as well as with the $x_{\text{HI}} \sim 0.5$ estimate from Jung et al. (2020) at $z \sim 7.6$.

The redshift at which hydrogen reionization completes z_R , defined as the redshift at which $x_{\text{HI}} \leq 1 \times 10^{-3}$ for the first time, is $z \sim 6.0$ for our best-fit model. After hydrogen reionization completes, our best-fit model results in an ionization fraction that falls below the estimate from Fan et al. (2006; reflected by the lower optical depth τ_{eff} in Figure 12). Nevertheless, our modified model (dashed blue line) shows better consistency with their estimate.

Later in cosmic history, high-energy radiation emitted by AGNs leads to the ionization of singly ionized helium (He II). For our best-fit model He II reionization starts at $z \sim 5$ and completes at $z \sim 3.0$, when the He II fraction reaches $x_{\text{HeII}} \leq 1 \times 10^{-3}$ for the first time. As the He II effective optical depth from our model is consistent with the observation from Worseck et al. (2019) for $2.4 \lesssim z \lesssim 2.9$, we argue that the end of He II reionization by $z \sim 2.9$ is suggested by their measurement.

Thomson scattering of the cosmic microwave background (CMB) by the free electrons in the IGM provides another diagnostic of the reionization history of the IGM. From the evolution of the electron density n_e given by the ionization state of hydrogen and helium from our models, we can compute the electron-scattering optical depth τ_e as

$$\tau_e(z) = \int_0^z \frac{c\sigma_T n_e(z)}{(1+z)H(z)} dz, \quad (8)$$

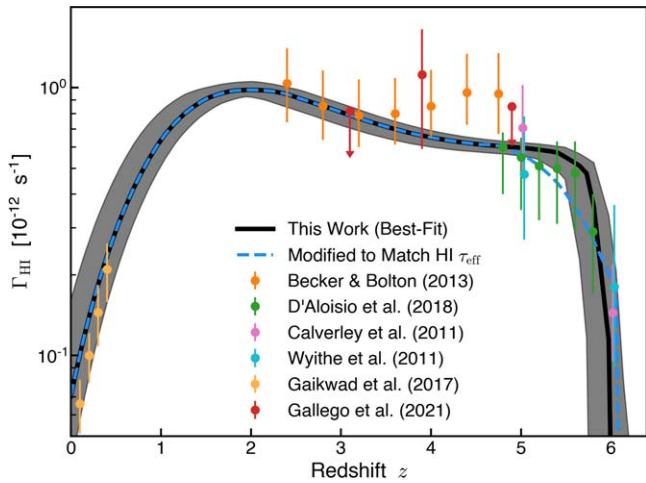


Figure 13. Evolution of the hydrogen photoionization rate Γ_{HI} from our best-fit determination and the 95% confidence interval (black line and shaded region). Data show observationally inferred photoionization rates measured by Calverley et al. (2011), Wyithe & Bolton (2011), Becker & Bolton (2013), Gaikwad et al. (2017), D’Aloisio et al. (2018), and Gallego et al. (2021). A modified model for Γ_{HI} designed to match the observational measurements of $\tau_{\text{eff,H}}$ from Bosman et al. (2018; see Figure 12) is shown as the dashed blue line. Our models agree well with the observationally inferred results, except for visible differences with the estimate from Becker & Bolton (2013) during $4 \lesssim z \lesssim 5$. These differences in Γ_{HI} reflect small differences between our best-fit model predictions for $\tau_{\text{eff,H}}$ and the observational $\tau_{\text{eff,H}}$ measurement by Becker et al. (2013) over this redshift range.

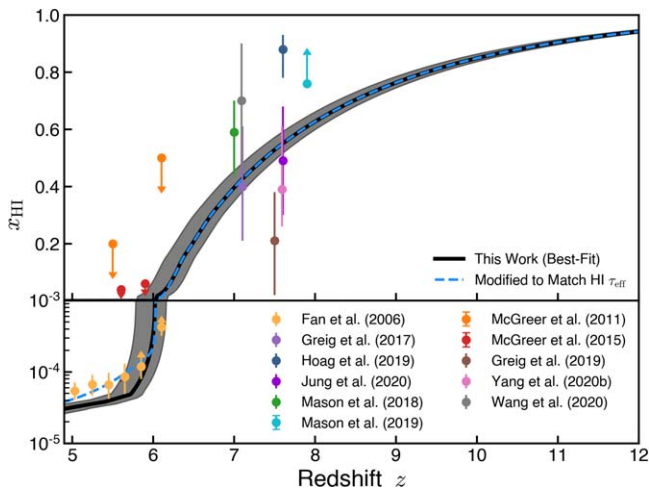


Figure 14. Redshift evolution of the volume-weighted neutral fraction of hydrogen for our best-fit model and the corresponding 95% confidence interval (black line and shaded region). Data points show the observational estimates reported in Fan et al. (2006), McGreer et al. (2011), McGreer et al. (2015), Greig et al. (2017, 2019), Mason et al. (2018, 2019), Hoag et al. (2019), Jung et al. (2020), Yang et al. (2020b), and Wang et al. (2020). For $z \gtrsim 7$ the observational estimates show a wide range of x_{HI} , from $x_{\text{HI}} \sim 0.2$ to $x_{\text{HI}} \sim 0.8$. Our models result in a $z \sim 7$ –8 neutral fraction of $x_{\text{HI}} \sim 0.4$ –0.5, consistent with the results from Greig et al. (2017), Jung et al. (2020), and Yang et al. (2020b). After hydrogen reionization completes at $z \lesssim 6.0$, our best-fit model shows an evolution of x_{HI} below the measurement by Fan et al. (2006). By modifying our best-fit photoionization rates to better match $\tau_{\text{eff,H}}$ (see Figure 12), we can also better match the x_{HI} data from Fan et al. (2006; dashed blue line).

where σ_{T} represents the Thomson scattering cross section. Figure 15 shows the electron-scattering optical depth τ_e from our best-fit model (black line; the shaded region shows the 95% confidence limit). Also shown are constraints from the Planck satellite (Planck Collaboration et al. 2020) and the recent

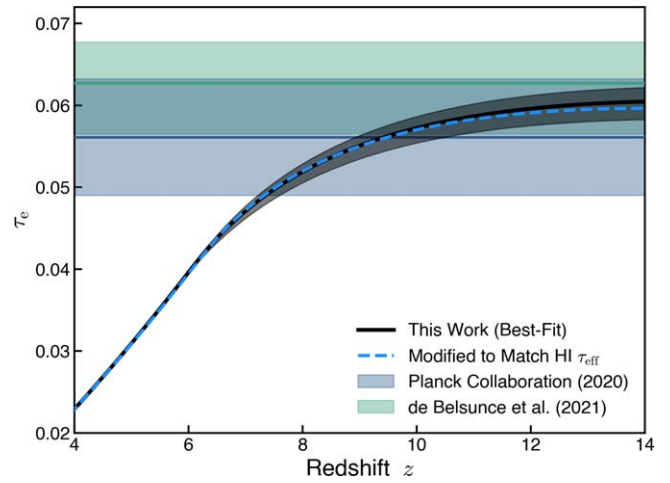


Figure 15. Thomson optical depth from electron scattering of the CMB τ_e from the best-fit model and the 95% confidence limit (black line and shaded bar) and our modified model to match the $z > 5 \tau_{\text{eff,H}}$ (dashed blue line). Also shown are the observational measurements from the Planck satellite presented in Planck Collaboration et al. (2020) and the constraint from de Belsunce et al. (2021). Our model results for τ_e lie within the Planck limits.

