

The effects of ultraviolet background correlations on Ly α forest flux statistics

Avery Meiksin¹^{*} and Martin White²

¹*Institute for Astronomy, University of Edinburgh, Blackford Hill, Edinburgh EH9 3HJ*

²*Departments of Astronomy and Physics, University of California, Berkeley, CA 94720, USA*

Accepted 2004 February 10. Received 2004 February 10; in original form 2003 July 14

ABSTRACT

We examine the possible effects of ultraviolet (UV) background fluctuations on the pixel flux power spectrum and autocorrelation function of the Ly α forest due to a finite number of sources in an attenuating medium. We consider scenarios in which quasi-stellar object (QSO) sources dominate the contribution to the UV background. To estimate their contribution, we use the QSO luminosity functions from the Two-Degree Field and the Sloan Digital Sky Survey. We estimate self-consistent values for the attenuation length at the Lyman edge through the intergalactic medium using particle mesh simulations of the Ly α forest, normalized by the measured mean Ly α flux. It is necessary to add the contribution from Lyman-limit systems based on their measured statistical properties because the simulations are unable to reproduce the measured abundances of these systems, suggesting that their formation may involve more complex hydrodynamical processes than simple collapse into dark matter haloes. We examine the convergence properties of the flux power spectrum and autocorrelation function under differing assumptions regarding pressure smoothing. On the basis of our tests, we recommend that simulations used to estimate the flux power spectrum at high redshifts ($z > 3$) be performed in a box of comoving size at least $25 h^{-1}$ Mpc on a side and a comoving force mesh as fine as $30\text{--}60 h^{-1}$ kpc. The requirements become more severe at lower redshifts. Even with these specifications, we find that the flux autocorrelation function does not converge to better than 10 per cent on scales exceeding 3 per cent of the box size without added pressure smoothing. Much better convergence properties are obtained for the *relative* effects of the UV background fluctuations on the flux power spectrum. The background fluctuations increase the large-scale power and suppress the power at intermediate scales. At $z \leq 4$, the effect is only at the few per cent level. By $z = 5.5$, however, the large-scale power is boosted by more than 10 per cent, growing rapidly at higher redshifts, and suggesting that the flux power spectrum may serve as a useful tool for distinguishing QSO-dominated UV background scenarios from those in which more abundant sources like galaxies dominate.

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

1 INTRODUCTION

In large part as a result of insights gained from hydrodynamical simulations of structure formation in the Universe (Cen et al. 1994; Zhang, Anninos & Norman 1995; Hernquist et al. 1996; Zhang et al. 1997; Bond & Wadsley 1997; Theuns, Leonard & Efstathiou 1998), the Ly α forest, as measured in high-redshift quasi-stellar object (QSO) spectra, has emerged as one of our premier probes of the Universe at high redshift. The results of simulations assuming early homogeneous H I and He II reionization reproduce the

cumulative flux distributions and H I column density distributions measured from high-resolution spectra to an accuracy of a few per cent (Meiksin, Bryan & Machacek 2001). More problematic are the predicted widths of the absorption features, which appear to require additional sources of broadening, perhaps late He II reionization, to match the measured widths (Theuns et al. 1999; Bryan & Machacek 2000; Meiksin et al. 2001).

Parallel to the development of the full hydrodynamical simulations of the intergalactic medium (IGM) are pseudo-hydrodynamical schemes using pure gravity (Petitjean, Mücke & Kates 1995; Croft et al. 1998; Gnedin & Hui 1998; Meiksin & White 2001). These simpler treatments are based on two key results from the full hydrodynamical computations. As a consequence of photoionization

*E-mail: aam@roe.ac.uk

equilibrium: (1) the gas follows a tight temperature–density relation for the absorbing material with the neutral hydrogen temperature proportional to the baryon density to the power of ~ 0.5 (Zhang et al. 1998; Schaye et al. 1999); and (2) the neutral hydrogen density closely traces the total matter density on the scales relevant to the forest ($0.1\text{--}10 h^{-1}$ Mpc) (Zhang et al. 1998). The structure in QSO absorption thus tracks, in a calculable way, slight fluctuations in the matter density of the Universe back along the line of sight to the QSO, with most of the absorption by the Ly α forest arising from gas with overdensities of a few times the mean density. Although these techniques are not sufficiently accurate to be used for detailed modelling of the flux and absorption-line properties of the IGM (Meiksin & White 2001), they provide a useful means of investigating convergence questions and the possible role of additional physical effects on the Ly α forest without the substantial overhead of the full hydrodynamical treatment. In particular, in recent years they have been increasingly relied on for predicting the flux power spectrum of the Ly α forest (Meiksin & White 2001; Zaldarriaga, Hui & Tegmark 2001; Croft et al. 2002b; Seljak, McDonald & Makarov 2003).

While much work has been done on the physics of the Ly α forest and the cosmology that may be extracted, the role of fluctuations in the UV background on the absorption properties of the Ly α forest are largely unexplored (see Zuo 1992a,b; Fardal & Shull 1993; Croft et al. 1999, 2002a; Gnedin & Hamilton 2002; Croft 2004). The effects of fluctuations in the UV background on the statistical properties of the Ly α forest were considered previously by Zuo (1992a,b) and Zuo & Phinney (1993). Zuo & Bond (1994) introduced the use of the flux transmission autocorrelation function for measuring the clustering of the absorbing gas, although they did not consider the effect of fluctuations in the UV background apart from a local proximity effect. As we shall show, the effects of UV background correlations arising from photon attenuation in the IGM are most conspicuous only at high redshifts ($z > 5$), where the precise level of absorption by the Ly α forest is most difficult to measure because of the extremely low flux values, requiring very high signal-to-noise ratio spectra.

Fardal & Shull (1993) explored the effects of UV background fluctuations on the Ly α forest along the lines of Zuo, but using Monte Carlo realizations of randomly distributed sources and clouds to estimate the effect of the background fluctuations. Croft et al. (1999) estimated the flux power spectrum resulting from UV background fluctuations generated by randomly distributed QSO sources in a uniform IGM, and found this to be small, at the few per cent level, compared with the flux power spectrum from the Ly α forest at $z = 2.5$. They found that allowing for QSO clustering had little additional effect. Croft et al. (2002a) accounted for the UV background fluctuations at $z = 3$ assuming the IGM was ionized by the galaxies in their simulation, and concluded again that the effect on the flux power spectrum is small. Gnedin & Hamilton (2002) examined the effect of a fluctuating UV background at $z = 4$ using hydrodynamical simulations with pseudo-radiative transfer, assuming that reionization is accomplished by galaxies in their simulation. They concluded that the fluctuations of the UV background affect the flux power spectrum at the per cent level, but cause an underestimate of the photoionization rate by $\mathcal{O}(20)$ per cent. A comparable boost in the photoionization rate is required if QSOs are the dominant source of the UV background at $z \geq 5$ (Meiksin & White 2003, hereafter Paper I). Croft (2004) has recently examined the effect of finite QSO lifetimes on UV background fluctuations, finding that the fluctuations induce a weak suppression in the flux power spectrum at very long wavelengths.

In this paper, we investigate the effect of the fluctuations on the two-point flux distribution (the power spectrum and autocorrelation function) of the Ly α forest. This extends our earlier work in which we examined the effects of fluctuations in the UV background due to source discreteness in an attenuating medium on the one-point distribution function (Paper I). In the next section we describe the simulations used. In Section 3, we discuss our model for the UV background fluctuations. The results are presented in Section 4, and we summarize our conclusions in Section 5. In Appendix A we discuss the numerical convergence tests that we have performed, and we describe our estimates of the mean Ly α flux in Appendix B.

2 NUMERICAL SIMULATIONS

To investigate the effects of UV background fluctuations on the power spectrum of the Ly α forest, we have performed simulations of five different cosmological models within the Λ CDM family. In each case we used pure particle mesh (PM) dark matter simulations, mimicking the temperature of the gas using a polytropic equation of state and assuming that gas and dark matter have the same spatial distribution. This has been shown (Petitjean et al. 1995; Croft et al. 1998; Gnedin & Hui 1998; Meiksin & White 2001) to produce results comparable to the full hydrodynamical simulations at the level of 10 per cent. The agreement is expected to be even better on the large scales of interest to this work. The parameters for the simulations are provided in Table 1. The models are all for a flat universe and are consistent with Ly α forest constraints on matter fluctuations σ_J on the Jeans scale (Meiksin et al. 2001). Models 1–3 were run previously to the *WMAP* results and were constructed to be consistent with existing large-scale structure, cluster abundance and cosmic microwave background (CMB) constraints at the time. Model 4 was later added explicitly to be consistent with *WMAP* results as well (Spergel et al. 2003; Hinshaw et al. 2003; Kogut et al. 2003). Although models 1–3 are not the best matches to the recent *WMAP* data, they are close to models that are good matches on the scales relevant to the Ly α forest. We provide in Table 2 a list of similar models with their χ^2 agreement to the *WMAP* TT data. Model C is the concordance model of Ostriker & Steinhardt (1995). A description of the parallel PM code used is given in Paper I. In each case we used 512^3 particles and a 1024^3 force mesh, in a cubic box with (comoving) side length $25 h^{-1}$ Mpc, except for model C, for which the comoving box size was $30 h^{-1}$ Mpc. The corresponding (comoving) force mesh size is $24\text{--}29 h^{-1}$ kpc. The force resolution of ~ 3 mesh zones corresponds to comoving scales of $\sim 70\text{--}90 h^{-1}$ kpc, or about 10 per cent of the Jeans wavelength

Table 1. Parameters for the cosmological models. Ω_m is the total mass density parameter; Ω_ν is the vacuum energy density parameter; Ω_b is the baryonic mass fraction; $h = H_0/100 \text{ km s}^{-1} \text{ Mpc}^{-1}$, where H_0 is the Hubble constant at $z = 0$; n is the slope of the primordial density perturbation power spectrum; σ_8 is the fluctuation normalization at $z = 0$ in a sphere of radius $8 h^{-1}$ Mpc; and σ_J is the fluctuation normalization at $z = 3$ filtered on the isothermal Jeans scale for gas at $T = 2 \times 10^4$ K, given by $\lambda_J \approx 108 \Omega_m^{-1/2} h^{-1}$ kpc (Meiksin et al. 2001).

Model	Ω_m	Ω_ν	$\Omega_b h^2$	h	n	σ_8	σ_J
1	0.30	0.70	0.020	0.70	1.05	0.97	1.59
2	0.35	0.65	0.020	0.70	0.95	0.88	1.35
3	0.40	0.60	0.020	0.55	1.10	0.90	1.57
4	0.30	0.70	0.022	0.70	0.95	0.92	1.32
C	0.30	0.70	0.018	0.67	1.00	0.90	1.28

Table 2. Parameters for cosmological models consistent with *WMAP* data. In addition to the parameters defined in Table 1: τ is the Thomson optical depth to the last scattering surface; and the values of $\chi^2 \approx -2 \log L$, where L is the likelihood, are from fits to the 899 data points in the *WMAP* TT angular power spectrum.

Ω_m	Ω_v	$\Omega_b h^2$	h	n	τ	σ_8	σ_7	χ^2
0.30	0.70	0.020	0.70	0.95	0.10	0.912	1.31	1020.5
0.30	0.70	0.020	0.70	0.95	0.15	0.958	1.38	1030.7
0.30	0.70	0.020	0.70	1.00	0.10	0.954	1.46	1040.8
0.30	0.70	0.020	0.70	1.00	0.15	1.000	1.54	1041.0
0.30	0.70	0.022	0.70	0.95	0.15	0.920	1.32	990.1
0.30	0.70	0.022	0.70	1.00	0.10	0.917	1.40	992.7
0.30	0.70	0.022	0.70	1.00	0.15	0.961	1.46	990.7
0.30	0.70	0.024	0.70	1.00	0.10	0.881	1.33	984.0
0.30	0.70	0.024	0.70	1.00	0.15	0.921	1.39	979.7
0.30	0.70	0.020	0.75	0.95	0.10	1.068	1.59	1212.1
0.35	0.65	0.020	0.65	0.95	0.10	0.901	1.32	1001.2
0.35	0.65	0.020	0.65	0.95	0.15	0.942	1.38	1008.2
0.35	0.65	0.020	0.65	1.00	0.10	0.940	1.47	1042.3
0.35	0.65	0.020	0.65	1.00	0.15	0.982	1.54	1039.1
0.35	0.65	0.022	0.65	0.95	0.15	0.907	1.32	986.2
0.35	0.65	0.022	0.65	1.00	0.10	0.903	1.40	1013.8
0.35	0.65	0.022	0.65	1.00	0.15	0.946	1.47	1008.1
0.35	0.65	0.024	0.65	1.00	0.10	0.866	1.34	1025.5
0.35	0.65	0.024	0.65	1.00	0.15	0.907	1.40	1017.6
0.35	0.65	0.024	0.70	0.95	0.10	0.987	1.49	1022.0

($2\pi/k_j$) at $z = 3$ for gas with $T = 2 \times 10^4$ K. Convergence tests, described in Appendix A, suggest that these scales are adequate for our purposes.

