



Future Constraints on the Reionization History and the Ionizing Sources from Gamma-Ray Burst Afterglows

Adam Lidz¹ , Tzu-Ching Chang^{2,3}, Lluís Mas-Ribas^{2,3} , and Guochao Sun³ ¹ Department of Physics & Astronomy, University of Pennsylvania, 209 South 33rd Street, Philadelphia, PA 19104, USA; alidz@sas.upenn.edu² Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA 91109, USA³ California Institute of Technology, 1200 East California Boulevard, Pasadena, CA 91125, USA

Received 2021 May 5; revised 2021 June 9; accepted 2021 June 11; published 2021 August 17

Abstract

We forecast the reionization history constraints, inferred from Ly α damping wing absorption features, for a future sample of ~ 20 $z \geq 6$ gamma-ray burst (GRB) afterglows. We describe each afterglow spectrum by a three-parameter model. First, L characterizes the size of the ionized region (the “bubble size”) around a GRB host halo. Second, $\langle x_{\text{H I}} \rangle$ is the volume-averaged neutral fraction outside of the ionized bubble around the GRB, which is approximated as spatially uniform. Finally, $N_{\text{H I}}$ denotes the column density of a local damped Ly α absorber (DLA) associated with the GRB host galaxy. The size distribution of ionized regions is extracted from a numerical simulation of reionization and evolves strongly across the epoch of reionization (EoR). The model DLA column densities follow the empirical distribution determined from current GRB afterglow spectra. We use a Fisher matrix formalism to forecast the $\langle x_{\text{H I}}(z) \rangle$ constraints that can be obtained from follow-up spectroscopy of afterglows with $S/N = 20$ per $R = 3000$ resolution element at the continuum. We find that the neutral fraction may be determined to better than 10%–15% (1σ) accuracy from this data across multiple independent redshift bins at $z \sim 6$ –10, spanning much of the EoR, although the precision degrades somewhat near the end of reionization. A more futuristic survey with 80 GRB afterglows at $z \geq 6$ can improve the precision here by a factor of 2 and extend measurements out to $z \sim 14$. We further discuss how these constraints may be combined with estimates of the escape fraction of ionizing photons derived from the DLA column density distribution toward GRBs extracted at slightly lower redshift, $z \sim 5$. This combination will help in testing whether we have an accurate census of the sources that reionized the universe.

Unified Astronomy Thesaurus concepts: [Reionization \(1383\)](#); [Early universe \(435\)](#)

1. Introduction

Gamma-ray bursts (GRBs) are, electromagnetically, the brightest explosions in the universe. The long-duration class of GRBs is thought to be powered by the collapse of massive rotating stars (Woosley 1993); as nuclear fuel is exhausted at the end of the star’s life, the core loses pressure support and implodes to form a black hole. In this process, jets of energetic particles are produced, which plow through the remaining stellar envelope. This relativistic outflow leads to highly beamed gamma-ray emission, while subsequent afterglows occur at longer wavelengths as the ejecta runs into surrounding gas in the circumburst and interstellar media (ISM). The incredible luminosity of these events makes long-duration GRBs detectable from even some of the earliest phases of cosmic history. Indeed, long-duration GRBs have been observed back into the epoch of reionization (EoR; Totani et al. 2006; Salvaterra et al. 2009; Tanvir et al. 2009; Cucchiara et al. 2011; Chornock et al. 2013; Melandri et al. 2015; Tanvir et al. 2018), when early generations of galaxies and/or accreting black holes emitted ionizing photons and gradually filled the universe with ionized gas (e.g., Loeb & Furlanetto 2013). Therefore, GRB afterglows can be exploited as bright backlights, with follow-up spectroscopy of the afterglows at high spectral resolution and signal-to-noise ratio (S/N) providing valuable constraints on the neutral fraction of the intergalactic medium (IGM; e.g., Totani et al. 2006; Hartoog et al. 2015) and the chemical enrichment history of the ISM in the GRB host galaxy (e.g., Lamb & Reichart 2000; Savaglio 2006).

A number of projects are being proposed to study reionization and metal enrichment using GRB afterglows. These include the

Gamow Explorer, which is expected to find roughly ~ 20 GRB afterglows in the EoR at $z \geq 6$ using an X-ray telescope with a wide field of view, exploiting that the GRB emission at high redshifts is largely shifted into the X-ray band (Ghirlanda et al. 2015; White 2020). The Gamow Explorer will then determine photometric redshifts for these sources using an onboard infrared telescope from the drop in flux blueward of Ly α at the source redshift. This will enable rapid targeted follow-up spectroscopic campaigns using the James Webb Space Telescope (JWST) in space or existing 8 m and future 30 m class telescopes from the ground. The Gamow Explorer is proposed to launch in 2028 with a nominal 3 yr mission lifetime. The detectable GRB rate during the EoR from the Gamow Explorer should improve on the Swift telescope’s (Gehrels et al. 2004) detection rate of such events by roughly an order of magnitude. In addition, HiZ-GUNDAM, which is similar in scope and timescale to the Gamow Explorer, is being planned in Japan (Yonetoku et al. 2014). In the slightly longer term, a larger European-led satellite, Theseus, is being proposed for launch in 2032 and should detect between 30 and 80 $z \geq 6$ GRBs (Amati et al. 2018; Stratta et al. 2018; Amati et al. 2021).

Motivated by these future prospects, the aim of the current paper is to forecast the reionization history constraints expected from these upcoming missions. The key spectral signature of the remaining neutral gas in the IGM is the damping wing of the Ly α line owing to its natural line width (Miralda-Escude 1998). This leads to spectrally extended absorption redward of Ly α at the source redshift, provided that significantly neutral gas—with a neutral hydrogen fraction of order unity—remains in the IGM at the GRB redshift.

In this regard, GRB afterglows offer a few potential advantages over quasars (Miralda-Escude 1998; Barkana & Loeb 2004), which have nevertheless provided our best backlights for studying the IGM thus far. First, the intrinsic unabsorbed spectrum of a GRB afterglow has a simple power-law form around Ly α at the source redshift. In contrast, in the case of quasars, the damping wing signature needs to be extracted in the presence of a Ly α emission line (e.g., Davies et al. 2018a). Second, the quasars themselves can ionize significant regions of neutral hydrogen around them. Furthermore, these sources likely reside in highly biased locations where surrounding galaxies form earlier than in typical regions of the universe. Reionization may hence occur earlier in quasar environments than in average parts of the universe (Lidz et al. 2007). Together, the ionizing impact of quasars and surrounding galaxies may lead to weaker damping wing signatures around quasars than around less biased sources such as long-duration GRBs, which are expected to trace more typical regions of massive star formation. This may make the damping wing signature more difficult to detect around quasar sources and also complicate the interpretation of quasar-based measurements that require modeling the relationship between the quasar environments and typical parts of the IGM. Third, the quasar luminosity function falls very steeply with increasing redshift, and current forecasts suggest that only a single quasar (with UV absolute magnitude $M_{\text{AB}} \leq -26$) is detectable in the observable universe at $z \gtrsim 9$ (Fan et al. 2019). The GRB afterglows provide the best candidates for observing damping wing signatures in the early phases of the EoR.

The GRB afterglow measurements are not without their challenges, however. First, one needs to detect the GRBs and their afterglows, rapidly determine photometric redshifts, and obtain high-quality follow-up spectroscopic observations of the fading afterglow. Second, current samples of GRB afterglows frequently show strong damped Ly α (DLA) absorption from neutral hydrogen in the ISM of the host galaxy (e.g., Tanvir et al. 2019). In some cases, the host DLA can swamp the signature of damping wing absorption from neutral gas in the IGM (Miralda-Escude 1998). Third, a challenge for interpreting the damping wing signatures of the IGM toward both GRB afterglows and quasar backlights is the patchiness of the reionization process (McQuinn et al. 2008; Mesinger & Furlanetto 2008). This leads to a significant sight line-to-sight line scatter in the neutral fraction toward these background sources. More importantly, it implies that the damping wing profile is at least partly sensitive to the precise distribution of neutral gas along each line of sight, and it is unclear how to best model the resulting spectra.

Here we consider a simple parameterization (see also, e.g., McQuinn et al. 2008) that partly accounts for the patchiness of the reionization process, includes the plausible impact of host DLAs, and allows for rapid, fully analytic forecasts of the reionization history constraints obtainable from future GRB afterglow spectral samples. At each redshift, afterglow spectra are described by a bubble size parameter, L ; a neutral fraction, $\langle x \rangle_{\text{H I}}$ (approximated as uniform outside the host bubble); and a host DLA column density, $N_{\text{H I}}$. The bubble sizes and column densities are drawn from probability distributions, inferred respectively from reionization simulations and observations, to allow for the likely broad sight line-to-sight line scatter in these quantities. In this framework, we use the Fisher matrix formalism to determine the neutral fraction constraints across

several redshift bins, marginalized over the bubble size and column density parameters.

In addition to providing backlights for studying the IGM damping wings, GRBs also provide valuable information about the escape fraction of ionizing photons (Chen et al. 2007; Tanvir et al. 2019) and the star formation rate density (SFRD) at high redshifts. Specifically, a plausible working assumption is that GRB afterglows provide lines of sight toward massive stars that are representative of the pathways traversed by ionizing photons as they propagate out into the IGM from the sources of reionization. The fraction of afterglows that have host columns less than roughly $\lesssim 10^{18} \text{ cm}^{-2}$, corresponding to Lyman continuum optical depths at the photoionization threshold of $\tau_{\text{H I}} \lesssim 6.3$, then provides an estimate of the escape fraction.⁴ Further, if the conversion between GRB rate and SFRD can be calibrated accurately at lower redshift, the redshift evolution of the GRB rate can be used to determine the SFRD during the EoR, including contributions from lower-mass host galaxies that are challenging to observe directly with other star formation indicators (e.g., Blain & Natarajan 2000; Kistler et al. 2008; Robertson & Ellis 2012; Vergani et al. 2015). Measurements of the neutral fraction can be combined with estimates of the escape fraction and internal or external determinations of the SFRD to address whether the neutral fraction constraints are consistent with known source populations. We consider this in the context of our Fisher matrix analyses.

The outline of this paper is as follows. In Section 2, we review the IGM damping wing signature and the impact of DLA host absorbers. We then describe the key ingredients of our model: the redshift distribution of GRB afterglows (Section 2.1), the reionization history model (Section 2.2), the bubble size distribution (Section 2.3), and the probability distribution of host column densities (Section 2.4). Section 3 describes our Fisher matrix calculations and characterizes parameter degeneracies, while Section 4 forecasts the resulting neutral fraction constraints. In this section, we further discuss the dependence of these constraints on the quality of follow-up spectroscopic observations and the possibilities for obtaining the requisite spectra. Section 5 discusses the prospects for determining the escape fraction of ionizing photons and the resulting implications. We conclude in Section 6.

2. Method and Model Parameters

We start by reviewing the physics of the Ly α damping wing, which is the key feature for probing the IGM neutral fraction with GRB afterglows (Miralda-Escude 1998). The Ly α optical depth close to line center is large even in the case of a highly ionized universe at the redshifts of interest, as evident in Equation (2) below. Consequently, the traditional Ly α forest blueward of Ly α at the source redshift is expected to be highly opaque at $z \gtrsim 6$ or so (as also indicated by current observations; e.g., Fan et al. 2006; Yang et al. 2020), regardless of the precise ionization state of the IGM. If significantly neutral gas—with a neutral hydrogen fraction of order unity—remains in the IGM, however, the broad damping wing of the Ly α line, owing to its natural line width, leads to an extended absorption feature redward of Ly α at the source redshift. In contrast, the

⁴ This is just a rough estimate. A more detailed determination of the escape fraction computes $\langle \exp(-\tau_{\text{H I}}) \rangle$, averaged over the sight lines in the sample and frequencies above the photoionization threshold, for some assumed ionizing spectral shape.

probability of absorption in the wing is negligible if the IGM is instead highly ionized.

In what follows, we ignore the impact of peculiar velocities, consider only the wing of the Ly α line (i.e., we do not treat absorption in the Doppler core), and approximate the gas density by its cosmic mean value. These should be good approximations for our goal of modeling the extended damping wing redward of line center. Further, we adopt the approximation to the line profile in Miralda-Escude (1998), which includes classical Rayleigh scattering and the Lorentzian form for the natural line width (see Lee 2013 for refinements, which we will neglect here, as they should be unimportant for our forecasts).

