COSMOLOGICAL IMPLICATIONS OF THE LYMAN ALPHA FOREST

In the previous lecture we have learnt how to 'translate' the shape of the absorption lines in the Lyα forest into physical parameters for the clouds producing them, parameters such as density, temperature, size and mass. In this lecture we are going to draw more broad-ranging conclusions about the intergalactic medium by analysing the distributions of these parameters and their evolution with redshift.

13.1 Absorption Distance

It is often the case that we wish to interpret some QSO absorption line statistic, for example the number of absorption lines per unit redshift interval, in terms of a more meaningful quantity, such as number of absorbers per unit volume.

Consider an ideal case where the absorbers (e.g. galaxies, or Lyα clouds) have a number density $n_{\rm a}$ (Mpc^{-3}) and present a cross section $\sigma_{\rm a}$ (Mpc^2). Then, the total number of absorber intersected along a unit *proper length* is $dN/dl_{\rm p} = n_{\rm a} \cdot \sigma_{\rm a}$ (Mpc^{-1}). But what we measure is the number of absorbers per unit redshift interval, which is related to the proper line element by:

$$\frac{dl_{\rm p}}{dz} = \frac{c}{H(z)\,(1+z)} \tag{13.1}$$

where $H(z)$ is given by our familiar expression:

$$H(z) = H_0 \left[\Omega_{\rm m,0} \cdot (1+z)^3 + \Omega_{\rm k,0} \cdot (1+z)^2 + \Omega_{\Lambda,0}\right]^{1/2} . \tag{13.2}$$

From the above it can be seen that, even if $n_{\rm a}$ and $\sigma_{\rm a}$ do not change with redshift, the number of absorbers per unit redshift interval, $d\mathcal{N}/dz$, will increase with increasing redshift as a result of the expansion of the Universe. Thus, if we want to test from our data whether the intrinsic properties of the absorbers evolve with cosmic time, we need to decouple

the cosmic expansion from the intrinsic evolution. This is achieved by introducing the *'absorption distance'*:

$$dX = \frac{H_0}{H(z)} \, (1+z)^2 \, dz \tag{13.3}$$

which has the property that absorbers with non-evolving $n_{\rm a}$ and $\sigma_{\rm a}$ will produce a constant number of absorption lines per unit absorption distance, i.e. $d\mathcal{N}/dX$ is constant for such absorbers.

13.2 H I Column Density Distribution

In the previous lecture we saw that it is possible to associate a column density of neutral hydrogen, $N_{\rm H\,I}$, with each Lyα cloud. Given the multitude of absorption lines in the Lyα forest, and the large number of quasar spectra now recorded at high resolution, you will not be surprised to learn that the statistics of the Lyα forest have been established with good precision. The frequency distribution of H I absorbers:

$$f(N_{\rm H\,I}, X) = \frac{\partial^2 \mathcal{N}}{\partial N_{\rm H\,I} \, \partial X} \tag{13.4}$$

describes the number of absorbers per unit absorption distance with column densities between $N_{\rm H\,I}$ and $N_{\rm H\,I} + dN_{\rm H\,I}$.

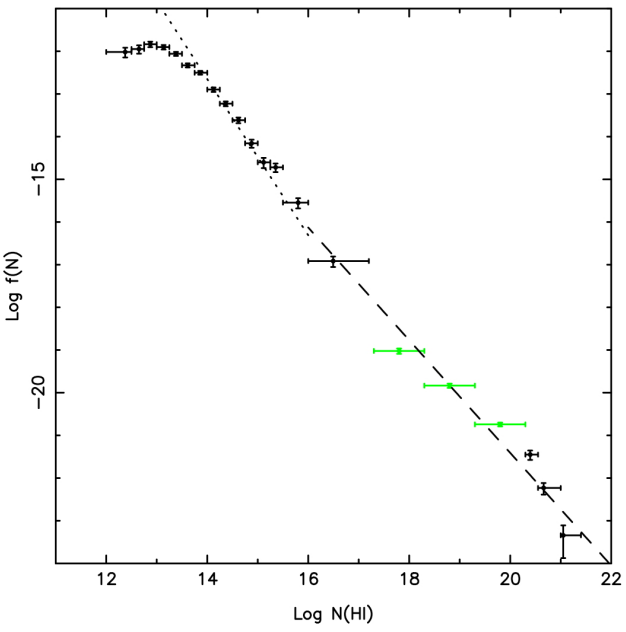

Figure 13.1: The distribution of H I column densities at $z \simeq 3$ from Peroux et al. (2003).

$f(N_{\rm H\,I}, X)$ spans ten orders of magnitude, from $N_{\rm H\,I} = 10^{12}\,{\rm cm}^{-2}$ to $10^{22}\,{\rm cm}^{-2}$ (see Figure 13.1). At present, these limits are entirely observational: Lyα absorbers with $N_{\rm H\,I} < 10^{12}\,{\rm cm}^{-2}$ undoubtedly exist but produce absorption lines too weak to be recognised in most spectra, while absorbers with $N_{\rm H\,I} > 10^{22}\,{\rm cm}^{-2}$ are so rare that even with the current large samples their statistics are poorly determined. Over the entire range of ten orders of magnitude, $f(N_{\rm H\,I}, X)$ can be approximately described by a single power law:

$$f(N_{\rm H\,I}, X) = B \cdot N_{\rm H\,I}^{-\beta} \tag{13.5}$$

with exponent $\beta \simeq 1.6$.

Fom the integral of $f(N_{\rm H\,I}, X)\, dN_{\rm H\,I}$ we can infer the total mass of neutral gas in the Universe expressed, as usual, as fraction of the critical density, $\rho_{\rm crit}$:

$$\Omega_{\rm H\,I+He\,I} = \frac{1}{\rho_{\rm crit}} \frac{H_0}{c} \frac{m_{\rm H}}{1-Y} \int N_{\rm H\,I} f(N_{\rm H\,I}, X)\, dN_{\rm H\,I} \tag{13.6}$$

where $m_{\rm H}$ is the mass of the hydrogen atom, and the factor of $(1-Y)$ in the denominator accounts for the primordial abundance of helium ($Y \simeq 0.25$ by mass).

