



The Impact of Cosmic Variance on Inferences of Global Neutral Fraction Derived from Ly α Luminosity Functions during Reionization

Sean Bruton¹ , Claudia Scarlata¹ , Francesco Haardt^{2,3} , Matthew J. Hayes⁴ , Charlotte Mason^{5,6} ,
Alexa M. Morales⁷ , and Andrei Mesinger⁸

¹ Minnesota Institute for Astrophysics, University of Minnesota, 116 Church St SE, Minneapolis, MN 55455, USA; bruto012@umn.edu

² National Institute of Nuclear Physics INFN, Milano—Bicocca, Piazza della Scienza 3, I-20126 Milano, Italy

³ DiSAT, Università dell’Insubria, via Valleggio 11, I-22100 Como, Italy

⁴ Department of Astronomy and Oskar Klein Centre for Cosmoparticle Physics, Stockholm University, AlbaNova University Centre, SE-10691, Stockholm, Sweden

⁵ Cosmic Dawn Center (DAWN), Denmark

⁶ Niels Bohr Institute, University of Copenhagen, Jagtvej 128, DK-2200 København N, Denmark

⁷ Department of Astronomy, The University of Texas at Austin, 2515 Speedway, Stop C1400, Austin, TX 78712, USA

⁸ Scuola Normale Superiore, Piazza dei Cavalieri 7, I-56126 Pisa, Italy

Received 2022 August 15; revised 2023 April 24; accepted 2023 April 28; published 2023 July 31

Abstract

We investigate the impact of field-to-field variation, deriving from cosmic variance, in measured Ly α emitter (LAE) luminosity functions (LFs) and this variation’s impact on inferences of the neutral fraction of the intergalactic medium (IGM) during reionization. We post-process a $z = 7$ IGM simulation to populate the dark matter halos with LAEs. These LAEs have realistic UV magnitudes, Ly α fluxes, and Ly α line profiles. We calculate the attenuation of Ly α emission in universes with varying IGM neutral fraction, \bar{x}_{HI} . In a $\bar{x}_{\text{HI}} = 0.3$ simulation, we perform 100 realizations of a mock 2 deg^2 survey with a redshift window $\Delta z = 0.5$ and flux limit $f_{\text{Ly}\alpha} > 1 \times 10^{-17} \text{ erg s}^{-1} \text{ cm}^{-2}$; such a survey is typical in depth and volume of the largest LAE surveys conducted today. For each realization, we compute the LAE LF and use it to recover the input \bar{x}_{HI} . Comparing the inferred values of \bar{x}_{HI} across the ensemble of the surveys, we find that cosmic variance, deriving from large-scale structure and variation in the neutral gas along the sightline, imposes a floor in the uncertainty of $\Delta \bar{x}_{\text{HI}} \sim 0.2$ when $\bar{x}_{\text{HI}} = 0.3$. We explore mitigation strategies to decrease this uncertainty, such as increasing the volume, decreasing the flux limit, or probing the volume with many independent fields. Increasing the area and/or depth of the survey does not mitigate the uncertainty, but composing a survey with many independent fields is effective. This finding highlights the best strategy for LAE surveys aiming at constraining the \bar{x}_{HI} of the universe during reionization.

Unified Astronomy Thesaurus concepts: [Reionization \(1383\)](#); [Lyman-alpha galaxies \(978\)](#); [Intergalactic medium \(813\)](#); [High-redshift galaxies \(734\)](#); [Early universe \(435\)](#)

1. Introduction

Reionization is the most recent phase change of the universe, in which the hydrogen in the intergalactic medium (IGM) went from being virtually fully neutral to nearly completely ionized. It is thought to have been complete within about the first billion years of the universe. It is expected to have been a patchy process, as ionized bubbles formed and grew around the first sources of ionizing radiation in the universe. Generally, it is thought that stars in the first galaxies provided the bulk of the necessary ionizing radiation to drive reionization (Bouwens et al. 2015a; Finkelstein et al. 2015; Robertson et al. 2015), though which galaxies dominate the ionizing photon budget is still being investigated.

Over the past several years, observations have been made which put constraints on the timeline of reionization using complementary methods sensitive to ionized and neutral hydrogen. Planck-Collaboration et al. (2020) use the scattering signature of electrons freed during reionization on the cosmic microwave background to constrain the midpoint on instantaneous reionization to redshift $z_{\text{re}} = 7.67 \pm 0.73$. The fraction of dark pixels in the Ly α and Ly β forests constrains

reionization to be very nearly complete at $z \sim 6$ (Becker et al. 2021; Qin et al. 2021; Bosman et al. 2022). Quasar spectra at $z > 7$ offer insight into the global neutral fraction (\bar{x}_{HI}) of the IGM throughout reionization based on the IGM’s Ly α damping wing imprints (Davies et al. 2018; Greig et al. 2019).

Galaxies, and in particular Ly α emitters (LAEs), are also used to constrain the IGM hydrogen neutral fraction and its evolution with time. Ly α is a resonant transition in hydrogen, and Ly α photons scatter multiple times even in small amounts of neutral gas (Dijkstra 2014, 2017). Thus, observations of Ly α are particularly useful to probe the HI spatial distribution, column density, and velocity fields. Observations of Ly α freely propagating through the IGM is indicative of a largely ionized IGM or a significantly redshifted Ly α line.

The LAE fraction, the fraction of galaxies which are found to emit Ly α , has been found to rapidly decline at $z > 6$ (Stark et al. 2010; Pentericci et al. 2011; Jung et al. 2018), and this decline is interpreted as a sign of increasing neutral hydrogen fraction in the IGM (Treu et al. 2013; Mesinger et al. 2015; Mason et al. 2018b). There are complications, however, in that the observed decreasing LAE fraction could also result from evolving galaxy properties, such as the escape fraction of Ly α photons (Dijkstra 2014), an increase in the number of Lyman limit systems (Bolton & Haehnelt 2013), or other considerations (see Finkelstein 2016 for a review). Still, Mesinger et al. (2015) and Mason et al. (2018b) argue that the increasing neutral fraction of the IGM is the most plausible explanation,

and in this context the rapid disappearance of Ly α at $z > 6$ paints a picture of late and rapid reionization (Ouchi et al. 2018; Hoag 2019; Mason et al. 2019; Yoshioka et al. 2022). Such a timeline for reionization may favor UV-bright, massive galaxies being the drivers of reionization (Naidu et al. 2020), as opposed to an undetected large population of faint objects (Finkelstein et al. 2019).

Many studies have also used the evolving luminosity function (LF) of LAEs to place constraints on reionization (Rhoads & Malhotra 2001; Malhotra & Rhoads 2004, 2006; Hu et al. 2019; Morales et al. 2021; Wold et al. 2022). However, the constraints on $\bar{x}_{\text{H I}}$ at $z = 7$ from LAE LFs are sometimes in slight tension. At $z = 7$, Zheng et al. (2017) find $\bar{x}_{\text{H I}} \sim 0.40\text{--}0.60$, Ota et al. (2017) find $\bar{x}_{\text{H I}} > 0.4$, Itoh et al. (2018) find $\bar{x}_{\text{H I}} = 0.25^{+0.25}_{-0.25}$, Hu et al. (2019) find $\bar{x}_{\text{H I}} \sim 0.2\text{--}0.4$, and Wold et al. (2022) find $\bar{x}_{\text{H I}} < 0.33$. More recently, Morales et al. (2021) compared observed Ly α LFs with models from an inhomogeneous reionization simulation to constrain $\bar{x}_{\text{H I}} \sim 0.08^{+0.08}_{-0.05}$, 0.28 ± 0.05 , and 0.69 ± 0.11 at redshifts 6.6, 7.0, and 7.3, respectively.

The moderate tension between these inferences on $\bar{x}_{\text{H I}}$ can perhaps be attributed to cosmic variance. There is stochasticity in the number of LAEs observed in a given survey arising from both large-scale structure and the inhomogeneity of reionization.

It is noteworthy that LAEs at $z > 6$ are fairly rare objects: dedicated narrowband surveys at $z \sim 7$ typically detect tens of objects in volumes $\sim 1 \times 10^6$ cMpc (Ota et al. 2017; Itoh et al. 2018; Hu et al. 2019; Wold et al. 2022). These small number statistics could easily lead to field-to-field variations in the measurement of the $z = 7$ LAE LF. This effect has been observed in features like the LAE LF’s “bright end bump” (Zheng et al. 2017; Hu et al. 2019; Wold et al. 2022), the result of very bright LAEs falling within the field of observation. The bright end bump is not observed in all LAE surveys at $z = 7$ (or even between different fields of the same survey), presumably either because of these objects’ inherent rarity or their concealment behind a neutral IGM.

This is the crux of the issue we will explore in this paper: there is a degeneracy between the stochasticity in the number of LAEs observed and the inferred global neutral fraction of the universe. Indeed, Mesinger & Furlanetto (2008a) showed that, owing to the inhomogeneity of reionization, there is intrinsic scatter in the inferred global neutral fraction of the universe from observations of the Ly α damping wing in high- z quasars or gamma-ray bursts (GRBs). McQuinn et al. (2008) similarly found that a single high- z GRB could not place a constraint on the global neutral fraction with an uncertainty better than $\Delta\bar{x}_{\text{H I}} \sim 0.3$. Mason et al. (2018a) demonstrated the same principle with regards to Ly α equivalent widths (EWs); inference on the global neutral fraction is inherently stochastic when using observations of the EWs of Ly α in UV-bright galaxies. Further, Mesinger & Furlanetto (2008a), McQuinn et al. (2008), and Mason et al. (2018a) all showed that the uncertainty on $\bar{x}_{\text{H I}}$ is a function of $\bar{x}_{\text{H I}}$ itself, $\Delta\bar{x}_{\text{H I}}(\bar{x}_{\text{H I}})$. Generally, the uncertainty tends to decrease as $\bar{x}_{\text{H I}}$ increases because the size distribution of ionized bubbles gets smaller, since they have not yet had time to grow and merge in early reionization (Mason et al. 2018a). We will demonstrate that there is a similar effect, a floor in the uncertainty on inferring $\bar{x}_{\text{H I}}$, when using the evolving Ly α LF, and explore mitigation strategies to reduce this uncertainty.

