The primordial deuterium abundance at $z_{abs} = 2.504$ from a high signal-to-noise spectrum of Q1009+2956

E. O. Zavarygin,[1]⋆ J. K. Webb,[1,2] V. Dumont[1,3] and S. Riemer-Sørensen[4,5]

[1]*School of Physics, University of New South Wales, Sydney, NSW 2052, Australia*
[2]*DAMTP, Centre for Mathematical Sciences, Wilberforce Road, Cambridge CB3 0WA, UK*
[3]*Department of Physics, University of California, Berkeley, CA 94720-7300, USA*
[4]*Institute of Theoretical Astrophysics, The University of Oslo, Oslo NO-0316, Norway*
[5]*ARC Centre of Excellence for All-sky Astrophysics (CAASTRO), University of Queensland, Brisbane, QLD 4072, Australia*

ABSTRACT

The spectrum of the $z_{em} = 2.63$ quasar Q1009+2956 has been observed extensively on the Keck telescope. The Lyman limit absorption system $z_{abs} = 2.504$ was previously used to measure D/H by Burles & Tytler using a spectrum with signal to noise (S/N) approximately 60 per pixel in the continuum near Ly α at $z_{abs} = 2.504$. The larger data set now available combines to form an exceptionally high S/N spectrum, around 147 per pixel. Several heavy element absorption lines are detected in this Lyman limit system, providing strong constraints on the kinematic structure. We explore a suite of absorption system models and find that the deuterium feature is likely to be contaminated by weak interloping Ly α absorption from a low column density H I cloud, reducing the expected D/H precision. We find D/H $= 2.48^{+0.41}_{-0.35} \times 10^{-5}$ for this system. Combining this new measurement with others from the literature and applying the method of Least Trimmed Squares to a statistical sample of 15 D/H measurements results in a 'reliable' sample of 13 values. This sample yields a primordial deuterium abundance of $(D/H)_p = (2.545 \pm 0.025) \times 10^{-5}$. The corresponding mean baryonic density of the Universe is $\Omega_b h^2 = 0.02174 \pm 0.00025$. The quasar absorption data are of the same precision as, and marginally inconsistent with, the 2015 CMB *Planck* (TT+lowP+lensing) measurement, $\Omega_b h^2 = 0.02226 \pm 0.00023$. Further quasar and more precise nuclear data are required to establish whether this is a random fluctuation.

Key words: ISM: clouds – quasars: absorption lines – quasars: individual: (Q1009+2956) – cosmology: observations.

1 INTRODUCTION

The primordial or big bang nucleosynthesis (BBN) epoch, a few minutes after the big bang, is the earliest time in the evolution of the Universe at which we are able to probe the physical properties of the Universe at high precision (see e.g. Iocco et al. 2009; Cyburt et al. 2016, and references therein). It was an epoch where a handful of light elements, such as H, He, Li, Be, were synthesized. Of these elements, deuterium is known as the best baryometer due to its sensitivity to, and monotonic dependence on, the baryon-to-photon ratio η_{10} or the baryon density Ω_b (since $\eta_{10} = 273.9 \, \Omega_b h^2$; Steigman 2006).

For years it has been known that the available D/H measurements show a large scatter relative to their uncertainties (e.g. Balashev et al. 2016). Whilst it is likely that this is simply a consequence of

underestimating systematic errors, it is important to establish that is the case in order to rule out cosmological inhomogeneities. A few quasars that have previously been used for D/H estimates have subsequently been observed extensively such that the newer and much higher quality data permit a more accurate spectral analysis. The higher signal to noise (S/N) permits more stringent checks for systematic problems in the data and allows more accurate modelling of the kinematic structure of the absorption system.

The importance of high S/N data has recently been demonstrated by Riemer-Sørensen et al. (2015). Their analysis of the $z_{abs} = 3.256$ absorption system towards PKS 1937−101 based on higher S/N data revealed a more complex kinematic structure than the previously published analysis (Crighton et al. 2004). The new measurement led to a higher value of D/H than the previous result and also led to a small reduction of the scatter. As another example, very recently Riemer-Sørensen et al. (2017) re-measured the $z_{abs} = 3.572$ absorption system towards PKS 1937−101 leading to a further slight reduction of the scatter. A further example is a systematic re-

⋆ E-mail: e.zavarygin@gmail.com (EOZ) jkw@phys.unsw.edu.au (JKW)

measurement of D/H in Damped Lyman Alpha systems by Cooke et al. (2014) and Cooke, Pettini & Steidel (2018).

Here, we present a re-measurement of the $z_{abs} = 2.504$ metal-poor Lyman limit system (LLS) towards the quasar Q1009+2956 (aka J1011+2941). Since the last analysis of this system by Burles & Tytler (1998) (hereafter BT) this quasar has been observed numerous times and the combined data set yields an effective S/N of 147 in the continuum for Ly α at $z_{abs} = 2.504$, compared to about 60 for the BT spectrum.

2 DATA REDUCTION

2.1 Data

The quasar Q1009+2956 ($z_{em} = 2.63$) was observed with the Keck telescope using the high-resolution Echelle spectrograph (HIRES, Vogt et al. 1994) during many independent programs. The raw science and calibration data were provided by the W. M. Keck Observatory Archive (KOA).[1] Out of all the available exposures we ignored a small number with low S/N because cosmic ray identification and removal during data reduction is less effective in those cases, and because the impact on the final data quality was negligible. The data set used in this study consists of 22 individual exposures with total observation time of 23.7 h. Information about individual exposures is presented in Table 1.

The HIRES data used by BT (six exposures, see their table 1) are not included in our data set. The reason for that is as follows. The first exposure (9000 s, 1995 December 28) does not have a ThAr frame for wavelength calibration. Reduction of the second exposure (7200 s, 1995 December 28) fails due to overlapping orders. The third exposure (4800 s, 1995 December 28) has very low S/N and thus was not included in our data set. Data of the remaining three exposures were not available from the archive or elsewhere. Therefore, the estimate of the primordial deuterium abundance given in this paper is based on a completely independent data set from BT.

2.2 Reduction

For data reduction, MAuna Kea Echelle Extraction (MAKEE) package[2] developed by Tom Barlow has been used. A standard data reduction procedure has been followed. This includes bias subtraction using a bias frame, tracing the object position on each CCD exposure using an image of either a bright star or a 'pinhole' quartz lamp exposure, optimal extraction of the science spectra for each Echelle order, flat-fielding using a quartz lamp frame, and wavelength calibration using a ThAr frame. For each science exposure the available calibration files closest in time and in setups were selected. MAKEE provides wavelength-calibrated spectra of each Echelle order in the heliocentric coordinate frame. There is a known issue with the air-to-vacuum correction formula in MAKEE described in section 2.9 of Murphy et al. (2001). We do therefore not apply this correction in MAKEE. In Appendix A, we describe all the specific data reduction issues and applied solutions for the data in Table 1.

Subsequent processing of the spectra, including the air-to-vacuum correction, has been done in UVES_POPLER[3] software (Murphy 2016). This software re-disperses spectra of Echelle orders on to a common vacuum heliocentric wavelength scale, combines the Echelle orders, and removes leftover cosmic rays by σ-clipping. The automatic procedure of cosmic ray clipping was checked by visual inspection. We also excluded all identified affected parts of the spectra (by blemishes of the CCD such as e.g. bad columns, bleeding regions etc., see Appendix A). Initial continuum estimation was done in UVES_POPLER. First, the spectra were divided into overlapping chunks of 10 000–20 000 km s^{-1} in the Ly α forest and 2000–2500 km s^{-1} in the red part, depending on S/N and setups. Then, the continuum in each chunk was fitted with low-order Chebyshev polynomial (fourth order at forest and sixth–eighth order in the red part, depending on S/N and setups). This automatic procedure was then checked by visual inspection to re-fit the continuum in places where the automatic procedure failed (mostly in the forest). Continuum across the high-order Lyman series transitions (shortwards of 3209 Å) where flux does not return to its initial position is estimated as a 1–2-order (depending on S/N and setups) polynomial fit to the absorption-free regions redwards of the Lyman Limit, 3209–3249 Å, extrapolated to lower wavelengths. Fig. 1 illustrates the continuum normalization results. The whole process results in a sample of 22 one-dimensional normalized spectra, with a few spectral gaps due to either gaps between CCD chips or bad regions having been masked out.

2.3 Velocity offsets and co-added spectra

Previous analyses show that high-resolution Echelle spectra may show velocity offsets between different exposures (e.g. Whitmore, Murphy & Griest 2010; Balashev et al. 2016) and long-range wavelength distortions within each exposure (e.g. Molaro et al. 2008; Rahmani et al. 2013; Whitmore & Murphy 2015; Dumont & Webb 2017). Whilst long-range distortion effects are important for varying constant analyses, they are unlikely to impact on the D/H measurement studied here so are ignored. We do however correct for velocity shifts between individual exposures.

In order to correct for possible velocity offsets between individual exposures, we use the following approach. Using purpose-written software[4] we determine velocity offsets between the normalized one-dimensional spectra using a method similar to one described by Evans & Murphy (2013). First, all exposures were convolved with a Gaussian (the assumed instrumental profile) whose values are given in Table 1. Then flux and error arrays of one, anchor, exposure are modelled with cubic splines to get continuous functions of wavelength and the constant velocity offset, Δv: $f_o(\lambda, \Delta v)$ and $\sigma_o(\lambda, \Delta v)$, respectively. The velocity offset between the jth exposure and the anchor exposure (subscript zero) is given by the minimum of the χ^2 function:

$$\chi^2(\Delta v) = \sum_i \frac{[f_o(\lambda, \Delta v) - f_j(i)]^2}{\sigma_o^2(\lambda, \Delta v) + \sigma_j^2(i)}, \tag{1}$$

where $f_j(i)$ and $\sigma_j(i)$ are flux and error arrays, i runs over all valid pixels unaffected by cosmic rays/CCD blemishes, not falling in gaps between the CCD chips and corresponding to a wavelength range covered by both the anchor and the examined spectrum. The wavelength in the anchor spectrum is calculated using a pixel-to-wavelength solution for the examined spectrum. The distribution of velocity offsets in our sample is illustrated in Fig. 2 showing that offsets between exposures up to 1.3 km s^{-1} appear to be present.

---

[1] https://koa.ipac.caltech.edu/cgi-bin/KOA/nph-KOAlogin
[2] http://www.astro.caltech.edu/~tb/makee/index.html
[3] http://astronomy.swin.edu.au/~mmurphy/UVES_popler/

[4] https://github.com/ezavarygin/voffset

Table 1. Journal for the Keck/HIRES observations of the Q1009+2956 used in this work.

Date	PI	Decker	Bin.[a]	Wavelengths[b] (Å)	Res. power	σ_v[c] (km s^{-1})	Exposure (ks)
1997-05-11	Tytler	C5	1	3160−4652	37000	3.4	3
1998-12-14	Tytler	C5	1	3099−4619	37000	3.4	8
1998-12-15	Tytler	C5	1	3099−4619	37000	3.4	6.3
1999-03-08	Tytler	C5	1	3121−4642	37000	3.4	7
1999-03-09	Tytler	C5	1	3121−4642	37000	3.4	7.2
2003-06-07	Sargent	C1	1	4605−7020	49000	2.6	3 + 3
2003-06-08	Sargent	C1	1	4612−7033	49000	2.6	3 + 3
2004-11-05	Tytler	C5	1	3075−5883	37000	3.4	3.6
2005-03-31	Tytler	C5	1	3129−5954	37000	3.4	5.4
2005-04-30	Sargent	C1	1	3157−5967	49000	2.6	3 + 3
2005-05-01	Sargent	C1	1	3157−5967	49000	2.6	3 + 3
2005-05-31	Steidel	C5	2	3138−5970	37000	3.4	1.8 + 1.8
2005-06-01	Steidel	C5	2	3138−5970	37000	3.4	1.8 + 1.8
2005-12-06	Tytler	C5	1	3076−5897	37000	3.4	6.4
2008-03-29	Tytler	C1	2	3129−5955	49000	2.6	3.6 + 3.6
Total							85.3

[a]CCD binning in the dispersion direction.

[b]The wavelength values are in the vacuum/heliocentric frame.

[c]Resolution in terms of $\sigma_v = \mathrm{FWHM}/2\sqrt{2\ln 2}$, where FWHM is a full width at half-maximum of the assumed Gaussian instrumental profile.
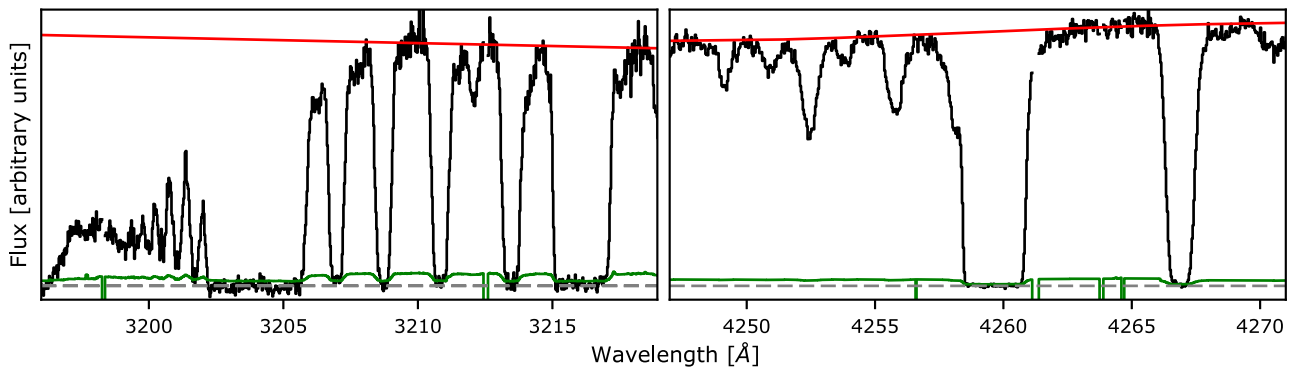
Figure 1. A reduced non-flux-calibrated spectrum and the initial continuum estimate for Q1009+2956. The spectrum shown is the exposure obtained on 2008 March 29 (starting time 06:40, UT). The other exposures are similar. The left-hand and right-hand panels show the Lyman limit and the Ly α regions of the $z_{abs} = 2.504$ LLS, respectively. The black histogram shows the flux in arbitrary units. The smooth red line on top of the flux indicates the initial continuum estimate. The horizontal dashed line shows a zero level. The green histogram near zero shows 1σ error with negative values being pixels clipped due to cosmic rays.