Entering the numerical values and with eq. 13.5 for $f(N_{\rm H\,I}, X)$, we have:

$$\Omega_{\rm H\,I+He\,I} = 1.8 \times 10^{-23} \int B N_{\rm H\,I}^{-\beta+1}\, dN_{\rm H\,I}\,. \tag{13.7}$$

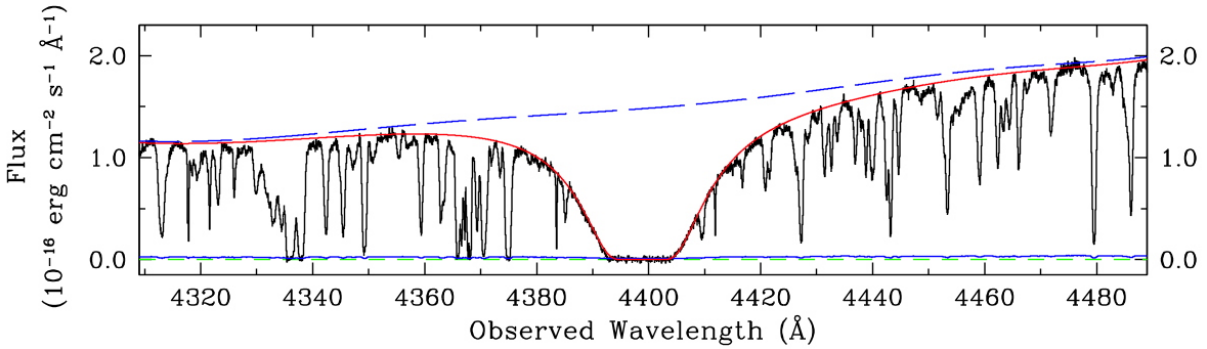

Figure 13.2: Portion of the high resolution spectrum of the QSO Q0913+072 recorded with the Ultraviolet and Visual Echelle Spectrograph on the Very Large Telescope of the European Southern Observatory. Clearly seen near $\lambda_{\rm obs} = 4000$ Å is a damped Lyα system with neutral hydrogen column density $N({\rm H\,I}) = 2.2 \times 10^{20}\,{\rm cm}^{-2}$ (red continuous line).

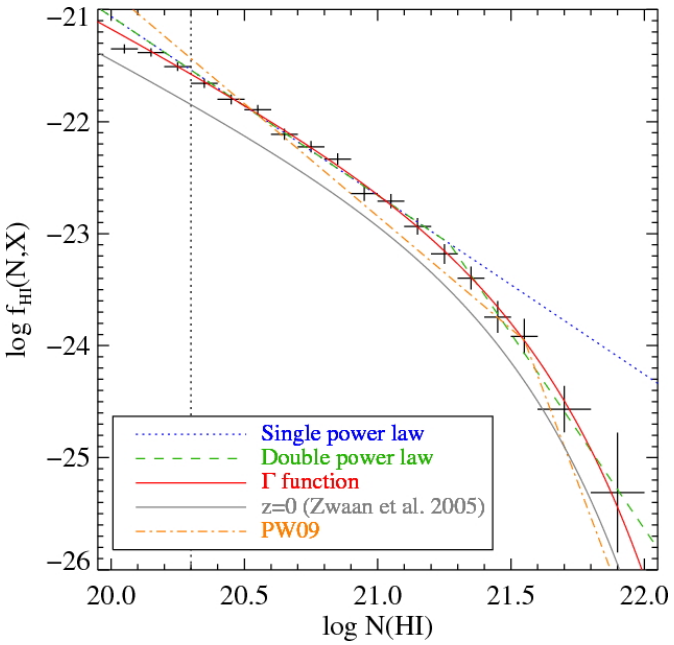

Figure 13.3: The distribution of H I column densities for DLAs, the absorbers with the largest values of $N(\mathrm{H\,I})$. The authors of this study (Noterdaeme et al. 2009) identified ~ 1000 DLAs with $z = 2.15 - 5.2$ from low-resolution spectra from the Sloan Digital Sky Survey. The single power-law has slope $\beta = 1.6$ for $\log[N(\mathrm{H\,I})/\mathrm{cm}^{-2}] \leq 21.3$.

Integrating, we obtain:

$$\Omega_{\mathrm{H\,I+He\,I}} = 1.8 \times 10^{-23} \frac{B}{2-\beta} \left[N_{\mathrm{H\,I}}^{2-\beta} \right]_{N_{\mathrm{H\,I\,min}}}^{N_{\mathrm{H\,I\,max}}} \tag{13.8}$$

For $\beta < 2$ the integral diverges at the high $N_{\mathrm{H\,I}}$ end. Now we see that it is the rarest, highest column density absorbers, the 'Damped Lyα Systems' (often referred to as DLAs) which dominate the mass budget of neutral gas in the Universe. Figure 13.2 shows an example of a damped Lyα system. Note that we do not need to know anything about the size, shapes, mass distribution (massive galaxies or just small, dense H I clouds?) of the absorbers in order to calculate the contribution they make to $\Omega_{\mathrm{H\,I}}$.

Figure 13.3 shows a recent determination of $f(N_{\mathrm{H\,I}}, X)$ for DLAs. A single power low with slope $\beta = 1.6$ is a good approximation for column densities $N(\mathrm{H\,I}) \leq 10^{21.3}\,\mathrm{cm}^{-2}$; absorption systems with larger values of $N(\mathrm{H\,I})$ become progressively rarer [$\beta = -3.5$ for $N(\mathrm{H\,I}) > 10^{21.3}\,\mathrm{cm}^{-2}$]. If, for simplicity, we integrate the single power law up to $N(\mathrm{H\,I}) = 10^{21.7}\,\mathrm{cm}^{-2}$, we obtain (from eq. 13.8):

$$\boldsymbol{\Omega_{\mathrm{H\,I+He\,I}}\,(z \simeq 2-4) \simeq 1 \times 10^{-3}}$$