Given a set of final particle positions and velocities, we compute the spectra as follows. First, the density and density-weighted line-of-sight velocity are computed on a grid (using CIC interpolation) and smoothed using fast Fourier transform (FFT) techniques. We try three different smoothing schemes, as described below and in Appendix A. The default scheme simply smoothes with a Gaussian of width one grid cell. This forms the fundamental data set. A regular grid of 32×32 sightlines is drawn through the box, parallel to the box sides. Along each sightline we integrate (in real space) to find $\tau(u)$ at a given velocity u . Specifically, we define

$$\tau(u) = \int dx A(x) \left[\frac{\rho(x)}{\bar{\rho}} \right]^2 T(x)^{-0.7} b^{-1} \exp \left[-\frac{(u - u_0)^2}{b^2} \right], \quad (1)$$

where $u_0 = xaHL_{\text{box}} + v_{\text{los}}$ and $b = \sqrt{2k_B T / m_H}$ is the Doppler parameter, where k_B is the Boltzmann constant and m_H is the mass of a hydrogen atom. The flux at velocity u is $\exp(-\tau)$. The integration variable x indicates the distance along the box in terms of the expansion velocity across the box. We apply periodic boundary conditions in performing the integration. The spectra are computed for the range of velocities across the box. In evaluating equation (1) we assume a power-law relation between the gas density and temperature

$$T \equiv T_0 \left(\frac{\rho}{\bar{\rho}} \right)^{\gamma-1}. \quad (2)$$

In practice, T_0 and γ will depend on the reionization history of the IGM, but we shall fix $T_0 = 2 \times 10^4$ K and $\gamma = 1.5$ throughout as described in Paper I.

In terms of the baryon density parameter Ω_b , hydrogen baryonic mass fraction X , metagalactic photoionization rate Γ_{-12} , in units of

Table 3. Some of the output times for the simulations. For each redshift, the mean Ly α flux, $\langle \exp(-\tau) \rangle$, the required H I photoionization rate Γ_{-12} , in units of 10^{-12} s^{-1} , to recover the mean flux assuming a homogeneous UV background, the attenuation length, r_0 , in comoving h^{-1} Mpc, and the mean number of sources per attenuation volume, N_0 , are listed for model 4. The last column, N_{box} , shows the number of sources in the simulation volume of side $25 h^{-1}$ Mpc (comoving). The two values at $z = 4.0$ correspond to two disparate estimates of the mean Ly α flux.

z	$\langle \exp(-\tau) \rangle$	Γ_{-12}	r_0	N_0	N_{box}
6.0	<0.006	<0.14	<23.0	3.9	1.2
5.5	$0.079^{+0.017}_{-0.013}$	$0.37^{+0.06}_{-0.05}$	77^{+13}_{-9}	130	1.0
5.0	$0.12^{+0.03}_{-0.04}$	$0.31^{+0.07}_{-0.09}$	85^{+17}_{-22}	200	1.2
4.0	0.36 ± 0.03	$0.43^{+0.06}_{-0.05}$	80 ± 5	240	1.7
4.0	0.47 ± 0.03	$0.76^{+0.12}_{-0.11}$	104 ± 6	510	1.7
3.89	0.48 ± 0.02	$0.68^{+0.08}_{-0.07}$	106 ± 5	560	1.8
3.0	0.70 ± 0.02	$0.88^{+0.14}_{-0.12}$	211 ± 11	5900	2.4
2.75	0.74 ± 0.04	$0.86^{+0.36}_{-0.24}$	250 ± 28	1.1×10^4	2.6

10^{-12} s^{-1} , and number f_e of electrons per hydrogen nucleus, the function $A(x)$ is given by

$$A(x) = 18.0 f_e \left(\frac{X}{0.76} \right) \left(\frac{\Omega_b h^2}{0.02} \right) \Gamma_{-12}^{-1}(x) L_{\text{box}} (1+z)^5, \quad (3)$$

where $\Gamma_{-12}(x)$ is evaluated at position x , L_{box} is in Mpc and b in equation (1) is in km s^{-1} . We normalize the spectra according to measurements of the mean absorbed flux $\langle \exp(-\tau) \rangle$ as reported in Fan et al. (2002, their fig. 1) for $z \geq 4$. We discuss the values adopted at other redshifts in Appendix B. Values for $\langle \exp(-\tau) \rangle$ at each of the output times of the simulation are provided in Table 3.

In order to compute the flux power spectra, it is necessary to filter the simulation data to reduce the effects of particle noise at the grid level. The minimal amount of filtering we consider is Gaussian smoothing with a width $\sigma = \Delta x$, where Δx is the width of a mesh cell. While this is appropriate for computing the power spectrum of the dark matter, it neglects any ‘natural’ filtering arising from pressure forces acting on the baryons, although it does include the effects of Doppler broadening on the spectra due to the thermal and peculiar velocities of the gas. If the effects of Doppler broadening exceed those of the pressure forces on the baryons, then the neglect of the pressure forces is justified. The effects of pressure forces on the flux power spectrum are most reliably predicted by performing full hydrodynamical simulations at both high spatial resolution and in a box of sufficient size to ensure convergence. Such calculations, however, are still prohibitively expensive to repeat for a range of models, and also because the origin of the line broadening is not fully understood (Theuns et al. 1999; Bryan & Machacek 2000; Meiksin et al. 2001), requiring additional modelling that affects both the shape and amplitude of the predicted flux power spectrum. As we are primarily interested in the *relative* effect of fluctuations in the UV background, we do not consider an exhaustive range of smoothing procedures, but only consider a few straightforward ones to judge how important smoothing is to our final results.

Following Gnedin & Hui (1998), in addition to the minimal smoothing models, we create models with additional smoothing on scales of the order of the Jeans length. Specifically we compute a comoving ‘Jeans scale’, $k_j = (3/2)^{1/2} c_s^{-1} H(z)/(1+z)$, where c_s is the sound speed using the assumed temperature and equation of state of the gas, $c_s = (\gamma k_B T / \bar{m})^{1/2}$, where $\bar{m} = 0.588 m_H$ is the mean mass per particle for fully ionized gas of cosmic abundances.

Gnedin & Hui argue for using a larger filter wavenumber $k_f = \eta k_j$, where k_f is determined from linear perturbation theory under differing assumptions for the evolution of the IGM temperature. There is thus some uncertainty in what value to use for η . If reionization was early, e.g. $z_{\text{ri}} > 18$, and assuming the IGM temperature remained isothermal thereafter, $\eta \sim 2$ for $3 < z < 5$. Assuming the post-reionization temperature decays linearly with the expansion factor, however, gives $\eta = 1.3\text{--}1.6$ over this redshift range. If reionization was late (or perhaps renewed), at e.g. $z_{\text{ri}} = 7$, then $\eta = 2.2\text{--}5.0$ for these two cases over $3 < z < 5$. Attempts to measure the temperature of the IGM from spectra of the Ly α forest, however, indicate a more complex temperature evolution, with, on average, an increasing temperature with expansion factor (Schaye et al. 2000). Assuming $T \sim (1+z)^{-1}$ gives, over the redshift range $3 < z < 5$, $\eta \lesssim 3$ for early reionization and $3 \lesssim \eta \lesssim 5$ for late reionization (see Appendix A). We consider the cases $\eta = 2$ and $\eta = 4$ here, noting that k_f is rather poorly determined from observations because of degeneracy with the effective dark matter spectral index, the IGM temperature and its history (Zaldarriaga et al. 2001).

We filter on the scale k_f in place of our one mesh cell filtering above. We consider two different types of smoothing: a Gaussian of width $\sigma = 2^{1/2}/k_f$, and a form motivated by linear perturbation theory, viz. $[1 + (k/k_f)^2]^{-1}$. When the Jeans filtering corresponds to filtering on scales smaller than a mesh cell, we default to the minimal smoothing case above (see Appendix A). We emphasize again that all of these approaches, like the pseudo-hydrodynamical approach employed by McDonald (2003), are ad hoc and can only be justified by comparison to true hydrodynamic simulations, but modify the power mostly on scales where thermal broadening is also very important. Unless stated otherwise, the following discussion refers to models computed with minimal Gaussian smoothing on the mesh scale.

3 UV BACKGROUND

3.1 UV intergalactic attenuation

For the uniform background case, normalizing the simulation results by the mean Ly α flux also fixes the attenuation properties of the IGM since the ionization structure of the IGM has been fixed by the normalization. This relation was exploited in Paper I to place constraints on the attenuation length r_0 and mean emissivity required to recover the measured values of the mean Ly α flux for $z > 4$. At lower redshifts, the simulation results are no longer sufficient for determining the Lyman-limit optical depth of the IGM. We find that the simulations fail to reproduce the measured number of Lyman-limit systems (LLSs) by a factor of several. A similar conclusion was reached by Gardner et al. (1997) on the basis of lower-resolution simulations than those we present. There, the authors attempted to correct for the missing LLSs by estimating the contribution from unresolved haloes on the basis of Press–Schechter theory. They found that they still fell well short of the observed number. Our simulations resolve haloes down to masses of about $10^8\text{--}10^9 M_\odot$, corresponding to circular velocities smaller than $20\text{--}25 \text{ km s}^{-1}$ at $z = 3$. The baryons in smaller-mass haloes would be heated to temperatures too high to remain bound after photoionization (Meiksin 1994), so that smaller-mass haloes would seem unlikely candidates for LLSs. We again find too few LLSs. This suggests that the origin of these systems may involve more complex hydrodynamical processes than simple collapse into dark matter haloes, such as Jeans fragmentation in the haloes (Meiksin 1994), ram-pressure stripping or galactic winds.

In order to estimate the attenuation length, we adopt the measured abundance of LLSs from Stengler-Larrea et al. (1995) for $\tau_L > 1$, $dN/dz = N_0(1+z)^\gamma$, with $N_0 = 0.25_{+0.17}^{-0.10}$ and $\gamma = 1.50 \pm 0.39$ for $0.3 \lesssim z \lesssim 4$. The contribution of the LLSs to the effective optical depth of the IGM at the Lyman edge is given by (Zuo 1992b)

$$\tau_{\text{eff}}^L = \int_{z_{\text{obs}}}^{z_{\text{em}}} dz \int_1^\infty d\tau_L \frac{\partial^2 N}{\partial z \partial \tau_L} \left\{ 1 - \exp \left[-\tau_L \left(\frac{1+z_{\text{obs}}}{1+z} \right)^3 \right] \right\}. \quad (4)$$

Here, the optical depth corresponds to the path between the source at z_{em} and the position of observation at z_{obs} . Because of the difficulty in measuring the high optical depths involved, the actual distribution of individual Lyman-limit optical depths is less certain. Adopting a power-law distribution $\partial^2 / \partial z \partial \tau_L = A(1+z)^\gamma \tau_L^{-\beta}$, and assuming a flat universe so that $dz/dr_p \approx (H_0 \Omega_m^{1/2}/c)(1+z)^{5/2}$ where r_p is proper length, we obtain

$$\frac{d\tau_{\text{eff}}^L}{dr_p} \approx \frac{H_0 \Omega_m^{1/2}}{c} N_0 (1+z)^{5/2+\gamma} [1 - e^{-1} + \Gamma(2-\beta, 1)], \quad (5)$$

where $\Gamma(a, x)$ is the incomplete Γ function. Adopting $\Omega_m = 0.3$, $h = 0.7$, $N_0 = 0.25$, $\gamma = 1.50$, we obtain $d\tau_{\text{eff}}^L/dr_p \approx 3.2 \times 10^{-5} T(1+z)^4$, where $T = 0.85$ for $\beta = 2$ and 0.91 for $\beta = 1.5$. In the limit of an infinitely steep distribution, so that all systems have $\tau_L = 1$, $T = 0.63$.

To estimate the total Lyman-limit optical depth through the IGM, we add the LLS contribution to the values we find from a given simulation for $z \leq 4$. Current measurements do not permit an estimate of the LLS contribution at higher redshifts. Given the uncertain origin of these systems, we conservatively do not include them in our estimates for $z > 4$. Doing so would result in a decrease of the attenuation length. We also assume the proper source emissivity does not evolve in our estimate of the attenuation length. Over the range of measured optical depths, this assumption only weakly affects our results.

