

Possible evidence for an inverted temperature–density relation in the intergalactic medium from the flux distribution of the Ly α forest

J. S. Bolton,^{1*} M. Viel,^{2,3} T.-S. Kim,⁴ M. G. Haehnelt⁵ and R. F. Carswell⁵

¹Max-Planck-Institut für Astrophysik, Karl-Schwarzschild Str. 1, 85748 Garching, Germany

²INAF – Osservatorio Astronomico di Trieste, Via G. B. Tiepolo 11, I-34131 Trieste, Italy

³INFN/National Institute for Nuclear Physics, Via Valerio 2, I-34127 Trieste, Italy

⁴Astrophysikalisches Institut Potsdam, An der Sternwarte 16, D-14482 Potsdam, Germany

⁵Institute of Astronomy, University of Cambridge, Madingley Road, Cambridge CB3 0HA

Accepted 2008 February 14. Received 2008 February 7; in original form 2007 November 13

ABSTRACT

We compare the improved measurement of the Ly α forest flux probability distribution at $1.7 < z < 3.2$ presented by Kim et al. to a large set of hydrodynamical simulations of the Ly α forest with different cosmological parameters and thermal histories. The simulations are in good agreement with the observational data if the temperature–density relation for the low-density intergalactic medium (IGM), $T = T_0 \Delta^{\gamma-1}$, is either close to isothermal or inverted ($\gamma < 1$). Our results suggest that the voids in the IGM may be significantly hotter and the thermal state of the low-density IGM may be substantially more complex than is usually assumed at these redshifts. We discuss radiative transfer effects which alter the spectral shape of ionizing radiation during the epoch of He II reionization as a possible physical mechanism for achieving an inverted temperature–density relation at $z \simeq 3$.

Key words: hydrodynamics – methods: numerical – intergalactic medium – quasars: absorption lines.

1 INTRODUCTION

Traditional analyses of the Ly α forest observed in the spectra of high-redshift quasi-stellar objects (QSOs) decompose the flux distribution into a series of discrete absorption profiles, generally characterized by a Voigt function (Hu et al. 1995; Lu et al. 1996; Kim et al. 1997; Kirkman & Tytler 1997; Kim, Cristiani & D’Odorico 2001). Voigt profiles provide an accurate description for an absorption line if the absorber is a localized gas cloud with a Gaussian velocity dispersion (Rauch 1998). In the now widely established paradigm for the origin of the Ly α forest most of the absorption is, however, caused by extended gas distributions broadened by the Hubble flow, removing much of the physical motivation for the decomposition of the flux distribution into Voigt profiles. Such a decomposition is nevertheless useful, but it is non-unique and the line fitting process is time consuming and somewhat subjective (see Kim et al. 2007, hereafter K07, for an exhaustive discussion). Alternative characterizations of the flux distribution based on pixel statistics (Jenkins & Ostriker 1991; Rauch et al. 1997; Gaztañaga & Croft 1999; McDonald et al. 2000; Theuns, Schaye & Haehnelt 2000; Meiksin, Bryan & Machacek 2001; Viel et al. 2004c; Lidz et al. 2006; Becker, Rauch & Sargent 2007) and wavelet analysis (Meiksin 2000; Theuns & Zaroubi 2000; Zaldarriaga 2002) have

also been developed. These measures are simple and – at least in principle – easy to compare to the same quantities extracted from theoretical models of the Ly α forest (Cen 1992; Miralda-Escudé et al. 1996; Bi & Davidsen 1997; Croft et al. 1999; Davé et al. 1999; Theuns et al. 1998; Jena et al. 2005). In particular, the power spectrum of the Ly α forest flux distribution has recently been successfully developed into a quantitative tool for measuring the matter power spectrum on scales of $0.5\text{--}40 h^{-1}$ Mpc (e.g. Croft et al. 2002; Viel, Haehnelt & Springel 2004b; McDonald et al. 2005).

The simplest pixel statistic is the Ly α flux probability distribution function (PDF), which is sensitive to the density distribution and thermal state of the intergalactic medium (IGM) (e.g. Becker et al. 2007). Several attempts have been made to obtain joint constraints on a variety of cosmological and astrophysical parameters using the flux PDF together with the flux power spectrum (Choudhury, Srianand & Padmanabhan 2001; Desjacques & Nusser 2005; Lidz et al. 2006; Desjacques, Nusser & Sheth 2007). However, the PDF is also sensitive to a range of systematic uncertainties, most notably the assumed continuum level, the noise properties of the spectra and contamination by metal absorption lines in the Ly α forest. In high resolution, high signal-to-noise ratio (S/N) data at $z < 3$, the first two uncertainties have a small impact on the shape of the PDF (K07). The $z < 3$ metal absorption in the Ly α forest, however, alters the shape and amplitude of the PDF significantly, especially at $0.2 < F < 0.8$. Since many numerical simulations of the Ly α forest do not include metals, and those that do must also correctly reproduce their

*E-mail: jsb@mpa-garching.mpg.de

uncertain distribution and abundance, this can lead to substantial systematic uncertainties when attempting to fit simulation data to the observed flux PDF of the Ly α forest.

K07 have completed a detailed analysis of the impact of metal absorption lines on the Ly α PDF at $1.7 < z < 3.2$. They presented new measurements of the PDF in which all identifiable metal contamination has been removed from the Ly α forest on a line-by-line basis. The metal cleaned PDF had a significantly different shape when compared directly to the metal contaminated sample. In this paper we shall investigate what can be inferred about the underlying matter distribution and the thermal state of the IGM from these new PDF measurements using state-of-the-art hydrodynamical simulations of the Ly α forest. Our approach is different to most recent studies of the flux PDF which have used either pseudo-hydrodynamical techniques, dark-matter-only simulations or semi-analytical models (Choudhury et al. 2001; Desjacques & Nusser 2005; Lidz et al. 2006; Becker et al. 2007) which model the effect of gas pressure on the underlying IGM density distribution in an approximate way.

The structure of this paper is as follows. We discuss the observational data and the hydrodynamical simulations in Section 2. In Section 3, we qualitatively compare the simulation data to the observations and determine the impact of various astrophysical and numerical effects on the simulated flux PDF. A quantitative comparison is undertaken in Section 4 and the best-fitting models to the K07 data are presented along with a comparison to other recent studies in the literature. We briefly discuss the implications of our results for measurements of the matter power spectrum from the Ly α forest flux power spectrum in Section 5 and conclude in Section 6.

2 OBSERVATIONS AND SIMULATIONS OF THE LY α FOREST

2.1 The observational data

The Ly α flux PDF at $1.7 < z < 3.2$ was recently measured by K07 using a set of 18 high resolution ($R \sim 45\,000$), high-S/N (≥ 30 –50) VLT/UVES spectra. Metal lines were removed from the Ly α forest by directly fitting all identifiable absorption features. This approach to metal absorption removal is different to other methods in the literature, which either account for metal contamination statistically (Tytler et al. 2004; Kirkman et al. 2005) or by excising heavily contaminated regions from the spectra (Rauch et al. 1997; McDonald et al. 2000; Lidz et al. 2006). Following this procedure, K07 obtained the Ly α flux PDF in three redshift bins at $\langle z \rangle = 2.07$, $\langle z \rangle = 2.52$ and $\langle z \rangle = 2.94$. We shall compare these new improved measurements to the PDF computed by numerical simulations of the Ly α forest. Further details regarding the observational data and its reduction may be found in K07.

2.2 Numerical simulations of the IGM

The simulations used in this study were run using the parallel Tree-SPH code GADGET-2 (Springel 2005), modified to compute the ionization state of the gas using the non-equilibrium ionization algorithm of Bolton et al. (2006). The fiducial simulation volume is a periodic box $20 h^{-1}$ comoving Mpc in length containing 2×256^3 gas and dark matter particles. The gravitational softening length was set to 1/30th of the mean linear interparticle spacing. The gas in the simulations is photoionized by a spatially uniform ultraviolet background (UVB) based on the galaxies and QSOs emission model of Haardt & Madau (2001). The mean free path for hydrogen

ionizing photons is expected to be much larger than the mean separation between ionizing sources at $z < 3$, and so the assumption of a uniform UVB at the hydrogen Lyman limit should be a reasonable one (Bolton & Haehnelt 2007b). Star formation (SF) is included using a simplified prescription which converts all gas particles with overdensity $\Delta > 10^3$ and temperature $T < 10^5$ K into collisionless stars, significantly speeding up the simulations. The runs were all started at $z = 99$, with initial conditions generated using the transfer function of Eisenstein & Hu (1999). The simulations explore a range of cosmological and astrophysical parameters consistent with other observational data, and are summarized fully in Table 1.

Constraints on the thermal state of the low-density IGM are usually expressed in terms of a polytropic temperature–density relation,¹ $T = T_0 \Delta^{\gamma-1}$, expected to arise for $\Delta = \rho/\langle \rho \rangle \leq 10$ when photoheating by the UVB is balanced by cooling due to adiabatic expansion (Hui & Gnedin 1997; Valageas, Schaeffer & Silk 2002). Measurements of the thermal state of the IGM at $1.7 < z < 3.2$ from Ly α forest data give temperatures in the range $10^{3.9} < T_0 < 10^{4.4}$ K and are consistent with slopes in the range of those expected within the standard paradigm for the photoheated IGM, $1 < \gamma < 1.6$ (Ricotti, Gnedin & Shull 2000; Schaye et al. 2000; McDonald et al. 2001). The fiducial model (20-256) was chosen to be representative of these constraints. In order to explore the effect of different temperature–density relations on the flux PDF, we also run six further simulations with different thermal histories. Four simulations (20-256t1, 20-256t2, 20-256g1 and 20-256g2) reproduce values of T_0 and γ consistent with the upper and lower end of the range quoted above. We have also simulated two further models with $\gamma < 1$ (20-256g3) and an additional scatter in the temperature–density relation (20-256xs). We shall discuss the motivation behind these latter choices further in Section 3. The different thermal histories are constructed by modifying the fiducial simulation He II photoheating rate, $\epsilon_{\text{He II}}^{\text{fid}}$, such that $\epsilon_{\text{He II}} = \alpha \Delta^\beta \epsilon_{\text{He II}}^{\text{fid}}$. In the case of the model with extra scatter (20-256xs), a random Gaussian distributed dispersion is applied to the He II photoheating rate. In this way we generate self-consistent thermal histories, in the sense that the thermal state of the IGM is correctly coupled to the gas pressure. Note, however, that a fully physically consistent treatment of the IGM thermal history at $z \simeq 3$ would require cosmological radiative transfer simulations with a level of detail beyond at least our current numerical capabilities (although see Bolton, Meiksin & White 2004; Maselli & Ferrara 2005; Tittley & Meiksin 2007). The values of α and β used in all the simulations are summarized in Table 1, along with the resulting parameters for the temperature–density relation in each simulation.²

The cosmological parameters used for our fiducial simulation (20-256) are consistent with the combined analysis of the third-year *Wilkinson Microwave Anisotropy Probe* (WMAP) and Ly α forest

¹ The temperature–density relation is also sometimes referred to as the IGM effective equation of state, or simply just the equation of state. Note, however, that the equation of state for this low-density gas in the correct sense is the ideal gas equation.