We post-process a cosmological inhomogeneous reionization simulation (Mesinger & Furlanetto 2007; Mesinger et al. 2011, 2016), populating the dark matter halos with simulated galaxies from empirical relations. We first verify our simulation against the intrinsic UV and Ly α LFs at $z \sim 6$, then model the Ly α LFs after IGM attenuation at $z \sim 7$. We perform mock observations on our simulation, reproducing the probed volumes of surveys which are carried out today. Across each of the many realizations of these mock observational programs, we measure the LAE LF and infer $\bar{x}_{\text{H I}}$. We investigate how the inferred $z = 7$ value of $\bar{x}_{\text{H I}}$ changes as a result of statistical variance in the observed LAE LFs.

In Section 2, we give an overview of the simulation and the post-processing we have performed. In Section 3, we explore the inherent uncertainty on the inferred global neutral fraction at $z \sim 7$ deriving from cosmic variance; we also explore if this uncertainty decreases by significantly increasing the survey area, depth, or using many independent fields. We conclude in Section 4. Throughout the paper we use the AB magnitude system and the Planck-Collaboration et al. (2016) cosmology with $(\Omega_{\Lambda}, \Omega_m, H_0) = (0.69, 0.31, 68 \text{ km s}^{-1} \text{ Mpc}^{-1})$.

2. The Reionization Simulation and Its Post-processing

In this section, we first summarize the simulation of dark matter halo masses, positions, and their Ly α optical depths. We then present an overview of our methods for assigning galaxy properties to the halos.

2.1. The 21cmFAST Simulation

We use nine custom simulations produced with the 21cmFASTv2 software (Mesinger & Furlanetto 2007; Mesinger et al. 2011, 2016). 21cmFASTv2 calculates the evolution of the hydrogen neutral fraction in the early universe using a semi-numerical approach. Inside the simulation cube, 1.6 cGpc on each side with a 1024^3 cell resolution, the code tracks dark matter and the phase of hydrogen in each cell while accounting for recombinations, photoheating star formation suppression, supernova feedback, and radiation. Dark matter halos are identified from the 3072^3 initial conditions, and mapped to Eulerian positions at $z = 7$ using perturbation theory (Mesinger & Furlanetto 2007). More detailed information can be found in Mesinger et al. (2016).

Each simulation contains the comoving Cartesian coordinate positions of halos with masses ranging from $10^{10.25}$ to $10^{12} M_{\odot}$. The halos start with only 14 discrete halo masses, evenly spread in log space. To avoid effects resulting from this discretized mass function, we redistribute the masses to produce a smooth distribution in log space, preserving each halo’s positions while approximating the original mass. The simulation also computes Ly α optical depths as a function of velocity offset from Ly α line center, $\tau(\Delta v)$, up to 525 km s^{-1} ($\Delta\lambda = 2.1 \text{ \AA}$) in steps of 75 km s^{-1} ($\Delta\lambda \sim 0.3 \text{ \AA}$). These optical depths are calculated by integrating the neutral hydrogen along a 300 comoving Mpc path; 300 cMpc was chosen experimentally, as it ensures convergence of the optical depth (Mesinger & Furlanetto 2008b). Intrinsically, the optical depths are smooth functions of velocity offset from the line center, but we only have coarsely sampled optical depths at particular velocity offsets. To mitigate any effect from this coarse sampling, we interpolate the optical depths across

velocity offsets when calculating the IGM attenuation on the Ly α line profile.

In total, we have nine $z = 7$ 21cmFASTv2 simulations with different global neutral fractions, $\bar{x}_{\text{H I}}$, in the range $0.01 < \bar{x}_{\text{H I}} < 0.92$. Each simulation takes the same $z = 7$ dark matter halo distribution and overlays on a different IGM ionization map, which corresponds to changing the ionization efficiency (McQuinn et al. 2007; Mason et al. 2018b). To get an intuitive picture, one may imagine that in a scenario with large $\bar{x}_{\text{H I}}$ (most IGM hydrogen neutral), any given halo will have a small ionized bubble around it, if any. In a low- $\bar{x}_{\text{H I}}$ scenario (most IGM hydrogen ionized), those same halos will have larger ionized bubbles around them, which may have begun to overlap with neighboring halos' bubbles. This means that each individual halo's optical depth, at a given velocity offset, smoothly increases as neutral fraction increases, which allows us to interpolate between the neutral fraction files, obtaining a finer grid in neutral fraction than our original nine simulations.

2.2. Propagating Halo Growth

In order to make predictions at $z = 7$, we need to calibrate the Ly α LF using observations at lower redshifts ($z \sim 6$), where there is evidence that reionization is very nearly complete (Qin et al. 2021). Though reionization may not be complete at $z \sim 6$, residual $\bar{x}_{\text{H I}}$ (i.e., if $\bar{x}_{\text{H I}} \sim 0.1$) has a very small effect on the optical depth of Ly α and so the LAE LF will not be significantly changed (Mason et al. 2018b). As redshift decreases, halos grow because they accrete matter through mergers. Thus, we need to forward propagate the halo mass distribution from $z = 7$ to lower redshifts for proper comparison to observations.

We use the median halo growth trajectories from Behroozi et al. (2019) to grow our halos. We interpolate between the trajectories to create a fine grid in halo mass, then assign each of the $z = 7$ halo masses to a particular growth curve. To assign a growth trajectory to a halo mass, we find the curve which has the $z = 7$ halo mass closest to the given halo mass in the simulation. Then, moving along the matched growth curves to the lower redshift of interest, we get the new masses. These new masses are then redistributed into a smooth log-space distribution.

2.3. Galaxy Properties

To achieve our goal of simulating LAEs, we need to assign intrinsic Ly α luminosities and Ly α line properties to the dark matter halos. Our road map will be as follows:

1. Assign absolute UV magnitudes (M_{UV}) using a M_h - M_{UV} relation (Section 2.3.1).
2. Assign a Ly α EW ($\text{EW}_{\text{Ly}\alpha}$) to each halo, using a UV-dependent EW probability distribution function (PDF; Section 2.3.2).
3. Assign Ly α line profiles, with line velocity offset and width dependent on the halo masses (Section 2.3.3).

Item 3 is required to compute the opacity to Ly α photons from the neutral IGM. In what follows, we first describe the methodology used to assign M_{UV} and Ly α EW to each halo and how this methodology was validated with real data. We then discuss the new prescription to assign the Ly α line profile and the comparison with observations.

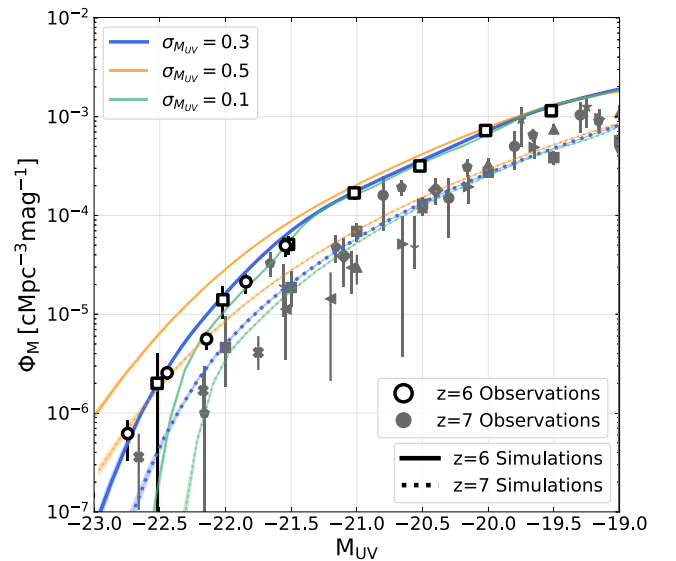


Figure 1. Three $z = 6$ UV LFs, resulting from three different M_h - M_{UV} scatter values, are plotted in solid green ($\sigma_{M_{\text{UV}}} = 0.1$), blue ($\sigma_{M_{\text{UV}}} = 0.3$), and orange ($\sigma_{M_{\text{UV}}} = 0.5$). We calculate the UV LF for each $\sigma_{M_{\text{UV}}}$ value across 50 different realizations of the simulation and show the region within which 95% of UV LFs fall with the shaded areas. The same color scheme is used for the $z = 7$ UV LFs, which are plotted as dotted lines. The $z = 6$ data are plotted as white markers with black borders; the squares are Bouwens et al. (2021) and the circles are Ono et al. (2018). The $z = 7$ data are shown in gray; the squares are Finkelstein et al. (2015), stars are Atek et al. (2015), circles are Bouwens et al. (2011), pentagons are Bouwens et al. (2015b), crosses are Bowler et al. (2014), diamonds are Castellano et al. (2010), down-pointing triangles are Livermore et al. (2017), up-pointing triangles are McLure et al. (2013), left-pointing triangles are Ouchi et al. (2009), right-pointing triangles are Schenker et al. (2013), and tri-down are Tilvi et al. (2013). The M_h - M_{UV} relation with a scatter of $\sigma_{M_{\text{UV}}} = 0.3$ reproduces the $z = 6$ and $z = 7$ data simultaneously. The bright end provides the demarcation between models with different $\sigma_{M_{\text{UV}}}$ values.

2.3.1. The M_h - M_{UV} Relation

We use the M_h - M_{UV} relation from Mason et al. (2015) to assign UV magnitudes to our halos. This relation has an in-built redshift dependence, since halo distributions and their growth rates change with redshift (Mason et al. 2015). We use relations tabulated at $z = 6$ and $z = 6.8$ in this paper. Considering our small range of redshifts of interest ($5.7 < z < 7$) and the relatively small evolution between the $z = 6$ and $z = 6.8$ M_h - M_{UV} relation, we do not derive a new M_h - M_{UV} relation for every redshift we consider. Instead, we apply the $z = 6$ relation if our simulation redshift is $5.7 < z < 6.4$ and the $z = 6.8$ relation if our simulation redshift is $6.4 < z < 7$.

The relations provided in Mason et al. (2015) do not include a prescription for the scatter, so we follow methods akin to Ren et al. (2019) and Whitley et al. (2020) and assume a normal scatter in M_{UV} . The specific value of the scatter is chosen so that the resulting UV LF reproduces the measured LFs at $z \sim 6$ and $z \sim 7$ simultaneously. In Figure 1, we show UV LF observations at $z = 6$ in black and $z = 7$ in gray. At both redshifts, we plot three UV LFs resulting from different scatter prescriptions: 0.1 in green, 0.3 in blue, and 0.5 in orange. Solid lines are used for $z = 6$, and dotted lines for $z = 7$. We apply each relation with its scatter 50 times and indicate the range within which 95% of the UV LFs fall with the shaded regions.