In Fig. 1, we show the redshift evolution of $d\tau_{\text{eff}}^L/dr_p$ from the simulation results for model 4, along with the range in the LLS contribution based on observations. We also show the evolution of the attenuation length, defined by $r_0 = (d\tau_{\text{eff}}^L/dr_p)^{-1}$, and compare our estimate with previous estimates based on direct Ly α forest line counts from Meiksin & Madau (1993) and Haardt & Madau (1996). We have included the LLS contribution to the attenuation lengths for $z \leq 4$. We also show the effect that including the LLS contribution, if extrapolated to $z \geq 5$, would have on r_0 by the open points at $z = 5$ and 5.5 . Our estimates for the attenuation length lie between the medium- and low-attenuation models of Meiksin & Madau (1993), and lie systematically above the estimate of Haardt & Madau (1996). The redshift dependence of the (proper) attenuation lengths we find is fitted by $r_0 \approx 6800 h^{-1} \text{ Mpc} (1+z)^{-3.5}$ over $2.75 \leq z \leq 5.5$ to ~ 20 per cent accuracy. The attenuation lengths allowing for an extrapolation of the LLS contribution to $z > 4$ are fitted over the same redshift range by $r_0 \approx 1.7 \times 10^4 h^{-1} \text{ Mpc} (1+z)^{-4.2}$ to ~ 10 per cent accuracy. We caution against extending the fits outside the specified redshift range. In particular, we find that the attenuation length decreases precipitously between $z = 5.5$ and 6 , a consequence of the observed sharp rise in the Ly α optical depth at these redshifts (Fan et al. 2002).

3.2 Source contribution

The QSO luminosity function has a slope at the bright end varying as $\phi(L) \propto L^{-\beta_1}$, where $\beta_1 = 3.41$ according to the Two-Degree Field (2dF) QSO sample for $z < 2.3$ (Boyle et al. 2000), and $\beta_1 = 2.58$




Figure 1. The (proper) gradient $d\tau_L/dr_p$ of the Lyman-limit optical depth (upper panel) and the (proper) attenuation length (lower panel) for model 4. The points are estimated from the simulation results using the observational constraints on the mean Ly α flux. The contribution from Lyman-limit systems is included in the values of the attenuation length for $z \leq 4$. The effect that Lyman-limit systems would have on the attenuation length at $z = 5$ and 5.5 is shown by the open points. The open square at $z = 4$ corresponds to an alternative lower estimate for the average absorbed flux (see text). The two solid curves in the upper panel show the upper and lower limits for the measured contribution from Lyman-limit systems, with the dashed extensions being extrapolations based on the lower-redshift measurements. The two dashed lines in the lower panel are based on the low- and high-attenuation models of Meiksin & Madau (1993), and the solid line is based on their medium-attenuation model. The dotted line is based on the attenuation estimate of Haardt & Madau (1996).

according to the Sloan Digital Sky Survey (SDSS) QSO sample for $z > 3.5$ (Fan et al. 2001), although $\beta_1 = 3.2$ is still permitted by the data. We use the procedure in Paper I for adjusting the parameters of the Boyle, Shanks & Peterson (1988) luminosity function,

$$\Phi(L) = \frac{\Phi^*}{[(L/L^*)^{\beta_1} + (L/L^*)^{\beta_2}]}, \quad (6)$$

viz. by fixing $\beta_2 = 1.58$, in accordance with Boyle et al. (2000), and matching the SDSS counts for $M_{1450} < -25.7$ for $z < 6$ and for $M_{1450} < -27.1$ at $z = 6$ (and allowing for differences in the assumed values of the Hubble constant). We show the predicted contribution of QSOs to Γ_{-12} in Fig. 2 for $\beta_1 = 3.2$, and compare these with the required rates from Table 3. We have followed the procedure of Paper I for estimating Γ_{-12} from the QSO emissivity, including hardening of the metagalactic radiation field on filtering through the IGM (Haardt & Madau 1996). Specifically, for $z \leq 3$ we assume a metagalactic spectral index $\alpha_{MG} = 0$, while for $3.89 \leq z \leq 5.5$ we take $\alpha_{MG} = -1.3$, except at $z = 5.0$ ($\alpha_{MG} = -0.5$) and $z = 6$ ($\alpha_{MG} = 0$). These values were chosen to recover the required photoionization rates within 2σ , though usually the addition of some diffuse radiation is required. The value of α_{MG} is uncertain and can only be computed fully self-consistently by solving the detailed radiative transfer problem through the high-redshift IGM, but we do not expect it to deviate much from the values we use. In any case, the spectral index only weakly affects the overall ionization rate after integrating over the photoelectric cross-




Figure 2. A comparison between the estimated H I metagalactic photoionization rate Γ_{-12} from QSO sources and the rate required to match the measured mean Ly α flux, as estimated for model 4. The filled data points with error bars are from Table 3. The open data point at $z = 4$ is for an alternative lower estimate of the mean Ly α flux (see text). The estimated rates for QSO sources, shown by the filled triangles, assume a QSO luminosity function with a steep-end slope of $\beta_1 = 3.2$. The inverted filled triangles include a 40 per cent boost, nearly the maximum possible contribution from diffuse recombination radiation emitted by the IGM. The open triangles and inverted triangles are analogous estimates assuming $\beta_1 = 3.41$. The filled square at $z = 4$ is for $\beta_1 = 3.2$ and the alternative lower mean Ly α flux considered. The open square includes a 40 per cent boost in this estimate due to diffuse radiation from the IGM. Also shown are the estimates from Haardt & Madau (1996) based on earlier QSO source counts, assuming an intrinsic QSO spectral index shortward of the Lyman edge of $\alpha_Q = 1.5$ (solid line) and $\alpha_Q = 1.8$ (dashed line), a larger amount of intergalactic attenuation at the Lyman edge, and including the contribution of diffuse recombination radiation from the IGM.

section. Significant diffuse radiation resulting from reradiation by the IGM through radiative recombinations is expected at these high redshifts (Meiksin & Madau 1993; Haardt & Madau 1996), with as much as a $(\alpha_A - \alpha_B)/\alpha_A \approx 40\text{--}50$ per cent boost in the total metagalactic photoionization rate over the direct contribution from the QSO sources. We show the enhanced rates for a 40 per cent boost in Fig. 2.

The QSO contribution to the metagalactic photoionization rate lies systematically below the required values at $z = 3\text{--}4$. Given the uncertainties in the required rates, as well as in the QSO counts at these redshifts, it is unclear how seriously the discrepancy should be taken. Based on a different estimate of the QSO counts, Haardt & Madau (1996) compute a photoionization rate that well matches the required rate at $z \approx 3$ (see Fig. 2), although their value lies much below the required rate by $z = 4$. [A required slower rate of decline at $z \gtrsim 3$ than predicted by Haardt & Madau was also found by Meiksin et al. (2001).] A boost in the IGM temperature at these redshifts, as expected for late He II reionization (Madau & Meiksin 1994; Reimers et al. 1997; Meiksin et al. 2001), would also help ease the discrepancy (because of the resulting reduced hydrogen recombination rate). Finally, if the luminosity from QSO sources is beamed, the estimated UV contribution would increase like the ratio of the UV beam solid angle to the optical beam solid angle, since the QSO counts are based on rest-frame optical or near-UV (longward of the Lyman edge) surveys. (In principle, this could also result in a decrease in the estimated QSO contribution to the UV background if the ratio were less than unity.) None the less, it is interesting in

this context that Steidel, Pettini & Adelberger (2001) report a large Lyman continuum emissivity from Lyman-break galaxies at $z \simeq 3.4$. Adjusting their value down by the factor $\Omega_m^{1/2}$ for model 4, the predicted photoionization rate at $z = 3$, assuming a non-evolving proper emissivity, is $\Gamma_{-12} = 4.4$, well in excess of the required value in Table 3. Within the context of our models, such a high rate would so reduce the neutral fraction of the H I that the IGM would become significantly more transparent to Ly α photons than measured. To match our estimated 3σ upper limit on the photoionization rate would require the Lyman continuum emitting phase of the Lyman-break galaxy population to last only over a redshift interval $\Delta z \approx 0.1$. Thus the required excess photoionization above the QSO contribution could be accounted for by the Lyman-break galaxies, but only if they are a bursting population at these redshifts. Based on population synthesis models of the galaxies, Ferguson, Dickinson & Papovich (2002) argue this indeed may be the case.

3.3 UV background fluctuations

We compute the autocorrelation function of the QSO contribution to the UV background, using the formalism of Zuo (1992b) (also see Paper I), for both $\beta_1 = 2.58$ and 3.2. The results are shown in Fig. 3. Formally the correlation function diverges at zero separation, since the variance in the radiation field diverges (Paper I; Zuo 1992a). The typical width of an absorption feature is 25 km s^{-1} , so, when the correlations become small on this scale, we expect them to have a negligible influence on the flux autocorrelations or power spectra. We find that, for $z > 3$, the correlations still exceed unity on this scale, so that we may expect the effect of UV background fluctuations to imprint a signature on the flux statistics of the Ly α forest. We note that, for $\beta_1 = 2.58$, additional sources (like galaxies) must be introduced to recover the required H I photoionization rate (Paper I),



Figure 3. The autocorrelation function in the UV background predicted from QSO sources with $\beta_1 = 2.58$ (upper panel) and $\beta = 3.20$ (lower panel). The solid curves, from top to bottom, correspond to the redshifts $z = 6, 5.5, 5, 4, 3.89, 3$ and 2.75 . The $z = 4$ solid curve is computed for a (comoving) attenuation length of $r_0 = 104 h^{-1} \text{ Mpc}$ and nearly coincides with the $z = 3.89$ result. Also shown is the result at $z = 4$ assuming $r_0 = 80 h^{-1} \text{ Mpc}$ (dashed line), which nearly coincides with the $z = 5$ result.

but that these sources would only dilute the correlations by a factor of 2–3.

The effect of a discrete source on the Ly α flux is shown in Fig. 4. Here we artificially suppose that the entire UV background flux comes from a single source located in the centre of the simulation box. For a line of sight that passes slightly off-centre the $A(x)$ and spectrum are as shown, along with the spectrum that would have resulted along this sightline if the UV background were completely uniform.

To model the fluctuations in the UV background from a ‘realistic’ population of discrete sources, we resort to a Monte Carlo procedure. We create random source lists, with luminosities drawn from the QSO luminosity function above. A random number of QSOs is chosen consistent with the luminosity density of QSOs with absolute magnitude at rest frame 1450\AA $M_{1450} < -18.5$. The luminosity of each QSO is drawn at random from the luminosity function, and the sources are distributed at random positions in a volume larger than the size of the simulation box. We choose the box size to be sufficiently large to include the full effects of attenuation but not so large that redshift or evolution effects become important, so that the box is not larger than the effective size over which sources contribute to the local UV intensity (no larger than l_U in Paper I).

The frequency-specific flux at the Lyman edge is then calculated at each position on the simulation grid from

$$f_L = \sum_i L_i \frac{\exp(-r_i/r_0)}{4\pi r_i^2}, \quad (7)$$

where the sum is over the sources, the specific luminosity of source i at the Lyman edge is L_i , and r_i is the distance of the i th source from the fiducial point in the simulation volume. To convert from the QSO luminosity at 1450\AA to the Lyman edge, we adopt a QSO spectral shape of the form $f_\nu \sim \nu^{-0.99}$ for $1050\text{--}1450\text{\AA}$ and $f_\nu \sim \nu^{-1.8}$ for $912\text{--}1050\text{\AA}$ (Zheng et al. 1997). Since our simulation box and the relevant attenuation lengths are small, we do not keep track of source lifetimes or light-cone effects (see Croft 2004). To compute the total rate Γ_{-12} of ionizing photons (in units of 10^{-12} s^{-1}) from $f_L(x)$, we adjust the spectral index of the metagalactic background to reproduce the target values of $\langle \exp(-\tau) \rangle$ in Table 3. As in the uniform UV background case, for $z \geq 3.89$ the required values are generally negative ($-1.3 \lesssim \alpha_{\text{MG}} \lesssim -0.5$), consistent with the expected hardening of the radiation field near the Lyman edge on being filtered through the IGM. Other sources, like galaxies, may of course contribute as well. Adjusting α_{MG} to reproduce the required ionization rate provides an upper limit to the possible effects of UV background correlations resulting from attenuation; the introduction of additional sources would tend to dilute these effects. We also consider cases with a diffuse radiation field, as is expected to arise from radiative recombinations within the IGM.

Once we have Γ_{-12} throughout the simulation volume, we compute $A(x)$. We ignore the fluctuations in the IGM temperature that would result, noting that this is another possible contribution to the flux power spectrum, but one again that requires detailed modelling of the reionization process. Since our simulation volume is small, we explicitly checked that imposing periodic boundary conditions on the flux distribution modified our results by much less than 1 per cent.