² The values of γ in the fiducial model are smaller than the value of $\gamma \simeq 1.6$ typically reached under the assumption that the low-density gas is optically thin and in photoionization equilibrium following reionization. We have not rescaled the He II photoheating rate in this instance. This difference is due to non-equilibrium ionization effects which occur during rapid changes in the UVB (e.g. Miralda-Escudé & Rees 1994; Haehnelt & Steinmetz 1998; Theuns et al. 1998). The smaller values of γ are due to the hardening of the UVB spectrum around $z \simeq 3.5$ due to the increased contribution of QSOs to the metagalactic UVB, leading to the reionization of the He II in the simulation.

Table 1. Parameters used in our suite of hydrodynamical simulations. In all models we assume a flat universe with $\Omega_\Lambda = 1 - \Omega_m$, $\Omega_m = 0.26$, $\Omega_b h^2 = 0.024$, $h = 0.72$ and $n = 0.95$. The temperature–density relation is estimated at each redshift by a least-squares fit to the volume weighted temperature–density plane at $0.1 \leq \Delta \leq 1$ and $T_0 \leq 10^5$ K. The T_0 values have been rounded to three significant figures.

Model	σ_8	T_0 (K) ($z = 2.07$)	T_0 (K) ($z = 2.52$)	T_0 (K) ($z = 2.94$)	γ ($z = 2.07$)	γ ($z = 2.52$)	γ ($z = 2.94$)	α	β	Notes
20-256	0.85	14 800	17 600	20 800	1.37	1.31	1.27	1.0	0.0	Fiducial model
20-100	0.85	16 800	19 600	22 700	1.40	1.35	1.30	1.0	0.0	Low resolution
20-200	0.85	15 100	18 000	21 300	1.38	1.33	1.29	1.0	0.0	Medium resolution
20-400	0.85	14 400	17 200	20 300	1.37	1.31	1.26	1.0	0.0	High resolution
40-200	0.85	16 900	19 700	22 700	1.42	1.36	1.31	1.0	0.0	Low-resolution/medium box
80-400	0.85	16 900	19 800	22 800	1.41	1.36	1.31	1.0	0.0	Low-resolution/large box
20-256m	0.85	14 800	17 700	20 900	1.37	1.31	1.27	1.0	0.0	Multiphase SF
20-256w	0.85	18 900	20 900	23 300	1.47	1.39	1.32	1.0	0.0	Strong winds
20-256t1	0.85	10 900	12 600	14 600	1.41	1.36	1.31	0.5	0.0	Low T_0
20-256t2	0.85	17 600	21 300	25 500	1.35	1.29	1.25	1.4	0.0	High T_0
20-256g1	0.85	15 300	18 300	21 700	1.59	1.56	1.53	1.0	0.5	High γ
20-256g2	0.85	14 100	16 700	19 800	1.16	1.08	1.02	1.0	-0.4	Low γ
20-256g3	0.85	12 600	14 700	17 200	0.67	0.54	0.44	1.0	-1.2	Inverted γ
20-256xs	0.85	15 900	18 800	21 800	1.38	1.32	1.27	1.0	0.0	Extra scatter
20-256s1	0.75	14 600	17 400	20 700	1.37	1.32	1.27	1.0	0.0	$\sigma_8 = 0.75$
20-256s2	0.95	15 000	17 700	21 000	1.37	1.31	1.26	1.0	0.0	$\sigma_8 = 0.95$
20-256s2g2	0.95	14 200	16 700	19 700	1.15	1.07	1.01	1.0	-0.4	$\sigma_8 = 0.95$, low γ

data. $(\Omega_m, \Omega_\Lambda, \Omega_b h^2, h, \sigma_8, n) = (0.26, 0.74, 0.024, 0.72, 0.85, 0.95)$ (Seljak, Slosar & McDonald 2006b; Viel, Haehnelt & Lewis 2006). We also run three additional simulations with different values for σ_8 : one which is consistent with the third-year *WMAP* data alone (20-256s1, Spergel et al. 2007) and another two consistent with the somewhat higher constraints from weak lensing and Ly α forest data (20-256s2 and 20-256s2g2, Lesgourgues et al. 2007). The third model (20-256s2g2) also has a slope for the temperature–density relation which is consistent with the lower end of the observational constraints. Two further models were run using the multiphase SF model of Springel & Hernquist (2003a). One model includes the multiphase SF with winds disabled (20-256m), while the other includes the effect of galactic winds with a velocity of 484 km s^{-1} (20-256w). This latter model is extreme in the sense that the kinetic energy of the winds is comparable to the energy output from supernovae in the model (Springel & Hernquist 2003b) and should therefore provide an upper limit on the impact galactic winds may have on the PDF.

Lastly, to check numerical convergence we run five additional simulations with different box sizes (40-200, 80-400) and mass resolutions (20-100, 20-200, 20-400). The parameters of these simulations are listed in Table 2 along with a summary of all the other models, while the parameters characterizing their resolution and box size are listed in Table 2.

Table 2. Mass resolution and box size of the five additional simulations used to check numerical convergence. The mass resolution and box size of the fiducial simulation is also listed for comparison at the bottom of the table.

Model	Box size (comoving $\text{Mpc } h^{-1}$)	Total particle number	Gas particle mass ($M_\odot h^{-1}$)
20-100	20	2×100^3	1.03×10^8
20-200	20	2×200^3	1.29×10^7
20-400	20	2×400^3	1.61×10^6
40-200	40	2×200^3	1.03×10^8
80-400	80	2×400^3	1.03×10^8
20-256	20	2×256^3	6.13×10^6

2.3 Synthetic Ly α spectra

Sets of synthetic Ly α spectra are constructed at $z = 2.07, 2.52$ and 2.94 using 1024 random lines of sight drawn from the simulations (e.g. Theuns et al. 1998). The line profile convolution is performed with the Voigt profile approximation of Tepper García (2006). We rescale the synthetic spectra to match the mean normalized flux, $\langle F \rangle$, in the redshift bins used by K07 to measure the flux PDF: $\langle F \rangle = [0.863, 0.797, 0.730]$ at $z = [2.07, 2.52, 2.94]$. These values correspond to an effective optical depth, $\tau_{\text{eff}} = -\ln(\langle F \rangle) = [0.147, 0.227, 0.315]$.

The raw spectra are then processed to resemble the observational data. The spectra are convolved with a Gaussian with FWHM = 7 km s^{-1} and are rebinned on to pixels of width 0.05 \AA . Noise is then added to the spectra in a similar manner to Rauch et al. (1997) and McDonald et al. (2000). Within each of the three observationally defined redshift bins for the PDF, the variance of the noise in the observed spectra is determined over flux ranges corresponding to each flux bin of the PDF. Note that we have used the same binning for the flux PDF as McDonald et al. (2000) and K07. The PDF bins have a width of $\Delta F = 0.05$, with the first and last bins centred on $F = 0$ and 1 , respectively. This gives 21 bins in total from $F = 0$ to 1 . Gaussian distributed noise consistent with the variances in each flux bin of the observational data is then added to the corresponding pixels in the synthetic spectra. This process mimics the noise properties of the observed spectra well (Rauch et al. 1997). Note that all pixels with flux levels smaller than $F = 0.025$ or greater than $F = 0.975$ have been allocated to the $F = 0$ or 1 PDF bins, respectively.

3 COMPARISON OF THE K07 PDF TO HYDRODYNAMICAL SIMULATIONS

3.1 The effect of resolution, box size and galactic winds

Fig. 1 displays the effect of resolution, box size and galactic winds on the simulated flux PDF. The observational data of K07 are shown by the open diamonds with error bars, and the shaded regions in all

