



HAL
open science

Constraining blazar heating with the $2 < z < 3$ Lyman- α forest

A. Lamberts, E. Puchwein, C. Pfrommer, P. Chang, M. Shalaby, A. Broderick, P. Tiede, G. Rudie

► **To cite this version:**

A. Lamberts, E. Puchwein, C. Pfrommer, P. Chang, M. Shalaby, et al.. Constraining blazar heating with the $2 < z < 3$ Lyman- α forest. *Monthly Notices of the Royal Astronomical Society*, 2022, 512 (2), pp.3045-3059. 10.1093/mnras/stac553 . hal-03574504

HAL Id: hal-03574504

<https://hal.science/hal-03574504>

Submitted on 23 Nov 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Constraining blazar heating with the $2 \lesssim z \lesssim 3$ Lyman- α forest

Astrid Lamberts^{1,2*}, Ewald Puchwein³, Christoph Pfrommer³, Philip Chang⁴,
Mohamad Shalaby³, Avery Broderick^{5,6}, Paul Tiede^{5,6}, Gwen Rudie⁷

¹Université Côte d’Azur, Observatoire de la Côte d’Azur, CNRS, Laboratoire Lagrange, Bd de l’Observatoire, CS 34229, 06304 Nice cedex 4, France.

²Université Côte d’Azur, Observatoire de la Côte d’Azur, CNRS, Laboratoire Artémis, Bd de l’Observatoire, CS 34229, 06304 Nice cedex 4, France.

³Leibniz Institute for Astrophysics, Potsdam (AIP), An der Sternwarte 16, D-14482 Potsdam, Germany

⁴Department of Physics, University of Wisconsin-Milwaukee, 3135 N. Maryland Ave., Milwaukee, WI 53211, USA

⁵Department of Physics and Astronomy, University of Waterloo, 200 University Avenue West, Waterloo, ON, N2L 3G1, Canada

⁶Perimeter Institute for Theoretical Physics, 31 Caroline Street North, Waterloo, ON, N2L 2Y5, Canada

⁷The Observatories of the Carnegie Institution for Science, 813 Santa Barbara Street, Pasadena, CA 91101, USA

Accepted XXX. Received YYY; in original form ZZZ

ABSTRACT

The intergalactic medium (IGM) acts like a calorimeter recording energy injection by cosmic structure formation, shocks and photoheating from stars and active galactic nuclei. It was recently proposed that spatially inhomogeneous TeV-blazars could significantly heat up the underdense IGM, resulting in patches of both cold and warm IGM around $z \approx 2 - 3$. The goal of this study is to compare predictions of different blazar heating models with recent observations of the IGM. We perform a set of cosmological simulations and carefully compute mock observables of the Lyman- α ($\text{Ly}\alpha$) forest. We perform a detailed assessment of different systematic uncertainties which typically impact this type of observables and find that they are smaller than the differences between our models. We find that our inhomogeneous blazar heating model is in good agreement with the $\text{Ly}\alpha$ line properties and the rescaled flux probability distribution function at high redshift ($2.5 < z < 3$) but that our blazar heating models are challenged by lower redshift data ($2 < z < 2.5$). Our results could be explained by HeII reionisation although state-of-the-art models fall short on providing enough heating to the low-density IGM, thus motivating further radiative transfer studies of inhomogeneous HeII reionisation. If blazars are indeed hosted by group-mass halos of $2 \times 10^{13} M_{\odot}$, a later onset of blazar heating in comparison to previous models would be favoured, which could bring our findings here in agreement with the evidence of blazar heating from local gamma-ray observations.

Key words: methods: numerical, data analysis - intergalactic medium - quasars: absorption lines

1 INTRODUCTION

Even at the present day, the majority of the baryons reside in the intergalactic medium (IGM) rather than within galaxies (Shull et al. 2012). As the main reservoir for baryons, the physical state of the IGM sets the initial conditions for galaxy formation. Due to its low density, the evolution of the IGM is mostly linear, closely follows the underlying dark matter and is directly influenced by fundamental cosmological parameters (Viel et al. 2004; Palanque-Delabrouille et al. 2013; Slosar et al. 2013; Palanque-Delabrouille et al. 2015).

Because of its linear nature, the IGM is an excellent calorimeter of energy injected by star formation and active galactic nuclei. More specifically, reionisation of H and HeI around $z \approx 10 - 5.3$ (Fan

et al. 2006; Kulkarni et al. 2019) and of HeII around $4 \lesssim z \lesssim 2.7$ provide most of the heat input in the IGM (McQuinn et al. 2009; Worseck et al. 2011, 2016). Subsequently, the thermal evolution of the IGM is set by photoheating following recombination and adiabatic cooling due to the Hubble expansion. The lowest density gas expands fastest and also receives the least photoheating as there are fewer recombinations that allow subsequent photoionisations. Together this yields, well after reionisation, a tight temperature-density relation $T = T_0 \Delta^{\gamma-1}$ for low density gas (Hui & Gnedin 1997; Puchwein et al. 2015; McQuinn & Upton Sanderbeck 2016) where T_0 is the temperature at the mean density, $\Delta = \rho/\bar{\rho} - 1$ is the baryon overdensity, and γ asymptotically approaches ≈ 1.6 . In this paper, we compare recent observational data with different models deviating from a tight power-law $T - \rho$ relation with $\gamma \approx 1.6$.

During and shortly after HeII reionisation, the thermal state

* E-mail: astrid.lamberts@oca.eu

of the IGM is more complex. Recent models of HeII reionisation, based on cosmological simulations including complete radiative transfer, indicate a patchy process yielding a broadening of the temperature-density distribution around and below mean density for $z \approx 3$ (McQuinn et al. 2009; Meiksin & Tittley 2012; Compostella et al. 2013; La Plante et al. 2017).

Nevertheless, the most dramatic impact on the low-density IGM could come from TeV-blazar heating (Chang et al. 2012; Pfrommer et al. 2012; Puchwein et al. 2012; Lamberts et al. 2015). This model is based on the electron/positron beams that result from pair-production from TeV gamma-rays from blazars on the extra-galactic background light.

The pair beams can be subject to plasma instabilities, which efficiently convert their energy into plasma modes within the IGM. If this energy is efficiently thermalised via, e.g., the nonlinear plasma mode interactions, the kinetic energy of the pair beams will ultimately be redistributed to the surrounding IGM (Broderick et al. 2012; Schlickeiser et al. 2013, 2012; Chang et al. 2014, 2016; Shalaby et al. 2017a,b; Rafiqhi et al. 2017; Vafin et al. 2018), but see Miniati & Elyiv (2013); Sironi & Giannios (2014). The resulting heating would only be limited by the number of TeV-photons, which makes it a competitive heating source in underdense regions of the IGM (Chang et al. 2012), where photoheating is insignificant as the recombination time is larger than the Hubble time. Assuming a uniform heating rate, blazar heating results in an inverted temperature-density distribution below the cosmic mean ($\gamma \leq 1$), with low density gas reaching 10^5K at $z = 3$ and $\approx 3 \times 10^5\text{K}$ in the present-day universe (see Fig. 12 in Puchwein et al. 2012, P12 hereafter).

Although uniform heating is a reasonable first order approximation, in Lamberts et al. (2015, hereafter Paper I) we showed that the heating is affected by the clustering of blazars. As a result, regions close to large overdense regions such as clusters or groups are receiving more heat than remote regions mostly surrounded by voids.¹ In our favoured model, where TeV blazars have the same bias as quasars, there are almost two orders of magnitude in temperature between the hottest and coldest gas between $z \approx 2 - 3$. However, the bulk of the gas follows a temperature consistent with the uniform model. By the present day, all regions would have been heated up and the uniform model provides a good description of the impact of blazar heating. The goal of this paper is to compare this blazar heating model with recent observational data.

Determining the thermal state of the IGM with observations is challenging, especially for the regions around or below the mean density. The IGM is mostly observed through absorption lines in the spectra of distant quasars, due to a tiny fraction of neutral hydrogen (Lynds 1971). The so-called Ly α forest can be observed with ground-based facilities for $z \geq 1.7$ while the *Hubble Space Telescope* is currently the only available facility for low redshift ($z \lesssim 0.5$) measurements. As such, we can only observe sections of the thermal history of the IGM.

A variety of statistics have been developed to analyse spectra from a wide range of instruments. However, deriving the physical parameters of the IGM from the different observables requires a careful calibration to cosmological simulations (see e.g. Rauch et al.

1997; Schaye et al. 2000; Becker & Bolton 2013; Bolton et al. 2017; Gaikwad et al. 2020a).

The derivation of the temperature is also very sensitive to the assumed reionisation and heating history due to a (partial) degeneracy between pressure smoothing and instantaneous temperature.

Because of the intrinsic observational challenges, numerical shortcomings and difficulty to establish the validity of the different statistics, as well as their sensitivity to different overdensities, accurate constraints on the IGM thermal state are difficult to obtain. While a consensus has arisen about the detection of heating due to HeII reionisation in mildly overdense gas (e.g. Becker & Bolton 2013; Gaikwad et al. 2020a), there is no observational consensus yet on the thermal state of the very low-density IGM in which the impact of blazar heating should be strongest.

