

A new measurement of the intergalactic temperature at $z \sim 2.55\text{--}2.95$

Alberto Rorai,^{1,2★} Robert F. Carswell,^{2★} Martin G. Haehnelt,^{1,2★} George D. Becker,³
James S. Bolton⁴ and Michael T. Murphy⁵

¹*Kavli Institute for Cosmology and Institute of Astronomy, Madingley Road, Cambridge CB3 0HA, UK*

²*Institute of Astronomy, Madingley Road, Cambridge CB3 0HA, UK*

³*Department of Physics and Astronomy, University of California, Riverside, 900 University Avenue, Riverside, CA 92521, USA*

⁴*School of Physics and Astronomy, University of Nottingham, University Park, Nottingham NG7 2RD, UK*

⁵*Centre for Astrophysics and Supercomputing, Swinburne University of Technology, Hawthorn, Victoria 3122, Australia*

Accepted 2017 November 2. Received 2017 October 4; in original form 2017 July 4

ABSTRACT

We present two measurements of the temperature–density relationship (TDR) of the intergalactic medium (IGM) in the redshift range $2.55 < z < 2.95$ using a sample of 13 high-quality quasar spectra and high resolution numerical simulations of the IGM. Our approach is based on fitting the neutral hydrogen column density $N_{\text{H I}}$ and the Doppler parameter b of the absorption lines in the Ly α forest. The first measurement is obtained using a novel Bayesian scheme that takes into account the statistical correlations between the parameters characterizing the lower cut-off of the b – $N_{\text{H I}}$ distribution and the power-law parameters T_0 and γ describing the TDR. This approach yields $T_0/10^3 \text{ K} = 15.6 \pm 4.4$ and $\gamma = 1.45 \pm 0.17$ independent of the assumed pressure smoothing of the small-scale density field. In order to explore the information contained in the overall b – $N_{\text{H I}}$ distribution rather than only the lower cut-off, we obtain a second measurement based on a similar Bayesian analysis of the median Doppler parameter for separate column-density ranges of the absorbers. In this case, we obtain $T_0/10^3 \text{ K} = 14.6 \pm 3.7$ and $\gamma = 1.37 \pm 0.17$ in good agreement with the first measurement. Our Bayesian analysis reveals strong anticorrelations between the inferred T_0 and γ for both methods as well as an anticorrelation of the inferred T_0 and the pressure smoothing length for the second method, suggesting that the measurement accuracy can in the latter case be substantially increased if independent constraints on the smoothing are obtained. Our results are in good agreement with other recent measurements of the thermal state of the IGM probing similar (over-)density ranges.

Key words: intergalactic medium – quasars: absorption lines.

1 INTRODUCTION

The intergalactic medium (IGM), containing the overwhelming majority of the Universe’s baryons, retains key information about the cosmic transformations that occurred during helium and hydrogen reionization (e.g. Hui & Gnedin 1997; Gnedin & Hui 1998; Theuns, Mo & Schaye 2001; Theuns et al. 2002; Hui & Haiman 2003). Many authors have studied the signature of thermal heating caused by the ionizing photons in the absorption profiles of the Ly α forest, with the goal of probing the temperature–density relation (TDR) of the IGM at $z = 2\text{--}4$ (e.g. Haehnelt & Steinmetz 1998; Schaye et al. 1999; McDonald et al. 2001; Zaldarriaga, Hui & Tegmark 2001; Lidz et al. 2010). These works have been motivated by the simple form that the low-density TDR should take according to theoretical pre-

dictions, and by the potential implications for the history of reionization. Analytical studies and hydrodynamics simulations have indicated that the low-density gas in the IGM should be concentrated in a narrow region of the temperature–density plane along a power law

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

where T_0 is the temperature at mean density, Δ is the density divided by the mean of the Universe and γ is the index of the power-law relation (Hui & Gnedin 1997). The shape of the TDR at different redshifts is dependent on the timing of reionization, on the nature of the sources and physical mechanisms responsible for the heating. If photoheating of residual neutral hydrogen is the dominant heat source, then it is predicted that $\gamma \approx 1.6$ well after reionization (Hui & Gnedin 1997; McQuinn & Upton Sanderbeck 2016).

Although several statistical analyses of the Ly α forest find values of γ close to this prediction (Rudie, Steidel & Pettini 2012; Bolton

* E-mail: arorai@ast.cam.ac.uk (AR); rfc@ast.cam.ac.uk (RFC); haehnelt@ast.cam.ac.uk (MGH)

Table 1. List of spectra analysed in this work. Listed, from left- to right-hand side, are the object name, redshift, Ly α absorption redshift range used, the instrument used for observation, the signal-to-noise range (per pixel) at the continuum level, the reduced data pixel size, and either the ESO program number (UVES) or the source of the data (HIRES).

Object	z_{em}	Δz	Source	S/N	pixel size	ESO program/reference
PKS 2126–158	3.28	0.0939	UVES	40–90	2.5 km s $^{-1}$	166.A-0106(A)
Q0347–383	3.23	0.1485	UVES	50–55	2.5 km s $^{-1}$	68.B-0115(A)
J134258–135559	3.21	0.2545	UVES	30–50	2.5 km s $^{-1}$	68.A-0492(A)
HS1425+6039	3.18	0.2356	HIRES	55–75	0.04 A	Sargent
Q0636+6801	3.17	0.3152	HIRES	35–60	0.04 A	Sargent
J210025–064146	3.14	0.1356	HIRES	20–25	1.3 km s $^{-1}$	KODIAQ O’Meara et al. (2015)
Q0420–388	3.12	0.2884	UVES	75–120	2.5 km s $^{-1}$	166.A-0106(A)
HE0940–1050	3.08	0.2910	UVES	35–70	2.5 km s $^{-1}$	166.A-0106(A)
GB1759+7539	3.05	0.1859	HIRES	27–32	2.0 km s $^{-1}$	Outram, Chaffee & Carswell (1999)
J013301–400628	3.02	0.1929	UVES	25–35	2.5 km s $^{-1}$	69.A-0613(A),073.A-0071(A),074.A-0306(A)
J040718–441013	3.02	0.0962	UVES	40–115	2.5 km s $^{-1}$	68.A-0361(A),68.A-0600(A),70.A-0017(A)
J224708–601545	3.00	0.3293	UVES	35–75	2.5 km s $^{-1}$	075A-158(B)
HE2347–2547	2.88	0.2205	UVES	95–125	2.5 km s $^{-1}$	077.A-0646(A),166.A-0106(A)

et al. 2014; Boera et al. 2016), other measurements, which consider the Ly α flux probability distribution function (PDF), have claimed evidence for an inverted ($\gamma < 1$) TDR (Bolton et al. 2008; Viel, Bolton & Haehnelt 2009; Calura et al. 2012; Garzilli et al. 2012) for which unconventional heating mechanisms such as blazar heating have been invoked as an explanation (Broderick, Chang & Pfrommer 2012; Chang, Broderick & Pfrommer 2012; Pfrommer, Chang & Broderick 2012; Puchwein et al. 2012; Lamberts et al. 2015). Some authors have suggested that this discrepancy could be ascribable to unaccounted effects from systematic uncertainties due to ‘continuum fitting’ of QSO absorption spectra necessary for the calculation of the flux PDF (Lee 2012) or to an overestimation of the statistical significance of the measurements (Rollinde et al. 2013). An alternative solution to reconcile the apparent discrepancies between the measurements and the expected thermal state of a photo-heated IGM was proposed in Rorai et al. (2017b), who analysed the PDF of the low-opacity pixels in a very high signal-to-noise quasar spectra in order to constrain the TDR in the low-density IGM. Rorai et al. (2017b) found that uncertainties in the continuum placement alone cannot explain the discrepancy with conventional models for the thermal state of the IGM. Instead, they found that a flat or inverse TDR (with high temperature in underdense regions) is indeed favoured by the PDF, though perhaps only at very low densities ($\Delta \lesssim 1$). They also showed that different Ly α forest statistics that give discrepant results, like the power spectrum or those based on line-fitting methods, are sensitive to disjoint density ranges ($\Delta \gtrsim 2$ –3). Rorai et al. (2017b) thus challenged the description of the low-density TDR as a single spatially invariant power law.

To investigate the low density TDR further, here we undertake a traditional Voigt-profile fitting decomposition of the Ly α forest absorption in the redshift range $2.55 < z < 2.95$ for a sample of 13 quasar spectra. We use a set of IGM models obtained by post-processing high resolution hydrodynamics simulations and generate model spectra with the same noise and resolution characteristics. We apply the same Voigt decomposition to these spectra so that they may then be compared directly with the telescope data. Following Schaye et al. (1999, 2000), Ricotti, Gnedin & Shull (2000), Rudie et al. (2012) and Bolton et al. (2014), we analyse the shape of the cut-off for narrow lines in the plane defined by the column density $N_{\text{H I}}$ and the line width b . Moreover, we introduce a Bayesian formalism to study not only the uncertainties on the inferred thermal parameters, including the pressure smoothing, but

also their degeneracies. We further develop a new technique based on the medians of the b distribution for separate column-density ranges, in order to exploit the information contained in the bulk of the distribution in the $N_{\text{H I}}-b$ plane.