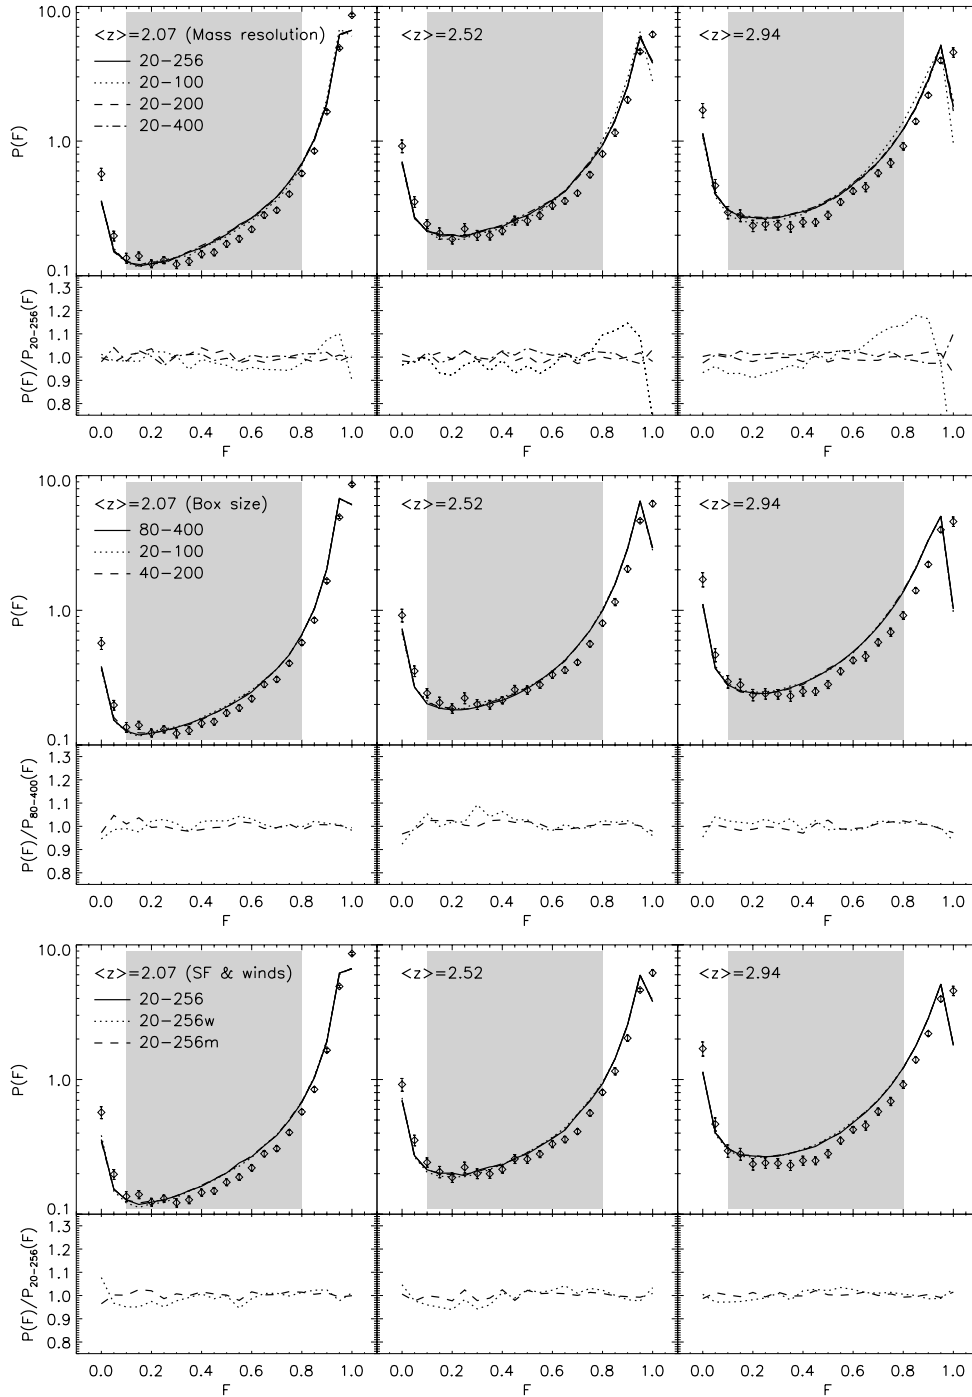


Figure 1. The impact of resolution, box size and galactic winds on the simulated flux PDF. Top row: The effect of mass resolution on the simulated flux PDF. The lower section in each panel displays the ratio of the flux PDFs to the fiducial model (20-256). Middle row: The effect of box size. The lower section in each panel displays the ratio of the flux PDFs to the model with the largest box size (80-400). Bottom row: The effect of the adopted SF model and galactic winds. The lower section in each panel again displays the ratio of the flux PDFs to the fiducial model (20-256). For comparison, the observational data of K07 are also shown (open diamonds with error bars) and the shaded grey region in each panel corresponds to the portion of the PDF which is least affected by noise and continuum uncertainties.

panels correspond to the flux range over which noise properties and continuum placement have a small impact on the PDF, $0.1 \leq F \leq 0.8$ (K07).

The top row in Fig. 1 displays the effect of mass resolution on the simulated PDF. The ratio of the three simulated PDFs with different mass resolutions to the fiducial model PDF (20-256, see Table 2 for details) is shown in the lower section in each of the panels. The

simulated PDF appears to be marginally converged and is generally to within a few per cent of the highest resolution run at all redshifts. As we shall see in the next section, these differences are small compared to those exhibited when other parameters are varied, most notably σ_8 and γ . Note, however, the flux PDF of the fiducial model does not agree well with the PDF from the observational data of K07. We shall return to this point later in this paper.

In the middle row, the effect of the simulation box size on the PDF is displayed. The lower sections of each panel show the ratio of the PDFs from the two smaller boxes to the PDF extracted from the largest box (80-400). All three simulations have the same mass resolution. Again, the differences in the simulated PDFs are very small. We therefore consider our fiducial simulation box size and mass resolution to be adequate for distinguishing between the various IGM models considered in this work.

The bottom row of panels in Fig. 1 shows a comparison of the PDFs computed from the fiducial model (20-256, solid curve) to the results from the simulations including the multiphase SF model of Springel & Hernquist (2003a) with (20-256w, dotted curve) and without (20-256m, dashed curve) winds. The lower sections in each panel again display the ratio of these PDFs to the fiducial model. The differences between the fiducial and multiphase SF model PDFs are within a few per cent, a result which gives us confidence the simple SF prescription used for the majority of the simulations in this work has little impact on the properties of the Ly α forest. This is because the highly overdense regions containing most of the SF correspond to only a very small proportion of the volume probed by the synthetic spectra. The difference between the model with strong galactic winds is also small, except at $F = 0$ in the $\langle z \rangle = 2.07$ bin. However, this is a rather extreme model and may overestimate the impact of winds on the Ly α forest. In comparison, after rescaling the mean flux of their synthetic spectra, Theuns et al. (2002c) find winds also have a very small effect on the Ly α forest. A similar result was found by Bruscoli et al. (2003). Bertone & White (2006) used a semi-analytical model to investigate the impact of galactic winds on the Ly α forest, and also found the impact of winds on the PDF is negligible. Our results are in qualitative agreement with these studies, and we expect the effect of galactic winds on the Ly α flux PDF to be small. Although we have not analysed the possible impact of outflows from rarer, more massive galaxies hosting active galactic nuclei (AGN) on the Ly α forest, these are expected to extend into a smaller volume of the IGM than the supernova driven galactic winds (Sijacki et al. 2007). AGN feedback is therefore also likely to have a negligible impact on the Ly α forest flux distribution.

3.2 The effect of cosmological and astrophysical parameters

The flux PDFs constructed from the simulations with different cosmological and astrophysical parameters are compared to the observational data of K07 (open diamonds with error bars) in Fig. 2. From the left- to right-hand side, in each of the four panels the effect of varying T_0 , σ_8 , γ and τ_{eff} on the PDF is shown. The solid curve in each panel corresponds to the PDF of the fiducial model (20-256). The lower segment of each panel shows the ratio of the PDFs indicated in the diagram to the fiducial model. Note again that except where indicated in the rightmost panels of Fig. 2, all the synthetic spectra are rescaled to have $\tau_{\text{eff}} = [0.147, 0.227, 0.315]$ at $\langle z \rangle = [2.07, 2.52, 2.94]$, corresponding to the effective optical depth in the redshift bins for which K07 have reported their measurement of the flux PDF.

At all redshifts the PDF is relatively insensitive to changing the IGM temperature at mean density, T_0 , for values within the observationally constrained range (Schaye et al. 2000). Higher temperatures slightly increase the amplitude of the PDF at $0.1 < F < 0.9$ and decrease the amplitude at $F = 0$ and 1. This behaviour is attributable to the Jeans smoothing effect of gas pressure on the spectra; higher temperatures increase the IGM gas pressure and smooth the IGM over larger scales. Combined with the increased Doppler widths in the higher temperature models, this redistributes the PDF by ef-

fectively erasing the small-scale power in the spectra. Varying σ_8 has a slightly larger effect. Models with smaller values of σ_8 tend to produce broader absorption lines which arise from the slightly smoother gas distribution at scales comparable to the Jeans length. The effect is to again erase small-scale power. The effect of σ_8 and temperature on the flux PDF are, however, small and not sufficient to account for the differences between the K07 data and the simulations. A more detailed discussion of the effect of these parameters on flux statistics can be found in Theuns et al. (2000).

Varying the temperature–density relation index, γ , has a much stronger effect on the PDF. Altering γ changes the temperature of the IGM in a density-dependent fashion; decreasing γ increases the IGM temperature below the mean IGM density and decreases the temperature above mean density. If we assume H I to be in photoionization equilibrium, the neutral hydrogen density scales approximately as $n_{\text{HI}} \propto T_0^{-0.7} \Gamma_{\text{HI}}^{-1} \Delta^{2-0.7(\gamma-1)}$, where Γ_{HI} is the photoionization rate per H I atom (e.g. Rauch 1998). Note that this is not the case for changes in T_0 only, where the gas temperature (and hence also the residual H I fraction in the IGM) are altered independently of the gas density.

These combined effects explain the changes seen in the PDF for different γ . As shown in Fig. 2, a lower value of γ decreases the amplitude of the PDF moderately over the flux range $0.1 < F < 0.9$, an effect which is compensated for by a strong increase of the amplitude over the flux ranges at $F = 0$ and 1. Absorption lines associated with moderately overdense regions become associated with colder, more neutral material and hence are narrower and have a higher opacity in the centre of the line. In contrast, underdense regions associated with transmission near the continuum level become hotter and therefore more highly ionized. Expressed in a different way, a larger γ tends to erase the small-scale power in the spectra by increasing the temperature of the gas associated with the most prominent absorption lines. Note that the effect of changing γ is therefore partially degenerate with changing T_0 and σ_8 within the observationally constrained ranges, although it has a significantly larger impact on the shape of the PDF. Fig. 2 suggests that a temperature–density relation which is close to isothermal is in better agreement with the K07 data.

Finally, the effect of varying τ_{eff} on the PDF is illustrated in the rightmost panels of Fig. 2. The value of τ_{eff} measured in observed spectra exhibits rather large fluctuations between different lines of sight, and it is rather sensitive to the number of strong absorption systems (Viel et al. 2004a) and the contribution of metal absorption (K07). It is therefore reasonable to consider τ_{eff} as an independent free parameter within the measurement errors of the observations. The values displayed correspond to a mean flux range of ± 0.020 around the fiducial values of $\langle F \rangle = [0.863, 0.797, 0.730]$ at $\langle z \rangle = [2.07, 2.52, 2.94]$. The PDF is very sensitive to even small changes of τ_{eff} but in a qualitatively different way to the other parameters. Altering τ_{eff} changes the shape of the PDF by increasing the PDF amplitude at one end and decreasing it at the other. This is in contrast to the effect of the other parameters, which when varied raise or lower the PDF at both ends simultaneously.