In this case, consider a significantly neutral portion of the IGM with a volume-averaged neutral fraction, $\langle x_{\text{HI}} \rangle$, between redshift z_{beg} and z_{end} along the line of sight to a GRB afterglow at some higher source redshift, z_s . For simplicity, in our modeling, we place each GRB at the center of its redshift bin and consider bins of width $\Delta z = 1$, as described further below. The optical depth in the damping wing of the Ly α line from this neutral patch—at wavelength offset $\Delta\lambda$ from Ly α at the source redshift—is denoted by $\tau_{\text{DW}}(\Delta\lambda)$. Hence, the observed wavelength is $\lambda_{\text{obs}} = \lambda_\alpha (1 + z_s) + \Delta\lambda$, where $\lambda_\alpha = 1215.67 \text{ \AA}$ is the rest-frame wavelength of the Ly α transition. We further denote the fractional wavelength offset as $\delta = \Delta\lambda / [\lambda_\alpha(1 + z_s)]$. The damping wing optical depth is then (see Miralda-Escude 1998; we mostly follow the notation in that work)

$$\tau_{\text{DW}}(\Delta\lambda) = \frac{\tau_{\text{GP}}(z_s)}{\pi} R_\alpha \langle x_{\text{HI}} \rangle (1 + \delta)^{3/2} \times [I(x_{\text{beg}}) - I(x_{\text{end}})]. \quad (1)$$

Here R_α is a dimensionless constant related to the natural line width of the Ly α line, with $R_\alpha = \Lambda_\alpha / (4\pi\nu_\alpha)$, where Λ_α is the Ly α decay constant and ν_α is the rest-frame frequency of the Ly α transition. Numerically, $R_\alpha = 2.02 \times 10^{-8}$. Here τ_{GP} denotes the Gunn–Peterson optical depth of fully neutral gas at the cosmic mean density (Gunn & Peterson 1965). The Gunn–Peterson optical depth is

$$\tau_{\text{GP}} = 4.0 \times 10^5 \left[\frac{1 + z_s}{7} \right]^{3/2} \left[\frac{\Omega_b h^2}{0.0225} \right] \left[\frac{\Omega_m h^2}{0.132} \right]^{-1/2} \times \left[\frac{1 - Y}{0.76} \right], \quad (2)$$

where Ω_m and Ω_b are the usual cosmological matter and baryon density parameters, h is the Hubble parameter in units of $100 \text{ km s}^{-1} \text{ Mpc}^{-1}$, and $1 - Y$ is the hydrogenic mass fraction. Finally, $x_{\text{beg}} = (1 + z_{\text{beg}}) / [(1 + z_s)(1 + \delta)]$, $x_{\text{end}} = (1 + z_{\text{end}}) / [(1 + z_s)(1 + \delta)]$, and $I(x)$ is an integral with

$$I(x_{\text{beg}}) - I(x_{\text{end}}) = \int_{x_{\text{end}}}^{x_{\text{beg}}} dx \frac{x^{9/2}}{(1 - x)^2}. \quad (3)$$

In what follows, we imagine that each GRB host is located within an ionized region, so z_{beg} is set by the distance between the GRB host halo and the first significantly neutral region encountered along the line of sight to the observer. As discussed further below, this distance is drawn from a distribution based on numerical simulations of reionization. Next, we approximate the exterior neutral fraction as uniform outside of the ionized bubble around the GRB host. Our baseline model further adopts the average neutral fraction $\langle x_{\text{HI}} \rangle$

at the center of the redshift bin in question and includes contributions only from neutral gas within the redshift bin. In other words, z_{end} in Equations (1) and (3) is set by the lower-redshift end of the redshift bin under consideration. Put differently, we neglect redshift evolution in the neutral fraction across each redshift bin and ignore contributions from neutral gas at lower redshifts outside of the bin. Although the precise assumptions here are arbitrary, and we will test their impact in what follows, note that the damping wing is fairly insensitive to neutral gas at redshifts much smaller than the source redshift (see Appendix B and Figure 15).

Throughout, we assume perfect knowledge of the GRB host redshift. We suppose that the GRB redshifts may be determined accurately through measuring the redshifts of host galaxy metal lines detected in absorption, perhaps associated with a local DLA (e.g., Totani et al. 2006). Alternatively, follow-up observations after the afterglow fades can be used to determine the redshift of the host galaxy using emission lines (e.g., Krühler et al. 2015).

2.1. Assumed Redshift Distribution

The first key ingredient in our forecasts is the assumed redshift distribution of GRB afterglows. Provided the GRB rate traces the SFRD, the number of GRBs per unit redshift may be written as (Kistler et al. 2008)

$$\frac{dN_{\text{GRB}}}{dz} = N_0 \frac{\dot{\rho}_*}{1 + z} \frac{dV_{\text{co}}}{dz}, \quad (4)$$

where $\dot{\rho}_*$ denotes the SFRD, i.e., the star formation rate per comoving volume per unit time; dV_{co}/dz is the comoving volume per unit redshift; and N_0 is a normalization constant (with units of observing time). The factor of $1 + z$ in the denominator accounts for time dilation. We adopt the empirical fitting formula for the SFRD from Madau & Dickinson (2014). We ignore any redshift dependence in the ability of the upcoming surveys to detect the GRBs and determine their redshifts in this equation, which should be a good approximation for the bright bursts targeted by the Gamow Explorer and related missions.

Our fiducial model fixes the normalization constant by requiring $N_{\text{GRB}}(\geq z = 6) = \int_6^\infty dz dN_{\text{GRB}}/dz = 20$. This is based on detailed projections of the number of high-redshift GRBs detectable with the Gamow Explorer (White 2020), including models for the GRB luminosity function, prompt emission and afterglow spectral shape, observing time, and Gamow instrumental sensitivity (Ghirlanda et al. 2015). The normalization of $dN_{\text{GRB}}/dz = 20$ reflects the expected number of GRBs detected over the nominal 3 yr lifetime of the Gamow mission. In order to explore some of the uncertainties involved, we also consider a model with an enhanced GRB rate at higher redshifts; in this alternate scenario, we multiply the left-hand side of Equation (4) by $(1 + z)^{1.5}$, normalized to the same dN_{GRB}/dz at $z = 6$. This case yields 31 GRBs at $z \geq 6$ for a 3 yr mission. This is intended to roughly represent scenarios where the SFRD estimates from Madau & Dickinson (2014) are incomplete at $z \gtrsim 6$, owing, for example, to star formation in faint galaxies or possible evolution in the GRB to star formation rate relationship. Finally, in order to explore the longer-term best-case scenario prospects, we also consider a case with $N_{\text{GRB}}(z \geq 6) = 80$ GRBs. This is at the upper end of

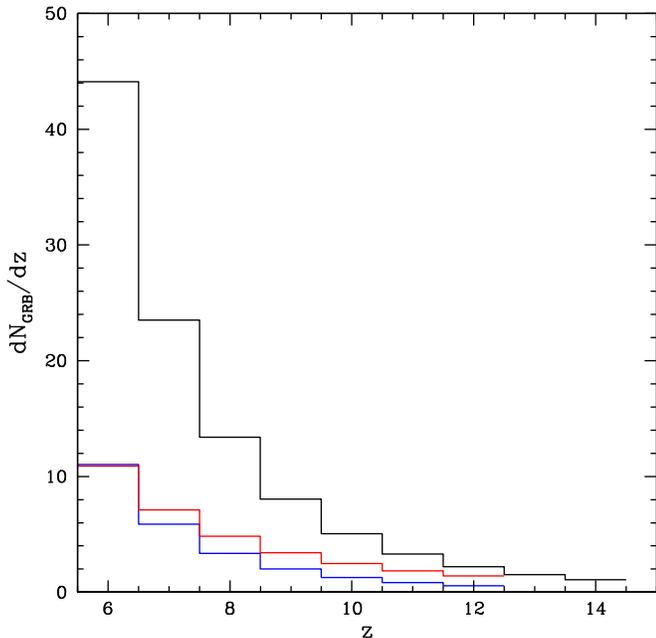


Figure 1. Assumed redshift distributions for the rate of long-duration GRBs during the EoR. The blue histogram shows our fiducial model in which the GRB rate traces the SFRD and the total number of GRBs above $z = 6$ is normalized to 20, as forecast for the Gamow Explorer over its nominal 3 yr mission lifetime. The red histogram shows an alternate case in which the GRB rate is enhanced at high redshift (see text), which is also considered in what follows. Finally, the black histogram shows the upper end of the predicted GRB rate for the planned Theseus mission, again for a 3 yr mission.

the range of 30–80 $z \geq 6$ GRBs forecast for the Theseus mission (Amati et al. 2018) over 3 yr.

The GRB redshift distributions in each of these three models are shown in Figure 1 for redshift bins of width $\Delta z = 1$. In what follows, we adopt the nearest integer to the average afterglow count forecast in each bin. In the blue and red histogram models, 11 afterglows lie in the redshift bin centered around $z_s = 6$, and at least one GRB is detectable out to the $z_s = 10$ bin in the fiducial case and $z_s = 12$ in the more optimistic scenario. In our fiducial reionization history, described below, these redshifts span most of the reionization process (see Figure 2) and should push past the highest redshift that will be probed using quasars, even with futuristic surveys, as mentioned in the Introduction. In the longer term, the optimistic upper end of the Theseus projections gives 44 GRBs in our $z_s = 6$ redshift bin, five in the $z_s = 10$ bin, and one, on average, in the $z_s = 14$ bin.⁵

2.2. Fiducial Reionization History

Our fiducial reionization history model is based on an approximate photon-counting description of reionization in which the volume-averaged ionization fraction is determined through the following equation (Shapiro & Giroux 1987;

⁵ Note that the total number of GRBs in our histogram exceeds 80 for Theseus because the lower end of our redshift $z_s = 6$ bin extends to $z_s = 5.5$. In our fiducial model, described in the next section, reionization completes at a redshift slightly below $z = 6$, as may, in fact, be suggested by observations (e.g., Lidz et al. 2007; Mesinger 2010; Malloy & Lidz 2015; Kulkarni et al. 2019; Keating et al. 2020).

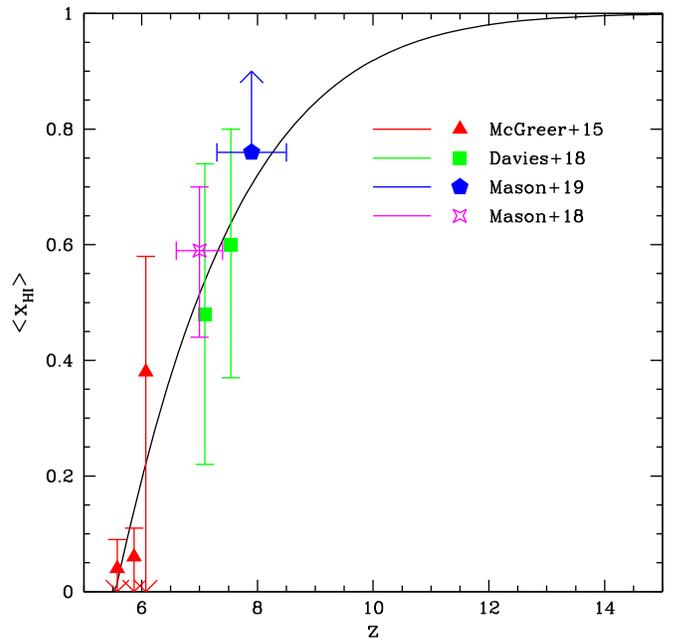


Figure 2. Fiducial model reionization history compared to some current observational constraints. The solid black curve shows our fiducial theoretical model described in the text below. The red triangles show upper limits on the neutral fraction from the dark pixel fraction in the Ly α and Ly β forests (McGreer et al. 2015). The green squares show estimates of the neutral fraction from damping wing measurements toward quasars (Davies et al. 2018b). Finally, the blue pentagon and magenta star, respectively, show a lower limit and a measurement of the neutral fraction using observations of Ly α emission lines toward Lyman-break galaxies (Mason et al. 2018, 2019). In all cases, the error bars are 1σ confidence estimates.