Different observational diagnostics suggest or disfavor the presence of blazar heating in the low density IGM. Based on the probability distribution function (PDF) of the transmitted Ly α flux, Kim et al. (2007); Bolton et al. (2008); Viel et al. (2009); Calura et al. (2012) find that a flat or inverted $T - \rho$ relation is in good agreement with the data. However, the flux PDF can be strongly affected by systematic errors in the continuum placement (Lee 2012) and sample variance (Rollinde et al. 2013). Analysis of the curvature of the Ly α spectrum (Becker et al. 2011; Boera et al. 2014) prefers a warmer temperature of the IGM. Unfortunately this method, which is based on the smoothness of the spectrum, is mostly sensitive to densities above the mean, especially at low redshift. Finally, Voigt profile fitting of the spectra yields the distribution of line-width (b) versus HI column density (N_{HI}). Using the lower envelope of the distribution as a proxy for the $T - \rho$ relation, Rudie et al. (2012b); Bolton et al. (2014, 2017) find no evidence of blazar heating, but are mostly sensitive to mildly overdense ($\Delta \gtrsim 1$) gas (also see the discussion in Rorai et al. 2018).

A strong case for blazar heating comes from a unique spectrum with high signal-to-noise (SNR) (Rorai et al. 2017). The high SNR allows the authors to rescale the optical depth to enhance the signal from low-density regions. At the considered redshifts ($2.5 \leq z \leq 3$), they do find that the high end of the transmitted flux probability distribution is better matched by a broken powerlaw for the temperature-density distribution, with an inverted slope at the low density end. Their model including temperature fluctuations at low densities produces a satisfactory match as well. The authors also perform a careful analysis of the power spectrum and lower envelope of the “ $b - N_{\text{HI}}$ ” distribution and show they are largely insensitive to the low-density regions.

In this paper, we compare our model for inhomogeneous blazar heating with observational data. Our work builds on the results of Puchwein et al. (2012) and includes more recent observational data. We start with a reminder of the main properties of inhomogeneous blazar heating (§2). We then describe the numerical simulations we use to model the IGM (§3) and the resulting observables we derive (§4). We then discuss the direct comparison with observations, suggest a unifying interpretation of Ly α and gamma-ray data (§5), and conclude (§6).

2 MAIN PROPERTIES OF INHOMOGENEOUS BLAZAR HEATING

Here we recall the main characteristics of the temperature-density distribution of the IGM under different heating assumptions, such as modelling or ignoring the spatial fluctuations in the blazar heating rate. We refer the reader to Paper I for a more detailed description.

¹ To ease comparison, we only varied the spatial heating rate in Paper I and kept the redshift evolution identical. In principle, a more highly biased population should also evolve later in a hierarchically growing universe, i.e., the blazar heating rate should have a different redshift evolution for the differently biased models. We postpone a study of this effect to future work.

Following [Chang et al. \(2012\)](#), we model the blazar uniform heating rate per comoving volume at a given redshift \dot{Q} by

$$\log_{10} \left(\frac{\dot{Q}(z)}{\dot{Q}(z=0)} \right) = 0.0315 \times [(1+z)^3 - 1] - 0.512 \times [(1+z)^2 - 1] + 2.27 \times [(1+z) - 1]. \quad (1)$$

As TeV blazar abundances at high redshift are observationally poorly constrained, we assume that the blazar luminosity density has a similar redshift evolution as the observed quasar luminosity density. The redshift dependence of Eq. (1) has been chosen to be consistent with the quasar luminosity functions of [Hopkins et al. \(2007\)](#). The $z = 0$ normalization of the TeV luminosity density can be obtained from *Fermi* observations of the local TeV blazar population (see [Broderick et al. 2012](#)). Here we specifically focus on the *intermediate* heating model of [Puchwein et al. \(2012\)](#) which assumes $\dot{Q}(z=0) = 1.08 \times 10^{-7} \text{ eV Gyr}^{-1} \text{ cm}^{-3}$. The homogeneous blazar heating model injects this energy with a spatially constant heating rate per unit volume. Our inhomogeneous blazar heating model has the same total amount of energy injected by blazars but accounts for regions receiving more or less heating according to their proximity to heating sources. This is modelled based on an analytic formalism relating the distribution of the heating rate to the underlying dark matter distribution. It results in the filtering function which removes small scale fluctuations and enhances fluctuations beyond $\approx 10 h^{-1} \text{ Mpc}$ at $z = 4$ and $\approx 40 h^{-1} \text{ Mpc}$ at $z = 2$. The shape and redshift evolution of the window function is set by the mean free path of the gamma rays combined with the bias of the heating sources. Roughly speaking, the heating rate is then given by the assumed blazar luminosity density (itself proportional to the matter density times a bias factor) convolved with a kernel that accounts for the $\propto r^{-2}$ dilution of the gamma ray flux as well as for the absorption with the appropriate mean free path (see Paper I for full details). While we presented two models in Paper I, with galaxy bias and quasar bias, in this work we focus on the model with quasar bias, which is probably more representative of the bias of TeV blazars and may even be a too conservative assumption (see our discussion in §5).

Figure 1 shows the temperature-density distribution according to our inhomogeneous heating model (two right-hand panels) and when blazar heating is absent (two left-hand panels). As time goes by, the cumulative impact of blazar heating increases, with a higher temperature difference with respect to the unheated case. Still, as some regions are too far from heating sources, they remain cold, as can be seen by the remnant cold gas, especially at $z = 2$. This is very different from the uniform heating model (shown by the black contours) and can potentially reconcile the blazar heating model with observations of absorption lines with Doppler parameter ($b \leq 20 \text{ km s}^{-1}$). At the redshifts most accessible with the Ly α forest, the temperature range covers almost two orders of magnitude at the lowest density and even around mean density, there is an important spread. At lower redshifts, the whole IGM gets heated up and for $z \leq 1$, the inhomogeneous model recovers the uniform model.

The physical size of the temperature fluctuations in the IGM is mostly set by the mean free path of TeV photons, and is a few tens of Mpc at $z = 3$ and increases up to $\approx \text{Gpc}$ in the present day universe. Fig. 2 shows a slice through the midplane of our simulation at $z = 2$ for the three heating models we considered. The regions influenced by blazar heating are typically tens of Mpc across. As such, our computational volume should be as large as possible to sample several regions in different thermal states. In the next section, we

describe the simulations we perform to model the IGM and how we postprocess them to extract mock observables. In section §4 we show how the thermal properties shown in Figs. 1 and 2 translate into observables.

3 METHODS

3.1 Cosmological simulations

Our simulations are very similar to the simulations presented in Paper I and are performed at higher resolution. We perform our simulations with GADGET-3, which is an upgraded version of GADGET-2 ([Springel 2005](#)). It is based on a Smoothed Particle Hydrodynamics (SPH) scheme and solves the gravitational evolution of gas and dark matter with a TREE-PM N -body method. The equations of hydrodynamics are solved with the entropy conserving scheme of [Springel & Hernquist \(2002\)](#). Our simulations use the cosmological parameters inferred from the *Planck* data combined with lensing, *WMAP* and high multipole measurements ([Planck Collaboration et al. 2014](#)): $\Omega_M = 0.305$, $\Omega_\Lambda = 0.694$, $\Omega_B = 0.0481$, $h = 0.679$, $\sigma_8 = 0.827$ and $n_s = 0.962$. These values are slightly updated with respect to the simulations presented in [Puchwein et al. \(2012\)](#) and [Lamberts et al. \(2015\)](#) but we do not expect any significant impact on the results presented here. In all cases, our simulations start with the same initial conditions at $z = 110$. The simulation domain has a comoving side length of $100 h^{-1} \text{ Mpc}$ and periodic boundary conditions. We use $N = 2 \times 1280^3$ particles, which yields a mass resolution of $m_{\text{gas}} = 1.5 \times 10^6 h^{-1} M_\odot$ and $m_{\text{DM}} = 7.2 \times 10^6 h^{-1} M_\odot$ for baryonic and dark matter particles, respectively. We used a comoving gravitational softening length of $3.9 h^{-1} \text{ kpc}$. The size of the box was chosen to cover the typical length scales of heating fluctuations, of tens of Mpc (see Fig. 2). We perform a set of simulations focusing on $z \geq 1.75$ for direct comparison with ground-based Ly α data.