3.3 Radiative transfer effects during reionization

Thus far we have assumed a tight, polytropic temperature–density relation for the low-density IGM, $T = T_0 \Delta^{\gamma-1}$, with $1 \leq \gamma \leq 1.6$, which in the standard paradigm for the thermal state of the photoionized IGM is expected to apply to the majority of the optically thin gas responsible for the absorption in the Ly α forest. The relation arises from the balance between photoheating and cooling by

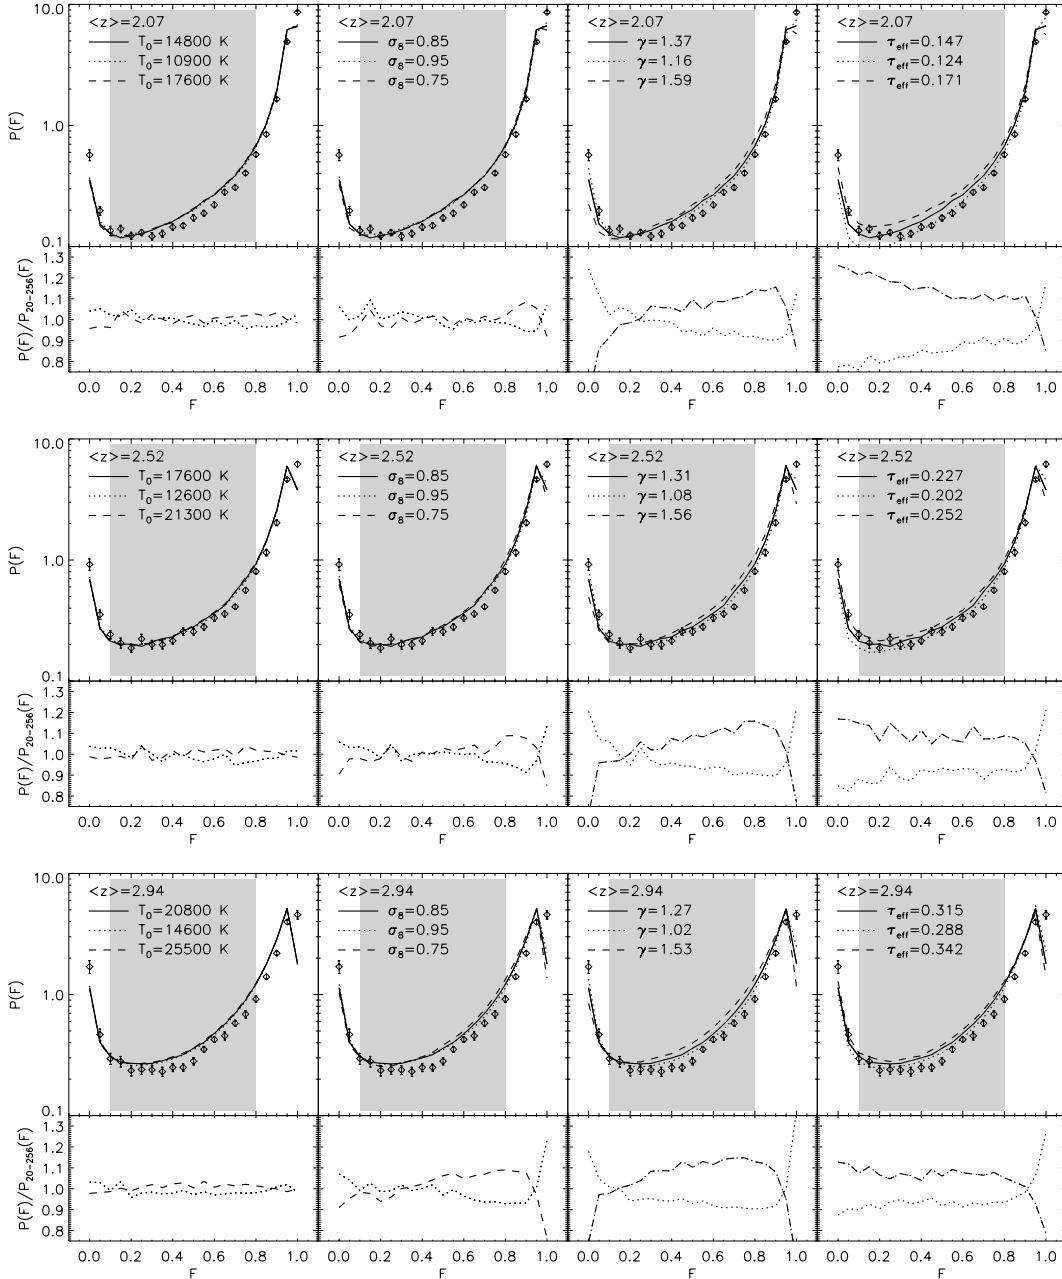


Figure 2. Comparison of the flux PDF computed from hydrodynamical Ly α forest simulations with different cosmological and astrophysical parameters to the observational data of K07. Top row: The observed PDF at $\langle z \rangle = 2.07$ (open diamonds with error bars). The solid curve in each panel shows the PDF derived from the fiducial model. From left- to right-hand side, the dotted and dashed curves demonstrate the effect on the PDF of varying T_0 , σ_8 , γ and τ_{eff} within the observationally determined boundaries. The lower section of each plot shows the ratio of these PDFs to the fiducial model. Middle row: As for top row but at $\langle z \rangle = 2.52$. Bottom row: As for top row but at $\langle z \rangle = 2.94$. The shaded grey region in each panel corresponds to the portion of the PDF which is least affected by noise and continuum uncertainties.

adiabatic expansion, with a slope which evolves from $\gamma \simeq 1$ following reionization to an asymptotic value of $\gamma \simeq 1.6$ (Hui & Gnedin 1997; Theuns et al. 1998; Valageas et al. 2002). Observational evidence for a boost in the IGM temperature and/or a flattening in the temperature–density relation at $z \simeq 3$ provides possible evidence for He II reionization (Ricotti et al. 2000; Schaye et al. 2000; Theuns et al. 2002b; Zaroubi et al. 2006). As we saw in the last section, the observed flux PDF indeed appears to prefer a thermal state where the voids in the IGM are relatively hot, although this model is still not in perfect agreement with the data. It is therefore possible that

a model with an inverted temperature–density relation with $\gamma < 1$ may instead provide a better fit to the data.

Is such a thermal state for the IGM plausible? During reionization the assumption of a tight temperature–density relation for the low-density IGM is likely to be a poor one. The filtering of ionizing radiation by optically thick gas alters the spectral shape of the UVB, giving rise to additional heating effects which raise the gas temperature and blur the dependence of the IGM temperature on the gas density (Abel & Haehnelt 1999; Bolton et al. 2004; Maselli & Ferrara 2005; Tittley & Meiksin 2007). The inhomogeneous

distribution of the ionizing sources themselves will also contribute to spatial variations in the IGM thermal state (Miralda-Escudé, Haehnelt & Rees 2000; Gleser et al. 2005; Furlanetto & Oh 2007). In addition, the cooling time-scale for the low-density gas responsible for the Ly α forest is of order a Hubble time, meaning that changes to the thermal state of the IGM will persist for some time after reionization (Miralda-Escudé & Rees 1994).

Radiative transfer effects on the thermal state of the IGM may thus result in an inversion of the temperature–density relation, such that $\gamma < 1$. As ionizing photons propagate through optically thick, moderately overdense regions in the IGM, the spectral shape of the emission is considerably hardened as lower energy photons are preferentially absorbed. Any subsequent underdense gas is then subjected to ionizing photons which have a substantially higher mean energy, boosting the temperature of the underdense region as it is photoionized and heated. One may envisage such a situation if ionizing sources are initially embedded in overdense, optically thick regions. State-of-the-art cosmological radiative transfer simulations indicate that this ‘inside-out’ topology may be appropriate for moderate overdensity and underdensity, at least during H I reionization (Iliev et al. 2006b). Such a radiative transfer induced temperature inversion effect was noted by Bolton et al. (2004), although this result was a specific case for a single density distribution. More generally, radiative transfer effects during inhomogeneous reionization are expected to produce a more complex, multiple valued relationship between temperature and density in the IGM. While we were in the final stages of preparing this manuscript Furlanetto & Oh (2007) presented a semi-analytical model aimed at describing the thermal state of the IGM during and after the inhomogeneous reionization of helium. While probably not fully realistic due to a lack of a proper modelling of the filtering and spectral hardening effects of radiative transfer, the model confirms that an initially inverted and multivalued temperature–density relation is indeed plausible. The exact character of these effects will also depend on the types of sources which reionize the IGM (Tittley & Meiksin 2007).

Consequently, adopting a temperature–density relation with $\gamma < 1$ may mimic some of the different heating effects expected due to inhomogeneous radiative transfer through an optically thick IGM. One may expect many underdense regions to exhibit significantly elevated temperatures following reionization (Abel & Haehnelt 1999; Bolton et al. 2004; Gleser et al. 2005; Tittley & Meiksin 2007). Helium in the IGM at $z \simeq 3$ exhibits patches which are still optically thick to photons with energies in excess of 4 Ry (e.g. Jakobsen et al. 1994; Heap et al. 2000; Reimers et al. 2005). Radiative transfer effects on the Ly α flux PDF may therefore be important at the redshifts we consider here, especially if the tail end of He II reionization occurs at $z \simeq 3.2$ (e.g. Songaila 1998; Schaye et al. 2000; Theuns et al. 2002b; Bernardi et al. 2003; Shull et al. 2004; Bolton et al. 2006; Faucher-Giguere et al. 2007). Furthermore, using a semi-analytical model for the Ly α forest based on density distributions drawn from hydrodynamical simulations, Becker et al. (2007) noted that an inverted temperature–density relation improved their fit to the flux PDF measured with an independent set of high-resolution QSO spectra at $1.7 < z < 5.8$.

The inhomogeneous spatial distribution of ionizing sources is, however, also expected to increase the *scatter* in the temperature–density relation (Theuns et al. 2002d). We therefore compare two further different models to the K07 data, one with an inverted temperature–density relation (20-256g3) and another with a moderately increased scatter in the temperature–density relation (20-256xs). For illustration, the volume weighted temperature–density planes for these two models at $z = 2.07$ are displayed in Fig. 3. The

fiducial model is shown in the rightmost panel. The dashed lines in each panel correspond to the temperature–density relations listed in Table 1.

Fig. 4 displays the comparison between the flux PDFs constructed from these models to the observational data. Note that we have again matched τ_{eff} in the simulated spectra to the values measured by K07. The PDFs for the fiducial model (solid curves) and the model with increased scatter in the temperature–density relation (dashed curves) are very similar. At all redshifts the flux PDFs exhibit differences of only a few per cent. A moderate increase in the scatter of the temperature–density relation appears unlikely to substantially affect the shape of the PDF. However, the temperature–density relation is still relatively well defined in this model. There is also no good reason to assume a Gaussian dispersion. In reality the scatter may be rather more inhomogeneous. We therefore should not rule out the possibility that more complex, multiple valued relationship between temperature and density fits the K07 data better. We now turn to the model with an inverted temperature–density relation (dotted curves). In this case the modification of the thermal state of the IGM makes a substantial difference to the shape of the PDF, as expected from the behaviour seen in Fig. 2. The inverted temperature–density relation brings the simulated and observed flux PDFs into excellent agreement. The previously most problematic points at the extreme ends of the PDF at flux levels close to zero and unity now show an excellent match between the simulations and data. Note again, however, that the IGM may nevertheless have a more complex thermal state following He II reionization, rather than actually having a well defined, inverted temperature–density relation. We now proceed to discuss the quality of the fits of different models more quantitatively in the next section.

