Section D: PROJECT DESCRIPTION

Now that the ΛCDM model is widely accepted, cosmology is advancing on two major fronts. One front uses precision observations to search for new fundamental physics in areas such as inflation, dark energy, CDM, and neutrino properties. The second front follows the formation and evolution of the objects and structures including the intergalactic medium (IGM), stars, black holes and galaxies. The Lyman-alpha forest (LYAF) absorption from the IGM provides unique information on both fronts.

The fluctuations in the density of the IGM are the smallest surviving relics from inflation. Even at moderate redshifts in the range 1.6 – 5, they remain close enough to linear density enhancements that hydrodynamic simulations can accurately follow their growth. At any one redshift, their amplitude depends on all the main cosmological parameters. The change in the amplitudes with redshift measures the growth of structure, at a rate that depends upon the equation of state of the dark energy. The LYAF is also sensitive to astrophysical events, especially ionization and heating by the ultraviolet background (UVB). There are additional localized effects from winds from galaxies that we can isolate and remove.

We propose an integrated set of observations, numerical simulations and analyses to extract cosmological and astrophysical information from the LYAF. We have been working on these topics for many years already, we have built many tools and our methods are working. This proposal is to continue refining these methods so we can extract valuable information from larger data sets.

Our ultimate goal is to try to measure the masses of the three light active neutrinos. To do this we must improve both the data and simulations that we use to measure the amplitude of the matter power spectrum on the small-scales probed by the LYAF. We choose this exciting and ambitious goal because going down this path is the best way to extract information in general from the IGM. Existing LYAF spectra may already be sufficient to measure the sum of their masses. Neutrino oscillation experiments have shown that the sum of the neutrino masses is > 0.05 eV. The best upper limits, all from cosmology, are now near 0.7 eV (WMAP 7yr + BAO +SN Komatsu et al. 10 table 2). A large mass near 0.7 eV would suppress the power in the LYAF range by about 30% at z=3.5, a large signal that we could detect in existing data when we have reduced the dominant systematic errors from the uncertainty in the IGM model. We will be able to detect masses over 75% of the 0.05 – 0.7 eV range, giving a significant chance that we will detect the mass. Pursuing this goal is the best way to prepare for the much larger samples of low-resolution spectra that the SDSS is obtaining. The Astro2010 decadal survey panel on cosmology and fundamental physics picked neutrino (astro)physics as one of four hot topics.

Our activities result in many valuable products:

1. We will continue observing to build a definitive set of high-resolution LYAF spectra to specify the physical properties of the IGM. High-resolution spectra are essential for three reasons. First, they resolve the Lyman-alpha (Lya) lines and give line widths and information on the IGM thermal history that is missing from lower resolution spectra (e.g. SDSS). With this information, we can ensure we are using physically realistic simulations. Second, they give the most accurate measurements of certain statistics of the LYAF, such as the mean amount of Lya absorption. Third, the main effect of neutrino mass on the power of the flux in the LYAF (Fig 6) is entirely at very small-scales that are sampled in only high-resolution spectra. We will also obtain lower resolution spectra to sample higher z (Keck LRIS, ESI) and lower z (Lick) than SDSS or BOSS.
2. We will prepare the high-resolution spectra by removing flux calibration errors and metal lines, and we will measure the continuum levels. From these prepared spectra, we will calculate complete and consistent sets of statistics, all from the same data. These new statistics will avoid the inconsistencies in published statistics that come from many different samples of spectra prepared in different ways.
3. We will run suites of large hydrodynamic simulations to find those that are consistent with the LYAF data over the range of z. These simulations give the set of cosmological and astrophysical parameters that describe the IGM. Parameters obtained from the IGM are independent of those from the CMB, galaxies or clusters. Although we now have simulations that match LYAF data, we

will need to refine this model as we include larger samples. We have experience with this multi-dimension matching that remains a major task.

4. Using the simulations to match the high-resolution and low-resolution spectra (from Lick, Keck and SDSS) we obtain the most accurate measurements of the amplitude and growth rate of small-scale primordial perturbations. This leads to the detection or upper limit on the neutrino masses.
5. We will be more accurate than past measurements because we will use about ten times more high-resolution spectra. Most information on the IGM physics and neutrino masses is only available at this high-resolution. We focus on the higher redshifts that most clearly show the neutrino mass (Fig 6). We prepare spectra by removing flux calibration errors and metal lines (not normally done) and we add continua (metals and continua are much larger problems with low-resolution spectra). We measure complete sets of statistics all from the same data, and we use large hydrodynamic simulations to match all the LYAF statistics simultaneously.
6. The key to measuring the neutrino masses, or any other information from the IGM, is to find a realistic model of the IGM. We will continue to observe and analyze absorption lines in pairs of QSOs that are within about 0.5 Mpc in the sky to explore the topology of the LYAF and to limit the influence of chemical and thermal feedback in the IGM from QSOs and galaxies.
7. The shape of the matter power spectrum (slope n, rolling of n) can be used to constrain the parameters of slow roll inflation (Viel, Weller & Haehnelt 2004b; Seljak et al. 2005).
8. The shape of the power on the small scales probed by HIRES is used to constrain the mass of warm dark matter particles that are leading candidates for part of the dark matter. Current limits from the forest are > 1.2 keV for early decoupled thermal relics and > 12 keV for sterile neutrinos (Viel et al. 08; Boyarsky et al. 09). Improved limits might allow us to rule out such dark matter (as prematurely claimed by Seljak 06), since limits on the X-ray flux from the annihilations show that the masses are less than about 10 - 40 keV. In models with supersymmetric gauge mediation, gravitinos must have a mass < 16 eV to be consistent with the small-scale power spectrum measured in the LYAF (Viel et al 05 PhysRevD 71 063534).
9. The simulations that match the LYAF data give the most accurate measurements of the intensity of the UVB (for H ionizations and somewhat for HeII). We need the UVB intensity to calculate the ionization of QSO absorbers. It constrains the total emissivity by all UV sources at earlier times, reduced by the amount of intervening absorption and modified by re-emission.
10. The same simulations will also give by far the most accurate measurement of the temperature of the IGM at $1.6 < z < 3.5$. The heat comes from the UVB, from shocks in structure formation and from any additional input. We need to know the temperature structure to find the matter power spectrum that gives the observed LYAF statistics. We will give upper limits on the extra heat input that others may use to limit late decaying particles (Sciama 94, Sethi 97), an early X-ray background (Ricotti & Ostriker 04) or early cosmic rays that Samui et al. 05 claim will escape into the IGM and may provide measurable heating down to redshift 2.

1. Sum of the masses of the three light active neutrinos

Neutrinos are amongst the commonest particles in the universe. They were created in abundance in the early universe, they help determine the numbers of nuclei made during Big Bang nucleosynthesis, and they have restricted the growth of galaxy sized density perturbations. They are created today in the cores of stars, in supernovae explosions, in cosmic ray interactions, and on earth in nuclear fission reactors and when accelerator beams strike targets. Yet they were not detected until 1953 because their interactions with other matter are exceedingly weak, being limited to gravity and the weak force.

While we have just seen two decades with momentous discoveries in neutrino physics, many properties of the neutrinos remain unknown (Haxton 08). The decay of the Z^0 has shown that there are only three active (interacting through the weak force) light neutrinos. The discovery that these neutrinos oscillate between their electron, mu and tau flavors as they travel, and hence they have mass, is amongst

the first direct signs of new physics beyond the Standard Model of particle physics. Oscillation experiments have measured the differences of the squares of the masses; $|\Delta m^2_{32}| = 2.43 \pm 0.13 \times 10^{-3}$ eV and $|\Delta m^2_{21}| = 7.59 \pm 0.21 \times 10^{-5}$ eV (Adamson et al. 2008; KamLAMD 2008). The larger of these differences requires that at least one neutrino has a mass of 0.05 eV of more, hence the sum of the masses of the three light neutrinos is $\Sigma m_\nu > 0.05$ eV. We do not know the masses of any of these neutrinos. We do not know the mass orderings of the neutrino eigenstates. There are two possible arrangements of the masses (Fig 2): the "normal" ordering $m_1 < m_2 < m_3$ and the "inverted" ordering $m_3 < m_1 < m_2$, depending on whether the m_3 is the largest or smallest mass. We do not know if the neutrinos are their own anti-particles (Majorana particles). We do not know if there are equal numbers of neutrinos and anti-neutrinos in the universe. We do not know the strength of CP violations for the neutrinos. If these violations are much stronger than amongst baryons, they might explain why there is more baryonic matter than antimatter, through the much favored leptogenesis mechanisms (Davidson et al. 08 Phys Rep 466 105).

