

Constraining the temperature–density relation of the intergalactic medium with the Lyman α and β forests

Elisa Boera,¹★ Michael T. Murphy,¹ George D. Becker^{2,3} and James S. Bolton⁴

¹Centre for Astrophysics and Supercomputing, Swinburne University of Technology, Hawthorn, Victoria 3122, Australia

²Space Telescope Science Institute, 3700 San Martin Dr, Baltimore, MD 21218, USA

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

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

Accepted 2015 October 28. Received 2015 October 27; in original form 2015 June 20

ABSTRACT

The post-reionization thermal state of the intergalactic medium is characterized by a power-law relationship between temperature and density, with a slope determined by the parameter γ . We describe a new method to measure γ using the ratio of flux curvature in the Lyman α and β forests. At a given redshift, this curvature ratio incorporates information from the different gas densities traced by Lyman α and β absorption. It is relatively simple and fast to compute and appears robust against several observational uncertainties. We apply this technique to a sample of 27 high-resolution quasar spectra from the Very Large Telescope. While promising statistical errors on γ appear to be achievable with these spectra, to reach its full potential, the dependence of the curvature ratio on the thermal state of the gas in the foreground Lyman α forest will require further, detailed forward modelling.

Key words: intergalactic medium – quasars: absorption lines – cosmology: observations.

1 INTRODUCTION

The intergalactic medium (IGM) at low densities, with overdensities $\Delta = \rho/\bar{\rho} \lesssim 10$ at $z \sim 2$ –4, where $\bar{\rho}$ is the gas mean density, has a long adiabatic cooling time that allows it to maintain a record of important events that affected its thermal history, such as H I and He II reionization (Hui & Gnedin 1997). The energy injected on relatively short time-scales during these epochs will increase the IGM temperature, but it should also change its temperature–density (T – ρ) relation (e.g. Ricotti, Gnedin & Shull 2000; Schaye et al. 2000). In the simplest scenario, the interplay between cooling and photoionization heating by the UV background (UVB) results in a well-defined T – ρ relation, $T(\Delta) = T_0 \Delta^\gamma$, with a power-law slope, $\gamma - 1$, and a temperature T_0 at the cosmic mean density (Hui & Gnedin 1997; McQuinn & Upton Sanderbeck 2015).

However, during and immediately after the H I and He II reionizations (at $z = 6$ –11 and 3–4, respectively, e.g. Syphers & Shull 2014; Worseck et al. 2014; McGreer, Mesinger & D’Odorico 2015) cosmological simulations predict that γ may vary, becoming multi-valued and spatially-dependent (e.g. Bolton, Meiksin & White 2004; McQuinn et al. 2009; Meiksin & Tittley 2012; Compostella, Cantalupo & Porciani 2013; Puchwein et al. 2015). Despite considerable recent improvements, accurately simulating the effect of reionization events on the IGM remain an open challenge. Precise observational constraints are therefore desirable.

The main laboratory to detect variations in the T – ρ relation has been the H I Lyman α (Ly α) forest in quasar spectra. Efforts to infer the thermal state of the IGM have used either line-profile decomposition to measure gas temperature as a function of column density (e.g. Ricotti et al. 2000; Schaye et al. 2000; McDonald et al. 2001; Rudie et al. 2013; Bolton et al. 2014) or a variety of statistical approaches which are valuable at higher redshifts, $z > 3$, where line fitting is problematic (e.g. Theuns et al. 2002; Becker, Rauch & Sargent 2007; Bolton et al. 2008; Lidz et al. 2010; Becker et al. 2011; Boera et al. 2014; Lee et al. 2015). While these methods probe wide redshift and density ranges ($z \approx 1.6$ –5, $\Delta \approx 0.3$ –8), large uncertainties, particularly in the measurements of γ , remain.

