

University of Massachusetts Amherst

ScholarWorks@UMass Amherst

Astronomy Department Faculty Publication
Series

Astronomy

2003

The Lyman-alpha Forest as a Cosmological Tool

DH Weinberg

R Dave

N Katz

University of Massachusetts - Amherst

J Kollmeier

Follow this and additional works at: https://scholarworks.umass.edu/astro_faculty_pubs



Part of the [Astrophysics and Astronomy Commons](#)

Recommended Citation

Weinberg, DH; Dave, R; Katz, N; and Kollmeier, J, "The Lyman-alpha Forest as a Cosmological Tool" (2003). *AIP Conference Proceedings*. 980.

<https://doi.org/10.1063/1.1581786>

This Article is brought to you for free and open access by the Astronomy at ScholarWorks@UMass Amherst. It has been accepted for inclusion in Astronomy Department Faculty Publication Series by an authorized administrator of ScholarWorks@UMass Amherst. For more information, please contact scholarworks@library.umass.edu.

The Lyman- α Forest As A Cosmological Tool

David H. Weinberg^{*}, Romeel Dave[†], Neal Katz^{**} and Juna A. Kollmeier[‡]

^{*}*The Ohio State University, Dept. of Astronomy, Columbus, OH 43210*

[†]*University of Arizona, Dept. of Astronomy, Tucson, AZ 85721*

^{**}*University of Massachusetts, Dept. of Physics and Astronomy, Amherst, MA, 91003*

[‡]*Ohio State University, Dept. of Astronomy, Columbus, OH 43210*

Abstract. We review recent developments in the theory of the Ly α forest and their implications for the role of the forest as a test of cosmological models. Simulations predict a relatively tight correlation between the local Ly α optical depth and the local gas or dark matter density. Statistical properties of the transmitted flux can constrain the amplitude and shape of the matter power spectrum at high redshift, test the assumption of Gaussian initial conditions, and probe the evolution of dark energy by measuring the Hubble parameter $H(z)$. Simulations predict increased Ly α absorption in the vicinity of galaxies, but observations show a Ly α deficit within $\Delta_r \sim 0.5h^{-1}$ Mpc (comoving). We investigate idealized models of “winds” and find that they must eliminate neutral hydrogen out to comoving radii $\sim 1.5h^{-1}$ Mpc to marginally explain the data. Winds of this magnitude suppress the flux power spectrum by ~ 0.1 dex but have little effect on the distribution function or threshold crossing frequency. In light of the stringent demands on winds, we consider the alternative possibility that extended Ly α emission from target galaxies replaces absorbed flux, but we conclude that this explanation is unlikely. Taking full advantage of the data coming from large telescopes and from the Sloan Digital Sky Survey will require more complete understanding of the galaxy proximity effect, careful attention to continuum determination, and more accurate numerical predictions, with the goal of reaching 5 – 10% precision on key cosmological quantities.

PHYSICS OF THE FOREST

The 1990s saw four epochal advances in our understanding of the Ly α forest. Spectra of quasar pairs showed coherence over scales of a hundred kpc and more, implying large sizes and thus low densities for the absorbing structures [1, 2, 3, 4]. Keck HIRES spectra of unprecedented resolution and signal-to-noise demonstrated the ubiquity of weakly fluctuating Ly α absorption in the high redshift universe [5], and they revealed the presence of metal lines associated with low column density hydrogen absorbers [6, 7]. Finally, and most directly relevant to this review, a combination of numerical simulations and related analytic models led to a compelling new physical picture of Ly α forest absorption [8, 9, 10, 11, 12, 13, 14, 15].

The basic numerical result is simple to summarize: given a cosmological scenario motivated by independent observations, 3-d simulations that incorporate gravity, gas dynamics, and photoionization by the UV background produce something very much like the observed Ly α forest, an outcome that requires no *ad hoc* adjustments to the model. The top three rows of Figure 1 illustrate this point, showing, respectively, the observed Ly α forest of the $z = 3.62$ quasar Q1422+231, expanded views of four selected regions of this spectrum, and simulated spectra of the same length along four randomly selected lines of sight through a cosmological simulation. The simulation uses smoothed

particle hydrodynamics (SPH) with 128^3 dark matter particles and 128^3 gas particles in a periodic cube of comoving size $11.111h^{-1}$ Mpc (1422 km s^{-1} at $z = 3$). It assumes a Λ CDM model (inflationary cold dark matter with a cosmological constant), with $\Omega_m = 0.4$, $\Omega_\Lambda = 0.6$, $h = 0.65$, $\Omega_b = 0.02h^{-2} = 0.0473$, inflationary spectral tilt $n = 0.95$, and a power spectrum normalization that corresponds to $\sigma_8 = 0.8$ at $z = 0$. Spectra are extracted from the $z = 3$ simulation output using the methods of [12].

Studies of the Ly α forest have traditionally focused on “lines” identified by a decomposition procedure and the “clouds” or “absorbers” that produce them. Hydrodynamic simulations show that typical marginally saturated absorption features at $z \sim 3$ arise in filamentary structures, which are analogous to (but smaller in scale than) today’s galaxy superclusters. However, in the simulations there is no sharp distinction between the “lines” and the “background,” and one can also characterize the Ly α forest as “Gunn-Peterson” [16] absorption produced by a smoothly fluctuating intergalactic medium [12, 17, 18, 19].