Madau et al. 1999):

$$\frac{d\langle x_i \rangle}{dt} = \zeta \frac{df_{\text{coll}}(>M_{\text{min}})}{dt} - \frac{\langle x_i \rangle}{\bar{t}_{\text{rec}}}. \quad (5)$$

Here the first term describes the rate of ionizing photon production, which is assumed to be proportional to the time derivative of the halo collapse fraction, above a minimum galaxy host halo mass, M_{min} . The second term accounts for recombinations, and \bar{t}_{rec} gives the average time for ionized gas to recombine, which may be written as $\bar{t}_{\text{rec}} = C/[\alpha_B(T)\langle n_e(z) \rangle]$, where C is a clumping factor, α_B is the case-B recombination coefficient of hydrogen, and $\langle n_e(z) \rangle$ is the proper number density of free electrons at redshift z . Our fiducial model adopts an ionizing efficiency parameter, $\zeta = 20$, a minimum galaxy host halo mass of $M_{\text{min}} = 10^9 M_\odot$, a gas temperature of $T = 2 \times 10^4$ K, and a clumping factor of $C = 3$ (see, e.g., Lidz 2016 and references therein for a discussion of these parameters). The halo collapse fraction is calculated using the Sheth–Tormen halo mass function (Sheth et al. 2001). The ionization fraction is fixed to unity below the redshift where the solution of Equation (5) first reaches this limiting value. The volume-averaged neutral fraction is determined through $\langle x_{\text{HI}}(z) \rangle = 1 - \langle x_i(z) \rangle$.

The resulting ionization history is shown in Figure 2. We emphasize that the fiducial model here is intended only as a plausible baseline for our ensuing forecasts and has not been designed or tuned to match current constraints in detail. It is nevertheless in good agreement with current measurements, as illustrated in the figure. Furthermore, the electron-scattering optical depth in our fiducial model is $\tau_e = 0.51$ (assuming helium

is singly ionized along with hydrogen and neglecting the likely small contribution from the second ionization of helium), consistent with recent Planck 2018 measurements; the fiducial model τ_e value differs from their TT, TE, EE, lowE, lensing + BAO constraint by less than 1σ (Aghanim et al. 2018).

2.3. Bubble Size Distribution

The next important ingredient in our model is to account for the ionized bubbles around GRB host halos. In general, these regions reflect the combined ionizing output of many surrounding galaxies, with the typical size scale growing as ionized bubbles expand and merge over the course of the reionization process (e.g., McQuinn et al. 2007b). In order to appreciate the role of the ionized bubbles here, consider the example of a GRB that occurs in a large ionized region. More precisely, the relevant thing here is the extent of the ionized skewer traversed along the line of sight from the GRB to the observer before encountering neutral gas in the IGM.⁶ In this case, photons initially just redward of Ly α at the source redshift, for instance, will redshift significantly before encountering neutral gas, and the source will show only weak damping wing absorption. In contrast, GRB afterglow photons originating in small bubbles or at the edge of larger ones will redshift less before running into neutral gas, and the damping wing feature will be stronger. Therefore, the size of the ionized bubbles around GRB host halos plays an important role in determining the strength of the damping wing absorption.

In order to account for this in our model, we turn to the reionization simulations of McQuinn et al. (2007a, 2008), which treat the radiative transfer of ionizing photons in a postprocessing step on top of an evolved N -body simulation of cosmological structure formation. From a number of simulation snapshots spanning much of the EoR, we extract sight lines in random directions toward dark matter halos with masses in the range $10^{10} M_\odot \leq M \leq 10^{11} M_\odot$ and record the comoving distance between each simulated dark matter halo and the first significantly neutral cell (defined by when the cell-averaged neutral fraction exceeds 0.9).⁷ These halos are well resolved in the simulations and plausible hosts for GRBs during the EoR. The reionization history in the simulations differs a little from our fiducial model, so we record the first-crossing distance as a function of the volume-averaged neutral fraction rather than redshift. In each GRB redshift bin, we then determine the mean ionized skewer size and the (1σ) scatter from the simulation output with an average neutral fraction closest to the fiducial model value at the desired redshift. Relying on simulation outputs at the correct neutral fraction but slightly different redshifts should be a good approximation; previous studies have shown that the spatial structure of reionization is sensitive mainly to the average ionization fraction and fairly independent of the precise redshift when a given stage of reionization is reached (McQuinn et al. 2007b). This first-crossing distance then sets the upper limit of the integral over redshift, z_{beg} , in Equations (1) and (3).⁸

⁶ We will use the terms “ionized bubble size” and “ionized skewer size” interchangeably, although “skewer” is a more appropriate term for the present application given that GRB afterglows provide a 1D probe through the IGM.

⁷ McQuinn et al. (2008) considered the same basic distribution but chose a different criterion for the first neutral cell and host halo mass range. Their results are hence slightly different than those here.

⁸ Toward the end of reionization in our model, the comoving distance to the first neutral region, which sets z_{beg} , may occasionally exceed the distance from the bin center to the lower-redshift edge of a bin, which sets z_{end} . In this case, the line of sight is considered to traverse entirely ionized gas and leads to negligibly small damping wing absorption in the model.

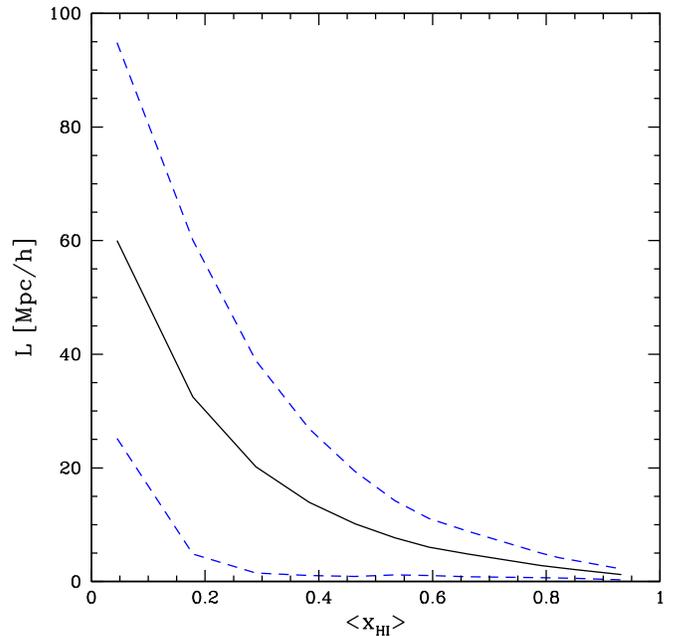


Figure 3. Mean bubble (or ionized skewer) size and its scatter—as a function of the volume-averaged neutral fraction—extracted from reionization simulations. The black solid line gives the average comoving distance between a GRB host and the first neutral region, while the blue dashed lines indicate the sight line-to-sight line scatter around the mean. The mean and scatter grow sharply as reionization proceeds and the average neutral fraction drops.

To briefly reiterate, the main approximation in our scheme is to ignore spatial variations in the neutral fraction outside of the first ionized bubble. We expect this to be most accurate in the early stages of reionization, when the neutral fraction is large and the size of typical ionized bubbles is small relative to the extent of the damping wing. Especially toward the end of reionization, a more elaborate parameterized model involving the size of the first neutral region, the extent of the second ionized bubble encountered, etc. may be required for more accurate results. We confine our current study, however, to the simpler three-parameter model.

Figure 3 shows the resulting distribution of ionized skewer sizes. The mean skewer size (in comoving units) evolves very strongly during the EoR, ranging from roughly ~ 1 Mpc/h when the universe is $\sim 95\%$ neutral, to ~ 10 Mpc/h when it is $\sim 50\%$ neutral, and to ~ 60 Mpc/h when about $\sim 5\%$ of the IGM volume is neutral. The scatter also grows strongly as reionization proceeds.

The size distribution of ionized skewers may be conveniently approximated by a lognormal distribution.⁹ Specifically,

$$\frac{dP_L}{dL} = \frac{1}{\sqrt{2\pi\bar{\sigma}^2}} \frac{1}{L} \exp\left[-\frac{(\ln L - \mu)^2}{2\bar{\sigma}^2}\right], \quad (6)$$

where L is the bubble size in units of Mpc/h. The lognormal parameters μ and $\bar{\sigma}$ are related to the mean ionized skewer size,

⁹ The lognormal distribution has been frequently used in the literature to describe the related 3D bubble size distribution (e.g., Dvorkin & Smith 2009). While we have not explicitly compared the lognormal PDF to the ionized skewer size distribution in the simulations, it matches the first two simulated moments by construction, and the simulated distribution is skew positive. We therefore expect a lognormal PDF to provide a good description and adopt it in what follows.

$\langle L \rangle$, and sight line-to-sight line scatter, σ_L , via

$$\mu = \ln \left[\frac{\langle L \rangle^2}{(\sigma_L^2 + \langle L \rangle^2)^{1/2}} \right] \quad (7)$$

and

$$\tilde{\sigma}^2 = \ln \left[\left(\frac{\sigma_L}{\langle L \rangle} \right)^2 + 1 \right]. \quad (8)$$

2.4. Contamination from Host DLAs

In addition to absorption from neutral hydrogen in the IGM, we must account for absorption from neutral hydrogen in the host galaxy. Empirically, as discussed further below, many GRB sight lines show strong DLA absorption. We assume that the DLA lies precisely at the redshift of the GRB host galaxy, which is likely an excellent approximation given that long-duration GRBs are expected to trace star-forming regions, and most star formation occurs in relatively small mass halos at the redshifts of interest. For instance, the circular velocity of a $10^{10} M_\odot$ halo at $z_s = 8$ is $v_{\text{circ}} = 74 \text{ km s}^{-1}$ (Barkana & Loeb 2001), so a DLA moving at the halo circular velocity relative to the host galaxy is offset in wavelength by only $\Delta\lambda/\lambda \sim v_{\text{circ}}/c \sim 2 \times 10^{-4}$, comparable to the spectral resolution assumed in what follows (Section 3). This is hence a negligible effect, although a DLA in a neighboring halo could, in principle, be more of a concern; an unlikely coincidence would, however, be required for the neighboring absorber to produce a large column density.

Since we are interested in the absorption in the wing of the line, we approximate the optical depth of the DLA at wavelength offset $\Delta\lambda$ using a Lorentzian profile,

$$\tau_{\text{DLA}}(\Delta\lambda) = N_{\text{HI}} \frac{\sigma_\alpha R_\alpha}{\pi (\Delta\lambda/\lambda)^2 + R_\alpha^2}, \quad (9)$$

where $\sigma_\alpha = 3\lambda_\alpha^2 \Lambda_\alpha / (8\pi\nu_\alpha) = 4.48 \times 10^{-18} \text{ cm}^2$ is the frequency-integrated cross section. The transmission through the IGM accounting for the IGM damping wing and the DLA is then given from Equations (1) and (9) as

$$f = e^{-\tau} = e^{-(\tau_{\text{DW}} + \tau_{\text{DLA}})}. \quad (10)$$

We will explore the redshift and parameter dependence of the transmission profiles in more detail in the next section, but Figure 4 gives an illustrative example for starters. Here we consider a GRB source at $z_s = 8$, in which case the IGM is 70% neutral in our fiducial model (Figure 2). The mean size of the ionized skewers at this neutral fraction is $\langle L \rangle = 4.8 \text{ Mpc}/h$, while the sight line-to-sight line dispersion is $\sigma_L = 4.0 \text{ Mpc}/h$ (Figure 3). In this example, the figure illustrates that the DLA dominates the absorption close to line center for the average ionized skewer size, but ionized skewers with size 1σ lower than the mean value produce comparable absorption at line center to the DLA. However, for any skewer size within the $\pm 1\sigma$ range around the mean in this example, we expect the IGM contribution to the damping wing to dominate over the DLA component far from line center (for $N_{\text{HI}} = 10^{20} \text{ cm}^{-2}$). That is, the IGM component, scaling roughly as $\tau_{\text{DW}} \propto 1/(\Delta\lambda)$ (Equation (1)),¹⁰ is more extended than the DLA contribution,