Using this result, we co-add[5] the 22 spectra, taking into account the offsets, into four spectra according to the decker and CCD binning in dispersion direction used at the exposures:

(i) C1, unbinned (C1 × 1)
(ii) C1, 2×binned (C1 × 2)
(iii) C5, unbinned (C5 × 1)
(iv) C5, 2×binned (C5 × 2).

Note, we also create a single co-added spectrum by combining the whole data set. This spectrum is only used to demonstrate the overall quality of the data in Fig. 3 and to calculate an effective S/N around the transitions used (Table 2). Apart from that, all final values and uncertainties stated in this work correspond to simultaneously fitting the four co-added spectra above.

---

[5]A purpose-written PYTHON code was used: https://github.com/ezavarygin/wspectrum
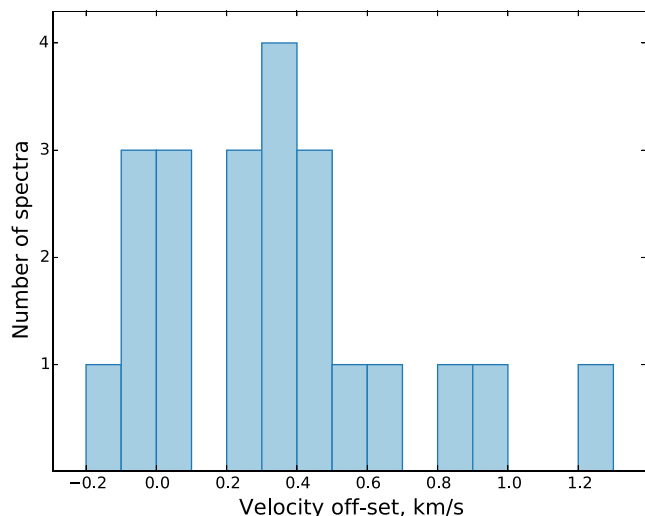
3 ANALYSIS

3.1 Fitting

To fit the spectra we use the VPFIT[6] software (Carswell & Webb 2014). It uses a standard procedure of Voigt profile fitting to find a solution for redshift, z, column density, N, and the Doppler b-parameter (with an option to solve for both temperature, T, and the turbulent b_{turb} parameter) and their uncertainties for each of the absorbers. VPFIT can also solve simultaneously for other parameters, including the local continuum and zero level. VPFIT minimizes a sum of the weighted squared residuals (χ^2) between the model profile convolved with the instrumental resolution and the observed spectra, summed over the four co-added spectra.

During the fitting process, we adopt the following assumptions. For all the models, redshifts of H I and D I are tied in each component. The D/H ratio is assumed to be constant over the complex.

---

[6]Version 10.2, https://www.ast.cam.ac.uk/~rfc/vpfit.html

Figure 2. Distribution of velocity offsets between all 22 exposures in the final sample. The offsets were calculated with respect to the exposure obtained on 2008 March 29 (starting time 06:40, UT).

The metallicity of this system is known to be low (BT), deuterium destruction is thus likely to be negligible, and the deuterium abundance should be very close to primordial in this system. In order to get a tighter constraint on the D/H ratio in the LLS, we solve for total H I and D I column densities rather than for column densities of the individual components. For models without metals we tie the Doppler parameters of H I and D I thermally (to be discussed in Section 3.6). When metals are used we solve for both b_{turb} and T.

To take into account any possible residual velocity offsets between independent spectra, a velocity shift free-parameter is included in the model (the best-fitting values lie close to zero: $\lesssim 0.1\,\mathrm{km\,s^{-1}}$). For transitions with saturated lines, a zero level is fitted, allowing correction for a possibly imperfect night-sky subtraction during the data reduction process.

3.2 Continuum fitting

Continuum fitting is a multistep procedure. During the spectral data reduction, the spectral orders from each individual Echelle CCD exposure are extracted to one-dimensional spectra (see also Section 2.2). The orders from each individual exposure are placed into a one-dimensional array that is continuum fitted using polynomial fitting to unabsorbed spectral regions, and then normalized to unit continuum. That process is repeated for each exposure/CCD so that multiple exposures can be combined using weighted addition to form a final co-added spectrum.

An example of the estimated continuum in the Ly α region for one of our 22 separate exposures is illustrated in the right-hand panel of Fig. 1. The left-hand panel of the same figure illustrates the Lyman limit region from the same 3600 s exposure (corresponding to the last entry in Table 1). Thus, continuum fitting has already played an important role in the reduction process prior to analysis, during which further refinement of the continuum takes place.

Then, during VPFIT, each spectral segment is assigned additional free parameters to allow further adjustment of the continuum, using a linear fit, with free parameters slope and normalization. If there are insufficient continuum regions flanking the absorption line segment being fitted, only the normalization is varied and the slope is

assumed to be zero. Importantly, the continuum parameters are adjusted simultaneously with all other model parameters. This means that the error estimates for the parameters of interest (in this case neutral hydrogen and deuterium column densities) properly take into account variation in the continuum parameters.

The reason for using a linear continuum fit simultaneous with profile modelling (rather than cubic or higher order) is that at the absorption profile modelling stage of the analysis, we are dealing with small wavelength segments (only a few angstrom). Non-linear continuum variations over such small scales are likely to be caused only by the presence of absorption lines and not by any other physical mechanism intrinsic to the quasar continuum. If weak blends are found to be present, these are most appropriately taken into account by including additional absorption components.

Fig. 4 illustrates the fitting regions used as horizontal blue solid lines near the base of the absorption features. Clearly estimating the continuum placement becomes less reliable as one approaches the Lyman limit. The lower panel of Fig. 4 shows a region covering Lyman 21–24 where blending of high-order Lyman lines results in there being no local absorption-free regions. In this region, the extrapolated continuum was allowed to float but with no variation to the slope.

Fig. 6 illustrates the final co-added spectra over the Ly α region for each of the four groupings of data described in Section 2.3. The spectra are normalized to unit continuum, shown as horizontal black dashed lines. The 1σ ranges of continuum parameters are illustrated as grey shaded areas. The continua in all four cases are very well determined so the grey shaded area is small. The modelling was carried out simultaneously to the four spectral groupings. However, Fig. 3 illustrates the overall data quality as the weighted co-addition of all four spectral groupings.

3.3 Lyman series

All the H I Lyman series transitions of the $z_{abs} = 2.504$ LLS down to the Lyman Limit have been covered by the exposures from our sample (Fig. 4). The following transitions were included in the fitting process: Ly α, Ly β, Ly γ, Ly 6, Ly 13, Ly 14, Ly 21–24 (Table 2). Whilst other Lyman series lines were available, blending with lower redshift Lyman forest lines was more severe, such that those lines did not help to constrain the parameters of interest. In addition to the main components of the LLS, there is one nearby component that has a fairly low H I column density. It becomes optically thin at Ly 5 and therefore does not show associated deuterium. This component lies $\sim +60\,\mathrm{km\,s^{-1}}$ from the main absorption complex. It is fitted simultaneously with other parameters in order to properly take into account any impact it may have on determining the parameters of interest.

An apparent deuterium feature is clearly seen at its expected position ($\sim -81.5\,\mathrm{km\,s^{-1}}$ bluewards from the centre of the H I line) of Ly α at $z_{abs} \approx 2.5036$. Due to blending with low-redshift Ly α forest lines the deuterium absorption is not visible in higher order Lyman series transitions.

3.4 Metals

Given line saturation, blending, and finite S/N, it is unlikely that one could reliably determine the velocity structure of the LLS based on the Lyman series alone. A more reliable estimate of velocity structure is obtained by simultaneously fitting both Lyman series lines and metal species.
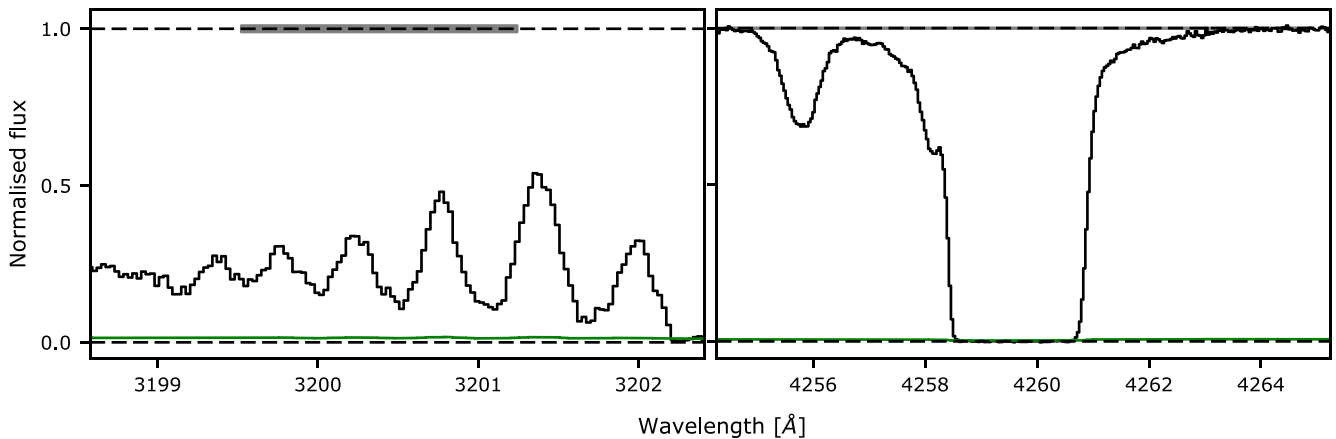
Figure 3. A weighted spectrum of the $z_{abs} = 2.504$ LLS towards Q1009+2956, co-added from all available exposures. The left-hand and right-hand panels illustrate the Lyman Limit and the Ly α regions, respectively. The black histogram is the flux, normalized to the best-fitting continuum corresponding to model 6a (see Sections 3.6 and 3.6.1). The grey-shaded area near $y = 1$ indicates the 1σ confidence region for the continuum fit. The green histogram near $y = 0$ is the 1σ error in the normalized flux. The horizontal dashed lines indicate the zero and normalized continuum levels. The figure illustrates the overall quality of the data and the continuum variation for model 6a.

Table 2. Transitions and spectral ranges used in this work.

Transition	Wavelength range[a] (Å)	Spectra[b]	Cont.[c]	S/N[d]
Ly α	4254.42−4265.00	1−4	2	147
Ly β	3591.10−3597.10	1−4	2	99
Ly γ	3403.95−3411.03	1−4	1	92
Ly 6	3259.30−3264.20	1−3	1	51
Ly 13	3209.81−3212.89	1−4	1	51
Ly 14	3208.08−3210.20	1−4	1	50
Ly 21−24	3199.50−3201.25	1−4	1	27
Si IV 1393	4882.30−4884.10	1−4	2	128
Si IV 1402	4913.90−4915.70	1−4	2	140
C II 1334	4674.80−4676.70	1−4	2	106
C III 977	3422.22−3425.60	1−4	1	80
C IV 1548	5423.20−5425.50	1−4	2	119
C IV 1550	5432.25−5434.60	1−3	2[e]	98

[a] Wavelength ranges used to fit the transition.

[b] Co-added spectra used in fitting out of four listed in Section 2.3. The 4th spectrum, C5 $\times$ 2, was excluded for Ly 6 and C IV 1550 due to very low S/N and/or cosmic ray contamination.

[c] Order of polynomials used to fit the continuum in the region: 1 – floating continuum, 2 – floating continuum with a non-zero slope.

[d] Average signal-to-noise ratio measured at continuum level around the transition. For the Ly 21−24 transitions the least absorbed pixels are used since the flux does not return to the continuum level.

[e] Except the C5 $\times$ 1 spectrum where the slope is not allowed to vary due to cosmic ray contamination.

The LLS at $z_{abs} = 2.504$ shows absorption by carbon and silicon from different ionization levels: C II, C III, C IV, Si II, Si III, and Si IV. Fig. 5 illustrates the relevant ionization potentials. Absorption lines from C II 1334, C III 977, C IV 1548 and 1550, Si II 1260, Si III 1206, Si IV 1393 and 1402 are detected. C III, C IV, and Si IV are all strong. Fitting each of these features independently shows consistent structures with at least two strong components. However, a two-component simultaneous fit to C III, C IV, and Si IV with tied redshifts and b-parameters does not yield a statistically acceptable fit. A three-component model does however fit well preferring a thermal broadening over the turbulent one (see Section 3.6). C II is also consistent with C III, C IV, and Si IV but is weak so provides little additional constraint. Fig. 10 illustrates the transitions used.

The Si III 1206 Å line was not used due to a problem with the rest-frame wavelength (see Section 3.4.1). Si II 1260 is very weak and is also badly blended so was discarded. Table 2 contains information about the transitions and wavelength ranges used to constrain the parameters of interest.

