

A SEARCH FOR OXYGEN IN THE LOW-DENSITY LYMAN- α FOREST USING THE SLOAN DIGITAL SKY SURVEY

MATTHEW M. PIERI¹, STEPHAN FRANK^{1,2}, SMITA MATHUR¹, DAVID H. WEINBERG¹,
DONALD G. YORK^{3,4}, AND BENJAMIN D. OPPENHEIMER⁵

Draft version April 28, 2010

ABSTRACT

We use 2167 Sloan Digital Sky Survey (SDSS) quasar spectra to search for low-density oxygen in the Intergalactic Medium. Oxygen absorption is detected on a pixel-by-pixel basis by its correlation with Ly α forest absorption. We have developed a novel Locally Calibrated Pixel (LCP) search method that uses adjacent regions of the spectrum to calibrate interlopers and spectral artifacts, which would otherwise limit the measurement of O VI absorption. Despite the challenges presented by searching for weak O VI within the Lyman- α forest in spectra of moderate resolution and signal-to-noise, we find a highly significant detection of absorption by oxygen at $2.7 < z < 3.2$ (the null hypothesis has a $\chi^2 = 80$ for 9 data points).

We interpret our results using synthetic spectra generated from a lognormal density field assuming a mixed quasar-galaxy photoionizing background (Haardt & Madau 2001) and that it dominates the ionization fraction of detected O VI. The LCP search data can be fit by a constant metallicity model with $[\text{O}/\text{H}] = -2.15^{+0.07}_{-0.09}$, but also by models in which low-density regions are unenriched and higher density regions have a higher metallicity. The density-dependent enrichment model by Aguirre et al. (2008) is also an acceptable fit. All our successful models have similar mass-weighted oxygen abundance, corresponding to $[(\text{O}/\text{H})_{\text{MW}}] = -2.45 \pm 0.06$. This result can be used to find the cosmic oxygen density in the Ly α forest, $\Omega_{\text{Oxy,IGM}} = 1.4(\pm 0.2) \times 10^{-6} \approx 3 \times 10^{-4} \Omega_b$. This is the tightest constraint on the mass-weighted mean oxygen abundance and the cosmic oxygen density in the Ly α forest to date and indicates that it contains $\approx 16\%$ of the total expected metal production by star formation up to $z = 3$ (Bouché et al. 2007).

Subject headings: galaxies: formation — intergalactic medium — quasars: absorption lines

1. INTRODUCTION

The Intergalactic Medium (IGM) is the spatially undulating reservoir of baryons from which all galaxies form. During the dark ages, before formation of the first stars, this reservoir was composed of elements in proportions arising from Big Bang Nucleosynthesis: hydrogen, helium and trace amounts of lithium and beryllium. Metals have been produced since the formation of the first stars and some of these metals were carried into the IGM. They are observed in the Lyman- α forest seen in quasar (QSO) absorption spectra at redshifts $z \lesssim 6$ (e.g. Meyer & York (1987); Cowie et al. (1995); Songaila & Cowie (1996); Ellison et al. (2000); Schaye et al. (2003); Pieri et al. (2006)). Progress has also been made reproducing this enrichment in models of galaxy formation with extragalactic winds (e.g. Mac Low & Ferrara (1999); Madau, Ferrara, & Rees (2001); Oppenheimer & Davé (2006); Pieri, Martel, & Grenon (2007); Pieri & Martel (2007)).

Much has been learned about the degree, extent and

evolution of metal enrichment of the IGM, but questions still remain unanswered. What is the precise level of enrichment in the forest, and is it consistent with that seen at the current epoch? Are underdense systems enriched? Are all enriched regions spatially associated with galaxies? In this paper, we attempt to place tighter limits on the level of enrichment in the IGM as a whole, more specifically the low-density regime ($1 \lesssim \rho/\bar{\rho} \lesssim 10$), using QSO spectra from the Sloan Digital Sky Survey (SDSS, York et al. 2000). We have searched for O VIa ($\lambda 1032\text{\AA}$) at redshifts $2.7 < z < 3.2$. At the densities analyzed, O VIa is expected to be the strongest heavy element transition in the optical band (Rauch et al. 1997; Hellsten et al. 1998). However, the line happens to reside in the same wavelength range as the Lyman series lines at other redshifts resulting in severe confusion as to the identity of the absorption. Any search for O VI must involve a sophisticated approach to overcoming this problem.

Fitting of narrow lines or lines that also show Ly α absorption have been performed by various authors (Carswell et al. 2002; Simcoe et al. 2002, 2004; Bergeron & Herbert-Fort 2005) in order to find individual O VI systems, where high resolution spectra are available. In contrast, pixel correlation searches (e.g. Cowie & Songaila 1998; Davé et al. 1998; Aguirre et al. 2002; Schaye et al. 2003; Pieri & Haehnelt 2004) are directed toward finding metals distributed throughout the forest without identifying individual systems. In the case of the search for O VI, they are limited by cosmic variance in the contaminating Lyman lines and noise. The

¹ Department of Astronomy, The Ohio State University, 140 West 18th Avenue, Columbus, OH 43210, USA; mpiri@astronomy.ohio-state.edu

² Observatoire Astronomique de Marseille-Provence, Pôle de l'Étoile Site de Château-Gombert, 38, rue Frédéric Joliot-Curie 13388 Marseille cedex 13, France

³ Department of Astronomy and Astrophysics, University of Chicago, Chicago, IL 60637, USA

⁴ Enrico Fermi Institute, University of Chicago, Chicago, IL 60637, USA

⁵ Astronomy Department, University of Arizona, Tucson, AZ 85721, USA

standard pixel (SP) correlation search (sometimes known as ‘pixel optical depth’) for O VI (e.g. Aguirre et al. 2002; Pieri & Haehnelt 2004; Aguirre et al. 2008, hereafter A08) is not, however, a suitable approach for the detection of metals in the SDSS Ly α forest because of the moderate resolution ($R = \lambda/\Delta\lambda = 1800$) and signal-to-noise ratio. A pixel correlation search is possible in principle, but modifications to the approach are required. We have developed a tailored pixel method, which we refer to as the Locally Calibrated Pixel (LCP) search method. Our statistical analysis of O VI absorption in the low density IGM complements the direct search for strong O VI absorption in the same spectra (Frank et al. 2009a,b).

While interpretation of pixel searches using high resolution, high S/N spectra is more straightforward, the small amount of high quality data leaves the conclusions subject to uncertainty owing to cosmic variance. The use of the large number of SDSS spectra (> 2000 used here) with the LCP search method offers a large reduction in this variance (and potentially an increase in overall S/N).

All pixel correlation searches require simulations in order to ascertain the gas mass densities of enriched systems, the number densities of ions required, the ionization fraction (and so metallicity of those systems), and the impact of observational limitations and errors. This is particularly important here as the resolution of the SDSS spectra introduces a significant uncertainty in the gas mass density of enriched systems. To address this issue and interpret our search results, we have simulated each SDSS spectrum used 10 times for various models of oxygen enrichment, totaling more than 4.5 million synthetic SDSS spectra.

Noise levels in individual SDSS spectra are higher than those of high resolution spectra used in previous studies, and with this (along with the lower resolution) comes more uncertainty in the continuum fitting. This problem becomes increasingly severe at the blue end of the spectrograph, where most of the signal of O VI lies. The correlation search technique that we outline below is calibrated to the flux level in the local spectrum to take into account variations in the continuum (hence the name ‘locally calibrated pixel search’).

This paper is set out as follows: in §2 we describe the data set used, in §3 we describe the LCP search method and results, in §4 we describe our interpretation of these results using simulated spectra, and in §5 we describe our measure of mass-weighted mean oxygen abundance and the cosmic oxygen density. This is followed by a discussion of the ionization corrections and other searches and, finally, the conclusions.

Following standard notation, we write $[O/H] = \log(O/H) - \log(O/H)_{\odot}$, where O/H is the ratio of oxygen to hydrogen (by number). We adopt $\log(O/H)_{\odot} = -3.13$ estimated for the solar envelope by Grevesse et al. (1996), for ease of comparison with other studies of oxygen enrichment of the IGM. While some analyses based on more detailed stellar atmosphere models have argued for lower $(O/H)_{\odot}$ (Asplund et al. 2005; Caffau et al. 2008), stellar interior models and helioseismology strongly support the Grevesse et al. abundance scale (Basu & Antia 2004; Delahaye & Pinsonneault 2006). For values of solar abundance taken from Asplund et al. (2005) and Caffau et al. (2008), scale our $[O/H]$ results by +0.21 and +0.11 respectively.

2. THE SAMPLE OF SPECTRA

The O VI doublet starts to become visible at the blue end of SDSS quasar spectra at redshifts of $z_{\text{abs}} \geq 2.7$. Hence, we retrieved QSO spectra from the SDSS Third Data Release (DR3; Abazajian et al. 2005) beyond this lower boundary redshift using the QSO Absorption Line Systems (QSOALS) project database (York et al. 2006).