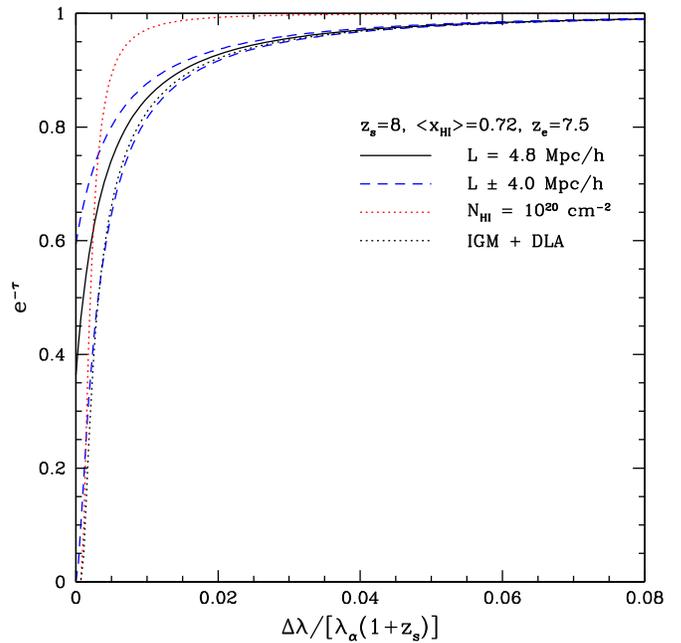


Figure 4. Example damping wing profile. The black solid line shows the IGM damping wing in our fiducial model at $z_s = 8$ for the mean ionized skewer size, while the blue dashed lines show the profiles for skewers of length $\pm 1\sigma$ above and below the mean. In each case, $z_{\text{end}} = 7.5$ (see text). The red dotted line shows the profile of a DLA at the host redshift with a column density of $N_{\text{HI}} = 10^{20} \text{ cm}^{-2}$. The black dotted line shows the transmission including both the DLA optical depth and the IGM contribution for the mean skewer size. The more gradual IGM profile is distinguishable from the DLA profile for the parameters shown, although the spatial variations in the ionized skewer size lead to fairly strong variations in the transmission close to line center.

which goes like $\tau_{\text{DLA}} \propto 1/(\Delta\lambda)^2$ (Equation (9)). As fleshed out further in Section 3 below, larger column density DLAs can nevertheless swamp the IGM damping wing contribution, especially at lower redshifts, where the neutral fraction is smaller, the universe is less dense, and the ionized skewers from the GRB host halo tend to be larger.

A key, yet uncertain, issue then regards the probability distribution of DLA column densities from local host absorption around GRBs during the EoR. The largest empirical sample currently addressing this question is published in Tanvir et al. (2019); these authors determined the column density distribution from a sample of 138 GRB afterglows ranging in redshift from $1.6 < z < 6.7$. The probability distribution from the tabulated data in this study is shown in Figure 5. Interestingly, Tanvir et al. (2019) derived an average ionizing photon escape fraction of only $\langle f_{\text{esc}} \rangle = 0.005$, with a 98% upper limit of $\langle f_{\text{esc}} \rangle \leq 0.015$. Indeed, the median column density in this sample is $N_{\text{HI}} = 10^{21.59} \text{ cm}^{-2}$, relatively few sight lines have $N_{\text{HI}} \leq 10^{20} \text{ cm}^{-2}$, and only a tiny fraction of sight lines have column densities low enough to allow appreciable ionizing photon escape fractions.

An important question, then, is how strongly the probability distribution of DLA column densities evolves with redshift. The current sample includes only seven lines of sight above $z > 5$, and while there is currently no evidence for redshift evolution in the column density probability distribution function (PDF; Tanvir et al. 2019), a larger data set will be required to test this more sharply. In fact, provided this distribution is indeed a reliable indicator of the escape fraction of ionizing photons, one can indirectly argue that the distribution must evolve strongly; otherwise, it is hard to understand how the universe could be

¹⁰ This follows from the small δ limit of Equation (1), provided that z_{end} and z_{beg} are close to z_s .

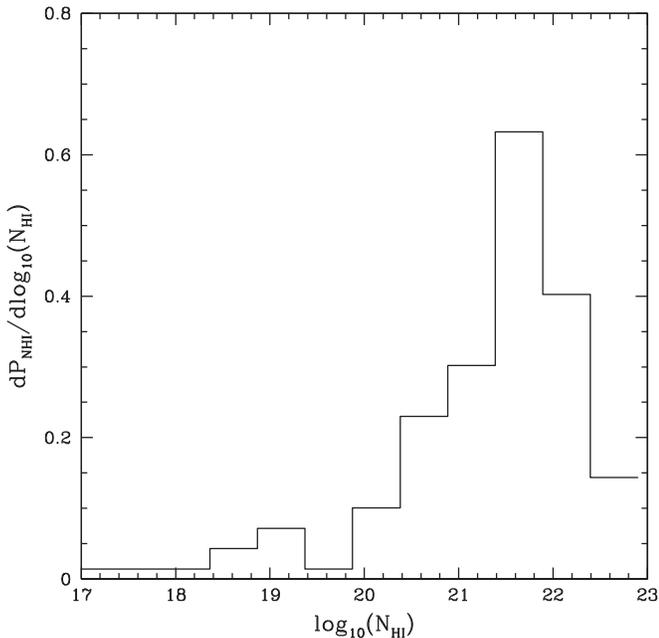


Figure 5. Probability distribution of DLA column density from the current sample of 138 GRB afterglows at $1.6 < z < 6.7$. This is based on the data tabulated in Tanvir et al. (2019).

reionized by $z \lesssim 5.5$ (as is observed) given the escape fraction estimate in Tanvir et al. (2019). For instance, escape fractions larger than roughly $f_{\text{esc}} \gtrsim 0.1$ are required to reionize the universe with star-forming galaxies, subject somewhat to uncertainties in faint-end extrapolations down the galaxy UV luminosity function, the ionizing spectral shape, and the clumping factor of the IGM (e.g., Bouwens et al. 2016). This is a factor of ~ 20 larger than the best-fit value from the current GRB samples, which are, however, concentrated at lower redshifts. The strong evolution implied here is at least broadly consistent with inferences of the escape fraction at $z \sim 2.5\text{--}5$ based on comparing galaxy UV luminosity functions to the ionizing emissivity extracted using the transmission through the Ly α forest at these redshifts, along with estimates of the mean free path of ionizing photons (Faucher-Giguere et al. 2008; Faucher-Giguere 2020). Nevertheless, our fiducial model adopts the current DLA column density PDF, since it is hard to predict how this may evolve toward higher redshift. We should keep in mind that this may be a pessimistic assumption (see also Section 4.1).

3. Fisher Matrix Formalism

In order to forecast the constraints on the ionization history, we adopt a Fisher matrix approach. The Fisher matrix calculations assume a quadratic expansion around the logarithm of the likelihood function, which in turn peaks for a set of specified fiducial model parameters. As shown in Appendix A, we can incorporate the expected broad distributions in the ionized skewer size and DLA column density PDFs by first computing a conditional Fisher matrix for sight lines with specified bubble size L and column density N_{HI} . We then find the expected error bar on the neutral fraction for that L and N_{HI} by inverting the conditional Fisher matrix. Finally, the error bar on the neutral fraction for an ensemble of GRB afterglow sight lines can be determined by computing the average of the neutral fraction’s inverse variance, weighted by the bubble size and column density PDFs (see

Equations (A9) and (A10)). The error bar in each redshift bin is scaled with $1/\sqrt{N_{\text{GRB}}(z_s)}$, where $N_{\text{GRB}}(z_s)$ is the number of independent afterglow spectra expected in the bin centered around source redshift z_s . For each sight line, we adopt a parameter vector \mathbf{q} with three components ($q_{x_{\text{HI}}}$, q_L , $q_{N_{\text{HI}}}$). Here each component of \mathbf{q} describes the fractional variation around the fiducial model value. For example, $q_{x_{\text{HI}}} = (\langle x_{\text{HI}} \rangle - \langle x_{\text{HI, fid}} \rangle) / \langle x_{\text{HI, fid}} \rangle$. Note also that the parameters q_L and $q_{N_{\text{HI}}}$ vary in meaning depending on the value of L and N_{HI} for each sight line.

The conditional Fisher matrix for a sight line with bubble size L and column density N_{HI} may be computed as

$$\tilde{F}_{ij}(L, N_{\text{HI}}) = \sum_k \frac{\partial f(k|L, N_{\text{HI}})}{\partial q_i} \frac{\partial f(k|L, N_{\text{HI}})}{\partial q_j} \times \frac{1}{\text{var}[f(k)]}. \quad (11)$$

Here the sum is over spectral pixels, which are taken to be uniformly spaced in $\ln(\lambda)$ with $\Delta\lambda/\lambda = 1/3000$; i.e., we consider a spectral resolution of $R = \lambda/\Delta\lambda = 3000$. We sum over all pixels out to 1500 Å redward (observed frame) of Ly α at the source redshift. Appendix B demonstrates that our results are insensitive to this choice (see Figure 16). In general, we assume that the noise in each afterglow spectrum is set by a combination of Poisson fluctuations in the photons from the source itself and a contribution from the sky background. In this case,

$$\text{var}[f(k)] = \frac{N_c f(k) + N_{\text{sky}}}{N_c^2}, \quad (12)$$

where N_c is the number of photons per spectral resolution element received from the source at the continuum, and N_{sky} is the same for the sky background. We adopt a flat continuum spectrum, which should be a very good approximation across the range of wavelengths considered in Equation (11).¹¹ The average number of photon counts in spectral pixel k is then given by $N_c f(k)$; i.e., the observed photon count is the number of photons at the continuum multiplied by the transmission. The S/N squared at the continuum is hence $S/N^2 = N_c^2 / (N_c + N_{\text{sky}})$.

Our baseline assumption is that the S/N at the continuum per spectral pixel is $S/N = 20$, and that $N_{\text{sky}} \gg N_c$. That is, we assume that the observations are in the sky background-dominated limit, where the source counts make a negligible contribution to the transmission variance. In this case, $\text{var}[f(k)] = N_{\text{sky}}/N_c^2 = 1/\text{SNR}^2$. We discuss the possibilities for obtaining our baseline S/N value in Section 4.2 and show that the sky background-dominated limit applies for ground-based observations but that space-based observations with the JWST will be in the source-dominated limit. We compare the sky background-limited and source-dominated cases in Appendix B and Figure 17.

We consider a redshift-independent S/N for the purposes of having a single baseline noise value; this helps us understand the overall requirements for spectroscopic follow-up observations. In practice, the sky background is a strongly increasing function of wavelength, so in practice, it will be hard to achieve the requisite S/N from the ground for higher-redshift bursts (see Section 4.2). However, in the source-dominated case

¹¹ Throughout, we assume that the unabsorbed continuum level is determined perfectly; i.e., we neglect continuum fitting errors in our analysis.