3.2 Modelling the thermal state of the IGM

As this work focuses on the IGM, we use a simplified model for star formation in our simulations, in which gas particles with density $\delta_{\text{gas}} \equiv \rho_{\text{gas}} / \bar{\rho}_{\text{baryon}} - 1 \geq 1000$ and $T \leq 10^5 \text{ K}$ are directly converted into stars ([Viel et al. 2004](#)). Although this yields inaccurate galaxy properties, it does not affect the low-density IGM and significantly speeds up the simulations ([Bolton et al. 2017](#)). Photoionisation and photoheating is based on a [Haardt & Madau \(2012\)](#) UV background and the assumption of ionization equilibrium. The simulations in Paper I and [Puchwein et al. \(2012\)](#) were based on the [Faucher-Giguère et al. \(2009\)](#) UV background, which in the absence of blazar heating overpredicts the effective optical depth in the Lyman- α forest at all redshifts. However, the optical depths are rescaled during postprocessing to match the mean transmitted flux from observations by [Becker et al. \(2013\)](#). This rescaling aims to correct for uncertainties in the assumed photoionisation rate and strongly reduces differences when comparing the observables resulting from simulations with different UV backgrounds. The main residual effect is that slight changes in the photoheating result in somewhat different IGM thermal histories, which in turn affect the pressure smoothing in the simulation and the thermal broadening of predicted absorption lines.

In Fig. 3, we compare the thermal evolution of the IGM (measured at mean cosmic baryon density) in our different simulations,




Figure 1. Volume-weighted temperature - density relation at redshifts $z = 2$ and 1 for the simulations with no blazar heating (two left-hand panels) and with inhomogeneous blazar heating (two right-hand panels). The overlying black contours show the corresponding $T - \rho$ relation for uniform blazar heating (Puchwein et al. 2012) for the same redshift range. The color scale is logarithmic.




Figure 2. Distribution of the temperature in the midplane of the simulation domain at $z = 2$ for the unheated case (left), inhomogeneous heating case (middle) and the uniform heating model (right). The white dashed line shows a line of sight that we illustrate in more detail in Fig. 4

as well as to results from the literature. We here perform three simulations: one without any blazar heating, one with uniform blazar heating and one with inhomogeneous blazar heating. Due to the different choice of UV background our simulation without blazar heating is slightly hotter than the run without blazar heating in Puchwein et al. (2012). The difference is, however, very small compared to the (potential) impact of blazar heating. The blazar heating in our uniform heating model is identical to that in the “intermediate heating” model presented in Puchwein et al. (2012). The blazar heating results in much higher IGM temperatures at low redshift. Note however that the relative importance of blazar heating also depends on density.

Finally, we show two simulations from the literature in Fig. 3, the 40-2048 simulation from the Sherwood simulation suite (Bolton et al. 2017) and a simulation of the same volume from the Sherwood-Relics project (Puchwein et al. 2022, in prep.; also see Gaikwad et al. 2020b; Puchwein et al. 2019). These simulations do not include heating by blazars but account for a larger (and likely more realistic) amount of heat injection by He II reionisation. These runs will be discussed further in Sec. 5.

3.3 Producing mock Ly α spectra

The main output of our simulations are the mock Ly α spectra throughout the entire simulation volume. To obtain these, we extract 100 lines of sight from each output. The lines of sight are randomly selected and aligned with the main axes of the simulation. We use the same lines of sight for all the different heating models, but change the random selection from one redshift to the other. For each line of sight, we determine the density, velocity, temperature and resulting




Figure 3. Median IGM temperature of gas at mean cosmic baryon density ($0.975 < \Delta < 1.025$) as a function of redshift in our simulations (solid lines). For comparison results from a previous study investigating blazar heating of the IGM are shown (Puchwein et al. 2012, dashed lines). Also indicated is the thermal evolution of the IGM in the Sherwood and Sherwood-relics simulations, which do not consider blazar heating but account for a larger amount of heating from HeII reionisation.

optical depth of HI along 2048 equally spaced bins. The resulting dataset is the basis of our subsequent analysis.

An example is shown in Fig. 4, which displays the normalized




Figure 4. Illustration of the transmitted HI Ly α flux at $z = 2$ in the different heating models along a line of sight which includes both warm and cold regions. The location of the line of sight is shown on Fig. 2, and is chosen in the midplane of the domain. The optical depths have been rescaled to the mean transmission at the considered redshift.

transmitted HI Ly α flux along one line of sight chosen to illustrate the differences in our heating models. This specific line of sight first crosses a region which is heated in both blazar-heated models (first half of the spectrum, see Fig. 2 for the associated temperature map) and then a region which is heated in the uniform model but remains cold in the biased heating model (second half of the spectrum). This effect is clearly visible in the spectrum, where the Ly α flux has an intermediate value between both extreme models.

In order to allow direct comparison of our synthetic spectra with observed data, we rebin the dataset to the spectral resolutions of the different spectrographs and add noise with the same properties as the observations. The outputs of our high-redshift simulations are separated by $\Delta z = 0.1$, starting at $z = 3$. This allows us to cover the complete pathlength of a photon from $z = 3$ down to $z = 1.8$. For each output, we compute a comoving line-of-sight length consistent with the line-of-sight length covered by the data (Fig. 1 in Rudie et al. (2013)).

Details of the exact computation are provided in the relevant comparisons with observations. In the next section, we show how the shape of the temperature-density distribution and the size of the temperature fluctuations affect the observable properties of the Ly α forest and compare with the observational data.

4 COMPARISON WITH OBSERVATIONS

We compare our simulations with different heating models with observations of the Ly α forest. In this section, we focus on statistics that are sensitive to low density regions and new observational data. In the Appendix, we provide comparisons with the simulations presented in Puchwein et al. (2012) for reference.

4.1 Rescaled flux PDF

Rorai et al. (2017) recently proposed a new method to probe the thermal state of the low-density IGM. In this method the optical

depth is rescaled by a factor A . As a result the flux probability distribution function (PDF) is more focused on the high-transmission regions, which are likely corresponding to the low-density regions. Additionally, Rorai et al. (2017) rescale the transmitted flux to the value of the 95th percentile of the distribution, thus reducing the impact of the error in the continuum placement. Their comparison with simulations shows a better agreement with models with additional heating and/or inhomogeneous heating. To compare with the rescaled flux PDF by Rorai et al. (2017), we perform the exact same steps as in their analysis.

As described in section 3 of Rorai et al. (2017), we first rescale the optical depth so that the mean transmitted flux F matches the observed value ($\langle F_{\text{obs}} \rangle = 0.371$) at the central redshift of the observation ($\langle z \rangle = 2.75$), using data from Becker & Bolton (2013). We perform this operation at once for all lines of sight resulting from the same simulation output. Then, as described in their section 4.1, we smooth the flux with a Gaussian of full width at half maximum of 7.2 km s^{-1} (to model the the line spread function) and rebin it into bins of $\Delta v = 2.5 \text{ km s}^{-1}$ the pixel resolution of the UVES spectrograph). We then add Gaussian noise with $\sigma = (\sigma_0^2 + F(\sigma_c^2 - \sigma_0^2))^{1/2}$ with $\sigma_0 = 0.0028$ and $\sigma_c = 0.0088$ (see Appendix F in Rorai et al. (2017)). We then rescale the transmitted flux to the transmitted flux in the 95th percentile. The latter is close to the peak of the flux PDF, and is therefore less noisy than the mean. This operation is performed independently for each line of sight, in chunks of size $10 h^{-1} \text{ Mpc}$. Finally, to enhance the impact of low density regions, we rescale the optical depth by a factor $A = 10$, which yields a rescaled transmitted flux $F_A = F^A$.

To reconstruct the pathlength of the Ly α absorbers in the observed line-of-sight, we stitch together subsections of lines of sight from all outputs between $z = 3$ and $z = 2.5$. For each output, the size of the subsection is set by the comoving distance Δl corresponding to $\Delta z = 0.1$ for the considered redshift. The first pixel of the subsection is randomly chosen along one of the lines of sight and we use the periodic boundary conditions to complete the line of sight if the edge of the computational box is reached before Δl is completed. The lines of sight are randomly chosen for each output. While this produces small discontinuities at the junctions of the subsections, it is a more realistic representation of the varying cosmic structure that would be encountered by a photon (see e.g. Hummels et al. 2017, for a discussion). This method implicitly assumes that the removal of contaminants in the observed sightline (such as metals) does not affect the redshift distribution of the Ly α absorbers along the path covered by the quasar sightline.

Figure 5 shows the PDF of F_A in our three different heating models, compared with the observed PDF. As the observed line of sight is unique, cosmic variance can significantly affect the resulting flux PDF (Rollinde et al. 2013; Rorai et al. 2017). To provide some measurement of the spread of the flux PDF, we recreate a total of one hundred lines of sight from each of our simulations (following the method described above). The colored shaded areas show the resulting one and two sigma deviations from the mean. While we are confident that we do not oversample specific locations of the simulation, in Paper I we showed that the typical length scale of heated/unheated regions is of a few tens of Mpc/h. As such, the number of regions sampled by our simulation is somewhat limited and we likely underestimate cosmic variance.

Our simulations show that the rescaled flux PDF of the Ly α forest cannot be reproduced by a model without additional heating before or around $2.5 < z < 3$. Inhomogeneous blazar heating does provide a flux PDF compatible with observations. The uniform model, where the same total amount of blazar heating is injected,

does not agree equally well with the data. The model lies beyond 2σ of a third of the data points. This model has less physical motivation, as blazar heating is naturally inhomogeneous, because of the biased distribution of the sources.