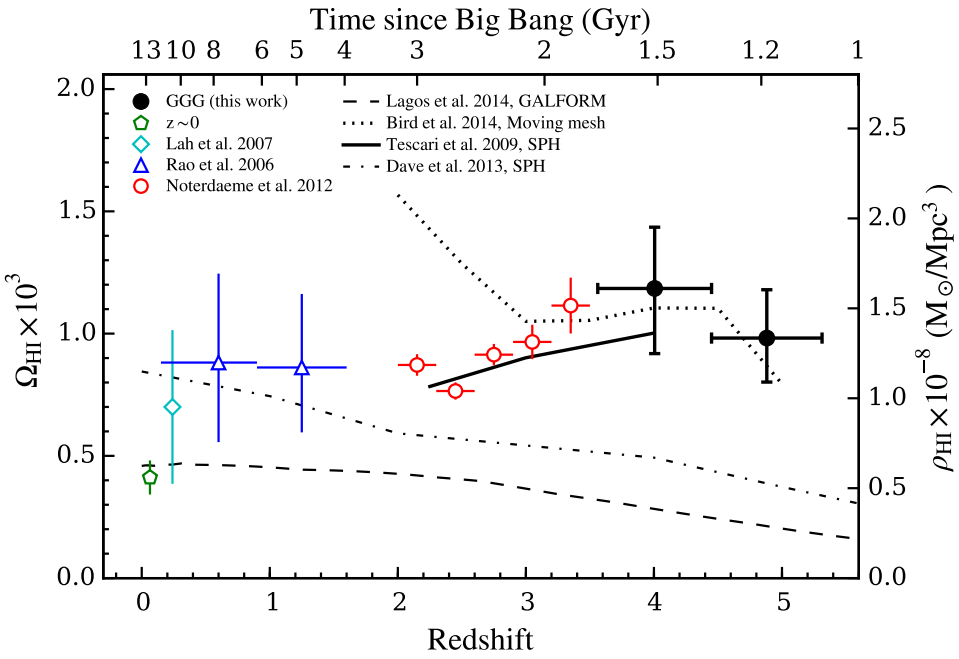

Figure 13.4: Estimates of $\Omega_{\rm H\,I}$ from $z = 0$ to 5 (reproduced from Crighton et al. 2015).

This value is comparable to, but somewhat lower than, $\Omega_{\rm stars} \sim 0.003$ (see Table 1.1). Determinations of $\Omega_{\rm H\,I}$ at lower and higher redshifts are more uncertain. To a first approximation, it appears that $\Omega_{\rm H\,I} \sim 1 \times 10^{-3}$ over most of the range of cosmic time probed (see Figure 13.4).

13.3 The Low-Density Lyα Forest: Repository of Most of the Baryons

The gas in DLAs is almost entirely neutral: when $N({\rm H\,I}) > 10^{20}\,{\rm cm}^{-2}$ the cloud is self-shielded from photons with sufficient energy to ionise H (i.e. photons with energy $h\nu > 13.6\,{\rm eV}$, corresponding to wavelengths $\lambda < 912\,$Å).

The converse is true at the other end of the H I column density distribution, where *neutral* hydrogen is a minor component: *most of the gas traced by low-column density Lyα forest lines is H II* (and therefore does not produce absorption lines).

The ratio between the densities of two successive ionisation stages (in this case $n_{\rm H\,I}$ and $n_{\rm H\,II}$), is given by the balance between number of ionisations and the number of recombinations per unit volume per unit time. In the

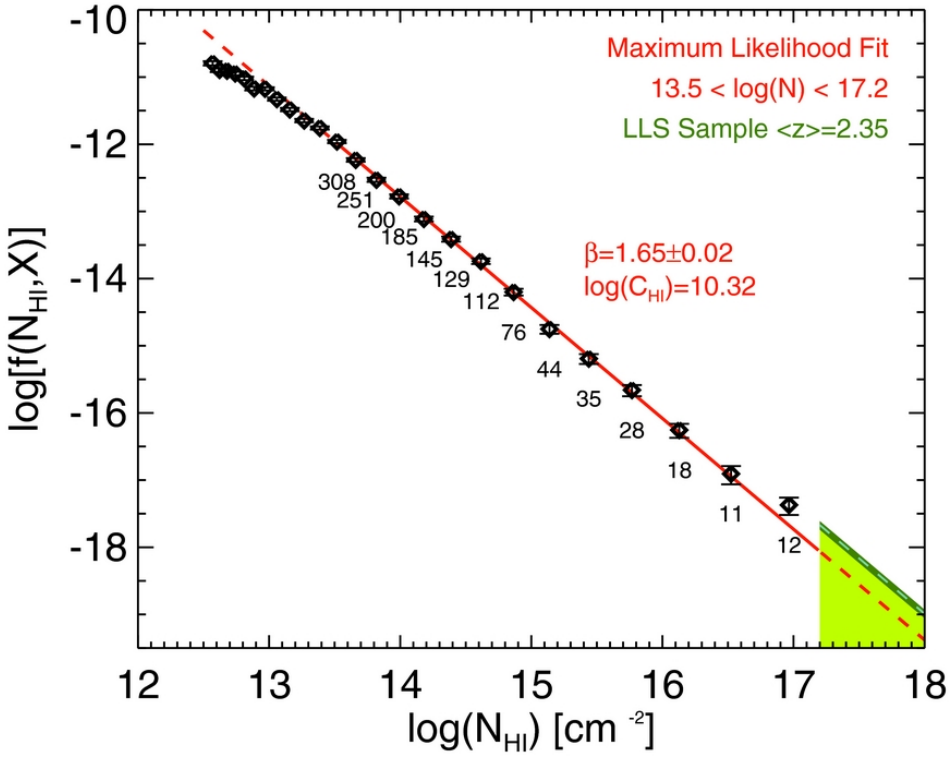

Figure 13.5: The distribution of H I column densities in the Lyα forest (reproduced from Rudie et al. 2013).

intergalactic medium we only need to consider two processes: photoionisation ($H^0 + h\nu \rightarrow H^+ + e^-$) and the reverse two-body radiative recombination. Thus:

$$\frac{n_{\rm H\,II}}{n_{\rm H\,I}} = \frac{\Gamma}{n_{\rm e}\,\alpha_{\rm H\,I}} \tag{13.9}$$

where Γ (s^{-1}) is the photoionization rate, $\alpha_{\rm H\,I} = f(T)$ ($cm^3\ s^{-1}$) is the recombination coefficient and $n_{\rm e} = n_{\rm H\,II} \simeq n_{\rm H}$ is the number density of free electrons. Entering values appropriate to the IGM at $z \sim 3$, we can rewrite eq. 13.9 as:

$$\frac{n_{\rm H\,I}}{n_{\rm H}} \simeq 0.46\, n_{\rm H}\, T_4^{-0.76}\, \Gamma_{12}^{-1} \tag{13.10}$$

where T_4 is the temperature in units of 10^4 K, and Γ_{12} is the photoionisation rate in units of $10^{-12}\,s^{-1}$. Note the following:

1. The density of neutral H, $n_{\rm H\,I}$, is proportional to the *square* of the gas density, $n_{\rm H}$.

2. In the previous lecture, we were given a scaling between $n_{\rm H}$ and $N({\rm H\,I})$:

$$N_{\rm H\,I} \simeq 2.3 \times 10^{13}\,{\rm cm}^{-2}\; T_4^{-0.26}\; \Gamma_{12}^{-1}\; \left(\frac{n_{\rm H}}{10^{-5}\,{\rm cm}^{-3}}\right)^{3/2} \left(\frac{f_{\rm g}}{0.16}\right)^{1/2} \qquad (13.11)$$