We run sufficient realizations of this procedure to ensure convergence (one to several hundred, depending on redshift), keeping track of the UV background and flux correlation functions and the flux power spectrum. This enables us to compute both the mean spectra and their correlations.




Figure 4. A simulated spectrum from a ‘typical’ sightline at $z = 5$ from model 1. The upper panel shows the spectrum with a uniform ionizing background ($A(x) \equiv \text{constant}$, dashed) and where the UV background comes from a single source situated at the centre of the box (dotted). The lower panel shows $0.1 \times \rho/\bar{\rho}$ (dotted) and $A(x)/\bar{A}$ (dashed) for the single source case.

4 RESULTS

4.1 Pixel flux distribution

The cumulative pixel flux distributions at $z = 5$ are shown in Fig. 5. The slopes of the distributions are determined essentially by the value of σ_J , and show a systematic trend of flattening with increasing σ_J . The same trend was found by Meiksin et al. (2001) at $z = 3$, who exploited it to place stringent constraints on σ_J from high spectral resolution measurements of the Ly α forest. These results corroborate the trend found there. For a single high-resolution spectrum of typically $N_{\text{pix}} = 10^4$ pixels, it is possible to distinguish the cumulative distributions only to differences of about $1/N_{\text{pix}}^{1/2} = 0.01$, or somewhat larger allowing for pixel flux correlations (Meiksin et al. 2001), so that, while models 1 and 3 could be distinguished from model 4, model 2 could not on the basis of a single spectrum.

The effect of fluctuations in the UV background on the cumulative flux distributions for model 4 at $z = 4, 5, 5.5$ and 6 is shown in Fig. 6. As found in Paper I, the UV background fluctuations produce only a small distortion on the pixel flux distribution. The distortions found here are somewhat larger than found in Paper I, because here we now also account for correlations in the UV background. As found in Paper I, the effects of the UV background fluctuations are generally smaller than or comparable to the uncertainty between models with differing values of σ_J . Only at the smallest flux values (< 0.1) would it be possible to distinguish the effects of the fluctuations from an uncertainty in the model if the model is known with sufficient accuracy. A comparable level of uncertainty is introduced by the uncertain equation of state of the gas (Paper I), but this may

be removed by adjusting the temperature to match the measured distributions of Doppler parameters or wavelet coefficients (Meiksin 2000; Meiksin et al. 2001).

As in Paper I, we find that the demands on the mean photoionization rate increase in the presence of UV background fluctuations. Allowing for the correlations in the UV background, however, reduces the effect, and it actually inverts by $z = 6$. Specifically, for $z \leq 5.5$ we find that the mean photoionization rate must increase by 5–10 per cent when the (correlated) UV background fluctuations are included in the models. The amount of the increase depends on the fractional contribution arising from a diffuse radiation field and on the redshift. At $z = 6$, the required mean photoionization rate is instead reduced by 10–20 per cent when the now strongly correlated UV background fluctuations are included.

4.2 Pixel flux power spectrum

We compute the flux power spectrum at wavenumber k for the spectra according to

$$P_F(k) = \langle |\widehat{\delta f}|^2 \rangle / \langle f \rangle^2, \quad (8)$$

where $\widehat{\delta f}$ denotes the Fourier transform of $\delta f = f - \langle f \rangle$. In order to assess the accuracy of our predicted flux power spectra, we compare the results for models 1 and 4 and the concordance model with the measurements of McDonald et al. (2000) at $z = 3.89, 3.00$ and 2.41, using spectra normalized to the identical mean Ly α flux values they report. The results are shown in Fig. 7. The same models are also compared with the data of Croft et al. (2002b) at $z = 2.47, 2.72$




Figure 5. (Top panel) The cumulative distributions of the pixel flux at $z = 5$ for the four models assuming no UV background fluctuations. The models are labelled as model 1 (dotted), model 2 (short-dashed), model 3 (dot-short-dashed) and model 4 (long-dashed). The distributions group according to the value of σ_J , with models 1 and 3 nearly overlapping and models 2 and 4 nearly overlapping, and flatten as σ_J increases. (Bottom panel) The differences between the cumulative flux distributions for models 1–3 and the distribution for model 4 assuming no UV background fluctuations at $z = 5$. (The line types correspond to those in the top panel.)




Figure 6. The differences between the cumulative flux distributions for model 4 with and without UV background fluctuations, at $z = 4$ (short-long-dashed), 5 (short-dashed), 5.5 (long-dashed) and 6 (solid). Also shown is the difference at $z = 6$ allowing for a diffuse component added to the fluctuating UV background (dot-short-dashed).

and 3.03 (using the simulation results at $z = 2.41, 2.75$ and 3.00). Since Croft et al. report no mean Ly α flux measurement for their spectra, we continue to adopt the values from McDonald et al. at $z = 2.41$ and 3.00 for our models, and the value 0.76 at $z = 2.75$, which we derive from published Keck spectra before allowing for a continuum offset (see Appendix B). The McDonald et al. and Croft et al. measurements agree reasonably well at $z = 2.7$ and 2.4. A systematic offset is apparent at $z = 2.4$, although this may be due to differences in the determination of the QSO continuum level,

and so some slight difference in the mean absorbed Ly α flux, to which the amplitude of the flux power spectrum is very sensitive (Croft et al.). The Croft et al. result at $z = 2.72$ is for their combined ‘fiducial sample’, resulting in a reduction in the errors compared with the five subsamples from which it is composed. We note that, according to equation (26) of Croft et al., based on their fiducial sample, model 1 is ruled out at the 3.9σ level, while model 4 and the concordance models are ruled out at only the 1.0σ and 1.9σ levels, respectively. At $z = 2.7$, we find model 1 fits about as well as model 4 and the concordance model. For $k < 0.015 \text{ km}^{-1}\text{s}$, the χ^2 for models 1, 4 and the concordance model, compared with the data of Croft et al., are, respectively, 13.8, 16.5 and 18.3. The respective probabilities $P(> \chi^2)$ for the eight k values are 0.087, 0.036 and 0.019, so that model 1 is actually preferred.

In order to understand the differences between our results and those of Croft et al. (2002b), we have performed a set of comparisons between the concordance model simulations and the measurements. The details of these comparisons are provided in Appendix A. The results exhibit a clear trend of diminishing χ^2 with increasing resolution, showing that the lower-resolution simulations have not yet matched the precision of the measured $P_F(k)$ values. Only at the highest resolution do the simulations begin to converge, giving values of sufficient precision to make a meaningful comparison with the measurements. The results also show that both the rate of convergence and the goodness of fit are highly sensitive to the amount of smoothing.

While we find the convergence properties improve with an increasing amount of Jeans smoothing, without direct comparisons to full hydrodynamical simulations of comparably high resolution, it is unclear how much artificial smoothing should be applied to the PM results, or, indeed, how accurately any smoothing scheme may mimic the true underlying hydrodynamical processes. A certain amount of smoothing will always be present in the spectra resulting from thermal broadening and peculiar velocities in the gas, which the PM simulations are able to account for. Indeed, there is a certain amount of degeneracy between the Jeans scale broadening and changes in the equation of state. The crucial question is whether the pressure forces missing in the PM simulations result in significantly less smoothing than should be physically present to obtain accurate predictions for $P_F(k)$. We are unable to address that question here, but we have demonstrated (see Appendix A) that the successful application of PM (or HPM) simulations, including assessments of convergence, for discriminating between cosmological models on the basis of the flux power spectrum relies critically on the adopted degree of artificial smoothing.

We compare our model predictions with the measured flux power spectra in Table 4, combining the data of McDonald et al. (2000) at $z = 2.41, 3.00$ and 3.89 and of Croft et al. (2002b) at $z = 2.72$, for $k < 0.015 \text{ km}^{-1}\text{s}$. The total number of degrees of freedom is 32. We have conservatively adopted mesh scale smoothing throughout. Within the context of mesh scale smoothing, the results we provide are upper bounds to the χ^2 for each model, since we have made no effort to improve the fits by making the several adjustments available. Specifically, we have not: (1) adjusted the mean Ly α flux to improve the fits; (2) adjusted the equation-of-state parameters T_0 and γ ; (3) accounted for the spread in the model predictions of $P_F(k)$; or (4) accounted for mode–mode correlations in the power for either the simulations or the data. We confine the comparison to $k < 0.015 \text{ km}^{-1}\text{s}$, since for this range the model predictions are more stable to changes in the box size, and the range avoids an anomalous low excursion point, with very small error bar, in the measured $P_F(k)$ at $z = 3$ for $k = 0.0179 \text{ km}^{-1}\text{s}$.




Figure 7. The flux power spectrum $P_F(k)$, as a function of rest-frame velocity wavenumber, for models 1 (cross) and 4 (square) and the concordance model (inverted Y), compared with the measurements in high-resolution spectra at the indicated redshifts. The solid data points are from McDonald et al. (2000) and the open data points from Croft et al. (2002b).

Table 4. Statistical comparison between the model predictions and measurements for the flux power spectrum, combining data at $z = 2.41, 2.72, 3.00$ and 3.89 . No adjustments to the equation of state or mean Ly α flux have been made. The number of measured flux power spectrum data points used is 32.

Model	χ^2	$P(> \chi^2)$
1	55.1	0.007
2	38.7	0.192
3	42.3	0.106
4	40.7	0.139
C	44.9	0.065

The effect of UV background fluctuations on the flux power spectrum is shown in Figs 8 and 9. Since the UV background fluctuations tend to suppress the mean flux somewhat (boost the mean optical depth) (Paper I), we have renormalized the uniform background case to the identical mean flux found for the fluctuating background case at each redshift to make a fair comparison. The effect of the fluctuations are most pronounced on large scales and small scales.

The small-scale peak is on the scale of the absorption linewidths and is due in part to a resetting of the mean ionization rate necessary to reproduce the same mean Ly α flux. At intermediate scales the power tends to be somewhat suppressed. The trend at low k of decreasing power with increasing k continues even to $z = 2.75$, although the effect is much reduced and essentially vanishes by $z = 2.41$. The rise at the lowest k value appears to be real, not an artefact of the finite box size, based on convergence tests applied to the concordance model (see Appendix A). At $z = 6$, we also show the small dilution effect of introducing an additional diffuse radiation background of 50 per cent of the average direct contribution from the QSO sources.

The ratios of the flux power spectra for models 1–3 to model 4 are shown in Fig. 10 at $z = 5$. The differences between the models are smaller than the effect induced by the UV background fluctuations on the largest scales (smallest wavenumbers). The effect of the fluctuations should be discernible in $z > 5$ spectra of sufficient signal-to-noise ratio within the current precision of the measured cosmological parameters. Because of the small flux values involved, however, it is unlikely that these tests could be performed at the currently achievable signal-to-noise ratios, except perhaps by going to higher orders in the Lyman series where the absorption will be less, although the contribution from the Ly α forest at the corresponding




Figure 8. The flux power spectrum $P_F(k)$ as a function of rest-frame velocity wavenumber, for model 4, with (square) and without (cross) UV background fluctuations. Also shown is $P_F(k)$ at $z = 6$ allowing for a slight suppression due to diffuse radiation (inverted Y). The results are shown at $z = 4, 5, 5.5$ and 6.

lower redshifts would have to be removed. These tests may best be performed with an extremely large telescope, providing novel science that would be made possible by such advanced technology.

4.3 Pixel flux autocorrelation function

We compute the flux autocorrelation function at a velocity separation v for the spectra according to

$$\xi_F(v) = \langle f(v')f(v' - v) \rangle / \langle f \rangle^2 - 1, \quad (9)$$

where $f(v)$ denotes the pixel flux at velocity v , and $\langle f \rangle$ is the measured mean flux at the relevant redshift. The results are shown for model 4 and the concordance model in Fig. 11 at $z = 2.75$. The data points are from Croft et al. (2002b) at $\langle z \rangle = 2.72$. The agreement is close at small separations. The small difference in amplitude may be due to a slight difference in the mean Ly α flux adopted for our models compared with that of Croft et al. (who have not derived an independent estimate of the mean Ly α flux from their spectra). We have made no attempt to adjust the adopted mean Ly α flux or equation of state to improve the match. At large separations the agreement is much worse, probably due to a lack of numerical convergence in the simulation results (see Appendix A).

We compare the flux autocorrelation function with and without UV background fluctuations in Fig. 12. The UV background fluctuations substantially boost the correlations on moderate to large scales. The effect is largest at $z = 6$ and decreases with decreasing redshift. The effect almost disappears by $z = 4$ for velocity separations smaller than 200 km s^{-1} . We note that the convergence of the autocorrelation function is not guaranteed for velocity separations exceeding $\sim 100\text{--}200 \text{ km s}^{-1}$, as discussed in Appendix A. At $z = 6$, we also show the small dilution effect of introducing an additional diffuse radiation background of 50 per cent of the average direct contribution from the QSO sources.