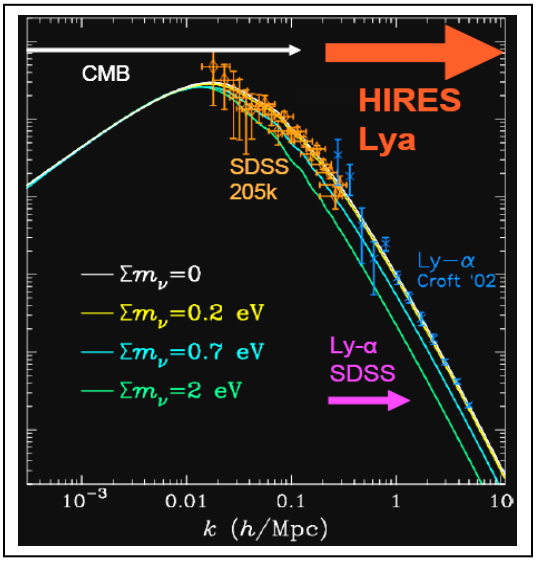

FIG 1 The vertical axis shows the power spectrum of matter, P(k) as a function of the wavenumber k. LYAF covers k = 0.2 – 10 (h/Mpc) all the way to the far right, which correspond to k=0.002 to 0.1 (s/km). HIRES spectra are most useful from 1 – 10 h/Mpc (0.01 – 0.1 s/km). The LYAF errors are larger on larger scales, (blue left of center) because there are fewer large modes in the spectra and because of continuum level errors. On the smallest scales on the far right random errors are small. In published LYAF results, systematic errors dominate on all scales. SDSS 205k is the power from galaxies at z = 0.1 (Tegmark et al. 04 ApJ 606 702) (modified from a figure by Kevork Abazajian).

From this long list of opportunities for discovery, the neutrino masses are especially interesting because they are the lowest masses of any particles, orders of magnitude below the second lightest particle, the electron. A major challenge of theories beyond the Standard Model of Particle Physics is to explain these tiny masses. In the popular seesaw mechanism, the neutrino masses are suppressed by M_D/M_R, where M_D is the Dirac mass typical of other fermions and M_R the Majorana mass that reflects a new very high-energy scale. Using the approximate neutrino mass from the oscillation measurements suggests that M_R is about 10^{15} GeV, similar to (an order of magnitude below) the Grand Unification Scale from supersymmetry. Hence, a cosmological measurement of the sum of the neutrino masses could illuminate physics beyond the Standard Model up near the GUT scale. A cosmological measurement of the sum of the neutrino masses would also greatly help in the interpretation of new neutrino double beta decay experiments (assuming they see a signal and that reactor experiments give the unknown third mixing angle θ_{13}). These theories also invoke new particles, sterile neutrinos that lack even weak interactions but might have a measurable effect on the early universe. The best laboratory upper limits on the neutrino masses give $\Sigma m_\nu < 6.6$ eV. The KATRIN tritium beta decay experiment has a design goal of 0.6 eV, similar to current cosmological limits. Cosmological measurements are the only known way to probe much lower masses.

We can detect the mass of neutrinos because they suppress growth of density perturbations on scales smaller than their free-streaming length. In the radiation-dominated era, perturbations on the sub-horizon scale grow logarithmically. Massive neutrinos suppress this growth lowering the power spectrum of matter in general. On scales smaller than the free streaming scale, the neutrinos stream out of high-density regions and into low-density regions. The lowering of the power spectrum amplitude is approximately

proportional to the sum of the masses, which is for z=0 $\Delta P/P = -8\Omega_\nu/\Omega_m = 0.03 - 0.4$ (Hu et al. 98) where $\Omega_\nu = \Sigma m_\nu/(94.22\ h^2 eV) = 0.0012 - 0.015$ for masses 0.05 – 0.7 eV.

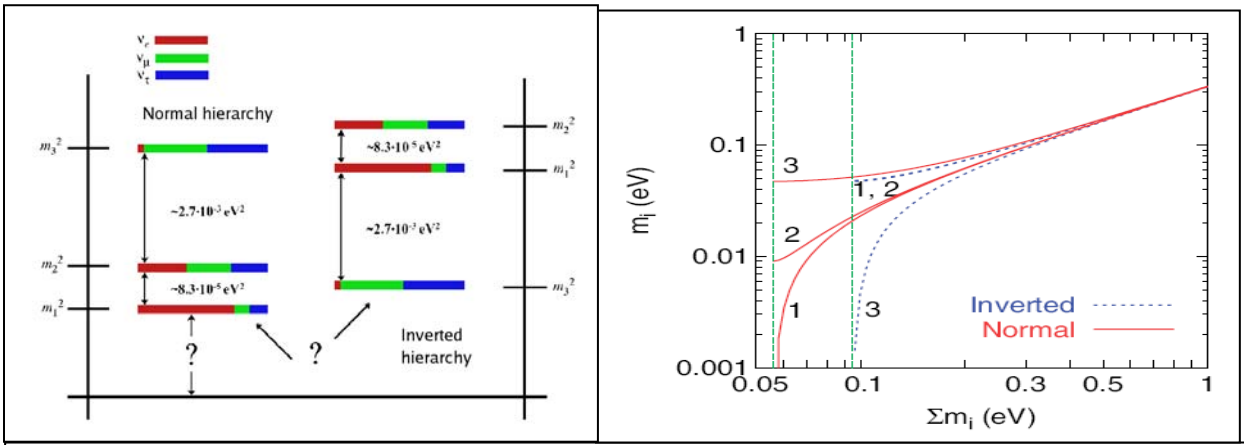

FIG 2 [**Left**] The electron, muon, and tau neutrinos are each a mixture of three "mass states" m_1, m_2, and m_3. Oscillations give two differences of the squares of the masses (vertical arrows). We do not know the masses (suggested by the left "?"). The ordering where m_3 is the most massive is the normal hierarchy on the left while the inverted hierarchy is on the right. (From Abdalla & Rawlings, 2007, MNRAS, 381, 1313.) [**Right**] Masses of the neutrino "mass states" 1, 2, and 3 as a function of the sum of the masses that we know is between 0.05 V and 0.7 eV. If measurements show that the sum of the masses is <0.1 eV, we will know that the hierarchy is normal and m_3 is near 0.05 eV. (From Lesgourgues & Pastor, 2006, Phys. Rep. 429, 307.)

The comoving free streaming wavenumber changes with z, and reaches a minimum when the neutrinos become non-relativistic, $k_{nr} \approx 0.018(m_\nu/1\ \mathrm{eV})^{1/2}(\Omega h^2)^{1/2}$ Mpc^{-1} (Agarwal & Feldman 2010). Even this minimum corresponds to relatively large scales, much larger than those probed by the LYAF. Hence, the best limits on the sum of the neutrino masses come from comparing the power on very large scales (from the CMB) to that on small-scales. The LYAF probes the smallest scales of any relevant measures and hence gives the most leverage because the power suppression by neutrinos continues to rise to the smallest scales (Fig 1). Other effects that may be competitive include lensing of the CMB (Planck $\Sigma m_\nu >$ 0.29 eV Shimon et al. 10, dePutter et al. 09), the galaxy distribution (SDSS, DES, BOSS 0.6 eV, LSST 0.1 eV Viel et al. 10, Lesgourgues et al. 08), galaxy lensing (LSST JDEM 0.05 eV), 21 cm galaxy surveys (LOFAR, MWA 0.1 eV; SKA 0.007 eV Mao et al. 08) and galaxy clusters (0.1 eV Lahav et al.10).