The spectra we used were fitted with an automatic continuum procedure developed by Arlin Crotts and used in the QSO absorption line data base of the SDSS QSOs (York et al. 2006). The continuum is derived by discarding absorption-contaminated pixels until the variance in the remaining pixels corresponds to that prescribed by the noise model for the measurement. This provides a smoothed estimate of the spectrum, forming a ‘continuum’ for the measurement of narrow lines over a range of pixels that varies with the large scale variations in the SDSS spectrum. At a minimum (e.g. near narrow emission lines from the QSO or the sky, or at the steep drop at the edge of a broad absorption line) the smoothing is over a 20 pixel region (between 18Å to 24Å in the forest used here). In smooth regions of the spectrum, outside the Lyman- α forest, the method produces continua accurate to about 2% when the photon noise is small enough to permit such precision. The method produces a reasonable estimate for the continuum in the Lyman- α forest for our purposes. Note that the rare cases of a Lyman- α BAL or the more common cases of an O VI BAL will be fitted away and this is a desirable feature for our purposes.

There are rarely rapid continuum changes to be expected over 30 Å, and certainly no significant non-linear trends. This gives us a good enough continuum for application of a locally calibrated continuum for the O VI absorption measurement, as described in §3. The continuum fitting method may subtract large-scale structures in Ly α absorption on scales $> 2800 \text{ km s}^{-1}$, but this effect will be minimal as indicated by the good agreement between the probability distribution function (PDF) of the Ly α optical depth in the data and the models as shown in Figure 4 below. This continuum fitting method results in a $\sim 16\%$ systematic underestimate of the continuum (and so an overestimate of the transmitted flux), measured as an equivalent underestimate of the mean flux decrement in the sample with respect to values in the literature (e.g. McDonald et al. 2006; Kim et al. 2007). We have taken this into account by introducing the same systematic to our simulations as described in §4.1.

We do not exclude BAL quasars from the sample used in this study. We have tested the impact of these BAL QSOs by performing the LCP search excluding the 430 QSOs that are identified by Gibson et al. (2009) as BAL QSOs and fall into our sample. The impact of this modification is much smaller than the 1σ error bars.

As can be seen in Figure 3 of Frank et al. (2009a), the majority of the sources have i band magnitudes between 19.5 and 20.5. The average signal-to-noise ratio of the spectra within the region of interest for us is only about 2.5, much less than the overall signal-to-noise ratio of the spectrum, since the flux levels within the Lyman- α forest are much lower than those beyond it. This figure also shows the distribution of these average signal-to-noise ratio values with the emission redshifts of the




FIG. 1.— Example of the noise characteristics of the sample spectra at 5000 Å. Plotted in black are the estimated noise levels in real flux, $\sigma(f_\lambda)$, versus the flux at 5000 Å. The noise can be separated into two components, the detector dependent read-out noise, σ_{ro} , and the flux dependent photon noise, σ_{ph} . The dotted red line shows the two parameter fit to the data, as described in the text. The dashed black lines indicate the location of the noise if it were purely dominated by the Poisson statistics of the photon noise with 5, 10 and 20 counts. Where the photon noise dominates, the counts are clearly above 20, and so this Poisson distribution is well described by a Gaussian approximation. Hence the noise distribution is Gaussian at all levels of flux. We have confirmed this for the complete wavelength range covered by SDSS.

sample sources. Particularly towards higher redshifts, it is apparent that the increasing density of the forest and the general faintness of the sources lead to a severe decrease in the signal-to-noise ratio.

In order to ensure that the noise characteristics of the sample can be adequately reproduced in the simulated spectra, we have performed an analysis of the behaviour of the noise in the full data set at every observed wavelength. The uncertainty in the flux measurement, $\sigma(f_\lambda, \lambda)$, is the sum in quadrature of two terms, the read-out noise of the detector system and the photon noise. The former is a Gaussian distribution with $\sigma_{ro}(\lambda) = k_1$, where k_1 is a constant at a particular wavelength, independent of the incident flux. The latter is a Poisson distribution with $\sigma_{ph}(f_\lambda, \lambda) = k_2 \times \sqrt{f_\lambda}$, where k_2 is a constant at a particular wavelength.

At each wavelength, we fit the distribution of the error estimate in the flux provided by the QSOALS database with this two parameter model, and hence are able to describe the complete noise characteristics of the underlying SDSS sample by the wavelength dependent $k_1(\lambda)$ and $k_2(\lambda)$. Figure 1 shows an example of the noise distribution as a function of the estimated flux, f_λ , at a specific wavelength. We have used this formalization of the noise characteristics to produce a realistic representation of observing conditions in synthetic spectra as described in §4.1. This figure also illustrates how the noise may always be characterized as a Gaussian distribution, since the (Gaussian) read out noise dominates where the Poisson nature of the photon noise is significant.

3. THE LOCALLY CALIBRATED PIXEL SEARCH

Pixel correlation searches for the detection and investigation of metal enrichment in the IGM have, until now, been limited to high resolution quasar absorption spec-

tra. In these spectra, the optical depth of every pixel of the Lyman- α forest can be characterized as a measure of the density in the IGM. Where the Ly α absorption is saturated (and so the measured Ly α optical depth measure is noisy), higher order Lyman lines can be used to reconstruct the Ly α optical depth.

Ionization species other than neutral hydrogen can produce their own forest of absorption. In particular there is a forest of five-times-ionized oxygen (O VI) and a forest of three-times-ionized carbon (C IV). These metal line forests can also be dealt with on a pixel-by-pixel basis at each redshift. Hence, one can collect Ly α and O VI or C IV pixel pairs, bin by Ly α optical depth, and one would see a clear trend in metal line optical depth (using only basic assumptions of ionization and metallicity, which we will return to throughout this paper).

Difficulties arise when the measured optical depth of the metal pixels is not entirely due to the intended species. Various techniques have been used in the SP search to minimize the degree of contaminating absorption, but it cannot be removed entirely, and it tends to dominate where absorption is weakest. These techniques involve correction of contaminating Lyman series lines in the O VI region of the spectra by removing expected higher order lines (based on the Ly α absorption at the same redshift and longer wavelength), or the use of both O VI lines to derive a minimally contaminated signal.

As is done in all pixel correlation searches in quasar spectra, we compared the absorption from a species at a fixed redshift with the absorption from another species at the same redshift. Where the absorption is correlated, a detection is found. In this analysis we search for O VIa ($\lambda 1032\text{\AA}$) absorption correlated with Lyman- α ($\lambda 1216\text{\AA}$) absorption and make no use of the O VI doublet line (O VIb) at $\lambda 1038\text{\AA}$. We produce these pixel pairs for a range of redshifts and on a grid set by the spectral binning of the O VIa pixels (the Ly α absorption is interpolated). Nearby pixels with a wavelength close to the O VIa wavelength are used to characterize continuum errors and contaminating absorption. This procedure is explained in more detail below.

We have simplified the algorithm in recognition of the fact that there are features of the SP search that are unsuitable in the context of our local calibration method. These techniques are directed toward subtracting contamination. We do not seek to remove contaminating absorption; as described above, we *characterize* it, and, as shall be seen below, the LCP analysis does this automatically.

The lower redshift limit for every quasar in the sample is set by the blue end of the spectrograph (around 3800\AA) or the Lyman- γ line at the quasar redshift, whichever results in a higher minimum redshift. We discard pixels with Lyman- γ as they provide little extra information but modify the distribution of absorbers in O VI signal enough to introduce unnecessary uncertainty. Note that this requirement also means that no absorption at the Lyman limit is possible. We rule out regions within 5000 km s^{-1} of the QSO emission redshift to eliminate most effects due to the QSO and its environment (Wild et al. 2008). This sets the upper redshift limit and also results in an O VI signal entirely within the Lyman- β forest (since the separation between O VIa and Ly β




FIG. 2.— Redshift distribution of pixel pairs available for the analysis in the SDSS sample. The lower envelope is set by the lowest redshift systems observable in O VI at the blue end of the spectrograph. The upper envelope is set by 5000 km s⁻¹ below the emission redshift of QSOs in the sample

absorption at the emission redshift is 1,804 km s⁻¹ in the O VI frame).

We combine the pixel pairs for each QSO into one combined sample for all the SDSS spectra in the required redshift range producing over a million pixel pairs. The redshift distribution of these data is shown in Figure 2.

3.1. H I Absorption

The determination of the degree of H I absorption is largely unchanged from the SP search. The optical depth to Lyman- α absorption $\tau_{\text{Ly}\alpha}$ is determined for each pixel. We make the same noise requirement as other searches: the flux transmission must be $\sigma_{\text{noise}}/2$ from either the zero level or the continuum, where σ_{noise} is the continuum scaled error estimation in the continuum fitted spectrum. This is a rather strict criterion in the context of the SDSS sample and discards $\approx 20\%$ of pixel pairs. We do not attempt to extend the analysis to the higher density regime, where Ly α absorption is saturated, by using higher order Lyman lines to derive the Ly α optical depth. This is a regime to which other methods are better suited (Frank et al. 2009a,b). It should also be noted we assume that all O VI is associated with H I absorption, but under most physical conditions this is a fair assumption.

The large amount of data at our disposal allows us to aggressively discard less reliable measures of $\tau_{\text{Ly}\alpha}$. Hence we require that $\sigma_{\text{noise}} < 0.1$, which discards a further 70% of the pixel pairs, so the feature we are searching *by* (Ly α) is as clean a signal as we can make it while preserving a large sample of data. As we shall show, the feature we are searching *for* (O VIa) is treated with no such noise requirements, and brute force averaging is used to provide a clean signal.