What makes the fluctuating IGM perspective especially powerful is the tight relation between the density and temperature of low density cosmic gas, $T \approx T_0(\rho/\bar{\rho})^\alpha$ with $\alpha \approx 0.6$, which emerges from the balance between photoionization heating and adiabatic cooling [20, 21]. The Ly α optical depth of photoionized gas is $\tau \propto n_{\text{HI}} \propto \rho^2 T^{-0.7} \Gamma^{-1}$, where $T^{-0.7}$ accounts for the temperature dependence of the recombination rate and Γ is the rate at which neutral atoms are ionized by the cosmic UV background. The temperature-density relation then leads to the “Fluctuating Gunn-Peterson Approximation” (FGPA), $F \equiv \exp(-\tau) = \exp[-A(\rho/\bar{\rho})^\beta]$, which relates the continuum normalized flux F to the local gas overdensity [17, 22, 18, 19]. The index β lies in the range $1.6 - 2$ depending on the thermal history of the gas, and the proportionality constant A is itself proportional to $\Omega_b^2 h^3 T_0^{-0.7} \Gamma^{-1}$. Thus, one can think of a Ly α forest spectrum as providing a 1-dimensional, non-linear map of the gas overdensity $\rho/\bar{\rho}$ along the line of sight. The map is smoothed by thermal broadening and distorted by peculiar velocities, but these effects are small enough to leave a tight correlation between Ly α optical depth and local gas density even in redshift space [17, 23]. Furthermore, pressure gradients in the diffuse, photoionized IGM are usually weak compared to gravitational forces, so the gas overdensity traces the dark matter overdensity fairly well.

The FGPA provides a valuable way of thinking about the information content of the Ly α forest, and it can be a useful calculational tool if one has a way to produce realizations of cosmic density and velocity fields. The analytic models of the Ly α forest by [8, 24, 13, 14] essentially combine the FGPA with a log-normal or Zel’dovich approximation model for creating non-linear density and velocity fields. Alternatively, one can run an inexpensive, gravity-only N-body simulation and assume that the diffuse gas traces the dark matter (or add an approximate treatment of gas pressure [25, 26, 27]). The dotted lines in the third row of Figure 1 show spectra extracted from a particle-mesh (PM) N-body simulation, run from the same initial conditions as the SPH simulation, illustrating both the effectiveness and the limitations of this approach. The two spectra trace each other over most of their length, with the largest breakdowns occurring in regions where shock heating has pushed the gas above the $T \propto \rho^\alpha$ temperature-density relation, which makes the optical depth in the SPH simulation lower than the FGPA predicts.