2. Results from Prior NSF Support

We discuss AST0507717 PI Tytler *Precision Measurements in the Intergalactic medium: Simulations and Observations* 08/01/05-07/31/08 and the follow-up *Integrating precision observations and simulations of the Intergalactic Medium* AST0808168 $799,158 08/01/08- 07/31/11 PI Tytler.

Paired QSOs. We obtained and analyzed the first large sample of absorption systems in paired QSOs with lines-of-sight separated by about 1 Mpc at 0.2 < z < 4. For our first paper Tytler et al. 09a we found 691 absorption systems in the spectra of 310 QSOs in 170 pairings, including 10 triples. These absorption systems show C IV or Mg II, visible in moderate resolution spectra. We since doubled the sample of close absorbers using spectra we obtained from Magellan. We now see 34 cases of absorption in one line-of-sight within 200 km/s (1 Mpc) of absorption in the paired line-of-sight. When we see absorption in one line-of-sight, the probability of also seeing absorption within about 500 km/s in the partner line-of-sight is at least 50% at < 100 kpc, declining rapidly to 23% at 100 – 200 kpc and 0.7% by 1 – 2 Mpc. The absorber-absorber correlation is primarily a probe of the large-scale distribution of metals around galaxies. Since the absorption redshifts have errors of about 23 km/s, we readily detect clustering on 0.5

Mpc scales. Redshift-space distortion (peculiar velocities) make the absorber-absorber correlation elongated along the line-of-sight, as in Fig 3 where we show our results from Lick, Keck and Magellan.

QSO MgII and CIV absorbers at $z = 2$ appear to arise in galaxies with a normal correlation function, normal systematic in-fall velocities and low random pair-wise velocity differences. The absorbing gas comes from a "cold" population most consistent with quiescent gas near blue galaxies. Blue galaxies have a pair-wise velocity dispersion of $\sigma_{12} = 200 – 400$ km/s while red galaxies have a dispersion of 500 – 800 km/s that is much larger than we see.

Absorption in gas flowing out from galaxies at a mean velocity of 250 km/s (1.2 Mpc on Fig 3) would produce much more extension in redshift space (smearing the clump in the lower left of Fig 3) than we see. The UV absorption from fast winds that Adelberger et al. 05 see in spectra of the UV luminous star forming regions of Lyman break galaxies (LBGs) is not representative of UV absorption that we see. Either the winds are confined to LBGs, or if they are common to all galaxies, they cannot extend to >40 kpc (where we detect the absorption) with large velocities, while continuing to make UV absorption that we can detect. The implication is that most metals were in the IGM at much earlier epochs, a claim earlier made by Rauch et al. 02 from the small shear velocities in the Lyman alpha forest. Much of this work (finding the absorbers, measuring the z, writing the analysis code) was done by undergraduates.

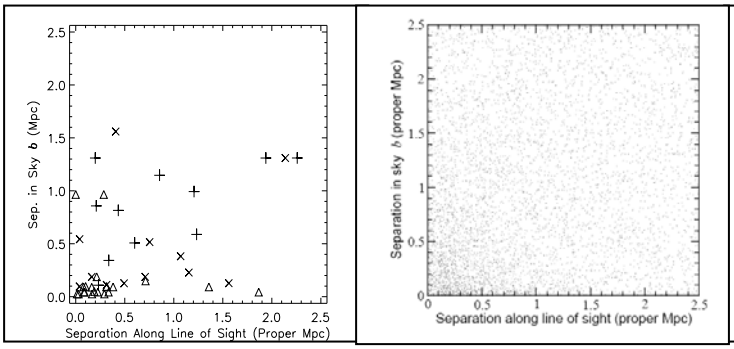

FIG 3. Evidence that QSO Mg II and C IV absorbers do not arise in fast moving winds. The separation of pairs of absorbers in the sky (vertical) and z (horizontal). One absorber at the origin, the partner is plotted. Left: data from our spectra of pairs of QSOs. Right: Monte-Carlo simulation with clustering, systematic infall and random pair-wise motions of 240 km/s.

We have also measured the amount of H I absorption (Kirkman & Tytler 08) and Mg II or C IV absorption (Tytler et al. 09a) around QSOs. For H I and Mg II, we see no change in the absorption in front of the QSOs and excess absorption out to 6 Mpc behind the QSOs. For CIV, the distribution is symmetric about the QSOs, using emission redshifts from the MgII emission line. We believe that the gas and galaxy density is symmetrically enhanced around the QSOs (Loeb & Eisenstein 95; Guimaraes et al. 07a). In front of the QSOs, the extra UV (that we also see) cancels the enhanced density leaving no change in the Lya and Mg II absorption, and extra C IV. Behind the QSOs, there is much reduced UV from the QSOs, so we see the excess density as more Lya and Mg II absorption. The QSO UV emission may be significantly anisotropic, with a 120 degree beam opening angle. A second possibility is that QSOs have had the UV luminosity we see for <1 Myr, and the gas behind the QSOs was exposed to a lower flux.

In his first year as an undergraduate David Gross led a project to measure the difference between the redshift given by the peak of the Mg II emission line (that we can readily observe for most QSOs with $z < 2.9$) and the narrow emission lines from [O II] and [O III]. These Oxygen lines give the systemic redshifts of the QSOs host galaxy, but they are rarely observed in the QSOs we use for LYAF statistics because they are shifted into the near IR. He found that Mg II is an excellent surrogate for the narrow O lines. The mean velocity of Mg II in the frame of [O II] is $–38 \pm 20$ km/s with a standard deviation of 343 km/s, while the mean Mg II velocity in the frame of [O III] is $–4 \pm 42$ km/s with a standard deviation of 336 km/s. These results show that on average the peak of Mg II is sufficiently close to the systemic velocity for our work, and though the dispersion is larger than ideal this is not enough to justify NIR spectra.

LYAF work. We have been working to find numerical simulations of the IGM that can precisely reproduce the statistical descriptions of the LYAF in QSO spectra. In Tytler et al. 09b we used the hydrodynamic code ENZO to run six simulations using the same cosmological and astrophysical parameters and cell size, but different box size. We found that larger boxes have fewer pixels with

significant absorption (flux < 0.96 where 1 is the continuum), they have more power on larger scales and they are hotter (longer modes give higher velocities and stronger shocks). These differences will be important when we compare to larger numbers of QSO spectra, e.g. the flux distribution from a small box will look like that with neutrino mass (top row of Fig 6).

We examined whether any of our standard simulations could match the published LYAF data at any z from 1.6 to 3.2. At each redshift, we examined simulations with arbitrary values for the astrophysical parameters that control the ionization and heating of the gas. We found that the best fits were worse at the higher redshifts, in part because the simulations that had the same power spectrum of the flux (**flux power**) as the QSO spectra had more Lya absorption than do QSO spectra. Fig 4 shows simulations can match several statistics of the LYAF at once, although in this case not the large-scale power from Kim et al. 04.

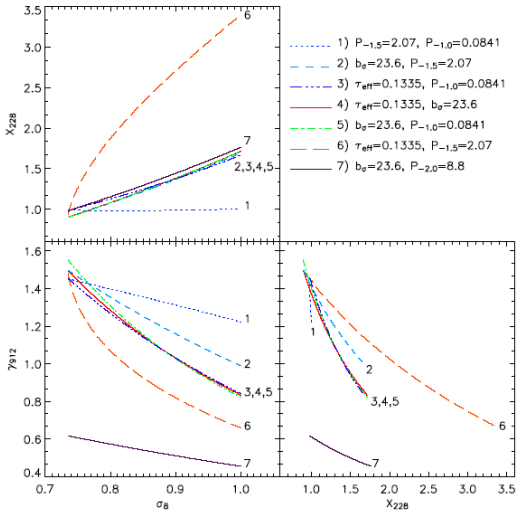

FIG 4 The parameters γ_{912} (H ionization), σ_8, and X_{228} (heat input - see Section 3.4) for simulations that match various pairs of LYAF statistics. The three panels are the sides of a 3D cube in which a simulation is a point. The lines show values that give simulated spectra which match various published statistics from QSO spectra at z=2. We show the power spectrum on three scales ($P_{-1.5}$ is is the power of the flux in the LYAF at wavenumber log k = -1.5 s/km), ($P_{-1.0}$ is at log k = -1.0 s/km), and ($P_{-2.0}$ is at log k = -2 s/km), b_σ is the typical b-value at z=2 in km/s and τ_{eff} is a measure of the mean absorption by Lya forest lines. Where the lines intersect (γ_{912}=1, σ_8=0.74 and X_{228}=1) we are matching 4 of these 5 statistics simultaneously. This family of simulations can match the power on the small and intermediate scales, but not the large-scale power (data from Kim et al. 2004) shown by line 7.