constraint from de Belsunce et al. (2021). Our result for $\tau_e = 0.60$ lies within the upper limit of the $\tau_e = 0.0540 \pm 0.0074$ constraint from Planck Collaboration et al. (2020) and in good agreement with the determination of $\tau_e = 0.0627_{-0.0058}^{+0.0050}$ from de Belsunce et al. (2021).

3.7. Modified UVB Rates for Matching the Observed High-redshift Hydrogen Effective Optical Depth

In Sections 3.4 and 3.5 we discuss how the IGM from our best-fit model is possibly too highly ionized after hydrogen reionization completes. The hydrogen effective optical depth $\tau_{\text{eff,H}}$ from the model is significantly lower compared with observations in the redshift range $5 \lesssim z \lesssim 5.8$ (see Figure 12). We can address this issue by decreasing the H I photoionization rate Γ_{HI} such that the sharp transition at $z \sim 5.8$ from the best-fit model is replaced by a more gradual increase of Γ_{HI} during the redshift range $4.8 \lesssim z \lesssim 6.0$ (dashed blue line in Figure 13). This alternative transition in Γ_{HI} was chosen such that the resulting evolution of H I $\tau_{\text{eff,H}}$ is consistent with the observations from Bosman et al. (2018).

Assuming that changes made to the photoionization rate Γ_{HI} correspond to a change of the mean free path of ionizing photons λ_{mfp} , then the He I photoionization rate Γ_{HeI} should also reflect the modification applied to Γ_{HI} . Correspondingly, we rescale the helium photoionization rate Γ_{HeI} such that the ratio $\Gamma_{\text{HI}}(z)/\Gamma_{\text{HeI}}(z)$ from the modified model matches the best-fit model.

Changing λ_{mfp} would also affect the photoheating rates \mathcal{H}_{HI} and \mathcal{H}_{HeI} . Assuming that the average energy of the ionizing photons remains the same in the modified model, we rescale the photoheating rates such that the ratios $\mathcal{H}_{\text{HI}}(z)/\Gamma_{\text{HI}}(z)$ and $\mathcal{H}_{\text{HeI}}(z)/\Gamma_{\text{HeI}}(z)$ match the best-fit model. Results from our modified model for photoheating and photoionization rates are shown in Figure 9 as dashed blue lines.

After hydrogen reionization completes at $z \lesssim 6.0$, hydrogen in the IGM is in photoionization equilibrium. During this epoch, decreasing the H I and He I photoionization rates effectively increases the neutral fraction of hydrogen and

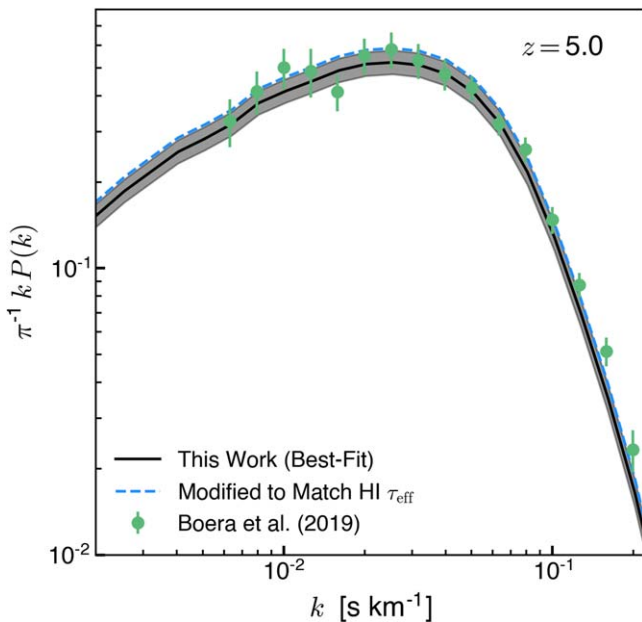


Figure 16. Power spectrum of the Ly α transmitted flux $P(k)$ at $z = 5$ from our best-fit model (black) and from our alternative model of the UVB (dashed blue), where the H I and He I photoionization and photoheating rates are modified in the redshift range $4.8 \lesssim z \lesssim 6.1$ such that $\tau_{\text{eff,H}}$ is consistent with the observation from Bosman et al. (2018). The effect on the power spectrum from the modified model is to increase $P(k)$ by a roughly constant factor of $\sim 12\%$ compared with the best-fit model due to the $\sim 6\%$ increase in the H I opacity at $z = 5$. Both models are consistent with the observation from Boera et al. (2019) for $k \lesssim 0.1 \text{ s km}^{-1}$.

helium. Consequently, the opacity of the IGM, quantified as the optical depth τ_{eff} , also increases during the redshift range. The temperature of the gas in the IGM is not strongly affected by the modified photoionization and photoheating rates because, in equilibrium, the gas temperature $T(z) \propto \mathcal{H}(z)/\Gamma(z)$ and this ratio is unchanged from the best-fit model.

The modified model only changes the photoionization and photoheating rates during the redshift range $4.8 \leq z \leq 6.1$. These changes result in an increase of $\tau_{\text{eff,H}}$ during $4.8 \leq z \leq 5.8$ and a decrease during $5.8 < z \leq 6.1$ but do not strongly affect the evolution of the gas temperature. For redshifts $z < 4.8$, the ionization fraction of hydrogen in the IGM in photoionization equilibrium is determined by the ratio of the photoionization rate to the recombination rate $x_{\text{HI}}(z) \propto \Gamma_{\text{HI}}(z)/\alpha_{\text{HI}}(z, T)$. The thermal evolutions resulting from the modified and best-fit models are very similar, and the rates Γ and \mathcal{H} at $z < 4.8$ are the same. Therefore, the evolution of the neutral fraction x_{HI} , the effective optical depth $\tau_{\text{eff,H}}$, and the Ly α power spectrum $P(k)$ resulting from the modified model is nearly unchanged from the best-fit model at redshifts $z < 4.8$.