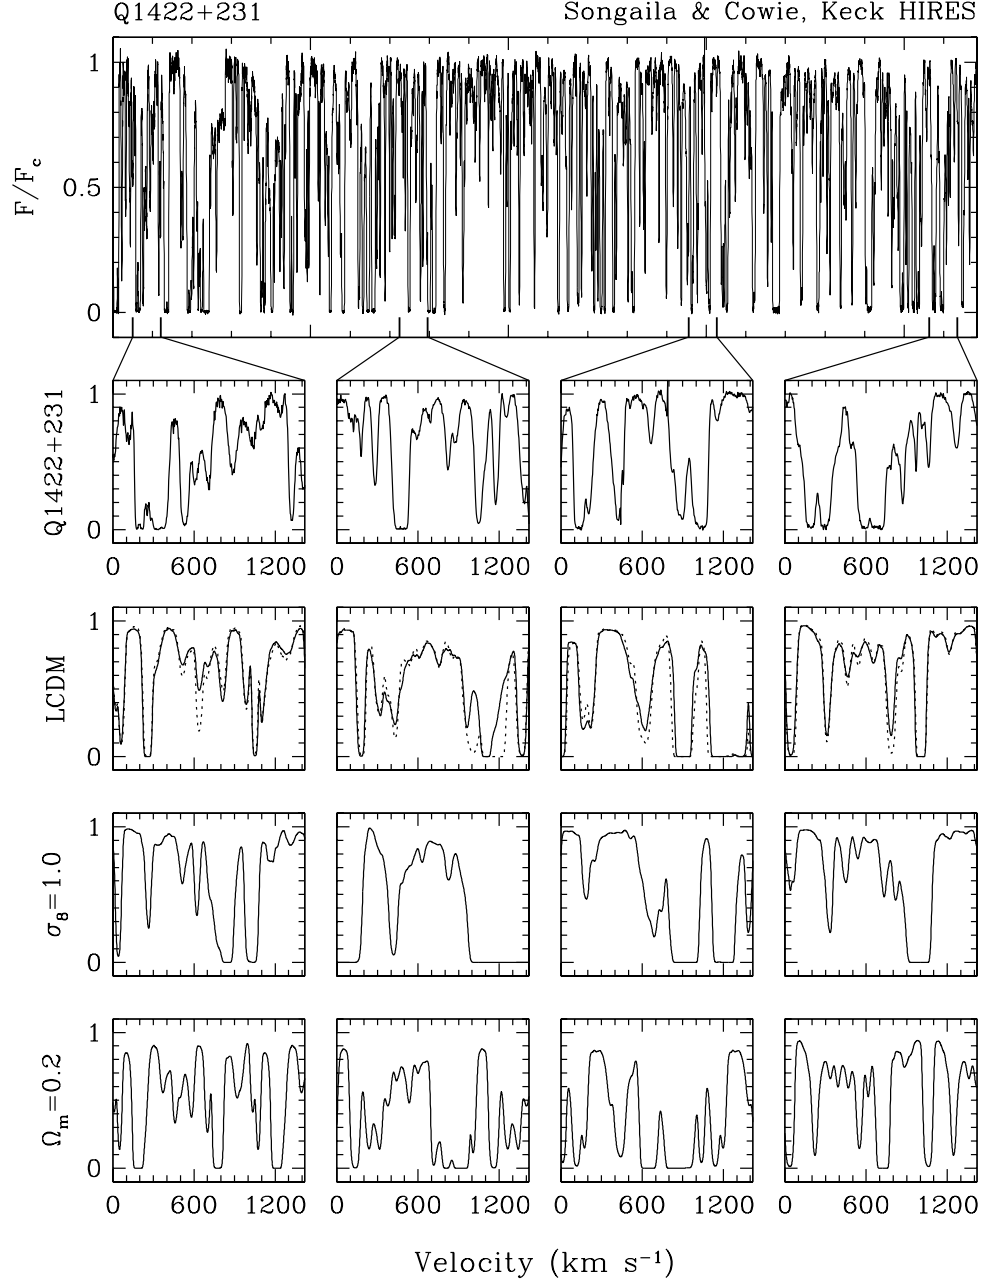


FIGURE 1. Simulating the $\text{Ly}\alpha$ forest. The top panel shows a continuum normalized, Keck HIRES spectrum of the $\text{Ly}\alpha$ forest region of the quasar Q1422+231 ($z = 3.62$), from [6]. The next row of panels shows blowups of four regions 1422 km s^{-1} in length. The third row shows simulated spectra extracted along four random lines of sight at $z = 3$ from an SPH simulation (solid) and a PM simulation (dotted) of the ΛCDM model, with $\Omega_m = 0.4$, $h = 0.65$, and $\sigma_8 = 0.8$ (at $z = 0$). The fourth and fifth rows show PM results for different cosmological parameter values: $\sigma_8 = 1.0$ and $\Omega_m = 0.2$, respectively. All simulated spectra are 1422 km s^{-1} in length.

The N-body+FGPA technique offers a convenient way to investigate the response of the Ly α forest to changes in cosmological parameters. Comparing the third and fourth rows of Figure 1 illustrates the effect of increasing the matter fluctuation amplitude by 20%, to $\sigma_8(z=0) = 1.0$, with the same Fourier phases in the initial conditions. The increased clustering of the high amplitude model is evident in the greater incidence of saturated absorption and in the merging of multiple small scale features into single larger scale features. The bottom row shows spectra from a model with $\sigma_8 = 0.8$ and the same linear power spectrum but a lower matter density, $\Omega_m = 0.2$. The reduction in Ω_m lowers the value of the Hubble parameter $H(z)$ at $z = 3$, and as a result features are more densely packed in redshift space, producing “choppier” spectra. In fact, our $11.111h^{-1}$ Mpc simulation cube is only 1024 km s^{-1} in length for $\Omega_m = 0.2$, and we have “padded” each spectrum to 1422 km s^{-1} by replicating the first 400 km s^{-1} . Reducing Ω_m also raises the fluctuation amplitude at $z = 3$ (since σ_8 is held fixed at $z = 0$ and there is less late time growth for lower Ω_m), but this effect is less important than the effect on the Hubble parameter.

The relative simplicity of the underlying physics and the existence of superb data at redshifts that are only sparsely probed by other observables make the Ly α forest a potentially powerful tool for testing cosmological models. Given independent estimates of Γ (e.g., from the quasar luminosity function or the proximity effect), one can use the mean opacity of the forest to obtain a lower limit to the cosmic baryon density [18, 28]. However, the observational constraints on Γ are loose, so for higher precision applications one must treat Γ as a free parameter and adjust it to match a single observable, usually taken to be the mean opacity. This normalization also absorbs uncertainties in Ω_b , T_0 , and h . We have followed this practice for each of the simulations in Figure 1 and will adopt it in all of our subsequent analyses. After normalizing to the mean opacity, the variation about the mean — i.e., the structure of the Ly α forest — is a prediction of the cosmological model, driven primarily by the structure in the underlying dark matter distribution. Ideally, one would like to calculate predictions for each cosmology using high resolution hydrodynamic simulations, but for some purposes other numerical or analytic approximations may be accurate enough. The adequacy of such approximations must be tested on a case-by-case basis, and the improving quality of the observational data places ever more stringent demands on the accuracy of the theoretical calculations.

CONTINUOUS FLUX STATISTICS

Because of the tight predicted correlation between the local Ly α optical depth and the local gas density, the most natural statistics for characterizing the Ly α forest are those that treat each quasar spectrum as a continuous 1-dimensional field, rather than a collection of discrete “lines.” We have reviewed this approach elsewhere [19, 23], and here we provide a brief recap with different examples and emphasis.