4 DETERMINATION OF THE BEST-FITTING MODELS TO THE K07 PDF

4.1 A quantitative comparison of the simulated and observed data

As should be clear from the previous section and analyses by other authors, there is considerable degeneracy in the dependence of the flux PDF on cosmological and astrophysical parameters. The quality of the observational data and the apparently good agreement with the flux PDF from our inverted temperature–density relation simulation motivated us, however, to perform a quantitative χ^2 analysis. For this we used the full covariance matrix derived from the observational data of K07.

As discussed by McDonald et al. (2000), using the covariance matrix is important; the errors bars on the PDF are highly correlated. However, the covariance matrices derived by K07 (their equation 1) are rather noisy for this purpose, preventing reliable inversion. We therefore follow the method suggested by Lidz et al. (2006), where the observed covariance matrix is regularized using the correlation coefficients, $r_s(i, j) = \text{cov}_s(i, j) / \sqrt{\text{cov}_s(i, i)\text{cov}_s(j, j)}$ computed for each set of synthetic spectra, such that $\text{cov}_d(i, j) = r_s(i, j) \sqrt{\text{cov}_d(i, i)\text{cov}_d(j, j)}$. Here cov_d and cov_s are the covariance matrices computed from the observed data and synthetic spectra, respectively. Note that this implicitly assumes the observed and simulated spectra have similar covariance properties.

We perform a simple χ^2 analysis on the K07 data for two different cases. In the first instance we adopt a conservative approach and analyse the PDF only over the flux range $0.1 \leq F \leq 0.8$, where the impact of noise properties and continuum uncertainties on the K07 PDF data are small. This is similar to the approach adopted

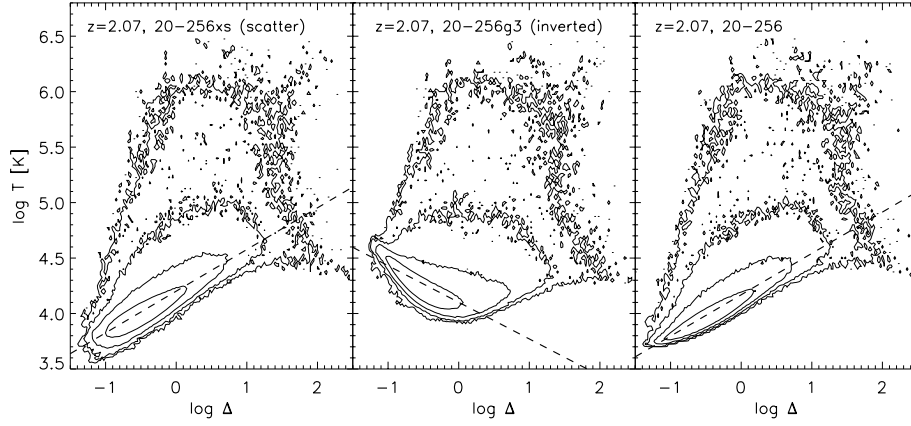


Figure 3. Contour plots of the volume weighted temperature–density plane at $z = 2.07$ in three of our hydrodynamical simulations. The number density of the data points increases by an order of magnitude with each contour level. The model in the left-hand panel (20-256xs) has a temperature–density relation which exhibits a moderately increased scatter, while the central panel displays the simulation (20-256g3) with an inverted temperature–density relation. For comparison, the temperature–density plane in the fiducial model (20-256) is shown in the right-hand panel. The dashed lines in each panel correspond to the temperature–density relations listed in Table 1.




Figure 4. Comparison of the observational data of K07 at $\langle z \rangle = [2.07, 2.52, 2.94]$ (open diamonds with error bars) to simulations with an inverted temperature–density relation (20-256g3, dotted curve) and a model with increased scatter (20-256xs, dashed curve). The solid curve in each upper panel shows the PDF derived from the fiducial model (20-256), while the lower panels display the ratio of the other PDFs to the fiducial model. The shaded grey region in each panel corresponds to the portion of the PDF which is least affected by noise and continuum uncertainties.

by Desjacques & Nusser (2005). However, as demonstrated in the last section, the extreme ends of the PDF at low and high flux levels contain important information on the thermal state of the IGM. For high-resolution, high-S/N data like that of K07, the uncertainties are expected to be relatively small over the entire range of the flux PDF. The continuum fits are generally accurate to within 1–2 per cent, and the typical S/N $\simeq 50$ (see K07 for a detailed discussion regarding the effect of these parameters on the PDF). We therefore also perform the χ^2 analysis over the entire PDF, $0 \leq F \leq 1$.

In our analysis we do not attempt a global minimization of χ^2 over the parameters we have varied in the simulations. We instead just tabulate the minimum reduced χ^2 values for our grid of models with different T_0 , σ_8 and γ . We vary τ_{eff} over an appropriate range, 25 per cent either side of the central values from the fit to the redshift evolution of the metal-cleaned Ly α effective optical depth reported by K07, $\tau_{\text{eff}} = 0.0023 \pm 0.0007(1+z)^{3.65 \pm 0.21}$. This scatter reflects the variance in the data. Continuum fitting uncertainties, although small

at the redshifts we consider here, are another source of error in constraints on this quantity. We shall discuss this in more detail shortly.

The results of the χ^2 analysis for the flux range $0.1 \leq F \leq 0.8$ are summarized in Table 3. In addition to the fiducial model (20-256) we tabulate the results for five other simulations with varying T_0 , γ and σ_8 . The models which have a low σ_8 , high T_0 , high γ and increased scatter in the temperature–density relation are omitted from Table 3, as all have reduced χ^2 values which are similar to or worse than the fiducial model. The inverted temperature–density relation model (20-256g3) provides the best fit to the K07 data with $\chi^2/\nu = [1.62, 1.44, 0.90]$ at $\langle z \rangle = [2.07, 2.52, 2.94]$ for $\nu = 11$ degrees of freedom, corresponding to probabilities of [9, 15, 54] per cent. These are therefore reasonable fits in a quantitative sense. We would, however, urge the reader to exercise some caution in interpreting the absolute χ^2 values; estimating the observational PDF errors and the covariance matrix is problematic. We may have underestimated the observational errors, especially at low flux levels. Note that the non-linear high-density regions responsible for the

Table 3. The minimum χ^2 values for the best-fitting models to the K07 PDF data at $\langle z \rangle = [2.07, 2.52, 2.94]$, fitted in the range $0.1 \leq F \leq 0.8$ only. The columns list the best reduced χ^2 (for $\nu = 15 - 4 = 11$ degrees of freedom) and the corresponding best-fitting value for the effective optical depth, τ_{fit} . This is expressed as a ratio with the mean value of τ_{eff} in each observational redshift bin. For reference these are $\tau_{\text{eff}} = [0.147, 0.227, 0.315]$ at $\langle z \rangle = [2.07, 2.52, 2.94]$.

Model	χ^2/ν ($z = 2.07$)	χ^2/ν ($z = 2.52$)	χ^2/ν ($z = 2.94$)	$\tau_{\text{fit}}/\tau_{\text{eff}}$ ($z = 2.07$)	$\tau_{\text{fit}}/\tau_{\text{eff}}$ ($z = 2.52$)	$\tau_{\text{fit}}/\tau_{\text{eff}}$ ($z = 2.94$)	Notes
20-256	4.57	2.84	2.03	0.840	0.857	0.817	Fiducial model
20-256t1	3.61	2.49	1.86	0.864	0.871	0.822	Low T_0
20-256s2	4.05	2.65	1.65	0.886	0.851	0.819	$\sigma_8 = 0.95$
20-256g2	2.34	1.90	1.34	0.898	0.927	0.849	Low γ
20-256s2g2	1.85	1.81	1.14	0.888	0.925	0.923	$\sigma_8 = 0.95$, low γ
20-256g3	1.62	1.44	0.90	0.935	1.025	1.013	Inverted γ

Table 4. The χ^2 analysis results for the best-fitting models to the whole PDF, $0 \leq F \leq 1$. As in Table 3, the columns list the best reduced χ^2 (for $\nu = 21 - 4 = 17$ degrees of freedom) and the corresponding best-fitting value for the effective optical depth, τ_{fit} , as a ratio of the mean observed value in each redshift bin. For reference these are $\tau_{\text{eff}} = [0.147, 0.227, 0.315]$ at $\langle z \rangle = [2.07, 2.52, 2.94]$.

Model	χ^2/ν ($z = 2.07$)	χ^2/ν ($z = 2.52$)	χ^2/ν ($z = 2.94$)	$\tau_{\text{fit}}/\tau_{\text{eff}}$ ($z = 2.07$)	$\tau_{\text{fit}}/\tau_{\text{eff}}$ ($z = 2.52$)	$\tau_{\text{fit}}/\tau_{\text{eff}}$ ($z = 2.94$)	Notes
20-256	11.15	15.28	13.47	0.781	0.833	0.823	Fiducial model
20-256t1	10.02	15.22	14.14	0.808	0.840	0.844	Low T_0
20-256s2	6.90	9.65	9.62	0.820	0.850	0.854	$\sigma_8 = 0.95$
20-256g2	5.03	8.14	8.03	0.884	0.914	0.902	Low γ
20-256s2g2	3.07	4.70	5.32	0.923	0.930	0.927	$\sigma_8 = 0.95$, low γ
20-256g3	3.06	1.56	1.82	1.057	1.029	1.025	Inverted γ

absorption at low flux levels are also the most sensitive to galactic winds and other uncertain details of the numerical simulations. It is nevertheless gratifying that the best-fitting values for τ_{eff} are all within 7 per cent of the mean values in each redshift bin, well within the variance of this quantity. The model with an almost isothermal temperature–density relation and $\sigma_8 = 0.95$ (20-256s2g2) is also in marginal agreement with the data. Note that although some of the remaining models are in reasonable agreement with the observed PDF for the highest redshift bin ($z = 2.94$), they are all in poor agreement if the full redshift range of the K07 data set is taken into account.

In Table 4 we list the minimum reduced χ^2 values for a comparison of the simulated and observed spectra over the full range of the PDF, $0 \leq F \leq 1$. The inverted temperature–density relation again provides the best fit to the K07 data with $\chi^2/\nu = [3.06, 1.56, 1.82]$ at $\langle z \rangle = [2.07, 2.52, 2.94]$ for $\nu = 17$ degrees of freedom. The formal

probabilities of the reduced χ^2 values are now rather low, with a very small probability at $\langle z \rangle = 2.07$ and [7.2] per cent probability at $\langle z \rangle = [2.52, 2.94]$. However, as already discussed, estimating the errors of the PDF is difficult. Rather than being evidence against this model this may suggest that we have somewhat underestimated the errors at the extreme ends of the PDF, especially at low flux levels. The best-fitting values for τ_{eff} are again within 6 per cent of the mean values in each redshift bin. All the other models are very bad fits.