3.4.1 Si III

Morton (2003) provides a Si III 1206 rest-frame wavelength of 1206.500 Å. Using this wavelength leads to an apparent shift of Si III 1206 redwards with respect to C III, C IV, and Si IV by $\sim$+ 2.5 km s^{-1}. The NIST spectral data base gives an observed wavelength of 1206.51 Å (Kramida et al. 2015). The 0.01 Å difference between the Morton and NIST wavelength corresponds to 2.49 km s^{-1} in the same direction. We conclude that the most likely explanation of the observed shift between Si III and other transitions is simply due to the fact that the true Si III 1206 wavelength is closer to the old experimental value of 1206.51 Å than the value listed in Morton (1206.500 Å). Due to this uncertainty, to avoid any possible bias we exclude Si III 1206 from our analysis – new experimental data for this transition are clearly needed.

Another potential explanation of the shift, in case the Morton value is correct, is that the Si III 1206 Å line is blended with some unidentified interloper. If so, this is unlikely to be hydrogen (even though the redshifted Si III 1206 line falls in the forest) because the observed line width is far smaller than typical forest b-parameters. We searched for potential interlopers at the redshifts of other absorption systems identified along the line of sight (and from the interstellar medium of our own Galaxy) but did not find any likely candidates (see Section 3.5).

3.5 Possible contamination

There is always a chance that a transition of interest is blended with absorption from another system at a different redshift. It is therefore important to try and identify all absorption systems present in the spectrum. This was done in two ways. First, a cross-correlation
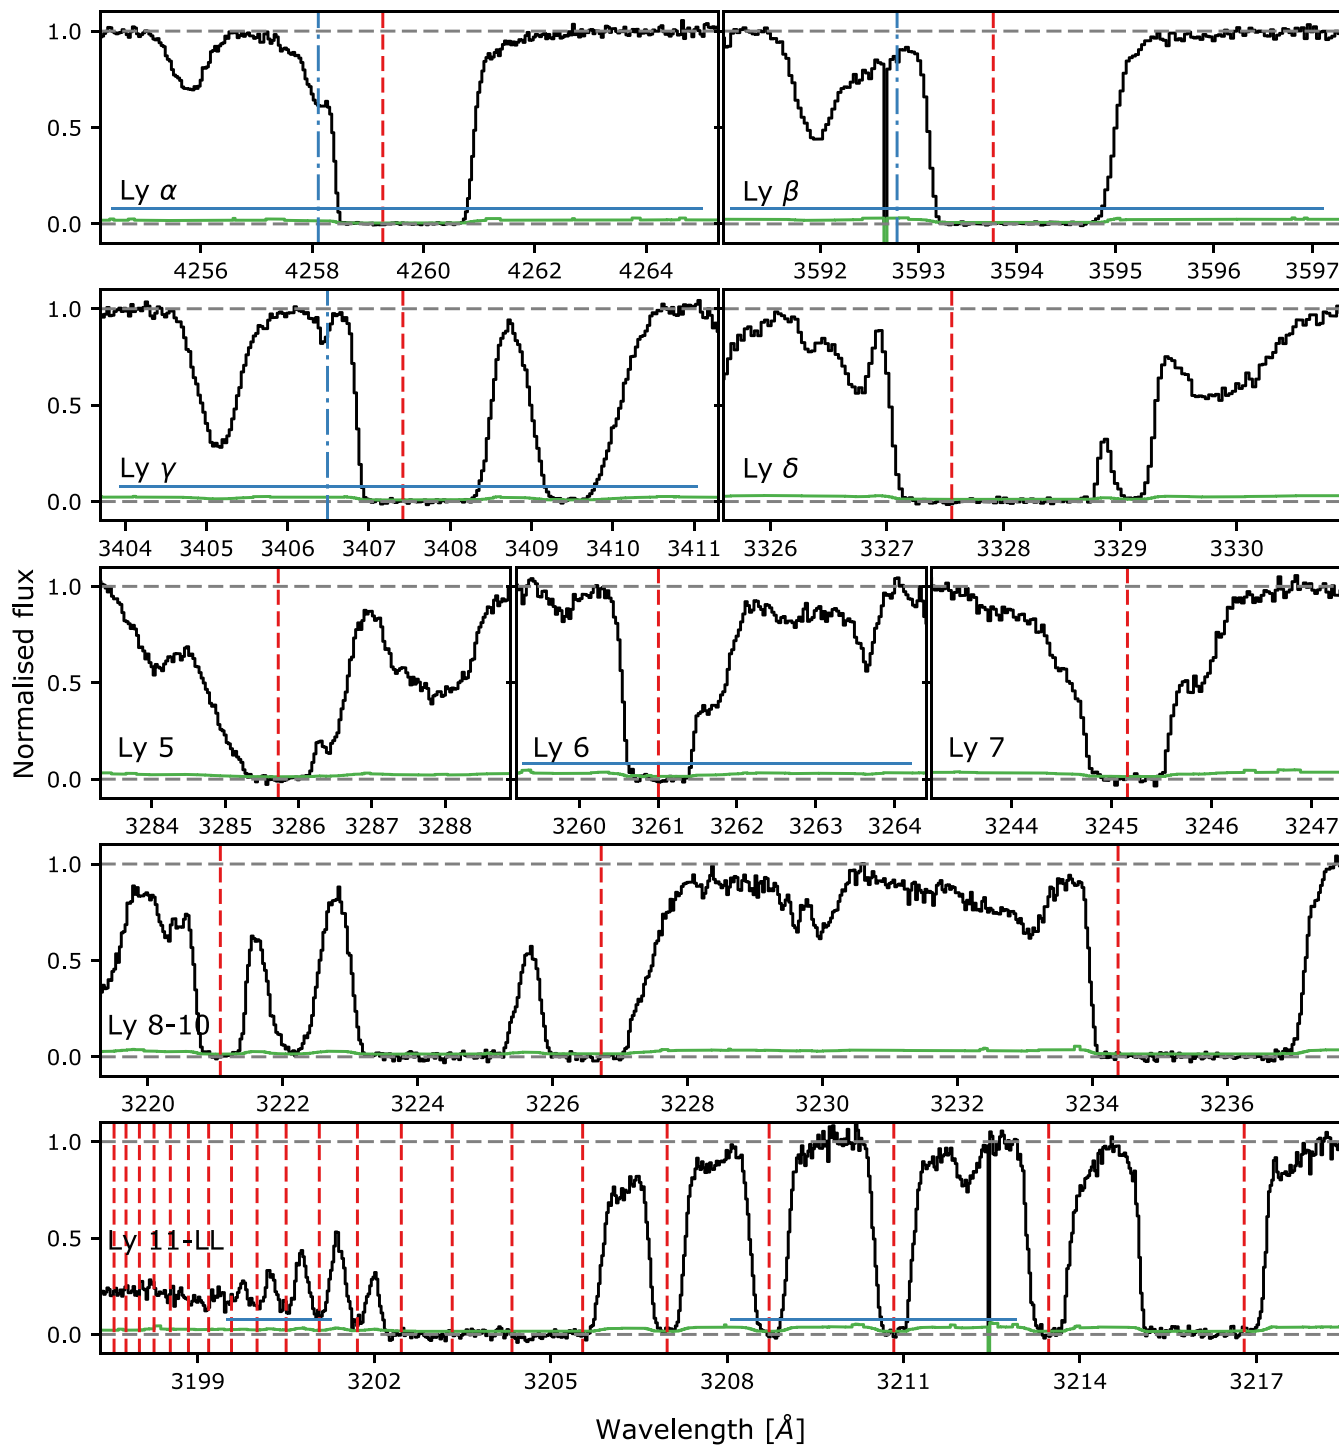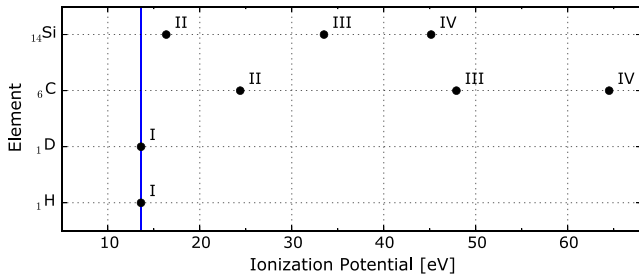
Figure 4. Lyman series transitions for the $z_{abs} = 2.504$ Lyman limit system towards Q1009+2956 for the co-added C1 × 2 spectrum. The green continuous line (around $y = 0$) illustrates the 1σ error array. The horizontal blue solid lines (at $y \simeq 0.1$) indicate wavelength regions used in fitting (see Table 2). The vertical red dashed lines indicate the centre of the H I Lyman series transitions in a one-component model (model 1a). The vertical blue dash–dotted line in the Ly α, Ly β, and Ly γ panels indicates the corresponding D I position. The sharp drops of the flux around 3592.7 Å in the top-right panel and 3212.5 Å in the bottom panel are pixels clipped due to cosmic rays.

method[7] was used to compare a list of standard metal lines with the list of the absorption lines detected in the spectrum. Secondly, a visual inspection of the spectrum was carried out, scanning in redshift, using a template of standard lines, to identify coincidences in velocity space for transitions in the same system. The whole process resulted in the identification of 24 absorption systems (including the $z_{abs} = 2.504$ system of interest) given in Appendix B.

[7]https://github.com/TrisD/Ab_Detect

Figure 5. Ionization potential values for the transitions detected in the $z_{abs} = 2.504$ Lyman limit absorption system. The vertical solid line indicates the ionization potential of H I, 13.5984 eV. The values were taken from the NIST Atomic Spectra Database (Kramida et al. 2015).

We then check whether any commonly detected transitions at each candidate redshift fall within $\pm 100\,\mathrm{km\,s^{-1}}$ of any transition in our $z_{abs} = 2.504$ system of interest[8]. Nevertheless, despite a careful search for potential contamination of this sort, none was found.

3.6 Fitting different models

Weak absorption lines due to metals with low ionization potentials (Si II, C II) are detected. We do not detect any neutral species, e.g. O I. The relative abundances of the various ionization stages, can vary significantly for each velocity component across the absorbing medium. Components detected in Si IV and C IV for example, may have weak (even undetectable) H I. This is of course taken into account in the modelling procedure because all column densities are free parameters.

BT concluded that the velocity dispersions of the H I and D I lines are dominated by thermal motions. We agree with this finding. In all models using metal transitions that we consider (see below), we fitted each of the main absorption components with additional free parameters corresponding to temperature and any possible turbulent component of the line broadening, b_{turb}. For all components, b_{turb} is small compared to the observed b-parameter (<5 per cent for H I), given by $b_{\mathrm{obs}}^2 = b_{\mathrm{therm}}^2 + b_{\mathrm{turb}}^2$. The full set of model parameters are provided as online MNRAS supplementary files and on GitHub (see links in Section 3.7). As a further check on this point, we fit metals both thermally and turbulently with a three-component model. A thermal fit is statistically preferable with $\chi^2/\mathrm{d.o.f.} = 0.678$ versus 0.763 for the turbulent fit. For models in which only the Lyman series is fitted (i.e. without simultaneously fitting to metal lines; see below), we therefore assume thermal broadening, thus avoiding degeneracy between T and b_{turb}.

In order to explore the possible impact on D/H of different assumptions about the velocity structure, we explore six different models for the absorption complex, where the H I components are tied (in redshift and b-parameters) to different combinations of metal lines:

(1) one-component model, fit to Lyman series only;
(2) two-component model, fit to Lyman series only;
(3) three-component model, fit to Lyman series only;
(4) two-component model, C II used;
(5) three-component model, C II used;
(6) three-component model, C II, C III, C IV, and Si IV used.

[8]The following routine has been used: https://github.com/vincentdumont/zblend

We assume throughout that D/H is the same across all components. The first three models use Lyman series transitions only. Since metals strongly support a thermal broadening, in these three models we tie H I and D I thermally. Even though model 1 is clearly unrealistic (at least two strong components are seen in the low ionization metal species), we include it for completeness and note that despite being too simple, it gives a reasonably consistent result for D/H (Section 3.7). Models 4 and 5 exclude the higher ionization species, including only C II 1334 as this line appears to be unblended. Model 6 is a simultaneous fit to four metal species plus the hydrogen and deuterium series. The parameters (z, b_{turb}, and T) of H I and D I were tied with those of metals in models 4 through 6.

The total H I column density is too low to explain the Lorentzian-like wings that are particularly visible in the wings of Ly α (see Fig. 6). These broad wings appear to be symmetric with respect to the profile associated only with the sharp central absorption feature. We therefore explore whether the broad residual absorption is most likely to be explained by multiple narrow blends in the Ly α wings or whether one large b-parameter, spread across the whole profile, provides the 'correct' model. These two groups of models are indicated as 1a−6a and 1b−6b, respectively.

3.6.1 Modelling the broad wings in H I Ly α using multiple narrow blends

We first assume that the broad wings are caused by multiple weak, narrow blends. The severe blending means that the line parameters for the multiple weak lines needed to model the wings adequately are poorly constrained. The blends in the red wing do not affect the D/H estimate substantially since the impact of these weak blends on $\log N$(H I) for the stronger central absorption is negligible. However, the broad residual absorption in the blue wing of Ly α can also be modelled using several narrow blends, one of which lies close to the D feature and which, unfortunately, significantly degrades the precision with which we can determine D/H for this system. Fig. 7 illustrates the details in the blue wing of H I Ly α for models 1a−6a.

The S/N of our spectra is much higher than the spectrum studied by BT (~147 versus 60 in the continuum level around Ly α). BT suggested that the deuterium line is blended with a Ly α forest line at approximately $-30\,\mathrm{km\,s^{-1}}$ of the D I Ly α line (see their fig. 8a). The BT model contains only this one blend in the vicinity of the D line. Using narrow lines, we found a best fit to the blue wing of Ly α using three components, illustrated in Fig. 7, which shows all of our six models.