Motivated by the agreement between the UV LFs with $\sigma_{M_{\text{UV}}} = 0.3$ and the observations at $z = 6$ and $z = 7$, we fix the value of $\sigma_{M_{\text{UV}}}$ to 0.3. The fact that we are able to reproduce UV LFs at two redshifts simultaneously with only one varying




Figure 2. The log 2D distribution of galaxies at $z = 6.6$ in the M_{UV} – $EW_{Ly\alpha}$ plane. The colored lines denoted in the legend indicate traces of constant limiting $Ly\alpha$ line flux in $\text{erg s}^{-1} \text{cm}^{-2}$. A survey with a given line-flux limit could detect LAEs which lie above the trace corresponding to its flux limit. Note that there is an abundance of galaxies at $EW_{Ly\alpha} = 0$ along the very bottom of the distribution, corresponding to the non- $Ly\alpha$ -emitting population.

parameter is encouraging; it validates both our halo growth prescription and the M_h – M_{UV} relation.

2.3.2. M_{UV} – $EW_{Ly\alpha}$ Distribution

The distribution of rest-frame $EW_{Ly\alpha}$ given M_{UV} takes the form

$$p(EW_{Ly\alpha}|M_{UV}) = \frac{A(M_{UV})}{W_c(M_{UV})} e^{-\frac{EW_{Ly\alpha}}{W_c(M_{UV})}} H(EW_{Ly\alpha}) + [1 - A(M_{UV})] \delta(EW_{Ly\alpha}) \quad (1)$$

from Mason et al. (2018b). $H(EW_{Ly\alpha})$ is the Heaviside step function and $\delta(EW_{Ly\alpha})$ is the Dirac delta function. Mason et al. (2018b) fit the parameters A and W_c using the De Barros et al. (2017) measurements of M_{UV} and $EW_{Ly\alpha}$. The population from De Barros et al. (2017) were photometrically selected using the Lyman break, so-called Lyman break galaxies (LBGs), and it is worth noting that not all LBGs show $Ly\alpha$ emission. The parameter A accounts for the fraction of galaxies which do not emit $Ly\alpha$, and W_c determines the exponential decline of the PDF toward larger $EW_{Ly\alpha}$. Both parameters are allowed to vary as a function of M_{UV} , and we take their value from Mason et al. (2018b). To give a sense of the resulting distribution, the log 2D distribution of the M_{UV} – $EW_{Ly\alpha}$ plane for $z = 6.6$ galaxies, along with traces for constant $Ly\alpha$ flux limits, are shown in Figure 2. UV-faint galaxies have a larger probability of taking on a large $EW_{Ly\alpha}$ value than UV-bright galaxies.

We are inherently assuming that the underlying $EW_{Ly\alpha}$ distributions are the same at $z \sim 6$, where the De Barros et al. (2017) observations are made, and at $z \sim 7$, where we will apply the $EW_{Ly\alpha}$ distribution. Mason et al. (2018b) argue that, while oversimplifying, this assumption is justified by the relatively short time frame; less than 200 Myr pass between $z = 6$ and $z = 7$. Further, this means that any evolution in the $EW_{Ly\alpha}$ distribution between $z = 6$ and $z = 7$ would be attributed to an evolving \bar{x}_{H1} in our model, as explored in Mason et al. (2018b).

2.3.3. $Ly\alpha$ Line Profile

Since the $Ly\alpha$ optical depth due to the neutral IGM depends on the wavelength offset from the line center, we need to assign a $Ly\alpha$ line profile that accounts for radiative-transfer effects in the galaxies’ interstellar medium (ISM) and circumgalactic medium (CGM). We will take a relatively simple approach in modeling the $Ly\alpha$ line profiles, modifying the methods in Mason et al. (2018b).

We assume that $z \sim 6$ $Ly\alpha$ line profiles are statistically the same as those at $z = 7$ when ignoring attenuation from a neutral IGM. Under this assumption, we take observed $Ly\alpha$ line properties at $z \sim 6$, correlate them with observable galaxy properties, and assign them to our $z = 7$ halos. The line profiles we will assign, then, have the effects of the ISM and CGM built in. Attenuation by the IGM could change the observed $Ly\alpha$ line profile statistics at $z = 7$.

We assume that the line profile modified by the ISM and CGM is a Gaussian, with an offset with respect to the rest-frame line center ($\Delta v_{Ly\alpha}$). Unlike Mason et al. (2018b), we assign the velocity offset with respect to the galaxy rest-frame using an empirical relationship from Hayes et al. (2021), which is explained further in the next subsection.

The line’s FWHM is set equal to $\Delta v_{Ly\alpha}$, as found in Verhamme et al. (2018). The $Ly\alpha$ line profiles are truncated blueward of the $Ly\alpha$ rest-frame line center because anything blueward of the $Ly\alpha$ line center will redshift into resonance and be scattered by residual neutral fraction. The $Ly\alpha$ line profile after leaving the ISM and CGM is

$$J_\alpha(\Delta v_{Ly\alpha}, M_h, v) \propto \begin{cases} \frac{1}{\sqrt{2\pi}\sigma_\alpha} e^{-\frac{(v-\Delta v_{Ly\alpha})^2}{2\sigma_\alpha^2}} & \text{if } v \geq 0 \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

Note, though, that roughly one quarter of ultraluminous ($\log(L_{Ly\alpha}/[\text{erg s}^{-1}]) > 43.5$) LAEs at $z > 6$ emit flux blueward of line center (Hu et al. 2016; Songaila et al. 2018; Meyer et al. 2021). These LAEs likely reside in large ionized bubbles in the IGM. These bubbles may have been carved out by ionizing radiation from the ultraluminous LAEs themselves, a population of nearby ionizing sources, an obscured quasar, or a combination of contributors (see Bagley et al. 2017 and Mason & Gronke 2020 for a discussion). Due to their rarity, we do not modify our line profiles to account for these ultraluminous galaxies.

Hayes et al. (2021) compiled 229 $Ly\alpha$ selected galaxies at $2.9 < z < 6.6$ from observations with the Very Large Telescope’s (VLT) Multi Unit Spectroscopic Explorer (MUSE) and 74 galaxies at $z < 0.44$ from Hubble’s Cosmic Origins Spectrograph and compared the two populations’ $Ly\alpha$ line profiles. The VLT galaxies’ $Ly\alpha$ velocity offsets relative to their systemic redshifts (i.e., the line profiles’ first moments) were found by calibrating the observations with low- z data and a model (Runnholm et al. 2021). These velocity offsets, $\Delta v_{Ly\alpha}$, and the $Ly\alpha$ luminosities are correlated. The data are described in more detail in Hayes et al. (2021). Note that these galaxies were selected with the $Ly\alpha$ line, while the De Barros et al. (2017) sample which provided the $EW_{Ly\alpha}$ distribution were selected with the Lyman break. While not all LBGs have $Ly\alpha$ emission, we have captured this behavior in the $EW_{Ly\alpha}$ distribution; those that do have $Ly\alpha$ emission may be detected in the Hayes et al. (2021) sample, and so these empirical velocity offsets from $Ly\alpha$ -selected galaxies may be applied to




Figure 3. We show the first moment of the Ly α lines vs. the log of the Ly α luminosity for the MUSE data at $z < 3.5$ (green) and $z > 5.5$ (red). Our linear fit is shown with the blue line, while the black hashed lines show the $1 - \sigma$ Gaussian scatter about the relation that we fit with `emcee`.

LBGs as well. Note, however, that we find that the transmission fraction of Ly α is largely insensitive to the exact velocity offset value in our model, provided that the line has a width and so some of the flux is away from line center. For example, if we set the velocity offsets to 0 km s^{-1} , drastically different from the typical values from our relation of $200\text{--}300 \text{ km s}^{-1}$, we see a maximum of a 6% difference in the Ly α LFs in the range from $42.4 < \log_{10}(L_{\text{Ly}\alpha}) < 43.7$ for $\bar{x}_{\text{H I}} = 0.3$ (this luminosity range was chosen because of the observational range of the Ly α LFs at $z = 7$; see Figure 5). For $\bar{x}_{\text{H I}} = 0.6$, this percent difference grows to 20%.

We use `emcee` to fit a linear relation to the observed high- z LAEs’ line profiles’ $\Delta v_{\text{Ly}\alpha}$ and the log of their Ly α luminosities. The best-fit linear relation is given by

$$\Delta v_{\text{Ly}\alpha} = -3156^{+882}_{-895} + 80^{+21}_{-21} \log_{10}(L_{\text{Ly}\alpha}) \quad (3)$$

with a Gaussian scatter of $13.5^{+6.1}_{-3.5} \text{ km s}^{-1}$. The best-fit relation, with the MUSE data, is shown in Figure 3. The $z < 3.5$ MUSE data are shown in orange, and the $z > 5.5$ data are shown in green. Both redshift ranges are consistent with the same linear relation, indicating little to no redshift evolution. This linear relation is used to center the Gaussian before truncation; a random Gaussian scatter of 13.5 km s^{-1} is added to the offsets.

As we will only be concerned with Ly α luminosities larger than $L_{\text{Ly}\alpha} \gtrsim 10^{42} \text{ erg s}^{-1}$, the fact that our relation will predict a negative velocity offset at Ly α luminosities $\lesssim 4 \times 10^{39} \text{ erg s}^{-1}$ is not an issue.

With each halo assigned a M_{UV} , $\text{EW}_{\text{Ly}\alpha}$, Ly α optical depth (τ_{IGM}), and Ly α line profile, we have everything we need to compute the transmitted fraction of Ly α flux. Explicitly, the fractional transmission through the IGM is

$$T = \frac{\int_0^\infty J_\alpha e^{-\tau_{\text{IGM}}} dv}{\int_0^\infty J_\alpha dv}. \quad (4)$$

2.4. Validation

With all the pieces required to simulate LAEs in place, we proceed to validate the recipes by comparing the simulation’s $z = 5.7$ Ly α LF to the observations. At this redshift, the dark-pixel fraction in the Ly α and Ly β forest indicates that reionization is nearly complete (Qin et al. 2021), but even




Figure 4. The $z = 5.7$ Ly α LF is shown in blue with constraints from Konno et al. (2018) and Ouchi et al. (2010). We see very good agreement between our simulation and the observational constraints.

with a neutral fraction of $\bar{x}_{\text{H I}} \sim 10\%$ the impact on Ly α transmission is negligible (Mason et al. 2018b). As such, this comparison will test only the $M_h\text{--}M_{\text{UV}}$ and $M_{\text{UV}}\text{--}\text{EW}_{\text{Ly}\alpha}$ relations.