This article is structured as follows. We start by presenting in Section 2 the sample of quasar spectra we use in our analysis and how these data are treated, in particular with respect to metal contaminants. We then describe the hydrodynamics simulations and the models to which we compare the data (Section 3). In Section 4, we explain how we analyse the statistical properties of Ly α lines in the forest to extract the information about the thermal properties of the IGM. The results of this analyses are illustrated in Section 5 and subsequently discussed in Section 6, where we also examine agreements and disagreements with previous studies. We draw our final conclusions in Section 7.

2 DATA

A sample of high signal-to-noise ratio quasar spectra from ESO Ultraviolet and Visual Echelle Spectrograph (UVES, Dekker et al. 2000) and Keck high resolution echelle spectrograph (HIRES, Vogt et al. 1994) archival data with coverage of the Ly α forest in the redshift range $2.55 < z < 2.95$ was selected. This redshift range is chosen to complement the reanalysis by Bolton et al. (2014) of the data presented by Rudie et al. (2012), which provided constraints on the TDR parameters at lower redshift. Note that the methods illustrated in this paper can, in principle, be applied to higher redshift data, but we found that at $z > 3.2$, the stronger blending of Ly α lines makes the decomposition into Voigt profiles increasingly ambiguous. To have a large segment of the chosen absorption redshift range clear of the quasar proximity region (chosen to be within 4000 km s^{-1} of the quasar redshift) requires that the emission redshift $z_{\text{em}} \gtrsim 2.85$, and the need to avoid Ly β blending with the Ly α forest leads to $z_{\text{em}} \lesssim 3.30$. A list of the 13 objects used, and the characteristics of the reduced spectra, is given in Table 1. The exposures for the nine UVES spectra were reduced using the European Southern Observatory (ESO) UVES CPL (Common Pipeline Language) software (v4.2.8) and combined using UVES_popl (Murphy 2016), as described in detail in Boera et al. (2014). The two spectra kindly provided by W. L. W. Sargent were reduced using MAKEE (Barlow 2008). Continuum estimates in the forest are based on fitting low order curves to the high points in the spectra.

Here, we wish to compare the results of fitting Voigt profiles to the Ly α forest using vPFIT (Carswell & Webb 2014) to the observational spectra with those from simulated data (see Section 3), so we have to be aware of some features in the data that cannot be reproduced in the simulations, and some restrictions the simulations may place on the way the observations are analysed. These are as follows:

(i) The resolution of the object spectra is not accurately known, since the observations were not always slit limited and nor would the seeing have been constant. Here, we assume a Gaussian resolution element with a full width at half-maximum (FWHM) of 6.5 km s^{-1} , which is a reasonable approximation for both the HIRES and UVES data. Most Ly α features have Doppler parameters $b \gtrsim 15 \text{ km s}^{-1}$, and for sample inclusion, we choose $b > 8 \text{ km s}^{-1}$ (corresponding to FWHM 13.3 km s^{-1}), so even 10 per cent uncertainties in the instrumental FWHM do not make a significant difference. We therefore convolve all simulated spectra with a Gaussian kernel with FWHM = 6.5 km s^{-1} .

(ii) The continuum estimates are based on large-scale high points in the observational data, and maybe in error. vPFIT allows a linear continuum adjustment as a function of wavelength over the fitting region, and inserts this adjustment automatically when the overall fit accuracy comes down to a specified level. To be consistent, this was used both for the observational data and for the simulations.

(iii) vPFIT occasionally introduces very large Doppler parameter lines, which appear to be better described as long range continuum adjustments. To remove these, we omitted features with Doppler parameters $> 100 \text{ km s}^{-1}$.

(iv) The zero level may be offset by a small amount in the observed spectra while it is accurately known in the simulated ones. A zero level adjustment was introduced during the automatic fitting process when the normalized χ^2 became < 5 if there were five or more contiguous pixels of the fitted profile below 5×10^{-3} of the continuum value. The details are given in Sections 6.5 and 7.4 of the vPFIT documentation (Carswell & Webb 2014).

(v) Heavy element lines contaminate the Ly α forest in the observational data, but are absent in the simulations. We identified as many as we could from systems that showed heavy element lines longward of the Ly α emission and then chose wavelength regions within the forest to avoid the stronger ones. Heavy element lines are usually narrow, so the 8 km s^{-1} threshold adopted for this analysis will remove most of the ones we have not identified. Also in the cases where metal lines are clustered in groups, they are still fitted as separate narrow components for the signal-to-noise ratios in our data sample.

(vi) The simulated spectra cover a fixed small range of just over 1000 km s^{-1} (or 15.5 \AA at redshift $z = 2.75$), and all lines in each of these spectra were fitted simultaneously. In the observed data, vPFIT was applied to regions of varying size, depending on the local line density and the positions of heavy element lines, but were chosen to be between 10 and 25 \AA long, with an average of $\sim 16 \text{ \AA}$.

(vii) The flux noise in the observational data is not constant from object to object, or even within an object spectrum, where it depends on the signal. The continuum level noise, σ_c , may be estimated by interpolating between regions of the spectrum where there is little absorption, and the zero level noise, σ_0 , by doing so between saturated Ly α lines. Then for the simulations the noise can be set using $\sigma = \sqrt{\sigma_0^2 + F(\sigma_c^2 - \sigma_0^2)}$, where F is the transmitted flux normalized by the continuum. We identified 53 spectral sections characterized by different (σ_c , σ_0) pairs. To account for this, the simulated sight-lines are divided into 53 subsets with path length proportional to the path lengths of the 53 sections. To each of them

we add flux-dependent Gaussian noise using the appropriate value of σ_c and σ_0 .

(viii) Since flux noise is estimated on a pixel-by-pixel basis, we rebin the simulated flux in pixels of 2.5 km s^{-1} , which is the pixel scale used in most of the data sample. For simplicity, in the few cases where the pixel size, Δv , of the data is not 2.5 km s^{-1} , σ_0 and σ_c of the corresponding simulated sets are rescaled by $\sqrt{\Delta v/2.5}$ before adding noise. We have checked, for a sample of spectra, that this rescaling procedure produces consistent fits with those in which the same pixel size and noise level as the data was used.

(ix) vPFIT adds as many components as necessary until a satisfactory fit is obtained. However, it may fail to converge to give an acceptable fit to either observational or model data, sometimes, for example, if there is a saturated Ly α line. Where this happens, the fits from those regions are omitted. Model spectral regions were chosen with noise characteristics mimicking the observational ones with acceptable fits and, in cases where convergence failed, different sightlines through the model were chosen until an acceptable fit was obtained.

All fitted components are characterized by their central redshift z , column density $N_{\text{H I}}$ (in cm^{-2}) and Doppler parameter b (in km s^{-1}), and vPFIT provides estimates and uncertainties for all these quantities. The observational data yielded 2271 fitted Ly α components in a total path length of $\Delta z = 2.788$. A potential problem with the approach taken to avoid metal line contamination is that possible C IV and Mg II doublets may occur without counterparts redwards of the Ly α emission line and would not be identified inside the forest. To assess the impact of these contaminants we have compiled a sample of lower redshift quasars ($2.13 < z_{\text{em}} < 2.54$) and identified C IV and Mg II doublets lying in the range $4316\text{--}4802 \text{ \AA}$, i.e. the range covered by the Ly α data (the same data were recently analysed by Kim et al. 2016). For consistency, we only consider doublets with no associated lines from lower ions (and longer wavelength transitions), which would have been detected in the first place by our identification process. The Voigt-profile fit parameters from this sample are then used to add the opacity of these contaminants to the simulated spectra and estimate the impact on our results. We have checked that by doing so the results of this paper are not significantly affected, once our chosen cuts on b and $N_{\text{H I}}$ are applied (see below).

3 MODELS

In order to predict the observed statistical properties of the Ly α forest, we used simulated spectra from the set of hydrodynamics simulations described in Becker et al. (2011, hereafter B11). The simulations were run using the parallel Tree-smoothed particle hydrodynamics (SPH) code GADGET-3, which is an updated version of the publicly available code GADGET-2 (Springel 2005). The fiducial simulation volume is a $10 \text{ Mpc } h^{-1}$ periodic box containing 2×512^3 gas and dark matter particles. This resolution is chosen specifically to resolve the Ly α forest at redshift $z \sim 4\text{--}5$ (Bolton & Becker 2009). The simulations were all started at $z = 99$, with initial conditions generated using the transfer function of Eisenstein & Hu (1999). The cosmological parameters are $\Omega_{\text{m}} = 0.26$, $\Omega_{\lambda} = 0.74$, $\Omega_b h^2 = 0.023$, $h = 0.72$, $\sigma_8 = 0.80$, $n_s = 0.96$, consistent with constraints of the cosmic microwave background from WMAP9 (Reichardt et al. 2009; Jarosik et al. 2011). The IGM is assumed to be of primordial composition with a helium fraction by mass of $Y = 0.24$ (Olive & Skillman 2004). The gravitational softening length was set to 1/30th of the mean linear interparticle spacing,

and star formation was included using a simplified prescription that converts all gas particles with overdensity $\Delta = \rho/\bar{\rho} > 10^3$ and temperature $T < 10^5$ K into collisionless stars. In this work, we will only use the outputs at $z = 2.735$.