We can estimate the impact of this blend on the D/H ratio precision by fixing the parameters of the blend and running VPFIT once again. For example, for three-component models 3a, 5a, and 6a, fixing the parameters of the blend leads to significantly reduced D/H uncertainties: from 0.5 to 0.027 dex, from 0.25 to 0.018 dex, and from 0.08 to 0.014 dex, respectively.

The blend parameters and therefore D/H are sensitive to the initial continuum level (Fig. 1). We include a floating continuum with a varying slope over the Ly α region in our models, allowing for fine-tuning of the continuum placement. The apparent blend close to the D line is not explained by a simple fluctuation in continuum fitting since it would require a 10–15 per cent variation (see Fig. 7) over the D I Ly α line. Such a strong variation over such a small scale (less than 2 Å or $140\,\mathrm{km\,s^{-1}}$) is unlikely.

We raise one further point concerning the apparent blending of the D I. We opted to interpret the 'extra' absorption as being due to an interloping H I line. The data are consistent with this interpreta-
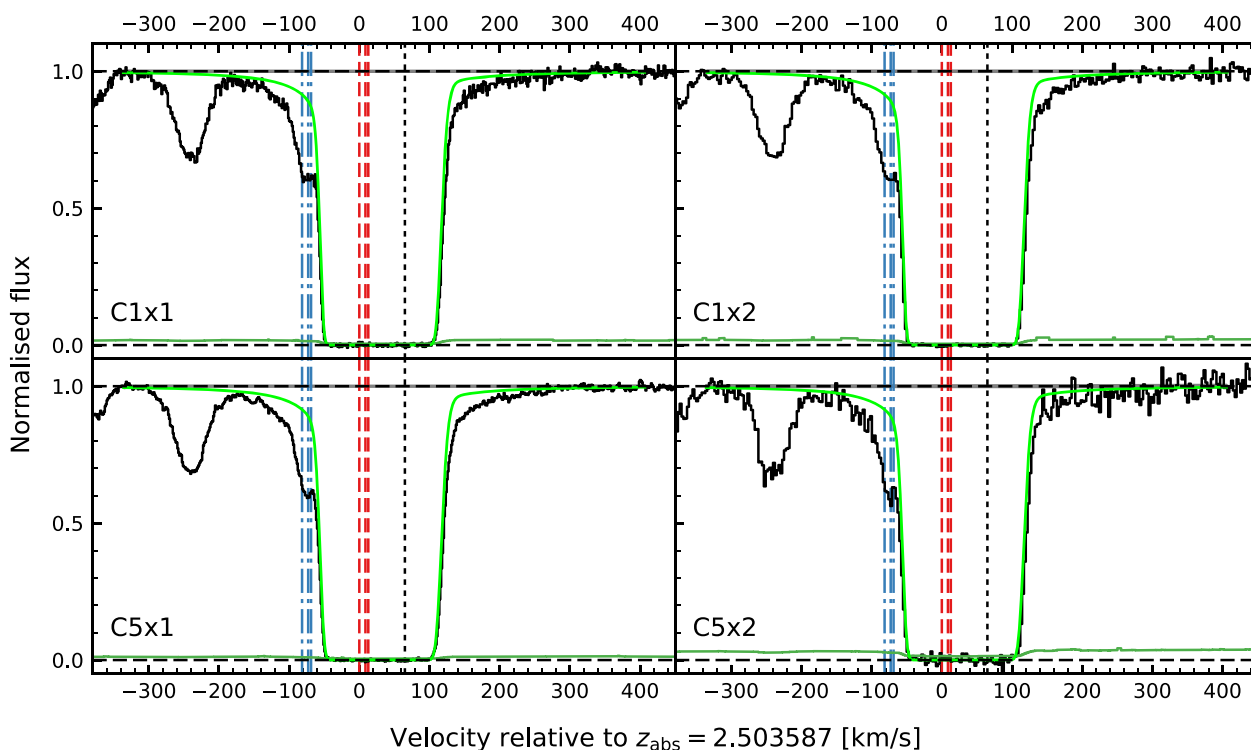
Figure 6. The Ly α transition for the $z_{abs} = 2.504$ Lyman limit absorption system for each of the four independent data sets, as indicated at the bottom left corner of each panel. The continuous green line is formed from the three H I components in model 6a (vertical red dashed lines) that give rise to the D I absorption feature (their expected position is indicated with vertical blue dash–dotted lines), plus one additional H I component (vertical black dotted line) that is required to account for the additional absorption redwards of the three main components. The 1σ error array is shown by green histogram close to $y = 0$. The grey (very narrow) shaded area near $y = 1$ corresponds to the $\pm 1\sigma$ uncertainty in the continuum. The purpose of this figure is to illustrate the additional broad absorption wings either side of the main, strong, H I components.

tion and indeed this is the most likely scenario. If we had instead interpreted the residual absorption as being due to D I, a corresponding additional H I component would obviously be required. If we maintain the assumption of constant D/H in all velocity components, the H I column density in this additional component would be too high (higher order Lyman lines are unsaturated at the corresponding velocity), thereby ruling out that interpretation. On the other hand, if we were to relax the assumption of constant D/H for all velocity components, the additional component in question would then permit a very high D/H in that specific component, requiring an inhomogeneous BBN explanation or subsequent inhomogeneous astration. Whilst we cannot rule that scenario out, we have not explored it in further detail and simply acknowledge this potential bias in interpreting the data.

3.6.2 Modelling the broad wings in H I Ly α using one single broad line

Given the symmetry of the additional absorption in the wings of the Ly α line, instead of using multiple narrow blends to accurately model the broad wings, we find that a single broad ($b \sim 105 \, \mathrm{km \, s^{-1}}$) line sitting on top of Ly α also provides a good fit to the data. In doing so, the parameters of the single broad blend are very well constrained, and the model yields an extremely high D/H precision of (see Section 3.7). Fig. 8 illustrates the details in the blue wing of H I Ly α for models 1b−6b.

3.7 Fitting results

Best-fitting values of the total H I and D I column densities and corresponding D/H ratios with 1σ confidence intervals for each of the models considered are presented in Tables 3 and 4. The total H I column density is very well constrained, is consistent amongst all our models and agrees with the constraints from BT. Fig. 9 illustrates the D/H values for our models.

Model 6a is illustrated in Figs 10 (a) and (b) and provides a velocity structure comparison between hydrogen and metal lines. The data plotted are C1 × 2 and C5 × 1, respectively. These two co-added spectra are the highest S/N. Models 1a−5a give similar-looking residuals (see Fig. 7). The blue wing of Ly α for models 1b−6b is shown in Fig. 8.

For model 6a, we demonstrate the justification for including the Ly 21–24 region in the modelling procedure. If we exclude the region entirely, and re-fit, we obtain $\log N_{total}(\mathrm{H \, I}) = 17.355 \pm 0.014$. This compares to $\log N_{total}(\mathrm{H \, I}) = 17.360 \pm 0.005$ when the region is included in the fit (with a local freely varying continuum). We thus see a factor of almost 3 improvement in the precision of the total neutral hydrogen column density measurement. The actual column densities in both cases are consistent and D/H does not alter significantly.

All VPFIT output files with the parameters and uncertainties for all the models considered are available in online supplementary files[9]

---

[9]MNRAS supplementary files url.
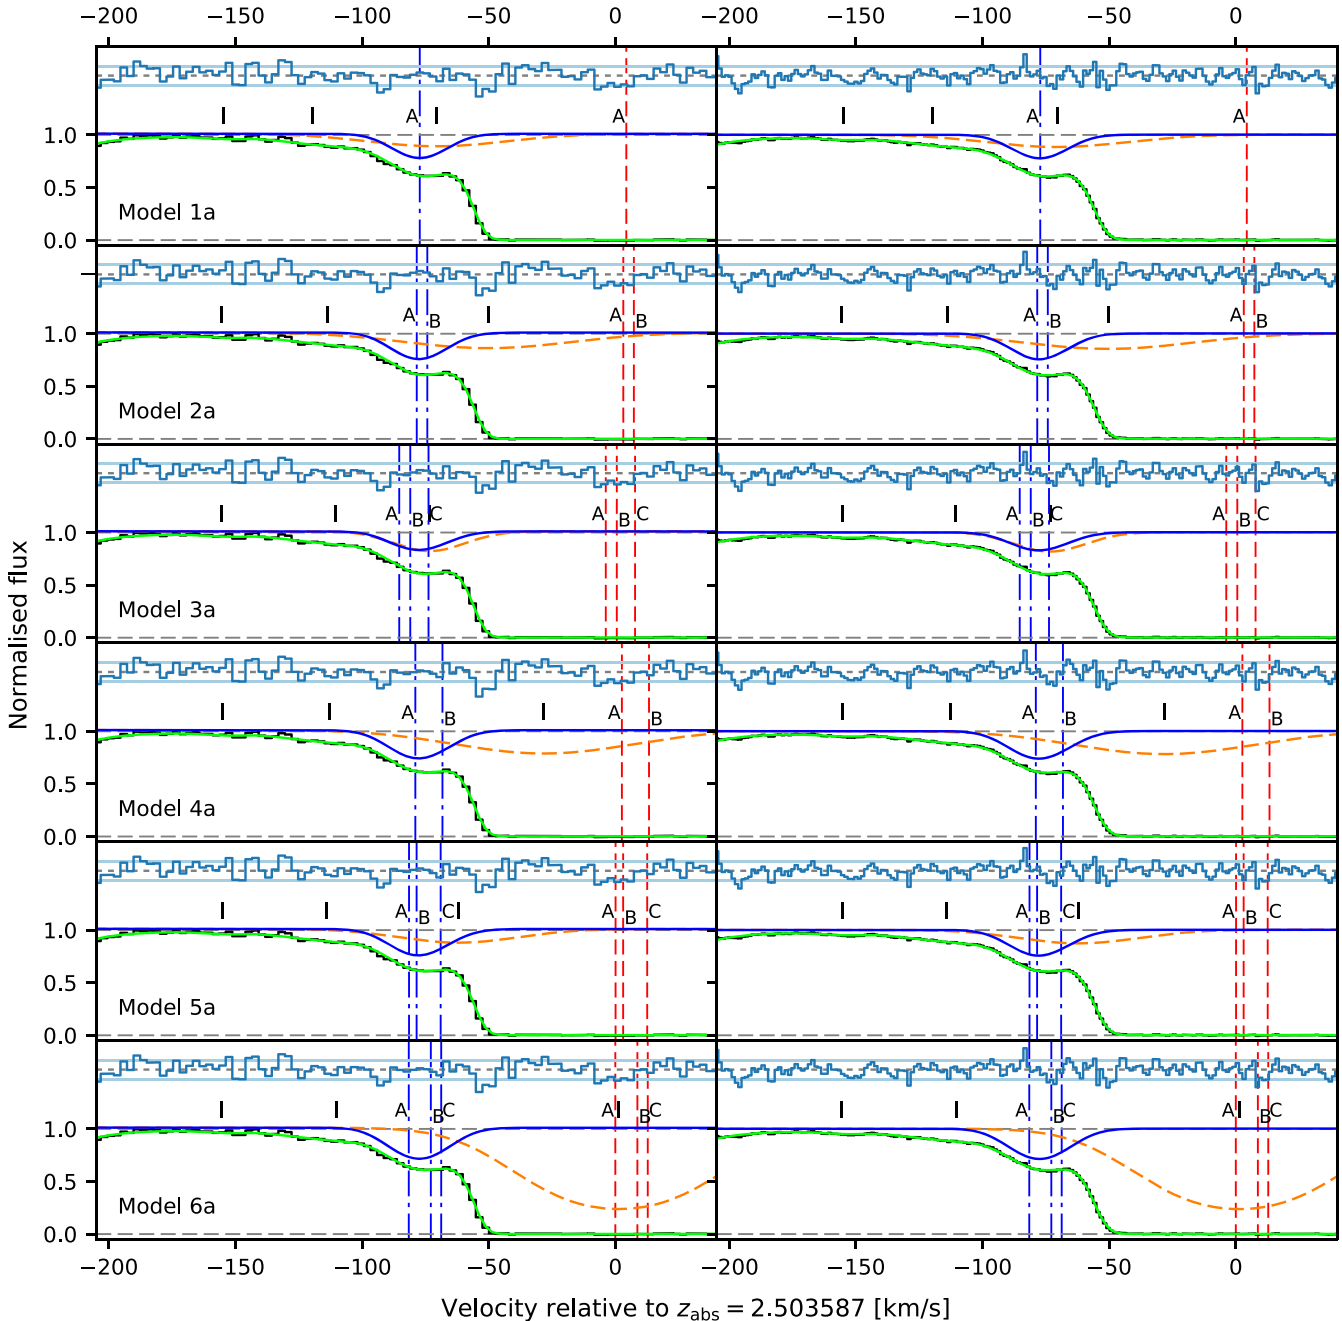
Figure 7. The overall absorption profile (green solid line) for the blue wing of the H I Ly α transition for the six considered models using multiple narrow blends (model 1a is at the top) described in Section 3.6.1. Left-hand and right-hand panels correspond to the C1 $\times$ 2 and C5 $\times$ 1 spectra (black histogram) respectively. These two co-added spectra illustrated are the highest S/N. However all four co-added spectra were fitted. The residuals with 1σ confidence range are shown in the blue histogram above each fit. The main absorption components are indicated with vertical red dashed lines with letters above. The positions of the corresponding deuterium transitions are indicated with the blue vertical dash–dotted lines. The ticks above each fit indicate absorption by additional H I clouds included in the model. The blue solid line illustrates the overall D I absorption. The orange dashed line shows the additional contaminating H I line that degrades the measurement precision on D I. The *x*-axis shows the velocity shift relative to component A for model 6a.

and also on GitHub[10]. A short description of the supplementary files is given in Appendix C.