The $z = 5.7$ Ly α LF is shown in Figure 4 along with observations from Hu et al. (2010), Ouchi et al. (2010), and Konno et al. (2018). The halos have been grown to $z = 5.7$, as explained in Section 2.2 and the UV luminosities are assigned from these evolved masses. The $\text{EW}_{\text{Ly}\alpha}$ are assigned via the relation described in Section 2.3.2. The resulting $z = 5.7$ Ly α LF is in very good agreement with observations, lending credit to the recipes we have used thus far.

This work will remain focused on the impact and mitigation of cosmic variance on the inference of $\bar{x}_{\text{H I}}$, particularly at $z = 7$. A reader interested in analytic modeling of the Ly α LFs at other redshifts should consult Morales et al. (2021), which uses similar methods as in this paper.

2.5. Full Volume $z = 7$ Ly α Luminosity Functions

We used the elements discussed in the previous sections to compute the $z = 7$ Ly α LFs from the full $\sim 4.1 \text{ cGpc}^3$ simulation volume, and for all values of $\bar{x}_{\text{H I}}$ ranging from 0.01 to 0.92, with a spacing of 0.03. Using the entire volume allows us to minimize shot noise in the range of $\bar{x}_{\text{H I}}$ and Ly α luminosity of interest. For each value of $\bar{x}_{\text{H I}}$, we recalculate the simulation 30 times. We take the median of the 30 realizations as the true model; this has the effect of further minimizing shot noise and sampling the sources of stochasticity from our empirical relations. The resulting Ly α LFs are our best estimate of the true LAE LF at $z = 7$ for various $\bar{x}_{\text{H I}}$ values. In Section 3.2, we will compare mock observations to these forward models to see how well the observations constrain $\bar{x}_{\text{H I}}$.

Figure 5 shows the LAE LF calculated from the full simulation volume for various $\bar{x}_{\text{H I}}$ values. As $\bar{x}_{\text{H I}}$ increases, attenuation of intrinsic Ly α emission increases and the resulting LAE LF is suppressed; between a fully ionized universe and a universe with $\bar{x}_{\text{H I}} \sim 0.85$, the suppression is about an order of magnitude. The observational data at $z = 7$ agree well with a fully ionized universe, which implies that in our simulation the evolution of the Ly α LF from $z = 5.7$ to $z = 7$ in observations can be fully explained via halo growth and changing intrinsic Ly α emission. Previous works, such as




Figure 5. The curves show the LAE LF after IGM attenuation for various global neutral fractions, \bar{x}_{HI} , at $z=7$. The shaded regions show the range in which 95% of 30 simulation realizations reside. Observational data from Ota et al. (2017), Itoh et al. (2018), and Hu et al. (2019) are shown with various markers. The data have been offset by 0.01 with respect to each other on the abscissa to aid with visualization. The observational data at $z=7$ is consistent with a completely or mostly ionized IGM, as demonstrated with the data’s agreement with the blue and orange curves. This, combined with the $z=5.7$ Ly α LF agreement, shows that the evolution of the $z=5.7$ to $z=7$ Ly α LF is explicable with smaller halos and lower intrinsic Ly α emission, rather than an increase in global neutral fraction in our simulation.

Dijkstra et al. (2014), have argued that at least some of the Ly α LF evolution likely comes from evolution of galaxy properties.

3. Cosmic Variance in Ly α Emitters

With our simulation constructed, we begin analysis to determine the importance of cosmic variance on the inference of global neutral fraction. As laid out in the Introduction (Section 1), Ly α photons are easily scattered by neutral hydrogen. This has the effect of decreasing the number of observed LAEs when the neutral fraction of HI increases. This decreases the volume density of observed LAEs. The differential change in the Ly α LF between redshifts has been used to infer the global neutral hydrogen fraction of the universe (Malhotra & Rhoads 2004; Ouchi et al. 2010; Kashikawa et al. 2011; Konno et al. 2014, 2018; Ota et al. 2017; Zheng et al. 2017; Itoh et al. 2018; Hu et al. 2019; Morales et al. 2021; Wold et al. 2022). Our goal is to quantify how cosmic variance impacts this inference and explore mitigation strategies in survey design with an updated, empirical LAE model.

3.1. A Nominal Ly α Emitter Survey

To gain some footing, we first consider a nominal survey, representative of volumes probed by typical surveys carried out by ground-based telescopes equipped with narrowband filters. We will “observe” our simulation with the specifics of the nominal survey and use Bayesian inference to estimate \bar{x}_{HI} from the “observed” Ly α LF.

We simulate a circular survey covering 2 deg^2 to a Ly α flux limit of $f_{\text{Ly}\alpha} > 1 \times 10^{-17} \text{ erg s}^{-1} \text{ cm}^{-2}$ (or $L_{\text{Ly}\alpha} \gtrsim 6 \times 10^{42} \text{ erg s}^{-1}$ at $z=7$). Such a survey pushes faint enough in flux/luminosity to detect LAEs fainter than the knee of the Ly α LF at $z=7$ (Ota et al. 2017; Itoh et al. 2018; Hu et al. 2019). This depth is comparable to those reached by narrowband surveys, and so provides an interesting baseline. Notably, we will consider a LAE “observed” simply if it falls into the geometric cutout and exceeds the

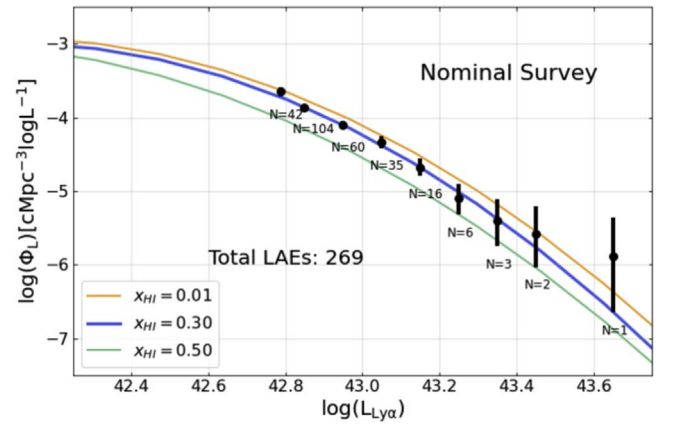


Figure 6. The black points show a LAE LF measured from a 2 deg^2 survey with perfect observations to a depth of $f_{\text{Ly}\alpha} > 1 \times 10^{-17} \text{ (erg s}^{-1} \text{ cm}^{-2})$ and a redshift window of $6.75 < z < 7.25$. The model LFs for $\bar{x}_{\text{HI}} = [0.01, 0.3, 0.5]$ are shown as the blue, black, and green curves, respectively. The error bars are assumed to be Poissonian for $N > 20$, otherwise they are tabulated from Gehrels (1986). The intrinsic scatter of the black points about the 0.3 model LF, the input neutral fraction, comes from cosmic variance.

minimum line-flux limit; in this way, our survey ignores any effects of contamination or incompleteness, and so any stochasticity on the inferred values of \bar{x}_{HI} between surveys is the result of stochasticity in the number of observed LAEs alone. We impose a redshift window of $\Delta z \sim 0.5$ centered at $z=7$ by selecting galaxies in a slice of 160 cMpc within the simulated cube. Such a survey has a volume of $7.58 \times 10^6 \text{ cMpc}^3$ at $z=7$, a volume comparable to $z=7$ narrowband surveys, which have larger areas but smaller redshift window (Ota et al. 2017; Itoh et al. 2018; Hu et al. 2019; Wold et al. 2022). One should keep in mind that the simulation is a snapshot at $z=7$: there is no differential evolution between the “front” end of our survey and the “back” end taken into account. This allows us isolate the effect of cosmic variance by treating the global neutral fraction as a constant over our survey window. We then have a well-defined global neutral fraction that should be inferred in our simulation from the mock observations. Note, however, that this lack of differential evolution is not important with regards to any specific LAE. The cumulative optical depth is dominated by the gas nearby the galaxy, so much so that the neutral patches with optical depth $\gtrsim 1$ are within $\lesssim 10 \text{ Mpc}$, corresponding to $\Delta z \approx 0.02$ at $z=7$ (Mesinger & Furlanetto 2008a). This implies that the differential evolution is unimportant when calculating the optical depth for a given LAE.

We create a $z=7$ simulation with a realistic input global neutral fraction of $\bar{x}_{\text{HI}} = 0.3$ (Morales et al. 2021) and observe an area in a randomized location within the cube. This randomized location can, in principle, have a different local neutral fraction than the input neutral fraction, but note that we are interested in inferring the latter rather than the former. This survey results in a total sample of 274 LAEs. The LAE LF for one specific realization is shown in Figure 6. We split the LAEs into bins of width $\Delta \log(L_{\text{Ly}\alpha}) = 0.1$. The model Ly α LFs (as measured from the entire simulation volume) for $\bar{x}_{\text{HI}} = 0.01, 0.30$ (input value), and 0.50 are shown as the blue, black, and green curves, respectively. The error bars associated with the simulated LF measurements come from Poisson statistics on the counts in each bin if $N_{\text{LAE}} > 20$. For smaller numbers of LAEs, we use the estimates tabulated in Gehrels (1986) for a small number Poisson errors.

Table 1
Beta Parameters

β_0	β_1	β_2	β_3	β_4	β_5	β_6	β_7	β_8
-2436.6	115.7	-2563.8	-1.4	-1.4	3.3	6146.7	-286.6	120.0

3.2. Inference on \bar{x}_{HI} from the Nominal Survey

In order to speed up the inference on \bar{x}_{HI} , we fit a function to the model LAE LFs from Section 2.5 as a function of \bar{x}_{HI} . More information on this analytical function can be found in the Appendix, along with the best-fit parameters in Table 1. Some example LAE LFs for different \bar{x}_{HI} are visualized in Figure 6.