Figure 2 shows the power spectrum of continuum normalized flux, converted from 1-d to 3-d as described by [22], and multiplied by k^3 to yield the variance per $\ln k$. The curves are noisy because they are based on a single $11.111h^{-1}$ Mpc comoving cube, but we use the same phases in each simulation, so the relative behavior should be minimally




FIGURE 2. The flux power spectrum. Heavy and light solid curves show the 3-d flux power spectrum from, respectively, the SPH simulation and a PM simulation with the same initial conditions. Dotted and dashed curves show results with lower and higher matter fluctuation amplitudes, corresponding to $\sigma_8 = 0.6$ and 1.0 at $z = 0$. All results are from a single $11.111h^{-1}$ Mpc cube realized with the same Fourier phases.

affected by noise. At scales $k > 0.03 \text{ (km s}^{-1}\text{)}^{-1}$, the power spectra turn over because of the combined effects of non-linearity, peculiar velocities, and thermal broadening. The predictions in this regime are also affected by the finite resolution of the simulations. At larger scales, the SPH and N-body+FGPA methods give similar but not identical results for the same cosmological model, and the amplitude of the flux power spectrum increases with the amplitude of the matter power spectrum, as expected based on the FGPA and on Figure 1.

On large scales, the shape of the 3-d flux power spectrum is similar to that of the linear matter power spectrum $P_m(k)$ [22, 29]. The close connection between the shape and amplitude of the flux power spectrum and the shape and amplitude of $P_m(k)$ is the basis of Croft et al.’s [22] method for recovering $P_m(k)$ from Ly α forest data. Applying this method to a sample of 30 Keck HIRES spectra and 23 Keck LRIS spectra yields a matter power spectrum in remarkably (or, perhaps, disappointingly) good agreement with the “concordance” Λ CDM model favored by CMB, supernova, weak lensing, and low- z large scale structure data ([29]; see [30] for an independent analysis of the same flux power spectrum and [31] for an independent comparison to other cosmological constraints). McDonald et al. [32] reach similar conclusions from a “forward” comparison of hydrodynamic simulation predictions to the flux power spectrum measured from eight HIRES spectra. The Ly α forest power spectrum tests the Λ CDM model in a previously unexplored regime of redshift and lengthscale. It confirms one of the scenario’s key predictions, a linear power spectrum that bends from the primeval k^n towards k^{n-4} on small scales. The implied constraints on cosmological parameter combinations complement those from the CMB and other data. The Sloan Digital Sky Survey (SDSS) has already obtained moderate resolution spectra of several thousand high redshift quasars, and these will soon provide measurements of the flux power spectrum with much greater precision on large scales.




FIGURE 3. PDF of the flux decrement $D = 1 - e^{-\tau}$, measured from simulated spectra smoothed with a Gaussian of $\text{FWHM}=8 \text{ km s}^{-1}$ (left) or 100 km s^{-1} (right). Line types as in Fig. 2.

Figure 3 shows a different statistic, the probability distribution function (PDF) of the flux decrement $D = 1 - F = 1 - e^{-\tau}$, for the same set of simulations. The measurements in the left hand panel are from simulated spectra smoothed with a Gaussian of $\text{FWHM} = 8 \text{ km s}^{-1}$, comparable to the resolution of Keck HIRES or VLT UVES spectra. The SPH and N-body+FGPA predictions agree well for $\sigma_8 = 0.8$. Models with higher fluctuation amplitude have a broader distribution of densities ρ and a correspondingly broader distribution of flux decrements D , with more saturated and low-opacity pixels and fewer pixels of intermediate opacity. The differences in the predicted PDFs for $D \sim 0.3 - 0.7$ are $\Delta \log P \sim 0.1$ — not enormous, but readily measurable at high statistical significance with reasonable observational samples. The fraction of saturated pixels (defined here by $D > 0.96$) increases from 7.2% to 8.3% to 9.1% as σ_8 goes from 0.6 to 0.8 to 1.0. There are larger differences in the number of nearly transparent pixels, but continuum fitting uncertainties make it difficult to measure the PDF accurately near $D = 0$. One can investigate the scale dependence of matter clustering by smoothing the spectra (or by observing them at lower spectral resolution), analogous to studying galaxy counts in cells of increasing size. Figure 3 (right) shows results for 100 km s^{-1} smoothing. The predicted PDFs are narrower, since smoothing drives pixel values towards the mean, but the dependence on σ_8 is similar. The saturated pixel fraction doubles, from 1.9% to 3.8%, as σ_8 rises from 0.6 to 1.

For cosmological models with Gaussian initial conditions, the shape of the flux PDF depends mainly on the amplitude of mass fluctuations, once one has chosen Γ to match the mean opacity, [33, 23]. The effective physical scale of this amplitude measurement is determined by the spectral smoothing or, for spectra that fully resolve the observed absorption features, by a combination of thermal broadening and gas pressure effects. Models with non-Gaussian initial conditions predict significantly different flux PDFs [23]. Observational tests to date show good agreement with models that have Gaussian initial conditions and a $P_m(k)$ amplitude compatible with ΛCDM predictions [18, 23, 32]. The uncertainties in these tests are comparable to the model differences in Figure 3. Reducing them requires careful attention to the effects of continuum fitting and noise and




FIGURE 4. Threshold crossing frequency in spectra with $\text{FWHM}=8 \text{ km s}^{-1}$ (left) and 100 km s^{-1} (right). Curves show the number of times per unit redshift that the spectrum crosses a threshold of flux decrement D as a function of the filling factor, the fraction $F(D)$ of pixels with flux decrement less than D . Heavy and light solid curves show the SPH and PM results, respectively, for $\Omega_m = 0.4$. Dotted and dashed lines show PM results for $\Omega_m = 0.3$ and 0.2 , respectively.

to the accuracy of the theoretical predictions, but it should be possible to obtain tight constraints on the normalization of the matter power spectrum, and some information on its shape by studying different smoothing scales. Consistency between results from the flux power spectrum and the flux PDF can be a sensitive diagnostic for primordial non-Gaussianity.