3.8 Model comparison: preferred number of velocity components

To establish the number of absorbing components favoured by the data, we use an Akaike information criterion (Akaike 1974) corrected for finite sample sizes (e.g. Burnham & Anderson 2002):

$$\text{AICc} = \chi^2 + 2p + \frac{2p(p+1)}{n-p-1}$$

[10]https://github.com/ezavarygin/q1009p2956
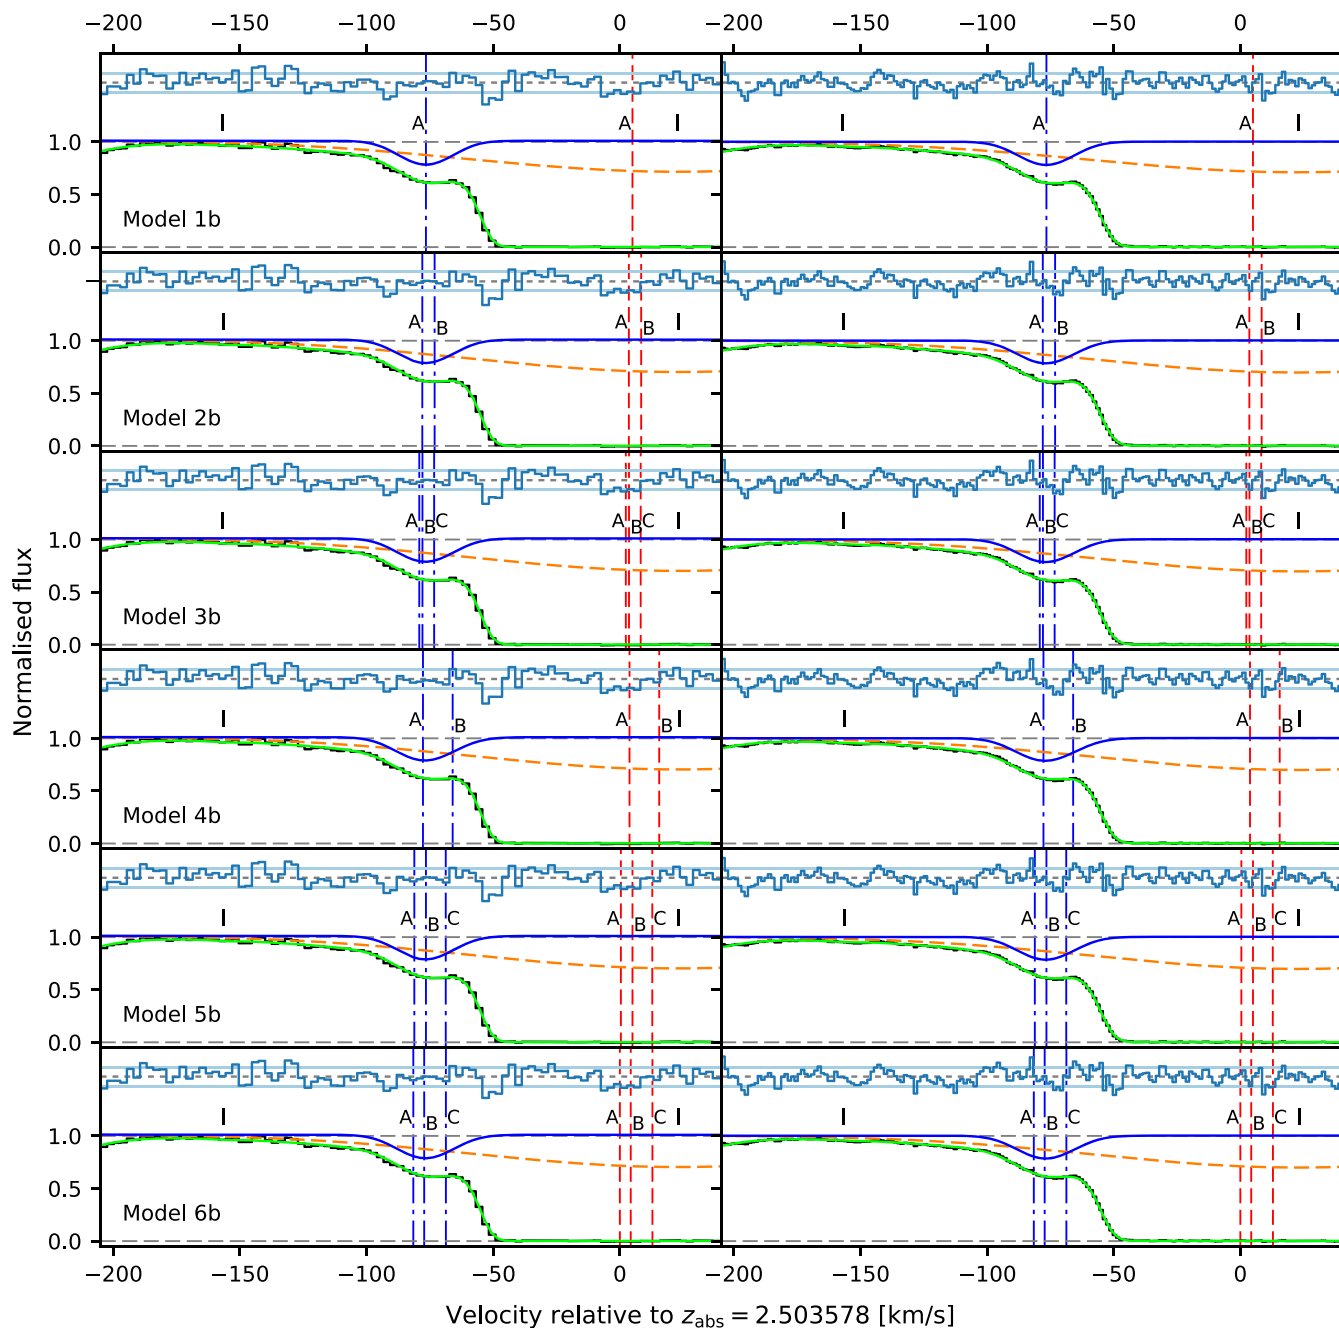
Figure 8. Same as Fig. 7 but for models 1b−6b using a single broad line (shown with the orange dashed line) to fit broad wings in H I Ly α.

where n and p are the sample size and the number of model parameters, respectively. Unlike to χ^2, AICc includes a penalty for including too many parameters (i.e. overfitting). The model with the smallest AICc is the one preferred by the observational data. One can of course only compare AICc values when using the same data sets, i.e. we can use AICc to compare models 1 through 3 or models 4 and 5 but not to compare e.g. models 3 and 5 (since an additional metal transition is used in the latter). We use the following scale of levels of empirical support of a given model (Burnham & Anderson 2002): once a model with the smallest AICc (AICc_{min}) is found, models that have ΔAICc = AICc − AICc_{min} greater than 10 are considered as having no support by the data and may be discarded. Models with $4 < \Delta$AICc < 7 and $0 < \Delta$AICc < 2

have 'considerably less' and 'substantial' empirical support, respectively.

Table 3 contains AICc values for models 1a to 6a and Table 4 gives the results for 1b to 6b. AICc differences are given between the given and the smallest ones within each of three model subsets: models 1–3, models 4–5, and model 6. We have avoided inter-comparing models a and b (in terms of AICc) because models b clearly have larger χ^2 values. However, this does not rule out models of this sort, i.e. models including a broad component to fit the Ly α wings since these models could be refined further, i.e. by including further blends.

Referring to Table 3 first, all six models appear to give an acceptable χ^2. However, based on the AICc values, a one-component

Table 3. The best-fitting D/H values for the models described in Section 3.6 using multiple narrow blends to fit the broad H I Ly α wings, as described in Section 3.6.1. All the VPFIT files with the best-fitting parameters and their uncertainties for each of the models are available in online supplementary files, given in the links in Section 3.7.

Model	No. components	Metals used	χ^2/d.o.f.	AICc	ΔAICc[a]	log N_{total}(H I)	log N_{total}(D I)	log D/H
1a	1	–	0.71911	4993.7	18.1	17.365 ± 0.005	12.70 ± 0.06	-4.67 ± 0.06
2a	2	–	0.71636	4980.2	4.6	17.363 ± 0.005	12.75 ± 0.07	-4.62 ± 0.07
3a	3	–	0.71500	4975.6	0	17.363 ± 0.005	12.6 ± 0.5	-4.8 ± 0.5
4a	2	C II	0.71400	5180.5	7.2	17.360 ± 0.005	12.81 ± 0.08	-4.55 ± 0.08
5a	3	C II	0.71190	5173.3	0	17.363 ± 0.005	12.77 ± 0.25	-4.60 ± 0.25
6a	3	C II, C III, C IV, Si IV	0.71520	6492.6	0	17.360 ± 0.005	12.86 ± 0.08	-4.50 ± 0.08
CMB prediction by Coc et al. (2015)								-4.611 ± 0.009

[a] ΔAICc = AICc − AICc$_{\text{min}}$, where AICc$_{\text{min}}$ corresponds to the model with the smallest AICc within each of the three subsets of models (no metals, C II, and all metals).

Table 4. Same as in Table 3 but with a single high b-parameter line used to model the broad wings of H I Ly α, as described in Section 3.6.2.

Model	No. components	Metals used	χ^2/d.o.f.	AICc	ΔAICc[a]	log N_{total}(H I)	log N_{total}(D I)	log D/H
1b	1	–	0.72705	5032.6	5.5	17.367 ± 0.005	12.688 ± 0.010	-4.679 ± 0.011
2b	2	–	0.72556	5027.1	0	17.363 ± 0.005	12.685 ± 0.011	-4.678 ± 0.012
3b	3	–	0.72592	5033.6	6.5	17.363 ± 0.005	12.687 ± 0.019	-4.677 ± 0.019
4b	2	C II	0.72344	5231.4	9.8	17.363 ± 0.005	12.704 ± 0.014	-4.659 ± 0.015
5b	3	C II	0.72095	5221.6	0	17.364 ± 0.005	12.687 ± 0.014	-4.677 ± 0.015
6b	3	C II, C III, C IV, Si IV	0.72228	6539.3	0	17.360 ± 0.005	12.700 ± 0.013	-4.660 ± 0.014
CMB prediction by Coc et al. (2015)								-4.611 ± 0.009