For comparison, we also show results for the Sherwood-Relics simulation (purple line, leftmost plot of Fig. 5), which does not include blazar heating, but predicts more heating from HeII reionization. The latter is a consequence of using a more accurate non-equilibrium ionization/heating solver and a UVB that produces a realistic reionization history (Puchwein et al. 2019). The increased HeII heating brings the rescaled flux PDF in the absence of blazar heating in somewhat better agreement with the data, but not to the extent that our biased blazar heating model does. We will discuss this further in Sec. 5.

4.2 Properties of absorption lines

The statistical properties of the line-width b and column density N_{HI} of the individual Ly α absorption lines carry information about the density and temperature of the IGM. Schaye et al. (1999) determined that it can be used as a proxy for the temperature-density relation in the low-density regime by assuming that, for a given column density, the smallest Doppler width is only set by thermal broadening while higher values result from turbulent motions within the absorbers. Intrinsically, this method relies on a unique relation between temperature and density, whereas the inhomogeneous blazar heating model shows a wide distribution (see Fig. 1). Although the complete $b - N_{\text{HI}}$ distribution carries information on the physical properties of the IGM, often only the lower envelope of the distribution is considered (but see the recent work by Hiss et al. 2019).

We compare our simulations with the data from Rudie et al. (2012a), based on 15 quasar absorption lines, as part of the Keck Baryonic Structure Survey (Steidel et al. 2010, KBSS). The sample covers a total pathlength $\Delta z = 8.27$ (comoving pathlength $\Delta l = 26.9$) between $2.02 \leq z \leq 2.84$ with a mean redshift $\bar{z} = 2.34$. The data was obtained with the HIRES spectrograph, with a FWHM $\approx 7 \text{ km s}^{-1}$ and typical signal-to-noise ratio $\approx 50 - 250$. These observations increase the number of high signal-to-noise Ly α absorbers in that redshift range by a factor of 2. Thanks to this large sample size, we can consider the complete 2D distribution of the absorbers in the $b - N_{\text{HI}}$ plane, rather than limiting ourselves to their 1D projections. We specifically focus on low column densities, $\log(N_{\text{HI}}) \leq 14$, corresponding to low density optically thin absorbers. Those are the most likely not associated with galaxies, but with the IGM (Rudie et al. 2013). In this section, we compare our simulations with the Rudie et al. (2012a) data, and follow their analysis as closely as possible.

Before extracting the physical information from our mock spectra, we add mock noise based on the noise distribution from the data. The average SNR is 100, but varies between 50 and 250, with the highest values at high redshift. Based on the histograms of the SNR of the individual pixels of each quasar, we determine the redshift-dependent probability distribution function of the SNR. For each simulated line-of-sight we then randomly choose the SNR from the distribution according to its redshift. Overall, the resulting line properties are, however, in agreement with the line properties from mock data obtained with a simpler, fixed SNR set to 100 (see the third column in Fig. 9).

In order to determine the line widths b and column densities N_{HI} of the absorbers, each of them has to be identified and individually fit. Rudie et al. (2012a) use the VPFIT software (Carswell &

Webb 2014), which automatically finds the properties ($z, b, \log N_{\text{HI}}$) of the absorbers using a χ^2 -reducing method. Lines with $b \leq 8$ or $b \geq 100$ or relative errors larger than 50 per cent for the properties of the absorbers are discarded from the initial sample. To enable the closest possible comparison, we also use VPFIT to analyse our lines of sight. Given the size of our simulations, each line of sight contains too many absorbers (≥ 150), making the convergence of VPFIT very slow. Following Bolton et al. (2017), we divide each line of sight into 5 equally-sized chunks and analyse each of them separately. At the end, we concatenate the absorbers from the 5 sections into one list. We remove the absorbers within $\Delta v = 500 \text{ km s}^{-1}$ of the section edges. Visual inspection determined that this is a satisfactory way to eliminate boundary effects while maintaining most of the sample. We have tested cuts up to 1000 km s^{-1} and find no significant difference in our results. We perform the same additional cuts as Rudie et al. (2012b) and also use their outlier removal procedure aimed at removing narrow metal lines from their sample. Although our mock spectra do not contain metal lines by construction, we perform the same outlier removal procedure to mimic accidental removal of hydrogen lines and find it has limited effect.

To allow for a statistical comparison with the data, we produce 100 mock samples, each with a pathlength distribution $\Delta l(z)$ based on Fig. 1 in Rudie et al. (2013), i.e. similar to the observed sample. To reconstruct each individual mock sample, we appropriately combine randomly selected individual absorbers using a redshift-dependent probability for an absorber to arise per unit path length from our different output redshifts.

In Fig. 6 we show the complete $b - N_{\text{HI}}$ distribution in the three heating models compared with the data (right panel). From these plots it is clear that the complete 2D distribution is different in the three heating models and that one may lose information by considering only its 1D projections. The broad temperature-density distribution in the inhomogeneous heating case is reflected in a broader $b - N_{\text{HI}}$ distribution at low column densities, hence representing an intermediate case between the unheated and the uniformly heated case. By eye, the Rudie et al. (2012a) sample looks closest to the unheated model. Although the inhomogeneous heating model seems to reproduce the lower envelope of the observed distribution, the peak of both blazar heated distributions show an offset.

Figures 7 and 8 show a quantitative comparison of the Rudie et al. (2013) data and the heating models. We show the complete dataset from Rudie et al. (2013) as well as the subset where likely metal lines have been removed. The impact of the metal removing procedure on the global line properties is small. We show the total number of lines as a function of their linewidth (Fig. 7) and column density (Fig. 8) on the top rows. These quantities may be subject to systematic uncertainties in the reconstruction of the pathlength but should be less sensitive to uncertainties in including narrow or low-column density lines in the fits. We also show the probability distribution functions of the linewidth and column density (bottom rows), which are unaffected by uncertainties on the pathlength but their normalization can be affected by numerical resolution limits, difficulties in reliably fitting narrow or low-column density lines, and errors in continuum placement. Line fitting algorithms sometimes fail to identify and/or model lines with a low column density ($\log(N_{\text{HI}}) < 12$). We have, hence, shaded the corresponding area in our plots and results should be considered with caution there.

The distribution of the linewidths shows that the heated models lack narrow lines (below 25 km s^{-1}), while all models have similar numbers of broader lines. This leads to a global shift of




Figure 5. Probability distribution function of the rescaled flux in the unheated (left), inhomogeneously heated (middle) and uniformly heated model (right) between $2.5 < z \leq 3$. The thick line represents the mean over the 100 lines of sight, and the dark and light shaded area the 1 and 2σ (standard deviation) uncertainty intervals around the mean (indicating both sample variance and the effects of noise). The observed PDF is given with the black dots (Rorai et al. 2017). It is worth keeping in mind that it is based on a single line of sight with exquisite data quality. The left panel includes the mean value of the rescaled flux PDF from the Sherwood-Relics simulation (purple line). Following Rorai et al. (2017), no error bars are added to the data points as uncertainties are model dependent and computed separately for each simulation.




Figure 6. Line-width versus column density distribution in the three different heating models, compared with the data from Rudie et al. (2012a). The redshift range extends from $z = 2.02$ to $z = 2.84$ with a mean value $\bar{z} = 2.34$.

the PDF towards broader lines for heated models (lower row). The unheated model reproduces both the normalised and un-normalised data, while the heated models fail at reproducing them. The distribution of column densities show that heated models underproduce the total number of lines by roughly 50% near the peak of the distribution and that they somewhat overproduce systems with higher column densities. The unheated model reproduces both the normalised and un-normalised data, while the heated models fail at reproducing them.

5 DISCUSSION

In P12 we showed that observations of the Lyman α forest such as the distribution of linewidths and column densities as well as the flux PDF and powerspectrum around redshift 3 were in good agreement with predictions accounting for blazar heating. In this paper, we find that i) a rescaled flux PDF from a single high-resolution spectrum from Rorai et al. (2017) prefers models with inhomogeneous blazar heating at $z \sim 2.5 - 3$ and that ii) the line-width column density distribution from Rudie et al. (2012a) is not compatible with blazar heating at $z \sim 2.3$. At first glance these facts seem in contradiction

with each other, but it is important to keep in mind that these two measurements probe not only different redshifts, but also different densities. The goal of this section is to explain how these statements can be reconciled and how systematic uncertainties typically impact this type of analyses. We will discuss whether heating with a different redshift and/or density dependence (e.g., inhomogeneous HeII photoheating) could play a role and how to reconcile these results with evidence for blazar heating from other channels such as the gamma-ray sky.

5.1 Systematic uncertainties and comparison with P12

Here, we detail redshift sampling effects, signal-to-noise models, and resolution effects, all of which are illustrated in the linewidth and column density PDFs shown in Fig. 9.