We went into ENZO and changed the adjustable heat input (per He II ionization) to become proportional to some power of the gas density, duplicating work by Bolton et al. 08 who suggested this ad hoc change might possibly emulate the heating with radiative transfer effects. This slightly changed the shape of the power spectrum of the LYAF flux, but it was not a major factor. McQuinn et al. 09 find that HeII ionization cannot produce the inverted temperature-density relation (low density gas hotter than high density) that Becker at al. 06 proposed and Bolton et al. 08 found gave a simulation that better matched the flux pdf measured by Kim et al. 07.

We have built a **pipeline** to increase the productivity and accuracy of our simulations and analyses. The pipeline generates hydrodynamic simulations, it generates synthetic spectra and it measures the statistics of the LYAF in these spectra. We also use the same pipeline to measure the statistics in the real spectra, where we must deal with the wavelengths that we reject because they contain metal lines or their errors are too large. We can now compare tens of simulations with data sets at many z.

Simultaneously with this work on the simulations, we have been working with HIRES spectra of a set of 30 QSOs that all have moderate to high SNR coverage of the LYAF. We have applied flux calibrations to all but a few of these spectra. We identified the metal lines both inside and outside of the LYAF and fit continua that we believe have errors of about 1%. With this data set, we then measure the entire set of LYAF statistics, all on the same block of data separated by SNR and z (first time this has been done). We found that these statistics are much closer to our simulations. One reason why our simulations did not match published LYAF statistics was that those statistics were measured in different samples of QSO spectra that had been prepared in different ways. Published LYAF statistics are in general not consistent.

Metal lines. Over the last few years, we have developed and applied methods to identify all the metal lines in HIRES spectra of 30 QSOs. The previously largest such set was 18 QSOs (Kim et al 07a). These spectra have enough resolution to readily show individual metals in the LYAF, and they have a range of

SNR, allowing us to explore how SNR changes the results. We find metal lines using various plots that highlight redshift systems together with a long and complex computer code that graduate student XiHeng Shi wrote to find possible redshift systems inside the LYAF. This code considers line positions, line widths and ionization. At our mean redshift (2.67) we find that metal lines that we identify absorb 1.5±0.2 % of the flux to the red of Lya emission but only 1.3±0.2 % of the flux inside the LYAF (1070 – 1190A away from emission lines where the continuum is less hard to estimate). We expect more, not less, absorption inside the LYAF because several ions (e.g. SiIII, FeIII) can absorb inside the LYAF but not outside. We estimate the total absorption due to these LYAF-only ions by multiplying their detected absorption by the fraction (absorption outside LYAF)/(absorption inside), determined for other ions that we can see on both sides of Lya. Outside the LYAF, CIV absorbs 50% of all the photons removed by metals, but this drops to only 20% inside the LYAF. Many of the metals lines that we fail to find inside the LYAF are individual CIV doublets that have no other lines to indicate their position. Combining this information, we estimate that metal lines (both identified and missed) absorb 1.9±0.3% of the flux inside the LYAF. This is considerably more accurate than our prior estimates (that others have been using) derived from lines outside the LYAF in lower resolution spectra.

The number of metal lines that we miss in the LYAF is a complex function of spectral resolution, SNR, the observed wavelength range, the precise absorption redshift, the number of lines identified in the redshift system and the complexity of the velocity structure. The bulk of the metal line absorption comes from the few redshift systems that have the largest column densities. These systems are easy to find in low-resolution spectra of wavelengths outside the LYAF. Once we know the velocity range, we can easily mask the expected wavelength of all their main lines.

We have checked the number of metals we can find in the LYAF by adding known lines into the LYAF of a real spectrum. We found that we could recover most of these lines (e.g. 17 of 20 CIV, 25/29 MgII, 26/34 Fe II, 9/10 Al III). We will use these fractions to tell us how many metal lines to add into simulated spectra. This part of our project (and some of that on metal lines) has been lead by Cameron Liang who is a first year undergraduate and minority student at UCSD. This was his first research project.

Lya lines with high HI columns. We have found and removed all the Lya lines due to absorbers with HI column densities log N(HI) > 17 (cm^{-2}) since such lines lie in regions with densities above those that we can accurately simulate in boxes large enough to correctly sample the LYAF. In Tytler et al. 2009b Fig 5 we show that our large simulations have about the correct number of absorbers with columns 13 – 17, but they systematically lack systems with log N(HI) > 17 cm^{-2}. We remove these absorbers by dividing the spectrum by the fit to the Lya line profile, and removing wavelengths where the fit absorbs > 50% of the flux. The division is necessary to deal with damping wings, while the cut is needed to avoid excessive noise from the cores of lines. We use (with great success) the same method for metal lines.

Remove the absorption from regions changed by galaxies. Simulations have shown that winds from galaxies have a very small effect on the LYAF statistics (Theuns et al. 02, Cen et al. 05). Observations also show that the IGM is quiescent, with very similar velocities in adjacent lines of sight. Yet we know that about half of absorbers with log N(HI) > 14.5 cm^{-2} show CIV absorption at similar velocities. We found minimal changes to LYAF statistics when we removed all lines with log N(HI)> 14.5 cm^{-2}, because these regions occupy only a tiny fraction of each spectrum.

We briefly mention **other work**. Misawa et al. 04, 07a (his PhD was jointly supervised by DT) are the first detailed studies of the properties of the H I in Lyman limits from 40 HIRES spectra. We found that Lya lines near to QSOs are wider by about 2 km/s, as if they are from hotter or faster moving gas. Misawa et al. 07b is a 172 page survey for intrinsic metal lines absorption in HIRES spectra of 37 QSOs. We find that at least 17% of C IV systems at 5000 – 70,000 km/s from the QSOs are intrinsic because they only partly cover the QSO continuum source. At least 35% of QSOs have one or more such systems. Janknecht et al. 06a explored the low redshift Lya forest. Reimers et al. 05, 06 use HIRES and HST spectra to measure the ionization of the IGM. Agafonova et al. 07 find that the UVB ionizing low column density metal lines systems is very hard and not of stellar origin. Tytler et al. 04a present Lick spectra of the forest that we use in Tytler et al. 04b to measure parameters that fit the Lya lines. Kirkman et al. 05

made the first calibrated measurement of the amount of absorption in the IGM using 27 HIRES spectra and Jena et al. 05 match the properties of the LYAF in supercomputer simulations and real spectra.

Student training. Our NSF activities brought many new people to astronomy research. Tytler advised 19 undergraduates in the last 4 years, including 8 women. These undergraduates worked significant hours and became regular members of our research group, with desks in an adjacent laboratory. They all worked on our main research topics, on projects that we will publish. The first research paper is an exciting milestone for all new researchers. Tytler and other team members work closely with the undergraduates, introducing them to every part of our research program including target selection, observing, data reduction, analysis, statistics, interpretation and papers. Most of our former undergraduate researchers have entered PhD programs. In 2006 Carl Melis entered UCLA (coauthor on 2 papers from our group, he completed his PhD after 3 yrs in 2009), in 2007 Sam Bockenhauer (on 1 paper) entered Stanford and Melissa Downey went to UCLA. Others from earlier years are also co-authors on papers that came out recently: Mark Gleed, Angela Chapman and Susan Hollywood each on one paper and Adam Orin who is of Latin ethnicity is on two papers. Nine graduate students have worked on this research: Tridi Jena, John O'Meara, Toru Misawa, Nao Suzuki, Pascal Paschos, Stephen Skory, Xiheng Shi, Aaron Day and Liliana Lopez.