The ratios of the flux autocorrelation function for models 1–3 to model 4 are shown in Fig. 10 at $z = 5$. The differences between the models are smaller than the effect induced by the UV background fluctuations for small to moderate scales ($\lesssim 200 \text{ km s}^{-1}$), suggesting that the effect of the fluctuations would not be masked by uncertainties in the cosmological model.

5 CONCLUSIONS

We have examined the effect of UV background fluctuations expected from QSO sources on the power spectrum and autocorrelation function of the Ly α forest. To do so, we first performed




Figure 9. The ratio of the flux power spectrum $P_F(k)$ with UV background fluctuations to the value without, as a function of rest-frame velocity wavenumber, for model 4. The results are shown at $z = 2.75, 3, 3.89$ and 4 . The effects of the UV background fluctuations are qualitatively similar to those found at higher redshift, shown in Fig. 8, though smaller in amplitude.

several sequences of tests to assess the simulation parameters required to reach convergence. We investigate the convergence for different amounts of smoothing meant to mimic hydrodynamical effects, almost all for a fixed equation of state. We note that varying the equation of state will affect the predicted flux power spectrum, and should be included when attempting to constrain cosmological models by comparison with the data. We have not performed an exhaustive search in smoothing algorithms and equations of state to achieve a best match to the data. Our emphasis in this paper is instead to explore the impact of these on the *relative* effect that UV background fluctuations are expected to have on the predicted flux power spectrum and autocorrelation function.

Our general conclusions concerning the convergence tests are as follows (where all length scales are comoving): (1) A spatial resolution of at least $60 h^{-1}$ kpc is required to converge to 10 per cent accuracy in the metagalactic photoionization background required to match the measured mean absorbed Ly α flux. (2) A spatial resolution of at least $60 h^{-1}$ kpc is required to converge to 10 per cent accuracy in the 1D flux power spectrum at low k values, and $30 h^{-1}$ kpc for 10 per cent accuracy near the peak in $kP_F(k)$ unless strong Jeans smoothing is applied. (3) Smoothing the density field on the

Jeans scale generally improves convergence, but flattens the flux power spectrum near its peak. (4) The level to which a model may be statistically rejected on the basis of current measurements of the 1D flux power spectrum is sensitive to the resolution, box size and degree of density field smoothing. We recommend a box size of at least $25 h^{-1}$ Mpc, and a force mesh of grid size $60 h^{-1}$ kpc when the density field is strongly smoothed on the Jeans scale and $30 h^{-1}$ kpc when smoothed on the mesh scale. (5) The flux autocorrelation function generally exhibits poor convergence properties, converging to better than 10 per cent only for separations smaller than 10 per cent of the box size. (6) We were not able to converge on the 1D underlying matter power spectrum on any scale for box sizes as large as $60 h^{-1}$ Mpc and force mesh resolutions as small as $30 h^{-1}$ kpc. A more detailed summary of our convergence tests may be found at the end of Appendix A.

We find significant degeneracy between the smoothing and the equation of state of the baryons. It is unclear how the degeneracy may be removed without performing full hydrodynamic simulations to infer the equation of state of the IGM by comparing with the measured small-scale fluctuations of the Ly α forest on the scale of the absorption features. A possibility is to use the measured Doppler

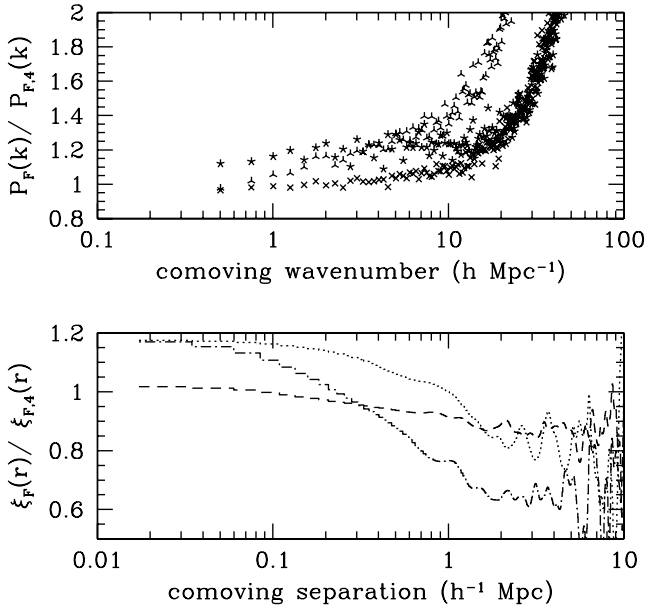


Figure 10. (Upper panel) The ratio of the flux power spectrum for model 1 (star), model 2 (cross) and model 3 (inverted Y) to that of model 4 at $z = 5$. (Lower panel) The ratio of the flux autocorrelation function for model 1 (dotted), model 2 (short-dashed) and model 3 (dot-short-dashed) to that of model 4 at $z = 5$. For model 4 at $z = 5$, $1 h^{-1}$ Mpc (comoving) corresponds to 134 km s^{-1} (proper), and $1 h \text{ Mpc}^{-1}$ (comoving) corresponds to $0.00741 \text{ km}^{-1} \text{ s}$ (proper).




Figure 11. The flux autocorrelation function $\xi_F(v)$ at $z = 2.7$, as a function of rest-frame velocity separation, for model 4 (solid) and the concordance model (dotted), assuming Gaussian smoothing on the mesh scale. Also shown are the concordance model using linear Jeans smoothing with $\eta = 2$ (long-dashed) and $\eta = 4$ (short-dashed). The models are compared with the measurements from Croft et al. (2002b).

parameters to constrain the equation of state (Schaye et al. 2000), but, without high-resolution simulations in box sizes of at least $\sim 25 h^{-1}$ Mpc, it is unclear whether or not the parameters deduced in this way are the optimal ones to use for predicting the flux power spectrum. One complication is the evident need for additional broadening at $z < 3.5$ beyond simple photoionization heating (Theuns et al. 1999; Bryan & Machacek 2000; Meiksin et al. 2001). The discrepancy in the predicted and measured linewidths suggests that




Figure 12. The flux autocorrelation function $\xi_F(v)$, as a function of rest-frame velocity separation, for model 4, with (solid) and without (dashed) UV background fluctuations. Also shown is $\xi_F(v)$ at $z = 6$ including UV background fluctuations, but allowing for a slight suppression due to additional diffuse radiation (dot-dashed line). The results are shown at $z = 4, 5, 5.5$ and 6 .

the underlying physics is still not fully understood, particularly for the moderate to low optical depth lines (Meiksin et al. 2001).

The simulations are normalized by the average Ly α flux transmitted through the IGM. We find that there is significantly more accord in the measurements of this parameter than previously recognized once existing constraints are interpreted on a consistent statistical basis, as discussed in Appendix B.

We assess the possibility of using measurements of the flux power spectrum for inferring the contribution of QSO sources to the meta-galactic UV ionizing background. Because of the sparsity of QSOs relative to alternative sources, like galaxies, the UV background contribution from QSOs will exhibit large, spatially correlated fluctuations. We model the fluctuations using Monte Carlo realizations of randomly distributed QSO sources based on QSO counts from the 2dF and SDSS QSO surveys. We estimate the attenuation lengths through the IGM self-consistently from our simulations, requiring a background ionization rate that recovers the measured mean Ly α flux. Because the simulations are not able to recover the measured number of Lyman-limit systems for $z < 4$, we add their contribution to the attenuation by hand.

The resulting fluctuations increase the required photoionization rates by 5–10 per cent over the uniform background case for $z \leq 5.5$ in order to match the measured mean Ly α fluxes. This is a somewhat smaller effect than was found in Paper I, where only the UV background fluctuations were accounted for and not their spatial correlations. By $z = 6$, we find the trend reverses in the presence of the now very strong UV background correlations: the required photoionization rate is reduced by 10–20 per cent compared with the uniform background case. By contrast, we find that the effect on the cumulative flux distribution increases when the correlations in the UV background fluctuations are included. Previously we found the fluctuations distorted the flux distributions by ~ 0.2 per

cent for $z \geq 5$. Allowing for the correlations increases the effect to 1–1.5 per cent, rendering them somewhat easier to detect.

We find that the magnitudes of the *relative* effects on the two-point flux statistics by fluctuations in the UV background are much better converged than the absolute predictions. The ratios of the predicted flux power spectra with and without UV background fluctuations are converged to better than 1 per cent on scales $k < 0.2 \text{ km}^{-1}$ s in our simulations, and are in agreement for Gaussian mesh, Jeans Gaussian and Jeans linear smoothings at this level over these scales. The ratio of the predicted autocorrelation function with and without UV background fluctuations, however, is converged to better than 10 per cent only for velocity separations $\lesssim 200 \text{ km s}^{-1}$.

We also find that the UV background fluctuations boost the flux power spectrum at large scales ($k < 0.01 \text{ km}^{-1}$ s) and small scales ($k \sim 0.2\text{--}0.4 \text{ km}^{-1}$ s, corresponding to the widths of individual absorption features), and suppress the power at intermediate scales. For $z \leq 4$, the effects are small, at the few per cent level. At higher redshifts, however, the large-scale effect grows due to the increasing strength of the UV background correlations, resulting in a boost in power of several per cent by $z > 5$. By $z = 6$, the background fluctuations produce an upturn in $k P_{\text{F}}(k)/\pi$ towards low k , resulting in a more than doubling in power and providing distinctive evidence for the presence of UV background fluctuations, as such an upturn would be unphysical from large-scale structure effects alone. We note that the boost may be suppressed at very large wavelengths if QSOs have lifetimes shorter than the light crossing time over an attenuation length (Croft 2004). Measuring the effects at $z > 5.5$ with the current level of flux detection is problematic because of the very high Ly α optical depths at these redshifts, but it may become feasible with future extremely large telescope technology, or possibly by searching for the signal at higher orders in the Lyman series where the optical depths are reduced.

ACKNOWLEDGMENTS

MW was supported by the NSF and NASA. Parts of this work were done on the IBM-SP at the National Energy Research Scientific Computing Center.