One way to improve this situation is to constrain the T – ρ relation by comparing Ly α and higher order Lyman series transitions, such as Ly β . Ly β lines of moderate optical depth ($\tau \sim 0.1$ –1.0) arise from higher overdensities where Ly α lines may be saturated; the Ly α -to- β optical depth ratio is $f_{\alpha\lambda_\alpha}/f_{\beta\lambda_\beta} = 6.24$ (proportional to the ratios of oscillator strengths and rest wavelengths). Indeed, using the Ly β forest in IGM temperature measurements has been suggested in several theoretical works (e.g. Dijkstra, Lidz & Hui 2004; Furlanetto & Oh 2009; Iršič & Viel 2014). However, so far no practical attempt has been made to directly measure γ from a joint Ly α and β forest analysis. One challenge is that the Ly β forest is entangled with lower-redshift, foreground Ly α absorption. Hereafter, we refer to the region between the Ly β and Ly γ emission lines, where the Ly β and foreground Ly α absorption occurs, as the Ly $\beta + \alpha$ region. However, assuming that the Ly β and foreground Ly α lines arise from physically uncorrelated IGM structures, a possible strategy to overcome this problem is

★E-mail: eboera@swin.edu.au

to statistically compare the properties of the Ly $\beta + \alpha$ and Ly α regions.

In this Letter, we present a new method to constrain the slope of the T - ρ relation using the two forest regions (Ly α and Ly $\beta + \alpha$) in 27 high-resolution quasar spectra. We use a statistic based on the flux curvature analysis of Becker et al. (2011) and Boera et al. (2014). These previous works demonstrated that the curvature method can measure the temperatures at the (redshift dependent) characteristic densities probed by the Ly α forest. However, as only a narrow density range is constrained, it has not yet been used to measure the slope of the T - ρ relation from data (but see Padmanabhan, Srianand & Choudhury 2015, for a recent theoretical analysis using the Ly α curvature distribution). Using hydrodynamical simulations, we show that the ratio between the curvatures of corresponding Ly α and Ly $\beta + \alpha$ forest regions (i.e. around the same redshift, where in the Ly $\beta + \alpha$ region the redshift refers to Ly β only) is sensitive to differences in the IGM thermal state between the two density regimes. Averaged over many lines of sight, this curvature ratio allows γ to be measured with little sensitivity to T_0 .

We demonstrate the potential for this technique using 27 quasar spectra spanning the redshift range $2 \lesssim z \lesssim 3.8$. That potential is currently limited by assumptions regarding the thermal state of the foreground Ly α forest; we propose how these can be overcome with future simulations and a refined data analysis approach.

2 THE OBSERVATIONAL DATA

The 27 quasar spectra were originally retrieved from the archive of the Ultraviolet and Visual Echelle Spectrograph (UVES) on the Very Large Telescope (VLT). They were selected on the basis of quasar redshift, wavelength coverage and signal-to-noise ratio from the sample of 60 spectra used in Boera et al. (2014) (hereafter B14). They have resolving power $R \sim 50\,000$ and continuum-to-noise $\geq 24 \text{ pix}^{-1}$ in the Ly α forest region. This level of spectral quality is necessary so that the curvature measurement is not dominated by noise and unidentified metal lines. Because our new method compares the curvature of the Ly α and $\beta + \alpha$ forest regions, we extended these same criteria to the Ly $\beta + \alpha$ region at $z \sim 2.0$ – 3.8 , reducing the sample from 60 to 27 spectra. The quasar sample details are provided in table 1 in the Supporting Information (hereafter SI).

To establish an initial continuum in the Ly $\beta + \alpha$ region, we applied the same automatic, piece-wise polynomial fitting algorithm in B14, with the same parameters, leaving the Ly α region unchanged. While manual refitting was necessary in some spectra for particular parts of the Ly $\beta + \alpha$ region, this initial continuum has little effect on the final curvature measurements because the spectra are subsequently renormalized based on a b-spline fit to the flux profile (see Section 4.1).