3.2. Measuring the Bulk O VI Absorption

Any measure of weak correlated O VI absorption must deal with the biggest uncertainty: the continuum level. The primary goal for the following technique is to remove the impact of errors in the continuum fitting from the derived signal. Rather than simply taking the transmitted flux at the position in the spectra corresponding to the required O VI absorption we determine the ratio of the transmitted flux at this point compared to that of a neighboring pixel. Where this pixel is sufficiently

close, this gives us a measure of the absorption with the local error in the continuum fitting subtracted. It also results in a subtraction of average levels of contaminating Ly α , Ly β and O VIb absorption along with a measure of nearby O VIa absorption. The ratio of transmitted flux, F , at the O VIa pixel to that on a locally calibrating pixel is

$$\frac{F_{\text{OVIa}}}{F_l} = \frac{f_c \exp[-\tau_{\text{OVIa}}] \exp[-\tau_u] + \text{noise}}{f_{c,l} \exp[-\tau_{\text{OVIa},l}] \exp[-\tau_{u,l}] + \text{noise}_l}, \quad (1)$$

where f_c is the flux level at the continuum, τ_{OVIa} is the optical depth to O VIa absorption, ‘noise’ is the Gaussian instrumental noise, the subscript ‘ l ’ indicates a local value and the subscript ‘ u ’ refers to uncorrelated contaminating absorption. This uncorrelated absorption is dominated by Ly α and Ly β absorption, but it also includes contamination by the second member of the O VI doublet ($\tau_u = \tau_{\text{Ly}\alpha} + \tau_{\text{Ly}\beta} + \tau_{\text{OVIb}}$), all at different redshifts. $\tau_{\text{OVIa},l}$ is the local O VIa absorption, which is also uncorrelated where the appropriate local pixels are used, as we shall argue in §3.4.

Multiple local pixels are used, and each one provides a new measure of spectral characteristics local to our required O VIa pixel. Operationally, we treat each local pixel as providing a new measure of F_{OVIa}/F_l and so a new pixel pair of $\tau_{\text{Ly}\alpha} - F_{\text{OVIa}}/F_l$. In this way, the final statistical error is dominated by the number of Ly α - O VI pixels and not by the noise in the local pixels.

3.3. The Aggregation of Pixel Pairs

Since we have a large sample of spectra with comparable characteristics and a limited redshift range we aggregate all acceptable pixel pairs from all our spectra. We use pixels in a redshift range of $2.7 < z < 3.2$ both to keep our data as homogeneous as possible and to limit ourselves to the subset that provides the strongest signal. This limits us to 2167 QSO spectra in our sample, out of the 5767 $z > 2.3$ quasar spectra in SDSS DR3.

Once we have a list of $\tau_{\text{Ly}\alpha} - F_{\text{OVIa}}/F_l$ pairs, as set out above, we bin them by their $\tau_{\text{Ly}\alpha}$ and take the median of the $\tau_{\text{Ly}\alpha}$ and the F_{OVIa}/F_l in order to produce our final search for O VIa. We use the median as an outlier- and noise-resistant measure of the typical absorption.

Since $\tau_{u,l}$ and τ_u are drawn randomly from the same distribution, and (where the local pixel is sufficiently close) $f_{c,l} \approx f_c$, the median(F_{OVIa}/F_l) is a quantity that characterizes correlated O VIa absorption normalized to a factor dependent on the mean O VIa optical depth. The precise level of this normalization is unimportant in the following results; as with other pixel correlation searches, a statistically significant correlation is required for a detection of metals. It is, however, broadly indicative of the average O VIa absorption level. It is also notable that this level is recovered in the null tests that follow.

3.4. Choice of Local Pixels

For local pixels we chose those that are i) sufficiently far away from the O VIa search pixel to avoid comparison with pixels within the same extended complex associated with the same enriched region, ii) sufficiently close to avoid producing spurious correlations with the Lyman- β and O VIb absorption, and iii) sufficiently close to measure the same continuum level. The second requirement

ensures that the third requirement is satisfied. Lyman- β is stronger and correlated over a larger velocity range compared with O VIb, hence we design our choice of local pixels for avoidance of this line. We also require that the choice of local pixels are mirrored on each side of the search pixel in order to cancel continuum level error trends across the line.

We take 18 local pixels in total for this analysis: the 9th to the 17th pixel from the search pixel at both lower and higher wavelengths. This results in separations of $\gtrsim 600 \text{ km s}^{-1}$ and $\lesssim 1200 \text{ km s}^{-1}$ from the O VIa and Ly β transitions, which is approximately $9 - 16 \text{ \AA}$ in the observed frame at $z = 3$. This provides $\approx 800,000$ pixel pairs in all the searches that follow.

We tested our choices of locally calibrating pixels by treating each of them individually as a ‘search target’ (searching for correlations with Ly α) and finding the expected null result. Pixels within 600 km s^{-1} of O VIa or Ly β , on the other hand, yielded signals of contaminating correlated absorption. The 600 km s^{-1} scale is characteristic of the winds from Lyman break galaxies (e.g. Pettini et al. 2001; Adelberger et al. 2003) and the maximum clustering scale found for C IV (Adelberger et al. 2003, 2005; Pieri et al. 2006; Scannapieco et al. 2006), so these results seem reasonable for O VI. Correlations in Ly β extend over greater scales than this due to large scale structure (McDonald et al. 2006), but such lines are somewhat stronger than expected here and our test searches confirm this. In principle C II($\lambda 1036 \text{ \AA}$) could produce correlated absorption within our calibrating pixel range, but in practice it appears to be too weak.

In test searches we detected the doublet line O VIb at a consistent level with that of the O VIa line. Since this line merely constitutes a parallel but less sensitive measure of O VI that does not contribute to the search algorithm, we have not used this information.

3.5. O VI Search Results

Once the pixel pairs required are combined and binned by their $\tau_{\text{Ly}\alpha}$, we find their medians to plot F_{OVI}/F_l vs. $\tau_{\text{Ly}\alpha}$ as shown in Figure 3. We required that more than 200 pixel pairs be available for a bin to be used. In all panels the error bars are produced by bootstrap resampling with 100 realizations and a bootstrap element of 1 SDSS QSO. The median is found for every realization, and the error bars are then taken to be the 1σ deviation in the distribution of medians.

Figure 3 shows F_{OVI}/F_l vs. $\tau_{\text{Ly}\alpha}$ using pixels in the redshift range $2.7 < z < 3.2$. There is a clear correlation between $\tau_{\text{Ly}\alpha}$ and F_{OVI}/F_l , and this correlation extends to all measured values of $\tau_{\text{Ly}\alpha}$. There is both excess O VI absorption in pixels with strong Ly α absorption and a deficit of O VI absorption ($F_{\text{OVI}}/F_l > 1$) in pixels where Ly α absorption is weak.

The null test shown in Figure 3 is performed using the same technique as set out above with the only modification that the restframe wavelength for correlations is no longer the O VIa restframe wavelength ($\lambda 1032 \text{ \AA}$) but $\lambda = 1047 \text{ \AA}$. As can be seen, the search approach passes the null test by recovering the expected result of $F_{1047} \approx F_l$, where neither an excess nor a deficit in absorption is seen as no signal is present.

Figure 3 is the main observational result of this pa-




FIG. 3.— The LCP search for O VIa in the SDSS Ly α forest from pixels with $2.7 < z < 3.2$. The red filled circles are the median F/F_l on and off the 1032 \AA O VIa resonance line vs the $\tau_{\text{Ly}\alpha}$ at the same redshift. Where F_{OVI}/F_l is below (above) 1 there is a greater (lower) than average degree of O VI absorption. A correlation in this plot is indicative of a detection of O VI. The black ‘+’s are the null case of searching on and off 1047 \AA .

per. In high resolution spectra there is an approximate one-to-one relation between $\tau_{\text{Ly}\alpha}$ and gas overdensity $\rho/\bar{\rho}$ (Rauch et al. 1997; Croft et al. 1998; Weinberg et al. 1998). At SDSS resolution this is no longer the case, so the interpretation of this result in terms of enrichment of the IGM must rely on synthetic spectra drawn from an underlying physical model.

4. INTERPRETATION USING SYNTHETIC SPECTRA

4.1. Production of Synthetic Spectra

We use an approach for the production of simulated line-of-sight density distributions that was set out in Pieri & Haehnelt (2004). We restate this method here, but for a full description refer to that paper. We use cosmological parameters consistent with WMAP5 (Komatsu et al. 2008) ($\Omega_m = 0.268$, $\Omega_\Lambda = 0.732$, $\Omega_b = 0.0441$, $\sigma_8 = 0.776$ and $H_0 = 70.4 \text{ km s}^{-1} \text{ Mpc}^{-1}$) in these simulations

We start with a 3D power spectrum of dark matter density fluctuations taken from Efstathiou et al. (1992). This is normalized using the cosmological parameter, σ_8 (the amplitude of clustering on an $8h^{-1} \text{ Mpc}$ scale at the present epoch), and the linear growth of structure at the required redshift. A Jeans length filtering ($\sim 0.9h^{-1} \text{ Mpc}$ comoving in the surveyed redshift range) to remove structure on small scales is used to describe baryonic pressure effects.