REFERENCES

- Bernardi M. et al., 2003, *AJ*, 125, 32
 Bond J. R., Wadsley J. W., 1997, in Petitjean P., Charlot S., eds, *Structure and Evolution of the Intergalactic Medium from QSO Absorption Line Systems*. Editions Frontières, Paris, p. 143
 Boyle B. J., Shanks T., Peterson B. A., 1988, *MNRAS*, 235, 935
 Boyle B. J., Shanks T., Croom S. M., Smith R. J., Miler L., Loaring N., Heymans C., 2000, *MNRAS*, 317, 1014
 Bryan G. L., Machacek M. E., 2000, *ApJ*, 534, 57
 Cen R., Miralda-Escudé J., Ostriker J. P., Rauch M., 1994, *ApJ*, 437, L9
 Croft R. A. C., 2004, *ApJ*, submitted (astro-ph/0310890)
 Croft R. A. C., Weinberg D. H., Katz N., Hernquist L., 1998, *ApJ*, 495, 44
 Croft R. A. C., Weinberg D. H., Pettini M., Hernquist L., Katz N., 1999, *ApJ*, 520, 1
 Croft R. A. C., Hernquist L., Springel V., Westover M., White M., 2002a, *ApJ*, 580, 634
 Croft R. A. C., Weinberg D. H., Bolte M., Burles S., Hernquist L., Katz N., Kirkman D., Tytler D., 2002b, *ApJ*, 581, 20
 Fan X. et al., 2001, *AJ*, 121, 54
 Fan X., Narayanan V. K., Strauss M. A., White R. L., Becker R. H., Pentericci L., Rix H.-W., 2002, *AJ*, 123, 1247
 Fardal M., Shull M., 1993, *ApJ*, 415, 524
 Ferguson H. C., Dickinson M., Papovich C., 2002, *ApJ*, 569, 65
 Gardner J. P., Katz N., Hernquist L., Weinberg D. H., 1997, *ApJ*, 484, 31
 Gnedin N. Y., Hamilton A. J. S., 2002, *MNRAS*, 334, 107
 Gnedin N. Y., Hui L., 1998, *MNRAS*, 296, 44
 Haardt F., Madau P., 1996, *ApJ*, 461, 20
 Hernquist L., Katz N., Weinberg D., Miralda-Escudé J., 1996, *ApJ*, 457, L51
 Hinshaw G. et al., 2003, *ApJS*, 148, 135
 Kim T.-S., Cristiani S., D'Odorico S., 2001, *A&A*, 373, 757
 Kogut A. et al., 2003, *ApJS*, 148, 161
 McDonald P., 2003, *ApJ*, 585, 34
 McDonald P., Miralda-Escudé J., Rauch M., Sargent W. L. W., Barlow T. A., Cen R., Ostriker J. P., 2000, *ApJ*, 543, 1
 Madau P., Meiksin A., 1994, *ApJ*, 433, L53
 Meiksin A., 1994, *ApJ*, 431, 109
 Meiksin A., 2000, *MNRAS*, 314, 566
 Meiksin A., Madau P., 1993, *ApJ*, 412, 34
 Meiksin A., White M., 2001, *MNRAS*, 324, 141
 Meiksin A., White M., 2003, *MNRAS*, 342, 1205 (Paper I)
 Meiksin A., Bryan G. L., Machacek M. E., 2001, *MNRAS*, 327, 296
 Ostriker J., Steinhardt P. J., 1995, *Nat*, 377, 600
 Peacock J. A., Dodds S. A., 1996, *MNRAS*, 280, 19
 Petitjean P., Mücke J. P., Kates R. E., 1995, *A&A*, 295, L9
 Press W. H., Rybicki G. B., Schneider D. P., 1993, *ApJ*, 414, 64 (PRS)
 Reimers D., Köhler S., Wisotzki L., Groote D., Rodriguez-Pascual P., Wamsteker W., 1997, *A&A*, 327, 890
 Schaye J., Theuns T., Leonard A., Efstathiou G., 1999, *MNRAS*, 310, 57
 Schaye J., Theuns T., Rauch M., Efstathiou G., Sargent W. L. W., 2000, *MNRAS*, 318, 817
 Seljak U., McDonald P., Makarov A., 2003, *MNRAS*, 342, L79
 Songaila A., Cowie L. L., 2002, *AJ*, 123, 2183
 Spergel D. N. et al., 2003, *ApJS*, 148, 175
 Steidel C. C., Sargent W. L. W., 1987, *ApJ*, 313, 171
 Steidel C. C., Pettini M., Adelberger K. L., 2001, *ApJ*, 546, 665
 Stengler-Larrea E. A. et al., 1995, *ApJ*, 444, 64
 Storrie-Lombardi L. J., Irwin M. J., McMahon R. G., 1996, *MNRAS*, 282, 1330
 Theuns T., Leonard A., Efstathiou G., 1998, *MNRAS*, 297, L49
 Theuns T., Leonard A., Schaye J., Efstathiou G., 1999, *MNRAS*, 303, L58
 Zaldarriaga M., Hui L., Tegmark M., 2001, *ApJ*, 557, 519
 Zhang Y., Anninos P., Norman M. L., 1995, *ApJ*, 453, L57
 Zhang Y., Anninos P., Norman M. L., Meiksin A., 1997, *ApJ*, 485, 496
 Zhang Y., Meiksin A., Anninos P., Norman M. L., 1998, *ApJ*, 495, 63
 Zheng W., Kriss G. A., Telfer R. C., Grimes J. P., Davidsen A. F., 1997, *ApJ*, 475, 469
 Zuo L., 1992a, *MNRAS*, 258, 36
 Zuo L., 1992b, *MNRAS*, 258, 45
 Zuo L., Bond J. R., 1994, *ApJ*, 423, 73
 Zuo L., Phinney E. S., 1993, *ApJ*, 418, 28

APPENDIX A: CONVERGENCE TESTS

Even including only gravitational physics, simulations of the Ly α forest are restricted by two size constraints. High spatial resolution is required to resolve the structure of the absorbers, while a large box size is required both to capture the large-scale power (which will affect the velocity widths of the absorption lines) and to obtain a fair sample of the universe. Although a small box may be adequate to resolve the line structure, if the box is too small the fluctuations corresponding to the scale of the box will become non-linear, in which case the results are no longer representative of the cosmological model simulated. These numerical effects have been investigated in different contexts by a number of authors (see references in McDonald 2003 for example).

To investigate these effects of numerical resolution further, we chose a relatively large box, $30 h^{-1}$ Mpc (comoving) on a side, and simulated the concordance cosmology using 128^3 , 256^3 or 512^3 particles (corresponding to 256^3 , 512^3 and 1024^3 mesh cells). At the high redshifts of interest to us, the fundamental mode of the




Figure A1. The (real-space) 1D matter power spectrum in the concordance cosmology at $z = 5$ from our resolution sequence with 128^3 , 256^3 and 512^3 particles and force meshes of 256^3 , 512^3 and 1024^3 . The solid line shows the predicted 1D spectrum, using the Peacock & Dodds (1996) fitting formulae for the 3D matter power spectrum, equation (A1), and a Gaussian filtering of three mesh cells (FWHM).

box is safely in the linear regime, suggesting that the simulation represents a fair sample of the universe. To isolate the effects of numerical resolution from sample variance, we sampled the initial conditions starting at low k and moving ‘outwards’ to higher k so that the initial phases and amplitudes of the common modes were identical from simulation to simulation.

The 3D matter power spectrum (not shown) exhibits the expected behaviour, with the higher-resolution simulations having progressively more small-scale power. At $z = 5$ the fundamental mode is safely in the linear regime, with $\Delta^2(k_{\text{fund}}) \simeq 0.04$, and the mesh scale is non-linear with $\Delta^2(k_{\text{mesh}}) > 10$. The fitting formula of Peacock & Dodds (1996) underestimates the non-linear power by $\mathcal{O}(50\text{ per cent})$. Fig. A1 shows the 1D matter power spectrum, in real space. This is an integral of the 3D spectrum over wavenumbers greater than the 1D wavenumber under consideration:

$$\frac{kP(k)}{\pi} = k \int_k^\infty \frac{dy}{y^2} \Delta_{3D}^2(y). \quad (\text{A1})$$

We find that the numerical results are relatively well fitted by the Peacock & Dodds (1996) formalism, if we filter the 3D power spectrum, $\Delta_{3D}^2(k)$, with a Gaussian filter of full width at half-maximum (FWHM) 3.2 mesh cells (plotted as the solid line for the highest-resolution run), to approximate the finite force resolution of the PM code. Note that missing high- k power in the lower-resolution simulations biases the 1D power spectrum low over the whole k range plotted since much of the long-wavelength power comes from aliasing of 3D short-wavelength power.

We show in Fig. A2 the power spectra of the flux at $z = 5$ produced by these three runs in a $30 h^{-1}$ Mpc box using a Gaussian smoothing of one mesh cell to define the density and velocity fields. The flux power spectrum converges to within 5 per cent for $k < 0.01 \text{ km}^{-1}\text{s}$, to 10 per cent for $k < 0.016 \text{ km}^{-1}\text{s}$ and to 20 per cent for $k < 0.032 \text{ km}^{-1}\text{s}$. At $z = 3$ (not shown here), the flux power spectrum converges to better than 12 per cent for $k < 0.032 \text{ km}^{-1}\text{s}$. From these results we infer that, for the relatively high redshifts of interest, a $30 h^{-1}$ Mpc box with 512^3 particles and a 1024^3 force mesh produces power spectra converged to better than the 10 per cent level for $k < 0.016 \text{ km}^{-1}\text{s}$.



Figure A2. The 1D flux power spectrum in the concordance cosmology at $z = 5$ from our resolution sequence with 128^3 , 256^3 and 512^3 particles and force meshes of 256^3 , 512^3 and 1024^3 . The box size in each run is $30 h^{-1}$ Mpc. These results indicate that our highest-resolution runs are converged at better than the 10 per cent level (in power) on scales $k < 0.016 \text{ km}^{-1}\text{s}$.

Similar conclusions were found by McDonald (2003), though our convergence is not quite at the level quoted therein. Part of the reason for this is the treatment of small-scale power in the simulations. To investigate this we modified our smoothing procedure following Gnedin & Hui (1998) to mimic the effects of thermal pressure on the baryonic component, using either Gaussian or linear smoothing as described in Section 2 above. Gnedin & Hui relate a filter scale k_f to the Jeans scale k_J using linear theory, giving $k_f = \eta k_J$, where η depends on the temperature history of the IGM after reionization. They provide expressions for $\eta(z)$ for the cases of an isothermal IGM and a temperature decaying linearly with the expansion factor. More generally, if the post-ionization temperature varies like $T \sim (1+z)^{-\alpha}$, the filter length at redshift z is given in an Einstein–de Sitter universe (which adequately represents a currently Λ -dominated universe at the early redshifts of interest) by (for $\alpha \neq -2$ or -2.5),

$$\frac{k_f^2(z)}{k_f^2(z)} = \frac{6}{(2\alpha + 4)(2\alpha + 5)} \left[1 + (2\alpha + 4) \left(\frac{1+z}{1+z_{ri}} \right)^{\alpha+5/2} - (2\alpha + 5) \left(\frac{1+z}{1+z_{ri}} \right)^{\alpha+2} \right], \quad (\text{A2})$$

where z_{ri} is the redshift of reionization. For an increasing IGM temperature with expansion factor corresponding to $\alpha = 1$,

$$\frac{k_f^2(z)}{k_f^2(z)} = \frac{1}{7} \left[1 + 6 \left(\frac{1+z}{1+z_{ri}} \right)^{7/2} - 7 \left(\frac{1+z}{1+z_{ri}} \right)^3 \right]. \quad (\text{A3})$$

Plausible values for η over $3 < z < 5$ range from $\eta \gtrsim 1$ at $z = 3$, for early reionization and a subsequently isothermal IGM, to $\eta = 5$ at $z = 5$, for late reionization and an IGM temperature increasing with expansion factor. We shall consider the values $\eta = 2$ and 4 here. We note that, for $\eta = 4$, k_f corresponds to scales below the minimal Gaussian grid smoothing we apply for the runs in a $30 h^{-1}$ Mpc box for force meshes with 256^3 and 512^3 zones, so these default to our minimal smoothing case, as do all our $60 h^{-1}$ Mpc box runs. For $\eta = 2$, only the 256^3 $30 h^{-1}$ Mpc box run and the 256^3 and 512^3 $60 h^{-1}$ Mpc box runs default to the minimal smoothing case.

We show in Fig. A3 that the Gaussian smoothing method used by Zaldarriaga et al. (2001), with $\eta = 2$, improves the




Figure A3. The 1D flux power spectrum in the concordance cosmology at $z = 5$ from a resolution sequence with 128^3 , 256^3 and 512^3 particles and force meshes of 256^3 , 512^3 and 1024^3 . The box size in each run is $30 h^{-1}$ Mpc. (Upper panel) The density and velocity fields are Jeans filtered by $\exp[-(k/k_f)^2]$ ($\eta = 2$; see text). (Lower panel) The density and velocity fields are Jeans filtered by $[1 + (k/k_f)^2]^{-1}$ ($\eta = 2$; see text).

convergence, resulting in convergence to 10 per cent over $k < 0.08 \text{ km}^{-1} \text{ s}$. This is to be expected since the filter scale here is close to the mesh size of our intermediate-resolution run. We found that pure linear theory smoothing introduced numerical artefacts unless it was supplemented with our standard single mesh cell Gaussian smoothing, so we show in Fig. A3 the results of this hybrid. For the smaller-resolution runs the Gaussian smoothing dominates, but at higher resolution the linear Jeans smoothing suppresses small-scale power as well, though the resulting convergence is not as good as in the Gaussian Jeans smoothing case. The flux power spectrum converges to 10 per cent over $k < 0.01 \text{ km}^{-1} \text{ s}$ and to 15 per cent over $k < 0.032 \text{ km}^{-1} \text{ s}$.

The convergence becomes fractionally better if we increase T_0 or decrease γ . It becomes only slightly worse if we lower T_0 to 10^4 K . A further sequence reducing the box size (not shown here) suggests that we can safely reduce the volume simulated to a box of side $25 h^{-1}$ Mpc if we are interested in the Ly α forest at high redshift. This makes the mesh scale slightly smaller, capturing the features in the forest slightly better, and thus we have used $25 h^{-1}$ Mpc throughout.

We show in Fig. A4 a convergence study at lower redshift ($z = 3$) in a larger box ($60 h^{-1}$ Mpc). The power spectrum in the larger box follows closely that of the smaller box with the identical physical resolution. It thus departs significantly from that in the $30 h^{-1}$ Mpc box at $k \simeq 0.01 \text{ km}^{-1} \text{ s}$ for the highest-resolution runs slightly earlier than might be expected from Fig. A2 due to the increased non-linear scale.

In Fig. A5, we show a convergence sequence at $z = 4$ for models including fluctuations in the UV background. The UV background fluctuations are based on the QSO counts with $\beta_1 = 3.2$ described in the text, with an added diffuse component increasing the ionization rate by 40 per cent to reproduce the measured average flux. The rise in $P_F(k)$ at the lowest k values appears to be real. The effect of the




Figure A4. The 1D flux power spectrum in the concordance cosmology at $z = 3$ from a resolution sequence with 128^3 , 256^3 and 512^3 particles and force meshes of 256^3 , 512^3 and 1024^3 . The box size in each run is $60 h^{-1}$ Mpc.