3 THE SIMULATIONS

We used the same hydrodynamical simulations in B14 to produce synthetic spectra for Ly $\beta + \alpha$ and Ly α in the redshift range $z = 2.0$ – 3.5 . The GADGET-3 smoothed-particle hydrodynamic simulations include dark matter and gas, with 2×512^3 particles and a gas particle mass of $9.2 \times 10^4 M_\odot$ in a periodic box of 10 comoving h^{-1} Mpc (see B14 for details). The gas, assumed to be optically thin, is in equilibrium with a spatially uniform UVB (Haardt & Madau 2001), but the photoheating rates have been rescaled so that the corresponding values of T_0 and γ vary between different simulations. The parameters characterizing the different models are summarized in table 2 in the SI. From each model, synthetic Ly α forest spec-

tra were generated for 1024 random lines of sight through each of six redshift snapshots over the range $z = 2.0$ – 3.5 . The Ly $\beta + \alpha$ spectra were produced by scaling the Ly α optical depths by a factor $f_\alpha \lambda_\alpha / f_\beta \lambda_\beta = 6.24$ and contaminating them with randomly selected foreground Ly α absorption from lower redshift outputs ($z = 1.6$ – 2.7) from the same simulation.

Finally, as in B14, we calibrated the synthetic spectra to match the properties of the real spectra (i.e. resolving power, pixel size and signal-to-noise ratio). The only difference is that, in this Letter, we scaled the effective optical depth of the synthetic spectra to match the recent results from Becker et al. (2013) rather than to a direct measurement from our spectra. As shown in Section 4.1, our measurements are relatively insensitive to this optical depth calibration.

Box size and mass resolution convergence tests are available in fig. 1 and table 3 of the SI.

4 THE CURVATURE RATIO METHOD

The curvature for a region of spectrum is defined as $\kappa \equiv F''(1 + F'^2)^{-3/2}$ (Becker et al. 2011), with the first and second derivatives of the normalized flux (i.e. transmission; F' , F'') taken with respect to wavelength or relative velocity. As demonstrated in previous works, the Ly α forest curvature is directly related to the IGM temperature at the characteristic overdensities probed by this absorption ($\Delta_\alpha \simeq 5$ – 2 for $z = 2$ – 3.5), regardless of γ . Because the median overdensity contributing to Ly β forest absorption is higher than that for Ly α (Furlanetto & Oh 2009), the mean absolute curvature computed from sections of Ly $\beta + \alpha$ forest will be, on average, a tracer of the IGM temperature in a higher density regime ($\Delta_\beta \simeq 8$ – 5 for $z = 2$ – 3.5). Therefore, the curvature ratio, at each redshift z ,

$$R_\kappa(z) \equiv \frac{\langle |\kappa_{\beta+\alpha}(z)| \rangle}{\langle |\kappa_\alpha(z)| \rangle}, \quad (1)$$

probes temperatures at two different gas densities and, consequently, is sensitive to γ . Here, the mean absolute curvatures for Ly $\beta + \alpha$ and Ly α are averaged over spectral sections of 10 comoving h^{-1} Mpc centered on the same redshift (corresponding to the simulation box size and, again, z in the Ly $\beta + \alpha$ region refers to Ly β absorption). Due to the presence of foreground Ly α absorption, $R_\kappa(z)$ also depends on the evolution of γ and T_0 with z (see Section 4.3).

4.1 Analysis of simulated spectra: the γ - $\log \langle R_\kappa \rangle$ relation

We find the connection between the curvature ratio, R_κ , and γ at the redshifts of the simulation snapshots ($z = [2.173, 2.355, 2.553, 2.760, 3.211, 3.457]$) by computing the mean R_κ over the 1024 simulated lines of sight for each thermal history, and fitting a simple function between $\log \langle R_\kappa \rangle$ and γ .