The increase in the hydrogen effective optical depth $\tau_{\text{eff,H}}$ during the redshift range $4.8 \leq z \leq 5.8$ in the modified model influences the Ly α power spectrum at this epoch. Given the available data, this modification only affects comparisons with the observed $P(k)$ at $z = 5.0$. Figure 16 shows $P(k)$ from the modified model (dashed blue) and best-fit model (black) at $z = 5$. Relative to the best-fit model, using the modified model results in a small increase ($\sim 12\%$) in $P(k)$ owing to the small increase ($\sim 6\%$) in $\tau_{\text{eff,H}}$. Either model shows consistency with the observational $P(k)$ measurement from Boera et al. (2019).

3.8. Limitations of the Model

For this work, we have modeled the evolution of the properties of the IGM using a spatially homogeneous ionizing background. Simulations of a more realistic, spatially inhomogeneous hydrogen reionization process show that spatial fluctuations in the temperature–density relation of the post-reionization IGM have a minor effect on the flux power spectrum (Keating et al. 2018) at $z \leq 5$, while the inhomogeneous UVB allows large islands of neutral hydrogen to persist up to redshift $z \leq 5.5$ and can reproduce the observed distribution of Ly α opacity (Kulkarni et al. 2019). Similarly, radiative transfer simulations of He II reionization show that the fluctuations in the ionization state of helium have a minor effect on observations of the hydrogen Ly α forest (La Plante et al. 2017; Upton Sanderbeck & Bird 2020). Not including the impact of galactic winds or AGN feedback on the forest is a conservative approach for simulations aimed at constraining effects that suppress small-scale power. AGN feedback in the form of heating or mass redistribution from small to large scales is also expected to suppress the 1D power spectrum on large scales and to have an increased effect at low redshifts (Viel et al. 2013b). Ignoring the impact of AGN feedback may lead to a few percent bias in the determination of cosmological and astrophysical parameters (Chabanier et al. 2020). This model uncertainty is comparable to the statistical uncertainties of the eBOSS data used in this work.

Another limitation of our method results from the UVB photoionization and photoheating rates used for our simulation grid being constructed from simple transformations of a template set of rates. We therefore do not probe the full range of ionization and thermal histories that could be allowed by the observations of the Ly α forest. However, our model produces statistical properties of the Ly α forest that agree with a wide range of observations and a thermal evolution of the IGM consistent with previous inferences. These features of our work represent a significant achievement enabled by the ability to explore a wide range of models for the UVB from self-consistently evolved simulations. We emphasize that with our computational capabilities, performing a very large number of simulations (e.g., thousands) is now a possibility. We therefore defer more flexible explorations of models for the heating and ionization from the UVB to future work.

In the approach used for this work, we modify the photoionization and photoheating jointly. This joint variation results in another important limitation of our study. The large scales of the power spectrum of the forest are sensitive to the ionization state of H I, which, in equilibrium, is set by the balance between photoionization and recombination. The large scales of $P(k)$ depend on the temperature of the gas through the recombination coefficient $\alpha(T) \propto T^{-0.72}$ but are mostly determined by the intensity of the photoionization rate Γ_{HI} . Since a large fraction of the data set used for our inference probes the large-scale $P(k)$, the best-fit photoheating rates are influenced by the determination of the best-fit photoionization rates. We have shown that the photoheating from our best-fit model is consistent with other estimates of the thermal state of the IGM determined independently. Nevertheless, the relatively small uncertainty in the thermal state parameters T_0 and γ from this work is in part a consequence of the well-constrained determination of the photoionization rate from the large-scale $P(k)$. In future work we will explore a more flexible approach in which the photoheating has some degree of freedom with

respect to the photoionization rate, such as using density-dependent UVB rates to better model an inhomogeneous reionization.

4. Summary

With the objective of finding a photoionization and photoheating history that results in properties of the IGM consistent with observations of the hydrogen and helium Ly α forest, we have used the GPU-native Cholla code to perform an unprecedented grid of more than 400 cosmological simulations spanning a variety of ionization and thermal histories of the IGM. These calculations extend our CHIPS suite of hydrodynamical simulations initially presented in Villasenor et al. (2021). We compare the properties of the Ly α forest from our simulations to several observational measurements to determine via a likelihood analysis the best-fit model for the photoionization and photoheating rates. From our best-fit model we have inferred the thermal history of the IGM and demonstrate consistency with recent estimates obtained from the properties of the Ly α forest. A summary of the efforts and conclusions from this work follows.

1. We present a direct extension of the CHIPS suite (Villasenor et al. 2021) consisting of a grid of 400 simulations ($L = 50 h^{-1}$ Mpc, $N = 1024^3$) that vary the spatially uniform photoionization and photoheating rates from the metagalactic UVB. The UVB rates applied for our grid use the Puchwein et al. (2019) model as a template and use four parameters that control a rescaling of the amplitude and redshift timing of the hydrogen and helium photoionization and photoheating rates.
2. The CHIPS simulations self-consistently evolve a wide range of ionization and thermal histories of the IGM. We compare the properties of the Ly α forest in the form of the power spectrum $P(k)$ of the hydrogen Ly α transmitted flux and the helium (He II) effective optical depth $\tau_{\text{eff,HeII}}$ from our simulations to several observational measurements covering the redshift range $2.2 \leq z \leq 5.0$ for $P(k)$ (Iršič et al. 2017a; Boera et al. 2019; Chabanier et al. 2019) and $2.4 \lesssim z \lesssim 2.9$ for $\tau_{\text{eff,HeII}}$ (Worseck et al. 2016).
3. We perform a Bayesian MCMC marginalization to determine the best-fit UVB model. The performance of each model in reproducing the observations is evaluated over the entire redshift evolution instead of comparing for each redshift bin independently. Additionally, our simulation grid naturally probes a large range of ionization histories that we match directly to evolution of the ionization state of hydrogen encoded in the power spectrum of the Ly α forest. We thereby avoid any need to rescale the optical depth from the simulations in post-processing to match the observed mean transmission of the forest, which is a common shortcoming of previous analyses.
4. Our approach does not require an assumption of a power-law relation for the density–temperature distribution of the gas, as the Ly α spectra are constructed from our self-consistently evolved simulations. We find that a single power law does not accurately describe the $\rho_{\text{gas}} - T$ distribution of the gas in the density range relevant to generating the signal of the Ly α forest.
5. From our analysis, we infer the evolution of the thermal state of the IGM. The temperature history of the IGM

shows a first temperature peak ($T_0 \simeq 1.3 \times 10^4$ K) due to hydrogen reionization at $z \simeq 6$. This peak is followed by an epoch of cooling due to adiabatic expansion of the universe until the onset of helium reionization from radiation emitted by AGNs. The ionization of helium leads to a second increase of the temperature until He II is fully ionized ($z \simeq 3$), resulting in a second peak of $T_0 \simeq 1.4 \times 10^4$ K. The second peak is followed by a second period of cooling from cosmic expansion. Our result is consistent with previous estimates from Gaikwad et al. (2020) and Gaikwad et al. (2021). We note that the method employed in this work, where we modify the UVB photoionization and photoheating rates by rescaling and shifting the model from Puchwein et al. (2019), limits the variation on the evolution of the thermal history of the IGM in our simulations. In future work we will allow for more flexibility in the photoheating history, which will result in a more complete sample of the IGM density–temperature distribution. The improved flexibility of the models may permit a better inference of the thermal history of the IGM, as for now our low-redshift ($z < 4$) constraints are largely informed by the ionization state of hydrogen, which likely results in an underestimated uncertainty in our $T_0 - \gamma$ evolution.