The gas in the simulations is assumed to be optically thin and in ionization equilibrium with a spatially uniform ultraviolet background (UVB). The UVB corresponds to the galaxies and quasars emission model of (Haardt & Madau 2001, hereafter **HM01**). Hydrogen is reionized at $z = 9$, and gas with $\Delta \lesssim 10$ subsequently follows a tight power-law TDR, $T = T_0 \Delta^\gamma - 1$, where T_0 is the temperature of the IGM at mean density (Hui & Gnedin 1997; Valageas, Schaeffer & Silk 2002). As in **B11**, the photo-heating rates from **HM01** are rescaled by different constant factors, in order to explore a variety of thermal histories. Here, we assume the photoheating rates $\epsilon_i = \xi \epsilon_i^{\text{HM01}}$, where ϵ_i^{HM01} are the **HM01** photoheating rates for species $i = [\text{H I}, \text{He I}, \text{He II}]$ and ξ is a free parameter. Note that, different from **B11**, we do not consider models where the heating rates are density-dependent. In fact, we vary ξ with the only purpose of varying the degree of pressure smoothing in the IGM, while the TDR is imposed in post-processing. In practice, we only use the hydrodynamics simulation to obtain realistic density and velocity fields. For this reason, we will refer to ξ as the ‘smoothing parameter’. We then impose a specific TDR on top of the density distribution, instead of assuming the temperature calculated in the original hydrodynamics simulation. This means that in our models the temperature is only a function of the density, strictly following equation (1) at all densities up to $\Delta = 10$. As done in Rorai et al. (2017b), we set the temperature to be constant at higher densities, i.e. $T(\Delta > 10) = T(\Delta = 10)$, in order to avoid unphysically high values. Note however that such densities correspond to strongly saturated Ly α absorbers, which are not used in our analysis (see below). We opt for this strategy in order to explore a wide range of parametrizations of the thermal state of the IGM, at the price of reducing the temperature–density diagram of the gas to a deterministic relation between T and ρ . In this work, we use a total of 107 models based on hydrodynamics simulations with $\xi = 0.3, 0.8$ and 1.45 . The grid of parameters spans values between 0.4 and 1.9 for γ and between 5000 and $35\,000$ K for T_0 .

Finally, we calculate the optical depth to Ly α photons for a set of 1024 synthetic spectra in each model, assuming that the gas is optically thin, taking into account peculiar motions and thermal broadening. We scale the UV background photoionization rate in order to match the observed mean flux of the forest at the central redshift of the sample¹ [$\bar{F}_{\text{obs}}(z = 2.75) = 0.7371$; Becker et al. 2013].

We stress that in this scheme the pressure smoothing and the temperature are set independently. Although not entirely physical, this allows us to separate the impact on the Ly α forest from instantaneous temperature, which depends mostly on the heating at the current redshift, from pressure smoothing, which is a result of the integrated interplay between pressure and gravity across the whole thermal history (Gnedin & Hui 1998).

3.1 Parametrization of the pressure smoothing

Varying the smoothing parameter ξ allows us to test the effect of different thermal histories on the structure of the IGM density field. In order to characterize it in a model-independent way, we adopt

¹ In reality the average transmission of the Ly α forest evolves throughout this redshift bin. We have verified, however, that modelling the forest using a single value for the mean flux over this redshift range has only a small effect on the results when compared to the statistical uncertainties (see appendix).

the definition proposed by Kulkarni et al. (2015) for the pressure smoothing length in hydrodynamics simulations λ_p . This is based on the *real-space* Ly α flux F_{real} , calculated as the transmitted flux of Ly α photons in the fluctuating Gunn-Petersson approximation

$$F_{\text{real}}(x) = \exp \left[-\frac{3\lambda_\alpha^3 \Lambda_\alpha}{8\pi H(z)} n_{\text{H I}} \right], \quad (2)$$

where $\lambda_\alpha = 1216 \text{ \AA}$ is the rest-frame Ly α wavelength, Λ_α is the Einstein A coefficient of the transition, $H(z)$ is the Hubble parameter and $n_{\text{H I}}(x)$ is the neutral hydrogen number density at the point x . Critically, F_{real} does not include the effects of thermal broadening or peculiar velocities on the transmitted flux. In the optically thin regime, where $n_{\text{H I}} \propto \rho^2$, this is a non-linear transformation of the density field that suppresses high densities, but preserves isotropy. In Kulkarni et al. (2015), it is shown that the F_{real} power spectrum is sensitive to the thermal history of the low-density IGM, and is well approximated by the function

$$\Delta_F(k) = Ak^n \exp[-(k\lambda_p)^2], \quad (3)$$

where k is the Fourier wavenumber and A , n and λ_p are the free parameters. We calculate the F_{real} field for our hydrodynamics simulations, using the $n_{\text{H I}}$ output at $z = 2.735$. Consistently to what we have done for the post-processed models described in the previous section, we rescale $n_{\text{H I}}$ so that the mean flux in a set of synthetic Ly α spectra extracted from the simulations matches the observed values at the correspondent redshift. We obtain λ_p from the fit of the power spectrum. This gives $\lambda_p = 69, 93$ and 108 kpc for $\xi = 0.3, 0.8$ and 1.45 , respectively. Here and elsewhere in this paper, we always express the smoothing length in comoving kpc.

4 METHODOLOGY

Simulated spectra for various IGM models were produced to be compared with the data, and we derived Doppler parameter b and column densities $N_{\text{H I}}$ for the Ly α lines using `VPFIT` in exactly the same way as for the telescope data (see Sections 2 and 3). The distribution of lines in the b – $N_{\text{H I}}$ plane forms the basis for our statistical analysis. For both the data and the simulations, there may be velocity structure which is not well represented by Voigt profiles. In the fitting process, this can result in a number of non-physical components, which can be quite close together, or in blended, broad but weak lines. A feature of these is that the error estimates tend to be large as a consequence of the presence of neighbouring systems. An example can be seen in Fig. 1, showing the results of applying the algorithm to a simulated spectrum. To remove many of these, and to prevent the b – $N_{\text{H I}}$ distributions being dominated by noise, we require the relative error estimate in the Doppler parameter to be smaller than 50 per cent, and the error on $\log N_{\text{H I}}$ smaller than 0.3. Additionally, for $\log N_{\text{H I}} > 14.5$, the Ly α line is usually saturated to a level where the column densities derived from fitting only the Ly α transition are unreliable. So we exclude systems with higher column densities from the analysis. Finally, for the highest b -values in the range, the line detection limit is $\log N_{\text{H I}} \sim 12.5$, so we adopt this as a lower limit for the comparison samples. To summarize, our statistical analysis is based on the absorption components with $8 < b < 100 \text{ km s}^{-1}$, $12.5 < \log(N_{\text{H I}}/\text{cm}^{-2}) < 14.5$, Doppler parameter relative error smaller than 50 per cent and \log column density error < 0.3 . With these restrictions, the observational data sample consists of 1625 points in the b – $N_{\text{H I}}$ space. The line distribution for the data and for three different models is shown in Fig. 2.




Figure 1. Decomposition of the Ly α forest in individual absorption lines. Top panel: the black histogram represents the Ly α absorption in a sight line from one of our models shown in velocity space. The average signal-to-noise level of this spectrum is 33.2. The red dashed curves are the individual fitted absorbers, whereas the total reconstructed absorption is shown as the blue dot–dashed curve. The red dotted line follows the continuum ($F = 1$). Bottom panel: the deviation of the reconstructed flux F_{fit} from the original spectrum F , compared with the assumed noise level for this synthetic spectrum (delimited by the red solid lines). Overall, the combined fit is a good approximation to the original spectrum, except in regions characterized by flat absorption profiles (for example around $v = 200\text{--}300\text{ km s}^{-1}$), which cannot be decomposed into individual lines. Note also that rather broad lines are sometimes required to reproduce the absorption profile, which are not obviously interpretable as single absorbers (e.g. in the case of the line centred around $v \sim 450\text{ km s}^{-1}$). The parameters of these broad lines are also particularly sensitive to systematic uncertainties due to the continuum placement.