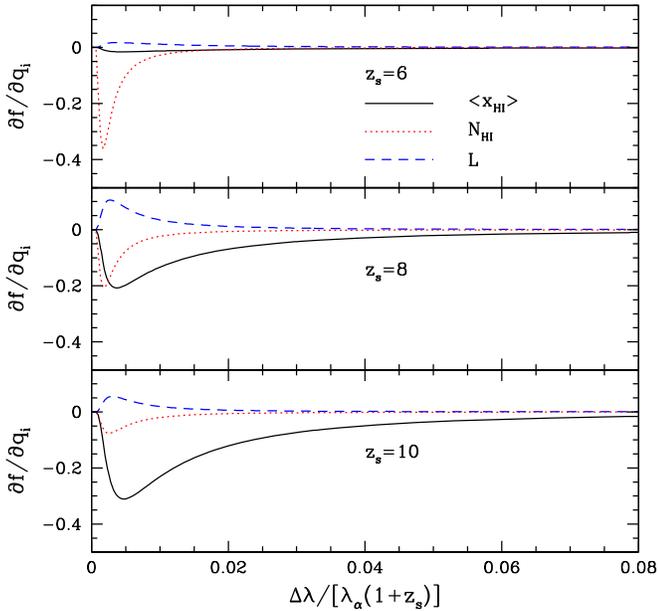


Figure 6. Examples of the derivatives that enter our Fisher matrix computations. Top panel: transmission derivatives at $z_s = 6$ for a fiducial model with $\langle x_{\text{HI}} \rangle = 0.21$, $L = \langle L \rangle = 32$ Mpc/h, and $N_{\text{HI}} = 10^{20}$ cm $^{-2}$. The black solid line shows the derivative with respect to $q_{x_{\text{HI}}}$, the red dotted line is the $q_{N_{\text{HI}}}$ derivative, and the blue dashed line gives the q_L derivative. Middle panel: same as top panel but for $z_s = 8$, $\langle x_{\text{HI}} \rangle = 0.72$, $L = \langle L \rangle = 4.8$ Mpc/h, and $N_{\text{HI}} = 10^{20}$ cm $^{-2}$. Bottom panel: same as upper two panels, but here we consider $z_s = 10$, $\langle x_{\text{HI}} \rangle = 0.92$, $L = \langle L \rangle = 1.2$ Mpc/h, and $N_{\text{HI}} = 10^{20}$ cm $^{-2}$.

applicable to JWST space-based observations, a redshift-independent S/N may be a decent approximation. This is the case because, given the afterglows observed at a fixed time interval after the explosion, higher-redshift ones are seen at an earlier time in the source frame—owing to cosmological time dilation—when they are brighter, compensating for the dimming from the increased luminosity distance (e.g., Barkana & Loeb 2004). We will consider variations around the S/N assumptions in Section 4.2.

3.1. Parameter Derivatives

It is instructive to first examine some examples of the transmission derivatives with respect to model parameters, as enter into Equation (11). Specifically, Figure 6 shows derivatives with respect to the neutral fraction, bubble size, and DLA column density parameters for the case of our fiducial reionization history with $L = \langle L \rangle$ and $N_{\text{HI}} = 10^{20}$ cm $^{-2}$ at each of $z_s = 6, 8,$ and 10 (top, middle, and bottom panels, respectively). In general, increasing the DLA column density leads to reduced transmission relatively close to the redshift of the source, while increasing the bubble size parameter has the opposite effect, as expected. Increasing the neutral fraction reduces the transmission, and the impact of the neutral fraction parameter is more extended in wavelength than either of the DLA or bubble size parameters, owing to the gradual falloff in the IGM damping wing profile of Equation (1). The fairly similar shape but opposite sign of the derivatives with respect to the bubble size and column density parameters suggests that there will be a fairly strong degeneracy direction between these parameters. That is, the enhanced transmission with increasing bubble size can be compensated for by raising the DLA column density.

Turning to the redshift dependence in Figure 6, one can see that the neutral fraction parameter becomes more important than the DLA and bubble size parameters as one moves to higher redshift. This results because the strength of the IGM damping wing grows as the universe becomes more neutral and denser toward higher redshift and as the bubble sizes shrink. Note further that the transmission changes less with increasing DLA column density toward high redshift because the derivatives are taken around a fiducial transmission profile, which is more absorbed at high redshift. For the example DLA column density and bubble sizes, the IGM damping wing signature is swamped at $z_s \sim 6$ but becomes more prominent at $z_s \gtrsim 8$. On the other hand, we expect to detect more GRB afterglows near $z_s \sim 6$ than at $z_s \sim 8$ – 10 (see Figure 1), so there is a trade-off between the weaker damping wing signatures toward low redshift and the more abundant data samples available there. Also, we should keep in mind that the IGM damping wing signature is more detectable toward sight lines with smaller ionized bubbles and lower column density DLAs (see Section 4).

3.2. Parameter Degeneracies

In order to quantify the ionization history constraints in detail, we must account for the full probability distributions in bubble size and DLA column density. First, however, it is helpful to consider parameter degeneracies for particular realizations of the bubble size and DLA column density. That is, we can invert the conditional Fisher matrix of Equation (11) and inspect the resulting parameter ellipses for particular values of L and N_{HI} . For simplicity of illustration, we suppose here that all of the sight lines in the redshift bin of interest have the specified bubble size and DLA column density. Here we take our fiducial GRB redshift distribution with 20 $z \geq 6$ GRBs (Figure 1).

Figure 7 shows an example case at $z_s = 6$ with $L = \langle L \rangle = 32$ Mpc/h and $N_{\text{HI}} = 10^{20}$ cm $^{-2}$. The first conclusion one can draw from this figure is that the fractional constraints on the DLA column density in this example are much tighter than those forecast for either the bubble size or the neutral fraction. This results because the derivative with respect to the DLA column density parameter is much stronger than the bubble size and neutral fraction parameter derivatives here, as previously illustrated in the top panel of Figure 6.¹² The strongest degeneracy direction here is between ionized skewer size and neutral fraction, since one can trade off for an increasing neutral fraction by boosting the bubble size. Similarly, the column density and bubble size parameters are positively correlated. Naively, the positive correlation between DLA column density and neutral fraction seen in the top left panel of the figure is surprising, but this results because the bubble size is also an important parameter in the problem; one can compensate for an increase in both the neutral fraction and DLA column density by increasing the bubble size. Indeed, if we fix the bubble size parameter entirely, we find the opposite degeneracy direction between $q_{N_{\text{HI}}}$ and $q_{x_{\text{HI}}}$, as expected. In spite of the weak effect from the IGM damping wing in this case, one would still expect an interesting upper bound on the neutral fraction after accumulating the 11 GRB afterglow

¹² Note that we do not incorporate a “physicality prior” on any of the parameters, although regions of parameter space with $q_i < -1$ are unphysical. This leads to conservative bounds.

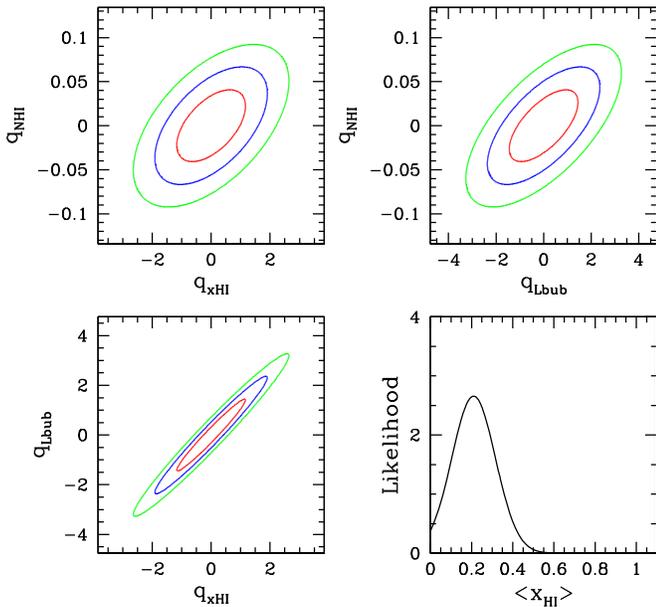


Figure 7. Parameter degeneracies for the simplified case that all of the sight lines in the $z_s = 6$ bins have bubble size $L = \langle L \rangle$ and $N_{\text{HI}} = 10^{20} \text{ cm}^{-2}$. The ellipses show 1σ (red), 2σ (blue), and 3σ (green) confidence intervals. The bottom right panel gives the resulting 1D likelihood, marginalized over bubble size and DLA column density parameters, for the average neutral fraction parameter in this redshift bin. We refine this error bar estimate in the next section by accounting for the full distribution of ionized skewer sizes and DLA column densities, but this simplified case nevertheless suffices for understanding the degeneracies involved.

spectra in this redshift bin, as illustrated in the bottom right panel of Figure 7. This panel shows the 1D likelihood for $\langle x_{\text{HI}} \rangle$ after marginalizing over the DLA column density and bubble size parameters. This estimate needs to be refined, however, to incorporate the DLA and skewer size distributions.

As an aside, we note that metal absorption lines—detected slightly redward of the damping wing feature in the same observations—may help in breaking the degeneracies shown. In particular, low-ionization metal lines should trace the same cold neutral gas in the ISM of the GRB host galaxy that produces the DLA contamination. Specifically, Figure 11 of Mas-Ribas et al. (2017) shows a strong trend between the equivalent width of low-ionization metal absorption lines and DLA column densities, as determined from stacking BOSS quasar spectra. Therefore, measurements of the equivalent widths of the low-ionization metal lines may provide independent constraints on the DLA column density parameter toward each GRB afterglow. Although we do not consider this further in what follows, it is an interesting direction for future work.

Figures 8 and 9 show analogous plots at $z_s = 8$ and 10, respectively. The main change relative to Figure 7 is that the neutral fraction parameter now has a stronger impact than either of the DLA and bubble size parameters. This reflects the increasing neutral fraction and smaller bubble sizes in our model at these redshifts, which make the IGM damping wing more prominent (see the middle and bottom panels of Figure 6). Even though we expect fewer GRB afterglows in these redshift bins than at $z_s \sim 6$ (for our fiducial GRB redshift distribution, there are three at $z_s = 8$ and one at $z_s = 10$; Figure 1), we expect tighter constraints on the neutral fraction. Indeed, the neutral fraction constraints in this simplified case (bottom right panels of Figures 8 and 9) look promising. In principle, the bubble size parameter itself, along with its

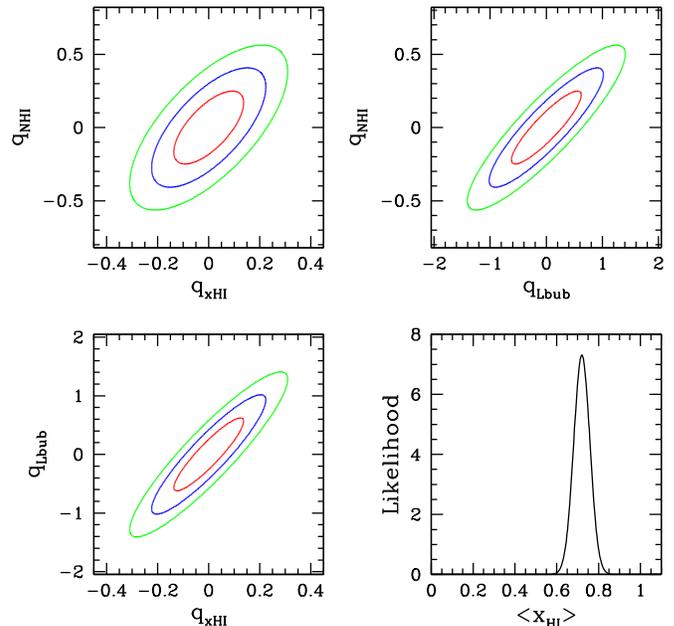


Figure 8. Identical to Figure 7 but for the $z_s = 8$ redshift bin.

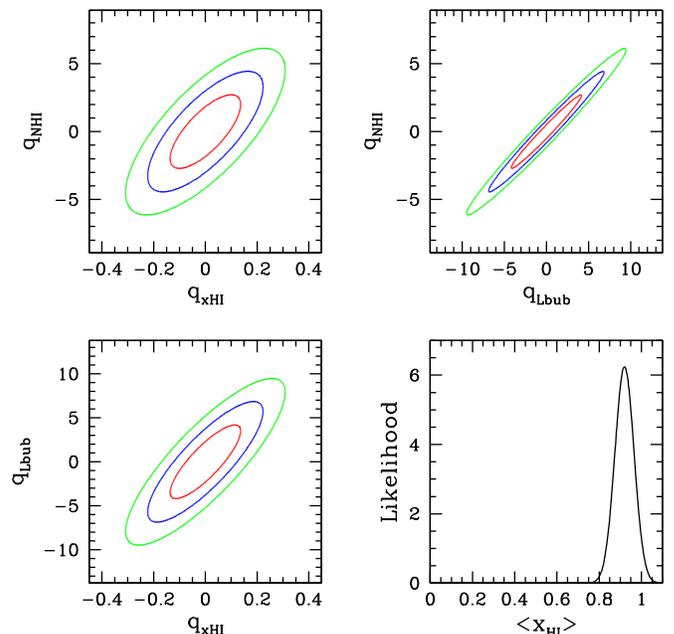


Figure 9. Identical to Figures 7 and 8 but for the $z_s = 10$ redshift bin.

distribution and redshift evolution, provides valuable information about reionization. However, the error bar forecasts on this parameter are fractionally large (Figures 7–9), so we do not consider the reionization information contained in this parameter further in what follows.