6. We compare the evolution of the hydrogen effective optical depth $\tau_{\text{eff,H}}$ from our best-fit model to several observational determinations. We find that after hydrogen reionization completes ($5 \lesssim z \lesssim 6$) the HI effective optical depth resulting from the model may underestimate the observations. We provide a modification to our best-fit model where the photoionization and photoheating rates are reduced during this epoch such that the evolution of $\tau_{\text{eff,H}}$ is consistent with measurements by Bosman et al. (2018). Additionally, the neutral fraction of hydrogen from the modified model shows consistency with the measurements by Fan et al. (2006) during this redshift interval.
7. The model for the photoionization and photoheating rates from the UVB obtained from our analysis shows consistency with the observations of the Ly α power spectrum and the effective optical depth from both hydrogen and helium (He II), the optical depth from the CMB probed by Planck Collaboration et al. (2020), and previous inferences of the thermal state of the IGM. This model can be applied in future cosmological simulations that aim to reproduce properties of the IGM consistent with the observed Ly α forest.

Our work shows that an exploration of the IGM properties from hundreds of self-consistently evolved models for the astrophysical processes that impact the gas in the medium is now possible by exploiting modern computational techniques on the world’s largest supercomputers. Using our efficient GPU-based code Cholla with Summit, we are able to run hundreds of cosmological simulations in just a few days using a small fraction of the system. We anticipate that when combined with the exquisite picture of the Ly α forest that experiments like DESI Collaboration et al. (2016) will provide, this capability will revolutionize future studies of the properties of the IGM. We can leverage next-generation exascale systems and simulate large volumes ($L \sim 50 h^{-1}$ Mpc) at high resolution ($N = 2048^3$) for thousands of models

describing the various astrophysical processes that affect the IGM with a range of cosmological parameters and study different models for the nature of dark matter and the mass hierarchy of neutrinos based on their impact on the small-scale power spectrum of the Ly α forest.

This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under contract DE-AC05-00OR22725, using Summit allocations CSC434 and AST169. An award of computer time was provided by the INCITE program, via projectAST175. We acknowledge use of the *lux* supercomputer at UC Santa Cruz, funded by NSF MRI grant AST1828315, and support from NASA TCAN grant 80NSSC21K0271. B.V. is supported in part by the UC MEXUS-CONACyT doctoral fellowship. B.E.R. acknowledges support from NASA contract NNG16PJ25C and grants 80NSSC18K0563 and 80NSSC22K0814. We acknowledge the comments and suggestions received from the anonymous referee, which helped improve the content and clarity of this work.

Software: Cholla (Schneider & Robertson 2015, <https://github.com/cholla-hydro/cholla>), Python (van Rossum & Drake Jr. 1995), Numpy (Van Der Walt et al. 2011), Matplotlib (Hunter 2007), MUSIC (Hahn & Abel 2011), GRACKLE (Smith et al. 2017).

Appendix A Resolution Convergence Analysis

To assess the possible impact of the simulation spatial resolution on our results, we compare the Ly α transmitted flux power spectrum measured from simulations with different resolutions. Each run was performed using the same box size ($L = 50 h^{-1}$ cMpc) for identical cosmological parameters (Planck Collaboration et al. 2020) and our best-fit determination for the photoionization and photoheating rates, and they

differ only in their grid resolution. Our comparison is made between three runs with sizes $N = 512^3$, $N = 1024^3$, and $N = 2048^3$ cells and dark matter particles, with comoving spatial resolutions of $\Delta x \simeq 98, 49,$ and $24 h^{-1}$ kpc, respectively. The initial conditions for the runs were generated to preserve common large-scale modes, such that the results from the simulations could be compared directly over shared spatial scales.

Figure 17 shows the power spectrum of the Ly α flux measured for our three simulations at redshifts $z = 2, 3, 4,$ and 5 . As shown, the structure of the Ly α forest becomes better resolved as the number of cells increases. The bottom panels present the fractional difference $\Delta P(k)/P(k)$ of the power spectrum measured from the $N = 512^3$ and $N = 1024^3$ simulations compared with the $N = 2048^3$ simulation on overlapping spatial scales. Our comparison shows that the effect of the decreased resolution is to increase the power on large scales ($k \lesssim 0.02 \text{ s km}^{-1}$) while the small-scale power is suppressed. For the low-resolution simulation ($N = 512^3$) the differences are significant, and on large scales the power spectrum is overestimated by $\sim 50\%$ at redshift $z = 5$. As the redshift decreases, the differences also decrease to $\sim 13\%$ by $z = 2$. On small scales, the power spectrum is suppressed by $20\% - 60\%$.

Our fiducial resolution for the CHIPS simulations was $N = 1024^3$. At this resolution we measure only small differences in the Ly α structure compared with the $N = 2048^3$ simulation, as on large spatial scales the power spectrum is overestimated by $\lesssim 7\%$, and for small scales ($0.03 \text{ s km}^{-1} \lesssim k \lesssim 0.2 \text{ s km}^{-1}$) we measure a suppression on $P(k)$ of $\lesssim 10\% - 25\%$. To account for the effect of resolution on simulations used to constrain the UVB model, we include a systematic uncertainty of the form $\sigma_{\text{res}} = \Delta P(k, z)$, where $\Delta P(k, z)$ is the redshift- and scale-dependent difference in the power spectrum measured between the $N = 1024^3$ and $N = 2048^3$ simulations.




Figure 17. Power spectrum of the Ly α transmitted flux $P(k)$ measured from simulations with different comoving spatial resolutions of $\Delta x \simeq 98, 49,$ and $24 h^{-1}$ kpc. The three simulations model an $L = 50 h^{-1}$ Mpc box with the Planck Collaboration et al. (2020) cosmology and apply our best-fit determination for the photoionization and photoheating rates. The bottom panels show the fractional difference in the power spectrum $\Delta P(k)/P(k)$ between the $N = 512^3$ and $N = 1024^3$ runs and the $N = 2048^3$ simulation. Low-resolution simulations show increased power on large scales ($k \lesssim 0.03 \text{ s km}^{-1}$) and suppressed structure in the small scales relative to higher-resolution simulations. For the intermediate-resolution simulation $N = 1024^3$, which corresponds to our fiducial CHIPS grid resolution, the differences in $P(k)$ with respect to the $N = 2048^3$ simulation are $\lesssim 7\%$ on the large scales and $\lesssim 10\%$ – 25% on the small scales. We account for this resolution effect during our inference procedure by adding a systematic error to the observational measurements of $P(k)$ in the form of $\sigma_{\text{res}} = \Delta P(k, z)$, where $\Delta P(k, z)$ is the redshift- and scale-dependent difference in the power spectrum measured from the $N = 1024^3$ run compared with the $N = 2048^3$ simulation.