Figure 2. Distributions in the $\log N_{\text{HI}}\text{--}b$ plane for the lines fitted in the quasar data sample (left-hand panel) and in three IGM models with $T_0 = 25000\text{ K}$, $\gamma = 1.6$ (centre left-hand panel), $T_0 = 15000\text{ K}$, $\gamma = 1.6$ (centre right-hand panel) and $T_0 = 15000\text{ K}$, $\gamma = 1$ (right-hand panel). All three models have smoothing parameter $\xi = 0.8$, corresponding to a smoothing length $\lambda_p = 93\text{ kpc}$ (comoving). The red solid lines represent the fitted power law to the cut-off of the distribution, which is sensitive to the thermal parameters of the IGM. The behaviour of these lines in the three models illustrates the sensitivity of the cut-off to the thermal parameters, which has been used in the past to constrain the thermal parameters (e.g. Schaye et al. 1999; Rudie et al. 2012), and the difference in T_0 between the second and third panels determines a clear change in the intercept of the fitted line. Conversely, the lower slope of the TDR used in the right-hand panel, which shows an isothermal model, makes the cut-off much flatter than in the other two cases.

4.1 The lower cut-off in the $b\text{--}N$ distribution

As noted previously (Schaye et al. 1999; Theuns, Schaye & Haehnelt 2000; McDonald et al. 2001), the $b\text{--}N$ distribution shows a pronounced cut-off at low values of b . This is usually interpreted as a signature of thermal broadening setting a lower limit to the absorption line widths. The position of this cut-off is dependent on the column density, suggesting that the temperature systematically varies with the density of the gas. This motivated several studies to use the slope and normalization of the lower of the $b\text{--}N$ distribution as a direct probe of the TDR of the IGM. Pressure smoothing also broadens the lines by increasing the physical size of absorbers, which increases their velocity width due to the Hubble flow. This effect can be taken into account by means of theoretical/analytical arguments or hydrodynamics simulations (Schaye et al. 1999; Theuns et al. 2000; Garzilli, Theuns & Schaye 2015).

The lower cut-off in the $b\text{--}N$ distribution is generally assumed to be a power-law relation between line width and column density,

$$b = b_0 \left(\frac{N_{\text{HI}}}{N_{\text{HI},0}} \right)^{\Gamma-1}, \quad (4)$$

where b_0 and Γ are the parameters connected with the TDR and $N_{\text{HI},0}$ is a reference value which is often chosen as the column density corresponding to gas at the mean density $\Delta = 1$ (see Bolton et al. 2014, for a discussion). In Fig. 2, we show the fitted cut-off for our observed absorption line sample and samples of simulated absorption lines for three different thermal models. In our analysis, we arbitrarily fix $N_{\text{HI},0} = 10^{13}\text{ cm}^{-2}$, but we will relax the assumptions on the functional dependencies between b_0 , Γ and the thermal parameters. The standard algorithm used to fit this cut-off, which we also adopt, was first introduced by Schaye et al. (1999) and is based on a recursive rejection process: The expression in




Figure 3. Calibrated relations between the parameters b_0 and Γ describing the normalization and slope of the lower cut-off in the b - N distribution and the thermal parameters T_0 and γ describing the TDR. Left-hand column: cut-off normalization $\log b_0$ at $\log N_{\text{H I}, 0}/\text{cm}^{-2} = 13$ as a function of $\log T_0$ (upper panel) and slope of the fitted Γ as a function of the index γ of the TDR (lower panel). Each point represents a different model in the grid and errors are estimated by bootstrapping the simulated line samples. The red bands delimit the values of $\log b_0$ and Γ measured from the data within the 1σ error. The calibrated function $\log b_0(\log T_0)$ is obtained via a linear fit of these points (dashed lines). The two relations shown in this plot are reasonably tight, but the level of scatter appears higher than what the error bars warrant, suggesting that dependencies on other parameters are not negligible. This can have an effect on the inference of the thermal parameters. Right-hand column: same as the left-hand panel, but using the generalized combination of thermal parameters described in Section 4.2. Each point represents the mean of the MCMC posterior distribution for the combinations used to describe $\log b_0$ (upper panel) and Γ (lower panel). The (small) horizontal error bars represent the propagated uncertainties on the combination coefficients as obtained from the posterior distribution. The dashed lines represent the identity. Adopting the generalized scheme substantially improves the fits: the chi-square for the four plots, divided by the number of points, is 3.49 (top left-hand panel), 2.70 (top right-hand panel), 4.16 (bottom left-hand panel) and 2.56 (bottom right-hand panel).

equation (4) is fitted to the line distribution to obtain $b_{\text{fit}}(N_{\text{H I}})$. We then calculate the mean absolute deviation $|\sigma_b|$ of the points from the fit in the b dimension, and discard all lines whose value of b is *greater* than the fitted relation by more than $|\sigma_b|$. The fit and the rejection of upper outliers are iteratively repeated until convergence, i.e. when all points lie below the upper mean deviation and the only outliers are below the fit. At that stage, the lines with $b < b_{\text{fit}}(N_{\text{H I}}) - |\sigma_b|$ are discarded and the fit is repeated one last time to define the final values of b_0 and Γ . The errors associated with $\log b_0$ (it is convenient to operate in logarithmic space) and Γ are estimated via a bootstrap technique applied to the line sample. In a first approximation, we follow the standard practice of considering $\log b_0$ and Γ as statistically independent and treat them separately. Correlations between these parameters are addressed in the next section.

One may notice from Fig. 2 that the lines fitted in the data sample (left-hand panel) present more outliers below the cut-off compared to the models (right-hand three plots). To understand what could generate this difference, we visually inspected all individual absorbers in the data with width lower than the cut-off by more than $\log b = 0.1$. Among 29 outliers, we found 7 that are compatible with being unidentified metal lines and other 7 that are Ly α absorbers

associated with metals at the same redshift. The latter are metal-enriched systems whose temperature might be affected by metal cooling, which is not accounted for in our models. We have tested the effect of removing these 14 lines and to re-apply the cut-off fitting procedure. This did not significantly change the estimated cut-off parameters, but it reduced the uncertainty on $\log b_0$ by about 20 per cent. For this reason, we consider keeping all the outliers to be the most conservative choice.

In the published literature, the relations between the narrow line cut-off parameters and the thermal parameters (T_0 , γ and in this work also λ_p) were based on theoretical arguments set out by Schaye et al. (1999) (see also Rudie et al. 2012) or calibrated with hydrodynamics simulations (Bolton et al. 2014). Here, we will first make the ansatz that such relations can be written in the form

$$\log T_0 = A + B \log b_0, \quad (5)$$

$$\gamma - 1 = C + D(\Gamma - 1), \quad (6)$$

where A , B , C , D are determined using the full set of models described in Section 3 (see Fig. 3). Analogous to previous works, we for now assume no dependency on the pressure smoothing length λ_p , which for standard hydrodynamics simulations implicitly assumes

that the thermal history is correct and hence consistently taken into account. Note that these relations are slightly different than what is usually assumed in other works. We will later check/verify the validity of this ansatz *a posteriori* using the model grid (see Fig. 3).

In the next section, we discuss how we abandon these assumptions in order to include a more general relation between the parameters describing the lower cut-off of the b - N distribution and parameters describing the thermal state and history of the IGM, including the smoothing length λ_P .

4.2 Generalizing the calibration of the cut-off parameters

The relations described by equations (5) and (6) are based on the heuristic argument that the renormalization of the TDR changes the line widths at all column densities by a similar proportion, at least for lines that fall near the cut-off. Although Fig. 3 suggests that such an approximation is reasonable, we would like to consider the possibility of more general dependencies in order to find a scheme to quantify the effect of the pressure smoothing on the line distribution.

We approach this problem by starting with the assumption that a generic observable ϕ can be approximated by a linear combination of the logarithm of the relevant parameters. More precisely, we define the combination ψ as

$$\psi = a + b \log T_0 + c(\gamma - 1) + d \log \lambda_P, \quad (7)$$

where a, b, c, d are free parameters. The choice of using logarithmic quantities (except for γ , which is an index) is equivalent to assuming a power-law relation with T_0 and λ_P . Note that changing the units in which T_0 and λ_P are expressed would only affect the offset a . In general, ϕ could be a more general function of the parameters than a linear relation, so the validity of this approximation will need to be verified *a posteriori*. If ψ is not a good description of the observable, or if there are statistical uncertainties in the models (e.g. due to finite box size), it is natural to expect some scatter. This can be accounted for by introducing some flexibility in choosing $\{a, b, c, d\}$.

A convenient way of doing so is to implement a likelihood probability for fitting ϕ with ψ and to include it in a Markov chain Monte Carlo (MCMC) algorithm in the space of coefficients. This can be combined with a likelihood of the measured observable ϕ_d when compared to the value of ψ for a given choice of the coefficients and of the thermal parameters. This allows us to draw quantitative inferences on the calibration coefficients $\{a, b, c, d\}$ and on the thermal parameters $\{T_0, \gamma, \lambda_P\}$ at the same time.