Two 1D power spectra are produced by integration of the 3D power spectrum (Kaiser & Peacock 1991), and

Gaussian random realizations of structure are generated. These 1D realizations of structure are combined to produce coupled density contrast and velocity fields. This velocity field, $v(x)$, is a non-dynamical approximation based on the linear density field, $\delta(x)$. For more details on this approach see Bi & Davidsen (1997) and references therein.

Further following the method of Bi & Davidsen (1997), we convert our linear density field to a lognormal distribution, using the mapping

$$\rho/\bar{\rho}(x) = \exp[\delta(x) - \langle \delta^2 \rangle / 2], \quad (2)$$

which they find reproduces the probability distribution function of non-linear structures. It should be noted that this method fails to reproduce the clustering of non-linear structures (Viel et al. 2002). The density field produced has a resolution of 0.01Mpc comoving, which substantially over-samples the resolution required by the SDSS sample but is adopted in order to ensure that the physics of the IGM is well described.

The temperature of the IGM in these simulations is set by the balance between photoionization heating and adiabatic cooling and follows the power-law relation

$$T = T_0(\rho/\bar{\rho})^{0.4} \quad (3)$$

(Hui & Gnedin 1997), where $\rho/\bar{\rho}$ is the overdensity and the temperature at mean density is $T_0 = 2 \times 10^4$ K. The temperature is required to calculate the Doppler parameter for this gas.

For a given set of cosmological parameters and a helium fraction of 0.242 by mass, it is trivial to calculate the hydrogen density. Oxygen is added in quantities as described in the following section. The ionization fractions of hydrogen and oxygen are calculated using CLOUDY version 08.00 (Ferland et al 1998) and the Haardt & Madau (2001) (hereafter HM01) quasar and galaxy UV background model (designated QG) with a 10% escape fraction from galaxies. While these calculations do take into account collisional ionization, the lack of heating other than via photoionization in our simulations makes the conditions for it to occur rare. As such, we predominately interpret the detectable O VI to be outside of this collisionally ionized, hot gas phase. This assumption is supported by observations that the bulk of previously detected O VI lines in the Ly α forest are photoionized (Carswell et al. 2002; Bergeron et al. 2002; Simcoe et al. 2004) based on their line widths. This approach provides fully self-consistent models since regions with $\rho/\bar{\rho} < 10$ will be overwhelmingly photoionized and we find this regime dominates our measured O VI (see §4.3).

We calculate the optical depths for Ly α , Ly β , O VI α and O VI β using

$$\tau(v) = \frac{\pi e^2 \lambda f}{m_e c H_z} \int \frac{n(v')}{b(v')} \exp\left[-\left(\frac{v' - v_0}{b(v')}\right)^2\right] dv', \quad (4)$$

where the transition dependent parameters are the rest frame wavelength λ , the oscillator strength f , the number density of absorbing particles n , and the Doppler parameter b . H_z is the redshift dependent Hubble parameter, v is the line-of-sight velocity, v_0 is the velocity of the Hubble flow, e is the electron charge, m_e is the electron mass, and c is the speed of light. This is integrated over the narrow range of velocities that contribute

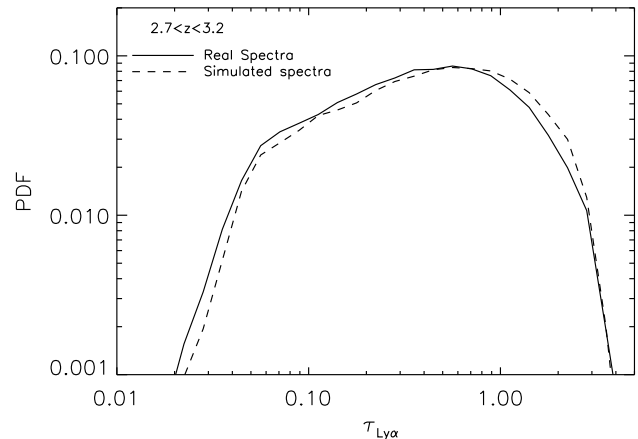


FIG. 4.— The PDF of Ly α optical depths in the observed sample (solid line) and one simulated suite (dashed line). The simulated sample was produced with a quasar and galaxy HM01 UV background, the intensity of which was renormalized to match the observed distribution giving an ionization rate of $\Gamma = 13 \times 10^{-13} \text{s}^{-1}$.

significantly to the calculation (effectively a few times the Doppler parameter).

We add instrumental broadening to match the SDSS spectral resolution. As stated in §2, we have performed an analysis of the characteristic $\sigma(f_\lambda, \lambda)$ in our sample, and we use this to add Gaussian noise to the spectra. In order to achieve this we use the continuum level in the real spectra to rescale the simulated transmission flux to a flux, f_λ . The photon noise is well described by a Gaussian distribution where it dominates, and the readout noise is also a Gaussian, so treating the combined noise $\sigma(f_\lambda, \lambda)$ as a Gaussian is a good approximation (§2). Finally, we rescale our transmitted fluxes up by 16% in line with the systematic underestimate of the continuum in our sample (§2).

The intensity of the UV background was renormalized (see §6.1) to provide a match of the PDF of optical depths derived from our synthetic spectra with that of the observed sample. We obtain agreement with an ionization rate of $\Gamma = 13 \times 10^{-13} \text{s}^{-1}$ as is shown in Figure 4. The quality of this agreement is sufficiently good to populate the bins in $\tau_{\text{Ly}\alpha}$ with statistics that are representative of the LCP search in the data.

4.2. The Monte Carlo Approach

For every oxygen abundance pattern we simulate a full suite of 2167 SDSS QSO spectra and do this ten times. For each one we perform an LCP search for O VI, and the resultant F_{OVI}/F_l vs. $\tau_{\text{Ly}\alpha}$ is averaged over the ten suites to obtain a model prediction (so that the quality of the fit between the model and observed LCP searches is dominated by the error in the observed search). We have used this set of 21,670 line-of-sight density distributions with varying levels of oxygen abundance to interpret our observations. Each of these lines-of-sight require the simulation of a 1D density distribution of length around $1000h^{-1} \text{Mpc}$. The production of this quantity of data is not practical with fully numerical simulation methods given current computer processor limitations. This underlines the value of the analytical technique described in the previous section. 212 models of oxygen enrichment have been tested, totaling more than 4.5 million synthetic spectra produced.

In the following analysis we do not include the mea-




FIG. 5.— The LCP search compared with models of density independent $[\text{O}/\text{H}]$. The red filled circles are the LCP search as in Figure 3. Four models are shown: the *solid* line is the null case where no oxygen is added, the *dashed* line shows the best fitting case of a constant $[\text{O}/\text{H}] = -2.15$. A further two models show oxygen metallicities that are ruled out: $[\text{O}/\text{H}] = -2.00$ (*dot-dashed* line) and $[\text{O}/\text{H}] = -2.30$ (*dotted* line). Note that the $\tau_{\text{Ly}\alpha} \approx 3.3$ point is not included in the fits. We draw detailed conclusions on metallicity constraints in §5.

sure at $\tau_{\text{Ly}\alpha} \approx 3.3$ (see Figure 3). The inclusion of this point results in poor agreement with the models. This is consistent with the findings of A08, that high $\tau_{\text{Ly}\alpha}$ absorption results from a different population of O VI absorbers that are most likely collisionally ionized.

4.3. Comparison of Real and Synthetic Spectra

We add oxygen with a constant $[\text{O}/\text{H}]$ to our synthetic spectra and have varied this $[\text{O}/\text{H}]$ to find a good fit to the observed search. Figure 5 shows a comparison of the observed search (filled circles) with four models. Once more we see that our LCP search responds well to the null test (solid line) by returning a result of $F_{\text{OVI}} \approx F_I$ when no oxygen is added (the *reduced* $\chi^2 = 10$). In the best fitting case, $[\text{O}/\text{H}] = -2.15$ (dashed line), we find good agreement with the observed result, and there is no need for the addition of more parameters in the fit. Models with constant $[\text{O}/\text{H}] = -2.0$ (dotted line) or $[\text{O}/\text{H}] = -2.3$ (dot-dashed line) are clearly ruled out and indicate that the LCP search in the SDSS sample is sensitive to small changes in the abundance of oxygen in the IGM.

We can quantify the quality of the fit by using the χ^2 test where the error bars are given by the bootstrapping of the observed data. We have done this for varying $[\text{O}/\text{H}]$ as shown in Figure 6 (solid line). The 1.5σ and 2σ level are shown as the dotted lines. It is clear that our best fitting model with $[\text{O}/\text{H}] = -2.15$ constitutes an adequate fit to the observed search. We can give this measure a 2σ confidence interval of $[\text{O}/\text{H}] = -2.15_{-0.09}^{+0.07}$. These models constitute the range of constant $[\text{O}/\text{H}]$ models favored by the LCP search. The quality of this fit is suggestive of an underlying density dependence to the $[\text{O}/\text{H}]$ signal (as found by A08), but the statistics do not warrant a two parameter fit.