Single R_κ values for each line of sight and redshift were obtained as defined in equation (1). To compute the mean absolute curvature of synthetic Ly α and Ly $\beta + \alpha$ sections, we adopt exactly the same algorithm and parameters described in Becker et al. (2011) and B14, an example of which is shown in Fig. 1 in each section of artificial spectrum, the mean absolute curvature is computed from a cubic b-spline fit which has been re-normalized by its maximum value in that interval. This approach is required to avoid systematic errors when determining the curvature from real spectra, so it must be applied to the synthetic spectra for consistency. The b-spline fit reduces the sensitivity of the curvature to noise and the re-normalization minimizes potential uncertainties arising from inconsistent continuum placement. Finally, only the pixels where




Figure 1. Curvature for simulated and real spectra. The top two panels display the simulated Ly α (blue solid line) and Ly $\beta + \alpha$ forests (green solid line) at $z = 2.76$ and the curvature from b-spline fits. The spectra are obtained by contaminating the corresponding Ly β forest (red dotted line) with a randomly chosen Ly α section at lower redshift (black dashed line). The lower two panels are as above but for a real Ly α and Ly $\beta + \alpha$ spectrum. The spectra (black lines) are plotted behind the b-spline fits. Shading shows a Ly $\beta + \alpha$ region contaminated by metal absorption (green dashed line); the corresponding part of the Ly α spectrum is also masked (black dashed line).

the re-normalized b-spline fit falls in the range 0.1–0.9 are used to measure the mean absolute curvature of each section. In this way we maximize the sensitivity to the signal and avoid possible systematic uncertainties: we exclude saturated pixels, which do not contain useful information, and pixels with little-to-no absorption whose curvature is near zero and uncertain.

In Fig. 2, we present the γ – $\log \langle R_\kappa \rangle$ relationship obtained by averaging the curvature ratio computed from the synthetic spectra with different thermal histories, i.e. different γ and T_0 values. At each redshift, models characterized by the same γ value, but with T_0 ranging between 5000–31 000 K, yield $\log \langle R_\kappa \rangle$ values that vary by only $\lesssim 0.01$ dex. Therefore, for each redshift, it is possible to fit a simple powerlaw ($a\gamma^b + c$) that connects the mean $\log \langle R_\kappa \rangle$ and γ independently of T_0 , at least in these simulations (see also table 4 and fig. 2 in the SI). As redshift decreases, the relation between $\log \langle R_\kappa \rangle$ and γ becomes steeper, as shown in Fig. 2. Given this tight correspondence, the curvature ratio represents an interesting tool to independently measure the slope of the IGM T – ρ relation.

The sensitivity of this nominal γ – $\log \langle R_\kappa \rangle$ relation to observational uncertainties in the spectra was tested as follows.

(i) Noise: the synthetic spectra from which we obtain the nominal relationship need to include noise at the same level as in the real spectra. If, as an extreme example, no noise was added, Fig. 2 shows the effect on the γ – $\log \langle R_\kappa \rangle$ relation: an ~ 5 per cent decrease in $\langle R_\kappa \rangle$ (~ 0.02 dex in $\log \langle R_\kappa \rangle$). This would cause an ~ 8 per cent underestimation of γ , comparable to the statistical errors in our observational sample (see Section 4.2). Therefore, any errors in how the noise properties of a sample like ours are incorporated into




Figure 2. Relationship between γ and the mean curvature ratio, $\log \langle R_\kappa \rangle$, from our nominal simulations at three different redshifts. For clarity, we show the values of $\log \langle R_\kappa \rangle$ from the different models and the relation without noise for $z = 2.760$ only (coloured points and dashed line, respectively). See fig. 2 and table 4 in the SI for the plots corresponding to the other redshifts and the fitting parameters.

the γ – $\log \langle R_\kappa \rangle$ relation should cause relatively small systematic uncertainties in γ measurements.