Combining 13.10 and 13.11, we now see that, for fiducial values of T_4, Γ_{12} and $f_{\rm g}$, Lyα clouds with $N({\rm H\,I}) \simeq 2 \times 10^{13}\,{\rm cm}^{-2}$ have a trifling neutral fraction $n_{\rm H\,I}/n_{\rm H} \simeq 5 \times 10^{-6}$. Even clouds with $N({\rm H\,I}) \simeq 2 \times 10^{16}\,{\rm cm}^{-2}$ are highly ionised: $n_{\rm H\,I}/n_{\rm H} \simeq 5 \times 10^{-4}$. Thus:

The gas producing the absorption lines in the Lyα forest is only a trace of the total density of gas present.

From the above it follows that, in order to calculate the total mass density in the Lyα forest, we need to include an additional term $n_{\rm H}/n_{\rm H\,I}$ in eq.13.6:

$$\Omega_{\rm gas} \equiv \Omega_{\rm H\,I} + \Omega_{\rm H\,II} = \frac{1}{\rho_{\rm crit}}\,\frac{H_0}{c}\,\frac{m_{\rm H}}{1-Y} \int N_{\rm H\,I}\,\frac{n_{\rm H}}{n_{\rm H\,I}}\,f(N_{\rm H\,I}, X)\,dN_{\rm H\,I} \qquad (13.12)$$

Using 13.10 and 13.11 to express $n_{\rm H}/n_{\rm H\,I}$ in terms of $N_{\rm H\,I}$ (with $f_{\rm g} = 0.16$ for simplicity), and entering the numerical values as before, we find:

$$\Omega_{\rm gas} = 3.1 \times 10^{-9}\,T_4^{0.59}\,\Gamma_{12}^{1/3} \int N_{\rm H\,I}^{1/3}\,f(N_{\rm H\,I}, X)\,dN_{\rm H\,I} \qquad (13.13)$$

If $f(N_{\rm H\,I}, X)$ is a power law with exponent β (eq. 13.5), we have:

$$\Omega_{\rm gas} = 3.1 \times 10^{-9}\,T_4^{0.59}\,\Gamma_{12}^{1/3}\,\frac{B}{4/3-\beta} \cdot \left[N_{\rm H\,I}^{4/3-\beta}\right]_{N_{\rm H\,I\,min}}^{N_{\rm H\,I\,max}} \qquad (13.14)$$

from which it can be seen that for $\beta > 4/3$ the integral diverges at the low $N_{\rm H\,I}$ end. Since $\beta \simeq 1.6$ (see Figure 13.5), we conclude that the largest contribution to $\Omega_{\rm gas}$ is from the lowest $N({\rm H\,I})$ Lyα clouds. Choosing $\log[N_{\rm H\,I}/{\rm cm}^{-2}] = 12-17$ as the integration limits (somewhat arbitrary, but $N_{\rm H\,I\,max}$ makes little difference and by $N_{\rm H\,I\,min} = 10^{12}\,{\rm cm}^{-2}$ the counts seem to turn over), and with $T_4, \Gamma_{12} = 1$, we find:

$$\boldsymbol{\Omega}_{\rm gas}\,(\log[N_{\rm H\,I}/{\rm cm}^{-2}] = 12-17,\; z \simeq 2-4) \simeq \mathbf{3} \times \mathbf{10^{-2}}$$

This value is comparable to $\Omega_{\rm b} \sim 0.05$ (see Table 1.1). Evidently, at $z \sim 3$ the Lyα forest accounts for a significant fraction of the baryons in the Universe.

13.4 The Proximity Effect and Metagalactic Ionising Background

The Lyα forest shows an obvious redshift evolution, in the sense that its opacity increases with redshift (Fig. 13.6). Qualitatively, this effect is readily apparent when we compare the Lyα forest in QSOs of increasing $z_{\rm em}$, as in Figure 13.7. The increase in opacity with redshift far exceeds that expected from cosmological effects alone, indicating an *intrinsic* evolution of the absorbers. This evolution is well understood in the context of structure formation: more of the volume of the Universe is at a density below the universal mean as time progresses (and z decreases; cfr. eq. 12.18), as demonstrated by the simulations reproduced in Figure 13.8. The changing intensity of the photoionising background radiation field (which we are going to discuss presently) also plays a role via the parameter Γ_{12} in eq. 12.18.

Even though the increase of Lyα line density with redshift is so obvious, it was not recognised at first because of a subtle complication that can mask the evolution if the redshift baseline is limited. The 'problem' is that we need quasars to probe the IGM. As we follow the Lyα forest to higher and higher redshift within the spectrum of a given QSO, the line density increases with z (as we have seen), up until a point where we

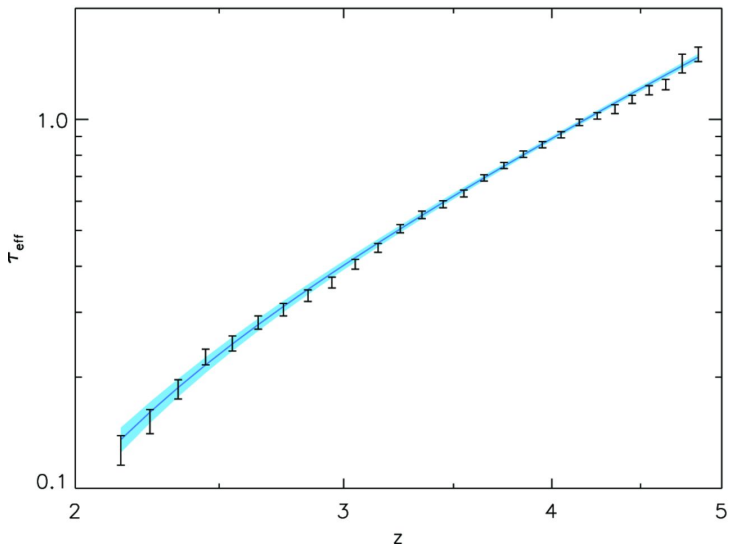

Figure 13.6: The evolving optical depth of the Lyα forest (reproduced from Becker et al. 2013).