In our analysis, the total likelihood is therefore composed of two parts. The first part quantifies how well the parameter combination ψ defined by a, b, c, d fits the points in the training grid of models. Given the observable ϕ , this can be written as

$$\log L_1 = - \sum_i \frac{[a + b \log T_{0,i} + c(\gamma_i - 1) + d \log \lambda_{P,i} - \phi_i]^2}{2\sigma_{\phi,i}^2}, \quad (8)$$

where the sum is performed over all the models in the grid and $\sigma_{\phi,i}^2$ is the uncertainty on ϕ_i , both relative to the i th model. An MCMC run that employs L_1 as likelihood suffices to estimate the posterior distribution for the calibration coefficients a, b, c, d , i.e. the range of parameter combinations that can be used to relate the observable ϕ to the IGM parameters via equation (7). In principle, this distribution could be used as a prior for a second MCMC run, which includes the thermal parameters. However, we find it more practical to combine $\log L_1$ with a likelihood where the same values

of a, b, c, d are employed to make a prediction for $\phi(T_0, \gamma, \lambda_P)$ and test it against the data. Assuming the error is Gaussian, this can be simply expressed as

$$\log L_2 = - \frac{[a + b \log T_0 + c(\gamma - 1) + d \log \lambda_P - \phi_d]^2}{2\sigma_{\phi,d}^2}, \quad (9)$$

where ϕ_d and $\sigma_{\phi,d}$ are estimated from the data. The two parts of the likelihood can then be summed to obtain the total likelihood which we use to run an MCMC in the 7-dimensional space of the calibration and thermal parameters $\{a, b, c, d, T_0, \gamma, \lambda_P\}$.

This scheme is computationally inexpensive and achieves several goals:

- (i) It generalizes the simple power-law calibration assumed in equation (5).
- (ii) It finds the parameter combinations to which the observable is sensitive, providing a way to quantitatively express its degeneracy direction.
- (iii) The uncertainty of this calibration can be determined based on the scatter of the simulated points and their uncertainties.
- (iv) It returns the constraints on the thermal parameters, marginalizing over the uncertainty on the coefficients of equation (7).

The main drawback is that the validity of equation (7) as an approximation of the observable ϕ must be verified *a posteriori*, using the estimated coefficients from the MCMC run.

In the case where the observable ϕ is not a scalar but a vector, it is possible to generalize the likelihood expressions 8 and 9 to multivariate Gaussian likelihoods. This will increase the dimensionality of the calculation because there must be four coefficients $\{a, b, c, d\}$ for each observable component. It will also be necessary to calculate the full covariance matrix of the different components of ϕ .

In this work, we have applied this formalism to the observables $(\log b_0, \Gamma)$ that describe the lower cut-off of the b - N distribution and to the differential median parameters (m_1, α) that we will define in the next section. In both cases, this involved an 11-parameter MCMC analysis and adopting a bivariate Gaussian likelihood, both for equation (8) and (9). The correlation between $\log b_0$ and Γ has been calculated by bootstrapping the sample of absorption lines, as it was done for the standard deviations.

4.3 Differential median of the line width distribution

Although the cut-off is certainly the most prominent feature of the distribution, it is reasonable to expect that the thermal properties of the IGM do not only affect the narrowest lines in the Ly α forest. Taking advantage of the additional information contained by the bulk of the line population is an obvious way of trying to refine the constraints achievable with a Voigt-profile decomposition approach. Some of our preliminary efforts to do so encountered systematic issues related to very broad lines ($b > 60 \text{ km s}^{-1}$). Our models were not able to reproduce the observed line distribution in that range, and, more problematically, inferences on the thermal parameters from various fitting schemes were strongly dependent on the way the upper b range was chosen or on the statistic employed to characterize the line distribution. We concluded that the broad lines are in most cases probably an artefact of the fitting procedure and do not represent physical properties of the IGM (see for example Fig. 1) and are therefore more subject to systematic errors, in particular, due to the uncertainty of the continuum placement.




Figure 4. Representation of the differential median statistic. The points in the four panels show the b - N distribution of the data and the simulation for the same three thermal models as in Fig. 2. The vertical dotted lines separate the two N_{HI} ranges for which the two medians m_1 and m_2 (shown in dashed red horizontal lines) are calculated. The blue lines describe the equation $b = m_1' - \alpha(\log N_{\text{HI}} - 13)$ chosen so that they horizontally divide both column-density ranges into two parts with an equal numbers of points (see text for details). m_1 and α are the two parameters used in the analysis. The plot illustrates their sensitivity to the thermal parameters γ and T_0 .

For this reason, we turned our attention to a statistical estimator which is not sensitive to the distribution of lines at extreme values of b . A natural choice is the median. In order to capture the column-density dependency of the Doppler-parameter distribution, we adopt the following approach. We define m_1 as the median of the line width distribution for $10^{12.5} < N_{\text{HI}} < 10^{13.5} \text{ cm}^{-2}$ and m_2 as the median for $10^{13.5} < N_{\text{HI}} < 10^{14.5} \text{ cm}^{-2}$. We then apply a linear transformation to $\log b$ parametrized as

$$\log b' = \log b + \alpha \log \left[\frac{N_{\text{HI}}}{10^{13} \text{ cm}^{-2}} \right], \quad (10)$$

where the transformation coefficient α is such that the medians m_1' and m_2' of the new quantity $\log b'$ are identical (for the same N_{HI} domains defined above). It is straightforward to determine α iteratively. This calculation is different from simply considering the differences of the median $m_2 - m_1$, because α depends on the positions of all individual lines in the plane and not just on the median of $\log b$ in the two parts. However, for simplicity, we will refer to this method as the ‘differential median’ method. We choose m_1 (the median calculated *before* the transformation) and α as our final observables to estimate the thermal parameters, to which we apply the analysis technique illustrated in the previous section. Analogously to the cut-off method, we calculate uncertainties and covariance for m_1 and α by bootstrapping the line sample. An illustration of our differential median statistic is shown in Fig. 4, both for the data and for the same three thermal models shown in Fig. 2. Its sensitivity to the parameters of the TDR is shown by the different slopes and normalizations of the lines for the different thermal models, which qualitatively have the same parameter dependencies as the cut-off.

5 RESULTS

In the left-hand panels of Fig. 3, we show the values of $\log b_0$ in the simulations as a function of the temperature at mean density and Γ as a function of the slope of the TDR γ . When fitting the linear relations in equation (5), we obtain $A = 2.06$, $B = 1.76$; $C = -0.07$ and $D = 3.28$ (corresponding to the black dashed lines). The cut-off-fitting algorithm applied to the data returns $\log b_0 = 1.186 \pm 0.084$ and $\Gamma = 0.180 \pm 0.062$, marked as the red shaded region in the plot. By applying the above coefficients to convert these numbers to a measurement of the TDR parameters, we get $T_0/10^3 \text{ K} = 14.3 \pm 5.0$

and $\gamma = 1.52 \pm 0.22$. The reported uncertainties take into account the propagated (small) errors of the cut-off calibration coefficients.

The generalization of the relation between the thermal parameters T_0 and γ , and the parameters characterizing the lower cut-off of the b - N distribution described in Section 4.2 is shown in the right-hand panels of Fig. 3. These plots are analogous to the left-hand panels except that the x -axes are combinations of the thermal parameters as defined in equation (7). The coefficients for equation (7) are picked from the mean values of the MCMC posterior distribution, applying the method described in Section 4.2 to a joint analysis of Γ and $\log b_0$. More precisely, the average combinations are $\psi_1 = -1.03 + 0.55 \log T_0 - 0.02(\gamma - 1) - 0.02 \log \lambda_p$ and $\psi_2 = 0.10 - 0.06 \log T_0 + 0.27(\gamma - 1) + 0.09 \log \lambda_p$. These relations do not represent a new ‘calibration’ in which the coefficients are free to vary consistently with the uncertainties on the values extracted from the simulations. However, showing that the relations between the observables and the typical combinations from the MCMC fall close to the identity relation (dashed line) is necessary to validate our approach. The constraints on the thermal parameters derived from the same MCMC analysis are shown in the green contour plots in Fig. 6. A degeneracy between the inferred values of T_0 and γ is evident, which is a consequence of taking the statistical correlation between the measured Γ and $\log b_0$ into account. The right-most panel demonstrates that the lower cut-off in the simulated b - N distribution is not sensitive to the smoothing length λ_p . When marginalized over all parameters of the MCMC analysis, including the coefficients for ψ_1 and ψ_2 , the values inferred for the TDR parameters are $T_0/10^3 \text{ K} = 15.6 \pm 4.4$ and $\gamma = 1.45 \pm 0.17$ (quoted as mean and standard deviation from the posterior distribution). The uncertainties are smaller than those inferred from the lower cut-off, and the two measurements are in good agreement.

The results obtained from the differential median technique are presented in Fig. 5. The two panels show that the quantities m_1 and α are reasonably approximated by a combination of the form of equation (7), although the scatter is not negligible. As in Fig. 3, the red bands trace the 1σ limits for the measured parameters: $m_1 = 1.446 \pm 0.020$ (left-hand panels) and $\alpha = 0.007 \pm 0.016$. The relations are well described by $\psi_1 = -0.433 + 0.256 \log T_0 + 0.038(\gamma - 1) + 0.409 \log \lambda_p$ and $\psi_2 = 0.285 - 0.141 \log T_0 - 0.194(\gamma - 1) + 0.200 \log \lambda_p$. The posterior distribution of the thermal parameters, shown as blue




Figure 5. Same as Fig. 3 (right-hand column), but for the differential median method, showing the relation between the low-column-density median m_1 and the differential coefficients α with the respective average parameter combinations obtained by the MCMC analysis, ψ_1 and ψ_2 .