We may not be able to constrain the density dependence of this oxygen enrichment, but we can place limits on the density range in which it resides. Hence we introduce a new parameter: a gas overdensity threshold, $(\rho/\bar{\rho})_{\text{cut}}$. Where $\rho/\bar{\rho} < (\rho/\bar{\rho})_{\text{cut}}$, we add no oxygen, and




FIG. 6.— The χ^2 test of fit (*solid* line) between LCP searches performed on our models with varying density independent $[\text{O}/\text{H}]$ and our observed LCP search. The *dotted* lines show the 1σ , 1.5σ and 2σ confidence levels indicating the range of $[\text{O}/\text{H}]$ that constitutes a good fit.

where $\rho/\bar{\rho} > (\rho/\bar{\rho})_{\text{cut}}$ we add the constant $[\text{O}/\text{H}]$ specified. Figure 7 shows a likelihood plot of the fit to the SDSS data set for 208 models (with two parameters and nine data points). Where $(\rho/\bar{\rho})_{\text{cut}} = 0$, no threshold is used and the models correspond to those used in Figure 6. There is little or no change in the quality of the fit going from $(\rho/\bar{\rho})_{\text{cut}} = 0$ to $(\rho/\bar{\rho})_{\text{cut}} = 1$, and this indicates that our LCP search for O VI in the SDSS QSO sample is not sensitive to metals in underdense regions of the Universe.

As the threshold is raised in our models a degeneracy becomes apparent, as seen in Figure 7: removing metals from low-density regions requires one to add more metals to higher density regions to compensate, and it is *still possible to obtain a good fit*. This indicates that our search is rather insensitive to the density of the systems enriched. As expected, the tight relationship between $\tau_{\text{Ly}\alpha}$ and overdensity seen in high resolution spectra (Rauch et al. 1997; Croft et al. 1998; Weinberg et al. 1998) breaks down due to the moderate resolution of the SDSS sample. Dense systems can always produce a low apparent $\tau_{\text{Ly}\alpha}$ by virtue of the instrumental broadening over multiple pixels, but the key question is *to what extent those systems can reproduce the observed signal?* The agreement begins to break down at $(\rho/\bar{\rho})_{\text{cut}} \approx 6$, where we have 13% confidence (1.5σ) in our fit to the observed result, and it breaks down completely at $(\rho/\bar{\rho})_{\text{cut}} \approx 8.5$, where we have a 5% confidence (2σ) in our fit. This breakdown occurs because there are proportionately too few higher density systems to obtain a fit to the search curve even with arbitrarily high metallicity. The best fit is obtained with $(\rho/\bar{\rho})_{\text{cut}} = 4$, where the fit has $\chi^2 = 9.4$ (23% confidence for two parameters and nine data points).

While Figure 7 shows that a range of two-parameter models can fit the data, we will show in §5 that the mass-weighted mean oxygen abundance is well constrained. To demonstrate this, we must first address the issue of scatter in the metallicity.

4.4. Scatter in the Metallicity

Schaye et al. (2003) (hereafter S03) and A08 found that a lognormal scatter in the metallicity is a good fit




FIG. 7.— The 1.5σ (dark red) and 2σ (orange) likelihood of fit between LCP searches performed on our models and our observed LCP search (with two parameters and nine data points). The models have a constant $[\text{O}/\text{H}]$ above a threshold gas overdensity, $(\rho/\bar{\rho})_{\text{cut}}$. Below this threshold no oxygen is added. The contours characterize the range of densities for which this search for O VI is sensitive. The likelihood is largely unaffected by the addition of a threshold at $\rho/\bar{\rho}_{\text{cut}} = 1$, indicating that the search is not sensitive to metals in underdense systems. Where a good fit can no longer be obtained for a particular value of $(\rho/\bar{\rho})_{\text{cut}}$, this constitutes the upper limit on the density range to which we are sensitive.

to all percentiles in their search. We assume the same scatter in order to have a better understanding of the nature of our search results. In particular, it is not clear if an $[\text{O}/\text{H}]$ model without scatter that fits the data gives a measure of the mean or median metallicity in the IGM.

A model with lognormal metallicity scatter, i.e Gaussian scatter with variance σ_s^2 in $\log Z = [\text{O}/\text{H}]$, has

$$\begin{aligned} \log Z_{\text{med}} &= \text{median}(\log Z) = \langle \log Z \rangle \\ &= \log \langle Z \rangle - \frac{\ln 10}{2} \sigma_s^2, \end{aligned} \quad (5)$$

where $\log Z_{\text{med}}$ and $\langle Z \rangle$ are the median and mean metallicity, respectively, and we adopt $\sigma_s^2 = 0.5$ (Schaye et al. 2003, eq. 7). The lognormal scatter is added on scales of $1.2h^{-1}\text{Mpc}$ comoving in regions where $-0.5 < \log(\rho/\bar{\rho}) < 2$. As equation 5 shows, lognormal models with different scatter, σ_s , have the same $\langle \log Z \rangle$ and the same Z_{med} but *different* $\langle Z \rangle$.

In order to test whether the LCP search for O VI is sensitive to the mean or median oxygen abundance we explore both interpretations. Figure 8 compares our best fitting constant $[\text{O}/\text{H}]$ model with no scatter (as shown in Figure 5) to two models with scatter, one that has the same $\langle Z \rangle$ and one that has the same Z_{med} . The good agreement in the former case indicates that our search for O VI is sensitive to the $\langle Z \rangle$ or the mean O/H. Furthermore, we have tested this hypothesis using the A08 density-dependent fit to O VI seen in high resolution spectra. They report a model fit

$$\begin{aligned} [\text{O}/\text{H}] &= [\text{O}/\text{C}] + [\text{C}/\text{H}] \\ &= -2.81_{-0.14}^{+0.15} + 0.08_{-0.06}^{+0.09}(z-3) \\ &\quad + 0.65_{-0.14}^{+0.1}(\log[\rho/\bar{\rho}] - 0.5) \end{aligned} \quad (6)$$

that is derived from O VI correlated with C IV and C IV correlated with H I. This model is a three parameter fit to the median $[\text{O}/\text{H}]$. Since it is fit to independent data we treat it as having nine degrees of freedom here (equal




FIG. 8.— The impact of scatter in our models and its consequences for the nature of the LCP search. The red filled circles are the LCP search as in Figure 3. The dashed line is the best fitting model with constant $[\text{O}/\text{H}]$ and no scatter as in Figure 5. We have added lognormal scatter to this model while keeping either the mean O/H fixed (*dotted* line) or the median O/H fixed (*dot-dashed* line). The agreement of the former but not the latter indicates that the scatter-free model fit yields a robust estimate of the mean O/H, not the median O/H (or the mean $[\text{O}/\text{H}]$). Note that the $\tau_{\text{Ly}\alpha} \approx 3.3$ point is not included in the fits.

to the number of data points). When we reproduce this model in our simulations including scatter and perform the LCP search, it provides a reasonable fit to the SDSS data (about 2σ). However, where we have simply added metals to our simulation at this median level without scatter our agreement is poor (*reduced* $\chi^2 = 6$). When we use the correction factor in equation 5 and add metals at the mean level instead, we obtain a reasonable fit once more (at about 1.5σ confidence).

5. MEASURING THE MEAN OXYGEN ABUNDANCE AND THE COSMIC OXYGEN DENSITY

Sections 4.3 and 4.4 present a variety of models that are consistent with our LCP search results, despite differences in density dependence and scatter. We can ask whether these models have a common feature of oxygen abundance that is well determined by our data. To this end we have calculated the volume-weighted and mass-weighted means of O/H for all models that fit the data: those within 1.5σ and 2σ contours, along with the A08 model. Specifically, we compute the means by summing O/H (not $[\text{O}/\text{H}]$) over all pixels in the simulated spectra with mass density in the range $-0.5 < \log[\rho/\bar{\rho}] < 2$, either weighting each pixel equally (volume-weighted) or weighting in proportion to mass density.

The mass-weighted mean oxygen abundance, $[(\text{O}/\text{H})_{\text{MW}}]$, for the ‘ridge line’ of best fitting models at every $(\rho/\bar{\rho})_{\text{cut}}$ is shown in Table 1. We list the models without scatter, but models with scatter yield the same results since (as shown in Figure 8) they must have the same mean oxygen abundance to fit the data. Taking models within the 1.5 and 2σ contours (rather than just this ridge line) provides a measure of both uncertainty in density dependence and observational uncertainty.

While the volume-weighted abundances span 0.4-1.3 dex in $[\text{O}/\text{H}]$, the successful models all have mass-weighted means in the range $-2.51 < [(\text{O}/\text{H})_{\text{MW}}] < -2.39$ (1.5σ) and $-2.56 < [(\text{O}/\text{H})_{\text{MW}}] < -2.10$ (2σ).

We conclude that the LCP measurements constrain the mass-weighted mean oxygen abundance of the IGM to be $[\langle\text{O}/\text{H}\rangle_{\text{MW}}] = -2.45 \pm 0.06$ (1.5σ) and $[\langle\text{O}/\text{H}\rangle_{\text{MW}}] = -2.33 \pm 0.23$ (2σ). We can compare this result directly with the mass-weighted mean in A08 of $[\langle\text{O}/\text{H}\rangle_{\text{MW}}] = -2.14 \pm 0.14$ (1σ), which is marginally consistent with our results. The volume-weighted mean oxygen abundance is more weakly constrained to $[\langle\text{O}/\text{H}\rangle_{\text{VW}}] = -3.06 \pm 0.22$ (1.5σ) and $[\langle\text{O}/\text{H}\rangle_{\text{VW}}] = -2.74 \pm 0.64$ (2σ). These error estimates do not include systematic uncertainties associated with the shape and intensity of the ionizing background (see §6.1) or uncertainties in the solar oxygen abundance. Using Caffau et al. (2008) values for the solar abundance of oxygen would give a mass-weighted $[\langle\text{O}/\text{H}\rangle_{\text{MW}}] = -2.34 \pm 0.06$ (1.5σ) and $[\langle\text{O}/\text{H}\rangle_{\text{MW}}] = -2.22 \pm 0.23$ (2σ).