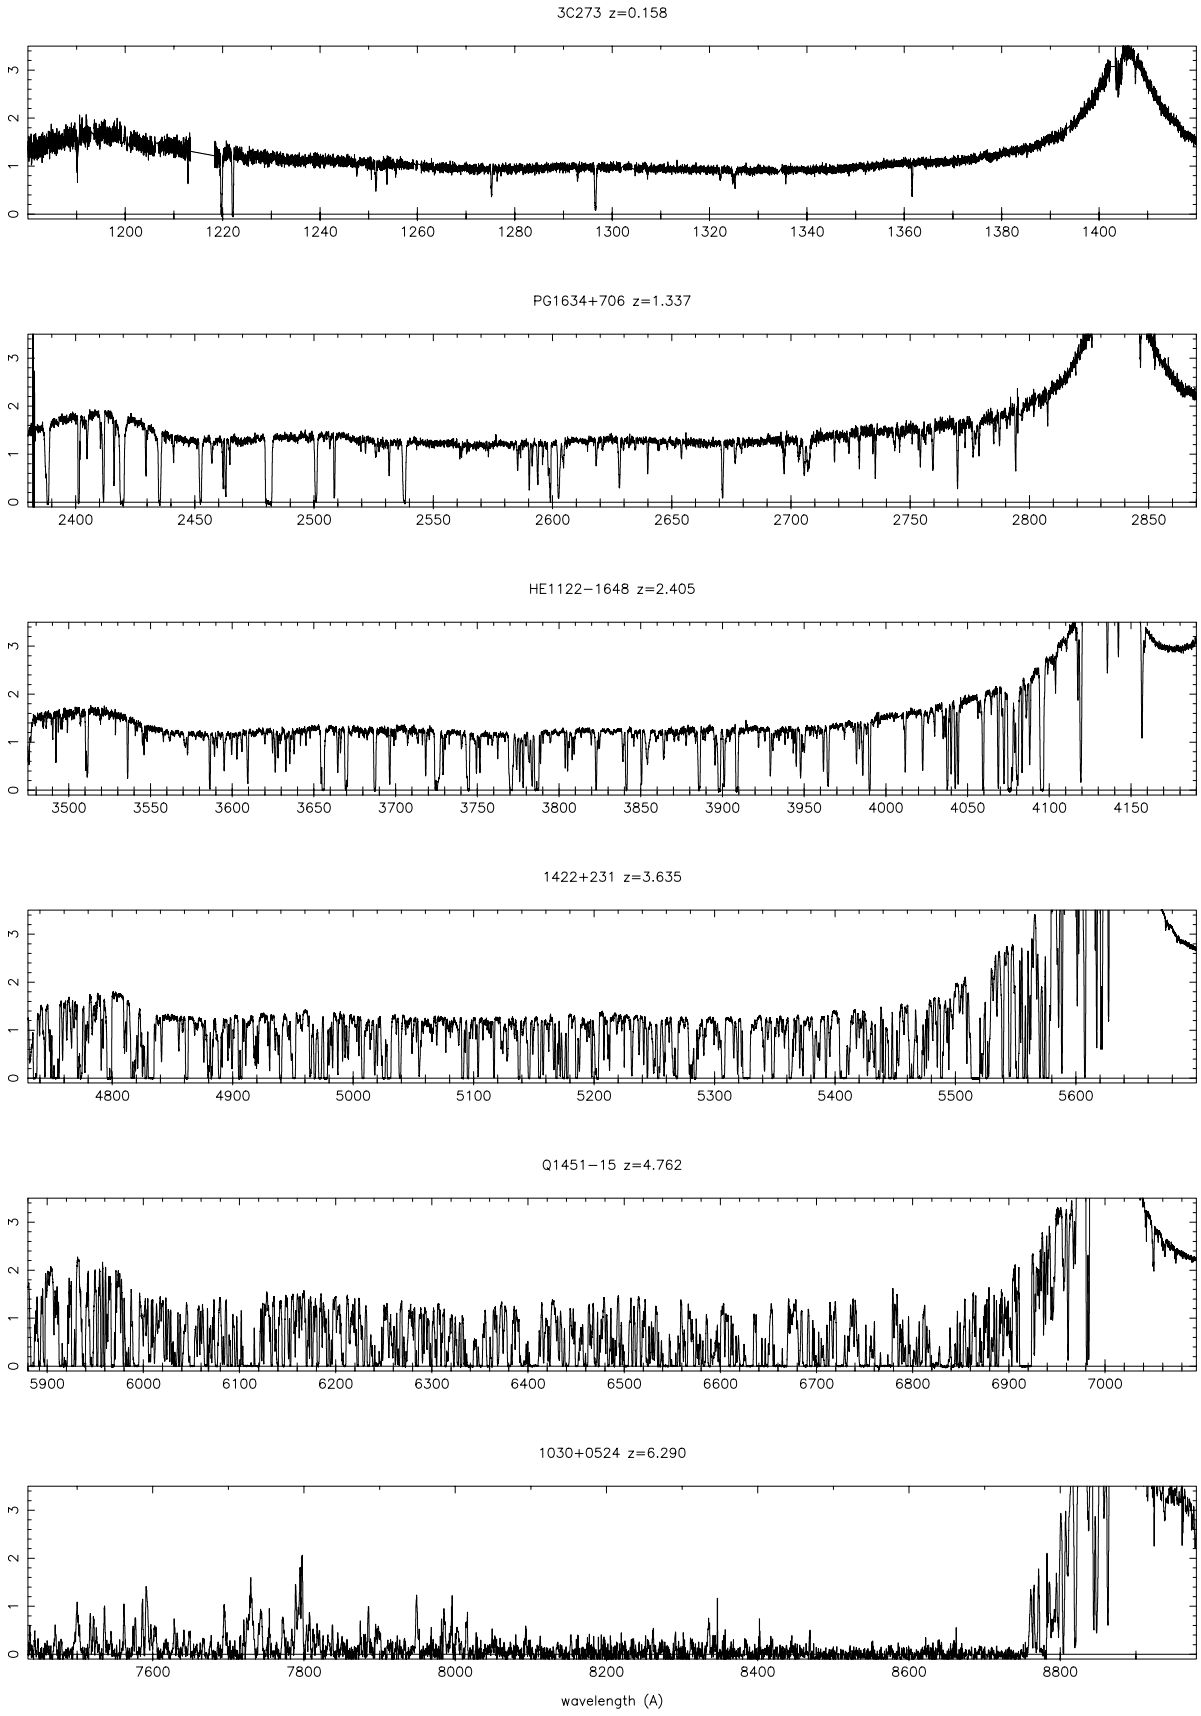

Figure 13.7: The density of absorption lines in the Lyα forest increases dramatically with increasing redshift. At $z > 6$ (bottom panel), the line density is so high that few transmission gaps, if any, can be found shortwards of the QSO Lyα emission line.