We use a Bayesian approach (implemented using the Markov Chain Monte Carlo, MCMC, fitting from the `pymc3` package) to estimate the best-fit hydrogen neutral fraction, \bar{x}_{HI} , given our model Ly α LFs (Section 2.5) and the simulated LF measurement (Section 3.1). The 2D surface fit in the Appendix is used to create the predicted LAE volume densities, μ , for a given neutral fraction. Then, the likelihood is given by

$$p(\phi_{\text{L}}|\bar{x}_{\text{HI}}, \tau) = \prod_i \sqrt{\frac{\tau}{2\pi}} \exp\left(-\frac{\tau}{2}(\phi_{\text{L},i} - \mu_i)^2\right), \quad (5)$$

where the subscript i denotes each of the measured luminosity bins. The unknown model parameters to be fit are the precision, τ , of the observed LAE LF data and \bar{x}_{HI} . We consider a gamma function prior on τ with $\alpha = \beta = 1$ and a flat prior on \bar{x}_{HI} .

The posterior distribution function for \bar{x}_{HI} is then given by

$$p(\bar{x}_{\text{HI}}|\phi_{\text{L}}) \propto \int p(\phi_{\text{L}}|\bar{x}_{\text{HI}}, \tau)p(\tau)p(\bar{x}_{\text{HI}})d\tau, \quad (6)$$

where $p(\tau)$ is the gamma prior on τ and $p(\bar{x}_{\text{HI}})$ is the uniform prior on \bar{x}_{HI} .

We note that we are assuming here that the precision of all of the LF measurements, τ , is the same. The results of our paper do not change if we instead assign adaptive luminosity bins to ensure that this assumption is true, that the precision between all the measurements is the same, i.e., the same number of LAEs go into each bin. We use fixed bin sizes in $\log(L_{\text{Ly}\alpha}/[\text{erg s}^{-1}])$, rather than these adaptive bins, for ease of comparison to measurements in the literature.

The posterior on \bar{x}_{HI} for the LAE LF shown in Figure 6 is displayed in Figure 7. The input neutral fraction is denoted with the vertical black dashed line. We can see that the width of the posterior is rather wide and the posterior PDF continues increasing toward low values of \bar{x}_{HI} . The posterior's median is indicated with the vertical dashed blue line—it underestimates the real neutral hydrogen fraction. Looking back at the LAE LF in Figure 6, the origin of this offset is apparent: more of the observed LF data points lie above the input \bar{x}_{HI} than below it, and a higher number density of LAEs leads to an inference of a lower \bar{x}_{HI} . In other words, the global neutral fraction is underestimated because the number of LAEs in the specific realization of the nominal area is higher than the predicted number, i.e., the randomized survey position fell on a slightly overdense region of the simulation (with respect to LAEs).

Motivated by this result, in the following sections we explore the range of LAE densities expected at $z = 7$, quantify how this range of environments propagates to a range on the inferred \bar{x}_{HI} , and then test methods to mitigate the effect of varying environments on the inference of \bar{x}_{HI} .




Figure 7. The posterior on \bar{x}_{HI} inferred from the LAE LF measured in Figure 6 is shown in blue. The simulation's input \bar{x}_{HI} , 0.3, is indicated with the vertical dashed black line. The posterior favors a slightly lower \bar{x}_{HI} than the true value because the measured region of space happened to have a slightly overdense number of LAEs. The large width of the posterior indicates that \bar{x}_{HI} is not very well constrained; 95% of the posterior is contained in the range $0.01 < \bar{x}_{\text{HI}} < 0.45$.

3.3. The Effect of Cosmic Variance on \bar{x}_{HI} Inference

We simulate 10,000 realizations of the nominal survey scattered randomly about the simulation volume, always with a global neutral fraction $\bar{x}_{\text{HI}} = 0.3$, and consider the total number of LAEs, N_{LAE} , in each realization which meet the detection threshold. The distribution of N_{LAE} is shown in blue in the left panel of Figure 8. For comparison, a Poisson distribution with the same population mean is shown in black; two particular survey realizations, one at the 2.5 and one at the 97.5 percentile of the distribution of N_{LAE} , are noted with green and orange vertical lines, respectively. We will look at these particular realizations to determine a conservative (i.e., covering 95% of surveys) absolute variance in \bar{x}_{HI} inferred from the different over- and underdense environments.

We calculate the posteriors on \bar{x}_{HI} from the LAE LF measured from the over- and underdense regions and show them in the right panel of Figure 8; orange shows the posterior on \bar{x}_{HI} inferred from the LAE-rich region, and green shows \bar{x}_{HI} inferred from the LAE-sparse region. The difference between the two medians is quite large, about $\Delta\bar{x}_{\text{HI}} \sim 0.19$. This difference is critically important, as it is impossible to know a priori if a particular survey has landed in a region which is overdense, underdense, or something close to the average density, and this has a direct effect on the inferred \bar{x}_{HI} . This issue is still present for studies which use the differential evolution of the LAE LF to infer \bar{x}_{HI} , because we cannot know if the comparative surveys are in over- or underdense regions of the universe with respect to LAEs.

This inference was done in a very ideal scenario, where there is no impurity, no incompleteness, and the model LAE LFs we

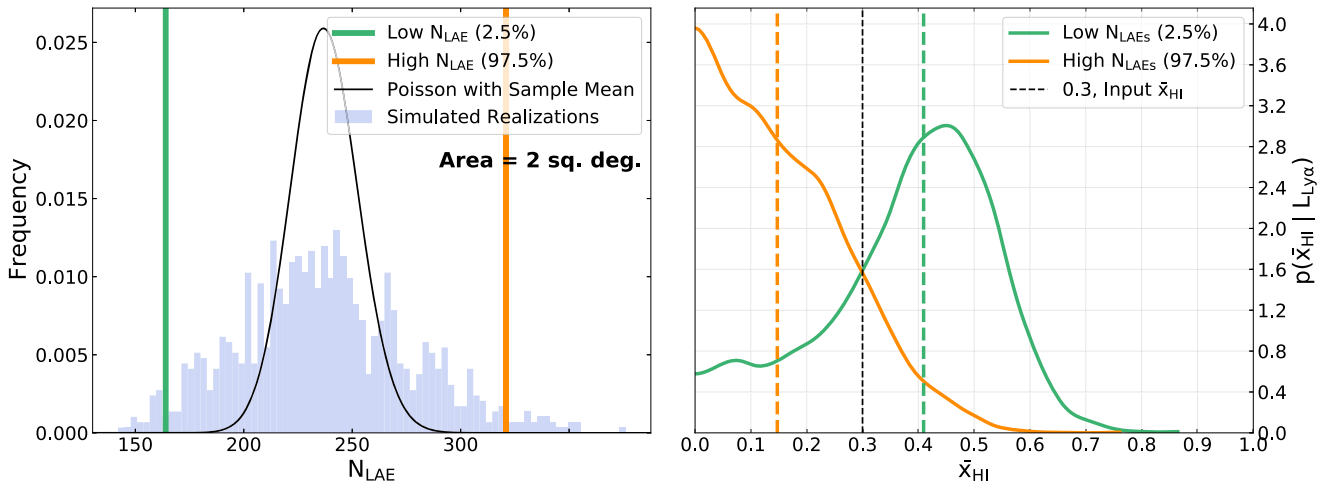


Figure 8. Left: the distribution of the number of LAEs observed, N_{LAE} , in 10,000 realizations of the nominal 2 deg^2 survey in a $\bar{x}_{\text{HI}} = 0.3$ at $z = 7$ universe in blue. We investigate one overdense region (orange vertical line) and one underdense region (green vertical line) further. A Poisson distribution with the same sample mean is shown in black—the excess variance over the Poisson distribution is due to the clumpy nature of the large-scale structure and the rarity of LAEs. Right: the posteriors on \bar{x}_{HI} for the extrema of the 95th percentile range of LAE densities in orange (high LAE density) and green (low LAE density). The simulation’s input \bar{x}_{HI} is shown with the vertical black dashed line. The difference between the medians of these distributions is $\Delta \sim 0.19$.

are comparing to our “observations” are from the same simulation, and so we know we have the physics in the modeling precisely correct. In reality, there would be additional uncertainty entering from the fact that comparison models will make assumptions on the underlying physics. Despite this extremely optimistic scenario, $\Delta \bar{x}_{\text{HI}} \sim 0.19$ at $z = 7$ when inferring from our nominal survey’s LAE LF. It is important to note that this uncertainty in \bar{x}_{HI} is a function of \bar{x}_{HI} itself, $\Delta \bar{x}_{\text{HI}}(\bar{x}_{\text{HI}})$. We find, in agreement with Mason et al. (2018a), that the uncertainty decreases as \bar{x}_{HI} increases, owing to a narrower distribution of ionized bubble sizes at high \bar{x}_{HI} values. We save a full quantification of $\Delta \bar{x}_{\text{HI}}(\bar{x}_{\text{HI}})$, akin to the work done for quasars and GRBs in Mesinger & Furlanetto (2008a), for future work, and instead take the uncertainty at $\bar{x}_{\text{HI}} = 0.3$ as a baseline and investigate mitigation strategies.

It is noteworthy that this nominal survey is already fairly large and pushes to moderate flux depths. Taking into account the considered redshift window, $6.75 < z < 7.25$, the survey’s volume, $7.58 \times 10^6 \text{ cMpc}^3$, is larger than many of the $z = 7$ LAE surveys observed to date, and comparable to Wold et al. (2022). Larger volumes, deeper surveys, or independent fields may mitigate this systematic uncertainty, so we now explore how $\Delta \bar{x}_{\text{HI}}$ changes when varying the survey strategy.

3.4. Mitigating the Systematics

We consider three scenarios to mitigate systematic uncertainty on the inference of \bar{x}_{HI} : increasing the survey area, increasing the survey depth, and considering a survey composed of many independent fields.

3.4.1. Increasing Area

To determine survey areas which may be effective at decreasing the effect of cosmic variance on the inference of \bar{x}_{HI} , we consider four increasingly large areas centered on the two positions identified in Figure 8, i.e., an overdense and an underdense region. We simulate areas of 2, 5, 10, and 20 deg^2 , while keeping the survey line-flux limit fixed to the nominal survey. For each survey, we use the same procedure described in Section 3.2 to infer the neutral hydrogen fraction. The

resulting posteriors for both the over- and underdense regions are shown in Figure 9.

Figure 9 shows that increasing the area by a factor of 5–10 has a tendency to modestly decrease the width of the posterior: in the LAE-sparse region, the standard deviation of the 2 deg^2 survey’s posterior is 0.15, while the 20 deg^2 survey’s standard deviation is 0.13.