Figure 4 shows the threshold crossing frequency, the number of times per unit redshift that the absorption spectrum crosses a decrement threshold D . Following [34], we plot dN_c/dz as a function of filling factor, the fraction $F(D)$ of pixels with flux decrement less than D , which cleanly separates the information in dN_c/dz from the information in the PDF and makes the model predictions nearly independent of the photoionization rate Γ . As one would expect from Figure 1, the threshold crossing frequency increases as Ω_m decreases because a lower Hubble parameter $H(z)$ “squeezes” redshift separations relative to comoving distances. The threshold crossing frequency drops slightly as the amplitude of $P_m(k)$ increases and gravitationally driven merging “smooths” structure [23], but the effect is small over the range $\sigma_8 = 0.6 - 1$. The crossing frequency also drops for redder matter power spectra, since these also lead to smoother structure [23]. For high spectral resolution, dN_c/dz also depends on the gas temperature, which determines the level of thermal broadening.

The great promise of the threshold crossing statistic is its potential for constraining the Hubble parameter at high redshift. In spatially flat models with a cosmological constant, the ratio $H(z)/H_0$ is determined by Ω_m . Alternatively, if Ω_m is known independently, the ratio $H(z)/H_0$ can constrain the equation of state of dark energy [35, 36]. Unfortunately, dN_c/dz also depends on other cosmological and IGM parameters, and it is numerically difficult to predict with high accuracy even when the model is fully specified. The difference between the PM and SPH predictions in Figure 4 is comparable to the Ω_m effects themselves, and the SPH result is still affected by the finite numerical resolution. On the observational side, accurate measurement of dN_c/dz requires high

signal-to-noise spectra. Furthermore, while moderate resolution spectra can provide useful diagnostics for the shape and amplitude of the matter power spectrum [23], the $H(z)$ application demands high spectral resolution, since smoothing over a scale that is fixed in redshift units erases the sensitivity to $H(z)$ (see Fig. 4, right). Exploiting dN_c/dz as a probe of dark energy thus represents a theoretical and observational challenge.

THE GALAXY PROXIMITY EFFECT

One can also use the observable correlations between Lyman Break Galaxies (LBGs) and the Ly α forest to study the environments of high redshift galaxies. Several groups have investigated this issue theoretically using simulations [37, 38, 39, 40], and Adelberger et al. have carried out an observational study using an LBG survey in fields probed by seven quasar lines of sight ([41], hereafter ASSP). On large scales — cubes of comoving size $\sim 13h^{-1}$ Mpc — ASSP find a clear correlation of galaxy overdensity with Ly α flux decrement, the expected signature of galaxy formation in overdense environments [37, 39]. However, the observed Ly α decrement *decreases* within $\Delta_r \approx 1h^{-1}$ Mpc of LBGs (comoving, redshift-space separation), where the simulations predict that absorption should be strongest [38, 39, 40].

Figure 5 illustrates this conflict. In the upper left panel, triangles show the predicted mean decrement in bins of Δ_r , while filled circles show the ASSP data points. The predictions come from an SPH simulation of a $22.222h^{-1}$ Mpc cube, and we select the 40 galaxies with the highest star formation rates to approximate the magnitude limit of ASSP’s spectroscopic survey (see [39] for details). Plausible random errors in galaxy redshifts can reduce the discrepancy at $\Delta_r \approx 1 - 2h^{-1}$ Mpc, but they cannot explain the results at the smallest scales, especially the innermost data point at $\Delta_r = 0 - 0.5h^{-1}$ Mpc, $\langle D \rangle = 0.11$ [38, 39, 42]. The right column of Figure 5 shows simulated spectra along 12 lines of sight that pass within $0.5h^{-1}$ Mpc of a target galaxy, centered on the galaxy redshift space position. All but one of these spectra have $D \geq 0.6$ at the galaxy redshift, and half have $D \geq 0.9$.

The obvious conclusion is that feedback from the observed galaxies reduces neutral hydrogen in their immediate surroundings. Local photoionization by the galaxies’ stars or active nuclei proves insufficient; this argument can be cast in general terms that seem difficult to escape [39]. The natural alternative is some form of supernova or AGN-driven wind [41, 38, 40]. Here (and in [42]) we have investigated highly idealized “wind” models in which we simply eliminate all neutral hydrogen in a sphere of comoving radius R_w around each target galaxy. Squares, pentagons, and hexagons in Figure 5 show results for $R_w = 0.75, 1.0$, and $1.5h^{-1}$ Mpc. Stars show a model where we place winds around all 641 resolved galaxies in the simulation volume (instead of the top 40 that constitute the “observed” sample) and scale the wind volume in proportion to the galaxy baryon mass, with a normalization $R_w = 1h^{-1}$ Mpc around the 40th-ranked galaxy. In the right column, dotted curves show spectra for the $R_w = 1.5h^{-1}$ Mpc model.