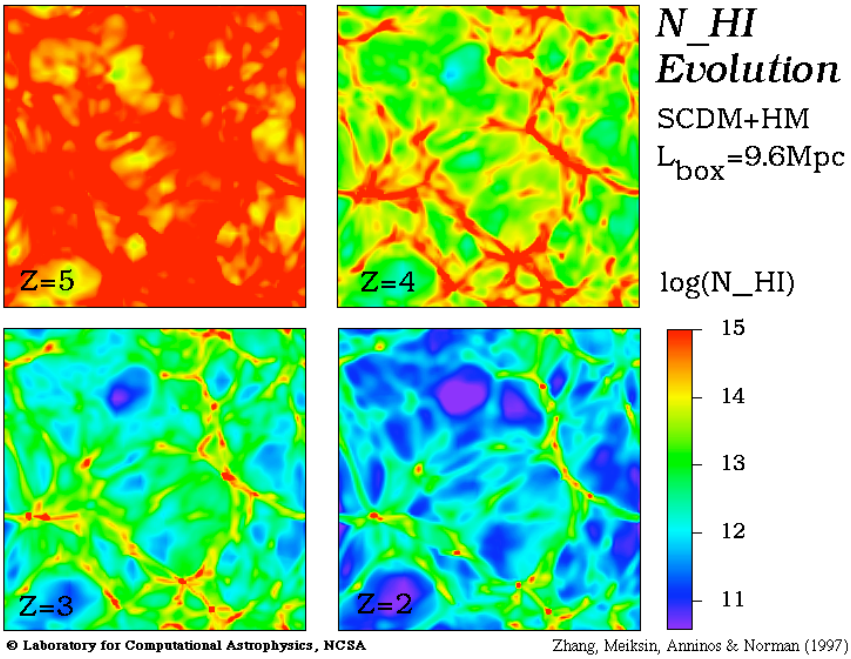

Figure 13.8: Hydrodynamical simulations of the growth of structure in the universe offer a natural explanation for the marked evolution in the density of absorption lines in the Lyα forest seen in the spectra of QSOs at different redshifts, as in Fig. 13.7.

get sufficiently close to the QSO itself that its radiation field is no longer insignificant relative to the ionising background. Within this QSO sphere of influence, the Lyα clouds are over-ionised, and this leads to an 'inverse' effect that weakens and eventually reverses the trend of increasing line density with redshift.

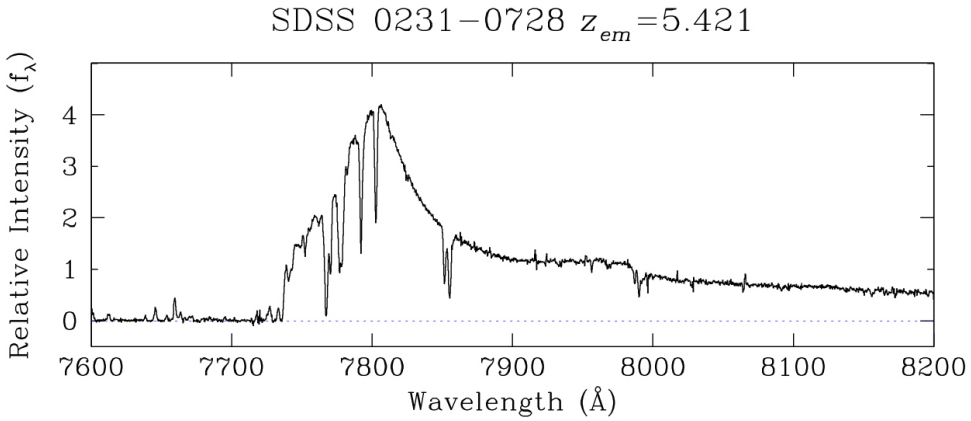

Figure 13.9: Keck spectrum of a $z_{\rm em} = 5.421$ QSO discovered by the *Sloan Digital Sky Survey*. At these high redshifts the line density in the Lyα forest is so high that all the lines merge together effectively extinguishing the QSO light. Only in the immediate vicinity of the QSO, where the QSO radiation field dominates over the ionising background, does the forest become sufficiently sparse again for individual Lyα absorption lines to be distinguished. This is the 'proximity' effect.

This 'proximity' effect, as it came to be known, gives us the means to measure the intensity of the ionising background, which is an important parameter, provided we know the systemic redshifts and the luminosities at the Lyman limit of the QSOs against which it is seen. Qualitatively, the effect can be recognised even 'by eye' in the spectra of the highest redshift QSOs, where the Lyα forest is so thick that it almost completely extinguishes the QSO light at wavelengths shorter than that of the Lyα emission line—see Fig. 13.9.

Quantitatively, we can compare the optical depth τ at any point along the line of sight to a quasar to the average value at that redshift, $\tau_{\rm forest}$ (i.e. the value that would be measured at the same location if the quasar were not there):

$$\tau = \tau_{\rm forest}\,[1+\omega(r)]^{-1} \tag{13.15}$$

where r is the proper distance from the quasar given by eq. 13.1, so that:

$$r \simeq \frac{c}{H(z)\,(1+z)}\,\Delta z \tag{13.16}$$

for a redshift difference $\Delta z = z_{\rm em} - z_{\rm abs}$. The parameter ω is defined as:

$$\omega(r) = \frac{\Gamma_{\rm QSO}(r)}{\Gamma_{\rm bkg}} \tag{13.17}$$

the ratio of the H I photoionisation rate due to photons impinging on a cloud from the quasar against which the Lyα forest is seen, to the photoionisation rate due to the uniform background of UV sources. It can be