The left panel of Fig. 9 illustrates how the linewidth and column density distributions evolve with redshift (darker red means increasing redshift). In the unheated model shown here, the whole column density distribution shifts by about half a dex between $z = 2$ and $z = 2.8$ (reflecting the significant evolution in the mean cosmic density and the opacity of the Lyman α forest), while the line width distribution is mostly unchanged (suggesting less evolution in




Figure 7. Normalised (bottom) and un-normalized (top) line-width distribution in the three different heating models, compared with the data from [Rudie et al. \(2012a\)](#). The black dots show the complete dataset and the cyan point shows the data resulting from the outlier removal procedure. The colored shaded areas show the one and two sigma confidence regions around the mean. The blue dashed lines shows the PDF with a resolution correction applied, which was derived from the Sherwood simulation suite (see §5). The redshift range extends from $z = 2.02$ to $z = 2.84$ with a mean value $\bar{z} = 2.34$ and the column density range is $[10^{11} - 10^{14.5}] \text{ cm}^{-2}$.




Figure 8. Normalised (bottom) and un-normalized (top) PDF of the column density between 10^{11} and $10^{14.5} \text{ cm}^{-2}$ in the three different heating models, compared with the data from [Rudie et al. \(2012a\)](#). The black dots show the complete dataset and cyan point shows the data resulting from the outlier removal procedure. The shaded areas show the one and two sigma confidence regions around the mean. The grey shaded area indicates regions where the line-fitting algorithms are less reliable, potentially affecting the results. The blue dashed lines shows the PDF with a resolution correction applied, which was derived from the Sherwood simulation suite (see §5). The redshift range extends from $z = 2.02$ to $z = 2.84$ with a mean value $\bar{z} = 2.34$.

absorber temperature²). In the inhomogeneous heating model (not shown here) the linewidth distribution is globally shifted by 5 km s^{-1} towards higher velocities at low redshift. These changes due to the redshift evolution are of comparable magnitude than those resulting from different heating models. Often observational results cover a range of redshifts, and the actual redshift distribution of the individual absorbers can vary significantly. Any comparison between observations and numerical simulations should therefore properly reconstruct the pathlength distribution, at the risk of otherwise introducing systematic shifts in the results.

In P12, the $z = 3$ simulated $b - N_{\text{HI}}$ distribution was compared with the observations from Kirkman & Tytler (1997) between $2.75 < z < 3.05$. The results are recalled here in the second panel of Fig. 9, showing that the observed distribution (purple dots) and uniform heating model (dashed blue line) are in good agreement. We also show that the simulations presented here (solid lines) are in good agreement with the simulations from P12 (dashed lines), for all the models. As such, the simulations presented here are fully compatible with the ones from P12. We show the less informative powerspectrum and flux PDF in Appendix A for completeness. The need for additional heating found in the $b - N_{\text{HI}}$ distribution around $z \approx 3$ is in agreement with our analysis of the rescaled flux PDF of Rorai et al. (2017), which covered the redshift interval $[2.5 - 3.0]$.

In this publication we considered a lower redshift region for the $b - N_{\text{HI}}$ distribution, focusing on the more recent data by Rudie et al. (2012b) between $2.0 < z < 2.8$. We modeled a fully reconstructed redshift distribution from the simulation and find that our blazar heated models do not agree with the data. If we restrict ourselves to the redshift range in which most of the data lies ($2.2 < z < 2.4$) and compare to the $z = 2.3$ snapshot, we find a qualitatively similar result as in our main analysis, as is shown in the third column of Fig. 9. While our results and the P12 results seem incompatible at first, this discrepancy can be explained by different redshifts in consideration. At high redshift ($z \approx 3$), additional heating of the IGM is needed, while at lower redshift ($z \approx 2.3$) additional heating is strongly disfavoured by the data. This point will be discussed further in Sec. 5.2.

The evolution of the SNR is another subtlety relating to redshift. Results are typically presented with a mean SNR but important variations can be present. As an example, the data in Rudie et al. (2012a) has a typical SNR of 100 but high redshift lines can have an SNR as high as 250 and low redshift lines can go as low as 50. We have modeled the redshift-dependent noise-distribution to match the observational data. The dashed lines in the third column of Fig. 9 show the distributions with the complete noise model in comparison with the average SNR = 100 model at $z = 2.3$. At this redshift, which is the mode of the redshift distribution, the difference is minimal and does not affect the comparison between simulations and observations presented here. However, we emphasize that this should be checked systematically, especially when one focuses on subsets of the data, which may have a systematically smaller/larger SNR than the bulk of the data.

Finally, the last row of Fig. 9 shows the impact of numerical resolution on the line parameters, based on the Sherwood simulation data at $z = 3$. Given the large size of the regions heated by blazars (see Fig. 2), our simulation domain covers $100 h^{-1} \text{ Mpc}$ on the side, which limits its effective mass resolution. Bolton et al. (2017) present a detailed study of resolution effects, based on their high

resolution data (2048^3 particles with a boxlength of $40 h^{-1} \text{ Mpc}$ on the side, dashed line) and a lower resolution simulation (512^3 particles, dash-dotted line), which has an effective resolution similar to ours. The increased resolution results in narrower lines, shifted by a few km s^{-1} and has almost no impact on the PDF of the column densities considered here. A similar study at $z = 2$ (not shown here) shows a smaller impact of resolution. Globally, resolution effects are smaller than the differences between our different heating models and are not expected to affect the conclusions of this work significantly.

Several additional systematic uncertainties can affect this work, and we have checked that their impact is smaller than the differences between the heating models. Our $100 h^{-1} \text{ Mpc}$ on the size simulation domain provides a good representation of large scale structure in the Universe, up to group sized objects, which could hosts some of the blazars we are considering, and the impact of cosmic variance should be limited. The 15 lines-of-sight composing the Rudie et al. (2012b) data also provide a good coverage of cosmic variance. On the other hand, the Rorai et al. (2017) study is only based on a single line of sight, and effects of cosmic variance may be important. To mitigate this we have bootstrapped our data and computed 100 mock spectra, allowing for a meaningful comparison. Additionally, as shown in Figs. 7-8, the uncertainty due to metal contamination in observational data is small. Also, we have tested that there are only minor differences in line properties between the line-fitting algorithm VPFIT which we have used here and AUTOVP (Davé et al. 1997) which was used in P12. Finally, errors in the continuum placement can affect the comparison between our models and the line properties from Rudie et al. (2012b). Given the high SNR of the observations, only absorbers with $\log(N_{\text{HI}}) < 12.5$ would be affected by this. This could potentially slightly skew the PDF of the linewidths and/or column densities but would let the un-normalized distributions mostly unaffected (top rows of Figs. 7-8). Specifically, the strong preference for the unheated model in the distribution of the column density will be unchanged. In conclusion, we find no systematic uncertainties which could qualitatively change our main conclusions.

5.2 Redshift evolution and physical mechanism of IGM heating

The seemingly contradictory results found in this study could be reconciled if significant heating of the low-density IGM occurs at $z \gtrsim 3$, and if the IGM cools down towards lower redshifts. This redshift trend is, however, unexpected for heating by TeV blazars (see Fig. 3).

Photoheating associated with the reionisation of HeII is instead expected to deposit energy into the IGM during the epoch of HeII reionisation ($z \gtrsim 2.7$, Worseck et al. 2016), while the IGM is expected to cool in the absence of blazar heating due to Hubble expansion and Compton cooling once helium is (almost) fully ionized. Compared to our simulations, homogeneous HeII heating is modelled more accurately in the Sherwood-Relics simulations, which use a non-equilibrium ionization/heating solver as well as a cosmic UV background that has been carefully matched to observational constraints on the reionisation history (Puchwein et al. 2019). The corresponding thermal evolution of the IGM is illustrated in Fig. 3. A temperature maximum is reached at $z \sim 3$, while the temperature falls off at lower redshifts. Indeed the thermal evolution at mean density is similar to our blazar-heated models at $z \gtrsim 3.5$ and close to our model without blazar heating at $z \lesssim 2$. At the redshifts probed by the rescaled flux PDF of Rorai et al. (2017), the temperature in this

² Note, however, that lines of fixed column density probe different overdensities at different redshifts.

model with increased HeII heating falls roughly half way between our models with and without blazar heating (although somewhat closer to the latter at the low densities, $\Delta \lesssim 0.5$, the rescaled flux PDF is most sensitive to). This is reflected in the rescaled flux PDF it predicts (gray curve in the left panel of Fig. 5). The deviation from the data is only about half of that in our simulation without blazar heating. Our inhomogeneously blazar-heated model is, however, in significantly better agreement with the data.