[a] Calculated with respect to the models with a broad line fitted.
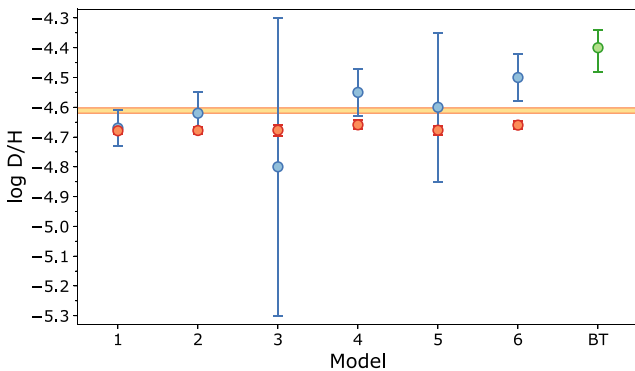
Figure 9. D/H values with 1σ confidence intervals (y-axis) for each of the six pairs of models considered (indicated on the x-axis). The blue points (larger error bars) correspond to models 1a−6a. The orange points (smaller error bars) correspond to models 1b−6b (see Sections 3.6.1 and 3.6.2). The BT measurement is shown at the far right. The horizontal band indicates the 1σ confidence interval for the CMB value from Coc et al. (2015).
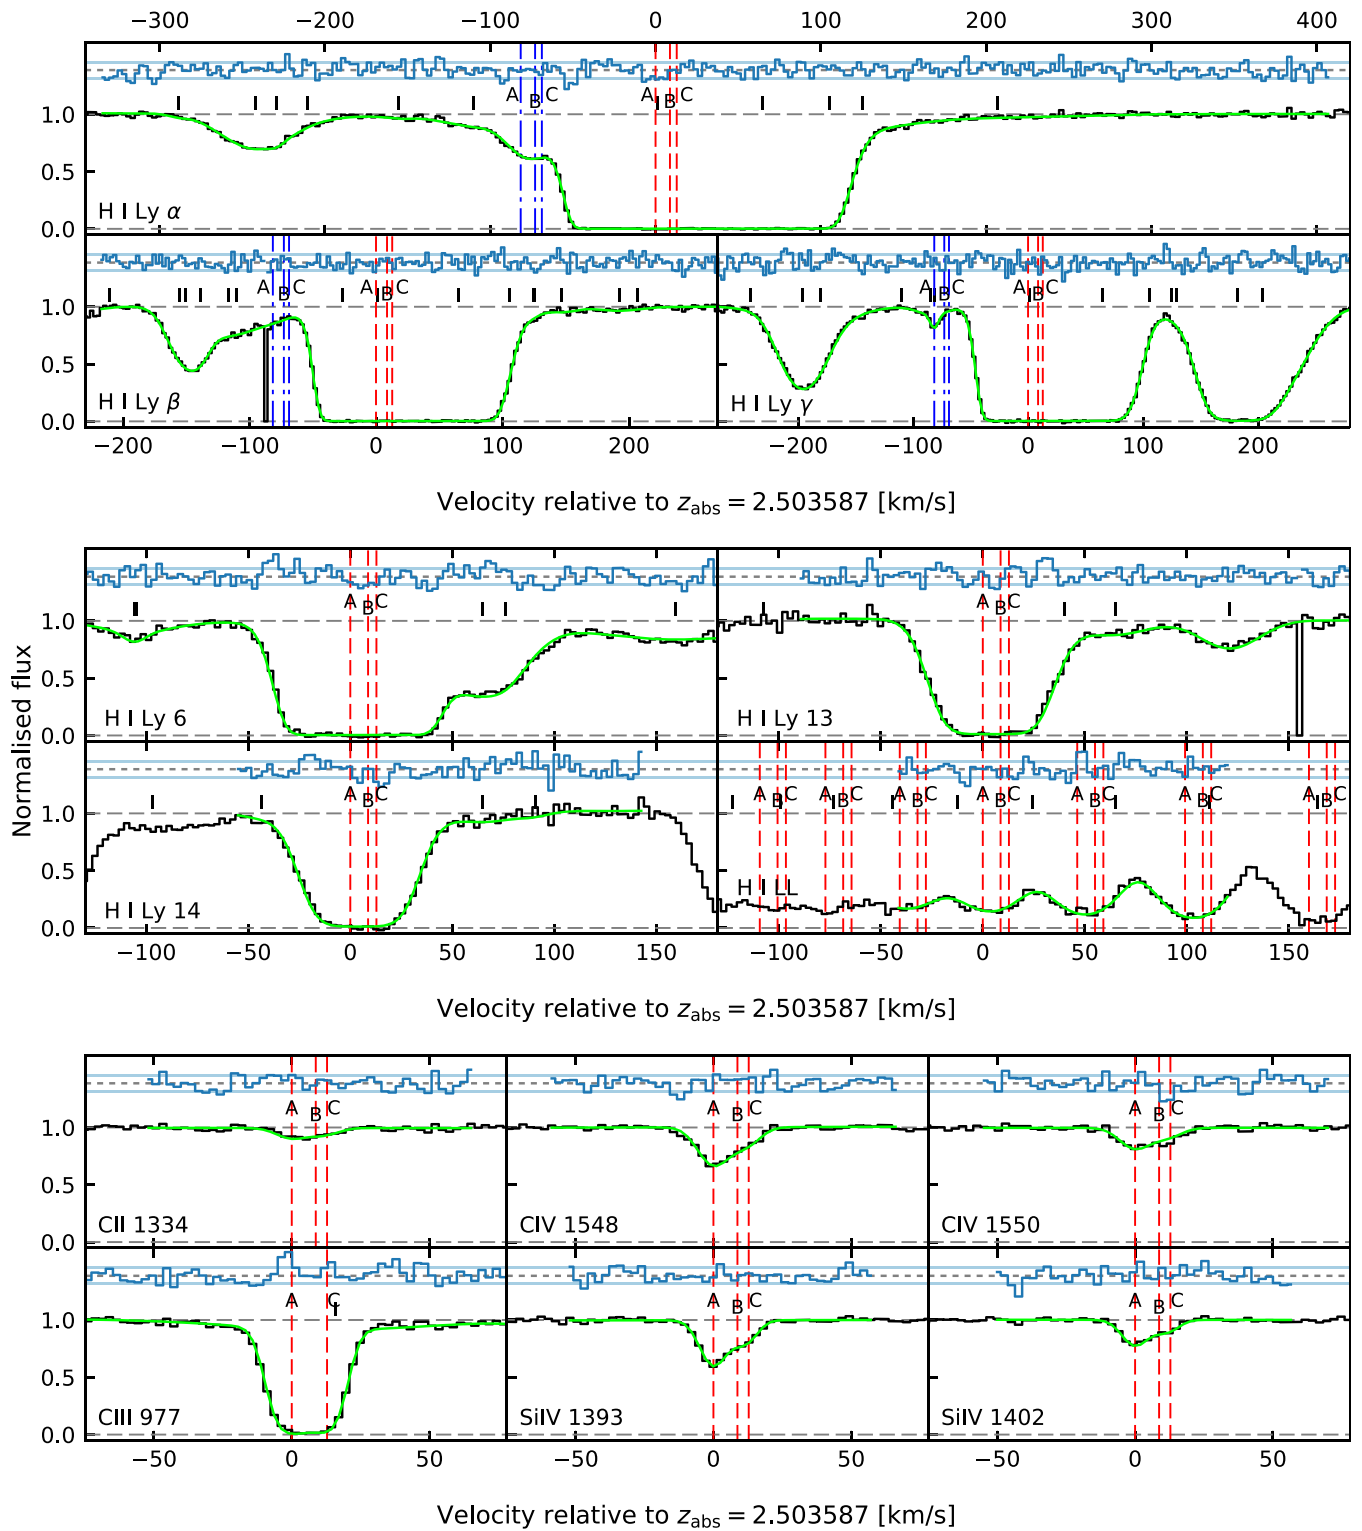
uncertainty on D/H. Hence, if we were to opt for a model that made no use of any metal line information, we would then have to conclude that without metals, one cannot extract D/H with high precision from this absorption system. However, metals are detected and it therefore is sensible to make optimal use of them to derive the best possible information on velocity structure.

Of the three-component models, model 6a uses the most metal lines to estimate the velocity structure and constrain gas temperature and turbulent motions. It is appropriate to use a mix of low and higher ionization metal species to constrain the velocity structure because H I column densities are fitted as unconstrained parameters. If any velocity component has ionization conditions such that all hydrogen is ionized, the H I column density in this component would be iteratively reduced below the threshold and the corresponding H I component would subsequently be rejected by VPFIT. In other words, we do not require that any particular velocity component necessarily exhibits H I absorption. If other velocity components were required that do not have corresponding metal lines, this would be indicated in the appropriate places by poor normalized residuals. This is not seen in the analysis.

We can compare models 3a and 6a. Each has the same number of absorbing components, but no metals were used in fitting model 3a (apart from having used the metal lines to show that all three main components are thermally broadened, see Section 3.6). Fig. 11 illustrates the parameter estimates and associated uncertainties for models 3a and 6a, showing the locations of the three main components, A, B, and C, for both models in the log$N$(H I)−z plane. Models 3a and 6a are consistent in both N and z, provided we associate the largest blue (uniformly shaded) rectangle with component C in model 6a. From Fig. 9, although the D/H error bars for models 3a and 6a overlap, the model 3a error bar is huge compared to model 6a, with a 0.3 dex difference for D/H (see Table 3). The much larger uncertainty on D/H in model 3a is because the total D I column den-

model (model 1a) is strongly disfavoured (ΔAICc = 18.1) with respect to the lowest AICc model (model 3a). Apart from this statistical test, multiple structures are seen in all the metals making a one-component model (model 1a) very unlikely. Two-component models 2a and 4a have 'considerably less' empirical support with respect to the three-component models 3a (ΔAICc = 4.6) and 5a (ΔAICc = 7.2), respectively. With this simple statistical test, we conclude that three-component models are favoured by our data.

Given the three three-component models (models 3a, 5a, and 6a) explored, what are the arguments, if any, to favour one above another? As Table 3 and Fig. 7 indicate, the interloping H I line blends particularly badly with D I for model 3a, leading to the large

(a) Overplotted with the C1×2 spectrum.

Figure 10. The overall absorption profile for model 6a (green solid line), the spectrum (black histogram), and residuals with 1σ confidence range (blue histogram above each fit). Figure (a) illustrates the C1 × 2 data and figure (b) shows the C5 × 1 data. The two co-added spectra illustrated are the highest S/N. However all four co-added spectra were fitted. The transition corresponding to each panel is indicated in the bottom left corner. The main components A, B, and C are indicated by vertical red dashed lines with letters above. The upper portion of the figure illustrates three H I panels, in which the positions of the corresponding deuterium transitions are indicated by blue vertical dash–dotted lines. All the Lyman limit transitions are shown in a single panel (H I LL) centred at Ly 23. The ticks above each fit indicate absorption by additional H I clouds included in the model. The x-axis is the velocity offset from component A. The absence of component B in C III 977 is because VPFIT iteratively reduced its column density until it fell below the acceptance criterion. The sharp drops of the flux at H I Ly β and Ly 13 of the C1 × 2 spectrum and C IV 1550 of the C5 × 1 spectrum are pixels clipped due to cosmic rays.
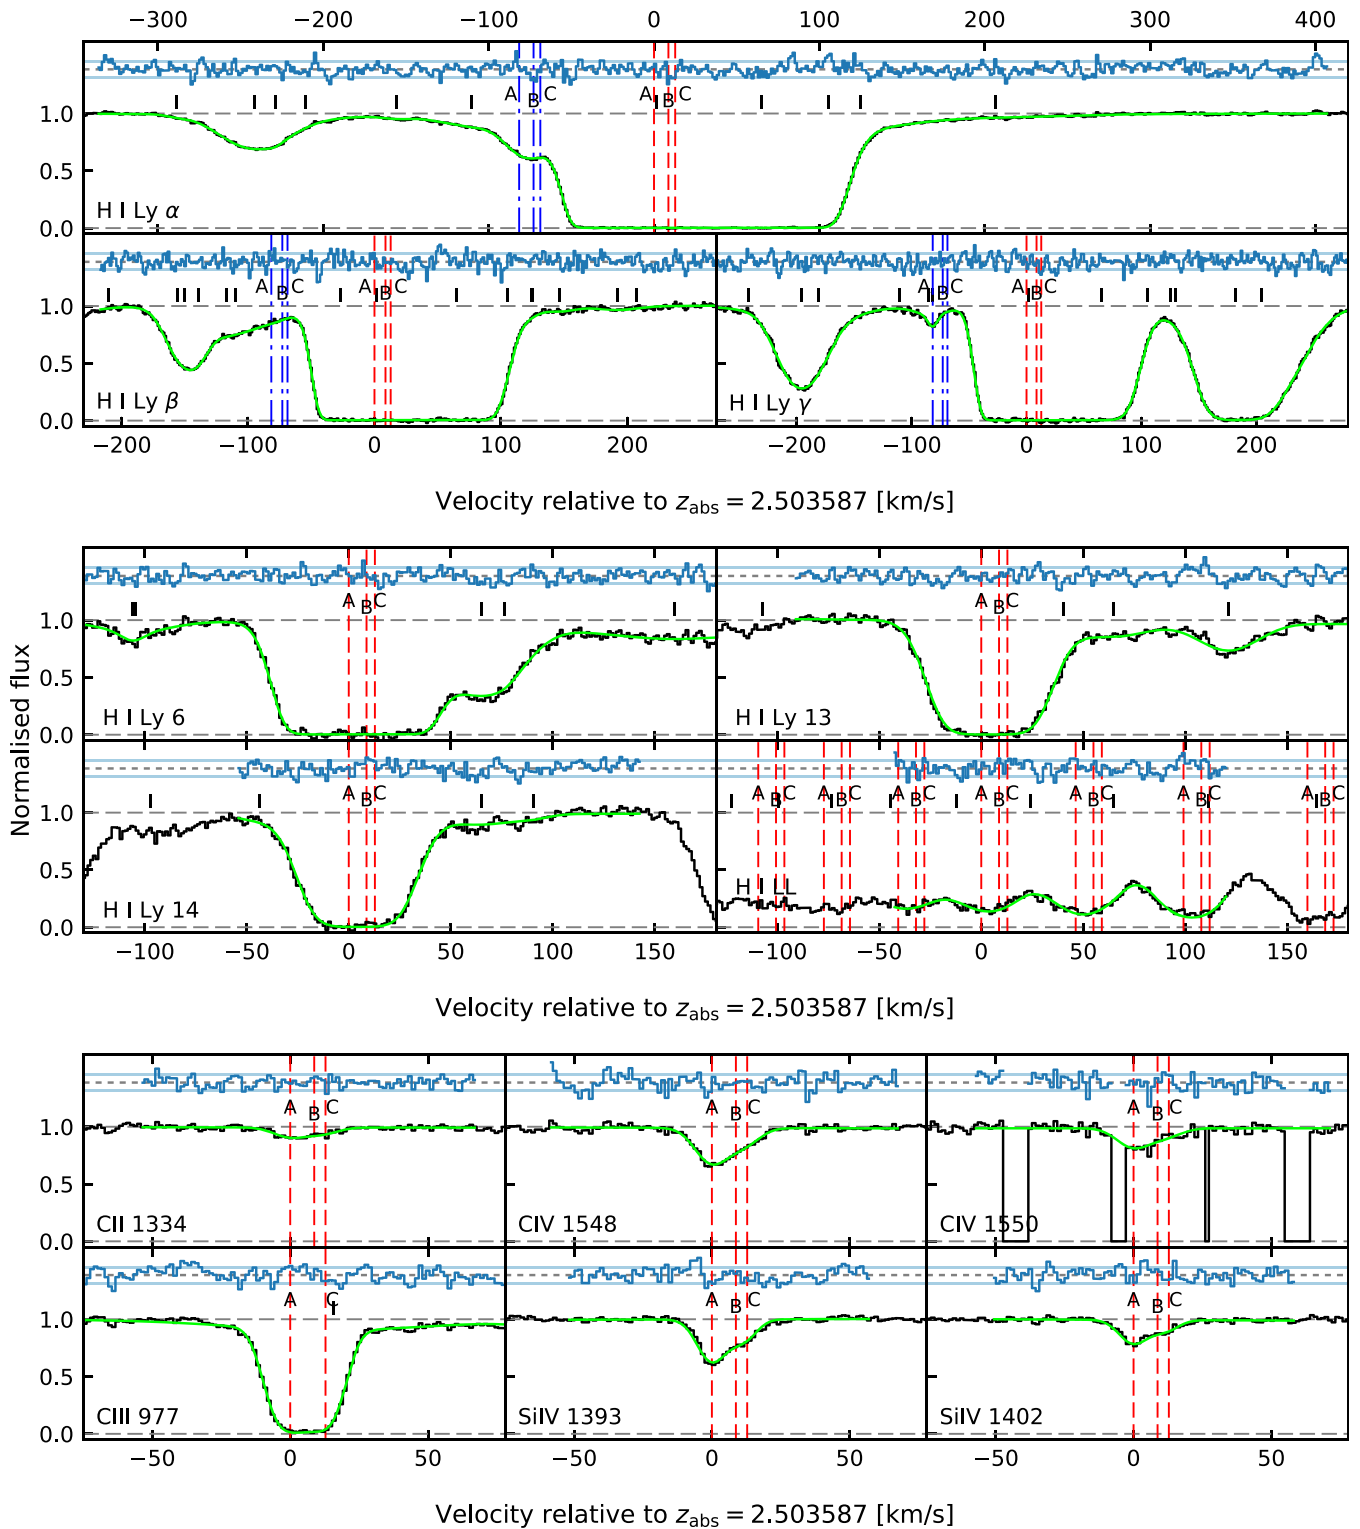
(b) Overplotted with the C5×1 spectrum.