The most important lesson from Figure 5 is that eliminating neutral gas to a *real space* distance R_w does not eliminate absorption within a *redshift space* separation R_w (e.g., to $\Delta V = \pm 200$ km s $^{-1}$ in the spectrum plots). Peculiar infall velocities allow gas




FIGURE 5. Influence of idealized “winds” on the conditional mean flux decrement and the flux power spectrum, based on an SPH simulation of a $22.222h^{-1}$ Mpc (comoving) cube. Triangles in the upper left panel show the mean flux decrement in pixels that lie at redshift-space separation Δ_r from one of the 40 brightest galaxies in the cube. Results are averaged over $0.5h^{-1}$ Mpc (comoving) bins of Δ_r . Squares, pentagons, and hexagons show the effect of removing all neutral hydrogen in spheres of comoving radius 0.75, 1.0, and $1.5h^{-1}$ Mpc around these galaxies before extracting spectra, and stars show a model with sphere volume proportional to galaxy baryon mass. Filled circles show the observational estimates of ASSP. Sample spectra on the right illustrate the effect qualitatively. Solid curves show spectra along 12 lines of sight selected to pass within $0.5h^{-1}$ Mpc ($24''$) of a target galaxy at redshift space position $\Delta V = 0$. Dotted curves show the corresponding spectra for $R_w = 1.5h^{-1}$ Mpc. The lower left panel shows the flux power spectrum for the unmodified SPH simulation (solid) and the various “wind” models (as marked).

at larger distances to produce absorption near the galaxy redshift, making the energetic requirements on any wind explanation of the observed Ly α deficit much more stringent. Only the mass-scaled and $R_w = 1.5h^{-1}$ Mpc models come close to matching the ASSP data. Winds that fully ionize or perfectly entrain gas to such large distances are not a natural outcome of hydrodynamic simulations with stellar feedback [43, 40], and even reaching $1.5h^{-1}$ Mpc in ~ 1 Gyr requires a sustained propagation speed $\sim 600 \text{ km s}^{-1}$.




FIGURE 6. Influence of winds on the flux decrement PDF (top, as in Fig. 3) and the threshold crossing frequency (bottom, as in Fig. 4). Results are shown for the $22.222h^{-1}$ Mpc SPH simulation with no modification (solid) and with the “wind” models illustrated in Fig. 5.

Given the challenges facing the wind explanation, it is worth considering the alternative possibility that extended $\text{Ly}\alpha$ emission from the target galaxies is “filling in” the corresponding region of the absorption spectrum. Steidel et al. [44] observed two extended $\text{Ly}\alpha$ “blobs” apparently associated with LBGs, with angular extents $\sim 15''$ and AB apparent magnitudes 21.02 and 21.14 in the $\text{Ly}\alpha$ band. Cooling radiation from gas settling into massive galaxies at $z = 3$ naturally produces $\text{Ly}\alpha$ fluxes of this order [45, 46]. The three quasars that contribute to ASSP’s innermost data point have G -band AB magnitudes of 20.1, 21.6, and 23.4, so if *all* of a target galaxy’s $\text{Ly}\alpha$ cooling radiation went down the slit it could potentially replace the quasar flux absorbed by the surrounding IGM. However, a $1.4''$ slit at an angular separation $\Delta\theta \sim 15 - 20''$ from a galaxy should intercept at most $\sim 1.4/2\Delta\theta \sim 0.03 - 0.05$ of the galaxy’s extended $\text{Ly}\alpha$ flux, at least on average, so this explanation seems to fail by 1 – 2 orders of magnitude. Furthermore, a fourth pair involving the $G = 17.8$ quasar Q0302-0019, shows no sign of absorption near the galaxy redshift, and in this case the quasar is clearly too bright for galaxy emission to compete with it. (This pair and two others are dropped from ASSP’s $\langle D \rangle$ calculation because of possible $\text{Ly}\beta$ contamination.) At this point, the $\text{Ly}\alpha$ emis-

sion explanation seems unlikely; it can be conclusively ruled out by observing more close pairs involving bright quasars or by obtaining symmetrically placed spectra away from the observed quasars to search for galaxy emission.

Assuming for now that winds are the correct explanation for the observed Ly α deficit, one can ask whether they completely spoil the picture painted in §§1 and 2. The lower left panel of Figure 5 shows the flux power spectrum for the various “wind” models. Eliminating neutral hydrogen to distances $R_w = 0.75$ or $1h^{-1}$ Mpc around bright galaxies has only a small impact on the flux power spectrum, because the filling factor of the “bubbles” is small and absorption close to the galaxies remains nearly saturated in any case. However, these models also do not explain the ASSP results. In the two more extreme models, winds suppress the flux power spectrum on large scales by 0.1 – 0.2 dex, comparable to the 0.1-dex 1σ uncertainty that [29] quote for the normalization of the matter power spectrum. Thus, winds of this magnitude could have systematic effects on $P_m(k)$ determinations that are significant relative to the present observational uncertainties. The influence of winds on the flux PDF and threshold crossing frequency is smaller, as shown in Figure 6. Even the more extreme models have a negligible impact on dN_c/dz , and they only slightly alter the shape of the PDF. The most significant effect is on the fraction of saturated pixels for 100 km s^{-1} smoothing, which drops from 2.5% for the no-wind case to 1.9% for $R_w = 1.5h^{-1}$ Mpc and 1.3% for the mass-scaled model.

PROSPECTS AND CHALLENGES

From the above discussion, it is clear that one immediate challenge is to better understand the galaxy proximity effect. Figures 5 and 6 show that winds can reach substantial distances from bright galaxies without having much impact on the global statistics of the Ly α forest (see also [43, 40]), but even the more extreme wind models considered in §3 do not reproduce the ASSP results very well. Simulations with more realistic wind physics [38, 43, 40] can shed further light on this problem, but major progress will have to await further observational studies, since current inferences rest crucially on a handful of galaxy-quasar pairs.