Figure A5. The ratios of the (upper panel) 1D flux power spectrum and (lower panel) flux autocorrelation function, with and without UV background fluctuations, in the concordance cosmology at $z = 4$ from a resolution sequence with comoving grid size $\Delta r = 58.6 h^{-1}$ kpc, in boxes of comoving sides $30 h^{-1}$ Mpc (cross) and $60 h^{-1}$ Mpc (inverted Y), and with a comoving grid size $\Delta r = 29.3 h^{-1}$ kpc in a box of comoving side $30 h^{-1}$ Mpc (star).

UV background fluctuations on the flux autocorrelation becomes poorly determined for velocity separations exceeding 300 km s^{-1} .

We show in Fig. A6 that the Gaussian and linear Jeans smoothing, with $\eta = 2$, improve the convergence of the flux autocorrelation function over the Gaussian single mesh scale smoothing. For 256^3 particles on a 512^3 force mesh, the Jeans smoothings give convergence to better than 5 per cent on comoving separations of $r < 10 h^{-1}$ Mpc, or a third of the box size. The single mesh Gaussian smoothing converges to only 15 per cent on these scales. Also shown is a comparison of the single mesh scale Gaussian smoothing for comoving box sizes of $30 h^{-1}$ Mpc and $60 h^{-1}$ Mpc, at two spatial resolutions, showing convergence to 10 per cent only for comoving separations of $r < 1 h^{-1}$ Mpc, or 3 per cent of the box size. We find




Figure A6. The error in the 1D flux autocorrelation function, as a function of comoving separation, in the concordance cosmology at $z = 4$ from (panels a–c) a resolution sequence in a box of comoving size $30 h^{-1}$ Mpc with 128^3 and 256^3 particles and force meshes of 256^3 and 512^3 , compared with 512^3 particles and a force mesh of 1024^3 , and from (panel d) a box size sequence comparing boxes of comoving size $30 h^{-1}$ Mpc with $60 h^{-1}$ Mpc for (comoving) force mesh resolutions of $\Delta r = 117 h^{-1}$ kpc and $58.6 h^{-1}$ kpc. (The Jeans smoothing cases are for $\eta = 2$.)

a similar result for the Gaussian Jeans smoothing case with $\eta = 2$, showing that the poor convergence with box size is not sensitive to the amount of smoothing. The results suggest instead that modes with wavelengths longer than the fundamental of the $30 h^{-1}$ Mpc box are still contributing substantially to the correlation function at large separations.

We note that the poor convergence found for the flux autocorrelation function does not conflict with the better convergence properties we obtain for the flux power spectrum in spite of the Fourier relation between the two. The reason is that the Fourier transform requires integrating either over vanishingly small k modes (to go from the power spectrum to the autocorrelation function) or infinitely long spatial extents (to go in the reverse direction). We find that, while the simulations are able to converge well on the smallest k modes available to the box, additional longer-wavelength modes are required to converge to the autocorrelation function to better than 10 per cent at separations exceeding a few per cent of the box size. This is perhaps not too surprising since the 1D power spectrum continues to climb at lower k modes than we have simulated, and these long-wavelength modes contribute substantially to the autocorrelation function at large separations. A further discussion of this point may be found in McDonald et al. (2000), who similarly found that the flux autocorrelation function exhibits poor convergence properties, which they attribute to the missing modes.

While the smoothing schemes introduced above generally improve convergence, this does not guarantee that they converge to the correct power spectra as would be computed using a full hydrodynamical scheme. A direct comparison is still prohibitively computationally expensive because of the wide range in length-scales required both to resolve the non-linear structures into which the bulk of the IGM collapses and to capture sufficient large-scale power to obtain an accurate computation of the flux power spectrum and flux autocorrelation function. The most we can demonstrate is that our main results are not excessively sensitive to the smoothing scheme.

In Fig. A7, we compare the three smoothing methods for the concordance model at $z = 4$, normalized to $\langle \exp(-\tau) \rangle = 0.47$. The power spectra for both Gaussian and linear Jeans smoothings (for fixed η) agree to within 1–2 per cent, although they are suppressed in power compared with the Gaussian single mesh smoothing by as much as 20 per cent near the peak for $\eta = 2$ and 8 per cent for $\eta = 4$. Comparable differences are found between the Jeans smoothed flux autocorrelation functions and the mesh scale Gaussian smoothing, with increased smoothing generally resulting in a flatter autocorrelation function. For simplicity, we retain the mesh scale Gaussian smoothing for the main results in the text, but make reference to alternative smoothings.

The different smoothing methods result in different rates of convergence of the metagalactic ionization rate Γ_{-12} required to match




Figure A7. The 1D flux power spectrum (upper panel) and flux autocorrelation function (lower panel) in the concordance cosmology at $z = 4$ for Gaussian smoothing on the mesh size, Jeans Gaussian smoothing and Jeans linear smoothing. The box size in each run is $30 h^{-1}$ Mpc.




Figure A8. The convergence of the required metagalactic ionization rate Γ_{-12} , in units of 10^{-12} s^{-1} , at $z = 5$ (upper panel) and $z = 4$ (lower panel) for force mesh (comoving) resolutions of $\Delta r = 29.3 h^{-1}$ kpc, $58.6 h^{-1}$ kpc and $117 h^{-1}$ kpc, for Gaussian smoothing on the mesh size (circle), Jeans Gaussian smoothing ($\eta = 2$, cross; $\eta = 4$, square) and Jeans linear smoothing ($\eta = 2$, inverted Y; $\eta = 2$, triangle). Results shown are for the concordance model. The comoving box size in each run is $30 h^{-1}$ Mpc.

the measured mean Ly α flux, as shown in Fig. A8. The two Jeans smoothing cases with $\eta = 2$ agree well with each other and show the least sensitivity to resolution, but the convergence becomes poorer for $\eta = 4$ and for the minimal mesh scale Gaussian smoothing case. For a force mesh of 1024^3 , the mesh scale smoothing and Jeans




Figure A9. The χ^2 per degree of freedom between the flux power spectrum $P_F(k)$ measured at $\langle z \rangle = 2.72$ by Croft et al. (2002b), and the predictions of the concordance model (at $z = 2.75$), as a function of (comoving) resolution. The open symbols use all data points with $k < 0.016 \text{ km}^{-1} \text{ s}$, and the filled symbols use all data points with $k < 0.032 \text{ km}^{-1} \text{ s}$. The results are shown for the $30 h^{-1}$ Mpc (comoving) box simulations, for Gaussian smoothing on the mesh scale (triangle), Gaussian Jeans smoothing ($\eta = 2$, inverted triangle) and linear Jeans smoothing ($\eta = 2$, square). (The stars are overlapping triangles.) Also shown is the Gaussian mesh size smoothing case for the $60 h^{-1}$ Mpc (comoving) box simulation for $k < 0.016 \text{ km}^{-1} \text{ s}$ (plus) and $k < 0.032 \text{ km}^{-1} \text{ s}$ (times), as well as the results using alternative equations of state (open and filled diamonds) for the $30 h^{-1}$ Mpc high-resolution box (see text). To help guide the eye, filled symbols are connected by solid lines and open symbols by dashed lines.

smoothing cases with $\eta = 2$ yield ionization rates differing by over 30 per cent at $z = 5$, though by less than 10 per cent by $z = 4$.

Finally, we treat the degree of convergence relative to the accuracy of the measurements of the flux power spectrum. We have performed a resolution sequence of comparisons between the concordance model simulations and the measurements, for a fixed equation of state (i.e. leaving T_0 and γ unchanged). In Fig. A9 we show how the reduced χ^2 varies with resolution, for both $k < 0.016 \text{ km}^{-1} \text{ s}$ and $k < 0.032 \text{ km}^{-1} \text{ s}$. [The simulation values are linearly interpolated to the k values of the measured $P_F(k)$ points.] We do the comparison for mesh scale smoothing, Gaussian Jeans smoothing and linear Jeans smoothing, for $\eta = 2$. For Jeans smoothing with $\eta = 4$, the smoothing defaults to the minimal Gaussian mesh scale smoothing except at the highest resolution, for which the resulting reduced χ^2 values are approximately double those of the minimal mesh scale smoothing case. A clear trend of diminishing χ^2 with increasing resolution is apparent in all cases, showing that the lower-resolution simulations have not yet matched the precision of the measured $P_F(k)$ values. Only at the highest resolution do the simulations begin to converge, giving values of sufficient precision to make a meaningful comparison with the measurements.

The decrease in χ^2 at higher resolution is noteworthy, as it suggests that a given model is a better fit to the data than would be believed on the basis of lower-resolution simulations. The larger χ^2 values found for the $60 h^{-1}$ Mpc box for $k < 0.032 \text{ km}^{-1} \text{ s}$ suggest that the simulations have still not quite converged on these scales, at least at a (comoving) resolution of $\Delta r = 58.6 h^{-1}$ kpc, particularly for $z < 3$. For comparison, the simulations used by Croft et al. (2002b) to normalize their model constraints were in a (comoving) box size of $27.77 h^{-1}$ Mpc with a force resolution of $109 h^{-1}$ kpc,

which for our models would be inadequate for converging to the correct χ^2 per degree of freedom except in the presence of strong Jeans smoothing.

The results in Fig. A9 also show that both the rate of convergence and the goodness of fit are highly sensitive to the amount of smoothing. Adopting Gaussian Jeans smoothing with $\eta = 2$ instead of smoothing on the mesh scale leads to a much improved rate of convergence, but to a significantly different value, resulting in $\chi^2 = 68.3$ for $k < 0.016 \text{ km}^{-1}\text{s}$ and $P(> \chi^2) = 1.0 \times 10^{-11}$. Adopting Gaussian Jeans smoothing with $\eta = 2$ gives instead $\chi^2 = 36.6$ for $k < 0.016 \text{ km}^{-1}\text{s}$ and $P(> \chi^2) = 1.4 \times 10^{-5}$. Both cases are in considerably poorer agreement with the data than found for the mesh scale smoothing case, which agrees with the measurements to within 2σ . The agreement is also much poorer than predicted by Croft et al. (2002b), although we have not attempted to adjust the mean flux or equation of state to improve the match. Similarly poorer results are found for linear Jeans smoothing. We also compute the effect of a change in the equation of state for the mesh scale smoothing case, adopting the values for T_0 and γ predicted by the late He II reionization model of Schaye et al. (2000). The agreement considerably worsens, as shown in Fig. A9, resulting in values comparable to the Jeans smoothing cases.

We conclude this appendix with a summary of our principal findings concerning convergence. (All length scales are comoving and velocity scales proper in the following.)

(i) To reach 10 per cent convergence in the required H I photoionization rate Γ_{-12} , a (comoving) spatial resolution of better than $60 h^{-1} \text{ kpc}$ is required for minimal smoothing on the mesh scale. The convergence is improved if sufficiently strong Jeans filtering is applied, although the converged value is sensitive to the amount of smoothing.

(ii) We were not able to reach convergence in the 1D dark matter power spectrum to better than 50 per cent on any scale. Box sizes exceeding $60 h^{-1} \text{ Mpc}$ and/or a spatial resolution of better than $30 h^{-1} \text{ kpc}$ appear to be required at $z = 5$. The convergence requirements are even more severe at lower redshifts. Fortunately, good convergence in the 1D matter power spectrum is not required for good convergence in the predicted flux power spectrum because thermal and Doppler broadening of the gas filters out the contribution of the high spatial frequency modes to the H I optical depth.

(iii) We show that the predicted flux power spectrum $P_F(k)$ is sensitive to the degree of smoothing. A crucial question is how much (or even whether) any additional smoothing to the gas variables should be applied beyond the inherent smoothing of the spectra due to Doppler and peculiar velocity broadening. For the minimal smoothing on the scale of a mesh cell, the minimum resolution and box size required to achieve 10 per cent convergence on (proper) scales $k < 0.016 \text{ km}^{-1}\text{s}$ is $60 h^{-1} \text{ kpc}$ in a $30 h^{-1} \text{ Mpc}$ box. We were not able to achieve convergence to 10 per cent at larger wavenumbers for $z = 5$, but could at $z = 3$. Allowing for additional smoothing improves the convergence. With the application of Gaussian Jeans smoothing, it is possible to achieve convergence on all scales $k < 0.08 \text{ km}^{-1}\text{s}$ to 10 per cent at $z = 5$ for a spatial resolution of at least $60 h^{-1} \text{ kpc}$ in a $30 h^{-1} \text{ Mpc}$ box. With the application of linear Jeans smoothing, we are able to achieve this level of accuracy only for $k \lesssim 0.01 \text{ km}^{-1}\text{s}$. The convergence requirements become more severe at lower redshifts.