With this measure of $[\langle\text{O}/\text{H}\rangle_{\text{MW}}]$, and extrapolating to all densities, we can calculate the cosmic oxygen density using

$$\Omega_{\text{Oxy,IGM}} \approx Y_{\text{H}} A_{\text{O}} \Omega_b 10^{[\langle\text{O}/\text{H}\rangle_{\text{MW}}]} (\text{O}/\text{H})_{\odot}, \quad (7)$$

where $Y_{\text{H}} = 0.76$ is the mass fraction of hydrogen and $A_{\text{O}} = 16$ is the atomic mass of oxygen. We obtain a value of $\Omega_{\text{Oxy,IGM}} = 1.4(\pm 0.2) \times 10^{-6}$ (1.5σ). This would be $\sim 5\%$ smaller if we limited ourselves to the density range $-0.5 < \log[\rho/\bar{\rho}] < 2$ (the range for which the mass-weighted mean was calculated). This shows that the calculation is not sensitive to the details of the density range integrated over, but this is not a statement about the metal enrichment of systems beyond our constraining density range. In particular, the population of strong absorbers, such as those resulting in the $\tau_{\text{Ly}\alpha} \approx 3.3$ point of our LCP, are better dealt with in directed searches (Frank et al. 2009a,b).

If we follow the procedure of Bouché et al. (2007) (and A08) and take $\Omega_{\text{Z,IGM}} \approx \Omega_{\text{Oxy,IGM}}/0.6$, then we obtain a total cosmic metal density of the IGM, $\Omega_{\text{Z,IGM}} \approx 2.3 \times 10^{-6}$. This is $\approx 16\%$ of their metal budget of $\Omega_{\text{Z}} = 1.5 \times 10^{-5}$ at $z = 3$ (from their figure 1 and based on integrated cosmic star formation with standard assumptions about about stellar IMF and yields). This metal content is over half that found in galaxies at $z = 2.5$ (30%) while sub-DLAs may contribute between 2 and 17% (Bouché et al. 2007). Summing these values would account for over half of the metals produced by star formation. The remainder could plausibly reside in intermediate densities (e.g. strong O VI absorbers, Frank et al. 2009b) or in a hidden warm-hot phase, but uncertainties in the calculated metal production and the observational estimates preclude strong statements about “missing” metals.

6. DISCUSSION

6.1. Ionization Corrections

We have obtained the most precise constraints to date on the mass-weighted mean O/H and Ω_{Oxy} , but our analysis rests on the assumption that the detected O VI is mostly ionized by a UV background with the spectral shape of the HM01 QG model. The use of this UV background model is motivated by the work of S03 and A08. They also use two other background models: one resulting from quasars only (“Q”) and a softened version of the QG model (“QGS”) with a lower flux above 4 Ryd.

TABLE 1
MASS-WEIGHTED OXYGEN ABUNDANCE IN SOME BEST FITTING MODELS

Model ($[\text{O}/\text{H}], (\rho/\bar{\rho})_{\text{cut}}$)	$[\langle\text{O}/\text{H}\rangle_{\text{MW}}]$	Goodness-of-fit
-2.15, 0	-2.15	1.8σ
-2.10, 1	-2.25	1.8σ
-2.00, 2	-2.39	1.5σ
-1.80, 3	-2.41	1.3σ
-1.65, 4	-2.45	1.2σ
-1.45, 5	-2.43	1.3σ
-1.80, 6	-2.43	1.4σ
-1.15, 7	-2.43	1.7σ
-0.95, 8.	-2.36	1.8σ
A08	-2.36	2σ

NOTE. — For models along the ridge-line of Figure 7, with enrichment density cutoffs but no scatter in metallicity at fixed density, we list the [O/H] corresponding to the mass-weighted mean oxygen abundance ratio $\langle\text{O}/\text{H}\rangle$ in column 2 and the acceptability of the χ^2 fit to the data in column 3.

They find that those two models are inconsistent with their relative abundance measures of [O/C] and [O/Si], while they obtain a good fit using the QG model. It is not clear to what extent small variations in the QG model can be tolerated, and what impact these would have on the quoted results. This is clearly an area that requires further investigation.

Our simulations assume a power-law equation of state and so under-produce gas at $T > 10^5\text{K}$, which is dominated by collisional ionization. This choice is justified in light of studies by Carswell et al. (2002), Bergeron et al. (2002) and Simcoe et al. (2004) who show that O VI lines are typically too narrow to be in this phase. It is also notable that O VI which arises from collisional ionization also tends to be strong due to the high O VI fraction in this regime, while the O VI absorption probed here is relatively weak (even compared to the studies listed above).

A08 find a substantial increase in [O/H], for a fixed [O/C] ratio, for systems that have $\rho/\bar{\rho} > 30$. Lower densities are well described by photoionization and a fixed [O/C] ratio in their models. They conclude that $\rho/\bar{\rho} > 30$ regions are in fact collisionally ionized (while systems with $\rho/\bar{\rho} < 30$ are photoionized). Simcoe et al. (2004) also find such a change in [O/C] and infer a change in ionization mechanism at higher densities. We find indications of the same effect for our measure at $\tau_{\text{Ly}\alpha} \approx 3.3$, which shows O VI absorption at a level far higher than for lower $\tau_{\text{Ly}\alpha}$. We have omitted this data point for our photoionization dominated analysis. We have only one point showing this effect, since we do not use higher order Lyman lines in our analysis and are therefore unable to cleanly isolate the highest density regions. We find a good fit at all detectable Ly α absorption levels below $\tau_{\text{Ly}\alpha} \approx 3.3$ using a constant [O/H] and assuming a photoionized medium.

One should note that A08 normalise their UV background to a level that differs from ours. This leads to an adjustment on the ionization correction. The ratio of optical depths is related to the metallicity by

$$\log\left[\frac{\tau_{\text{OVI}}}{\tau_{\text{Ly}\alpha}}\right] = \log\left[\frac{n_{\text{OVI}}}{n_{\text{O}}}\frac{n_{\text{H}}}{n_{\text{HI}}}\right] + [\text{O}/\text{H}] + \log(\text{O}/\text{H})_{\odot} + S, \quad (8)$$

where $S = \log[(f\lambda)_{\text{OVI}}/(f\lambda)_{\text{Ly}\alpha}]$. The first term on the




FIG. 9.— The ionization correction parameter $(n_{\text{OVI}}/n_{\text{HI}})/(n_{\text{OI}}/n_{\text{HI}})$ in terms of the UV background intensity (normalized to the HM01 level) and ionization rate at $z = 3$. This is calculated for 3 densities that span our range of sensitivity: mean density (*solid* line), $\log[\rho/\bar{\rho}] = 0.5$ (*dotted* line), and $\log[\rho/\bar{\rho}] = 1$ (*dashed* line). The thin vertical lines show the UVB intensity used in A08 and here.

right hand side of this equation is the ionization correction term, and this is plotted against UV background intensity (normalized to the HM01 level) and the ionization rate in Figure 9. This was calculated assuming the equation of state in equation 3 and $z = 3$ for three overdensities: $\log[\rho/\bar{\rho}] = 0, 0.5$ and 1 that span our sensitivity range. Our chosen UV background intensity of $1.35 \times \text{HM01}$ ($\Gamma = 13 \times 10^{-13} \text{s}^{-1}$) is shown along with the A08 value of $0.48 \times \text{HM01}$ ($\Gamma = 4.6 \times 10^{-13} \text{s}^{-1}$). For mean density the ionization correction is largely unchanged between our simulations and those of A08, however, our models would result in systematically lower $[\text{O}/\text{H}]$ resulting from oxygen in systems of $\log[\rho/\bar{\rho}] > 0.5$ by between 0.3 and 1.1 dex. This might go some way to explaining the 0.2 dex lower value we find for $[\text{O}/\text{H}]$. Another source of potential discrepancy is differences in the density distributions between the studies. Despite these issues, we are in broad agreement with a metal budget fraction of $\sim 15 - 35\%$ in A08 for $z = 2 - 3$.

The success of the LCP search method for detection of O VI in the SDSS spectra may indicate that it is appealing for other searches for low level absorption in the IGM, where large numbers of moderate resolution spectra with moderate signal-to-noise and a substantial uncertainty in the continuum are available. There are a number of other species that may be detectable in the SDSS sample with a similar approach. Correlation searches in spectra from other instruments such as FUSE and STIS, or even HIRES and UVES, may benefit from this kind of analysis, which is resistant to errors in the continuum. Modulation of the number of local pixels used may also be a useful approach to measuring clustering of absorbers.