3. LYAF Spectra and Simulations

We describe the simulations and spectra that we have or will obtain and their preparation.

3.1. QSO Spectra

We will make extensive use of high-resolution spectra with FWHM 8 km/s from Keck+HIRES. The high-resolution spectra are the only ones to sample the small-scales that we need to break degeneracy between the small-scale mass power and the IGM temperature distribution. We are part way through obtaining HIRES spectra with 8 km/s FWHM of 300 QSOs for this program. These complement the HIRES spectra of 120 QSOs that we obtained during our search for QSOs that show deuterium. About 35 have SNR > 30 per 2.1 km/s pixel, 10 have SNR > 50. There are >12 QSOs contributing LYAF data at all $2.3 < z < 3.5$, with a peak of 35 at $z=2.8$. Some of them have already been used for other purposes (Crotts et al. 97; Burles et al. 98c; Fang et al. 98; Weinberg et al. 99; Pando et al. 02; Croft et al. 02; Prochaska et al. 01, 02ab, Telfer et al. 02; Jamkhedkar et al. 03; Misawa et al. 04 07a,b). We easily remove bias in selecting for D since we know which redshift system caused us to look at each QSO, and these all have HI column densities that we remove from all spectra. We will also use publicly available high-resolution spectra from the VLT UVES if we can apply adequate relative flux calibration.

We have intermediate-resolution spectra (FWHM 55 km/s) from the Keck ESI spectrograph of 100 QSOs mostly at $z > 3$. We will use ESI (or Magellan MagE in the south) to observe higher z QSOs that are too faint for HIRES. We have low-resolution spectra (FWHM 250 km/s) from the Lick Kast spectrograph of 200 QSOs at $1.6 < z < 2.4$. These all cover the LYAF to $z = 1.6$, below the SDSS limit of 2.0 – 2.2. Tytler et al. 04a published 77 of these. We also have Keck LRIS spectra (86 km/s FWHM) of pairs of QSOs with 0.5 – 3 arc-minutes separation at $z = 2$.

We will use SDSS spectra. DR7 contains spectra of 10,317 QSOs with $z > 2.55$ that completely sample rest wavelengths 1070-1170 that we use for LYAF statistics, and another 2072 with $z > 2.35$ that partly cover this region. For comparison, McDonald et al 05 used 3100 SDSS spectra from DR3. The SDSS III BOSS survey is now obtaining spectra of the LYAF of about 150,000 QSOs, at a resolution of 132 – 192 km/s, slightly better than SDSS I and extending to 3600 A. BOSS spectra will be released as follows: DR9 in July 2012 (spectra obtained before July 2011), DR10 in July 2013 (spectra before July 2012) and DR12 Dec 2014 the final release of 150,000 QSO spectra.

3.2. Preparation and Measurement of High-resolution QSO Spectra

We will calibrate and clean the real spectra to try to remove problems that we cannot simulate, and we will add to the simulated spectra representations of problems with the real data that we cannot remove.

These calibrations are the key to precision measurements with the LYAF and they will take most of our time. We have years of experience from our D/H work. Here we give an overview of these techniques; for detailed discussions see Tytler et al. 04ab, 09ab, Kirkman et al. 97, 05, 08 and Suzuki 03, 05.

Relative flux calibration. We have a sophisticated system that allows us to apply relative flux calibration to HIRES spectra (Suzuki et al. 2003). We have measured 1-2% accuracy in the relative flux over 200 Angstroms (several spectral orders) and a factor of up to ten times better than in some of our earlier efforts using flux standard stars. Accurate flux calibration allows us to get the accurate continuum levels needed for accurate measures of the mean amount of absorption and to measure flux power over much larger scales than is possible with typical spectra. Without this calibration, the flux power measurement is limited to about ¼ of a HIRES order, about 15 A, or 1000 km/s. We have applied this calibration to about 60 HIRES spectra and we need to complete this work for the rest. We will need to obtain Lick Kast spectra of about 200 QSOs for these calibrations. We also need to quantify the distribution of the relative flux errors (amplitude, correlation length etc), to better understand how they will change the flux power.

Spectral resolution. Our flux calibration process involves degrading high-resolution spectra with well-known resolution to the resolution of other spectra with reliable flux calibrations. In the process, we get accurate measurements of the spectral resolution in the lower resolution (e.g. Lick or SDSS) spectra. We need to know this for statistics that depend on the resolution, such as the power spectrum, and the fraction of pixels with a given flux (**flux pdf**) where the flux is from 0 to 1 at the continuum level.

Combining spectra and measuring the SNR. Combining spectra is tricky when the relative flux calibration is in error, since they then differ because one or both are in error. We need better algorithms.

Improved continuum fits. The accuracy of the continuum is a major source of error in the measurement of the mean amount of absorption and the flux pdf, especially when automated methods are used that cannot identify emission line shapes in low SNR spectra. We have shown (Suzuki et al. 05) that continua derived from principle component spectra make excellent emission lines shapes but fail in the overall flux level in the LYAF. We intend to try to fix this problem by adding parameters that will allow for errors in the flux calibration (extinction that night) and the temperature structure of the accretion disk that gives the continuum in the LYAF. We will also improve the method by changing the set of spectra that we use to find the principal components. We have used 50 QSO spectra from HST FOS that are at low z and hence have few absorption lines that we removed by hand. We will instead use HIRES spectra of QSOs at z=2. Their continua are much better known because of the higher SNR (per Angstrom) and resolution. They will also have overall UV spectral shapes that are more representative of all the QSOs we use in this project. We can further improve the continuum accuracy by using only those QSOs with the simplest continua, with the least emission lines in the LYAF (Suzuki 06). We will check our continuum fitting on a sample of 100 QSOs at z = 2 (few Lya lines) that we will observe at high SNR. As with the flux calibration, we will quantify the continuum level errors, their amplitude and correlation length, for QSOs with various shapes of emission lines, all as a function of SNR, z and resolution.

Finding and masking metals and strong Lya lines. We already explained how we find and remove metal lines in the LYAF and the high column density Lya lines. Tytler et al. (2004b) found that one spectrum with metals and the Lya of LLS removed has the same amount of information on the LYAF (on large scales of 150 Mpc or Δz=0.1 at z=2) as 9 with the mean amount removed. Kim et al. (2004a) showed that metals are about 50% of the signal in the flux power on small scales. The value depends on the SNR since photon noise is important on these small-scales. In our own work with 30 HIRES spectra, we find that the metal lines increase the power in the LYAF to 5-10% of the power from the LYAF (excluding metals) over large-scales (k=0.006 – 0.06 s/km). This means that we need to measure the power from the metals with an error of about 5 – 10% if we want to know the power of the LYAF without the metals to 1%. This will be relatively easy; we expect an error of <10% from only about 50 QSO spectra. As the wavenumber increases from k = 0.1 to 0.2 s/km the power from the metal lines rises steeply to 10 times more than that from the LYAF alone (the HIRES resolution corresponds to k=0.4 s/km), hence we do not expect much useful information from k > 0.1 s/km (off the right edge of Fig 1).

Mean amount of Lya absorption. The mean amount of Lya absorption helps us distinguish IGM models. Tytler et al. 04ab estimated that about 5 high-resolution spectra (free of metals and Lya with log N(HI) > 17 cm^{-2}) are enough to give a random error on the mean amount of absorption of 0.01 at z=1.95. Hence, flux calibration and continuum level errors, and not the number of spectra, are the dominant issues for higher accuracy measurement of the mean amount of Lya absorption. We will automatically obtain better accuracy on the mean amount of absorption once we have improved the flux calibration, the continuum fits and the removal of metal lines and high N(HI) lines.