(ii) Effective optical depth: changing the effective optical depth by 10 per cent alters $\log \langle R_\kappa \rangle$ by 0.01 dex and, consequently, γ by ~ 4 per cent, again well within the statistical uncertainties in our observational sample. This point is particularly promising because the curvature for the Ly α forest alone (and the Ly $\beta + \alpha$ forest alone) is considerably more sensitive to this quantity, as explored in B14.

However, while the nominal γ – $\log \langle R_\kappa \rangle$ relation is robust to these observational aspects, systematic uncertainties in the evolution of the IGM thermal state must also be considered (see Section 4.3).

4.2 The observed curvature ratio

To apply the method to the 27 real quasar spectra we compute the curvature ratio in sections of 10 comoving h^{-1} Mpc of metal-free Ly α and corresponding Ly $\beta + \alpha$ forest regions. Narrow metal lines ($b \lesssim 15$ km s $^{-1}$) represent a potentially serious source of systematic errors in any measure of forest absorption and should be avoided. With this aim we ‘clean’ the spectra, using the exact procedure detailed in B14, extended to the Ly $\beta + \alpha$ region: metal absorbers redward of the Ly α emission line are identified and all strong metal transitions at their redshifts are masked out, followed by a by-eye check of the remaining forest.

While the metal correction produces spectra that are reasonably free of contaminants, this procedure reduces the quantity of information available in different sections in a non-uniform way, introducing a possible source of bias. Because the curvature ratio traces differences in the absorption features of two different regions of the same observed spectrum, avoiding systematic effects requires that they cover the same absorption redshift range. Therefore, before measuring R_κ from the real data, we mask out regions of the Ly $\beta + \alpha$ forest corresponding to any range masked from the Ly α forest, and vice versa. Fig. 1 shows an example of this masking procedure. Finally, possible edge effects are avoided: we do not include the 4 pixels closest to the edge of any masked region in the curvature ratio calculation.

Fig. 3 presents the curvature ratio results from our observational sample. The statistical uncertainty in an individual $\log R_\kappa$ measurement, computed from a single pair of Ly $\beta + \alpha$ and Ly α




Figure 3. Curvature ratio, R_κ from each pair of Ly $\beta + \alpha$ and Ly α sections from 27 UVES/VLT quasar spectra (black circles). The measurement of $\log \langle R_\kappa \rangle$, averaged within each redshift bin, is shown with 1σ bootstrap errors both before (red squares) and after (green points) the metal masking procedure. The expected evolution of $\log \langle R_\kappa \rangle$, from our nominal simulations with relatively constant γ , is presented for $\gamma \sim 1.5$ (solid blue line) and $\gamma \sim 1.0$ (dashed blue line). The redshift bins are $(z_{\min}, \bar{z}, z_{\max}) = (2.0, 2.27, 2.5), (2.5, 2.86, 3.1), (3.1, 3.36, 3.74)$, and the (metal-masked) $\log \langle R_\kappa \rangle$ measurements are $-0.110 \pm 0.026, -0.019 \pm 0.030$ and 0.017 ± 0.032 , respectively.

spectral sections, is negligible compared to the much larger variance among different measurements. Therefore, in each of three broad redshift bins, which include >30 individual measurements, we calculate the mean $\log \langle R_\kappa \rangle$ and its uncertainty using a bootstrap technique using the measurements in the bin. The width and roughly Gaussian shape of the $\log R_\kappa$ distribution within each bin is also reproduced by our simulated spectra, providing some confidence that the simulations adequately describe the statistical properties of the observed forest absorption (see fig. 3 in the SI).

There is evidence for a mild evolution in $\log \langle R_\kappa \rangle$ as a function of redshift. A Spearman rank correlation test reveals a positive correlation ($r \approx 0.26$) with an associated probability of ≈ 0.001 . For a constant γ , this increase in $\log \langle R_\kappa \rangle$ is consistent with the expected increase in $\log \langle R_\kappa \rangle$ at increasing redshifts seen in Fig. 2. In particular, our $\log \langle R_\kappa \rangle$ measurements in Fig. 3 show good agreement with the expected evolution for a model with $\gamma \sim 1.5$, i.e. a scenario consistent with the recent result of $\gamma = 1.54 \pm 0.11$ at $z \sim 2.4$ from Bolton et al. (2014). The absence of a strong change in the curvature ratio in the redshift range considered is consistent with the assumption in the nominal simulations that γ remains relatively constant over redshifts $z = 2-3.5$ (but see discussion on systematic uncertainties below).