We take the extremely optimistic case, an observationally perfect 20 deg^2 survey with $\Delta z = 0.5$ and a luminosity limit $L_{\text{Ly}\alpha} > 6 \times 10^{42} \text{ erg s}^{-1}$, and investigate it further.

We compute 100 realizations of this survey, randomly positioning them within the simulated volume, and infer \bar{x}_{HI} from each. We calculate the median for each of these 100 posteriors and plot the distribution in Figure 10. For a baseline comparison, we repeat the same procedure with the nominal 2 deg^2 survey (which has the same flux limit as this 20 deg^2 survey); 100 realizations are created and each is used to make an inference on \bar{x}_{HI} . The medians of the 100 nominal survey posteriors are plotted in Figure 10 as the underlying black distribution. Ninety-five percent of the nominal survey’s medians are within a range $\Delta \bar{x}_{\text{HI}} = 0.25$, while the 20 deg^2 medians have 95% of their realizations within $\Delta \bar{x}_{\text{HI}} = 0.21$. The width of the distribution of these medians gives us a sense of the systematic uncertainty in \bar{x}_{HI} arising from cosmic variance between the two survey strategies; there is an extremely modest tightening of posteriors for the 20 deg^2 surveys compared to the 2 deg^2 surveys, which does not seem to be statistically significant. It is important to emphasize that the narrowing of the posteriors in Figure 9 was for one particular survey only, whereas Figure 10 shows that the uncertainty arising from cosmic variance is not improved by increasing the area. This is demonstrated by the fact that the distributions of the medians of the 100 posteriors for the 2 and 20 deg^2 surveys are approximately equivalent.

Thus, even increasing our area by a factor of $10\times$ to 20 deg^2 , considering a fairly wide redshift window, and detecting LAEs fainter than the LF knee at $z = 7$ is not enough to significantly reduce the uncertainty on \bar{x}_{HI} . For surveys larger than 20 deg^2 , we begin to get to the point where even our extremely large simulation cannot provide a large number of independent fields




Figure 9. We increase the area of the nominal surveys on the previously considered density extrema locations while keeping the luminosity limit and redshift window constant. The posteriors in the low- N_{LAE} region flatten out at a nonzero value because of a single bright LAE which happened to fall within the survey. Increasing the area to 20 deg^2 centers moves the peak of the posterior toward the true value of \bar{x}_{HI} , so we will further investigate 20 deg^2 surveys to determine the intrinsic spread in inferred \bar{x}_{HI} from cosmic variance.




Figure 10. The distribution of 100 posterior medians for the 20 deg^2 survey (blue) and nominal 2 deg^2 survey (black) for MCMC analyses. Both of these surveys have $\Delta z = 0.5$ and a luminosity limit $L_{\text{Ly}\alpha} > 6 \times 10^{42} \text{ (erg s}^{-1}\text{)}$. The area of the 20 deg^2 survey was chosen because it seemed to remove the offset in inferred \bar{x}_{HI} with respect to the true value seen in Figure 9. The scatter in the posterior medians is consistent between the two survey strategies, indicating that even increasing the area of the survey by a factor of 10 is not enough to overcome the fact that the uncertainty on \bar{x}_{HI} is dominated by cosmic variance and imposes an uncertainty $\Delta \bar{x}_{\text{HI}} \sim 0.2$.

of view. We thus cannot conclusively say what volume is necessary to drive the uncertainty on \bar{x}_{HI} below ~ 0.2 for a survey with flux limit $f_{\text{Ly}\alpha} > 1 \times 10^{-17} \text{ erg s}^{-1} \text{ cm}^{-2}$; we can only say that it is larger than $7.58 \times 10^7 \text{ cMpc}^3$, a significantly larger volume than those which have been probed in LAE surveys so far. Evidently, simply taking a survey with a larger area is not enough to mitigate the effect of cosmic variance.

3.4.2. Increasing Depth

One may wonder what effect additional depth would have on the inference of \bar{x}_{HI} , motivated by the fact that less luminous LAEs will generally occupy less massive halos, and so be less

biased, and thus be less sensitive to cosmic variance. To investigate this idea, we repeat the exercise from Section 3.4.1, but modify the survey: we instead consider a 10 deg^2 survey with a flux limit of $f_{\text{Ly}\alpha} > 3 \times 10^{-18} \text{ erg s}^{-1} \text{ cm}^{-2}$ (or $L_{\text{Ly}\alpha} \gtrsim 1.8 \times 10^{42} \text{ erg s}^{-1}$ at $z=7$). This survey's flux limit is about $3.3\times$ deeper than the nominal survey and it is $5\times$ larger, so it would take about $55\times$ as long to execute. Going deeper in flux results in a drastic increase in the number of LAEs observed due to the shape of the LF; going from a flux limit of $f_{\text{Ly}\alpha} > 1 \times 10^{-17} \text{ erg s}^{-1} \text{ cm}^{-2}$ to $f_{\text{Ly}\alpha} > 3 \times 10^{-18} \text{ erg s}^{-1} \text{ cm}^{-2}$ increases the number of observed LAEs by about a factor of 10. Following the steps from the previous section, we calculate the posteriors on \bar{x}_{HI} from 100 realizations of such a survey; again, the distribution of the medians for the 100 surveys are shown in Figure 11. Again, the distribution for the 100 nominal surveys is shown in black as a baseline to compare against.

The width which encompass 95% of the posterior medians for the deep survey is virtually the same as the nominal survey, $\Delta \bar{x}_{\text{HI}} = 0.23$. Thus, even pushing $\sim 3.3\times$ deeper in flux and increasing the area by $5\times$ over the nominal survey, requiring $\sim 55\times$ more exposure time, has not decreased the systematic uncertainty.

3.4.3. Independent Fields

Lastly, we consider a survey composed of many independent fields, the motivation being that each field will probe a different environment for the LAEs, averaging out high- and low-density environments.

Further motivation can be seen examining Figure 12. The left panel shows the distributions of the number of LAEs per deg^2 observed by 1000 surveys of areas 2 and 20 deg^2 (and to depth $L_{\text{Ly}\alpha} \gtrsim 6 \times 10^{42} \text{ erg s}^{-1}$). The blue dashed line indicates a 2 deg^2 survey realization centered on an underdense region (at the 2.5 percentile) and the red dashed line shows a 20 deg^2 survey centered on that same position. As expected, the distributions of the number of LAEs per square degree




Figure 11. The distribution of 100 posterior medians for the deep ($L_{\text{Ly}\alpha} > 1.8 \times 10^{42} [\text{erg s}^{-1}]$) 10 deg^2 surveys (blue) and the posterior medians of the 100 nominal surveys (black). For both survey strategies $\Delta z = 0.5$. The widths of the distributions of both survey strategies is about the same: 95% of the nominal survey’s medians are within a range $\Delta \bar{x}_{\text{HI}} = 0.25$, while the deep 10 deg^2 medians have 95% of their realizations within $\Delta \bar{x}_{\text{HI}} = 0.23$. This indicates that pushing to deeper flux limits is not a viable strategy in mitigating the effect of cosmic variance.

converges toward the overall simulation mean as the survey area increases; the 20 deg^2 surveys’ distribution (red) is far narrower than the 2 deg^2 surveys’ (blue).

However, the imprint of the underdense region identified with a 2 deg^2 survey is still dominating the number of LAEs observed in the 20 deg^2 survey when compared to the parent distribution, as shown in the right panel. The 20 deg^2 survey centered on the underdense region (dashed red vertical line) is extremely close to the parent distribution’s 2.5 percentile (solid black vertical line).

This is indicative of the fact that adding area to a given survey has a tendency to add the average number of LAEs per square degree, by definition. This will have no effect on moving a given realization toward the center of the distribution of N_{LAE} for the parent distribution of that survey size; only adding an underdense or overdense region changes its position relative to the parent distribution. Then, a survey centered on an underdense (overdense) region will continue to exhibit low (high) LAE counts relative to other similar surveys, even as the area increases drastically.

Instead, we simulate a survey of 100 independent 0.2 deg^2 fields, totaling 20 deg^2 , to a depth of $f_{\text{Ly}\alpha} > 3 \times 10^{-18} \text{ erg s}^{-1} \text{ cm}^{-2}$, the same total area as our largest survey and as deep as our deepest simulated survey. All of these surveys are simulated in a simulation with global neutral fraction $\bar{x}_{\text{HI}} = 0.3$. Each of the 100 0.2 deg^2 fields has the inference done on them separately, and the resulting 100 posteriors are multiplied together to form a joint posterior on \bar{x}_{HI} .

Running a survey in this particular manner is computationally expensive, and so we limit ourselves to 90 realizations of such a survey; we show the distribution of the posterior medians in Figure 13. The width which encompasses 95% of the surveys has decreased, $\Delta \bar{x}_{\text{HI}} = 0.05$. Probing many independent fields is the only effective way to decrease the systematic uncertainty introduced by cosmic variance, decreasing the 95% width by about a factor of 4. Note that the inferred global neutral fraction is biased slightly high by about 0.03. This is because the analytical fit to the model $\text{Ly}\alpha$ LFs is only accurate to 4%. In practice, it would be possible to remove this

bias by bypassing the step of fitting the LFs analytically; however, for our analysis we need to compare many observing strategies and do hundreds of instances of inference for each. The analytic function speeds up the inference process and is the only way to complete our tests in a feasible amount of time.

This method has an additional advantage: Mesinger & Furlanetto (2008b) showed that counts-in-cell statistics can be an effective way to use the enhanced clustering of LAEs during reionization to constrain \bar{x}_{HI} . While the counts-in-cells method does not strictly require that the fields be discontinuous, a survey strategy such as ours has the additional benefit of overcoming any effect of cosmic variance, which could impact a clustering analysis as well. Thus, such a survey would have two methods of constraining \bar{x}_{HI} , both of which are robust against cosmic variance, and which could be checked against each other for consistency.

3.4.4. The Implication of the Mitigation Techniques

We are left to conclude that the only effective way to break through the floor in the uncertainty in \bar{x}_{HI} arising from the cosmic variance of LAEs is with surveys composed of many independent fields. Drastically increasing the survey area and pushing to much deeper fluxes alone proved ineffective. Of note, the survey composed of independent fields, with a flux limit of $f_{\text{Ly}\alpha} > 3 \times 10^{-18} \text{ erg s}^{-1} \text{ cm}^{-2}$ and an area of 20 deg^2 would take ~ 110 times as long to carry out as the nominal survey, which is currently representative of the largest dedicated LAE surveys to date. Thus, this serves more as a proof of concept than a realistic goal for a survey design in the immediate future.