Figure 10 *– continued*

sity is poorly constrained, and this is because of the dramatically worse constraint on kinematic structure available when fitting the Lyman series alone. Out of models 1a to 6a, we therefore argue that 6a is the most compelling.

We now turn to the second set of models, summarized in Table 4 and illustrated (for Ly α) in Fig. 8. Models 1b to 6b were motivated by the apparent symmetry of the residual absorption seen in both wings of the Ly α profile, hinting at the possible presence of one
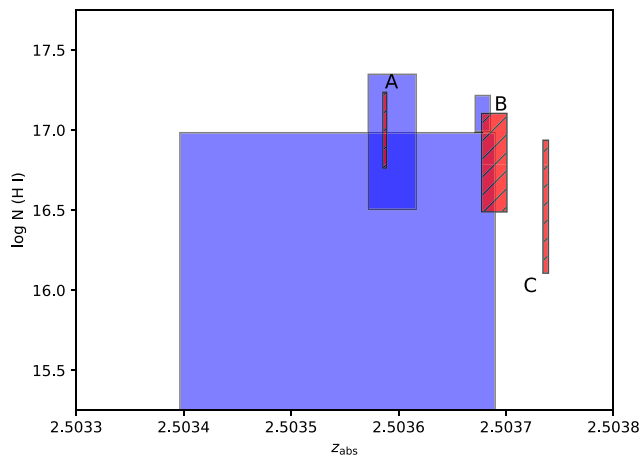
Figure 11. Comparison between the $N(\text{H\,\textsc{i}})$ and z parameter estimates for models 3a (derived by fitting to the Lyman series alone) and 6a (with metals used). Each rectangle illustrates the parameter range for one absorption component (three main components in each model). The red hatched rectangles illustrate model 6a. The labelling A, B, and C assumes that we can associate the largest blue (continuous shaded) rectangle with the right red (hatched) rectangle, and that the remaining two pairs of rectangles are matched to their closest neighbour. The two models are shown to be consistent with each other, although model 3a clearly is considerably less-well determined.

very broad interloper. Such a high b-parameter would be unusual and inconsistent with the general b distribution. Nevertheless, the data itself suggests the possibility, as discussed in Section 3.6.2. The one broad interloper in these models replaces four narrow interlopers used in models 1a to 6a. The interesting aspect of these models is that the broad interloper no longer causes such severe degeneracy between a narrow blend and the deuterium feature. This is reflected by the dramatically smaller uncertainties on D/H for this set of models compared to models 1a to 6a.

Whilst all six models, 1b to 6b, yield consistent D/H values, for the same reasons outlined previously, we identify model 6b as the most reliable in this set. Since models 6a and 6b make use of the same data sets, we can compare AICc values, and see that, formally, 6a is the preferred model. However, there is some degree of arbitrariness in that choice as it would be relatively straightforward to make slight refinements to the model (i.e. by including one additional interloper) to reduce both χ^2 and AICc. For that reason we think it reasonable to identify model 6b as an acceptable solution and use it to estimate a systematic contribution to the final error budget and our final D/H.

3.9 Systematic uncertainty and the final D/H estimate

In order to estimate the systematic error associated with the choice of the velocity structure (i.e. the blending model), we adopt a systematic error corresponding to half the difference between the model 6a and 6b results, $\sigma_{sys}(\log D/H) = 0.08$. Since this gives a crude estimate of the total rather than the systematic error, using this as the systematic component is conservative. This additional error is random in nature, so we add it in quadrature to the statistical error for models 6a and 6b. Our final D/H for the LLS analysed in this paper is a weighted mean of both values,

$$\log(D/H) = -4.606 \pm 0.066$$

or

$$D/H = 2.48^{+0.41}_{-0.35} \times 10^{-5}$$

in excellent agreement with the CMB value from Coc et al. (2015) of D/H $= (2.45 \pm 0.05) \times 10^{-5}$.

4 DISCUSSION

4.1 Comparison with previous measurement of BT

In the previous analysis of this LLS, BT explored six different models. However, no metals were used to constrain the velocity structure. Their preferred model is a three-component model, as is ours (although we note that their model and ours differ not only in terms of additional blends impacting on the D line, as we have described in Section 3.6.1, but also elsewhere – compare BT's fig. 8 with our Fig. 10).

BT give a value of $\log D/H = -4.40^{+0.06}_{-0.08}$. In some of our models, the fitting process positioned a blend close to the D\,\textsc{i} line, degrading the precision with which one can measure $N(\text{D\,\textsc{i}})$ and hence D/H. BT's models do not include this blend. This is why our final error D/H estimate is not a dramatic improvement on BT's, despite the far better S/N. As discussed in Section 3.6.1, if we artificially fix the parameters of this blend, the severe impact on the apparent D/H precision becomes apparent; for our model 3a we would derive an uncertainty on $\log D/H$ of 0.027 and for model 6a, an uncertainty on $\log D/H$ of 0.014.

4.2 The primordial deuterium abundance from a sample of measurements

Table 5 presents 15 D/H measurements from the literature. This list is the same as the 13 listed in table 4 of Riemer-Sørensen et al. (2017) plus two further systems: Srianand et al. (2010) and Cooke et al. (2016). It is well known that the apparent scatter in D/H measurements exceeds that expected on the basis of the statistical error estimates. Table 5 includes two 'low' D/H estimates that one might opt (subjectively) to exclude as being deviant: Srianand et al. (2010) and Pettini & Bowen (2001). However we choose here to apply a modified Least Trimmed Squares procedure (LTS; Rousseeuw 1984) to estimate D/H from the sample, which allows us to avoid such subjective decisions as to which measurements are 'good' and which are unreliable.

The procedure we adopt is as follows. As an initial step, we adopt an LTS trimming fraction of 0.85, i.e. we allow LTS to discard 15 per cent of the points (in this case 2). We thus allow LTS to select the 'best' $k = 13$ from $n = 15$ values. LTS identifies entries 7 (Srianand et al. 2010) and 15 (Pettini & Bowen 2001) in Table 5 as the most discrepant. To allow for excess scatter in the sample, above the statistical error, at each LTS iteration (that is, for each combination of k from n), we iteratively increase all error bars by adding an additional term in quadrature, σ_{rand}, such that the normalized χ^2 over all points is unity. The 'successful' combination of 13 points is taken to be that having the smallest value of σ_{rand}. Interestingly, this value turns out to be zero, i.e. the scatter amongst the remaining 13 values is consistent, on average, with the published errors. The final result, a weighted mean of these 13 D/H values, is

$$(D/H)_p = (2.545 \pm 0.025) \times 10^{-5}. \tag{2}$$

The value above, based on a relatively small sample of quasar observations, is seen to be in good agreement with the CMB values derived by Coc et al. (2015), D/H $= (2.45 \pm 0.05) \times 10^{-5}$, and

Table 5. D/H measurements used to estimate $\Omega_b h^2$. A Least Trimmed Squares procedure was used to eliminate outliers (rejection of 2 points, being 15 per cent of the total set of 15). LTS rejected the measurements of J1337+3152 and Q2206−199. See Section 4.2 for details.

Quasar	z_{em}	z_{abs}	$\log N(\text{H\,\textsc{i}})$	[X/H], X	D\,\textsc{i}/H\,\textsc{i} $(\times 10^5)$	References
HS 0105+1619	2.65	2.537	19.426 ± 0.006	−1.77 O	$2.58^{+0.16}_{-0.15}$	Cooke et al. (2014)
J0407−4410	3.02	2.621	20.45 ± 0.10	−1.99 O	$2.8^{+0.8}_{-0.6}$	Noterdaeme et al. (2012)
Q0913+072	2.79	2.618	20.312 ± 0.008	−2.40 O	$2.53^{+0.11}_{-0.10}$	Cooke et al. (2014)
Q1009+2956	2.63	2.504	17.362 ± 0.005^a	−2.5 Si^b	$2.48^{+0.41}_{-0.35}$	This work
J1134+5742	3.52	3.411	17.95 ± 0.05	< -4.2 Si	$2.0^{+0.7}_{-0.5}$	Fumagalli, O'Meara & Prochaska (2011)
Q1243+3047	2.56	2.526	19.761 ± 0.026	−2.77 O	2.39 ± 0.08	Cooke et al. (2018)
J1337+3152	3.17	3.168	20.41 ± 0.15	−2.68 Si	$1.2^{+0.5}_{-0.3}$	Srianand et al. (2010)
SDSS J135803.97+034936.0	2.89	2.853	20.524 ± 0.006	−2.80 O	2.62 ± 0.07	Cooke et al. (2016)
J1358+6522	3.17	3.067	20.495 ± 0.008	−2.33 O	2.58 ± 0.07	Cooke et al. (2014)
J1419+0829	3.03	3.050	20.392 ± 0.003	−1.92 O	2.51 ± 0.05	Cooke et al. (2014)
J1444+2919	2.66	2.437	19.983 ± 0.010	−2.04 O	$1.97^{+0.33}_{-0.28}$	Balashev et al. (2016)
J1558−0031	2.82	2.702	20.75 ± 0.03	−1.55 O	$2.40^{+0.15}_{-0.14}$	Cooke et al. (2014)
Q1937−1009^c	3.79	3.256	18.09 ± 0.03	−1.87 O	$2.45^{+0.30}_{-0.27}$	Riemer-Sørensen et al. (2015)
		3.572	17.925 ± 0.006	−2.26 O	2.62 ± 0.05	Riemer-Sørensen et al. (2017)
Q2206−199	2.56	2.076	20.436 ± 0.008	−2.04^d O	1.65 ± 0.35	Pettini & Bowen (2001)
CMB prediction					2.45 ± 0.05	Coc et al. (2015)

^{a}Mean of the values for models 6a and 6b .

^{b}Measured by BT.

^{c}There are two absorption systems on the sight-line towards Q1937−1009 with identified D\,\textsc{i} lines.

^{d}Measured by Pettini et al. (2008).

by Marcucci et al. (2016), D/H $(2.49 \pm 0.03 \pm 0.03) \times 10^{-5}$. The differences between these two CMB values arise from the different nuclear reaction rates used to determine the primordial abundances (see Cooke et al. 2016; Riemer-Sørensen & Sem Jenssen 2017, for a comprehensive discussion). The precision of the quasar data is now approaching the point such that observations based on quasar absorption systems may provide a method of determining nuclear rates.

4.3 The baryon density of the Universe

Using the fitting formula for the BBN calculations (Coc et al. 2015, private communication):

$$(\text{D/H})_p = (2.45 \pm 0.04) \times 10^{-5} \left(\frac{\Omega_b h^2}{0.02225} \right)^{-1.657},$$

we infer the baryon density of the Universe, the statistical uncertainty, and the uncertainty from nuclear data parameters:

$$\Omega_b h^2 = 0.02174 \pm 0.00013_{\text{qso}} \pm 0.00021_{\text{nucl}}.$$

The overall uncertainty is now dominated by the nuclear data imprecision rather than low statistics of the high-redshift quasar data. Comparing our result to the *Planck* 2015 TT+lowP+lensing value (Planck Collaboration XIII 2016, see their table 4),

$$\Omega_b h^2 = 0.02226 \pm 0.00023,$$

i.e. there is a marginal tension at the 1.6σ level with the result we present in this paper.

5 CONCLUSIONS

The Lyman limit system at $z_{abs} = 2.504$ towards Q1009+2956 (or J1011+2941) was previously used to measure the primordial deuterium abundance by BT. It has long been considered one of the more robust measurements. We embarked on a re-measurement of

this system because of newer, far higher S/N data. Whilst our a priori expectation was a correspondingly more precise measurement, the higher S/N data in fact revealed a more complex velocity structure than had previously been thought. The consequence is lower precision than expected. Further, the higher S/N of the newer data (up to 147 per pixel in the continuum level of Ly α, compared to a peak S/N of about 60 in the BT data) requires us to consider a broader range of absorption system models. Taking into account both statistical uncertainties and the additional uncertainties associated with velocity structure ambiguity, our final result for this LLS is

$$\text{D/H} = 2.48^{+0.41}_{-0.35} \times 10^{-5}$$

in excellent agreement with the CMB value.

A weighted mean of 13 D/H values from the literature (including the result reported here) gives a primordial D/H value of

$$(\text{D/H})_p = (2.545 \pm 0.025) \times 10^{-5}.$$

This leads to a baryon density of the Universe of

$$\Omega_b h^2 = 0.02174 \pm 0.00013_{\text{qso}} \pm 0.00021_{\text{nucl}}$$

marginally inconsistent with the *Planck* CMB data. Further quasar data as well as experimental improvement in the nuclear data are required to establish whether the marginal discrepancy is a random fluctuation or has a physical origin.