Fig. 3 also shows $\log \langle R_\kappa \rangle$ computed without first masking the metal absorption lines. Even though the effect of metal contamination is important when measuring the curvature of the Ly α forest alone (B14), it is similar in the corresponding sections of Ly $\beta + \alpha$ forest, so the curvature ratio is less sensitive to this correction. The results in Fig. 3 show that, even *without* applying the metal correction, the bias introduced in $\log \langle R_\kappa \rangle$ is ~ 0.05 dex. Therefore, possible errors in the metal masking procedure will introduce small uncertainties in γ compared to our statistical ones.

4.3 Modelling uncertainties

Our nominal simulations assume only mild evolution in T_0 ($\Delta T_0 \sim 2000$ K) and γ ($\Delta\gamma \lesssim 0.02$) between the Ly β and the foreground

Ly α redshifts (see fig. 1 in Becker et al. 2011). However, if these parameters vary more drastically on short time-scales due to e.g. blazar heating (Puchwein et al. 2012) or non-equilibrium photoionization effects during He II reionization (Puchwein et al. 2015), these assumptions may have important consequences for the γ - $\log \langle R_\kappa \rangle$ relation and, therefore, preclude a final conversion of our $\log \langle R_\kappa \rangle$ measurements to γ values with formal error bars.

We can construct a toy model to investigate how sensitive the γ - $\log \langle R_\kappa \rangle$ relation is to rapid evolution in T_0 and γ by altering these parameters in the foreground Ly α forest via a simple post-processing of the simulated spectra. For a given simulation model, at a given redshift, new synthetic spectra are extracted after imposing a new one-to-one power-law T - ρ relationship. Note, this is an approximation, as it removes the natural dispersion in the temperatures at a given overdensity from shock-heating/radiative cooling. However, in this way we can easily modify the T_0 and γ parameters, and their evolution, to explore the effect on the γ - $\log \langle R_\kappa \rangle$ without running new hydrodynamical simulations. The largest effects on the γ - $\log \langle R_\kappa \rangle$ relation were found in the following two tests.

(i) Rapid evolution in T_0 : for the redshift range $z = 2-3.5$, T_0 evolves in our nominal simulations such that the temperature difference between the foreground and the Ly β redshift is small, $\Delta T_0 \equiv T_0(\text{Ly } \beta) - T_0(\text{foreground}) \lesssim 2000$ K. To test the effect of much stronger variations in the temperature at the mean density, we modified the values of T_0 in each of our nominal simulations, for the foreground Ly α forest only, using the post-processing approach. We then computed the change in $\log \langle R_\kappa \rangle$ compared to the nominal values, $\Delta \log \langle R_\kappa \rangle$. Fig. 4 shows the direct relationship between ΔT_0 and $\Delta \log \langle R_\kappa \rangle$. For example, if T_0 changes by a further ≈ 5000 K between $z = 3.2$ (Ly β) and $z = 2.6$ (foreground), we would expect a systematic error in our measurement of $\log \langle R_\kappa \rangle$ of ≈ 0.03 , similar to the statistical error per redshift bin derived from our 27 quasar spectra.

(ii) Rapid evolution in γ : similar to the ΔT_0 test above, we emulated rapid changes in γ over short time-scales, $\Delta z \approx 0.6$, by post-processing the foreground Ly α forest only. Again, we find a strong correlation between the change in foreground γ and $\log \langle R_\kappa \rangle$; for example, a decrease in the foreground γ by 0.15 implies $\Delta \log \langle R_\kappa \rangle \approx 0.03$, again equivalent to the statistical uncertainties in our three redshift bins.