(iv) We find that introducing Jeans Gaussian or Jeans linear smoothing significantly flattens the predicted flux power spectrum near its peak. Without full hydrodynamical simulations at compa-

table resolution to our highest-resolution simulations in a box size of at least $25 h^{-1} \text{ Mpc}$, it is not possible to assess whether or not the additional Jeans smoothing, while it improves the convergence properties, also improves the accuracy of the flux power spectrum as predicted from pure gravity simulations.

(v) We were unable to obtain convergence in the flux autocorrelation function $\xi_F(k)$ to better than 10 per cent on spatial separations exceeding ~ 3 per cent of the box size for the minimal mesh scale smoothing case. Allowing for Jeans smoothing improves the convergence for a fixed box size, reaching a level of 5 per cent convergence over a third of the box size for sufficiently strong smoothing. Doubling the box size, however, produces significant changes in the autocorrelation function, so that it is unclear whether convergence at better than 10 per cent is possible on separations exceeding ~ 3 per cent of the box size even for strong Jeans smoothing. Comparisons with the data for this statistic for separations exceeding $100\text{--}200 \text{ km s}^{-1}$ appear problematic.

(vi) A criterion of the level of simulation convergence needed to make a meaningful comparison with observations is the convergence in χ^2 between the simulation predictions and the measured values. We find that χ^2 for the comparison between the predicted and measured $P_F(k)$ continues to decrease with increasing resolution. For our minimal mesh scale smoothing case, we find a reduction in χ^2 at $z = 2.7$ by a factor of 2.5 on going from a resolution of $\sim 60 h^{-1} \text{ kpc}$ to $\sim 30 h^{-1} \text{ kpc}$ in a $30 h^{-1} \text{ Mpc}$ box, and an increase by ~ 50 per cent on doubling the box size to $60 h^{-1} \text{ Mpc}$ at a fixed spatial resolution of $\sim 60 h^{-1} \text{ kpc}$. Allowing for strong Jeans smoothing ($\eta = 2$) improves the convergence in the reduced χ^2 , but results in converged values differing from weak Jeans smoothing ($\eta = 4$) or minimal mesh scale smoothing by a factor of 2–4. We find that, while the concordance model would be very strongly rejected on the basis of moderate-resolution simulations using the minimal amount of mesh scale smoothing, the model is rejected at only the $\sim 2\sigma$ level in our highest-resolution simulations. By contrast, when Jeans smoothing is applied, the model is very strongly rejected at all resolutions. We have not explored variations in the equation of state that could potentially improve (or degrade) the agreement.

APPENDIX B: MEAN FLUX

The simulations leave unspecified the ionization level of the IGM. As a consequence, the flux normalization of the spectra is undetermined. We fix the normalization of the simulations by adjusting A in equation (3) to match the measured mean Ly α flux values as given by Fan et al. (2002) in their fig. 1 for $z \geq 4$, as tabulated in Table 3. For $z < 3.89$, we estimate the mean flux from a set of published Keck HIRES spectra, as described below.

The measurements of the mean flux have proven difficult, largely as a result of uncertainties in the QSO continuum across the absorbing region. Deriving error bars on the estimates is even more difficult because of difficulties in measuring the error on the continuum estimate, and because of flux correlations between pixels. (Press et al. 1993, hereafter PRS) used bootstrap resampling to estimate the errors on their fit to the mean Ly α optical depth $\bar{\tau}_\alpha$, obtaining $\bar{\tau}_\alpha = A(1+z)^{1+\gamma}$, with $\gamma = 2.46 \pm 0.37$ and $A = 0.0175 - 0.0056\gamma \pm 0.0002$. The interpretation of the error is not completely clear. Seljak et al. (2003) argue that, if the errors from PRS are interpreted as Gaussian, then the PRS fit results in both a bias and large error estimate for $\bar{\tau}_\alpha$. Seljak et al. find at $z = 2.72$, $\langle \bar{\tau}_\alpha \rangle = 0.29 \pm 0.14$, where $\langle \dots \rangle$ denotes an average over Gaussian distributions for A and γ , while using the mean values of A and γ would instead give $\bar{\tau}_\alpha = 0.35$. The error they quote, however, is the statistical rms on $\bar{\tau}_\alpha$,




Figure B1. The probability distribution of the mean absorbed Ly α flux at $z = 2.72$ according to the estimate of PRS. The long vertical tick mark indicates the average mean flux.

which carries little probabilistic meaning since the resulting probability distribution of $\bar{\tau}_\alpha$ is highly skewed. Moreover, they allowed for unphysical values of $\bar{\tau}_\alpha < 0$ in their averaging. Truncating the distribution to $\bar{\tau}_\alpha \geq 0$ yields an rms of 0.09 instead, which we compute from a Monte Carlo calculation. In terms of an equivalent Gaussian 1σ error spread, defined so that the probabilities of obtaining a value 1σ above or 1σ below the mean are each 0.16, the equivalent 1σ error spread in $\bar{\tau}_\alpha$ is $\langle \bar{\tau}_\alpha \rangle = 0.30^{+0.06}_{-0.05}$ (after truncating to $\bar{\tau}_\alpha \geq 0$).

Relating the mean optical depth to the mean absorbed flux and its error is not straightforward. Instead it is better to start with the distributions for A and γ and estimate $\langle \exp(-\bar{\tau}_\alpha) \rangle$ directly. The resulting probability distribution for the mean flux at $z = 2.72$ is shown in Fig. B1. The peak of the distribution is at $\exp(-\bar{\tau}_\alpha) = 0.71^{+0.03}_{-0.02}$, where the error bars refer to the range in the mean flux enclosing 68 per cent of the probability. The mean is $\langle \exp(-\bar{\tau}_\alpha) \rangle = 0.75 \pm 0.07$, where the error is the rms, and the distribution was truncated to $\bar{\tau}_\alpha \geq 0$. Adopting the 1σ equivalent probabilistic interpretation of the error, as above, gives instead a mean of $\langle \exp(-\bar{\tau}_\alpha) \rangle = 0.75^{+0.04}_{-0.05}$. Because of the skewed nature of the probability distribution, the estimate of the error for any quantity depending on the mean flux must take into account the full probability distribution of the mean flux.

For $z < 4$, we estimate the mean Ly α flux from the published Keck HIRES spectra used by Meiksin et al. (2001) for testing models of the Ly α forest. These spectra cover the redshift range $2.5 < z < 3.7$. We perform the estimate for each spectrum over a narrow redshift interval ($\Delta z = 0.1$ – 0.3) centred at each redshift considered, and average the results, adopting the error in the mean as our error estimate. We also allow for small systematic offsets in the continuum level (typically 1–4 per cent), as provided by Meiksin et al. for the two best-fitting models to the flux distribution (which are Λ CDM models similar to those considered in this paper). We estimate the error in the systematic offset from the spread between the best-fitting models (1–2 per cent of the continuum level), and add this linearly to the error in the mean. (Since the error is systematic rather than statistical, adding the error in quadrature would be inappropriate.)

For $z = 3.89$, we adopt the value of McDonald et al. (2000) of 0.48 ± 0.02 . For $z = 4.0$, we use the best fit of Songaila & Cowie (2002) to the measured redshift trend (their equation 19), to obtain a value for the mean flux of 0.47. We also linearly extrapolate the values of Songaila & Cowie (2002) at $\langle z \rangle = 4.34$ and 4.07 down to $z = 4.0$, obtaining a mean flux of 0.36 ± 0.03 . The difference appears

Table B1. Estimates for $\langle \exp(-\tau) \rangle$ using fits to $\bar{\tau}_\alpha$ from Press et al. (1993) (PRS), Zhang et al. (1997) as fitted to the D_A values of Steidel & Sargent (1987) (SS87), and Kim et al. (2001) (KCD). Also shown are the values adopted in this paper.

z	PRS	SS87	KCD	this paper
4.0	$0.46^{+0.06}_{-0.08}$	0.48 ± 0.05	0.49 ± 0.01	0.47 ± 0.03
3.89	$0.48^{+0.05}_{-0.08}$	0.51 ± 0.05	0.51 ± 0.01	0.48 ± 0.02
3.0	$0.69^{+0.04}_{-0.05}$	0.71 ± 0.03	0.71 ± 0.01	0.70 ± 0.02
2.75	$0.74^{+0.04}_{-0.05}$	0.76 ± 0.03	0.76 ± 0.01	0.74 ± 0.04
2.41	$0.80^{+0.03}_{-0.04}$	0.82 ± 0.02	0.82 ± 0.01	–

to be due to a single outlying measurement that substantially lowers the mean. We shall consider both values, noting that the larger one better matches the trend in Table B1.

For comparison, we also show the estimates, using the above procedure as applied to the fit to $\bar{\tau}_\alpha$ of PRS, based on similar estimates from Kim, Cristiani & D’Odorico (2001) with $A = 0.0144 - 0.00471\gamma$ and $\gamma = 2.43 \pm 0.17$ (which we infer from their error estimates for A and γ), and the fit by Zhang et al. (1997) to the flux decrement values (D_A) of Steidel & Sargent (1987), $A = 0.0028 \pm 0.0004$, $\gamma = 2.46$ (fixed). The errors provided correspond to the equivalent 1σ Gaussian errors, as described above.

The results are tabulated in Table B1. All the estimates are found to be in accord, including those of PRS. This is noteworthy in that the estimates that had been made using the best value for $\bar{\tau}_\alpha$ from PRS generally gave values significantly higher than measured using high-resolution spectra. Estimating the mean flux including the errors as we have done shows that these results are actually in agreement with the estimates from the higher-resolution data. It is also noteworthy that the estimates based on the moderate-resolution data of Steidel & Sargent (1987) are nearly identical to those using the $\bar{\tau}_\alpha$ fit of Kim et al. (2001), which are based on a compilation of several high-resolution spectra. The agreement in the determinations of the mean flux between the moderate- and high-resolution spectra suggests that moderate-resolution spectra, like those of the Sloan QSO survey, should, in principle, be adequate for measuring the mean flux. An attempt to measure the mean flux from the Sloan survey was made by Bernardi et al. (2003). The values found at $z = 2.41$, 2.75 , 3.0 and 3.89 are, respectively (as given in their fig. 21), $\langle \exp(-\tau) \rangle = 0.74 \pm 0.03$, 0.71 ± 0.01 , 0.65 ± 0.01 and 0.36 ± 0.02 . These all lie systematically lower than the values listed in Table B1. The reasons for the discrepancy are unclear.

Since the mean flux is the quantity of interest for normalizing the Ly α forest simulations, we use direct measurements of the mean flux to normalize our simulations rather than estimates of the mean Ly α optical depth. Historically, $\bar{\tau}_\alpha$ was of interest in the context of models for which the absorption was produced by discrete absorption clouds rather than modulations in a continuous medium. For the purposes of studies of the IGM under the current continuum model of its structure, $\bar{\tau}_\alpha$ has lost most of its significance; the mean flux is almost always to be preferred. An exception is in estimating the contribution of known discrete systems to the mean flux. Of particular interest is the contribution from damped Ly α absorbers (DLAs). The contribution of discrete absorption systems to $\bar{\tau}_\alpha$ is given by (e.g. Meiksin & Madau 1993; PRS)

$$\bar{\tau}_{\text{eff}} = \frac{1+z}{\lambda_\alpha} \int dN_{\text{HI}} \frac{\partial^2 N}{\partial N_{\text{HI}} \partial z} w(N_{\text{HI}}), \quad (\text{B1})$$

where $w(N_{\text{HI}})$ is the equivalent width corresponding to a system with H I column density N_{HI} , and λ_α is the (rest) wavelength of

the Ly α transition. Here, $\partial^2 N / (\partial N_{\text{HI}} \partial z)$ is the H I column density distribution at redshift z , assumed to vary slowly over the width of an absorption feature. The N_{HI} distribution is estimated to be a power law $\propto N_{\text{HI}}^{-\beta}$ with $\beta = 1.5$, with a possible high N_{HI} cut-off (Storrie-Lombardi, Irwin & McMahon 1996). Since DLAs are on the square-root part of the curve of growth ($w \propto N_{\text{HI}}^{1/2}$), $\bar{\tau}_{\text{DLA}}$ is logarithmically divergent at the upper end. Even if the distribution steepens at the upper end, as suggested by Storrie-Lombardi et al., the statis-

tics are so poor that the contribution of DLAs to $\bar{\tau}_\alpha$ is highly uncertain. Seljak et al. (2003) claim to have adjusted their error estimate of $\bar{\tau}_\alpha$ to allow for DLAs, although they give no details of their procedure. In view of the uncertainty in $\bar{\tau}_{\text{DLA}}$, the most reliable approach is to use only spectral regions that are clean of DLA contamination to estimate the mean flux.

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