Studies with Keck HIRES and VLT UVES have provided a (still growing) trove of high quality data on the Ly α forest at $z \sim 2 - 4$. The fluctuating IGM perspective, furthermore, shows that large samples of moderate resolution spectra can be a powerful resource for studying structure on large scales, since continuous flux statistics do not require resolution of individual “lines.” Such samples can be assembled quickly with large telescopes, and the SDSS is producing an enormous sample at resolution $R \sim 2000$ in the course of its normal operations. One of the observational frontiers is the use of correlations across multiple lines of sight. Quasar pair studies provided the first decisive evidence for large coherence scales of absorbing structures [1, 2, 3, 4], but larger samples allow more ambitious goals, such as using the Alcock-Paczynski test [47] to measure spacetime geometry [48, 49, 50], improving measurements of the flux power spectrum with cross correlations [51], and mapping large scale 3-dimensional structure at high redshift [52, 53].

Analyses of existing data have already led to an important cosmological conclusion,

namely that models with matter clustering similar to that of “concordance” Λ CDM at $z \sim 3$ are consistent with the observed Ly α forest while models with substantially different clustering amplitudes or suppression of small scale power are not [18, 54, 55, 23, 32, 56, 57, 58, 26, 29, 30]. They have also provided constraints on the temperature of the diffuse IGM [59, 60, 61] and indications of helium reionization at $z \sim 3.2$ [61, 62, 63, 64]. However, to do justice to the quality and quantity of data and keep pace with the tightening observational constraints from other observables, we must play for higher stakes. It looks possible in principle to achieve precision of 5 – 10% on quantities like the matter fluctuation amplitude and $H(z)$, but even without the potential complications of galaxy feedback, inferences at this level require more extensive theoretical modeling. Many effects that are unimportant at the 25% level — differences between approximate methods and full hydrodynamics, numerical resolution and box size limitations, spatial fluctuations in IGM temperature, details of continuum determination — may become critical at the 5 – 10% level. Despite these challenges, the Ly α forest is the most promising tool we have for precision cosmological measurements at $z \sim 2 - 4$. These measurements might, in the end, simply confirm the cosmological model favored by other data, but complementary constraints have the potential to break parameter degeneracies and thereby reveal subtle quantitative discrepancies. These in turn could yield insight into the nature of dark energy, the mechanisms of inflation, or some other fundamental aspect of our universe.

REFERENCES

1. Bechtold, J., Crotts, A. P. S., Duncan, R. C., & Fang, Y. 1994, ApJ, 437, L83
2. Dinshaw, N., Impey, C. D., Foltz, C. B., Weymann, R. J., & Chaffee, F. H. 1994, ApJ, 437, L87
3. Dinshaw, N., Foltz, C. B., Impey, C. D., Weymann, R. J., & Morris, S. L. 1995, Nature, 373, 223
4. Crotts, A. P. S., & Fang, Y. 1998, ApJ, 502, 16
5. Hu, E.M., Kim, T.S., Cowie, L.L., Songaila, A., & Rauch, M. 1995, AJ, 110, 1526
6. Songaila, A. & Cowie, L.L. 1996, AJ, 112, 335
7. Ellison, S. L., Lewis, G. F., Pettini, M., Chaffee, F. H., & Irwin, M. J. 1999, ApJ, 520, 456
8. Bi, H.G., 1993, ApJ, 405, 479
9. Cen, R., Miralda-Escudé, J., Ostriker, J.P., & Rauch, M. 1994, ApJ, 437, L9
10. Petitjean, P., Mückel, J. P., & Kates, R. E. 1995, A&A, 295, L9
11. Zhang, Y., Anninos, P., & Norman, M.L. 1995, ApJ, 453, L57
12. Hernquist L., Katz, N., Weinberg, D.H., & Miralda-Escudé, J. 1996, ApJ, 457, L51
13. Bi, H.G., & Davidsen, A. 1997, ApJ, 479, 523
14. Hui, L., Gnedin, N., & Zhang, Y. 1997, ApJ, 486, 599
15. Theuns, T., Leonard, A., Efstathiou, G., Pearce, F. R., & Thomas, P. A. 1998, MNRAS, 301, 478
16. Gunn, J.E., & Peterson, B.A. 1965, ApJ, 142, 1633
17. Croft, R.A.C., Weinberg, D.H., Katz, N., Hernquist, L., 1997, ApJ, 488, 532
18. Rauch, M., 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
19. Weinberg, D. H., Katz, N., & Hernquist, L. 1998, in ASP Conference Series 148, Origins, eds. C. E. Woodward, J. M. Shull, & H. Thronson, (ASP: San Francisco), 21, astro-ph/9708213
20. Hui, L., & Gnedin, N. 1997, MNRAS, 292, 27
21. Weinberg, D. H., Hernquist, L., & Katz, N. 1997, ApJ, 477, 8
22. Croft, R. A. C., Weinberg, D. H., Katz, N., & Hernquist, L. 1998, ApJ, 495, 44
23. Weinberg, D. H., et al. 1999, in Evolution of Large Scale Structure: From Recombination to Garching, eds. A.J. Banday, R. K. Sheth, & L. N. Da Costa, (Twin Press: Vledder NL), 346 astro-ph/9810142
24. Bi, H., Ge, J., & Fang, L.-Z. 1995, ApJ, 452, 90