On the other hand, we found that the γ - $\log \langle R_\kappa \rangle$ relation was insensitive to variations in the integrated thermal history in the

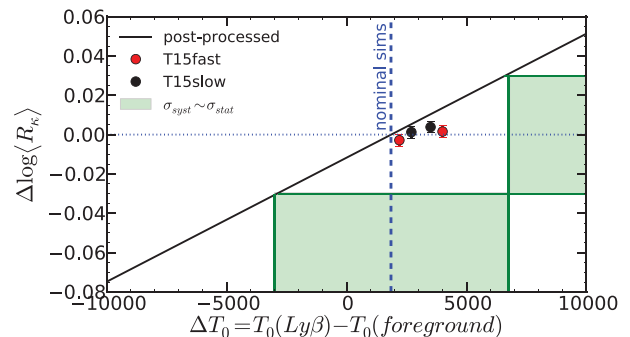


Figure 4. Expected relation between the variation in $\log \langle R_\kappa \rangle$ and the temperature change, ΔT_0 between the Ly β redshift and foreground Ly α forest (black curve, see text for details). The green shading indicates the typical statistical uncertainty in our measured $\log \langle R_\kappa \rangle$ values. The ΔT_0 in our nominal simulations is typically $\lesssim 2000$ K (dashed line) and the specific results for the self-consistent ‘T15fast’ and ‘T15slow’ simulations with their statistical errors are shown for $z = 2.7$ and 3.2 .

simulations (i.e. Jeans smoothing effects) when the instantaneous T_0 and γ are held fixed (see fig. 4 in the SI).

While highlighting the potential importance of assumptions for the foreground Ly α forest, the above tests rely on a rather simplistic toy model that does not reproduce self-consistently the evolution of the complex relationships between physical parameters. Ideally, the sensitivity of the γ - $\log \langle R_\kappa \rangle$ relation to different physical assumptions and thermal histories would need to be tested with additional self-consistent simulations in the redshift range of interest. Our simulation suite does offer one such self-consistent test in the case of strong T_0 evolution: we used the ‘T15fast’ and ‘T15slow’ simulations of Becker et al. (2011) (see their fig. 8) to mimic a possible ~ 5000 K heating event from He II reionization at $z > 3$ such that between the Ly β and foreground redshifts there was a typical decrease $\Delta T_0 \approx 2000$ – 4000 K. In both simulations we find that $\log \langle R_\kappa \rangle$, varies by $\lesssim 0.01$ compared to the nominal simulations (see Fig. 4), which seems less sensitive to substantial evolution in T_0 than implied by our simplistic toy model. However, differences in the pressure smoothing scale in these models (which may be acting to improve the agreement) prevent us from reliably estimating the systematic uncertainties involved without further self-consistent tests. Nevertheless, these results suggest that the curvature ratio is a promising alternative tool to measure the density dependence of the IGM thermal state.

5 CONCLUSIONS

We have presented a new approach to constraining the slope of the T - ρ relation using the ratio of curvatures, $\langle R_\kappa \rangle$, in the Ly α and β forests. The statistic appears robust against observational uncertainties in the noise level, metal contamination and effective optical depth, and is relatively simple and fast to compute. We measure $\langle R_\kappa \rangle$ in 27 VLT/UVES quasar spectra, achieving ≈ 6 per cent statistical accuracy in three redshift bins covering $z = 2.0$ – 3.5 . In the absence of any other systematics, this translates to a $\lesssim 10$ per cent uncertainty in γ , suggesting that $\langle R_\kappa \rangle$ should be a useful tool to constrain the slope of the T - ρ relation, possibly competitive with recent attempts to measure γ using line decomposition (Rudie et al. 2013; Bolton et al. 2014).