Appendix B

Effect of Rescaling the HI Effective Optical Depth on the Ly α Flux Power Spectrum

The power spectrum of the Ly α transmitted flux $P(k)$ is computed from flux fluctuations $\delta_F = (F - \langle F \rangle) / \langle F \rangle$. The power spectrum is sensitive to changes on the ionization state of hydrogen in the IGM, which in turn changes the effective optical depth $\tau_{\text{eff,H}}$ and the mean transmitted flux $\langle F \rangle = \exp(-\tau_{\text{eff,H}})$. To estimate how changes in the overall ionization state of the IGM affect the power spectrum of the Ly α flux, we can rescale the optical depth of the simulated skewers and remeasure $P(k)$. We rescale by a constant factor tuned such that the effective optical depth measured from the rescaled skewers follows $\widetilde{\tau}_{\text{eff,H}} = (1 + \alpha)\tau_{\text{eff,H}}$, where $\tau_{\text{eff,H}}$

is the original effective optical depth obtained from the simulated skewers. From the rescaled skewers, we compute the corresponding fluctuations of the transmitted flux $\widetilde{\delta}_F = (\widetilde{F} - \langle \widetilde{F} \rangle) / \langle \widetilde{F} \rangle$, where $\langle \widetilde{F} \rangle = \exp(-\widetilde{\tau}_{\text{eff,H}})$. Finally, from $\widetilde{\delta}_F$ we compute the mean flux power spectrum $\widetilde{P}(k)$ for the rescaled sample.

Figure 18 shows the fractional difference of the flux power spectrum $\Delta P(k)/P(k) = \widetilde{P}(k)/P(k) - 1$ measured between the rescaled skewers and the original sample for several values in the range $\alpha \in [-0.3, 0.3]$. Because of the nonlinear relation between the optical depth τ and the transmitted flux $F = \exp(-\tau)$, rescaling the effective optical depth $\tau_{\text{eff,H}}$ in the skewer sample to higher values $\alpha > 0$ has the effect of increasing the overall normalization of $P(k)$ on most of the




Figure 18. Consequences of rescaling the effective optical depth for the power spectrum of the Ly α transmitted flux at redshifts $z = 2, 3, 4,$ and 5 . Shown is the fractional difference $\Delta P(k)/P(k)$ after rescaling the optical depth along the skewer sample from our simulations by a constant factor such that $\widetilde{\tau}_{\text{eff,H}} = (1 + \alpha)\tau_{\text{eff}}$ for α in the range $[-3, 3]$. Rescaling the optical depth along the skewers such that $\tau_{\text{eff,H}}$ increases (decreases) has the effect of increasing (decreasing) $P(k)$. On scales in the range $0.002 \text{ s km}^{-1} \lesssim k \lesssim 0.1 \text{ s km}^{-1}$ the change induced on $P(k)$ is almost uniform, while for the smallest scales $k \gtrsim 0.1 \text{ s km}^{-1}$ the effect is redshift and scale dependent.

scales relevant for this work, namely, $0.002 \text{ s km}^{-1} \lesssim k \lesssim 0.1 \text{ s km}^{-1}$. In a similar way, decreasing $\tau_{\text{eff,H}}$ decreases the normalization of $P(k)$ at these scales. For smaller scales $k > 0.1 \text{ s km}^{-1}$ the effects are redshift dependent and we find that increasing (decreasing) $\tau_{\text{eff,H}}$ tends to also increase (decrease) $P(k)$ for $z \gtrsim 3.5$, while it has the opposite effect for $z \lesssim 3.5$ as $P(k)$ decreases (increases) when $\tau_{\text{eff,H}}$ is increased (decreased).

This study shows that the Ly α power spectrum itself is sensitive to the hydrogen effective optical depth, and for this reason we do not include the observational measurements of $\tau_{\text{eff,H}}$ for our inference of the UVB model presented in this work.

Appendix C

Covariance Matrices of the Transmitted Flux Power Spectrum from the Simulations

In Section 2.9 we present the likelihood function employed for our MCMC analysis (Equation (5)). When comparing the power spectrum of the Ly α transmitted flux from the

simulations to the observational measurements, we employ the covariance matrices of $P(k)$ reported by the observational works (Iršič et al. 2017a; Boera et al. 2019; Chabanier et al. 2019). In this appendix, we quantify the effect on the covariance of the simulated $P(k)$ from variations in our model parameters.

Figure 19 shows the normalized covariance of $P(k)$ at $z = 4.6$ for simulations with different values for the parameters β_{H} (top panels) and Δz_{H} (bottom panels). Decreasing the parameter β_{H} increases the Ly α opacity of the IGM, which increases the normalization of $P(k)$ (see Appendix B). The increase of $P(k)$ also increases its covariance on roughly all scales. We measure small element-wise differences < 0.1 in the normalized covariance matrices across simulations that vary β_{H} , while for simulations with different Δz_{H} the impact is minimal and results in only < 0.03 element-wise differences. Figure 20 presents the covariance matrix of $P(k)$ at $z = 3.0$ for simulations that vary the parameters β_{He} (top panels) and Δz_{He} (bottom panels). Here we also measure the impact to be small with differences < 0.05 .




Figure 19. Normalized covariance matrix of the Ly α transmitted flux power spectrum at $z = 4.6$ measured from simulations that vary the parameters β_{HI} (top panels) and Δz_{HI} (bottom panels) independently. The structure of the covariance is maintained across the simulations. Decreasing the parameter β_{HI} increases the normalization of $P(k)$ and its covariance on roughly all scales. We measure small element-wise differences < 0.1 in the normalized covariance across simulations with different β_{HI} . The effect of changing Δz_{HI} is minimal, with element-wise differences < 0.03 .




Figure 20. Normalized covariance matrix of the Ly α transmitted flux power spectrum at $z = 3.0$ measured from simulations that vary the parameters β_{HeI} (top panels) and Δz_{HeI} (bottom panels) independently. The structure of the covariance is maintained across the simulations. Changes in β_{HeI} and Δz_{HeI} cause small variations in the normalized covariance matrix; we measure only small element-wise differences < 0.05 over these simulations.