Set of LYAF statistics. In addition to the mean absorption (or 1- mean flux), the flux pdf and the flux power, we also measure the Lya column density N distribution, the b-value distribution and the b-N correlation (Kirkman & Tytler 97), the flux pdf (McDonald et al. 00), wavelet coefficients (Meiksin 00; Theuns & Zaroubi 00; Pando et al. 02) and a void statistic modified from Duncan, Ostriker & Bajtlik 89.

3.3. Preparation and Measurement of Low-resolution Spectra

Here we explain how we will prepare and use low-resolution spectra from both Lick and SDSS DR7 and other samples (BOSS) as they become available. Since the SDSS spectra are numerous, they give small random errors, but they are more prone to systematic errors because they have lower resolution.

Flux calibration. Any systematic error in the flux calibration of the low-resolution spectra will appear as an error in the mean amount of absorption at that wavelength. It is less likely to appear in the continuum level because a given observed wavelength is sampled by all the different rest wavelengths in the LYAF when we include QSOs with a full range of emission redshifts. One way to check for such errors is to examine the mean flux at the same observed wavelengths in QSOs with lower emission redshifts that do not cover the LYAF. We will also check flux calibration comparing spectra from Lick and SDSS.

Continuum level. We will use the methods we described for high-resolution spectra to apply the continuum to low-resolution spectra. We will explicitly check the accuracy of the continua on the low-resolution spectra by comparison with high-resolution spectra of the same QSOs. We will use these results to improve the algorithm that we use to find the continua on the low-resolution spectra. At worst we anticipate about 10% errors at any one wavelength in the continua on individual low-resolution spectra, but these will be reduced (leaving mostly systematic offsets) using the new PCA templates that we will calculate from HIRES spectra. We will look for systematic errors in the continua on the low-resolution spectra by examining whether the mean amount of absorption depends on the width or strength of the emission lines of the QSOs, the mean slope of the spectrum and the SNR of the spectrum.

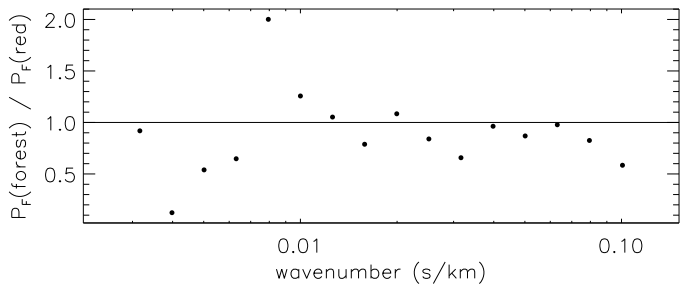

FIG 5 Two ways of estimating the flux power caused by metal lines. From30 HIRES spectra we measure P(forest) = P(Lya + metals) – P(Lya only), while P(red) is power of metal lines outside LYAF. Values near 1 show agreement. The points fluctuate most on large scales (left) where the continuum is less accurate and we see shot noise from the few metals systems. We use 0.01 < k < 0.1 (s/km).

Metal lines. We have a plan to use the thousands of SDSS spectra to give much better random errors on the mean amount of absorption by metal lines at wavelengths just larger than those in the LYAF. This is the next step in our long-term program to improve the accuracy of the measurement of the mean absorption due to the Lya lines alone. With SDSS spectra, we will have small enough random errors to be able to measure the mean amount of absorption due to specific lines that have rest wavelengths larger than Lya (e.g. SiIV doublet 1402, 1393, OI 1302, Si II 1260). We will also calculate how the absorption by these lines changes with redshift. This allows us to predict the absorption by these lines (and other lines of the same ions) inside the LYAF. We still need the high-resolution spectra to tell us the amount of

absorption by lines of ions that appear only in the LYAF (about 20% of the total metal absorption in the LYAF). Frank et al. 2010ab show that some strong metal lines inside the LYAF can be found in SDSS spectra with high SNR, but too few lines are found (about one OVI per QSO) for our purpose.

We will calculate the power of the metal lines to the red of Lya. McDonald et al. 05a (Eqn 4) used this power (from 1268-1380 A) as a surrogate for the power of the metal lines inside the LYAF. They took the P(Lya lines alone) ≈ P(all absorption in the LYAF) – P(metals outside the LYAF), and they used the autocorrelation of the flux in the LYAF to estimate the effect of the Si III 1206 absorption (Eqn 14,15). We will check the accuracy of this method by comparing it to the results we will obtain from high-resolution spectra where we can directly measure the power of the LYAF excluding wavelengths that contain metal absorption. In Fig. 5, we show an example of this using 320k pixels from 30 HIRES spectra. The power of the metal lines to the red of Lya emission is indeed close to the power that (different) metal lines add inside the LYAF. The two will not be identical because there are metal lines with rest wavelengths < 1216 A that we see only inside the LYAF. If these two estimates differ in larger samples, then the SDSS spectra may have to rely on HIRES estimates of the power from metals.

Lya lines with HI columns log N(HI) > 17 cm^{-2}. In low-resolution spectra, we can find all the DLAs with very high log N(HI) > 20.3 cm^{-2}, but we will miss most with columns 17—19 cm^{-2}. McDonald et al. 05b discuss the effects of DLAs alone. We will add high column lines, with realistic clustered components into the simulated spectra. The number of lines we add will change with the SNR, redshifts, wavelengths and resolution of the data sample. To do this we must first measure their frequency in HIRES spectra.

3.4. Simulations

We have been using ENZO simulations for a decade now. The generic simulation is fully hydrodynamic, assumes an evolving uniform UVB, and calculates non-equilibrium ionization, heating and cooling (Zhang et al. 1997, 1998; Anninos et al. 1997, Bryan et al. 1999, Machacek et al. 2000). Each fluid element in the simulation experiences exactly the same intensity and spectrum of ionizing radiation, which changes with redshift consistent with evolving populations of galaxies and quasars (Haardt and Madau 2001). There is no mechanical or chemical feedback. Simulations in the ENZO code use uniform Eulerian grids as large as 2048^3. The heating is entirely from the UVB that we parameterize with two numbers: the ionization rate for H γ_{912}, and the heat input per He II ionization X_{228}. We measure both in units of the amount expected (from Haardt & Madau 01). The heating is self-consistent for an optically thin universe. The lack of optical depth means that there is less absorption of photons that can ionize He II at z=2-4 than in the real universe. This leads to absorption lines that are slightly too narrow. We increase X_{228} about 10% to boost the heating rate of He II to compensate. Compared to a simulation with optical depth and RT, this will place too much heat in the low density IGM and not enough in high-density regions that are optically thick to He II ionizing photons. We used a suite of 40 of these simulations in Tytler et al. 2004b and Jena et al. 2005. We now have more than 144 with a wider range of input parameters and resolution. We have sufficient supercomputer time to run hundreds of these simulations. GADGET is an example of a different type of cosmological code that uses smoothed particle hydrodynamics (Springel V., Yoshida N., White S. D. M., 2001, New Astronomy, 6, 51). It is a Lagrangian code, unlike the Eulerian Enzo, and both codes give very similar results for the LYAF.

3.5. How Simulations Estimate the Effect of Neutrino Mass

There are several ways to include neutrino mass in simulations. Masses of 0.05 – 1eV correspond to free streaming lengths of 200 – 30 Mpc and thermal velocities of 3000 – 450 km/s, large enough to prevent neutrinos from clustering with CDM and baryons and leaving them in the linear regime. Hence, Agarwal & Feldman 10 ignore non-linear neutrino perturbations and they include the neutrinos only in the initial conditions to their numerical simulations. ENZO can already include neutrino mass in its initial conditions. A second method is to add a linear gird for the neutrinos, which is simple and runs quickly but does not follow non-linear effects. A third method that does allow for non-linear corrections is by treating

neutrinos as a new dark matter field with different velocity dispersion. Both methods 2 and 3 were implemented by Brandbyge & Hannestad 09 and Viel et al. 10. We expect that we will need to implement the particle method or perhaps one of the hybrid methods discussed by Brandbyge & Hannestad 10.