The primary goal of this Letter is to introduce the $\langle R_\kappa \rangle$ statistic. However, we emphasize that the thermal state of the foreground Ly α forest may complicate a direct translation of $\langle R_\kappa \rangle$ into a constraint on γ . The simulations used in this Letter have simple thermal histories where both T_0 and γ remain relatively constant over $2 \lesssim z \lesssim 5$. The impact on $\langle R_\kappa \rangle$ for some deviations from these histories are explored in Section 4.3; however, we have not attempted to capture the full range of possible T_0 and γ evolutions expected for non-equilibrium photoheating during He II reionization (e.g. McQuinn et al. 2009; Haardt & Madau 2012; Puchwein et al. 2015) or more exotic models incorporating e.g. blazar heating (Puchwein et al. 2012). We intend to more fully explore these issues in a future work. If the major systematic uncertainties can be marginalized over, or if $\langle R_\kappa \rangle$ can be combined with other statistics to simultaneously determine $T_0(z)$ and $\gamma(z)$, then the curvature ratio promises to be an effective tool for constraining the thermal history of the IGM over a large redshift range.

ACKNOWLEDGEMENTS

The simulations were performed using the Darwin Supercomputer of the University of Cambridge High Performance Computing Service, provided by Dell Inc. using Strategic Research Infrastructure

Funding from the Higher Education Funding Council for England. MTM thanks the Australian Research Council for *Discovery Project* grant DP130100568, GDB acknowledges support from the Kavli Foundation and JSB acknowledges the support of a Royal Society University Research Fellowship.

REFERENCES

- Becker G. D., Rauch M., Sargent W. L. W., 2007, *ApJ*, 662, 72
 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 (B14)
 Bolton J., Meiksin A., White M., 2004, *MNRAS*, 348, L43
 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
 Compostella M., Cantalupo S., Porciani C., 2013, *MNRAS*, 435, 3169
 Dijkstra M., Lidz A., Hui L., 2004, *ApJ*, 605, 7
 Furlanetto S. R., Oh S. P., 2009, *ApJ*, 701, 94
 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
 Haardt F., Madau P., 2012, *ApJ*, 746, 125
 Hui L., Gnedin N. Y., 1997, *MNRAS*, 292, 27
 Iršič V., Viel M., 2014, *J. Cosmol. Astropart. Phys.*, 12, 24
 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
 McGreer I. D., Mesinger A., D’Odorico V., 2015, *MNRAS*, 447, 499
 McQuinn M., Upton Sanderbeck P., 2015, preprint ([arXiv:1505.07875](https://arxiv.org/abs/1505.07875))
 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
 Padmanabhan H., Srianand R., Choudhury T. R., 2015, *MNRAS*, 450, L29
 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
 Ricotti M., Gnedin N. Y., Shull J. M., 2000, *ApJ*, 534, 41
 Rudie G. C., Steidel C. C., Shapley A. E., Pettini M., 2013, *ApJ*, 769, 146
 Schaye J., Theuns T., Rauch M., Efstathiou G., Sargent W. L. W., 2000, *MNRAS*, 318, 817
 Syphers D., Shull J. M., 2014, *ApJ*, 784, 42
 Theuns T., Zaroubi S., Kim T.-S., Tzanavaris P., Carswell R. F., 2002, *MNRAS*, 332, 367
 Worseck G., Prochaska J. X., Hennawi J. F., McQuinn M., 2014, preprint ([arXiv:1405.7405](https://arxiv.org/abs/1405.7405))

SUPPORTING INFORMATION

Additional Supporting Information may be found in the online version of this article:

Tables 1, 2, 3 and 4, and Figures 1, 2, 3 and 4. (<http://www.mnras.oxfordjournals.org/lookup/suppl/doi:10.1093/mnras/456/l79-183/-/DC1>). Please note: Oxford University Press is not responsible for the content or functionality of any supporting materials supplied by the authors. Any queries (other than missing material) should be directed to the corresponding author for the article.

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