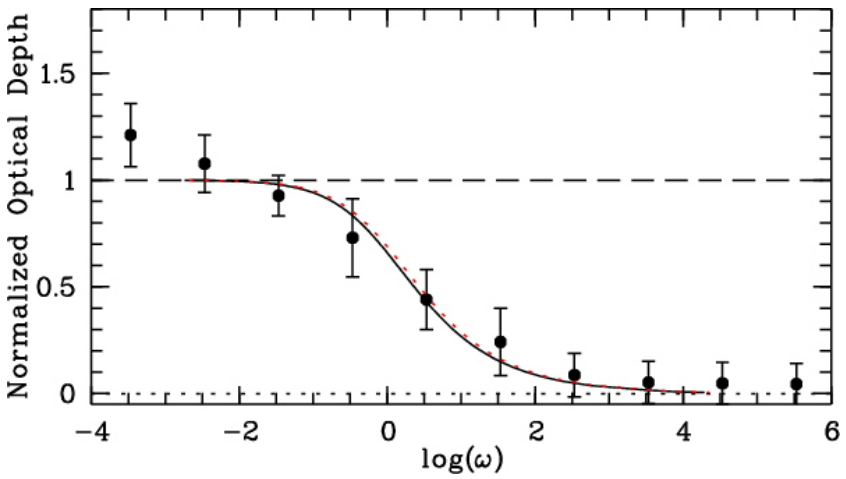

Figure 13.10: Normalised Lyα forest optical depth vs. ω (defined in eq. 13.17) for a combined sample of 40 QSO sightlines analysed by Dall'Aglio et al. (2008). The horizontal long-dash line illustrates the case of no proximity effect.

seen from eq. 13.15 that we recover $\tau_{\rm forest}$ at large values of r, while $\tau \to 0$ very close to the QSO, where $\Gamma_{\rm QSO} \gg \Gamma_{\rm bkg}$ (see Figure 13.10).

If we define $R_{\rm eq}$ as the distance from the quasar where $\Gamma_{\rm QSO}(R_{\rm eq}) = \Gamma_{\rm bkg}$, we can rewrite eq.13.15 as:

$$\tau = \tau_{\rm forest} \left[1 + \left(\frac{r}{R_{\rm eq}} \right)^{-2} \right]^{-1} \tag{13.18}$$

Thus, by measuring $R_{\rm eq}$ we can deduce the intensity of the UV background if the flux of ionising photons from the quasar is known, assuming a $1/r^2$ fall-off:

$$F_{\nu_0}^{\rm QSO}(r) = \frac{L_{\nu_0}^{\rm QSO}}{4\pi r^2} \tag{13.19}$$

where ν_0 is the frequency of the Lyman limit.

Typically, $R_{\rm eq} \sim$ a few Mpc. There are a number of caveats to this method: L_{ν_0} is determined from an extrapolation of the QSO spectral slope at longer wavelengths, and in any case may not be constant (most QSOs vary); $z_{\rm em}^{\rm QSO}$ may not be known accurately (not all emission lines give the same $z_{\rm em}$); the environments of QSOs may be overdense compared to more typical volumes

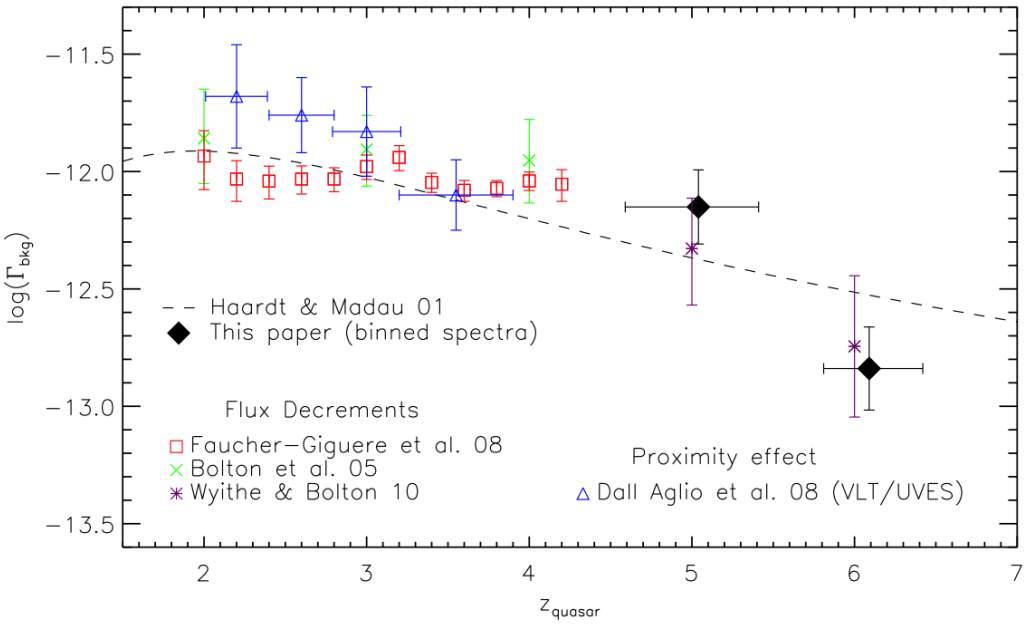

Figure 13.11: Redshift evolution of the H I photoionisation rate from the UV metagalactic background (reproduced from Calverley et al. 2011).

of the IGM. All of these effects have been modelled in an attempt to estimate the magnitude of the bias they may introduce in the determination of $\Gamma_{\rm bkg}$.

13.4.1 Redshift Evolution of the H I Ionising Background

Figure 13.11 shows recent determinations of $\Gamma_{\rm bkg}$. At $z = 2-4$, $\Gamma_{\rm bkg} \simeq 1 \times 10^{12}\,{\rm s}^{-1}$ (that's why in Lecture 12 we scaled the physical properties of the Lyα clouds relative to $\Gamma_{12} = 1$).

By $z = 0$, the intensity of the UV background has decreased by a factor of ~ 20. A recent determination of $\Gamma_{\rm bkg}$ between $z = 0$ and 0.5 finds $\Gamma_{\rm bkg}(z) = (4.6 \times 10^{-14}\,{\rm s}^{-1})\,(1+z)^{4.4}$ (Shull et al. 2015). This steep decrease of $\Gamma_{\rm bkg}$ with decreasing redshift from $z = 2$ reflects the steep fall-off with cosmic time in the comoving density of bright quasars, which are presumed to be the dominant source of H ionising photons (see Figure 13.12, left panel).