3.6. Realistic Simulated Spectra.

We will make simulated spectra that are visually almost indistinguishable from the real ones. Starting with a LYAF spectrum, we add emission lines obtained using principal component spectra described in Suzuki et al. (2006). We add metal absorption lines (representing those that we failed to remove from spectra), photon noise, flux calibration errors and non-Poisson pixel errors, such as poorly removed cosmic rays. We make a set of many simulated spectra to match each real spectrum. We give the simulated spectra SNR distributions like those of the real spectra. This is complicated because the SNR varies systematically with wavelength, and across each spectral order and in and out of spectral absorption lines. The noise level in the core of an absorption line depends on the sky brightness that also changes with wavelength. We have not yet implemented all details that we will need to work with the flux pdf and power of low SNR spectra. We bin the simulated spectra to the pixel size of the real spectra, and we fit continua to the simulated spectra when required (Tytler et al. 2004b).

4. How we Measure the Sum of the Neutrino Masses

We want to know the amplitude of the power spectrum of the matter in the IGM (**the matter power**) because this is sensitive to the neutrino masses. We do not observe the matter power, but we can measure various statistics of the flux in the LYAF, including the power spectrum of the flux (**the flux power**). All the LYAF statistics respond to the matter power in some way, hence we look for simulations that match all statistics simultaneously. Here we will discuss only two statistics, the flux power and the flux pdf.

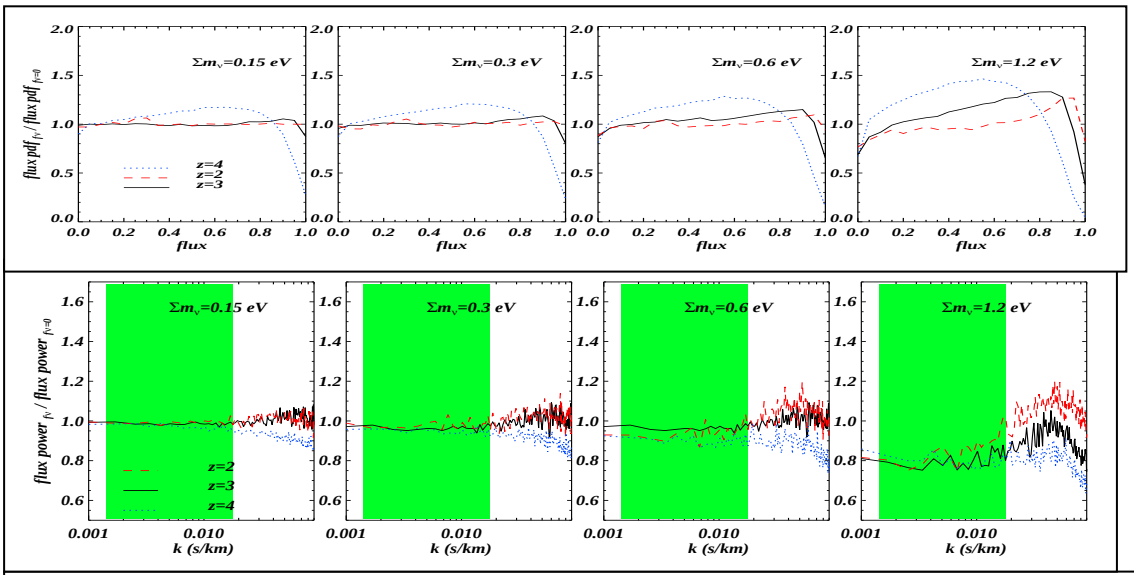

FIG 6. There are 4 panels in each row. Sum of the masses of the light neutrinos increases from 0.15 eV(left), 0.3, 0.6 to 1.2eV (right). Each panel shows z=4 (blue dots), 3 (red dash) and 2 (black). Top row show the LYAF flux pdf divided by the same without neutrinos. The smallest changes are at flux = 0 (left) and the largest on the right (flux = 1) at z=4. Bottom row shows flux power with neutrinos divided by the same without neutrinos. The green region shows the wavenumbers from SDSS spectra. With HIRES we cover from <0.01 to 0.1 just off the right side of each panel. Figs 13 & 15 from Viel et al. 2010 JCAP 10 1475

The matter and flux power are distinctly different and do not have a simple one-to-one correspondence because (Doppler) velocity changes can smear out fluctuations in spectra. Fig 6 from Viel et al. 10 shows how two statistics of the LYAF change with neutrino mass. The top row of panels show

how the flux pdf (fraction of pixels with a given flux, measured from 0 to 1) changes with neutrino mass increasing from 0.15 eV (top left), 0.3 eV, 0.6 eV to 1.2 eV (top right). The changes are many times larger at z=4 (faint dotted blue curve) than at z=3 (grey solid line) or 2 (red dashed), showing that higher z spectra contain much more information. For this reason we concentrate on high z, and we will obtain new ESI or LRIS spectra of QSOs that are at high z but too faint for HIRES. The changes are also largest at fluxes near to one, which means that spectra with higher SNR will be valuable. The changes decline slowly with neutrino mass, such that the change is less than a factor of two (at z=4) as the neutrino mass is halved. This will help us detect the neutrino mass.

The second row of panels shows the changes in the flux power (with neutrino mass, divided by without mass) for the same neutrino masses 0.16, 0.3, 0.6 and 1.2 eV. Again the changes are much larger at z = 4 than at 3 or 2. The green vertical bands show the wavenumbers resolved SDSS DR3. Except for the highest masses the changes are minimal. This figure shows a key advantage of high-resolution spectra, where we readily sample $0.01 < k < 0.1$ (s/km), all of the plots to right of the green bands. By far the strongest signal is at these small scales, and at high z. Adjustments to σ_8 can hide > ½ of the changes produced by the neutrino mass, hence we will use the most accurate σ_8 values from other measurements.

4.1. Errors and Sample Size

We used the power we have already measured (in 30 QSOs) to estimate that LYAF spectra of 167 QSOs will give the LYAF flux power (excluding metal lines and high N(HI) Lya) with a random error of 2% in each of eleven k values from $0.01 < k < 0.1$ (s/km). This is assuming that all the QSOs have the same emission redshift (near 3.0) and we obtain a redshift path of 0.4 from each. From the lower left panel of Fig 5, the change in the flux power is about 2% at z=2, 4% at z=3 and 10% at z=4. Hence, a 2% change in flux power at k = 0.07 (s/km) should give small enough random errors to detect a mass of 0.16 eV at z = 3 or 4, and it might be enough at even z = 2. For any QSO spectrum, we also obtain other statistics. From the same sample of 167 HIRES spectra, we will know the flux pdf with an error of 1.5% at intermediate fluxes 0.2 – 0.8. We intend to have HIRES spectra of >300 QSOs for this project.

The lower limit that we can set on the sum of the masses of the light neutrinos will be set by systematic errors. Our work with 30 HIRES spectra has shown that we need only of order 167 HIRES spectra to detect neutrino mass down to 0.16 eV, provided the systematic errors are held less than the random errors. Hence, the key task for us is to keep the systematic errors small. This is exactly where the high-resolution spectra have major advantages (compared to low resolution spectra) and this is why we have discussed the various calibration issues in detail above.

Viel et al. 04a give a table summarizing the errors in their measurement of amplitude of the mass power from the Croft et al. 02 and Kim et al. 04 samples: 8% is from the mean absorption, 5% from the temperature density relation, 5% from the method and 8% from differences between simulations. We avoid the last three by using fully hydrodynamic simulations and ensuring that these match a full set of statistics including the line widths. The current mean absorption errors are ½ those used by Viel et al. 04a and the work discussed above will give another factor of 2, leaving a 2% error, similar in size to the random error 167 QSO spectra. As we discussed on p6, the error in the flux power from missed metal lines from 167 QSOs will be smaller than <0.5%. The mean amount of absorption, and hence the flux calibration and the continuum will be the main errors from this work.