25. Gnedin, N. Y., & Hui, L. 1998, MNRAS, 296, 44
26. Zaldarriaga, M., Hui, L., & Tegmark, M. 2001, ApJ, 557, 519
27. Viel, M., Matarrese, S., Mo, H. J., Theuns, T., & Haehnelt, M. G. 2002, MNRAS, 336, 685
28. Weinberg, D.H., Miralda-Escudé, J., Hernquist, L., & Katz, N., 1997, ApJ, 490, 564
29. Croft, R. A. C., Weinberg, D. H., Bolte, M., Burles, S., Hernquist, L., Katz, N., Kirkman, D., Tytler, D. 2002, ApJ, 581, 20
30. Gnedin, N. Y. & Hamilton, A. J. S. 2002, MNRAS, 334, 107
31. Tegmark, M. & Zaldarriaga, M. 2002, PRD, 66, 103508
32. McDonald, P., Miralda-Escudé, J., Rauch, M., Sargent, W. L. W., Barlow, T. A., Cen, R., & Ostriker, J. P. 2000, ApJ, 543, 1
33. Cen, R. 1997, ApJ, 479, L85
34. Miralda-Escudé J., Cen R., Ostriker, J.P., & Rauch, M. 1996, ApJ, 471, 582
35. Kujat, J., Linn, A. M., Scherrer, R. J., & Weinberg, D. H. 2002, ApJ, 572, 1
36. Viel, M., Matarrese, S., Theuns, T., Munshi, D., & Wang, Y. 2002, MNRAS, submitted, astro-ph/0212241
37. McDonald, P., Miralda-Escudé, J., & Cen, R. 2002, ApJ, 580, 42
38. Croft, R. A. C., Hernquist, L., Springel, V., Westover, M., White, M. 2002, ApJ, 580, 634
39. Kollmeier, J.A., Weinberg, D.H., Davé, R., Katz, N., 2002, ApJ, submitted, astro-ph/0209563
40. Bruscoli, M., Ferrara, A., Marri, S., Schneider, R., Maselli, A., Rollinde, E., Aracil, B., MNRAS, submitted, astro-ph/0212126
41. Adelberger, K.L., Steidel, C.C., Shapley, A.E., Pettini, M. 2002, ApJ, in press, astro-ph/0210314
42. Kollmeier, J.A., Weinberg, D.H., Davé, R., & Katz, N., these proceedings
43. Theuns, T., Viel, M., Kay, S., Schaye, J., Carswell, R. F., & Tzanavaris, P. 2002, ApJ, 578, L5
44. Steidel, C. C., Adelberger, K. L., Shapley, A. E., Pettini, M., Dickinson, M., & Giavalisco, M. 2000, ApJ, 532, 170
45. Haiman, Z., Spaans, M., & Quataert, E. 2000, ApJ, 537, L5
46. Fardal, M. A., Katz, N., Gardner, J. P., Hernquist, L., Weinberg, D. H. & Davé, R. 2001, ApJ, 562, 605
47. Alcock, C., & Paczyński, B. 1979, Nature, 281, 358
48. Hui, L., Stebbins, A., & Burles, S. 1999, ApJ, 511, 5
49. McDonald, P. & Miralda-Escudé, J. 1999, ApJ, 518, 24
50. McDonald, P. 2001, ApJ, submitted, astro-ph/0108064
51. Viel, M., Matarrese, S., Mo, H. J., Haehnelt, M. G., & Theuns, T. 2002, MNRAS, 329, 848
52. Liske, J., Webb, J. K., Williger, G. M., Fernández-Soto, A., & Carswell, R. F. 2000, MNRAS, 311, 657
53. Rollinde, E., Petitjean, P., Pichon, C., Colombi, S., Aracil, B., D’Odorico, V., & Haehnelt, M. G. 2002, MNRAS, submitted
54. Croft, R. A. C., Weinberg, D. H., Pettini, M., Katz, N., & Hernquist, L. 1999, ApJ, 520, 1
55. Theuns, T., Leonard, A., Schaye, J., & Efstathiou, G. 1999, MNRAS, 303, L58
56. Narayanan, V. K., Spergel, D. N., Davé, R., & Ma, C. 2000, ApJ, 543, L103
57. Theuns, T., Schaye, J., & Haehnelt, M. G. 2000, MNRAS, 315, 600
58. Meiksin, A., Bryan, G., & Machacek, M. 2001, MNRAS, 327, 296
59. Ricotti, M., Gnedin, N. Y. & Shull, J. M. 2000, ApJ, 534, 41
60. McDonald, P., Miralda-Escudé, J., Rauch, M., Sargent, W. L. W., Barlow, T. A., & Cen, R. 2000, ApJ, 562, 52
61. Schaye, J., Theuns, T., Rauch, M., Efstathiou, G. & Sargent, W. L. W. 2000, MNRAS, 318, 817
62. Theuns, T., Zaroubi, S., Kim, T., Tzanavaris, P., & Carswell, R. F. 2002, MNRAS, 332, 367
63. Theuns, T., Bernardi, M., Frieman, J., Hewett, P., Schaye, J., Sheth, R. K., & Subbarao, M. 2002, ApJ, 574, L111
64. Bernardi, M., et al. 2003, AJ, 125, 32