4. Full Ionization History Forecasts

We next compute full ionization history forecasts, assuming the bubble size distribution of Equations (6)–(8) and the DLA column density PDF of Figure 5. The required computations here are described by Equations (A9) and (A10) of Appendix A and amount to averaging the inverse variance of the neutral fraction over the L, N_{HI} PDFs.

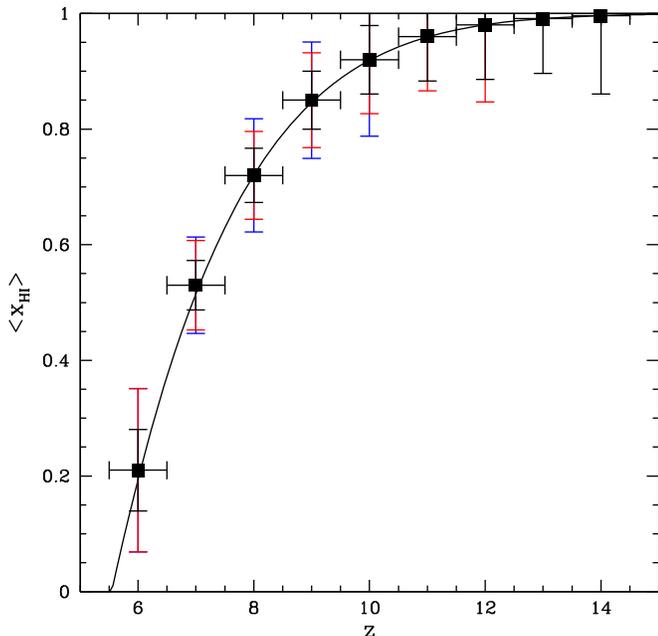


Figure 10. Forecast constraints on the reionization history of the universe. The blue 1σ error bars adopt the fiducial GRB redshift distribution of Figure 1, while the red error bars show the more optimistic case from the same figure. Finally, the black points and error bars show the constraints for the Theseus-inspired ~ 80 $z \geq 6$ GRBs sample. Each of these results incorporates the ionized bubble size distribution and DLA column density PDFs discussed in the text. Each afterglow spectrum is assumed to have an $S/N = 20$ per $R = 3000$ resolution element at the continuum. Note that in the red/blue models, the error bars are identical in the $z_s = 6$ redshift bin, since these cases give an identical GRB count in this bin. As discussed in Section 2.1, the number of GRBs assumed here corresponds to estimates for the nominal 3 yr lifetimes of the upcoming missions.

The results of these calculations are shown in Figure 10. These appear quite encouraging; provided the spectroscopic follow-up observations we considered are feasible, a future GRB mission like the Gamow Explorer should help to determine the redshift evolution of the neutral fraction during cosmic reionization in some detail. Fractionally, the error bar on the neutral fraction is largest toward the end of reionization; in the $z_s = 6$ bin, we forecast a fractional error of 70% at 1σ . In the other redshift bins, the error bars are in the 10%–15% range in our fiducial model. The absolute errors range from $\sigma_{x_{\text{HI}}} = 0.08$ to 0.14 with the largest values in the first and last redshift bins. The error bars are still smaller in the case of the optimistic GRB redshift distribution model, and two additional redshift bins are populated in this case—out to $z_s = 12$ —at which point the fiducial neutral fraction is around $\langle x_{\text{HI}} \rangle = 0.98$.

Finally, the black points and error bars show the more futuristic case, in which the Theseus mission obtains high- S/N follow-up spectra for 80 $z \geq 6$ GRB afterglows. In this scenario, the statistical precision of the error bars is further improved, with the error bar forecasts approaching 5% level constraints across many redshift bins in the EoR, degrading mainly toward low redshift (where the damping wing feature is typically somewhat feeble) and at very high redshift, where there are few GRBs. Nevertheless, a project like this could potentially determine the neutral fraction from $z_s \sim 6$ to 14 with good statistical precision.

These forecasts consider the average expected number of GRBs in each redshift bin across the nominal 3 yr mission

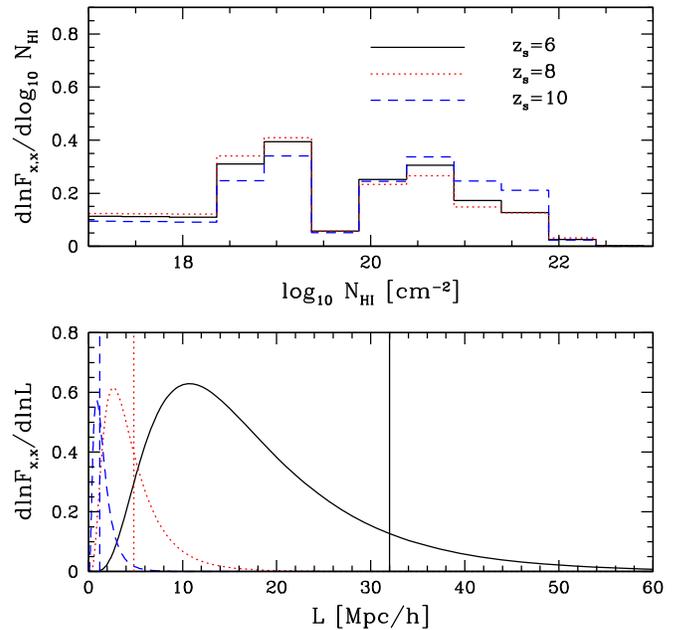


Figure 11. Fisher information vs. DLA column density and ionized skewer size. Top panel: contributions to the neutral fraction Fisher information of Equation (A9) as a function of DLA column density at each of $z_s = 6, 8,$ and 10 . The bimodal nature arises because the N_{HI} PDF is peaked at high column densities (see Figure 5), but the sight lines with low column density systems contain more information about the neutral fraction. The histograms are somewhat noisy because they rely on empirical estimates of the N_{HI} PDF. Bottom panel: same for the bubble size. The vertical lines indicate the average bubble size at each redshift. The Fisher information here is dominated by sight lines with bubble sizes smaller than the cosmic mean, with the results turning over as small bubbles become rare. Here $F_{x,x}$ in the legend is shorthand for $F_{(x_{\text{HI}}), (x_{\text{HI}})}$.

lifetimes of the upcoming missions. If these lifetimes can be extended, then the error bars in each redshift bin will shrink as one over the square root of the number of GRBs in the bin. In addition, higher-redshift bins will become populated, on average. For example, if the Gamow Explorer mission lifetime can be doubled, then the expected neutral fraction error in each bin would go down by about 30%, and measurements out to $z = 12$ would become possible, on average.

4.1. Contributions to the Fisher Information

In order to unpack these results a little further, it helps to examine the integrand of the Fisher information for the neutral fraction in Equation (A9). Specifically, Figure 11 plots both $d \ln F_{(x_{\text{HI}}), (x_{\text{HI}})} / d \ln L$ (bottom panel) and $d \ln F_{(x_{\text{HI}}), (x_{\text{HI}})} / d \log_{10} N_{\text{HI}}$ (top panel) as a function of L and $\log_{10} N_{\text{HI}}$, respectively. Put differently, the figure shows how much of the neutral fraction inverse variance is contributed by sight lines with various bubble sizes and column densities.

The bottom panel illustrates that ionized regions smaller than the mean skewer size dominate the Fisher information at each redshift. The information grows toward small bubble sizes because the damping wing is stronger for small bubbles and then turns over as sufficiently small ionized regions become rare. This scale grows as reionization proceeds.

Likewise, the top panel shows that the Fisher information peaks for rare sight lines with low column density DLAs, although there is still appreciable information in the bins with $20 \lesssim \log_{10}(N_{\text{HI}}/\text{cm}^{-2}) \lesssim 22$. This reflects a trade-off between the fact that sight lines with lower column densities contain more

information about the neutral fraction but are rarer than the high column density sight lines closer to the peak of the fiducial N_{HI} distribution (Figure 5), which tend to swamp the IGM damping wing signature. The histograms in Figure 11 are somewhat noisy because we rely on an empirical model for the N_{HI} PDF. For instance, the dip just below $\log_{10}(N_{\text{HI}}/\text{cm}^{-2}) = 20$ reflects a corresponding drop in the observed N_{HI} PDF (Figure 5).

Another important implication of this figure is that the constraint forecasts may improve appreciably if the DLA column density PDF evolves strongly with redshift, shifting toward lower column densities, as may be expected if this distribution is directly related to the ionizing photon escape fraction (see Section 2.4). For example, the actual Fisher information in the column density bin around $\log_{10}(N_{\text{HI}}/\text{cm}^{-2}) = 18.6$ is a factor of ~ 7 larger than the information in the bin centered on $\log_{10}(N_{\text{HI}}/\text{cm}^{-2}) = 20.6$; they only make comparable contributions to $d \ln F_{\langle x_{\text{HI}} \rangle, \langle x_{\text{HI}} \rangle} / d \log_{10} N_{\text{HI}}$ because the lower column density bin here is on the rare tail of the assumed N_{HI} PDF. Note that a factor of ~ 7 gain in the Fisher information translates into a smaller error bar on the neutral fraction by a factor of $\sim \sqrt{7.8} = 2.8$, so this could be significant.

4.2. Dependence on Spectral S/N

Here we consider the prospects for obtaining spectral S/Ns of ~ 20 per $R = 3000$ resolution element at the continuum, as adopted in the calculations of Figure 10. We caution that our goal here is only to obtain rough estimates and explore the dependence on source redshift, flux density, sky background, and observing time. More detailed calculations will be required to cement the observing strategy and improve on our estimates, which are likely accurate only to within a factor of ~ 2 in S/N. We also quantify the impact of varying the S/N around the fiducial value of $S/N = 20$. To estimate the expected S/N of an observation, we first calculate the number of photons received by a telescope in each spectral resolution element from the source and/or night-sky background:

$$N_{c,\text{sky}} = \frac{f_{\nu}^{c,\text{sky}}}{h_p} \epsilon \frac{\pi D_{\text{tel}}^2 \Delta \lambda}{4 \lambda} t_{\text{obs}}. \quad (13)$$

Here N_c denotes the number of photon counts from the source at the continuum, while f_{ν}^c is the source flux density at the continuum and an observed wavelength of $\lambda_{\text{obs}} = \lambda_{\alpha}(1 + z_s)$.¹³ The quantities N_{sky} and f_{ν}^{sky} denote, respectively, the number of photons and the flux density received from the sky background. In this equation, we approximate the telescope aperture as circular, so that the collecting area is $\pi D_{\text{tel}}^2/4$, where D_{tel} is the diameter of the telescope, ϵ is an efficiency factor to account for photon losses, $\Delta \lambda/\lambda$ is the size of the spectral pixels, t_{obs} is the observing time, and h_p is Planck's constant. We further denote $N_{\text{sky}} = \alpha N_c$, noting that $\alpha \gg 1$ for ground-based observations.

In order to estimate the sky background for ground-based observations, we use the Very Large Telescope (VLT) X-Shooter exposure time calculator, which returns the specific intensity of the sky background in different frequency bands.¹⁴ Specifically,

this gives (59, 340, 650) $\mu\text{Jy arcsec}^{-2}$ in the (I , Y , J) bands, which, respectively, include Ly α from source redshifts of $z_s = (6, 7-8, 9-10)$. These values then determine the flux density of the sky background, further assuming $1''$ spatial pixels for ground-based observations. It should be possible to improve on this seeing-limited case using adaptive optics, but we do not consider this further in what follows. The sky background is dominated by atmospheric emission, or ‘‘airglow,’’ and is a strong function of wavelength; this makes the highest-redshift GRB measurements challenging from the ground. From space, with the JWST, the sky background is expected to be dominated by zodiacal light (see below), and such observations will be in the source-dominated regime, i.e., $N_c \gg N_{\text{sky}}$, $\alpha \sim 0$.