4.2. Why we will obtain an Improved Limit on the Neutrino Masses.

The best measurements of the matter power from the LYAF include Croft et al. 02b using 30 of our (uncalibrated) HIRES spectra, Kim et al. 04a,b using 27 high-resolution VLT+UVES spectra, Viel et al. 08 using 55 HIRES spectra. The main result from the SDSS QSO spectra is McDonald et al. 05a who used 3100 SDSS spectra plus 7 HIRES spectra. They needed high-resolution to help specify the IGM model, but they used only 7 HIRES spectra because of calibration incompatibilities with the other samples. While the SDSS has more spectra, giving smaller random errors, these spectra do not resolve Lya lines and they alone do not contain the information needed to measure the power of the matter.

Systematic errors, and not random errors, dominate all measurement of the matter power from the LYAF. The best way to reduce systematic errors is to use carefully calibrated high-resolution spectra. This is the main reason why we will obtain improved limits on the matter power on the small-scales sampled by the LYAF, and hence better limits on the neutrino masses.

The dominant errors in all power measurements (and neutrino mass limits) from the LYAF are systematic errors associated with the uncertainty over the IGM model. Our high-resolution spectra reveal 20 times smaller-scales (than SDSS spectra) and fully resolve the Lya lines. We can then find the physical description of the IGM that matches the line shapes. Without the accurate model of the IGM that we specify with high resolution spectra, neutrino mass limits estimates from both high and low resolution LYAF spectra will be limited by systematic errors.

For example, we measure and use the Lya line widths (**b-values**). Nearly all past work to estimate the matter power (e.g. Kim et al. 04ab, and the SDSS work by McDonald et al. 05) fitted the power alone and not the full set of LYAF statistics. "Nuisance parameters" (such as the slope of an idealized power law temperature – density correlation) were introduced to try to model the (unknown) possible states for the IGM. These parameters were allowed to adjust to best match the flux power. This gave an IGM that was clearly too hot, hotter than allowed by the b-values (Viel & Haehnelt 06, Bolton et al. 08, Peeples et al. 10a). The published measurements of the small-scale power and the neutrino mass limits from the LYAF do not use realistic models for the IGM and their values should not be trusted.

In addition to line widths, we use a full set of all the common LYAF statistics, all calculated from the same set of spectra. We avoid the problems that come with using a bag of statistics each from a different sample. We have been biasing our observing towards high redshift QSOs because there is much more information on the neutrino mass at z = 4 than at 3. In total, we will have 5 – 10 times more high-resolution spectra than past samples. High-resolution spectra are also the key to neutrino mass because Viel et al. 10 (Fig 6) show that almost all the information on the neutrino mass (in the LYAF flux power) is at small-scales that are sampled by HIRES but not by the SDSS resolution. In this sense, our 300 high-resolution spectra will contain more information on neutrino mass that thousands of SDSS spectra, but clearly the high and low-resolution data sets are most powerful when used together.

5. Work on Paired QSOs to check the IGM model

This work checks that the simulations are a realistic physical model for the IGM. There are three separate ways we use the spectra of pairs, as in prior work. The observing is efficient because we can use many of the same QSO spectra for more than one part and we are able to observe many pairs from Lick.

First, we will obtain high-resolution spectra of the LYAF in several pairs of QSOs (we have one already) separated by about 0.5 Mpc, wider than those studied by Rauch 02. The correlations between these sight lines (while far from either QSO) provides a critical test of the physical conditions in simulations that match individual spectra. This 3D information on the IGM can resolve ambiguities that arise when we examine only single sight lines. For example, the cross-correlations are much more sensitive to the size and shape of the absorbing regions than are single spectra. In single spectra, we cannot distinguish between many small absorbing regions and a few large ones, both of which make the same number of absorption lines. Peeples et al. 10b demonstrate that line of sight broadening of Lya lines comes from the gas temperature, while the coherence in the sky is from pressure support.

Second, we examine how the UV from individual QSOs changes the Lya and metal absorption lines that we see near to the QSOs in the plane of the sky. We will observe tens of pairs of QSOs for this. We will examine the ionization of those absorbers to check our earlier finding that the UV from QSOs might be anisotropic or short lived, changes that would alter the reionization of He II that is ongoing at z=3.

Third, we explore winds from galaxies by examining how metal absorption systems (far from any QSOs) differ in QSO sight lines separated by <100 kpc. In Tytler et al. 09b, we discovered that these absorbers show the same velocity in both lines of sight, as if they arise in a cold population with minimal pair-wise velocity dispersion (difference in velocity, after subtracting the mean Hubble flow and infall). This indicates absorption in blue rather than red galaxies. The bulk of the absorption that we see cannot

come from winds flowing rapidly out of galaxies. This suggests that most metals were in place long before z=2. Recently Martin et al. 10 have confirmed this result. We will extend this work to much weaker lines that are more likely to arise farther from galaxies. Our goal here is to use simulations and spectra to improve our determination of the fraction of the volume that has been affected by winds.

6. Broader Impact: Training and Development

We will continue our established and successful program (described under prior work) of integrating undergraduate students into our main NSF funded research. Some students come from the general introductory courses that Tytler teaches each year, while others have heard by word of mouth. Most years we have about four new students. The budget below includes money for one undergraduate per year.

As the diversity officer for the UCSD physics dept., Tytler is trying to increase the enrollment of underrepresented minorities in our PhD program. In the fall of 2008, he contacted physics students and department heads at California State Universities and elsewhere. We now have some valuable connections and our minority enrollment increased from zero (2008) to 2 out of 38 (2009), and 1 in 2010, a small start. We are especially interested in continuing to supporting Liliana Lopez on this grant. She contacted Tytler in 2008 as she was completing her MSc in astronomy as SFSU. UCSD admitted her in the fall of 2009 and she began working with us in April 2010. On a different tack, Tytler attended a Math for America meeting in San Diego in October 2010, and will be helping interview teacher candidates for the (NSF sponsored) Math for America program.

We have demonstrated in the above report on prior work that we are routinely integrating undergraduates, including many women and a few minorities into our NSF sponsored research. We now have years of experience and tools that make this activity an outstanding success from both our perspective and that of the students we train.

We use GUI driven software that allows all students, even those with no prior computer experience, to do publication quality work after about one week of training. The software allows students to see, manipulate and extract physical parameters from QSO spectra. The tool is ideal for introducing the difference between real features, photon noise and systematic errors. It is also a teaching tool, helps them learn about the interstellar medium, the photoionization of gas, element abundances, turbulent motions and thermal broadening, all the main topics in astronomical spectroscopy.

Tytler has been the UC champion for remote observing since the beginning of the Keck Observatory in 1994. He ensured that UCSD was the first campus in the UC system after the UCSC prototype to have a dedicated remote observing room. Kirkman and Tytler acted as beta-testers. We now use this regularly with great success to train our graduates and a few selected undergraduates in observing. Those in our team with classes the next day can stay until midnight while those in the general astronomy introductory classes that Tytler teaches come to watch for an hour. They see how we obtain the spectra we use for this research from Keck. Starting in 2007, we became the first UC campus to do regular remote observing at Lick. This is a major advance because we have 60 nights per year at Lick giving many more opportunities to see observing. Remote observing brings the excitement of observing to many more students. In November 2008 we installed new large screen monitors and a video link display (Polycom HD system) for Keck and Lick remote observing. These take advantage of our large room that can hold 10 students.

In Jan 2009, Tytler made the first UC use of remote observing with the Lick 40-inch as a part of the general introductory course (for non-science majors) that he teaches. Each student was asked to choose a target to observe (asteroid, comet, planetary nebula, galaxy, one brought the star named after her), and they were asked to come to one of 12 evening sessions with the RA, Dec, likely hour angle and description of their target. This was the most fun we have had in years, giving one-on-one contact, showing the students how we observe. Some students ask if they can come to all the observing sessions. We have now done this for three classes and will continue this during the course of this proposal. In 2010 we hosted Prof Michael Burin from California State University San Marcos for a night of observing of galaxies. His class of 12 students came in 2010 with target coordinates and they left with images of galaxies from which they measured light distributions. We will do this again each year of this proposal.