In Fig. 5, we show the best fits to the K07 flux PDF over the range $0.1 \leq F \leq 0.8$ (solid curves) and $0 \leq F \leq 1$ (dotted curves). The models are almost indistinguishable to the eye in the two highest redshift bins, although one may clearly perceive the difference between the different PDFs at $\langle z \rangle = 2.07$. The poor fit to the K07 data at $\langle z \rangle = 2.07$ for $0 \leq F \leq 1$ is mainly due to the influence of the data




Figure 5. The best fit to the PDF data of K07 (open diamonds with error bars) for the inverted temperature–density relation model. The solid curves show the best fits over the flux range $0.1 \leq F \leq 0.8$, corresponding to the data in the shaded portion of the plots, while the dotted curves show the best χ^2 fits to the entire PDF. The best-fitting values of the effective optical depth for the models are also displayed in each panel.

point at $F = 0$; the agreement between the observed and simulated PDF at $F > 0.8$ is very good. Considering the results presented in Fig. 1, a model with strong galactic winds may also alleviate some of this discrepancy. Even though we have attempted to accurately model the noise properties of the observed spectra, differences in the noise level in the spectra may also play a small role in this disagreement (K07). In any case, the improved agreement between the observational data and the Ly α forest simulations with the inverted temperature–density relation is impressive.

Finally, we note that the best-fitting τ_{eff} values for the inverted temperature–density relation model correspond to a photoionization per hydrogen atom of $\Gamma_{\text{HI}} = [1.43, 1.05, 0.97] \times 10^{-12} \text{ s}^{-1}$ for $0.1 \leq F \leq 0.8$ and $\Gamma_{\text{HI}} = [1.14, 1.05, 0.95] \times 10^{-12} \text{ s}^{-1}$ for $0 \leq F \leq 1$ at $\langle z \rangle = [2.07, 2.52, 2.94]$. These values are in excellent agreement with the results of Bolton et al. (2005) which are based on independent measurements of the Ly α forest opacity (Schaye et al. 2003).

4.2 The effect of continuum uncertainties on the χ^2 analysis

We now consider how important the continuum placement on the observed spectra is for the quality of the fit of the flux PDF at $0 \leq F \leq 1$. K07 discussed the effect of varying the continuum on the observational data in detail. For their typical estimated continuum uncertainty of ~ 1 –2 per cent, the changes in the PDF are small and are most evident at $F > 0.8$. Tytler et al. (2004) and Kirkman et al. (2005) also estimated continuum uncertainties to be of order 2 per cent based on an analysis of a large set of moderate-resolution Ly α forest spectra at $z < 3.2$. Faucher-Giguere et al. (2007) estimate the continuum placement to be somewhat larger with a systematic bias towards an underestimate of the continuum level by as much as 4 per cent at $z = 3$ and 1 per cent at $z = 2$ in a recent measurement of the Ly α effective optical depth using high-resolution data. Note, however, that while a systematic misplacement of the continuum appears likely with decreasing mean flux level (and thus increasing redshift) it will also depend strongly on the shape of the Ly α flux PDF at high flux levels and thus on the temperature of the gas in low-density regions of the IGM.

We have checked for the presence of such a potential bias by comparing all our Ly α forest models to the K07 data with synthetic spectra for which we have raised or lowered the continuum by 1.5 and 3 per cent. Raising the continuum on the synthetic spectra significantly worsens the agreement between the K07 data and all the simulations. Lowering the continuum on the synthetic spectra, mimicking a continuum which has been placed too low on the observational data, also strongly increases the χ^2 values, although the models with either an inverted or isothermal temperature–density relation are still slightly preferred in some redshift bins. Misplacing the continuum changes the flux PDF in a complex way. The fact that the agreement between the observed data and all models worsens significantly on varying the continuum suggests that continuum fitting errors for the high-quality data of K07 are under control at a level of 1–2 per cent, consistent with the estimates of the continuum uncertainties by K07. Altering the continuum placement on the synthetic spectra does not alter our conclusions regarding the best-fitting model to the K07 data.

4.3 Comparison to other studies

Several other detailed studies of the Ly α flux PDF are presented within the existing literature, and we now discuss some of these results in comparison to this work.

McDonald et al. (2000) (hereafter M00) presented a measurement of the flux PDF at $\langle z \rangle = [2.41, 3.00, 3.89]$ from a sample of eight high-resolution QSO spectra. A detailed comparison between the widely used M00 measurement and the measurement used in this work is given in K07. Here we comment on the interpretation of the M00 data when compared to hydrodynamical Ly α forest simulations. M00 compare their observational data to the L10 Eulerian hydrodynamical simulation of Miralda-Escudé et al. (1996). The simulation parameters were $\Omega_{\text{m}} = 0.4$, $\Omega_{\Lambda} = 0.6$, $\Omega_{\text{b}} h^2 = 0.015$, $h = 0.65$, $\sigma_8 = 0.79$, $n = 0.95$, $T_0 = 16\,000 \text{ K}$ and $\gamma \simeq 1.3$ at $z = 3$, with a box size of $10 h^{-1}$ comoving Mpc and a 288^3 grid. Performing a χ^2 analysis with the covariance matrix, M00 obtain $\chi^2/\nu = [4.4, 1.7, 1.1]$ for 19 degrees of freedom at $\langle z \rangle = [2.41, 3.00, 3.89]$. Note, however, that they do not minimize over τ_{eff} as we do. Instead, M00 rescale the synthetic spectra to match the measured effective optical depth of the observational data, $\tau_{\text{eff}} = [0.201, 0.380, 0.744]$. The fit is performed over the range $0 \leq F \leq 1$, and they attribute the disagreement at $\langle z \rangle = 2.41$ as being due to the effect of continuum uncertainties on the PDF.

As discussed in the last section this is in contrast with our findings: using updated cosmological parameters we find the K07 PDF, which is systematically lower at $0.2 < F < 0.7$ than the M00 data, strongly favours an inverted temperature–density relation when fitting over $0 \leq F \leq 1$. It is not obvious if this difference is due to the different shape of the measured PDFs, differences between the model parameters or differences in the numerical simulations. To investigate this further we have attempted to fit the flux PDF of our simulations at $z = 2.94$ to the M00 measurement of the flux PDF at $\langle z \rangle = 3$. The small difference in redshift should not be important here. When minimizing over τ_{eff} we find that all of the models listed in Table 3 are in good agreement with the M00 data for $0.1 \leq F \leq 0.8$ (all have $\chi^2/\nu \leq 1.04$). The inverted temperature–density relation model still provides the best reduced χ^2 ($\chi^2/\nu = 0.41$ for $\tau_{\text{eff}} = 0.411$) perhaps suggesting that the errors in this flux range have been overestimated by M00. When performing the χ^2 analysis over the entire M00 PDF, we were at first not able to obtain an acceptable fit to the M00 data. Note, however, that M00 also adjusted the normalization of their synthetic spectra to take into account possible continuum errors of the observational data. They achieved this by identifying the maximum transmitted flux in each synthetic line of sight, F_{c} , as the continuum flux and then proceeded to divide the flux in all other pixels in the line of sight by F_{c} .

We therefore again raise or lower the continuum in our synthetic spectra by 1.5 per cent, this time for comparison to the M00 data at $\langle z \rangle = 3$. As for the K07 data, raising the continuum on the synthetic spectra significantly worsens the agreement between the M00 data and all our models. Lowering the continuum in this instance, however, improves the agreement between the M00 data and simulations remarkably. Assuming our synthetic spectra provide an accurate representation of the Ly α forest, this suggests that the continuum in the absorption spectra on which the M00 measurement is based may indeed be placed somewhat low. If we lower the continuum level by 1.5 per cent, we find the model with a temperature–density relation close to isothermal and a larger $\sigma_8 = 0.95$ (20-256s2g2) provides the best fit to the M00 data, with $\chi^2/\nu = 1.08$ for an effective optical depth $\tau_{\text{eff}} = 0.385$, which is pleasingly close to the M00 value. The inverted temperature–density relation model (20-256g3) provides the next best fit with $\chi^2/\nu = 1.91$ for $\tau_{\text{eff}} = 0.418$. This is in contrast to M00, who achieved a reasonable fit to the data using different (now somewhat outdated) cosmological parameters and a standard temperature–density relation. The different conclusion we reach in this work is therefore probably largely due to the updated

cosmological parameters and perhaps also numerical differences in the simulations. In summary, we conclude that the M00 PDF measurements are consistent with our numerical simulations if the continuum is lowered by 1.5 per cent. The M00 measurement also appears to favour a flattened or inverted temperature–density relation at $z \simeq 3$, consistent with our results based on the independent K07 data set.

A more recent joint analysis of the best-fitting model to the M00 PDF and flux power spectrum was undertaken by Desjacques & Nusser (2005) at $z = 3$. They used N -body simulations with 256^3 dark matter particles and a box size of $25 h^{-1}$ comoving Mpc. Synthetic Ly α forest spectra were constructed by smoothing the dark matter density and velocity fields with a Gaussian filter over a characteristic scale, which was left as an extra free parameter in their analysis. Unfortunately, this extra free parameter complicates our comparison to this work as it is degenerate with other relevant quantities, most notably γ . Consequently, Desjacques & Nusser (2005) do not advocate a value for this quantity. On the other hand, Desjacques & Nusser (2005) also found that their joint analysis of the M00 flux power spectrum and PDF (fitted over the range $0.05 \leq F \leq 0.75$) may favour a lower value of σ_8 than previous analyses of the power spectrum alone (Viel et al. 2004a; McDonald et al. 2005), although this result should again be degenerate with the thermal state of the IGM. We instead find that the K07 PDF measurements on their own favour value of σ_8 towards the upper end of current observational constraints if the temperature–density relation is close to isothermal.

Lidz et al. (2006) compare synthetic spectra drawn from pseudo-hydrodynamical simulations to the M00 observational data, but instead use a different estimator for the flux PDF. This estimator is derived from data which is smoothed on two different scales, thereby having the advantage of rendering the observational data insensitive to the normalization *and* shape of the continuum, but also to structure on small scales in the spectra. This latter effect may be advantageous when using simulations which are not fully resolved, but it may also erase important information in the observational data, especially with regard to the thermal state of the IGM. Lidz et al. (2006) find their simulations favour a temperature–density relation with $\gamma \geq 1.32$ (2σ) at $z = 2.72$. However, given the approximate treatment of smoothing due to gas pressure Lidz et al. (2006) adopt in their simulations, as well as their different observational estimator for the PDF, the significance of this constraint on comparison to our results is difficult to judge.