Using Equation (13), we can estimate the spectral S/N (squared), accounting for Poisson fluctuations in the source and sky background photon counts as

$$(S/N)^2 = \frac{1}{\alpha + 1} \frac{f_{\nu}^c}{h_p} \epsilon \frac{\pi D_{\text{tel}}^2 \Delta \lambda}{4 \lambda} t_{\text{obs}}. \quad (14)$$

We neglect readout noise and dark current. Plugging in some characteristic numbers, we find

$$(S/N)^2 = 20^2 \left[\frac{f_{\nu}^c}{10 \mu\text{Jy}} \right] \left[\frac{35}{1 + \alpha} \right] \left[\frac{\epsilon}{0.25} \right] \left[\frac{\Delta \lambda/\lambda}{1/3000} \right] \times \left[\frac{D_{\text{tel}}}{8\text{m}} \right]^2 \left[\frac{t_{\text{obs}}}{6.1 \text{ hr}} \right]. \quad (15)$$

That is, we find that an 8 m telescope can achieve an S/N of 20 on a 10 μJy GRB afterglow in $t_{\text{obs}} = 6.1$ hr of observing time at a spectral resolution of $R = 3000$. Here we use $\alpha = 34$, appropriate for ground-based observations in the Y band of a source at $z_s \sim 8$, where $f_{\nu}^{\text{sky}} = 340 \mu\text{Jy}$, as determined from the VLT exposure time calculator in footnote 14.¹⁵ In other words, an S/N of 20 could be achieved on a 10 μJy afterglow at $z_s = 8$ from the ground on an 8 m telescope in a single night of observations. Although this is a fairly bright burst, note that the Totani et al. (2006) spectrum, for example, was observed 3.4 days after the burst and so could have been detected as a 10 μJy afterglow if it had been followed up more promptly, a little less than a day after the burst (assuming a typical $\sim t^{-1}$ decay). This can also be seen directly in the light curve of this source presented in the discovery paper by Haislip et al. (2006).

Moreover, in the future, it will be possible to exploit the power of 30 m class telescopes on the ground and the JWST in space. For example, Equation (15) shows that, using a 30 m telescope, one could observe a burst comparable to the Totani et al. (2006) one, with a flux of $f_{\nu}^c = 3 \mu\text{Jy}$ at an S/N of 20 in 4.7 hr, assuming the $z_s = 8$ sky background numbers. As mentioned earlier, adaptive optics might help to reduce the sky background contribution and required observing time. Using the JWST (with $D_{\text{tel}} = 6.5$ m), an $f_{\nu}^c = 2 \mu\text{Jy}$ burst observation would reach an S/N of 20 after 1.3 hr of observing time.¹⁶ Since the JWST

¹³ Throughout this work, we quote the S/N at the continuum and near Ly α at the source redshift. The true S/N in the spectrum will be smaller because of absorption in the damping wing, as properly accounted for in our Fisher matrix calculations.

¹⁴ <https://www.eso.org/observing/etc/bin/gen/form?INS.NAME=X-SHOOTER+INS.MODE=spectro>. We use the default fixed sky model parameters, a Moon phase parameter FLI = 0.5, and airmass = 1.5. We further assume that the sky background is comparable at other good observing sites.

¹⁵ We tested the accuracy of Equation (15) by comparing with the afterglow spectrum in Totani et al. (2006). This was a 2.6 μJy burst at $z = 6.3$, observed for 4 hr. Adopting a sky background estimate of $\sim 80 \mu\text{Jy}$, in between the I - and J -band values above, we estimate an S/N of 8 from Equation (15), which appears consistent with the published spectrum at the continuum near Ly α .

¹⁶ For the case of JWST, we consider NIRSPEC observations with a pixel size of 0.01 arcsec^2 and the highest spectral resolution mode with $R = 2700$. We use the sky background data from <https://jwst-docs.stsci.edu/jwst-observatory-functionality/jwst-background-model>. We assume the same efficiency factor, $\epsilon = 0.25$, as for the ground-based observations.

Table 1
Exposure Time Required to Obtain an $S/N = 20$ at Rest Frame 1216 \AA in the GRB Afterglow Spectra

z	6			8			10		
f_ν^c	1 μJy	5 μJy	10 μJy	1 μJy	5 μJy	10 μJy	1 μJy	5 μJy	10 μJy
f_ν^{sky} (ground)	59 μJy (<i>I</i> band)			340 μJy (<i>Y</i> band)			650 μJy (<i>J</i> band)		
8 m	4.4 days	4.5 hr	1.2 hr	24.8 days	1.0 day	6.1 hr	47.3 days	1.9 days	11.2 hr
30 m	7.5 hr	19.2 minutes	5.1 minutes	1.8 days	1.7 hr	26.0 minutes	3.4 days	3.2 hr	47.7 minutes
f_ν^{sky} (space)	0.07 μJy			0.07 μJy			0.06 μJy		
6.5 m JWST	2.4 hr	28.4 minutes	14.2 minutes	2.4 hr	28.4 minutes	14.2 minutes	2.4 hr	28.4 minutes	14.2 minutes

Note. These are only intended as rough estimates.

observations are source count-limited—and not by the wavelength-dependent sky background—the S/N here only depends on the source flux density, with time dilation also counteracting the dimming from increasing luminosity distance toward higher redshifts, as mentioned earlier. The JWST will therefore play a crucial role in obtaining high- S/N afterglow spectra of the most distant GRBs at $z \gtrsim 9$, where airglow makes ground-based observation especially challenging, although the prospects here with 30 m telescopes toward relatively bright afterglows are also fairly promising.

Table 1 provides a further summary of the exposure time requirements for obtaining an S/N of 20. Here we show examples at redshifts $z = 6, 8,$ and 10 for afterglows with specific flux densities of $f_\nu^c = 1, 5,$ and $10 \mu\text{Jy}$. We show results for 8 and 30 m telescopes from the ground and space-based observations with the JWST.

While these estimates suggest that there are multiple paths forward for obtaining the requisite $S/N \sim 20$ follow-up spectra, further work is needed to refine our S/N estimates and assess what fraction of the GRBs detectable with the Gamow Explorer, HiZ-GUNDAM, and the Theseus mission will meet these demands. This depends on the flux distribution of observable GRBs and how rapidly the follow-up spectroscopic observations can be executed.

We can also illustrate how our Fisher matrix forecasts depend on the quality of the afterglow follow-up observations. This is shown in Figure 12 and demonstrates the improvements that may be possible if the S/N is boosted to 30. It also shows how much the constraints degrade if the typical follow-up spectrum has $S/N = 10$ per $R = 3000$ resolution element.

5. Constructing a Census of the Ionizing Sources

In this section, we briefly illustrate how the combination of GRB afterglow constraints on both the escape fraction and the ionization history can help assess whether we have an accurate census of the ionizing source populations. The redshift distribution of long-duration GRBs also offers a handle on the evolution of the SFRD, provided evolution in the GRB to star formation rate is understood and/or empirically well calibrated. Essentially, the source term in Equation (5) is determined (at each redshift) by the SFRD, the escape fraction, and the ionizing spectrum of the sources. The recombination term in Equation (5) is generally less important, since the recombination time is fairly long near the cosmic mean density, and current simulation work suggests a fairly small clumping factor, $C \sim 2-3$ (e.g., McQuinn et al. 2011), during most of the EoR. Hence, if the ionizing spectrum may be constrained with, e.g., JWST observations, GRB afterglow data may allow one to

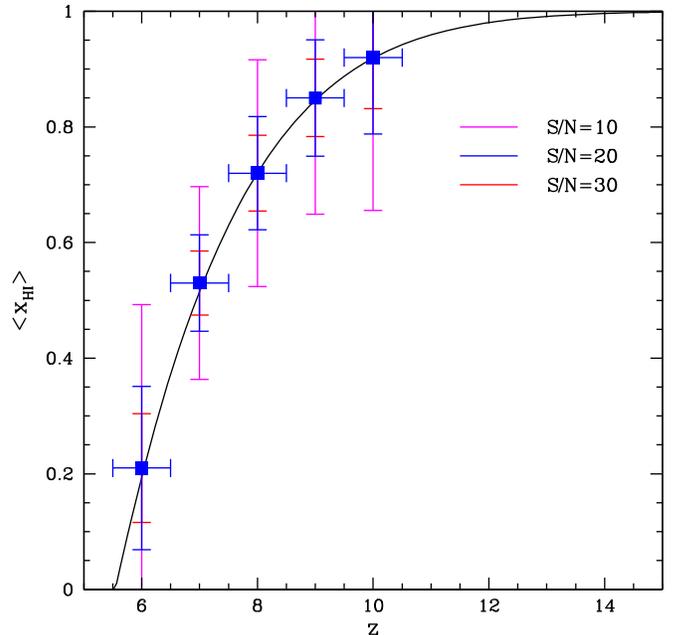


Figure 12. Forecast constraints on the reionization history of the universe vs. spectral S/N . This is identical to Figure 10, except here we vary the S/N of the spectral pixels from $S/N = 10$ (magenta error bars), to 20 (blue points, identical to the blue point in Figure 2), to 30 (red error bars). In all cases, we assume a spectral resolution of $R = 3000$ and the fiducial GRB redshift distribution.

predict the ionization history from the observed SFRD and escape fraction; this may be compared with the damping wing measurements of the reionization history.¹⁷ The SFRD determinations may, of course, be compared with those from JWST observations, which probe the bright end of the galaxy populations. Line-intensity mapping offers another potential handle on the cumulative impact of faint source populations (Kovetz et al. 2017; Sun et al. 2021). Taken together, these observations should provide a valuable test of whether we have an accurate census of the sources that reionized the universe.

The escape fraction determinations clearly play a key role here, because this quantity is highly uncertain both theoretically and empirically, and there is relatively little hope (other than through GRB afterglows) for direct constraints during the EoR, especially in the small-mass galaxies that may dominate the

¹⁷ One caveat here is that Equation (5) assumes a simplified form for the source term, where, for instance, the ionizing efficiency factor is independent of host halo mass and redshift. A more complex model will be required to interpret the data, but we focus on this simple model here for the purpose of illustration.

reionization process. A fairly large sample of afterglows with follow-up spectroscopy is, however, required to determine the escape fraction. The data required motivate follow-up spectroscopy of afterglows at slightly lower redshift, e.g., near $z \sim 5$; the higher SFRD rate and slightly less stringent requirements for detecting these more nearby sources will yield larger samples than the 20 $z_s \geq 6$ afterglows expected for the Gamow Explorer.

This postreionization sample avoids any damping wing contribution from the IGM; the escape fraction is then determined by fitting for the neutral hydrogen column density associated with each GRB host galaxy and averaging over the sample of sight lines (e.g., footnote 4, Chen et al. 2007; Tanvir et al. 2019). The likely scenario here is that a fraction $1 - f_{\text{esc}}$ of the sight lines have DLAs and hence effectively vanishing ionizing photon leakage, while the f_{esc} of the afterglows shows negligible neutral hydrogen absorption, i.e., no DLAs or even Lyman-limit systems. In this case, the requirements for follow-up spectroscopy are likely less stringent than in the case of measuring the full shape of the IGM damping wing. In a first pass, one mainly wants to determine whether or not a given sight line has a DLA. For estimating the escape fraction, it is not necessary to determine the precise DLA column density, since sight lines with DLAs give negligible escape fractions in any case. In the case of sight lines that do not show DLAs, further follow-up may be required to determine whether there is a Lyman-limit system and the associated column density.