It is worth keeping in mind that an uncertainty of $\Delta \bar{x}_{\text{HI}} = 0.2$, as we have demonstrated is characteristic when the universe has $\bar{x}_{\text{HI}} = 0.3$, is still a competitive constraint. Particularly, if such measurements are made at a range of redshifts, we may be able to construct a fairly constraining timeline of reionization; thus, the LAE LF will continue to be a useful tool in constraining reionization, in addition to other techniques, such as damping wings in quasars and GRBs, LAE clustering, LAE fraction, and $\text{Ly}\alpha$ EW distributions.

4. Conclusions

We have post-processed a $z = 7$ comoving 4.1 Gpc^3 simulation to create a mock LAE catalog using empirical relations. Our simulation simultaneously reproduces the $z = 6$ and $z = 7$ UV LFs and the $z = 5.7$ $\text{Ly}\alpha$ LF.

We produce $z = 7$ $\text{Ly}\alpha$ LF models, ranging from $42 < \log(L_{\text{Ly}\alpha}/[\text{erg s}^{-1}]) < 44$, as a function of \bar{x}_{HI} . We created mock surveys with different observation strategies and compare the measured $\text{Ly}\alpha$ LFs to these models, inferring posteriors on \bar{x}_{HI} . The observed LAE LF varies, deriving from the large-scale structure and stochasticity in the amount of neutral gas along the sightline between the LAEs and the observer, and the shape and position of the peak of the posterior distribution function on \bar{x}_{HI} changes in turn.

We show that the precision of \bar{x}_{HI} estimates drawn from an LAE LF derived from a single 2 deg^2 field with $f_{\text{Ly}\alpha} > 1 \times 10^{-17} \text{ erg s}^{-1} \text{ cm}^{-2}$ is limited by the cosmic variance of LAEs. We find an uncertainty floor of $\Delta \bar{x}_{\text{HI}} \sim 0.2$ when $\bar{x}_{\text{HI}} = 0.3$. Note that a 2 deg^2 survey with a redshift window of $\Delta z = 0.5$ has a volume of $7.58 \times 10^6 \text{ cMpc}^3$, about the same size as existing narrowband LAE surveys.

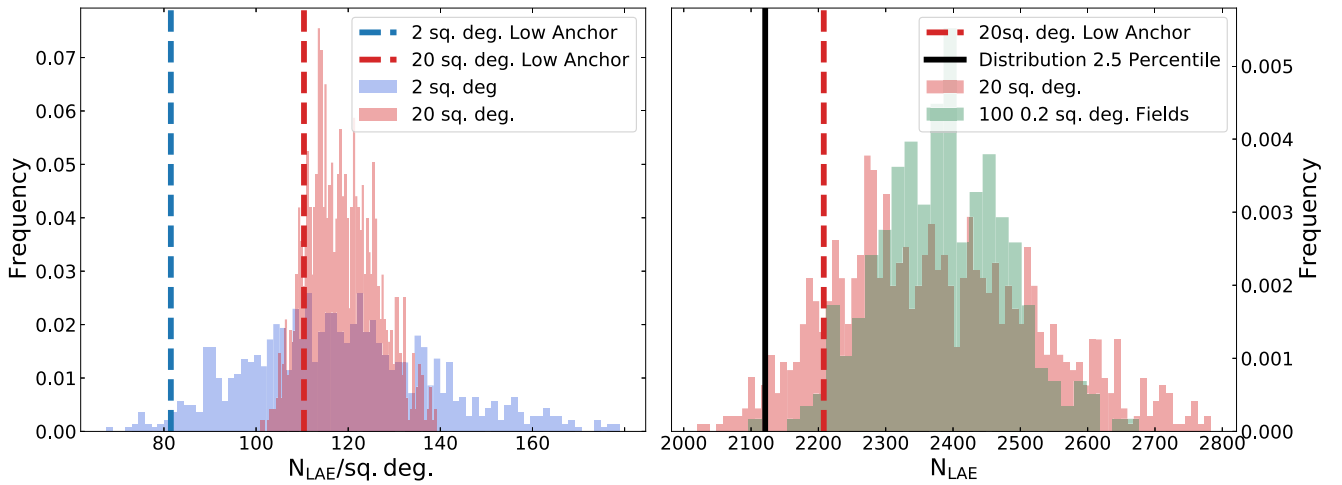


Figure 12. The left panel shows the distribution of the number of LAEs per square degree observed by 1000 surveys of areas 2 and 20 deg^2 (and to depth $L_{\text{Ly}\alpha} \gtrsim 6 \times 10^{42} \text{ erg s}^{-1}$) in blue and red, respectively. Surveys for each area, centered on the underdense region identified in the nominal 2 deg^2 survey, are indicated with the dashed vertical lines of the same colors. The distributions of N_{LAE} per square degree tend to narrow as survey area increases. The right panel shows the distribution of the N_{LAE} observed in 1000 realizations of the 20 deg^2 . The underlying green distribution shows the distribution of N_{LAE} in 300 surveys of 100 0.2 deg^2 fields, totaling 20 deg^2 . This survey strategy will be discussed more in Section 3.4.3; note that the distribution of N_{LAE} is narrower for this strategy. The dashed red vertical line again indicates the underdense region selected with the 2 deg^2 survey, and the distribution’s 2.5 percentile is indicated with the solid black vertical line. The imprint of the underdense region identified with a 2 deg^2 survey is still dominating the 20 deg^2 survey’s total number of LAEs compared to other 20 deg^2 surveys, illustrated by the fact that the red dashed line does not move much toward the distribution’s mean.

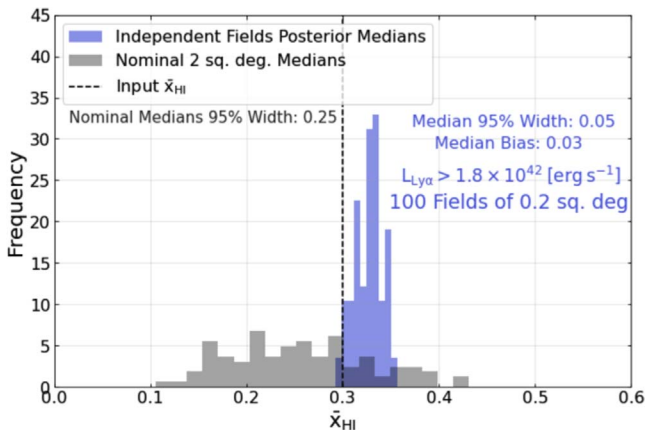


Figure 13. The distribution of posterior medians for 90 surveys composed of 100 0.2 deg^2 fields to a depth of $L_{\text{Ly}\alpha} \gtrsim 1.8 \times 10^{42} \text{ erg s}^{-1}$ is shown in blue. The distribution of the 100 nominal survey’s medians is shown in black. Ninety-five percent of the nominal survey’s medians are within a range $\Delta \bar{x}_{\text{HI}} = 0.25$, while the 90 surveys composed of 100 independent fields have 95% of their realizations within $\Delta \bar{x}_{\text{HI}} = 0.05$. This strategy is effective in reducing the systematic uncertainty on \bar{x}_{HI} .

We investigated three methods to push this floor in the uncertainty down: (i) increasing the area covered by a factor of 10, (ii) pushing to $\sim 3.3\times$ deeper flux limits, and (iii) breaking up the survey into independent fields. For the added area and deeper survey strategies, the variance in the medians of the posteriors on \bar{x}_{HI} remains virtually unchanged from the nominal 2 deg^2 survey. This demonstrates that the \bar{x}_{HI} inferred from observed LFs in contiguous fields are largely at the whim of what cosmic environment one’s survey happened to land in.

However, a large, deep survey composed of smaller independent fields proved effective in reducing the systematic uncertainty on \bar{x}_{HI} . Under this strategy, we see the width of 95% of the posterior medians of \bar{x}_{HI} decrease from $\Delta \bar{x}_{\text{HI}} \sim 0.2$ to $\Delta \bar{x}_{\text{HI}} \sim 0.05$. Thus, probing multiple independent fields is critically important in constraining \bar{x}_{HI} .

This result demonstrates that surveys of LAEs aiming at constraining \bar{x}_{HI} should adopt a strategy similar to the one proposed in Section 3.4.3 in order to minimize the systematic uncertainty on \bar{x}_{HI} . The Parallel Application of Slitless Spectroscopy to Analyze Galaxy Evolution (PASSAGE; PI: M. Malkan) JWST survey, as a pure parallel survey, follows this strategy, though JWST’s small field of view implies that the total area will be much smaller than the survey proposed in Section 3.4.3. PASSAGE is being carried out during Cycle 1 of JWST observations, and its observations of high- z LAES will allow the inference of new information about the timeline of reionization. The Nancy Grace Roman Space Telescope, with its large field of view, is a very good candidate to run a survey analogous to the one discussed in Section 3.4.3, though with a higher flux limit.

Acknowledgments

The authors would like to thank Masami Ouchi and Kazuaki Ota for sharing their Ly α LF observations. Also, a thanks to Steven L Finkelstein for sharing his conglomerated UV LF data, a very useful resource. Lastly, a very special thanks to Micaela Bagley, who provided useful discussions at key points throughout this project. S.B. and C.S. were partially supported by a Jet Propulsion Lab/NASA grant to the University of Minnesota (RSA 1677370/NNN12AA01C). M.H. is supported by the Knut and Alice Wallenberg Foundation. C.M. acknowledges support by NASA Headquarters through a NASA Hubble Fellowship grant No. HST-HF2-51413.001-A awarded by the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Inc., for NASA, under contract NAS5-26555, and support by the VILLUM FONDEN under grant No. 37459. The Cosmic Dawn Center (DAWN) is funded by the Danish National Research Foundation under grant No. DNR140.

Software: Astropy (Robitaille et al. 2013), SciPy (Oliphant 2007), NumPy (van der Walt et al. 2011), Matplotlib (Hunter 2007), emcee (Foreman-Mackey

et al. 2013), `pymc3` (Salvatier et al. 2016), and `arviz` (Kumar et al. 2019).