Appendix D

Colder Intermediate-redshift IGM from Reduced Photoheating

In Section 3.3 we discuss how our best-fit model results in a warmer IGM compared to the estimates from Boera et al. (2019) during the interval $4.2 \lesssim z \lesssim 5.0$, as the temperature T_0 from our model is $\sim 1\sigma$ higher compared to their result. We

explore scenarios where the intermediate-redshift IGM is cooled relative to our best-fit model by decreasing the HI and HeI photoheating rates in the redshift range $4.2 \lesssim z \lesssim 6.2$. The modified photoheating rates are shown in Figure 21 (middle and right panels), along with the fractional differences relative to the best-fit model shown in the respective bottom




Figure 21. Evolution of the IGM temperature T_0 (left panel) from models of the UVB where the HI and He I photoheating rates have been reduced in the interval $4.2 \leq z \leq 6.2$ relative to our best-fit model (middle and right panels). The fractional differences of T_0 and the heating rates \mathcal{H}_{HI} and \mathcal{H}_{HeI} with respect to the best-fit model are shown in the bottom panels. The reduced photoheating rates decrease T_0 for $z < 6.2$, but the change is most significant for $3.5 \lesssim z \lesssim 6.0$. At $z \lesssim 3.5$ the impact on T_0 is minimal, as heating from He II reionization dominates. A reduction of $\sim 80\%$ in the photoheating rates at $z \sim 6.0$ causes a decrease in T_0 of $\sim 20\%$ at $z \sim 5.0$. For $z \lesssim 3.5$ the reduced photoheating has a minimal impact on T_0 of $\lesssim 5\%$.

panels. To compute the history of T_0 for the reduced photoheating models, we integrate the evolution of the temperature of a single cell at $\rho_{\text{gas}} = \bar{\rho}$ following the method from Hui & Gnedin (1997; see Section 2 of their work for a detailed description). The resulting evolution of T_0 for the different models is presented in the left panel of Figure 21. We show that reducing the HI and HeI photoheating rates by $\sim 80\%$ at $z \sim 6$ results in a colder IGM where T_0 is reduced by $\sim 20\%$ at $z \sim 5$ such that $T_0 \sim 8 \times 10^3$ K for $4.2 \lesssim z \lesssim 5.0$ agrees well with the estimate from Boera et al. (2019). However, we find that for such a scenario T_0 at $z \sim 5.4$ is lower than the inference from Gaikwad et al. (2020). This conflict exhibits some degree of tension between the estimates at $z \sim 5.0$ and $z \sim 5.4$ from Boera et al. (2019) and Gaikwad et al. (2020), respectively.

The photoheating \mathcal{H} and the photoionization Γ rate from the UVB are given by the intensity of the background radiation $J(\nu, z)$ as

$$\begin{aligned} \Gamma(z) &= \int_{\nu_0}^{\infty} \frac{4\pi J(\nu, z)}{h\nu} \sigma(\nu) d\nu, \\ \mathcal{H}(z) &= \int_{\nu_0}^{\infty} \frac{4\pi J(\nu, z)}{h\nu} (h\nu - h\nu_0) \sigma(\nu) d\nu, \end{aligned} \quad (\text{D1})$$

where ν_0 and $\sigma(\nu)$ are the threshold frequency and photoionization cross section, respectively. Consider power-law models for the cross section and the intensity of the radiation at wavelengths $\lambda > 912 \text{ \AA}$, which can be written as $\sigma(\nu) = \sigma_0(\nu/\nu_0)^\phi$ and $J(\nu) = (\nu/\nu_0)^\alpha$, with indices $\phi < 0$ and $\alpha < 0$. Physically, reducing the photoheating rate relative to the photoionization rate can be achieved by changing the spectral index of the ionizing radiation α . By solving the integrals in Equation (D1) assuming these power-law models and evaluating the fractional change in the photoionization $\Delta\Gamma/\Gamma$ and photoheating $\Delta\mathcal{H}/\mathcal{H}$ for a change in the spectral

index $\Delta\alpha$, we find that the following relation is satisfied:

$$\Delta\alpha = (1 + \alpha + \phi) \frac{\frac{\Delta\Gamma}{\Gamma} - \frac{\Delta\mathcal{H}}{\mathcal{H}}}{1 + \frac{\Delta\mathcal{H}}{\mathcal{H}}}. \quad (\text{D2})$$

Equation (D2) relates the change of the spectral index of the radiation necessary to produce some variation of the photoionization and photoheating from a given UVB model. By applying Equation (D2), we can modify the photoheating relative to the photoionization of a UVB model within a physically plausible range for the index α . In future work, we will explore which variations in the IGM temperature T_0 from changes of the photoheating rate match the observed hydrogen effective optical depth at $z > 5$ while using physically plausible source populations.

Appendix E Accuracy of the Power-law Fit to the Density–Temperature Distribution of the Gas in Our Simulations

A common method to infer the thermal state of the IGM from observations of the Ly α forest involves marginalizing over the thermal properties T_0 and γ in the approximate power-law density–temperature relation $T = T_0(\rho_{\text{gas}}/\bar{\rho})^{\gamma-1}$ (Bolton et al. 2014; Nasir et al. 2016; Hiss et al. 2018; Boera et al. 2019; Walther et al. 2019; Gaikwad et al. 2021). The density of the IGM gas that contributes to the majority of the Ly α forest signal lies in the range $-1 \leq \log_{10}(\rho_{\text{gas}}/\bar{\rho}) \leq 1$. From our simulations we find that a single power law fails to reproduce the density–temperature distribution of the gas over this density interval. The left panels of Figure 22 show the density–temperature distribution of the gas in one of our simulations and the corresponding power-law fit to the distribution over the density range $-1 \leq \log_{10}(\rho_{\text{gas}}/\bar{\rho}) \leq 1$ at redshift $z = 3$ (top) and $z = 4$ (bottom). The deviations of the gas temperature in the simulation relative to the power-law fits are presented in the right panels, showing that the fractional differences $\Delta T/T$ from