6.2. Continuum Correction Uncertainty

We systematically underestimate our continuum fitting such that the continua must be corrected up (and the transmitted flux down) by 16% to match the mean flux measured in studies using high resolution spectra. The uncertainty in those measurements may lead to sys-

tematic errors in our measurements. We use $\langle F \rangle = 0.68 \pm 0.04$ at $z = 3$ from McDonald et al. (2006) and those 1σ error estimates correspond to allowed corrections of 11-25%. This would lead to a 15% offset in $\langle \tau_{\text{Ly}\alpha} \rangle$, which would be corrected by a 15% offset in the recombination rate. This change in recombination rate results in change to the H I fraction which is independent of density for the regime probed here, leaving the Ly α forest largely unchanged. There is a small distortion due to the fact that an error which is a multiple in flux would be corrected by a factor which is a multiple in optical depth, but crucially most of the signal in the LCP search comes from higher optical depths for which the correction is even smaller. Also large errors of this sort would be evident in the comparison between the PDF of the flux in the simulations and the data in Figure 4.

A systematic error in the chosen Γ (based on an error in the mean flux in the literature) would also impact upon the O VI fraction. This error corresponds to half a minor tick interval in the upper horizontal axis of Figure 9 and so ionization correction parameter would be largely unchanged over the density range of interest. One can conclude that in the context of (related) uncertainties in the shape and amplitude of the UV background and so Γ , these issues are a sub-dominant systematic uncertainty.

6.3. Other Searches for O VI in the IGM

Fox et al. (2008) argue that the lack of velocity dependence in the column densities (N) of their proximate ($\delta v < 8000 \text{ km s}^{-1}$ from z_{QSO}) sample indicates that their 26 weak O VI absorbers must be collisionally ionized. The majority of these absorbers likely correspond to systems of $\rho/\bar{\rho} > 10$ based on their associated H I column densities and using the $N - \rho/\bar{\rho}$ relation from Schaye (2001). Hence they are sampling predominantly stronger absorbers than those considered here, so this work is not at odds with our choice of a photoionization model for our calculation of $[\text{O}/\text{H}]$.

Our results appear consistent with other comparable measures of $[\text{O}/\text{H}]$ in the IGM, given the differences of method and thus the systems sampled. Carswell et al. (2002) also use high resolution spectra to search for O VI, but unlike Aguirre et al. (2008) they use Voigt profile fitting of detected lines. They argue that these lines at $z = 2$ are photoionized (and so are comparable with those we find), with abundances $10^{-3} - 10^{-2}$ of solar.

Telfer et al. (2002) use stacked O V in FUSE spectra to determine $[\text{O}/\text{H}]$ and obtain a range of -2.2 to -1.3. Their higher value is partly caused by a combination of their use of a ‘‘Q’’ UV background model and by their lower redshift range ($1.6 < z < 2.9$). The main cause for the discrepancy is likely their use of the 78th percentile of their stack (in order to maximize signal-to-noise), while we use the 50th. Aguirre et al. (2008) argue that this difference likely corresponds to a 0.5 dex adjustment, bringing the results of Telfer et al. (2002) broadly into line with ours.

Bergeron & Herbert-Fort (2005) search for O VI lines at somewhat lower redshift range ($2 < z < 2.6$) than ours. They divide these lines as ‘metal-poor’ and ‘metal-rich’. This differs greatly from our analysis and so is difficult to compare. In particular they find the $[\text{O}/\text{H}]$ of systems with detected O VI absorption while we calculate the average metallicity over all Ly α absorption.

Their ‘metal-poor’ sample shows $-3 < [\text{O}/\text{H}] < -1$ and a cosmic oxygen density $\Omega_{\text{Oxy}} \approx 2.3(\pm_{0.6}^{2.1}) \times 10^{-6}$ using a photoionization model similar to the ‘‘Q’’ model described above, which is consistent with our results, but the numbers are not directly comparable.

Frank et al. (2009a) have searched the same SDSS dataset with the goal of identifying strong O VI absorbers, requiring additionally associated Ly α and Ly β absorption. The approach used is complementary to this work, as they can only retrieve strong absorbers that are typically saturated in both these Lyman lines. Hence, these absorbers are a different population from the ones probed here with higher density and (perhaps) collisional ionization. Frank et al. (2009b) provide a measure of the cosmic oxygen density contributed by these systems of $\Omega_{\text{OVI}} \geq 1.9 \times 10^{-8} h^{-1}$ (which corresponds to a firm lower limit of $\Omega_{\text{Oxy}} \geq 1 \times 10^{-7} h^{-1}$ assuming the maximum O VI ionization fraction of 20%) at absorber redshifts $2.8 \leq z_{\text{abs}} \leq 3.2$.

Our measure of the fraction of the estimated metal budget detected at $z = 3$ is in line with values calculated by Bouché et al. (2007) using results from various papers and various metals. They infer a ‘forest’ contribution of $< 15\%$ at $z = 2$. One study they quote is particularly notable: Simcoe et al. (2004) find $\Omega_{\text{Oxy}} \approx 2 \times 10^{-6}$ using a ‘hard’ (‘Q’ like) UV background. Bouché et al. (2007) reanalyze these results in the context of a ‘softer’ (‘QG’ like) UV background model and find $\Omega_{\text{Oxy}} \approx 5.0 \times 10^{-6}$ at $z = 2.5$, and so a fraction of the metal budget of $\approx 30\%$. This measure is consistent with our findings given redshift evolution and systematic uncertainties associated with search methods and ionization corrections. Simcoe et al. (2004) find a mean $[\text{O}/\text{H}] = -2.85$ for their lines (which is effectively a volume-weighted mean) for the 70% of systems in which they detected oxygen, which is also consistent with our findings (setting the other 30% to zero metallicity would result in a volume-weighted mean of $[(\text{O}/\text{H})_{\text{VW}}] = -3.00$).

7. CONCLUSIONS

We have developed a new method to detect O VI absorption in low-density regions with a data set that might initially seem ill suited to this purpose, the moderate resolution ($R = 1800$), moderate signal-to-noise ratio SDSS spectra of high-redshift QSOs. Not only have we successfully detected O VI with high significance (the null case is ruled out with *reduced* $\chi^2 = 10$), but we have placed tighter limits on the mass-weighted mean oxygen abundance of low-density regions than ever before. The success of this method hinges on the sheer size of the SDSS sample of spectra and the use of local, uncorrelated pixels near the O VI absorber as a way of characterizing both the continuum fitting error and the degree of contaminating absorption. This locally calibrated pixel (LCP) search for O VI seen in SDSS DR3 is the main result of the paper and the high precision of this new tracer for metals in the IGM is clear. Modeling has been necessary to interpret these results and conclusions about oxygen abundance the lognormal random fields have been used to produce more than 4.5 million synthetic spectra.

Our search at $2.7 < z < 3.2$ is not sensitive to the density dependence of the $[\text{O}/\text{H}]$ because of spectral resolution constraints, but using detailed modeling we are able to measure the oxygen abundance IGM densities of

$1 \lesssim \rho/\bar{\rho} \lesssim 9$ and can conclude the following:

- A model with constant $[\text{O}/\text{H}] = -2.15_{-0.09}^{+0.07}$ provides an adequate fit to the data (at the 2σ level), but a fit can also be obtained by removing metals from lower density regions and elevating the metallicity in regions enriched.
- The addition of lognormal metallicity scatter leaves the quality of fit of these models largely unchanged, where the scatter free model is treated as a measure of the mean $[\text{O}/\text{H}]$ (or $\langle Z \rangle$) and *not* the mean $[\text{O}/\text{H}]$ (or $\langle \log Z \rangle$).
- The mass-weighted mean oxygen abundance is nearly constant among our viable models. As a result we have been able to place the tightest constraints thus far on this quantity, with $[(\text{O}/\text{H})_{\text{MW}}] = -2.45 \pm 0.06$ (at 1.5σ confidence) computed over the density range $-0.5 < \log[\rho/\bar{\rho}] < 2$. This value is based on a Grevesse et al. (1996) solar abundance of oxygen. A Caffau et al. (2008) solar abundance would give a abundance $[(\text{O}/\text{H})_{\text{MW}}] = -2.34 \pm 0.06$.
- The models that fit the data have volume-weighted mean, $[(\text{O}/\text{H})_{\text{VW}}] = -3.01 \pm 0.33$ (at 1.5σ confidence).
- We calculate the cosmic density of oxygen, $\Omega_{\text{Oxy,IGM}} = 1.4(\pm 0.2) \times 10^{-6}$. In the context of the work by Bouché et al. (2007) this constitutes a total metal contribution of $\Omega_{\text{Z,IGM}} \approx 2.3 \times 10^{-6}$, which is $\approx 16\%$ of their estimated metal budget at $z = 3$.

This novel LCP search method need not be limited to searches for metals in the IGM. It could be adapted to spectral searches wherever bulk, low-level correlated absorption is expected and continuum errors and contaminating absorbers are the limiting factor. Little prior knowledge of these errors and uncertainties is required, only that they are uncorrelated.

We thank Arlin Crofts, J. Xavier Prochaska and Romeel Davé for their useful comments.

Funding for the SDSS and SDSS-II has been provided by the Alfred P. Sloan Foundation, the Participating Institutions, the National Science Foundation, the U.S. Department of Energy, the National Aeronautics and Space Administration, the Japanese Monbukagakusho, the Max Planck Society, and the Higher Education Funding Council for England. The SDSS Web Site is <http://www.sdss.org/>.