TRO), through project number CE110001020. EOZ is grateful for the hospitality of the Institute of Theoretical Astrophysics at the University of Oslo during his visit in 2016 January–February. JKW thanks the Centre for Mathematical Sciences in Cambridge for an enjoyable and productive Sabbatical visit. This work is based on observations collected with the Keck Observatory Archive (KOA), which is operated by the W. M. Keck Observatory and the NASA Exoplanet Science Institute (NExScI), under contract with the National Aeronautics and Space Administration. A significant part of the analysis and figures were done using publicly available PYTHON packages: MATPLOTLIB (Hunter 2007), NUMPY (Van Der Walt, Colbert & Varoquaux 2011), SCIPY (Oliphant 2007), ASTROPY (Astropy Collaboration 2013), BARAK[11] by Neil Crighton, and QSCAN (Dumont 2017). This research has made use of NASA's Astrophysics Data System. EOZ is also grateful to GitHub for the Student Developer Pack.

REFERENCES

Akaike H., 1974, IEEE Trans. Autom. Control, 19, 716

Astropy Collaboration, 2013, A&A, 558, A33

Balashev S. A., Zavarygin E. O., Ivanchik A. V., Telikova K. N., Varshalovich D. A., 2016, MNRAS, 458, 2188

Burles S., Tytler D., 1998, ApJ, 507, 732

Burnham K. P., Anderson D. R., 2002, Model Selection and Multimodel Inference: A Practical Information-Theoretic Approach. Springer New York

Carswell R. F., Webb J. K., 2014, Astrophysics Source Code Library, record ascl:1408.015

Coc A., Petitjean P., Uzan J.-P., Vangioni E., Descouvemont P., Iliadis C., Longland R., 2015, Phys. Rev. D, 92, 123526

Cooke R. J., Pettini M., Jorgenson R. A., Murphy M. T., Steidel C. C., 2014, ApJ, 781, 31

Cooke R. J., Pettini M., Nollett K. M., Jorgenson R., 2016, ApJ, 830, 148

Cooke R., Pettini M., Steidel C. C., 2018, ApJ, 855, 102

Crighton N. H. M., Webb J. K., Ortiz-Gil A., Fernández-Soto A., 2004, MNRAS, 355, 1042

Cyburt R. H., Fields B. D., Olive K. A., Yeh T.-H., 2016, Rev. Mod. Phys., 88, 015004

Dumont V., 2017, qscan: Quasar Spectra Scanning Tool, https://doi.org/10.5281/zenodo.437903

Dumont V., Webb J. K., 2017, MNRAS, 468, 1568

Evans T. M., Murphy M. T., 2013, ApJ, 778, 173

Fumagalli M., O'Meara J. M., Prochaska J. X., 2011, Science, 334, 1245

Hunter J. D., 2007, Comput. Sci. Eng., 9, 90

Iocco F., Mangano G., Miele G., Pisanti O., Serpico P. D., 2009, Phys. Rep., 472, 1

Kramida A., Yu. Ralchenko Reader J. NIST ASD Team 2015, NIST Atomic Spectra Database (ver. 5.3), [Online]. National Institute of Standards and Technology, Gaithersburg, MD. Available: http://physics.nist.gov/asd [2017 May 25].

Marcucci L. E., Mangano G., Kievsky A., Viviani M., 2016, Phys. Rev. Lett., 116, 102501

Molaro P., Levshakov S. A., Monai S., Centurión M., Bonifacio P., D'Odorico S., Monaco L., 2008, A&A, 481, 559

Morton D. C., 2003, ApJS, 149, 205

Murphy M. T., 2016, UVES_popler: POst-PipeLine Echelle Reduction software, https://doi.org/10.5281/zenodo.56158

Murphy M. T., Webb J. K., Flambaum V. V., Churchill C. W., Prochaska J. X., 2001, MNRAS, 327, 1223

Noterdaeme P., López S., Dumont V., Ledoux C., Molaro P., Petitjean P., 2012, A&A, 542, L33

Oliphant T. E., 2007, Comput. Sci. Eng., 9, 10

Pettini M., Bowen D. V., 2001, ApJ, 560, 41

Pettini M., Zych B. J., Steidel C. C., Chaffee F. H., 2008, MNRAS, 385, 2011

Planck Collaboration XIII, 2016, A&A, 594, A13

Rahmani H. et al., 2013, MNRAS, 435, 861

Riemer-Sørensen S., Sem Jenssen E., 2017, Universe, 3, 44

Riemer-Sørensen S. et al., 2015, MNRAS, 447, 2925

Riemer-Sørensen S., Kotuš S., Webb J. K., Ali K., Dumont V., Murphy M. T., Carswell R. F., 2017, MNRAS, 468, 3239

Rousseeuw P. J., 1984, J. Am. Stat. Assoc., 79, 871

Srianand R., Gupta N., Petitjean P., Noterdaeme P., Ledoux C., 2010, MNRAS, 405, 1888

Steigman G., 2006, J. Cosmol. Astropart. Phys., 10, 16

Van Der Walt S., Colbert S. C., Varoquaux G., 2011, Comput. Sci. Eng., 13, 22

Vogt S. S., 1994, HIRES user's manual (UCO/Lick Observatory technical report), http://www2.keck.hawaii.edu/inst/hires/manual1.pdf

Vogt S. S. et al., 1994, in Crawford D. L., Craine E. R., eds, SPIE Conf. Ser. Vol. 2198, Instrumentation in Astronomy VIII. SPIE, Bellingham, p. 362

Vogt S. S., Hill G., Kibrick R., 2008, HIRES CCD Upgrade User Notes, http://www2.keck.hawaii.edu/inst/hires/usernotes.pdf

Whitmore J. B., Murphy M. T., 2015, MNRAS, 447, 446

Whitmore J. B., Murphy M. T., Griest K., 2010, ApJ, 723, 89

SUPPORTING INFORMATION

Supplementary data are available at *MNRAS* online.

Please note: Oxford University Press is not responsible for the content or functionality of any supporting materials supplied by the authors. Any queries (other than missing material) should be directed to the corresponding author for the article.

APPENDIX A: HIRES DATA REDUCTION ISSUES

A1 Wavelength distortions

A1.1 Night-sky line correction

There is an option in MAKEE to apply an offset to the wavelength array (in velocity space) to correct any possible wavelength shift between the ThAr and science exposures. In order to determine the offset that generally seems to be between 0.1 and 0.4 pixels (unbinned), MAKEE uses the night-sky lines in the background of the object exposure. The problem arises for the current HIRES detector consisting of three CCDs. Since each of the three CCD exposures is reduced separately, offsets are measured independently. Normally, for wavelength coverage of about 3100−6000 Å there are enough sky lines to determine the offset on the third (red) CCD chip only. To avoid introduction of an intra-chip velocity offset we used the following approach. We, first, reduced all three CCD separately, giving MAKEE a chance to determine the offset. Then, in case the offset was different for different chips (or if the offset was determined only for the red chip as was usually the case), we calculated a weighted mean value of the offset. Finally, we repeated the reduction specifying the offset manually.

A1.2 Air to vacuum correction

As found in Murphy et al. (2001), MAKEE applies an inaccurate, air to vacuum conversion formula. Instead of using the Edlen formula (as used for the wavelength-calibration ThAr lines), MAKEE adopts a

[11] https://pypi.python.org/pypi/Barak/0.3.2

Cauchy dispersion relation (for details see Paragraph 2.9 in Murphy et al. 2001). This produces a wavelength-dependent distortion. To resolve this issue, during the data reduction procedure with MAKEE, we prevented any air to vacuum conversion being applied and instead a correction was applied (at the order-combination stage) in UVES_POPLER, which uses the correct Edlen formula.

A2 Blemishes on the CCDs

A2.1 Exposure time dependence of bad regions on the legacy CCD

It is generally known fact that the legacy 'Tektronix' CCD of HIRES had a few prominent cosmetic defects such as a large felt-tip pin mark near the centre of the CCD (Vogt 1994), three 'bleeding' bad regions, a 'hot' corner and a few bad columns. MAKEE attempts to mask these regions by default. However, we found that the bleeding regions and the hot corner increase in size in both spatial and dispersion directions with increasing exposure time. Occasionally, they extend outside of the masked regions and affect adjacent parts of the spectrum as well as adjacent orders. To tackle this problem we visually inspected all the legacy exposures and excluded all these affected regions.

A2.2 Bad columns on the new CCD

In contrast to the legacy CCD of HIRES, the current one (upgraded in August 2004) has less cosmetic defects (see item 7 in Vogt, Hill & Kibrick 2008). The most problematic ones are bad partial columns that, in contrast to the legacy CCD, go along the Echelle orders. If an Echelle order happens to be placed on a bad column, significant part of its spectrum is affected and the absorption line profiles in this region become unreliable. It can also mimic a fluctuation of the continuum. In turn, this can affect the normalization process in the vicinity of this region. In order to avoid any possible bias we identified the problematic regions by visual inspection of the flat-field spectra and then excluded all affected regions. We managed to identify four bad columns[12] on the first (blue) chip, one bad column on the second (green) chip, and two bad columns on the third (red) chip. All these columns start from the blue edge (in dispersion direction) of the chips and cover up to three-quarters of the CCDs.

A3 Scattered light

Most HIRES images of bright objects (quartz lamp, bright star) are affected by scattered light (see e.g. Vogt 1994; Vogt et al. 2008). Having visually inspected all the flat-field frames we used in this work using ds9[13], we identified at least three scattered light features in the legacy data and at least seven features in the current CCD exposures. Occasionally, this light affects the spectrum significantly, in our case – exposures from 2004 November 5 and 2005 March 31. The scattered light on the first (blue) chip, where the deuterium lines fall, causes an increase of the flux in the flat-field images by up to 10 per cent introducing a systematic bias during the flat-fielding process. To avoid any possible bias we completely excluded the affected parts of these spectra from the analysis.

[12]There is also a fifth column. However, it is located close to the blue edge of the chip in spatial direction so that no orders are usually extracted from this part of the CCD.
[13]http://ds9.si.edu/site/Home.html

Table B1. List of absorption systems and corresponding species tentatively identified in the spectrum of Q1009+2956.

Redshift	Species[a]
-0.0001^b	Na I, Ca II, Ti II
0.8599	Fe II, Mg II
1.1117	C IV (?), Al II, Al III, Fe II, Mg I (?), Mg II
1.1146	C IV, Mg II
1.1707	C IV
1.1729	C IV, Al II (?), Al III, Mg II
1.2709	C IV, Al II, Al III, Mg II, Fe II
1.5211	C IV, Si IV
1.5583	C IV (?), Si IV (?)
1.5986	C II, C IV, Si II, Si IV, Al II, Al III, Fe II, N V (?)
1.9061	C IV
2.1437	C IV, Si III (?), Si IV (?)
2.2062	C IV
2.2073	C IV
2.2533	C III (?), C IV, Si II, Si III, Si IV, Al II (?)
2.3612	C IV
2.4070	C II, C III, C IV, Si II, Si III, Si IV, Al II, Al III, Fe II, Ni III (?)
2.4271	C III, C IV, Si IV
2.4290	C II, C III, C IV, Si II, Si III, Si IV, Al II, Al III, Ni III (?)
2.5036	C II, C III, C IV, Si II, Si III, Si IV
2.5533	C IV
2.5726	C III (?), C IV
2.6059	C IV, Al II (?)
2.6490	C III, C IV, Si IV, N V

[a]Question marks indicate possible detections. Usually, it corresponds to species with only one transition available, which is also blended significantly.
[b]Absorption by the interstellar medium.

APPENDIX B: ABSORPTION SYSTEMS TOWARDS Q1009+2956

In Table B1, we provide a list of absorption systems and corresponding species identified in the spectrum of Q1009+2956. The technique used in the identification process is described in Section 3.5.

APPENDIX C: DESCRIPTION OF THE SUPPLEMENTARY FILES

All the VPFIT files with the best-fitting parameters and their uncertainties for each tested model, the co-added spectra and other supplementary files are available in online supplementary materials and on GitHub (the links are given in Section 3.7).

abs_syst/, All the supplementary files come in four separate folders: data/, vpfit/, and voffset/:

(i) data/ contains four text files representing four co-added spectra (see Section 2.3). Each text file has three columns: wavelength in vacuum-heliocentric frame, normalized flux,). Each text file has three columns: wavelength in vacuum-heliocentric frame, normalized flux,

(ii) vpfit/ contains an atom.dat file with atomic data used, vp_setup.dat VPFIT setting files and resulted VPFIT output files for all the considered models. The latter contain all parameters for each model. Numeration of the models matches the one given in Sections 3.6, 3.6.1, and 3.6.2. In addition, the metals/ sub-folder contains VPFIT output and vp_setup.dat files for the thermal and turbulent fits of metal lines mentioned in Section 3.6. We refer an

interested reader to the VPFIT manual[14] for details on the format of these files.

(iii) voffset/ contains a file with velocity offsets in km s^{-1} for each individual exposure with respect to the one obtained on 2008 March 29 for each individual exposure with respect to the one obtained on 29 March2008 UT). See Section 2.3 for description of how these offsets were calculated. Names of the exposures are given in the KOA format[15].

(iv) abs_syst/ contains a list of identified species per absorption system described in Section 3.5. This is just an electronic version of Table B1.

This paper has been typeset from a TEX/LATEX file prepared by the author.

[14] Available at https://www.ast.cam.ac.uk/~rfc/vpfit.html
[15] Described at https://www2.keck.hawaii.edu/koa/public/faq/koa_faq.php#U9