Figure 6. Constraints on the thermal parameters from the two analyses presented in this work. In green we show the 1σ (dark) and 2σ (light) confidence levels in the T_0 – γ and T_0 – λ_p planes, derived from the cut-off method. The analysis takes into account the statistical correlation between the uncertainties of the two observables, $\log b_0$ and Γ . Mainly due to this correlation, there is a strong degeneracy between the estimated values of T_0 and γ , although the marginalized uncertainty of the two parameters of the TDR is comparable to the one achieved in the standard method. The right-hand panel shows that the temperature T_0 inferred from the lower cut-off of the b – N distribution is not sensitive to the smoothing scale λ_p of the numerical simulations used for calibration. The blue contours are the 1σ (dark) and 2σ (light) confidence levels for the differential median method. Similar to the analysis based on the lower cut-off of the b – N distribution, there is a strong degeneracy between the inferred T_0 and γ ; however, for this method, the temperature is also strongly degenerate with the smoothing length λ_p . The range for λ_p shown in this plot is consistent with recent measurements from Rorai et al. (2017a) at $z \in 2.2$ – 2.7 (shown as a red circle with errorbars) and $z \in 2.7$ – 3.3 (black, the vertical position of these two points is arbitrary). The vertical dashed line shows the value of λ_p measured in the non-equilibrium hydrodynamics model by Puchwein et al. (2015) at $z = 2.79$. Stronger constraints on the pressure smoothing will eventually help break the degeneracy shown in this plot and significantly improve the precision achievable with this technique.

contours in Fig. 6, reveals degeneracies between the inferred values of all three considered parameters. In particular, and in contrast to the cut-off method, there is a significant sensitivity to the pressure smoothing length λ_p . This confirms that the overall b – N distribution contains information about the spatial smoothing of the IGM gas, as already noted by Garzilli et al. (2015). The marginalized uncertainties for the TDR parameters are in this case $T_0/10^3 \text{ K} = 14.6 \pm 3.7 \text{ K}$ and $\gamma = 1.37 \pm 0.17$, in good agreement with those inferred from the lower cut-off of the b – N distribution. The fact that using the full distribution does not substantially improve the accuracy of the constraints, compared to using just the lower cut-off, is due to the degeneracy with λ_p , which was negligible for the lower cut-off method. On the other hand, the emergence of such degeneracy implies that the precision can be substantially improved by combining these results with independent measurement of the smoothing length. At the moment, the only constraints available on λ_p are

those obtained by the analysis of correlated Ly α absorption in close quasar pairs (Rorai, Hennawi & White 2013; Rorai et al. 2017a). The horizontal bars in Fig. 6 represent the result from Rorai et al. (2017a) in the redshift bins 2.2–2.7 (red) and 2.7–3.3 (black), both overlapping with the redshift range analysed in the present work. The nominal value of the λ_p measurement has been decreased by 10 per cent, in order to match the definition of pressure smoothing as given by the F_{real} cut-off (see section 3.1 and fig. S11 in Rorai et al. 2017a). As it can be seen, the whole range for λ_p considered in this paper is consistent with either of the two data points. More precise measurements of the smoothing, or an analysis conducted on the same redshift limits used in this work, are required in order to break the degeneracy and improve the accuracy on T_0 and γ achievable with the differential median technique.

An alternative way of visualizing our results is to look at the constraints in the temperature–density plane. This does not require us

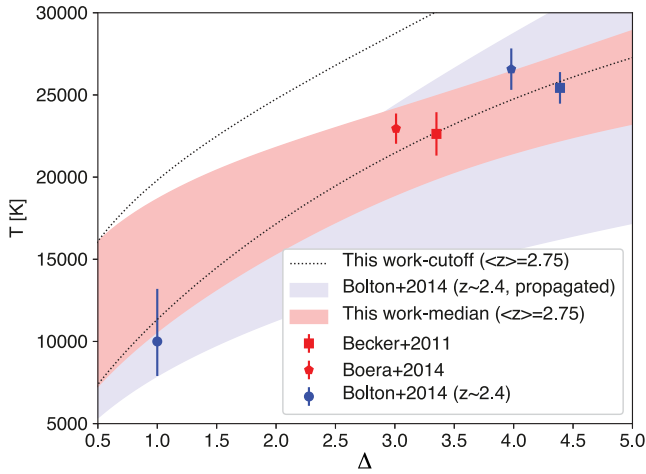


Figure 7. Constraints on the TDR in the temperature–density plane. The black dashed lines show the 16th and 84th percentile of the temperature posterior distribution as a function of density, obtained from our analysis of the lower cut-off of the b – N distribution. Analogous limits for the differential median method are shown by the red shaded area. The two measurements are in good agreement with each other. The red square and pentagon report the measurements of $T(\Delta)$ at the same redshift from BB11 and Boera et al. (2014), respectively. For comparison, we also report the results for T_0 at $z = 2.4$ from Bolton et al. (2014) (blue circle) and the propagated 1σ limits at all densities considering the measurement of γ from the same work (light blue shaded area). The blue square and pentagons are the values of $T(\Delta)$ obtained in BB11 and Boera et al. (2014) at this redshift.

to reduce the posterior distribution to two parameters with relative uncertainties. It also allows a straightforward comparison with the measurement of the temperature at overdensity $\bar{\Delta}$ obtained with the ‘curvature’ statistic in BB11 and Boera et al. (2014), without requiring error propagation by which some information might be lost. The way we do this is the following: given the posterior distribution for T_0 and γ obtained from one of our MCMCs, we calculate the marginalized distribution of $T_0\Delta^{\gamma-1}$ for the range of overdensities Δ we are interested in. We then plot the 16th and the 84th percentiles as a function of Δ as the 1σ limits in the temperature–density plane. The results for the various techniques described in this paper are shown in Fig. 7. The black dotted lines show the 16th and 84th percentiles of the temperature distributions obtained from the lower cut-off MCMC analysis, whereas the red shaded region shows the same for the differential median method. Note that the curvature results do not include uncertainties related to pressure smoothing. There is good agreement between the two measurements, and there is also good consistency with the results from BB11 and Boera et al. (2014) at the same redshift (red square and pentagon, respectively), although our data set partially overlap with the sample used in Boera et al. (2014). Note that the uncertainties of the temperature measured with the median technique are lower at mild overdensities ($\Delta \sim 2.5$ – 3.5) than at the mean density. This is a consequence of the particular degeneracy direction in the T_0 – γ plane of Fig. 5.

For reference, we also show an analogous comparison at $z = 2.4$ between the results of BB11 (blue square), Boera et al. (2014) (blue pentagon) and b – N_{H} cut-off results from Bolton et al. (2014) (blue circle). These limits assume that T_0 and γ are uncorrelated, which is likely incorrect given the results of our analysis at slightly higher redshifts. The light blue shaded area in the background represent the propagated 1σ limits on $T(\Delta)$, assuming the measured value and uncertainties of T_0 and γ from Bolton et al. (2014).

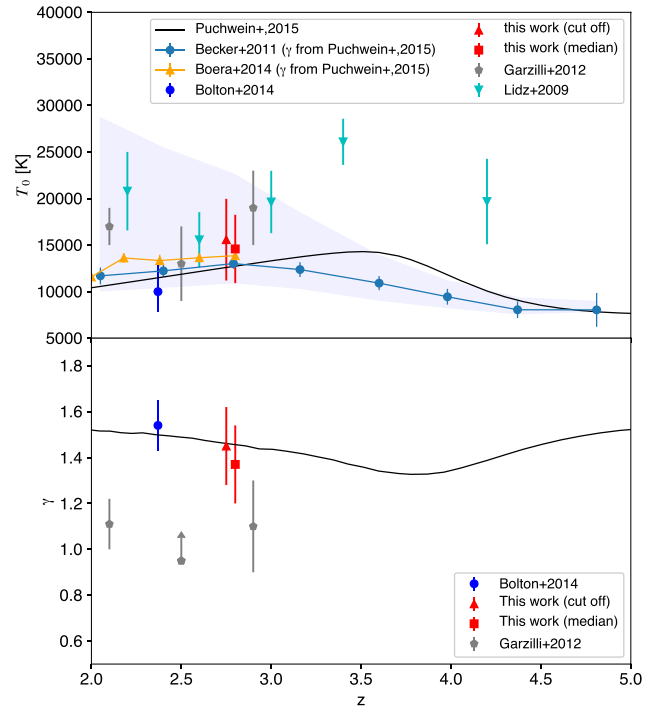


Figure 8. A summary of recent constraints on the TDR parameters from the literature, as a function of redshift. Our constraints on T_0 and γ derived from the lower cut-off of the b – N distribution (red triangles) and the differential median (red squares) methods are compared to the hydrodynamics model of Puchwein et al. (2015), black line) and the extrapolated values of T_0 from the curvature measurement of BB11 (dark blue connected circles). The latter values are obtained from the temperatures at a redshift-dependent overdensity $\bar{\Delta}(z)$ by assuming the value of γ from the Puchwein et al. (2015) model (black line in the lower panel). The light-blue shaded area represents the range of values extrapolated from BB11 assuming a flat prior on γ in the interval 1–1.6. We also plot the constraints from Bolton et al. (2014, blue circles) at somewhat lower redshift. The orange triangles are the values of T_0 extrapolated from a recalibrated version of the curvature measurement from Boera et al. (2014) and Boera (private communication), obtained in the same way as for BB11. The cyan triangles show the constraints on T_0 from Lidz et al. (2010) and the grey pentagons those from Garzilli et al. (2012), both obtained with a wavelet analysis (combined with the flux PDF in the latter). The grey pentagons in the lower panel are also from the PDF-Wavelet combined analysis of Garzilli et al. (2012). The tension between the points at $z = 2.9$ with our measurements is discussed in the text.