The SDSS is managed by the Astrophysical Research Consortium for the Participating Institutions. The Participating Institutions are the American Museum of Natural History, Astrophysical Institute Potsdam, University of Basel, University of Cambridge, Case Western Reserve University, University of Chicago, Drexel University, Fermilab, the Institute for Advanced Study, the Japan Participation Group, Johns Hopkins University, the Joint Institute for Nuclear Astrophysics, the Kavli Institute for Particle Astrophysics and Cosmology, the

Korean Scientist Group, the Chinese Academy of Sciences (LAMOST), Los Alamos National Laboratory, the Max-Planck-Institute for Astronomy (MPIA), the Max-Planck-Institute for Astrophysics (MPA), New Mexico

State University, Ohio State University, University of Pittsburgh, University of Portsmouth, Princeton University, the United States Naval Observatory, and the University of Washington.

REFERENCES

- Abazajian, K. et al. 2005, AJ, 129, 1755
 Adelberger, K. L., Steidel, C. C., Shapley, A. E., & Pettini, M. 2003, ApJ, 584, 45
 Adelberger, K. L., Shapley, A. E., Steidel, C. C., Pettini, M., Erb, D. K., & Reddy, N. A. 2005, ApJ, 629, 636
 Aguirre A., Schaye J., & Theuns T., 2002, ApJ, 576, 1
 Aguirre A., Dow-Hygelund C., Schaye J., & Theuns, T. 2008, ApJ, 689, 851 (A08)
 Asplund, M., Grevesse, N., & Sauval, A. J. 2005 ASP Conf. Ser., ed Barnes, III, T. G. & Bash, F. N.
 Basu, S. and Antia, H. M. 2004, ApJ, 606, L85
 Bergeron J., Aracil B., Petitjean P., & Pichon, C. 2002, A&A, 396, L11
 Bi H. G., Boerner G., & Chu Y., 1992, A&A, 266, 1
 Bi H., & Davidsen A. F., 1997, ApJ, 479, 523
 Bergeron, J., & Herbert-Fort, S. 2005, (astro-ph/0506700)
 Bouché N., Lehnert M., Aguirre A., Péroux, C., & Bergeron J. 2007, 378, 525
 Carswell, B., Schaye, J., & Kim, T.-S., 2002, ApJ, 578, 43
 Caffau, E., Ludwig, H.-G., Steffen, M., Ayres, T. R., Bonifacio, P., Cayrel, R., Freytag, B., & Plez, B. 2008, A&A, 488, 1031
 Cowie, L. L., & Songaila, A., Kim, T.-S., & Hu, E., M. 1995, AJ, 109, 1522
 Cowie, L. L., & Songaila, A. 1998, Nature, 394, 44
 Croft, R. A. C., Weinberg, D. H., Katz, N., & Hernquist, L. 1998, ApJ, 495, 44
 Davé, R., Hellsten, U., Hernquist, L., Katz, N., & Weinberg, D. H. 1998, ApJ, 509, 661
 Delahaye, F., & Pinsonneault, M. H. 2006, ApJ, 649, 529
 Efsthathiou G., Bond J. R., & White S. D. M., 1992, MNRAS, 258, 1P
 Ellison, S. L., Songaila, A., Schaye, J., & Pettini, M. 2000, ApJ, 120, 1175
 Ferland, G. J., Korista, K. T., Verner, D. A., Ferguson, J. W., Kingdon, J. B., & Verner, E.M. 1998, PASP, 110, 761
 Fox, A., Bergeron, J., & Petitjean, P. 2008, MNRAS, 388, 1557
 Frank, S., Mathur, S., York, D., & Pieri M. M. 2009, submitted to AJ, (astro-ph/0707.1700)
 Frank, S., Mathur, S., York, D. & Pieri M. M. 2009, submitted to AJ, (astro-ph/0806.3071)
 Gibson, R. R., Jiang, L., Brandt, W. N., Hall, P. B., Shen, Y., Wu, J., Anderson, S. F., Schneider, D. P., Vanden Berk, D., Gallagher, S. C., Fan, X., & York, D. G. 2009, ApJ, 692, 758
 Grevesse, N., Noels, A., & Sauval, A. J. 1996, in ASP Conf. Ser., ed. S. S. Holt and G. Sonneborn, 99, 117
 Haardt, F., & Madau, P. 2001, in Clusters of Galaxies and the High Redshift Universe Observed in X-rays, ed. Neumann, D. M. & Tran, J. T. V., (astro-ph/0106018) (HM01)
 Hellsten U., Hernquist L., Katz N., & Weinberg D. H., 1998, ApJ, 499, 172
 Hui, L., & Gnedin 1997, MNRAS, 292, 27
 Kaiser N., & Peacock J. A. 1991, ApJ, 379, 482
 Kim, T.-S., Bolton, J. S., Viel, M., Haehnelt, M. G., & Carswell, R. F. 2007, MNRAS, 382, 1674
 Komatsu, E., Dunkley, J., Nolta, M. R., Bennett, C. L., Gold, B., Hinshaw, G., Jarosik, N., Larson, D., Limon, M., Page, L., Spergel, D. N., Halpern, M., Hill, R. S., Kogut, A., Meyer, S. S., Tucker, G. S., Weiland, J. L., Wollack, E. & Wright, E. L. 2008, submitted to ApJS, (astro-ph/0803.0547)
 Mac Low, M.-M., & Ferrara, A. 1999, ApJ, 513, 142
 Madau, P., Ferrara, A., & Rees, M. J. 2001, ApJ, 555, 92
 Meyer, D. M., & York, D. G. 1987, ApJ, 315, L5
 McDonald, P., Miralda-Escudé, J., Rauch, M., Sargent, W. L. W., Barlow, T. A., Cen, R., & Ostriker, J. P. 2000, ApJ, 543, 1
 McDonald, P., Seljak, U., Burles, S., Schlegel, D. J., Weinberg, D. H., Cen, R., Shih, D., Schaye, J., Schneider, D. P., Bahcall, N. A., Briggs, J. W., Brinkmann, J., Brunner, R. J., Fukugita, M., Gunn, J. E., Ivezić, Ž., Kent, S., Lupton, R. H., & Vanden Berk, D. E. 2006, ApJS, 163, 80
 Oppenheimer, B. D., & Davé, R. 2006, MNRAS, 373, 1265
 Pettini, M., Shapley, A. E., Steidel, C. C., Cuby, J., Dickinson, M., Moorwood, A. F. M., Adelberger, K. L., & Gialalisco, M. 2001, ApJ, 554, 981
 Pieri, M. M., & Haehnelt, M. G. 2004, MNRAS, 347, 985
 Pieri, M. M., Schaye, J., & Aguirre, A., ApJ, 638, 45
 Pieri, M. M., Martel, H., & Grenon, C. 2007, ApJ, 658, 36
 Pieri, M. M. & Martel, H. 2007, ApJ, 662, 36, L7
 Rauch M., Haehnelt M. G., & Steinmetz M., 1997, ApJ, 481, 601
 Rauch, M. and Miralda-Escudé, J., Sargent, W. L. W., Barlow, T. A., Weinberg, D. H., Hernquist, L., Katz, N., Cen, R., & Ostriker, J. P. 1997, ApJ, 489, 7
 Scannapieco, E., Pichon, C., Aracil, B., Petitjean, P., Thacker, R. J., Pogosyan, D., Bergeron, J., & Couchman, H. M. P. 2005, MNRAS, 365, 615
 Schaye, J., Aguirre, A., Kim, T.-S., Theuns, T., Rauch, M., & Sargent, W. L. W. 2003, ApJ, 596, 768 (S03)
 Schaye, J. 2001, ApJ, 559, 507
 Simcoe R. A, Sargent W. L. W., & Rauch, M. 2002, ApJ, 578, 737
 Simcoe R. A, Sargent W. L. W., & Rauch, M. 2004, ApJ, 606, 92
 Songaila, A., & Cowie, L. L. 1996, AJ, 112, 335
 Telfer, R. T., Kriss, G., Zheng, W., Davidsen, A. F., & Tytler, D. 2002, ApJ, 579, 500
 Viel, M., Matarrese, S., Mo, H. J., Theuns, T., & Haehnelt, M. G. 2002, /mnras, 336, 685
 Weinberg, D. H., Katz, N., & Hernquist, L. 1998, ASP Conf. Ser., ed Woodward, C. E., Shull, J. M., & Thronson, H. A.
 Wild, V., Kauffmann, G., White, S., York, D., Lehnert, M., Heckman, T., Hall, P. B., Khare, P., Lundgren, B. Schneider, D. P., & Vanden Berk, D. 2008, MNRAS, 338, 227
 Yip, C. W., Connolly, A. J., Vanden Berk, D. E., Ma, Z., Frieman, J. A., SubbaRao, M., Szalay, A. S., Richards, G. T., Hall, P. B., Schneider, D. P., Hopkins, A. M., Trump, J., & Brinkmann, J. 2004, AJ, 128, 2603
 York, D. G., et al. 2000, AJ, 120, 1579
 York, D. G., Khare, P., Vanden Berk, D., Kulkarni, V. P., Crotts, A. P. S., Lauroesch, J. T., Richards, G. T., Schneider, D. P., Welty, D. E., Alsayyad, Y., Kumar, A., Lundgren, B., Shanidze, N., Smith, T., Vanlandingham, J., Baugher, B., Hall, P. B., Jenkins, E. B., Menard, B., Rao, S., Tumlinson, J., Turnshek, D., Yip, C.-W., & Brinkmann, J. 2006, MNRAS, 367, 945