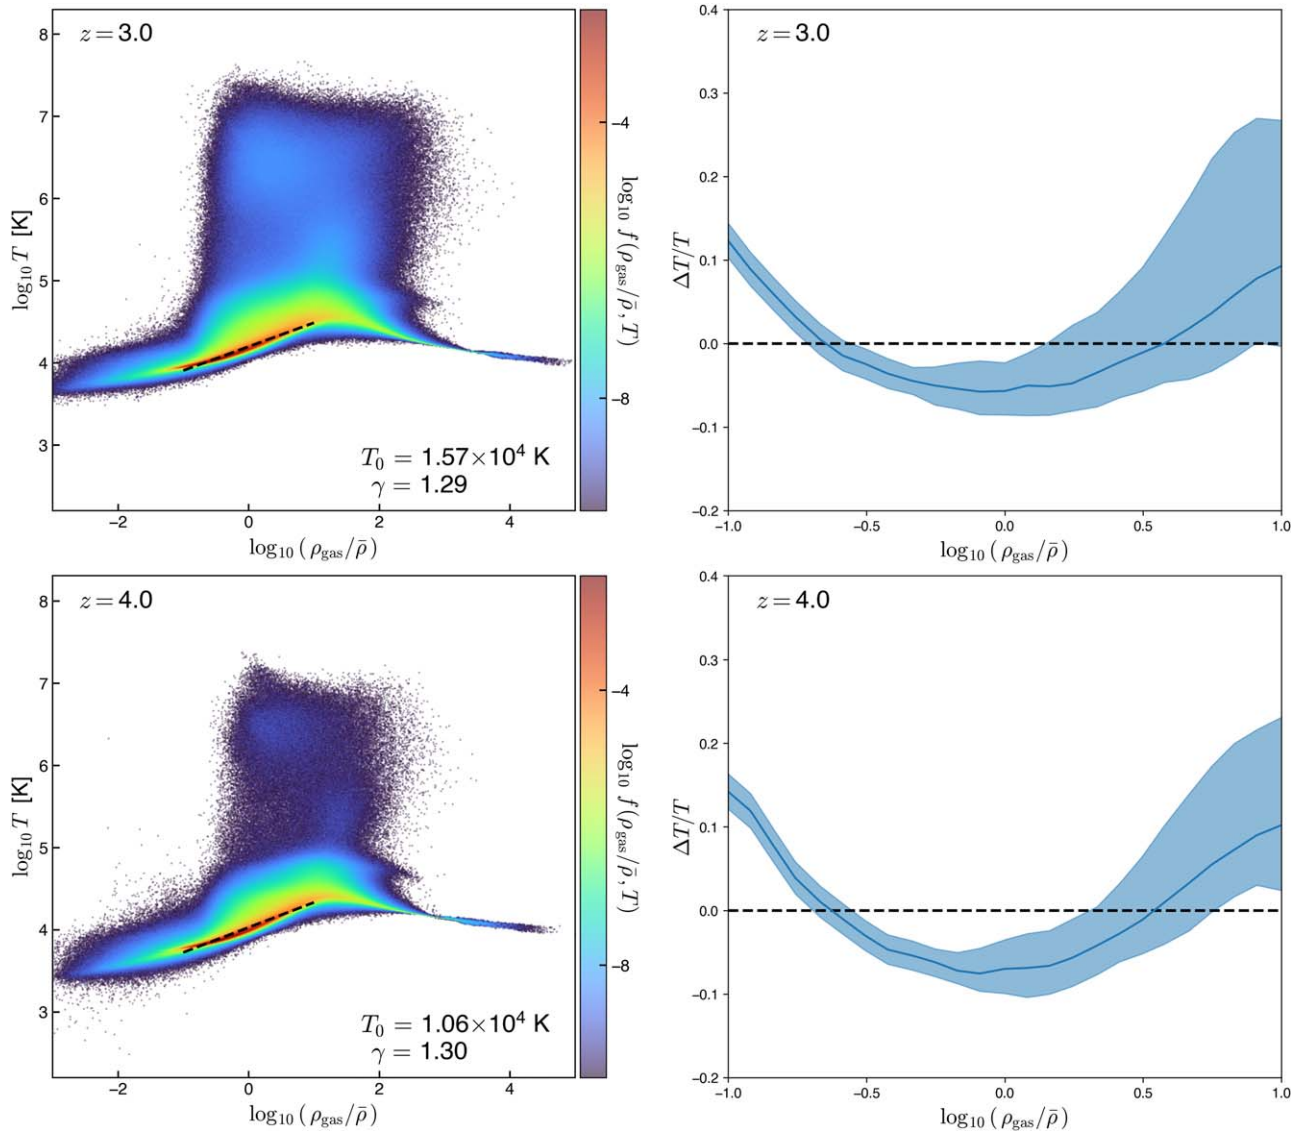


Figure 22. Density–temperature distribution of the IGM gas (left column) from one of our simulations at redshift $z = 3$ (top) and $z = 4$ (bottom). A power-law fit of the form $T = T_0(\rho_{\text{gas}}/\bar{\rho})^{\gamma-1}$ over the range $-1 \leq \log_{10}(\rho_{\text{gas}}/\bar{\rho}) \leq 1$ is shown (black dashed lines). The right column shows deviations of the density–temperature distribution with respect to the power-law fit over the fitted range. The blue region corresponds to the 68% highest probability interval for the temperature as a function of the overdensity $\rho_{\text{gas}}/\bar{\rho}$. The differences between the distribution of the gas relative to the power-law fit can be as large as $\sim 15\%$.

the density–temperature distribution in the simulation with respect to the power-law fit can be as large as $\sim 15\%$.

ORCID iDs

Bruno Villasenor <https://orcid.org/0000-0002-7460-8129>

Brant Robertson <https://orcid.org/0000-0002-4271-0364>

Piero Madau <https://orcid.org/0000-0002-6336-3293>

Evan Schneider <https://orcid.org/0000-0001-9735-7484>

References

- Becker, G. D., & Bolton, J. S. 2013, *MNRAS*, 436, 1023
- Becker, G. D., D’Aloisio, A., Christenson, H. M., et al. 2021, *MNRAS*, 508, 1853
- Becker, G. D., Hewett, P. C., Worseck, G., & Prochaska, J. X. 2013, *MNRAS*, 430, 2067
- Becker, R. H., Fan, X., White, R. L., et al. 2001, *AJ*, 122, 2850
- Bird, S., Rogers, K. K., Peiris, H. V., et al. 2019, *JCAP*, 2019, 050
- Boera, E., Becker, G. D., Bolton, J. S., & Nasir, F. 2019, *ApJ*, 872, 101
- Bolton, J. S., Becker, G. D., Haehnelt, M. G., & Viel, M. 2014, *MNRAS*, 438, 2499
- Bond, J. R., Kofman, L., & Pogosyan, D. 1996, *Natur*, 380, 603
- Bosman, S. E. I., Fan, X., Jiang, L., et al. 2018, *MNRAS*, 479, 1055
- Calverley, A. P., Becker, G. D., Haehnelt, M. G., & Bolton, J. S. 2011, *MNRAS*, 412, 2543
- Cen, R., Miralda-Escudé, J., Ostriker, J. P., & Rauch, M. 1994, *ApJL*, 437, L9
- Chabanier, S., Bournaud, F., Dubois, Y., et al. 2020, *MNRAS*, 495, 1825
- Chabanier, S., Palanque-Delabrouille, N., Yèche, C., et al. 2019, *JCAP*, 2019, 017
- Colella, P., & Woodward, P. R. 1984, *JCoPh*, 54, 174
- Croft, R. A. C., Weinberg, D. H., Katz, N., & Hernquist, L. 1998, *ApJ*, 495, 44
- D’Aloisio, A., McQuinn, M., Davies, F. B., & Furlanetto, S. R. 2018, *MNRAS*, 473, 560
- Davies, F. B., Hennawi, J. F., Bañados, E., et al. 2018, *ApJ*, 864, 142
- de Belsunce, R., Gratton, S., Coulton, W., & Efstathiou, G. 2021, *MNRAS*, 507, 1072
- DESI Collaboration, Aghamousa, A., Aguilar, J., et al. 2016, arXiv:1611.00036
- Eilers, A.-C., Davies, F. B., & Hennawi, J. F. 2018, *ApJ*, 864, 53
- Fan, X., Strauss, M. A., Becker, R. H., et al. 2006, *AJ*, 132, 117
- Faucher-Giguère, C.-A. 2020, *MNRAS*, 493, 1614