6 DISCUSSION

A more comprehensive comparison of the main results of this work with recent constraints on the thermal state of the IGM from the literature is presented in Fig. 8, where we show the evolution of T_0 (upper panel) and γ (lower panel) as a function of redshift. The red triangles correspond to the constraints from our fit to the lower cut-off of the b – N distribution (Section 4.2), whereas the red squares are those obtained from the differential median method (Section 4.3). The black solid lines are the predictions of a recent hydrodynamics simulation from Puchwein et al. (2015), where the UV background is assumed to follow the model by Haardt & Madau (2012). This simulation fits well both the observational points from this work and those reported by Bolton et al. (2014) (blue circles). If we assume a slope of the TDR as in the Puchwein et al. (2015) simulation, (i.e. the black line in the lower panel), we can extrapolate the curvature measurements of $T(\Delta)$ (BB11, Boera et al. 2014) to mean density.

The corresponding values of T_0 are shown as the dark blue connected circles for BB11, and the orange connected triangles for an updated version of the results of Boera et al. (2014) and Boera, private communication. The good agreement of our measurements of T_0 and the evolution of T_0 inferred from the extrapolation of the temperature measurements with the curvature method to mean density based on the γ values predicted by the Puchwein et al. (2015) simulation is a non-trivial result that suggests that our understanding of the thermal state IGM is converging from multiple different approaches. Convergence of results is also suggested by the measurements of T_0 obtained with a wavelet-analysis technique by Lidz et al. (2010) (cyan triangles) and Garzilli et al. (2012) (grey pentagons). The discrepancy with their measurements in the redshift bins close to that of our measurement is less than 1σ .

In the lower panel, the measurement of γ at $z = 2.9$ derived from a joint analysis of the flux PDF and the wavelet coefficients (Garzilli et al. 2012) is clearly in tension with our work presented here and the analysis of Bolton et al. (2014). As discussed earlier, other measurements based on the flux PDF (e.g. Viel et al. 2009; Calura et al. 2012) also have claimed that an ‘inverted’ TDR ($\gamma < 1$) is required to match the data (but see Lee 2012; Rollinde et al. 2013; Lee et al. 2015). This apparent discrepancy has been recognized for some time and was recently discussed in considerable detail by Rorai et al. (2017b). By analysing a high signal-to-noise quasar spectrum at $z = 3$, Rorai et al. (2017b) confirmed that the flux PDF constrains the TDR to be flat or to rise towards low density; however, as they point out, this is true only for gas around and below the mean density, to which the PDF statistic is most sensitive. Rorai et al. (2017b) further show that techniques for measuring the IGM temperature that rely on quantifying the smoothness of the absorption profile, such as those discussed in this paper or the curvature method developed by B11, are mainly sensitive to overdense gas, i.e. $\Delta > 1$. This has also been highlighted in Bolton et al. (2014), where they show the relation between column density and optical depth-weighted density (see fig.1 in their paper). The fact that our work presented here constrains a spatially invariant TDR to have $\gamma > 1$ corroborates the conclusion of Rorai et al. (2017b) that a single, spatially invariant power law is not able to describe the TDR of the low density IGM. This may perhaps be explained by simulations of He II reionization where radiative transfer effects have been implemented (Abel & Haehnelt 1999; Paschos et al. 2007; McQuinn et al. 2009; Compostella, Cantalupo & Porciani 2013; La Plante et al. 2017), in which the lower densities show a bimodal temperature distribution with increased temperatures and a flattening of the relation between temperature and density in regions where helium has most recently become doubly ionized.

Finally, we note that in a recent work, Garzilli et al. (2015) found that the cut-off of the line distribution is significantly sensitive to the pressure smoothing, different from what our analysis shows (see Fig. 6). This could be possibly due to the differences in the fitting algorithms, both for the individual Ly α lines and for the cut-off, but it is most likely related to the particular range of column density we have selected. Fig. 8 in their paper suggests that the effect of pressure smoothing on the low- b lines is most prominent for lines with low column densities, whereas here we have considered only those with $N_{\text{H I}} > 10^{12.5} \text{ cm}^{-2}$, for which the line width is dominated by thermal broadening, as argued also by Bolton et al. (2014).

7 CONCLUSIONS

We have analysed the Ly α forest of a sample of 13 high-resolution quasar spectra in the redshift range $2.55 < z < 2.95$ with the help of

high resolution numerical simulations of the IGM. The continuum-normalized spectra were decomposed into individual H I absorbers using vPFIT, and we have carefully identified and excluded regions potentially contaminated by metal lines. We have then used the lower cut-off in the b - $N_{\text{H I}}$ distribution and a newly introduced statistic based on the medians of the line-width distribution in separate column density ranges to obtain two new measurements of the temperature of the IGM. In both cases, we employed Bayesian MCMC techniques to constrain thermal parameters, using a grid of thermal models where the TDR has been imposed in post processing. Our results can be summarized as follows.

(i) Fitting the lower cut-off of the b - N distribution in the standard way gives $T_0/10^3 \text{ K} = 14.3 \pm 5.0 \text{ K}$ and $\gamma = 1.52 \pm 0.22$.

(ii) The parameters describing the lower cut-off and that describing the TDR are strongly correlated in simulations, but there is significant scatter due to the residual dependence on other thermal parameters (in particular, the smoothing length λ_p).

(iii) We have introduced a new calibration of the temperature measurement based on a linear combination of the thermal parameters (in logarithmic space) whose coefficients are left free to vary. For this new calibration, we have implemented an MCMC analysis in which the calibration coefficients are estimated together with the thermal quantities T_0 , γ and λ_p .

(iv) Including the correlation between the statistical uncertainties of $\log b_0$ and Γ in our likelihood, analysis gives $T_0/10^3 \text{ K} = 15.6 \pm 4.4 \text{ K}$ and $\gamma = 1.45 \pm 0.17$. Taking the correlation between T_0 and γ shows that the inferred values for the two parameters are strongly degenerate. Conversely, no sensitivity is found to the smoothing length λ_p .

(v) We have introduced an alternative statistical estimator for the thermal parameters which is based on the medians m_1 and m_2 of the line-width distribution for $N < 10^{13.5}$ and $N > 10^{13.5} \text{ cm}^{-2}$, respectively. For this, we have defined the transformation $\log b' = \log b + \alpha(\log N_{\text{H I}} - 13)$, with α chosen such that $m_1 = m_2$ for b' .

(vi) By applying the parameters m_1 and α to the same MCMC technique used for the lower cut-off method, we obtain $T_0/10^3 \text{ K} = 14.6 \pm 3.7 \text{ K}$ and $\gamma = 1.37 \pm 0.17$, in good agreement with our measurement from the lower cut-off. For the measurement based on the differential median method, we find a strong degeneracy with the smoothing length λ_p , suggesting, unsurprisingly, that the overall b - N distributions contains information about the small-scale spatial structure of the IGM.

(vii) When we use the posterior distribution for the TDR parameters to infer the temperature at $\Delta \approx 3$, we obtain values consistent with the measurements at the same redshift from BB11 and Boera et al. (2014).

(viii) Our measurements are also in good agreement with the theoretical predictions from the hydrodynamics model of Puchwein et al. (2015) and with the measurements of T_0 at similar redshifts from Lidz et al. (2010) and Garzilli et al. (2012).

(ix) Our constraints on γ are in disagreement with claims of an inverted TDR based on the flux PDF, similar to published results using statistics based on the smoothness of the absorption. Our findings further corroborate those of Rorai et al. (2017b), who argued that this is due to the different overdensity range probed by the PDF.

The work presented here marks a further step towards a consistent characterization of the thermal state of the IGM at $z \lesssim 3$. The agreement with other recent measurements at the same redshift and with high resolution hydrodynamics models is an encouraging sign that our understanding of the thermal state of the IGM is converging. The techniques developed and presented here can be easily applied to

data sets at other redshifts and should improve the well-established approach of characterizing the thermal state of the IGM with help of the $b-N$ distribution of Ly α forest absorbers. Our analysis takes into account the uncertainties in the calibration between the lower cut-off (or the median b -parameters in different column density ranges) and the thermal parameters, as well as parameter correlations and second-order dependencies which were previously neglected. The degeneracies we find, especially in our analysis based on the differential median, suggest that a joint analysis with different statistics constraining the pressure smoothing length will significantly improve the precision of current constraints on the thermal state of the IGM.