The density of bright QSOs also falls at $z > 3$. Thus, the epoch $z \sim 2$–3 seems to be the time when both the cosmic star formation rate and the accretion rate to massive black holes (that fuels the QSO luminosity) were at their peaks. However, the decrease in the comoving density of bright

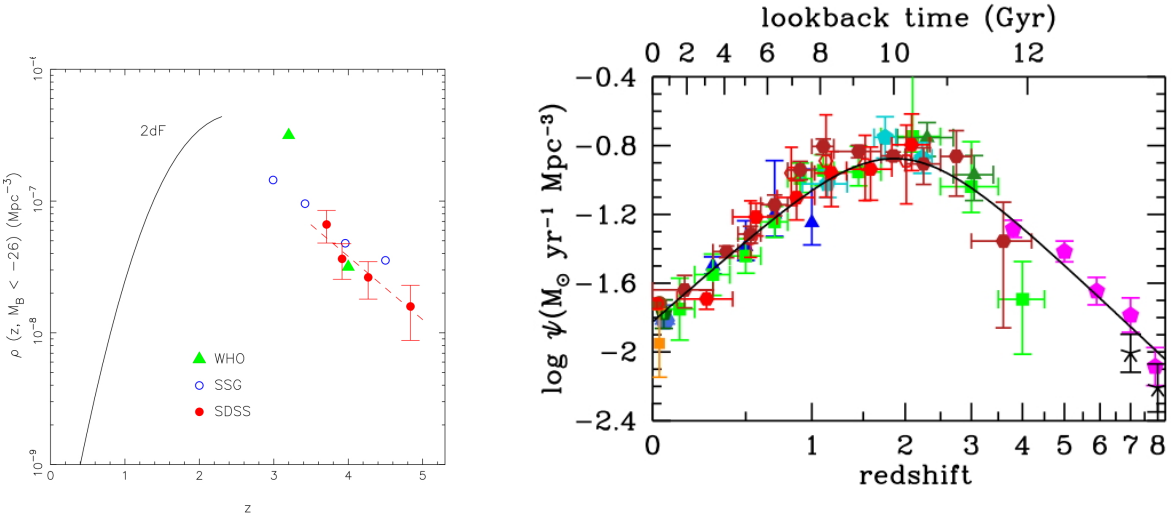

Figure 13.12: *Left:* Redshift evolution of the comoving density of bright QSOs, (reproduced from Fan et al. 2001). *Right:* Redshift evolution of the cosmic star formation rate density ($M_\odot\,{\rm yr}^{-1}\,{\rm Mpc}^{-3}$) (reproduced from Madau & Dickinson 2014).

QSOs at $z > 3$ is much steeper than the decrease in the cosmic star formation rate (see Figure 13.12). Extrapolating to $z \sim 10$, this would imply that *stars*, rather than Active Galactic Nuclei, were primarily responsible for reionising the Universe. Detail consideration of the UV photon budget finds that massive stars do produce sufficient Lyman continuum photons to ionise the IGM; the issue is what fraction of these photons manage to escape into the intergalactic medium, rather than being absorbed in the gas-rich environments where the stars are born. This is still an open question.

The ionisation potential of He^+ is $4\times$ greater than that of H at 54.4 eV; corresponding to photon wavelengths $\lambda < 228$ Å. Even the most massive stars do not produce many photons at such short wavelengths. To a first approximation, stars emit as blackbodies and the hottest stars known (on the main sequence) have effective temperatures $T_{\rm eff} \simeq 50\,000$ K. Their blackbody emission peaks at $\lambda = 580$ Å and falls off exponentially at shorter wavelengths (Wien's law).

On the other hand, the spectra of AGN are well approximated by power laws: for a typical QSO, $f_\nu \propto \nu^{-1.75}$ shortwards of 1200 Å [usually referred to as the extreme ultraviolet (EUV) wavelength range]. Thus, it is thought that He in the IGM became full ionised at later epochs than H, between $z \sim 4$ and 3, when the comoving density of QSOs peaked (Figure 13.12).

13.5 Summary: What We Have Learnt from the Lyman Alpha Forest

- The absorption spectra of quasars are a window on the intergalactic medium, a constituent of the Universe that holds most of the baryons and yet is unobservable by any other means (at least with current technologies).

- The 'clouds' giving rise to the Lyα forest are large structures, $L \simeq 100\,\mathrm{kpc}\,(N_{\mathrm{H\,I}}/10^{14}\,\mathrm{cm}^{-2})^{-1/3}$, consisting of gas of moderate overdensity (compared to the cosmic mean), $(1+\delta) \sim (N_{\mathrm{H\,I}}/3{\times}10^{13}\,\mathrm{cm}^{-2})^{2/3}$, and with masses comparable to today's dwarf galaxies, $M \sim 10^9 M_\odot\,(N_{\mathrm{H\,I}}/10^{14}\,\mathrm{cm}^{-2})^{-1/3}$. These scalings apply at $z \sim 3$.

- The H I gas producing the absorption lines in the Lyα forest is only a trace of the total amount of gas present, because the clouds are highly ionised by the metagalactic UV background: $n_{\mathrm{H\,I}}/n_{\mathrm{H}} \sim 5 \times 10^{-6}\,(N_{\mathrm{H\,I}}/2 \times 10^{13}\,\mathrm{cm}^{-2})^{2/3}$ (again at $z \sim 3$).

- The frequency distribution of H I column densities spans ten orders of magnitude, from $N_{\mathrm{H\,I}} = 10^{12}\,\mathrm{cm}^{-2}$ to $10^{22}\,\mathrm{cm}^{-2}$; over this range it can be approximated by a single power law: $f(N_{\mathrm{H\,I}}, X) \propto N_{\mathrm{H\,I}}^{-1.6}$, where X is the absorption distance. The absorbers with the highest H I column densities, the DLAs, account for most of the *neutral gas*, but it is the absorbers with the lowest column densities that account for most of the gas (neutral + ionised).

  $\Omega_{\mathrm{DLA}} \simeq 1 \times 10^{-3}$, of the same order as the fraction of baryons in stars and stellar remnants today, $\Omega_{\mathrm{stars}} \sim 0.003$.

  $\Omega_{\mathrm{H\,I+H\,II}} \sim 0.03$; for comparison $\Omega_{\mathrm{baryons}} \simeq 0.05$.

- The density of Lyα lines per unit redshift exhibits a marked redshift evolution, in the sense that opacity of the Lyα forest increases with redshift. This evolution reflects the fact that larger volumes of the Universe are at densities below the cosmic mean as time progresses (and therefore do not produce detectable Lyα absorption lines).

- In the proximity of QSOs, Lyα clouds are *over*ionised compared to an average location in the intergalactic medium. By measuring the extent of this QSO 'sphere of influence', it is possible to estimate the

intensity of the metagalactic UV background. The intensity of the background was highest at redshift $z \sim 2$–3, when the space density of QSOs was at its maximum. Moving to higher redshifts, the intensity of the background falls off more slowly than QSO numbers, suggesting that the first generations of *stars*, rather than Active Galactic Nuclei, were responsible for reionising the Universe at $z \sim 10$.