Nevertheless, these findings suggest that a model in which HeII heating occurs at similar redshifts as in the Sherwood-Relics simulations, but with an enhanced amplitude, might match all the data sets considered in this work. Additional heating would be primarily needed in the underdense gas probed by the rescaled flux PDF, while the temperature of denser gas probed by other techniques seems to be in better agreement with observational constraints (see, e.g., Gaikwad et al. 2020a). A much larger amount of HeII heating is, however, not expected in most models, e.g., all the models considered in Puchwein et al. (2019) and Upton Sanderbeck & Bird (2020) predict temperatures at mean density that are at least a few thousand Kelvin lower at $z \sim 3$ than the blazar-heated models considered here. In principle, radiative transfer effects during HeII reionization can specifically boost the temperature in the lowest density regions as they are often reionized last and have hence the least time to cool afterwards. The most extreme radiative transfer model in La Plante et al. (2017) (their simulation H5) comes very close in the temperature at mean density and has a fairly flat (but not inverted) temperature-density relation at lower densities. It is, hence, at this point not entirely clear whether the high IGM temperatures at low densities favoured by the rescaled flux PDF data can be realistically reached with an efficient HeII heating, or whether an additional heating channel such as provided by TeV blazars is needed.

Fig. 3 illustrates that the difference in IGM temperature between models with and without blazar heating increases strongly at $z \lesssim 2.5$. Furthermore, blazar heating has the largest impact at the lowest densities (see Fig. 1). Hence, we would ideally like to probe the IGM thermal state in underdense regions at lower redshifts. Unfortunately, the hydrogen Ly α forest becomes almost completely transparent there. Prominent Ly α absorbers correspond to increasingly dense systems at lower redshifts. Getting the maximum sensitivity to low density regions requires very high signal-to-noise data and analyzing regions that are almost transparent, which is essentially what was done in the analysis of Rorai et al. (2017) for higher redshifts. Repeating such an analysis with similarly good data at $z \sim 2$ to 2.5 should be a very promising way of constraining blazar heating with increased sensitivity. It would also be important to reduce the cosmic variance in the observed data set by using more lines of sight, compared to the single one that was studied in Rorai et al. (2017). In Fig. 10, we show predictions of the rescaled flux PDF for this redshift range the data is provided in Table B1 in the Appendix for possible comparison with future observations or simulations. As expected the difference between the different heating models is larger at these lower redshifts. Data from future observational studies could be compared to this.

Another interesting approach would be to use HeII rather than hydrogen for absorption line studies of the thermal state of the IGM. The Ly α line of HeII also produces a forest of absorption lines, however, with a different redshift evolution and density dependence. Despite the lower abundance of helium compared to hydrogen, the HeII Ly α forest is more opaque due to the smaller number of photons that can ionize HeII, as well as due to the faster recombination of HeII. Unfortunately, the HeII Ly α forest needs to be observed from space and requires very bright background quasars without

intervening Lyman limit systems that would absorb the relevant part of the spectrum. There is, hence, much less data available with typically lower signal-to-noise ratios. Nevertheless, it would be interesting to explore in a future work to what extent it can constrain blazar heating of the IGM.

5.3 A unified picture of blazar heating: Ly α forest, gamma-ray sky and clustering

Blazar heating results from heating of the IGM by plasma instabilities of electron-positron beams created by pair production from TeV gamma-rays. There are strong lines of evidence from gamma-ray observations that plasma instabilities drain the kinetic energy of electron-positron pairs that are produced by very high energy photon scattering off of ~ 1 eV photons from the extragalactic background light. In particular, an otherwise expected component of secondary inverse Compton-scattered gamma-rays is not observed, neither directly toward the BL-Lac sources nor when observing potential pair beams from the side. This limits intergalactic magnetic fields, which have been invoked to explain a deflection of the secondary component out of the beam direction, both from below and above, respectively. Limits of $B \gtrsim 10^{-14} - 10^{-15}$ G (see for instance Neronov & Vovk 2010; Ackermann et al. 2018) and $B \lesssim 10^{-15} - 10^{-16}$ G (Broderick et al. 2018; Tiede et al. 2020) have been derived, where the latter use a novel analysis that explores the anisotropic nature of the gamma-ray halos (Broderick et al. 2016; Tiede et al. 2017). These mutually exclusive constraints preclude magnetic deflection of pairs as a mechanism of suppression of the secondary inverse Compton emission and suggest that plasma instabilities may instead be at work.

In addition, the principle plasma instability that is driven by these pair beams, the so-called ‘‘oblique instability’’ (Bret et al. 2004; Broderick et al. 2012; Chang et al. 2016) appears to be robust in the conditions of the intergalactic medium against both nonlinear effects (Chang et al. 2014) and large scale density gradients (Shalaby et al. 2018, 2020). This is not to say that the principle instability that is driven by pair beams is the ‘‘oblique instability’’, but generic arguments for the suppression of plasma instabilities are erroneous (Miniati & Elyiv 2013).

The results of the observed *lack* of blazar heating in the Ly α forest at $z \lesssim 2.5$ are in tension with the lines of evidence from the gamma-ray sky. One possibility is that the plasma unstable modes that are driven by these pair beams do not ultimately end up as thermal heat, but remain in small scale magnetic fields and plasma waves. The extremely-long-term evolution of these modes have never been studied and much of our intuition of the behavior of these plasma instabilities are interpolated from a much more prosaic region of parameter space.

Another possibility is that the redshift evolution for blazar heating is different from what we have assumed in this work and that the blazar luminosity function may drop off faster toward high redshift, $z \gtrsim 1.5 - 2$. In our previous work, we rescaled the quasar luminosity function to produce the blazar luminosity function by noting that the blazar luminosity distribution appears to be a rescaling of the quasar luminosity distribution at $z = 0.1$ (Broderick et al. 2012; Chang et al. 2012). However, the evidence for the presence of IGM beam-plasma instabilities and the absence of the re-processed inverse Compton-scattered GeV gamma-ray photons accumulates mostly at low redshift ($z \lesssim 1.5$), whereas the evidence from the Ly α is from high redshift ($z \gtrsim 2.5$).

In particular, the number of resolved TeV blazars probe very low redshifts ($z \lesssim 0.5$). We demonstrate in Broderick et al. (2014)




Figure 9. Illustration of systematic effects impacting line width (top row) and column density (bottom row) probability distribution functions. First column: redshift evolution, from light to deeper red in the unheated model. Second column: high redshift simulations show consistency with the P12 results and the heated model matches the Kirkman & Tytler (1997) data. Third column: low-redshift simulations and impact of noise model. The unheated model is preferred by the Rudie et al. (2012b) data. Final column: Impact of increased resolution illustrated by the Sherwood simulations.




Figure 10. Predicted probability distribution function of the rescaled flux in the unheated (left), inhomogeneously heated (middle) and uniformly heated model (right) between $2 < z \leq 2.5$. The thick line represents the mean over the 100 lines of sight, and the red and pink area the 1 and 2 σ uncertainty around the mean. The corresponding numerical values are provided in Tab. B1 in the Appendix. The dashed line in the center and right plot show the mean in the unheated model for easier comparison.

that a positively evolving distribution of hard gamma-ray blazars (which are similar to high-frequency synchrotron peaked BL Lacs) in combination with the presence of virulent plasma beam instabilities that preempt the inverse Compton cascade provides an excellent match to the observational data, in particular the hard gamma-ray blazar redshift distribution and $\log N$ - $\log S$ distribution (where N and S are the number and flux of nearby hard gamma-ray-bright blazars in the *Fermi* gamma-ray band, respectively). Moreover, more than 90 per cent of the isotropic extragalactic gamma-ray background at energies $E_\gamma > 10$ GeV (Ackermann et al. 2015) is accounted for by hard gamma-ray blazars at redshifts $z \lesssim 2$

(see figure 9 of Broderick et al. 2014) with the remainder likely contributed by starburst galaxies (Ackermann et al. 2012). The low-energy ($E_\gamma < 10$ GeV) isotropic extragalactic gamma-ray background is likely dominated by unresolved soft gamma-ray blazars and flat-spectrum radio quasars (FSRQs), which are known to show a positive redshift evolution with increasing z (Ajello et al. 2014).

However, this fortuitous relationship between the blazar and quasar luminosity distribution (Broderick et al. 2012) may only exist at low redshift ($z \lesssim 1.5$) and diverge at higher redshift. One line of argument that this might be the case is that the analysis of the two-point correlation function of BL-Lacs and FSRQs from *Fermi*

suggest that these objects occupy dark matter halos of $\sim 2 \times 10^{13} M_{\odot}$ at least up to $z \sim 1$ (Allevato et al. 2014). This is inline with the halos hosting radio loud SDSS quasars between $z = 0.3 - 1.3$ (Shen et al. 2009; Krumpe et al. 2012). On the other hand, optical quasars are associated with dark matter haloes of $\sim 10^{12} M_{\odot}$ (Shen et al. 2009; Allevato et al. 2014). Because more massive dark matter halos that host BL-Lacs and FSRQs form later in a hierarchically growing universe (Wechsler et al. 2002) and in consequence are much rarer at higher redshift, this would suggest that TeV emission and, hence, blazar heating falls off rapidly at higher redshift.