Lastly, and as discussed previously, Becker et al. (2007) found that an inverted temperature–density relation may be required to adequately fit the flux PDFs measured from an independent set of high-resolution QSO spectra over a wide redshift range, $1.7 < z < 5.8$. Our results are also consistent with this possibility. Note, however, that Becker et al. (2007) use a semi-analytical model for the Ly α forest based on the IGM density distributions drawn from a hydrodynamical simulation. This model does not self-consistently model Jeans smoothing effects, and a direct comparison to their results is again not possible.

5 IMPLICATIONS FOR THE LY α FOREST FLUX POWER SPECTRUM

The Ly α flux PDF is of course not the only diagnostic of the thermal state of the IGM. There should be good prospects to break some of the degeneracies discussed here and further tighten constraints by a combined analysis with other Ly α flux statistics. Assumptions about the IGM thermal state are of particular relevance for mea-




Figure 6. The ratio of the Ly α flux power spectrum computed from synthetic spectra drawn from the models indicated on the panel at $z = 2.94$ to the flux power spectrum of the fiducial model (20-256, $\sigma_8 = 0.85$, $\gamma = 1.27$). The shaded region corresponds to the range of wavenumbers used by Viel et al. (2004b) to infer the amplitude and shape of the matter power spectrum from the Ly α flux power spectrum, $0.003 < k(\text{s km}^{-1}) < 0.03$. An inverted temperature–density relation increases the power within this wavenumber interval, mimicking the effect of a larger amplitude σ_8 of the underlying matter power spectrum.

surements of the matter power spectrum from the Ly α flux power spectrum. Such analyses have so far not allowed for the possibility of an inverted temperature–density relation and have marginalized over what was believed to be the plausible range for the slope, $1 < \gamma < 1.6$. Due to the temperature dependence of the recombination coefficient a smaller γ should translate into a steeper power law relation between neutral hydrogen density and matter density. For measurements of the matter power spectrum a smaller assumed γ should therefore result in a smaller amplitude for the matter power spectrum. In this section we therefore briefly discuss the implications of an inverted temperature–density relation for the Ly α forest flux power spectrum. Note this is a qualitative comparison only; a detailed joint analysis of the Ly α flux PDF and flux power spectrum with a larger suite of numerical simulations with inverted temperature–density relations will be required to examine this more accurately, which is beyond the scope of this paper.

Fig. 6 illustrates the effect of an inverted temperature–density relation on the flux power spectrum at $z = 2.94$. The solid curve shows the ratio of the flux power spectrum, $P_F(k)$, computed from the model with the inverted temperature–density relation to the fiducial model power spectrum, $P_{F,20-256}(k)$. The shaded region corresponds to the range of wavenumbers used by Viel et al. (2004b) to infer the matter power spectrum from the Ly α forest flux power spectrum. The flux power spectrum of the simulations with the inverted temperature–density relation clearly exhibits more power over this range on comparison to the fiducial model, and this increases noticeably towards larger wavenumbers. As discussed in Section 3.2, this is because a larger γ increases the temperature and hence the ionized fraction of the gas associated with the most prominent absorption features, thereby suppressing predominantly the small-scale power. The dotted curve shows the same comparison for the isothermal temperature–density relation model. The effect is again to increase power at all scales, although to a somewhat lesser extent. The dashed

and dot-dashed curves also show the comparison for two models with different matter power spectrum normalizations, $\sigma_8 = 0.75$ and 0.95 , respectively. As expected, a lower σ_8 suppresses power and a larger σ_8 increases power over the relevant wavenumber range. An inverted or isothermal temperature–density relation therefore mimics some of the effects a larger normalization for the matter power spectrum has on the Ly α forest flux power spectrum.

This has potentially important implications for inferring the matter power spectrum from the Ly α forest. The measurements of the matter power spectrum from Ly α forest data are affected by nonlinearities as well as by complex degeneracies between the thermal state of the IGM, the density distribution and the assumed effective optical depth. However, if we naively extrapolate the results of Viel et al. (2004b) to smaller values of γ we would expect a 5–10 per cent lower amplitude σ_8 for the slope of $\gamma \sim 0.5$ favoured by the analysis of the Ly α flux PDF presented here. A 5–10 per cent lower amplitude for the inferred matter power spectrum would largely alleviate the strain between the measurements from Ly α forest data and the WMAP 3-year results (Viel & Haehnelt 2006). Unfortunately it would also substantially weaken the strong upper limits on the neutrino mass from the joint analysis of Ly α forest and CMB data (Seljak et al. 2006a; Gratton, Lewis & Efstathiou 2007).

6 CONCLUSIONS

We have compared the improved measurements of the Ly α flux PDF at $1.7 < z < 3.2$ by K07 obtained from a full Voigt profile analysis of 18 high-resolution, high-S/N UVES/VLT spectra to a large set of hydrodynamical simulations of the IGM. Our main findings are as follows.

(i) The shape of the Ly α flux PDF depends only weakly on the fluctuation amplitude of the matter power spectrum characterized by σ_8 and on the normalization of the temperature–density relation for the low-density IGM, $T = T_0 \Delta^\gamma$ (Hui & Gnedin 1997; Valageas et al. 2002). A moderately increased scatter in the temperature–density relation, which may result from inhomogeneous radiative transfer effects during the epoch of He II reionization, also has little effect on the shape of the flux PDF. The shape of the Ly α flux PDF is mainly determined by the effective optical depth τ_{eff} and by the slope of the temperature–density relation for the low-density IGM. The extreme ends of the PDF at flux levels of zero and unity are particularly sensitive to the slope of the temperature–density relation.

(ii) A comparison of the full range of the Ly α flux PDF to hydrodynamical simulations favours an inverted temperature–density relation. The flux PDFs of our synthetic spectra with an inverted temperature–density relation, such that $\gamma = [0.67, 0.54, 0.44]$ at $\langle z \rangle = [2.07, 2.52, 2.94]$, provide an excellent fit to the observed metal-cleaned flux PDF measured by K07. This suggests that the voids in the IGM are significantly hotter than previously assumed.

(iii) If we restrict our analysis of the PDF to flux levels $0.1 \leq F \leq 0.8$, which are least affected by continuum fitting, noise, galactic winds and the uncertainties of how to simulate the relevant physics in dense regions of the IGM, the comparison of observed and synthetic spectra still favours a model where the temperature–density relation is either nearly isothermal or inverted.

(iv) Raising or lowering the assumed continuum level of our synthetic spectra in order to mimic the possibility of a systematic over or underestimate of the continuum level in the observed spectra rapidly worsens the quality of the fit to the K07 PDF for offsets larger than 1.5 per cent. This suggests that continuum fitting errors

for the high-quality data of K07 are under control at a level of 1–2 per cent, consistent with the estimates of the continuum uncertainties by K07.

(v) A comparison of our simulations with the full range of the Ly α flux PDF measured by M00 at $\langle z \rangle = 3$ from a smaller observational sample gives consistent results if the continuum on the synthetic spectra is lowered by 1.5 per cent. Using the currently preferred cosmological parameters, the M00 data also favour a nearly isothermal or inverted temperature–density relation.

(vi) An inverted temperature–density relation has important implications when inferring the underlying matter power spectrum from the Ly α forest flux distribution. The effect of an inverted temperature–density relation on the flux power spectrum is similar to the effect of a larger normalization, σ_8 , of the matter power spectrum. This could alleviate some of the existing strain between the measurements of σ_8 from Ly α forest data and the WMAP 3-year results (Viel & Haehnelt 2006).

The preference of the data for an inverted temperature–density relation would, if consolidated, be a rather profound result. It would suggest that the widely accepted paradigm of a tight IGM temperature–density relation at low densities, evolving from being almost isothermal following hydrogen reionization to a relation where the temperature rises with increasing density (Hui & Gnedin 1997; Theuns et al. 1998; Valageas et al. 2002), is too simplistic.

Perhaps the most plausible explanation for an inverted temperature–density relation is He II reionization. There is evidence to suggest this may end around $z \sim 3$ (Songaila 1998; Schaye et al. 2000; Theuns et al. 2002a; Shull et al. 2004; Bolton et al. 2006; Faucher-Giguere et al. 2007), and thus lies at the upper end of the redshift range considered in this work. Changes in the photoheating rate which occur as ionizing photons propagate through optically thick regions can significantly raise the temperature of the IGM (Abel & Haehnelt 1999). First attempts at modelling inhomogeneous He II reionization and/or the associated radiative transfer effects also suggest a more complex and possibly multivalued and inverted temperature–density relation is plausible (Bolton et al. 2004; Gleser et al. 2005; Furlanetto & Oh 2007; Tittley & Meiksin 2007). Therefore, while a well-defined, inverted temperature–density relation may not faithfully represent the true thermal state of the IGM, a model with $\gamma < 1$ should mimic some of the expected inhomogeneous radiative transfer induced effects on the IGM temperature following He II reionization. Within this framework, this model provides a good fit to the K07 observational data and suggests that the thermal state of the IGM may be significantly more complex at $z \simeq 3$ than previously thought. Further investigation of the thermal state of the IGM around $z \simeq 3$ is therefore highly desirable. Note, however, that studies using wavelets have failed to detect the expected fluctuations in the IGM temperature at these redshifts (Theuns et al. 2002d; Zaldarriaga 2002).

While radiative transfer effects and the inhomogeneity of He II reionization provide a plausible explanation for an increased temperature in underdense regions of the IGM, it will require full numerical radiative transfer simulations to confirm the validity of this picture. Helium reionization is driven by hard photons emitted by rare, luminous and most likely short-lived sources (QSOs) under conditions where photon mean free paths are significantly longer and recombination time-scales shorter (if compared to the Hubble time) than for hydrogen reionization at higher redshifts. Current 3D cosmological radiative transfer codes which incorporate multiple sources are designed primarily to investigate the topology of H I reionization, and are less reliable when modelling the resulting

thermal state of the IGM (Iliev et al. 2006a). Most detailed studies of radiative transfer effects on the IGM temperature therefore employ either reduced dimensionality (Bolton et al. 2004; Bolton & Haehnelt 2007a) or neglect the modelling of multiple sources (Maselli & Ferrara 2005; Tittley & Meiksin 2007). Furthermore, the physical scales involved during He II reionization are much larger than those during H I reionization at $z > 6$, placing even greater demands on the required computational resources (although see Sokasian, Abel & Henquist 2002; Paschos et al. 2007). Radiative transfer simulations which accurately track the temperature evolution are thus very challenging. Developments in this area will be key to interpreting the increasing amount of high-quality Ly α forest data.