ACKNOWLEDGEMENTS

We are grateful to W.L.W. Sargent for kindly providing two of the spectra analysed in this work, Volker Springel for making GADGET-3 available and Girish Kulkarni for sharing the code for the calculation of λ_P in SPH simulations. We thank the anonymous referee for useful comments and suggestions, which greatly improved this manuscript. MH acknowledges support by ERC ADVANCED GRANT 320596 ‘The Emergence of Structure during the epoch of Reionization’. GDB was supported by the National Science Foundation through grant AST-1615814. JSB acknowledges the support of a Royal Society University Research Fellowship. MTM thanks the Australian Research Council for *Discovery Project* grant DP130100568. This work made use of the DiRAC High Performance Computing System (HPCS) and the COSMOS shared memory service at the University of Cambridge. These are operated on behalf of the STFC DiRAC HPC facility. This equipment is funded by BIS National E-infrastructure capital grant ST/J005673/1 and STFC grants ST/H008586/1, ST/K00333X/1.

REFERENCES

Abel T., Haehnelt M. G., 1999, *ApJ*, 520, L13
 Barlow T., 2008, MAKEE: MAuna Kea Echelle Extraction, available at: <http://www.astro.caltech.edu/~tb/makee>
 Becker G. D., Bolton J. S., Haehnelt M. G., Sargent W. L. W., 2011, *MNRAS*, 410, 1096 (B11)
 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
 Boera E., Murphy M. T., Becker G. D., Bolton J. S., 2016, *MNRAS*, 456, L79
 Bolton J. S., Becker G. D., 2009, *MNRAS*, 398, L26
 Bolton J. S., Viel M., Kim T., 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
 Broderick A. E., Chang P., Pfrommer C., 2012, *ApJ*, 752, 22
 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, VPFIT: Voigt profile fitting program, Astrophysics Source Code Library, record ascl:1408.015
 Chang P., Broderick A. E., Pfrommer C., 2012, *ApJ*, 752, 23
 Compostella M., Cantalupo S., Porciani C., 2013, *MNRAS*, 435, 3169
 Dekker H., D’Odorico S., Kaufer A., Delabre B., Kotzlowski H., 2000, in Iye M., Moorwood A. F., eds, Proc. SPIE Conf. Ser. Vol. 4008, Optical and IR Telescope Instrumentation and Detectors. SPIE, Bellingham, p. 534
 Eisenstein D. J., Hu W., 1999, *ApJ*, 511, 5
 Garzilli A., Bolton J. S., Kim T.-S., Leach S., Viel M., 2012, *MNRAS*, 424, 1723

Garzilli A., Theuns T., Schaye J., 2015, *MNRAS*, 450, 1465
 Gnedin N. Y., Hui L., 1998, *MNRAS*, 296, 44
 Haardt F., Madau P., 2001, in Neumann D. M., Tran J. T. V., eds, Clusters of Galaxies and the High Redshift Universe Observed in X-rays. CEA, Saclay, p. 64
 Haardt F., Madau P., 2012, *ApJ*, 746, 125
 Haehnelt M. G., Steinmetz M., 1998, *MNRAS*, 298, L21
 Hui L., Gnedin N. Y., 1997, *MNRAS*, 292, 27
 Hui L., Haiman Z., 2003, *ApJ*, 596, 9
 Jarosik N. et al., 2011, *ApJS*, 192, 14
 Kim T.-S., Carswell R. F., Mongardi C., Partl A. M., Mücke J. P., Barai P., Cristiani S., 2016, *MNRAS*, 457, 2005
 Kulkarni G., Hennawi J. F., Oñorbe J., Rorai A., Springel V., 2015, *ApJ*, 812, 30
 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
 Lee K.-G. et al., 2015, *ApJ*, 799, 196
 Lidz A., Faucher-Giguère C.-A., Dall’Aglio A., McQuinn M., Fechner C., Zaldarriaga M., Hernquist L., Dutta S., 2010, *ApJ*, 718, 199
 McDonald P., Miralda-Escudé J., Rauch M., Sargent W. L. W., Barlow T. A., Cen R., 2001, *ApJ*, 562, 52
 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
 Murphy M., 2016, UVES_popler: UVES_popler: POst-PipeLine Echelle Reduction software, available at: <https://doi.org/10.5281/zenodo.56158>
 O’Meara J. M. et al., 2015, *AJ*, 150, 111
 Olive K. A., Skillman E. D., 2004, *ApJ*, 617, 29
 Outram P. J., Chaffee F. H., Carswell R. F., 1999, *MNRAS*, 310, 289
 Paschos P., Norman M. L., Bordner J. O., Harkness R., 2007, preprint (arXiv:0711.1904)
 Pfrommer C., Chang P., Broderick A. E., 2012, *ApJ*, 752, 24
 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
 Richardt C. L. et al., 2009, *ApJ*, 694, 1200
 Ricotti M., Gnedin N. Y., Shull J. M., 2000, *ApJ*, 534, 41
 Rollinde E., Theuns T., Schaye J., Pâris I., Petitjean P., 2013, *MNRAS*, 428, 540
 Rorai A., Hennawi J. F., White M., 2013, *ApJ*, 775, 81
 Rorai A. et al., 2017a, *Science*, 356, 418
 Rorai A. et al., 2017b, *MNRAS*, 466, 2690
 Rudie G. C., Steidel C. C., Pettini M., 2012, *ApJ*, 757, L30
 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
 Springel V., 2005, *MNRAS*, 364, 1105
 Theuns T., Schaye J., Haehnelt M. G., 2000, *MNRAS*, 315, 600
 Theuns T., Mo H. J., Schaye J., 2001, *MNRAS*, 321, 450
 Theuns T., Schaye J., Zaroubi S., Kim T., Tzanavaris P., Carswell B., 2002, *ApJ*, 567, L103
 Valageas P., Schaeffer R., Silk J., 2002, *A&A*, 388, 741
 Viel M., Bolton J. S., Haehnelt M. G., 2009, *MNRAS*, 399, L39
 Vogt S. S. et al., 1994, in Crawford D. L., Craine E. R., eds, Proc. SPIE Conf. Ser. Vol. 2198, Instrumentation in Astronomy VIII. SPIE, Bellingham, p. 362
 Zaldarriaga M., Hui L., Tegmark M., 2001, *ApJ*, 557, 519

APPENDIX: EFFECT OF TEMPERATURE AND OPTICAL DEPTH VARIATION

To understand the effect of an evolving optical depth and a possible variation in temperature, we conducted the following test:




Figure A1. Effect of mixed temperatures or effective optical depths on the cut off and on the differential median. Left-hand panel: a combination of three models with $T_0 = 10\,000$ K (blue diamonds and lines), $T_0 = 15\,000$ K (black crosses and lines) and $T_0 = 20\,000$ K (red squares and lines). The differential median is represented as in Fig. 4, by dashed lines for the individual models and by a thick solid black line for the combination of all of them. The same is done for the lower cut-off. Right-hand panel: same as the left-hand panel, but for a combination of four models with different effective optical depth $\tau = 0.33$ (blue) $\tau = 0.30$ (black) $\tau = 0.29$ (red) $\tau = 0.27$ (magenta).

(i) We created a model which is a mixture of four models with different effective optical depth, $\tau_{\text{eff}} = -\log \bar{F} = 0.33, 0.3, 0.29, 0.27$. This values encompass the evolution of the mean transmission across our redshift bin.

(ii) We created another model which is a mixture of three models with different temperatures at mean density $T_0 = 10\,000, 15\,000, 20\,000$ K. This mixture could be interpreted either as a (rather extreme) temperature evolution with redshift or as spatial fluctuations.

We then calculated the cut-off and the differential median statistic for the two mixed models and for their individual components. The results are shown in Fig. A1.

In the case of the optical depth mixture (right-hand panel), the differences between the total model (black solid lines) and the four components (coloured dashed lines) are minimal, especially for the differential median statistic. One of the four cut-offs slightly deviates from the others at low column densities, but given that there is no clear trend and given the magnitude of the error (the

outlier model has an uncertainty of $\Delta \log b_0 = 0.065$, four times larger than the other models), we consider it a statistical fluctuation due to the noise/bootstrap realizations.

In the case of the temperature mixture (left-hand panel), it appears evident that the differential median of the total model (black solid) depends on the average temperature, i.e. coincides with the middle of the three components, whereas the cut-off picks the coldest of the three. This suggests that if there is significant temperature evolution or fluctuation, measurements based on the cut-off will be biased towards the lower values, differently from the median statistic. However, our measurements based on the two methods are in good agreement, and if anything, the cut-off analysis points towards slightly higher temperatures, which leads us to conclude that this effect is negligible at this level of precision.

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