6 CONCLUSIONS

We have performed an in-depth comparison between mock data from large-scale cosmological simulations from different blazar heating models and observations of the Ly α forest between $2 < z < 3$. In addition to previous work by Puchwein et al. (2012), we have considered a more physically motivated inhomogeneous blazar heating model aside from the uniform heating model. This work was also motivated by a recently published very high resolution rescaled flux PDF for $2.5 < z < 3$ (Rorai et al. 2017) and a set of high resolution quasar absorption lines around $\langle z \rangle = 2.3$ (Rudie et al. 2013). The latter enables us to compare our simulations with lower redshift observations than was done in P12. We find that:

- Inhomogeneous blazar heating is in very good agreement with observations of the low-density IGM at redshift $2.5 < z < 3$. We base this on a comparison between our simulations and a very high resolution spectrum where the global optical depth was rescaled to enhance the contribution from underdense regions (Rorai et al. 2017), which are most sensitive to blazar heating. Additional rescaling of the global transmitted flux also allowed mitigating continuum placement uncertainties, thereby making this measurement possible.

- The linewidth and column density distributions determined by Rudie et al. (2013) at redshift $\langle z \rangle = 2.3$ are incompatible with our blazar heated models and agree well with our model without additional heating.

- At first glance, our results seem in tension with P12, which concluded in favor of the presence of blazar heating. This can be explained by the different redshifts that were being considered in both publications. Most of the data favoring blazar heating in P12 was at $z > 2.5$, while we perform a more detailed comparison to lower redshift ($z < 2.5$) data here. We show that both results are not incompatible and that carefully matching the redshift sampling of predictions to the observational data is important.

- We discuss whether heating from HeII reionisation alone could explain the rescaled flux PDF data. While the redshift evolution required for the heating is consistent with expectations for HeII photoheating, the amount of heating is insufficient for most models of HeII reionisation. It may work, however, for the most extreme models present in the literature, in particular if radiative transfer effects can enhance heating of low density regions.

- Overall, the comparisons from this paper call for a modification of blazar heating of the IGM as described in Broderick et al. (2012); Chang et al. (2012); Pfrommer et al. (2012). However, observations of the gamma-ray sky and theoretical work on plasma instabilities does provide evidence of plasma instabilities at work. Whether the energy is eventually transferred through other channels than heating or whether the evolutionary history of blazar heating needs to be revised remains a question. Clustering studies of blazars provide a hint that the blazar luminosity redshift distribution decreases

faster towards high redshift in ~ 2 comparison to the quasar luminosity distribution, thus providing circumstantial evidence that the latter explanation may at least explain part of the effect found in this study.

Our thorough comparison between mock data and observations has provided us with new insight on blazar heating and the thermal history of the IGM in a broader sense. A firmer answer on (a potentially delayed) blazar heating will rely on additional data, covering a wider range of redshifts. We particularly advocate for an analysis similar to Rorai et al. (2017) focused on lower redshift ($z < 2.5$). Studies of the HeII Lyman- α forest may be another interesting avenue to explore to get a better handle on the thermal state of the very low density IGM. In all cases, we find that discriminating comparisons between simulations and observations are only possible when both datasets are well described and documented, thus reducing systematic uncertainties (e.g. redshift sampling, signal-to-noise model...).

ACKNOWLEDGEMENTS

The authors warmly thank James Bolton for his insightful comments on this work, as well as for making line lists from the Sherwood simulations available to us. A.L. thank Alberto Rorai for many clarifications he provided on his work. CP and MS acknowledge support by the European Research Council under ERC-CoG grant CRAGSMAN-646955. PC acknowledges support by the NASA ATP program through NASA grant NNH17ZDA001N-ATP. This work was supported in part by Perimeter Institute for Theoretical Physics. Research at Perimeter Institute is supported by the Government of Canada through the Department of Innovation, Science and Economic Development Canada and by the Province of Ontario through the Ministry of Economic Development, Job Creation and Trade. PT receives support from the Natural Science and Engineering Research Council through the Alexander Graham Bell CGS-D scholarship. A.E.B. thanks the Delaney Family for their generous financial support via the Delaney Family John A. Wheeler Chair at Perimeter Institute. A.E.B. receives additional financial support from the Natural Sciences and Engineering Research Council of Canada through a Discovery Grant. The Sherwood and Sherwood-Relics simulations were performed using the Curie and Irene supercomputers at the Tres Grand Centre de Calcul (TGCC), and the DiRAC Data Analytic system at the University of Cambridge. The latter is operated by the University of Cambridge High Performance Computing Service on behalf of the STFC DiRAC HPC Facility (www.dirac.ac.uk). This equipment was funded by BIS National E-infrastructure capital grant (ST/K001590/1), STFC capital grants ST/H008861/1 and ST/H00887X/1, and STFC DiRAC Operations grant ST/K00333X/1. DiRAC is part of the National E-Infrastructure.

DATA AVAILABILITY

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

REFERENCES

- Ackermann M., et al., 2012, *ApJ*, **755**, 164
 Ackermann M., et al., 2015, *ApJ*, **799**, 86
 Ackermann M., et al., 2018, *ApJS*, **237**, 32