Finally, we must also not lose sight of the possibility that other processes may contribute to the heating of the IGM (e.g. Madau & Efstathiou 1999; Nath, Sethi & Shchekinov 1999; Inoue & Kamaya 2003; Ricotti & Ostriker 2004; Samui, Subramanian & Srianand 2005). Our results may instead indicate that an alternative modification to the current models for the Ly α forest and the thermal state of the IGM is required.

ACKNOWLEDGMENTS

We thank Volker Springel for his advice and for making GADGET-2 available, and Benedetta Ciardi and Simon White for comments on the draft manuscript. We also thank the anonymous referee for a positive and helpful report. This research was conducted in cooperation with SGI/Intel utilizing the Altix 3700 supercomputer COSMOS at the Department of Applied Mathematics and Theoretical Physics in Cambridge. COSMOS is a UK-CCC facility which is supported by HEFCE and STFC/PPARC.

REFERENCES

- Abel T., Haehnelt M. G., 1999, *ApJ*, 520, L13
 Becker G. D., Rauch M., Sargent W. L. W., 2007, *ApJ*, 662, 72
 Bernardi M. et al. 2003, *AJ*, 125, 32
 Bertone S., White S. D. M., 2006, *MNRAS*, 367, 247
 Bi H., Davidsen A. F., 1997, *ApJ*, 479, 523
 Bolton J., Meiksin A., White M., 2004, *MNRAS*, 348, L43
 Bolton J. S., Haehnelt M. G., 2007a, *MNRAS*, 374, 493
 Bolton J. S., Haehnelt M. G., 2007b, *MNRAS*, 382, 325
 Bolton J. S., Haehnelt M. G., Viel M., Springel V., 2005, *MNRAS*, 357, 1178
 Bolton J. S., Haehnelt M. G., Viel M., Carswell R. F., 2006, *MNRAS*, 366, 1378
 Bruscoli M., Ferrara A., Marri S., Schneider R., Maselli A., Rollinde E., Aracil B., 2003, *MNRAS*, 343, L41
 Cen R., 1992, *ApJS*, 78, 341
 Choudhury T. R., Srianand R., Padmanabhan T., 2001, *ApJ*, 559, 29
 Croft R. A. C., Weinberg D. H., Pettini M., Hernquist L., Katz N., 1999, *ApJ*, 520, 1
 Croft R. A. C., Weinberg D. H., Bolte M., Burles S., Hernquist L., Katz N., Kirkman D., Tytler D., 2002, *ApJ*, 581, 20
 Davé R., Hernquist L., Katz N., Weinberg D. H., 1999, *ApJ*, 511, 521
 Desjacques V., Nusser A., 2005, *MNRAS*, 361, 1257
 Desjacques V., Nusser A., Sheth R. K., 2007, *MNRAS*, 374, 206
 Eisenstein D. J., Hu W., 1999, *ApJ*, 511, 5
 Faucher-Giguere C.-A., Prochaska J. X., Lidz A., Hernquist L., Zaldarriaga M., 2007, *ApJ*, preprint (arXiv:0709.2382)
 Furlanetto S., Oh S. P., 2007, *ApJ*, preprint (arXiv:0711.0751)
 Gaztañaga E., Croft R. A. C., 1999, *MNRAS*, 309, 885
 Gleser L., Nusser A., Benson A. J., Ohno H., Sugiyama N., 2005, *MNRAS*, 361, 1399
 Gratton S., Lewis A., Efstathiou G., 2007, preprint (arXiv:0705.3100)
 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*, preprint (arXiv:astro-ph/0106018)
 Haehnelt M. G., Steinmetz M., 1998, *MNRAS*, 298, L21
 Heap S. R., Williger G. M., Smette A., Hubeny I., Sahu M. S., Jenkins E. B., Tripp T. M., Winkler J. N., 2000, *ApJ*, 534, 69
 Hu E. M., Kim T.-S., Cowie L. L., Songaila A., Rauch M., 1995, *AJ*, 110, 1526
 Hui L., Gnedin N. Y., 1997, *MNRAS*, 292, 27
 Iliev I. T. et al., 2006a, *MNRAS*, 371, 1057
 Iliev I. T., Mellema G., Pen U.-L., Merz H., Shapiro P. R., Alvarez M. A., 2006b, *MNRAS*, 369, 1625
 Inoue A. K., Kamaya H., 2003, *MNRAS*, 341, L7
 Jakobsen P., Boksenberg A., Deharveng J. M., Greenfield P., Jedrzejewski R., Paresce F., 1994, *Nat*, 370, 35
 Jena T. et al., 2005, *MNRAS*, 361, 70
 Jenkins E. B., Ostriker J. P., 1991, *ApJ*, 376, 33
 Kim T.-S., Hu E. M., Cowie L. L., Songaila A., 1997, *AJ*, 114, 1
 Kim T.-S., Cristiani S., D'Odorico S., 2001, *A&A*, 373, 757
 Kim T.-S., Bolton J. S., Viel M., Haehnelt M. G., Carswell R. F., 2007, *MNRAS*, 382, 1657 (K07)
 Kirkman D., Tytler D., 1997, *ApJ*, 484, 672
 Kirkman D. et al., 2005, *MNRAS*, 360, 1373
 Lesgourgues J., Viel M., Haehnelt M. G., Massey R., 2007, *J. Cosmol. Astropart. Phys.*, 11, 8
 Lidz A., Heitmann K., Hui L., Habib S., Rauch M., Sargent W. L. W., 2006, *ApJ*, 638, 27
 Lu L., Sargent W. L. W., Womble D. S., Takada-Hidai M., 1996, *ApJ*, 472, 509
 Madau P., Efstathiou G., 1999, *ApJ*, 517, L9
 Maselli A., Ferrara A., 2005, *MNRAS*, 364, 1429
 McDonald P., Miralda-Escudé J., Rauch M., Sargent W. L. W., Barlow T. A., Cen R., Ostriker J. P., 2000, *ApJ*, 543, 1 (M00)
 McDonald P., Miralda-Escudé J., Rauch M., Sargent W. L. W., Barlow T. A., Cen R., 2001, *ApJ*, 562, 52
 McDonald P. et al., 2005, *ApJ*, 635, 761
 Meiksin A., 2000, *MNRAS*, 314, 566
 Meiksin A., Bryan G., Machacek M., 2001, *MNRAS*, 327, 296
 Miralda-Escudé J., Rees M. J., 1994, *MNRAS*, 266, 343
 Miralda-Escudé J., Cen R., Ostriker J. P., Rauch M., 1996, *ApJ*, 471, 582
 Miralda-Escudé J., Haehnelt M., Rees M. J., 2000, *ApJ*, 530, 1
 Nath B. B., Sethi S. K., Shchekinov Y., 1999, *MNRAS*, 303, 1
 Paschos P., Norman M. L., Bordner J. O., Harkness R., 2007, preprint (arXiv:0711.1904)
 Rauch M., 1998, *ARA&A*, 36, 267
 Rauch M. et al. 1997, *ApJ*, 489, 7
 Reimers D., Fechner C., Hagen H.-J., Jakobsen P., Tytler D., Kirkman D., 2005, *A&A*, 442, 63
 Ricotti M., Ostriker J. P., 2004, *MNRAS*, 352, 547
 Ricotti M., Gnedin N. Y., Shull J. M., 2000, *ApJ*, 534, 41
 Samui S., Subramanian K., Srianand R., 2005, in *International Cosmic Ray Conference*, p. 215, preprint (arXiv:astro-ph/0505590)
 Schaye J., Theuns T., Rauch M., Efstathiou G., Sargent W. L. W., 2000, *MNRAS*, 318, 817
 Schaye J., Aguirre A., Kim T., Theuns T., Rauch M., Sargent W. L. W., 2003, *ApJ*, 596, 768
 Seljak U., Makarov A., McDonald P., Trac H., 2006a, *Phys. Rev. Lett.*, 97, 191303
 Seljak U., Slosar A., McDonald P., 2006b, *J. Cosmol. Astropart. Phys.*, 10, 14
 Shull J. M., Tumlinson J., Giroux M. L., Kriss G. A., Reimers D., 2004, *ApJ*, 600, 570
 Sijacki D., Springel V., di Matteo T., Hernquist L., 2007, *MNRAS*, 380, 877
 Sokasian A., Abel T., Hernquist L., 2002, *MNRAS*, 332, 601
 Songaila A., 1998, *AJ*, 115, 2184
 Spergel, D. N. et al., 2007, *ApJS*, 170, 377
 Springel V., 2005, *MNRAS*, 364, 1105

- Springel V., Hernquist L., 2003a, MNRAS, 339, 289
 Springel V., Hernquist L., 2003b, MNRAS, 339, 312
 Tepper García T., 2006, MNRAS, 369, 2025
 Theuns T., Zaroubi S., 2000, MNRAS, 317, 989
 Theuns T., Leonard A., Efstathiou G., Pearce F. R., Thomas P. A., 1998, MNRAS, 301, 478
 Theuns T., Schaye J., Haehnelt M. G., 2000, MNRAS, 315, 600
 Theuns T., Bernardi M., Frieman J., Hewett P., Schaye J., Sheth R. K., Subbarao M., 2002a, ApJ, 574, L111
 Theuns T., Schaye J., Zaroubi S., Kim T., Tzanavaris P., Carswell B., 2002b, ApJ, 567, L103
 Theuns T., Viel M., Kay S., Schaye J., Carswell R. F., Tzanavaris P., 2002c, ApJ, 578, L5
 Theuns T., Zaroubi S., Kim T.-S., Tzanavaris P., Carswell R. F., 2002d, MNRAS, 332, 367
 Tittley E. R., Meiksin A., 2007, MNRAS, 380, 1369
 Tytler D. et al., 2004, ApJ, 617, 1
 Valageas, P., Schaeffer, R., Silk, J., 2002, A&A, 388, 741
 Viel M., Haehnelt M. G., 2006, MNRAS, 365, 231
 Viel M., Haehnelt M. G., Carswell R. F., Kim T.-S., 2004a, MNRAS, 349, L33
 Viel M., Haehnelt M. G., Springel V., 2004b, MNRAS, 354, 684
 Viel M., Matarrese S., Heavens A., Haehnelt M. G., Kim T.-S., Springel V., Hernquist L., 2004c, MNRAS, 347, L26
 Viel M., Haehnelt M. G., Lewis A., 2006, MNRAS, 370, L51
 Zaldarriaga M., 2002, ApJ, 564, 153
 Zaroubi S., Viel M., Nusser A., Haehnelt M., Kim T.-S., 2006, MNRAS, 369, 734

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