- Gaikwad, P., Khaire, V., Choudhury, T. R., & Srianand, R. 2017, *MNRAS*, **466**, 838
- Gaikwad, P., Rauch, M., Haehnelt, M. G., et al. 2020, *MNRAS*, **494**, 5091
- Gaikwad, P., Srianand, R., Haehnelt, M. G., & Choudhury, T. R. 2021, *MNRAS*, **506**, 4389
- Gallego, S. G., Cantalupo, S., Sarpas, S., et al. 2021, *MNRAS*, **504**, 16
- Greig, B., Mesinger, A., & Bañados, E. 2019, *MNRAS*, **484**, 5094
- Greig, B., Mesinger, A., Haiman, Z., & Simcoe, R. A. 2017, *MNRAS*, **466**, 4239
- Haardt, F., & Madau, P. 2012, *ApJ*, **746**, 125
- Hahn, O., & Abel, T. 2011, *MNRAS*, **415**, 2101
- Hernquist, L., Katz, N., Weinberg, D. H., & Miralda-Escudé, J. 1996, *ApJL*, **457**, L51
- Hiss, H., Walther, M., Hennawi, J. F., et al. 2018, *ApJ*, **865**, 42
- Ho, M.-F., Bird, S., & Shelton, C. R. 2022, *MNRAS*, **509**, 2551
- Hoag, A., Bradač, M., Huang, K., et al. 2019, *ApJ*, **878**, 12
- Hui, L., & Gnedin, N. Y. 1997, *MNRAS*, **292**, 27
- Hunter, J. D. 2007, *CSE*, **9**, 90
- Iršič, V., Viel, M., Berg, T. A. M., et al. 2017a, *MNRAS*, **466**, 4332
- Iršič, V., Viel, M., Haehnelt, M. G., et al. 2017b, *PhRvD*, **96**, 023522
- Jung, I., Finkelstein, S. L., Dickinson, M., et al. 2020, *ApJ*, **904**, 144
- Keating, L. C., Puchwein, E., & Haehnelt, M. G. 2018, *MNRAS*, **477**, 5501
- Khaire, V., & Srianand, R. 2019, *MNRAS*, **484**, 4174
- Kulkarni, G., Keating, L. C., Haehnelt, M. G., et al. 2019, *MNRAS*, **485**, L24
- La Plante, P., Trac, H., Croft, R., & Cen, R. 2017, *ApJ*, **841**, 87
- Madau, P., & Haardt, F. 2015, *ApJL*, **813**, L8
- Mason, C. A., Fontana, A., Treu, T., et al. 2019, *MNRAS*, **485**, 3947
- Mason, C. A., Treu, T., Dijkstra, M., et al. 2018, *ApJ*, **856**, 2
- McGreer, I. D., Mesinger, A., & D’Odorico, V. 2015, *MNRAS*, **447**, 499
- McGreer, I. D., Mesinger, A., & Fan, X. 2011, *MNRAS*, **415**, 3237
- McQuinn, M. 2016, *ARA&A*, **54**, 313
- Meiksin, A. A. 2009, *RvMP*, **81**, 1405
- Miralda-Escudé, J., & Rees, M. J. 1994, *MNRAS*, **266**, 343
- Nasir, F., Bolton, J. S., & Becker, G. D. 2016, *MNRAS*, **463**, 2335
- Oñorbe, J., Hennawi, J. F., & Lukić, Z. 2017, *ApJ*, **837**, 106
- Pieri, M. M., Bonoli, S., Chaves-Montero, J., et al. 2016, in SF2A-2016: Proc. of the Annual meeting of the French Society of Astronomy and Astrophysics, ed. C. Reylé et al. (Lyon: Centre de Recherche Astrophysique de Lyon), 259
- Planck Collaboration, Aghanim, N., Akrami, Y., et al. 2020, *A&A*, **641**, A6
- Puchwein, E., Bolton, J. S., Haehnelt, M. G., et al. 2015, *MNRAS*, **450**, 4081
- Puchwein, E., Haardt, F., Haehnelt, M. G., & Madau, P. 2019, *MNRAS*, **485**, 47
- Qin, Y., Mesinger, A., Bosman, S. E. I., & Viel, M. 2021, *MNRAS*, **506**, 2390
- Robertson, B. E., Ellis, R. S., Furlanetto, S. R., & Dunlop, J. S. 2015, *ApJL*, **802**, L19
- Schneider, E. E., & Robertson, B. E. 2015, *ApJS*, **217**, 24
- Slosar, A., Font-Ribera, A., Pieri, M. M., et al. 2011, *JCAP*, **2011**, 001
- Smith, B. D., Bryan, G. L., Glover, S. C. O., et al. 2017, *MNRAS*, **466**, 2217
- Upton Sanderbeck, P., & Bird, S. 2020, *MNRAS*, **496**, 4372
- Upton Sanderbeck, P. R., D’Aloisio, A., & McQuinn, M. J. 2016, *MNRAS*, **460**, 1885
- Van Der Walt, S., Colbert, S. C., & Varoquaux, G. 2011, *CSE*, **13**, 22
- van Rossum, G., & Drake, F. L., Jr. 1995, Python Reference Manual (Amsterdam: Centrum voor Wiskunde en Informatica Amsterdam)
- Viel, M., Becker, G. D., Bolton, J. S., & Haehnelt, M. G. 2013a, *PhRvD*, **88**, 043502
- Viel, M., Schaye, J., & Booth, C. M. 2013b, *MNRAS*, **429**, 1734
- Villasenor, B., Robertson, B., Madau, P., & Schneider, E. 2021, *ApJ*, **912**, 138
- Walther, M., Armengaud, E., Ravoux, C., et al. 2021, *JCAP*, **2021**, 059
- Walther, M., Hennawi, J. F., Hiss, H., et al. 2018, *ApJ*, **852**, 22
- Walther, M., Oñorbe, J., Hennawi, J. F., & Lukić, Z. 2019, *ApJ*, **872**, 13
- Wang, F., Davies, F. B., Yang, J., et al. 2020, *ApJ*, **896**, 23
- Wolfson, M., Hennawi, J. F., Davies, F. B., et al. 2021, *MNRAS*, **508**, 5493
- Worseck, G., Davies, F. B., Hennawi, J. F., & Prochaska, J. X. 2019, *ApJ*, **875**, 111
- Worseck, G., Prochaska, J. X., Hennawi, J. F., & McQuinn, M. 2016, *ApJ*, **825**, 144
- Wyithe, J. S. B., & Bolton, J. S. 2011, *MNRAS*, **412**, 1926
- Yang, J., Wang, F., Fan, X., et al. 2020a, *ApJ*, **904**, 26
- Yang, J., Wang, F., Fan, X., et al. 2020b, *ApJL*, **897**, L14