- Ajello M., et al., 2014, *ApJ*, **780**, 73
- Allevato V., Finoguenov A., Cappelluti N., 2014, *ApJ*, **797**, 96
- Becker G. D., Bolton J. S., 2013, *MNRAS*, **436**, 1023
- Becker G. D., Bolton J. S., Haehnelt M. G., Sargent W. L. W., 2011, *MNRAS*, **410**, 1096
- Becker G. D., Hewett P. C., Worseck G., Prochaska J. X., 2013, *MNRAS*, **430**, 2067
- Boera E., Murphy M. T., Becker G. D., Bolton J. S., 2014, *MNRAS*, **441**, 1916
- Bolton J. S., Viel M., Kim T.-S., Haehnelt M. G., Carswell R. F., 2008, *MNRAS*, **386**, 1131
- Bolton J. S., Becker G. D., Haehnelt M. G., Viel M., 2014, *MNRAS*, **438**, 2499
- Bolton J. S., Puchwein E., Sijacki D., Haehnelt M. G., Kim T.-S., Meiksin A., Regan J. A., Viel M., 2017, *MNRAS*, **464**, 897
- Bret A., Firpo M.-C., Deutsch C., 2004, *Phys. Rev. E*, **70**, 046401
- Broderick A. E., Chang P., Pfrommer C., 2012, *ApJ*, **752**, 22
- Broderick A. E., Pfrommer C., Puchwein E., Chang P., 2014, *ApJ*, **790**, 137
- Broderick A. E., Tiede P., Shalaby M., Pfrommer C., Puchwein E., Chang P., Lamberts A., 2016, *ApJ*, **832**, 109
- Broderick A. E., Tiede P., Chang P., Lamberts A., Pfrommer C., Puchwein E., Shalaby M., Werhahn M., 2018, *ApJ*, **868**, 87
- Calura F., Tescari E., D'Odorico V., Viel M., Cristiani S., Kim T.-S., Bolton J. S., 2012, *MNRAS*, **422**, 3019
- Carswell R. F., Webb J. K., 2014, Astrophysics Source Code Library, Chang P., Broderick A. E., Pfrommer C., 2012, *ApJ*, **752**, 23
- Chang P., Broderick A. E., Pfrommer C., Puchwein E., Lamberts A., Shalaby M., 2014, *ApJ*, **797**, 110
- Chang P., Broderick A. E., Pfrommer C., Puchwein E., Lamberts A., Shalaby M., Vasil G., 2016, *ApJ*, **833**, 118
- Compostella M., Cantalupo S., Porciani C., 2013, *MNRAS*, **435**, 3169
- Davé R., Hernquist L., Weinberg D. H., Katz N., 1997, *ApJ*, **477**, 21
- Fan X., Carilli C. L., Keating B., 2006, *ARA&A*, **44**, 415
- Faucher-Giguère C.-A., Lidz A., Zaldarriaga M., Hernquist L., 2009, *ApJ*, **703**, 1416
- Gaikwad P., Srianand R., Haehnelt M. G., Choudhury T. R., 2020a, arXiv e-prints, p. [arXiv:2009.00016](https://arxiv.org/abs/2009.00016)
- Gaikwad P., et al., 2020b, *MNRAS*, **494**, 5091
- Haardt F., Madau P., 2012, *ApJ*, **746**, 125
- Hiss H., Walther M., Oñorbe J., Hennawi J. F., 2019, arXiv e-prints, Hopkins P. F., Richards G. T., Hernquist L., 2007, *ApJ*, **654**, 731
- Hui L., Gnedin N. Y., 1997, *MNRAS*, **292**, 27
- Hummels C. B., Smith B. D., Silvia D. W., 2017, *ApJ*, **847**, 59
- Kim T.-S., Bolton J. S., Viel M., Haehnelt M. G., Carswell R. F., 2007, *MNRAS*, **382**, 1657
- Kirkman D., Tytler D., 1997, *ApJ*, **484**, 672
- Krumpe M., Miyaji T., Coil A. L., Aceves H., 2012, *ApJ*, **746**, 1
- Kulkarni G., Keating L. C., Haehnelt M. G., Bosman S. E. I., Puchwein E., Chardin J., Aubert D., 2019, *MNRAS*, **485**, L24
- La Plante P., Trac H., Croft R., Cen R., 2017, *ApJ*, **841**, 87
- Lamberts A., Chang P., Pfrommer C., Puchwein E., Broderick A. E., Shalaby M., 2015, *ApJ*, **811**, 19
- Lee K.-G., 2012, *ApJ*, **753**, 136
- Lynds R., 1971, *ApJ*, **164**, L73
- McQuinn M., Upton Sanderbeck P. R., 2016, *MNRAS*, **456**, 47
- McQuinn M., Lidz A., Zaldarriaga M., Hernquist L., Hopkins P. F., Dutta S., Faucher-Giguère C.-A., 2009, *ApJ*, **694**, 842
- Meiksin A., Tittley E. R., 2012, *MNRAS*, **423**, 7
- Miniati F., Elyiv A., 2013, *ApJ*, **770**, 54
- Neronov A., Vovk I., 2010, *Science*, **328**, 73
- Palanque-Delabrouille N., et al., 2013, *A&A*, **559**, A85
- Palanque-Delabrouille N., et al., 2015, *JCAP*, **11**, 011
- Pfrommer C., Chang P., Broderick A. E., 2012, *ApJ*, **752**, 24
- Planck Collaboration et al., 2014, *A&A*, **571**, A16
- Puchwein E., Pfrommer C., Springel V., Broderick A. E., Chang P., 2012, *MNRAS*, **423**, 149
- Puchwein E., Bolton J. S., Haehnelt M. G., Madau P., Becker G. D., Haardt F., 2015, *MNRAS*, **450**, 4081
- Puchwein E., Haardt F., Haehnelt M. G., Madau P., 2019, *MNRAS*, **485**, 47
- Rafighi I., Vafin S., Pohl M., Niemiec J., 2017, *A&A*, **607**, A112
- Rauch M., et al., 1997, *ApJ*, **489**, 7
- Rollinde E., Theuns T., Schaye J., Pâris I., Petitjean P., 2013, *MNRAS*, **428**, 540
- Rorai A., et al., 2017, *MNRAS*, **466**, 2690
- Rorai A., Carswell R. F., Haehnelt M. G., Becker G. D., Bolton J. S., Murphy M. T., 2018, *MNRAS*, **474**, 2871
- Rudie G. C., et al., 2012a, *ApJ*, **750**, 67
- Rudie G. C., Steidel C. C., Pettini M., 2012b, *ApJ*, **757**, L30
- Rudie G. C., Steidel C. C., Shapley A. E., Pettini M., 2013, *ApJ*, **769**, 146
- Schaye J., Theuns T., Leonard A., Efstathiou G., 1999, *MNRAS*, **310**, 57
- Schaye J., Theuns T., Rauch M., Efstathiou G., Sargent W. L. W., 2000, *MNRAS*, **318**, 817
- Schlickeiser R., Ibscher D., Supsar M., 2012, *ApJ*, **758**, 102
- Schlickeiser R., Krakau S., Supsar M., 2013, *ApJ*, **777**, 49
- Shalaby M., Broderick A. E., Chang P., Pfrommer C., Lamberts A., Puchwein E., 2017a, *ApJ*, **841**, 52
- Shalaby M., Broderick A. E., Chang P., Pfrommer C., Lamberts A., Puchwein E., 2017b, *ApJ*, **848**, 81
- Shalaby M., Broderick A. E., Chang P., Pfrommer C., Lamberts A., Puchwein E., 2018, *ApJ*, **859**, 45
- Shalaby M., Broderick A. E., Chang P., Pfrommer C., Puchwein E., Lamberts A., 2020, *Journal of Plasma Physics*, **86**, 535860201
- Shen Y., et al., 2009, *ApJ*, **697**, 1656
- Shull J. M., Smith B. D., Danforth C. W., 2012, *ApJ*, **759**, 23
- Sironi L., Giannios D., 2014, *ApJ*, **787**, 49
- Slosar A., et al., 2013, *JCAP*, **4**, 026
- Springel V., 2005, *MNRAS*, **364**, 1105
- Springel V., Hernquist L., 2002, *MNRAS*, **333**, 649
- Steidel C. C., Erb D. K., Shapley A. E., Pettini M., Reddy N., Bogosavljević M., Rudie G. C., Rakic O., 2010, *ApJ*, **717**, 289
- Tiede P., Broderick A. E., Shalaby M., Pfrommer C., Puchwein E., Chang P., Lamberts A., 2017, *ApJ*, **850**, 157
- Tiede P., Broderick A. E., Shalaby M., Pfrommer C., Puchwein E., Chang P., Lamberts A., 2020, *ApJ*, **892**, 123
- Upton Sanderbeck P., Bird S., 2020, arXiv e-prints, p. [arXiv:2002.05733](https://arxiv.org/abs/2002.05733)
- Vafin S., Rafighi I., Pohl M., Niemiec J., 2018, *ApJ*, **857**, 43
- Viel M., Haehnelt M. G., Springel V., 2004, *MNRAS*, **354**, 684
- Viel M., Bolton J. S., Haehnelt M. G., 2009, *MNRAS*, **399**, L39
- Wechsler R. H., Bullock J. S., Primack J. R., Kravtsov A. V., Dekel A., 2002, *ApJ*, **568**, 52
- Worseck G., et al., 2011, *ApJ*, **733**, L24
- Worseck G., Prochaska J. X., Hennawi J. F., McQuinn M., 2016, *ApJ*, **825**, 144

APPENDIX A: COMPARISON WITH PUCHWEIN ET AL, 2012

Here we compare the transmitted flux PDF and power spectrum from the simulations presented here with previous work presented in P12. Given the systematic uncertainties which can affect the interpretation of both observational and simulated data (see §5) we only compare simulations with each other and do not attempt to compare with observational data here. To allow for an exact comparison, we use outputs at the exact same redshift from both simulation sets and only present the redshift range relevant to the work presented here.

Figure A1 shows the one-dimensional power spectrum of the Ly α forest in the simulations, similarly to the right panel in Fig. 8 of P12. Both sets of simulations show good agreement for the different models although slightly more power is present at small scales in P12. This is likely due to the higher resolution of the P12 simulations. Figure A2 shows the transmitted flux PDF in both sets of




Figure A1. Power spectra $P_{\text{ID}}(k)k/\pi$ of transmitted flux contrast $\exp(-\tau)/\langle \exp(-\tau) \rangle - 1$ in the three heating models presented here (solid lines) and the uniform and unheated models presented in P12 at redshift $z = 2.5$ (dashed lines). Aside from rescaling to the mean transmitted flux (Becker et al. 2013), no treatment has been applied to the simulations and a comparison with observations is not straightforward.

simulations, similarly to the right panel in Fig. 7 of P12. Both sets show a very good agreement.

These comparisons show that the simulations presented here are very consistent with those presented in P12. However, this work and P12 lead to different conclusions regarding the blazar heating of the low density IGM. This emphasizes the need to use more sensitive observational constraints (such as the linewidth and column density distributions or rescaled flux PDF) and to consider a wide redshift range to understand the thermal evolution of the IGM.

APPENDIX B: PREDICTED RESCALED FLUX PDF

Table B1 shows the rescaled ($A=10$) flux pdf between $z = 2 - 2/5$ predicted by various heating models.

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




Figure A2. Regular ($A = 1$) flux PDF in the three heating models presented here (solid lines) and the uniform and unheated model presented in P12 at redshifts $z = 2.5$ and 3. Aside from rescaling to the mean transmitted flux (Becker et al. 2013), no treatment has been applied to the simulations and a comparison with observations is not straightforward.

bin center	PDF _{noheat}	σ_{noheat}	PDF _{biased}	σ_{biased}	PDF _{uniform}	σ_{uniform}
-9×10^{-3}	4.47	0.31	4.77	0.326	4.75	0.346
0.05	1.11	0.09	0.92	0.078	0.85	0.084
0.11	0.73	0.055	0.57	0.051	0.54	0.052
0.17	0.64	0.059	0.48	0.049	0.45	0.044
0.23	0.6	0.051	0.45	0.041	0.42	0.041
0.29	0.59	0.053	0.45	0.045	0.41	0.036
0.35	0.59	0.05	0.45	0.046	0.42	0.04
0.41	0.61	0.051	0.47	0.046	0.43	0.042
0.47	0.64	0.051	0.49	0.044	0.46	0.037
0.52	0.66	0.051	0.53	0.045	0.52	0.044
0.60	0.69	0.048	0.6	0.047	0.58	0.047
0.66	0.74	0.057	0.69	0.053	0.68	0.053
0.72	0.79	0.061	0.81	0.056	0.82	0.057
0.77	0.82	0.065	0.97	0.067	1.02	0.075
0.84	0.8	0.071	1.1	0.085	1.20	0.110
0.90	0.74	0.063	1.09	0.083	1.19	0.096
0.96	0.6	0.054	0.88	0.055	0.94	0.06
1.02	0.43	0.032	0.55	0.025	0.57	0.027

Table B1. Numerical values for Fig. 10 presenting the rescaled flux PDF ($A=10$) for $2 < z \leq 2.5$ and its standard deviation in the three heating models.