Appendix Luminosity Function Analytical Fit

We assume the following 2D functional form, which reproduces the logs of the modeled Ly α LFs to 4% accuracy with a median -0.1% difference across all LAE LF bins and neutral fractions:

$$\log_{10}(\Phi) = \beta_0 + \beta_1 x + \beta_2 y + \beta_3 x^2 + \beta_4 x^2 y + \beta_5 x^2 y^2 + \beta_6 y^2 + \beta_7 x y^2 + \beta_8 x y, \quad (\text{A1})$$

where x is the log of the Ly α luminosity and y is $\bar{x}_{\text{H I}}$. The LFs are in units of $\text{cMpc}^{-3} \log L^{-1}$ and the luminosities are in units of erg s^{-1} . The β coefficients are reported in Table 1 and are unitless. The fit is valid in the range $0.01 < \bar{x}_{\text{H I}} < 0.91$ and $42 < \log(L_{\text{Ly}\alpha}/[\text{erg s}^{-1}]) < 44$.

ORCID iDs

Sean Bruton  <https://orcid.org/0000-0002-6503-5218>
 Claudia Scarlata  <https://orcid.org/0000-0002-9136-8876>
 Francesco Haardt  <https://orcid.org/0000-0003-3291-3704>
 Matthew J. Hayes  <https://orcid.org/0000-0001-8587-218X>
 Charlotte Mason  <https://orcid.org/0000-0002-3407-1785>
 Alexa M. Morales  <https://orcid.org/0000-0003-4965-0402>
 Andrei Mesinger  <https://orcid.org/0000-0003-3374-1772>

References

- Atek, H., Richard, J., Jauzac, M., et al. 2015, *ApJ*, 814, 69
 Bagley, M. B., Scarlata, C., Henry, A., et al. 2017, *ApJ*, 837, 11
 Becker, G. D., D’Aloisio, A., Christenson, H. M., et al. 2021, *MNRAS*, 508, 1853
 Behroozi, P., Wechsler, R., Hearin, A., & Conroy, C. 2019, *MNRAS*, 488, 3143
 Bolton, J. S., & Haehnelt, M. G. 2013, *MNRAS*, 429, 1695
 Bosman, S. E. I., Davies, F. B., Becker, G. D., et al. 2022, *MNRAS*, 514, 55
 Bouwens, R. J., Illingworth, G. D., Oesch, P. A., et al. 2011, *ApJ*, 737, 90
 Bouwens, R. J., Illingworth, G. D., Oesch, P. A., et al. 2015a, *ApJ*, 811, 140
 Bouwens, R. J., Illingworth, G. D., Oesch, P. A., et al. 2015b, *ApJ*, 803, 34
 Bouwens, R. J., Oesch, P. A., Stefanon, M., et al. 2021, *AJ*, 162, 47
 Bowler, R. A. A., Dunlop, J. S., McLure, R. J., et al. 2014, *MNRAS*, 440, 2810
 Castellano, M., Fontana, A., Boutsia, K., et al. 2010, *A&A*, 511, A20
 Davies, F. B., Hennawi, J. F., Bañados, E., et al. 2018, *ApJ*, 864, 142
 De Barros, S., Pentericci, L., Vanzella, E., et al. 2017, *A&A*, 608, A123
 Dijkstra, M. 2014, *PASA*, 31, e040
 Dijkstra, M. 2017, arXiv:1704.03416
 Dijkstra, M., Wyithe, S., Haiman, Z., Mesinger, A., & Pentericci, L. 2014, *MNRAS*, 440, 3309
 Finkelstein, S. L. 2016, *PASA*, 33, e037
 Finkelstein, S. L., D’Aloisio, A., Paardekooper, J.-P., et al. 2019, *ApJ*, 879, 36
 Finkelstein, S. L., Ryan, R. E., Jr., Papovich, C., et al. 2015, *ApJ*, 810, 71
 Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, *PASP*, 125, 306
 Gehrels, N. 1986, *ApJ*, 303, 336
 Greig, B., Mesinger, A., & Bañados, E. 2019, *MNRAS*, 484, 5094
 Hayes, M. J., Runnholm, A., Gronke, M., & Scarlata, C. 2021, *ApJ*, 908, 36
 Hoag, A., Bradač, M., & Huang, K. 2019, *ApJ*, 878, 12
 Hu, E. M., Cowie, L. L., Barger, A. J., et al. 2010, *ApJ*, 725, 394
 Hu, E. M., Cowie, L. L., Songaila, A., et al. 2016, *ApJL*, 825, L7
 Hu, W., Wang, J., Zheng, Z.-Y., et al. 2019, *ApJ*, 886, 90
 Hunter, J. D. 2007, *CSE*, 9, 90
 Itoh, R., Ouchi, M., Zhang, H., et al. 2018, *ApJ*, 867, 46
 Jung, I., Finkelstein, S. L., Livermore, R. C., et al. 2018, *ApJ*, 864, 103
 Kashikawa, N., Shimasaku, K., Matsuda, Y., et al. 2011, *ApJ*, 734, 119
 Konno, A., Ouchi, M., Ono, Y., et al. 2014, *ApJ*, 797, 16
 Konno, A., Ouchi, M., Shibuya, T., et al. 2018, *PASJ*, 70, S16
 Kumar, R., Carroll, C., Hartikainen, A., & Martin, O. 2019, *JOSS*, 4, 1143
 Livermore, R. C., Finkelstein, S. L., & Lotz, J. M. 2017, *ApJ*, 835, 113
 Malhotra, S., & Rhoads, J. 2004, *ApJL*, 617, L5
 Malhotra, S., & Rhoads, J. 2006, *ApJL*, 647, L95
 Mason, C., Trenti, M., & Treu, T. 2015, *ApJ*, 813, 21
 Mason, C. A., Fontana, A., Treu, T., et al. 2019, *MNRAS*, 485, 3947
 Mason, C. A., & Gronke, M. 2020, *MNRAS*, 499, 1395
 Mason, C. A., Treu, T., de Barros, S., et al. 2018a, *ApJL*, 857, L11
 Mason, C. A., Treu, T., Dijkstra, M., et al. 2018b, *ApJ*, 856, 2
 McLure, R. J., Dunlop, J. S., Bowler, R. A. A., et al. 2013, *MNRAS*, 432, 2696
 McQuinn, M., Lidz, A., Zahn, O., et al. 2007, *MNRAS*, 377, 1043
 McQuinn, M., Lidz, A., Zaldarriaga, M., Hernquist, L., & Dutta, S. 2008, *MNRAS*, 388, 1101
 Mesinger, A., Aykutalp, A., Vanzella, E., et al. 2015, *MNRAS*, 446, 566
 Mesinger, A., & Furlanetto, S. 2007, *ApJ*, 669, 663
 Mesinger, A., & Furlanetto, S. 2008a, *MNRAS*, 385, 1348
 Mesinger, A., & Furlanetto, S. 2008b, *MNRAS*, 386, 1990
 Mesinger, A., Furlanetto, S., & Cen, R. 2011, *MNRAS*, 411, 955
 Mesinger, A., Greig, B., & Sobacchi, E. 2016, *MNRAS*, 459, 2342
 Meyer, R. A., Laporte, N., Ellis, R. S., Verhamme, A., & Garel, T. 2021, *MNRAS*, 500, 558
 Morales, A. M., Mason, C. A., Bruton, S., et al. 2021, *ApJ*, 919, 120
 Naidu, R. P., Tacchella, S., Mason, C. A., et al. 2020, *ApJ*, 892, 109
 Oliphant, T. E. 2007, *CSE*, 9, 10
 Ono, Y., Ouchi, M., Harikane, Y., et al. 2018, *PASJ*, 70, S10
 Ota, K., Iye, M., Kashikawa, N., et al. 2017, *ApJ*, 844, 85
 Ouchi, M., Harikane, Y., Shibuya, T., et al. 2018, *PASJ*, 70, S13
 Ouchi, M., Mobasher, B., Shimasaku, K., et al. 2009, *ApJ*, 706, 1136
 Ouchi, M., Shimasaku, K., Furusawa, H., et al. 2010, *ApJ*, 723, 869
 Pentericci, L., Fontana, A., Vanzella, E., et al. 2011, *ApJ*, 743, 132
 Planck-Collaboration, Ade, P. A. R., Aghanim, N., et al. 2016, *A&A*, 594, A13
 Planck-Collaboration, Aghanim, N., Akrami, Y., et al. 2020, *A&A*, 641, A6
 Qin, Y., Mesinger, A., Bosman, S. E. I., & Viel, M. 2021, *MNRAS*, 506, 2390
 Ren, K., Trenti, M., & Mason, C. A. 2019, *ApJ*, 878, 114
 Rhoads, J. E., & Malhotra, S. 2001, *ApJL*, 563, L5
 Robertson, B. E., Ellis, R. S., Furlanetto, S. R., & Dunlop, J. S. 2015, *ApJL*, 802, L19
 Robitaille, T. P., Tollerud, E. J., Greenfield, P., et al. 2013, *A&A*, 558, A33
 Runnholm, A., Gronke, M., & Hayes, M. 2021, *PASP*, 133, 034507
 Salvatier, J., Wiecki, T. V., & Fonnesbeck, C. 2016, *PeerJ Comput. Sci.*, 2, e55
 Schenker, M. A., Robertson, B. E., Ellis, R. S., et al. 2013, *ApJ*, 768, 196
 Songaila, A., Hu, E. M., Barger, A. J., et al. 2018, *ApJ*, 859, 91
 Stark, D. P., Ellis, R. S., Chiu, K., Ouchi, M., & Bunker, A. 2010, *MNRAS*, 408, 1628
 Tilvi, V., Papovich, C., Tran, K.-V. H., et al. 2013, *ApJ*, 768, 56
 Treu, T., Schmidt, K. B., Trenti, M., Bradley, L. D., & Stiavelli, M. 2013, *ApJL*, 775, L29
 van der Walt, S., Colbert, S. C., & Varoquaux, G. 2011, *CSE*, 13, 22
 Verhamme, A., Garel, T., Ventou, E., et al. 2018, *MNRAS Lett.*, 478, L60
 Whittler, L. R., Mason, C. A., Ren, K., et al. 2020, *MNRAS*, 495, 3602
 Wold, I. G. B., Malhotra, S., Rhoads, J., et al. 2022, *ApJ*, 927, 36
 Yoshioka, T., Kashikawa, N., Inoue, A. K., et al. 2022, *ApJ*, 927, 32
 Zheng, Z.-Y., Wang, J., Rhoads, J., et al. 2017, *ApJL*, 842, L22