5.1. Prospects for Measuring the Escape Fraction

For example, suppose that the escape fraction is $f_{\text{esc}} = 0.1$ and that roughly 100% of ionizing photons escape along a fraction f_{esc} of sight lines, while no ionizing photons escape from the remaining $1 - f_{\text{esc}}$ of afterglows. In this case, on average, one expects $N_{\text{esc}} = f_{\text{esc}} N_{\text{GRB}}$ sight lines to show escaping ionizing photons in a sample of N_{GRB} afterglows. Assuming Poisson fluctuations in N_{esc} around this average, the fractional error on the escape fraction from the sample is $\sigma_{f_{\text{esc}}}/f_{\text{esc}} \sim 1/\sqrt{N_{\text{esc}}} \sim 1/\sqrt{f_{\text{esc}} N_{\text{GRB}}}$. That is, in a sample of $N_{\text{GRB}} = 100$ GRB afterglows, we expect $N_{\text{esc}} = 10$ sight lines to show escaping ionizing photons and $\sim 30\%$ errors on the escape fraction estimate. Here we assume that the errors are dominated by Poisson fluctuations in the number of sight lines with escaping ionizing photons and ignore additional errors that may arise from imperfect column density determinations, for example. Note that the escape fraction assumed here implies strong redshift evolution from the Tanvir et al. (2019) sample, but this may be required to reionize the universe with star-forming galaxies (e.g., Bouwens et al. 2016), as discussed in Section 2.4. It would be interesting to split future samples by redshift and stellar mass (obtained from follow-up host galaxy observations) to explore any correlations or trends in these quantities.

Even a modest $\sim 30\%$ level constraint on the escape fraction would be interesting and helpful. In order to illustrate this, Figure 13 considers an alternate model for the ionizing sources, in which our fiducial minimum galaxy-hosting halo mass is boosted from $M_{\text{min}} = 10^9$ to $5 \times 10^9 M_{\odot}$. This is intended to loosely mimic scenarios in which there are no additional faint galaxy populations below the threshold of UV luminosity function measurements; that is, it is a case where bright galaxy populations produce most of the ionizing photons, termed “reionization by oligarchs” by Naidu et al. (2020). Although different than our

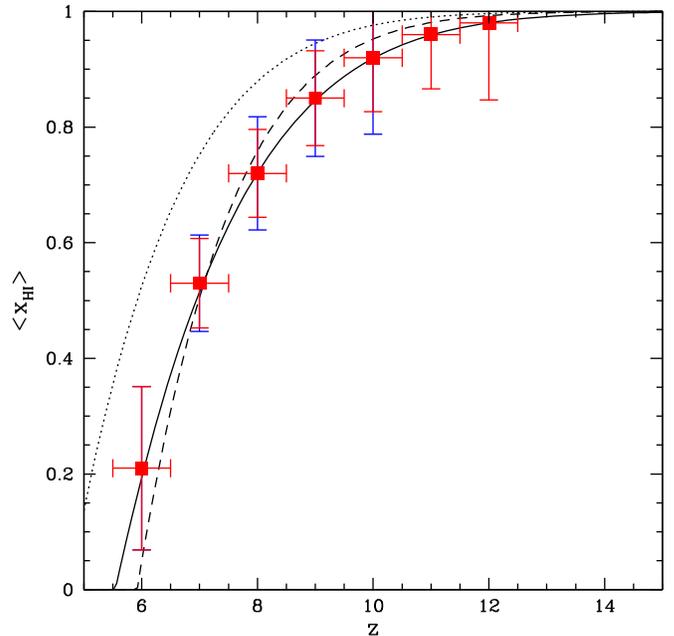


Figure 13. Illustration of how joint constraints on the ionization history and escape fraction can improve our understanding of ionizing source populations. The red and blue error bars are the forecast constraints on the reionization history from Figure 10. The black solid line shows the fiducial model (with a minimum host halo mass of $M_{\text{min}} = 10^9 M_{\odot}$). The black dotted line shows the impact of increasing the minimum host halo mass in the model to $M_{\text{min}} = 5 \times 10^9 M_{\odot}$ while fixing the escape fraction. The black dashed line also adopts $M_{\text{min}} = 5 \times 10^9 M_{\odot}$ but increases the escape fraction by a factor of 2, which largely compensates for the delay in reionization with this M_{min} . A measurement of the escape fraction can help rule out this case, which would otherwise be allowed by the ionization history measurements and determinations of the bright end of the UV luminosity function.

fiducial model, this type of scenario is also in agreement with current data, as discussed in that work. Adopting the same value of the escape fraction, or more precisely, the same ionizing efficiency factor in Equation (5), gives the dotted black curve in Figure 13. This delays the reionization history of the universe relative to our fiducial model and would be ruled out by the damping wing measurements given the forecasts of the previous section. In fact, this case is already moderately disfavored by current observations (e.g., Figure 2) but suffices for illustration.

However, one can compensate for this delay by increasing the escape fraction. In particular, the dashed line in Figure 13 boosts the escape fraction by a factor of 2, and this largely counteracts the effects from increasing the minimum mass.¹⁸ Put differently, there are degeneracies between the minimum host halo mass and the escape fraction (see also, for example, Figure 9 of Sun & Furlanetto 2016). This can be broken by adding the GRB determinations of the escape fraction, which—in our fiducial model—would argue against the factor of 2 boost in escape fraction. Without the escape fraction measurements, one could be spuriously led into thinking that the brighter sources reionize the universe. Hence, the combination of ionization history constraints and escape fraction measurements can play an important role in determining whether we have accurately identified the source populations responsible for reionization. Furthermore, the comparison between the

¹⁸ Formally, note that the dashed line case would also be disfavored at moderate statistical significance by the ionization history constraints. We sharpen this in the Fisher analysis below (Section 5.2).

redshift distribution of the GRBs and the SFRD determined from bright populations with, e.g., the JWST might provide independent evidence for important contributions from faint source populations, since long-duration GRBs may be produced in such galaxies, while these may remain below the JWST flux limits.

5.2. Source Parameter Constraints

We can sharpen this argument by further forecasting the error bars on the parameters of the ionizing sources, which may be obtained from both the GRB afterglow constraints on the ionization history and the escape fraction determinations. To do this, we perform a Fisher matrix forecast for the ionizing source parameters in Equation (5), given the error bars previously forecast for $\langle x_{\text{HI}}(z) \rangle$. Here, for simplicity, we consider a two-parameter model, characterized by $q_\zeta = (\zeta - \zeta_{\text{fid}})/\zeta_{\text{fid}}$ and $q_{M_{\text{min}}} = (M_{\text{min}} - M_{\text{min, fid}})/M_{\text{min, fid}}$. That is, we fix the clumping factor at $C = 3$ and ignore any redshift or mass dependence in these parameters. Our fiducial parameter values are $\zeta_{\text{fid}} = 20$ and $M_{\text{min}} = 10^9 M_\odot$. The Fisher matrix for these parameters may be computed as

$$F'_{ij} = \sum_k \frac{\partial \langle x_{\text{HI}}(z_k) \rangle}{\partial q_i} \frac{\partial \langle x_{\text{HI}}(z_k) \rangle}{\partial q_j} \frac{1}{\text{var}[\langle x_{\text{HI}}(z_k) \rangle]}, \quad (16)$$

where the prime notation distinguishes this Fisher matrix from the one in Equation (11). The sum runs over the redshift bins. The parameter derivatives are computed from the ionization equation (Equation (5)), and the variance of the neutral fraction estimates in each bin follow from the calculations in Section 4. Finally, we can add the independent information provided by an escape fraction estimate. Note that, in principle, the ionizing efficiency parameter depends on the overall star formation efficiency and the ionizing spectrum, as well as the escape fraction. For simplicity, we suppose that in the JWST era, the escape fraction is the main unknown, so that an $\sim 30\%$ error in the escape fraction translates into an $\sim 30\%$ error in the ionizing efficiency. In this case, assuming a fiducial escape fraction of 10% and a sample of 100 GRBs determines ζ with a fractional accuracy of $\sim 1/\sqrt{10}$, as discussed earlier.

The resulting constraint ellipses are shown in Figure 14 for each of the GRB redshift distributions in Figure 1, which result in the corresponding neutral fraction error bars in Figure 10.¹⁹ As anticipated previously, the ionization history measurements alone leave a strong positive degeneracy between q_ζ and $q_{M_{\text{min}}}$. However, the f_{esc} measurement can help break this degeneracy. Quantitatively, in our baseline $N_{\text{GRB}}(z \geq 6) = 20$ model, we find that the fractional M_{min} error is 74% for a joint constraint but only 200% without the information from f_{esc} . For the more optimistic GRB redshift distribution with $N_{\text{GRB}}(z \geq 6) = 31$, the corresponding numbers are 69% and 150%. In the Theseus-inspired scenario, with $N_{\text{GRB}}(z \geq 6) = 80$, these numbers improve to 59% and 92%, respectively.²⁰ A good deal of the constraining power is driven by the bursts at high redshift; for instance, if we only include the bursts at $z \geq 9$ in the Theseus case, the error bars on the minimum mass only increase by 8%.

¹⁹ As in our previous results, we do not adopt positivity priors on the source parameters, so $q_i \leq -1$ in some cases.

²⁰ In the Theseus case, one might imagine that this more futuristic mission could obtain a larger set of afterglows for the escape fraction measurement, which would tighten these limits.

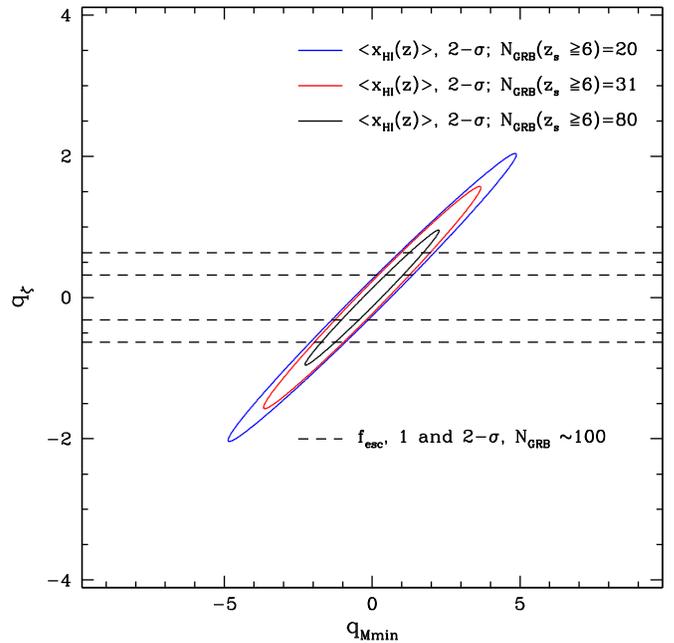


Figure 14. Constraints on the ionizing source properties from the reionization history and escape fraction measurements. The blue, red, and black contours show 2σ constraints on the ionizing efficiency parameter, ζ , and the minimum host halo mass of the ionizing sources, M_{min} , from the ionization history constraints for the fiducial, optimistic, and Theseus-inspired GRB samples, respectively. The black dashed lines show 1σ and 2σ constraints on the escape fraction from a sample of 100 afterglows, assuming that the uncertainty in the ionizing efficiency parameter is dominated by that in the escape fraction (see text). The combination of these measurements is helpful for determining the M_{min} parameter.

The high-redshift bursts play a valuable role here because the halo mass function becomes exponentially sensitive to M_{min} , as the galaxy-hosting halos form only in the rare high- σ peaks of the mass distribution in these highest-redshift bins. The numbers quoted are 1σ constraints. Hence, the escape fraction limits do help noticeably, and they may be still more valuable if f_{esc} is larger than our fiducial value of $f_{\text{esc}} = 0.1$, which would improve the escape fraction measurement error.

Further, we have adopted a very simple model for the ionizing sources in this illustration. In more realistic cases, the parameters may depend on halo mass, redshift, or other galaxy host properties. It will be important to combine ionization history constraints with a full suite of measurements to help pin down the properties of the ionizing sources. Determinations of the escape fraction and any measurements or limits on the redshift dependence of this quantity and/or correlations with host properties should be extremely valuable here. Although our simple illustration shows improvements on the M_{min} constraint from an f_{esc} determination, this may in fact understate the value of such measurements.

6. Conclusions

We have forecast the constraints on the reionization history for upcoming GRB afterglow samples during the EoR. Assuming 20 $z \geq 6$ GRB afterglow spectra with a typical S/N of 20 per $R = 3000$ spectral resolution element at the continuum, we find that the volume-averaged neutral fraction can be determined to 10%–15% accuracy from $z \sim 7$ –10 and with $\sim 70\%$ errors near $z = 6$. Although there appear to be reasonable prospects for obtaining the requisite follow-up

observations (see Section 4.2), more work is required to assess the number of GRBs that will meet these requirements in different redshift bins.

The ionization history forecasts represent a significant improvement on current constraints (e.g., Figure 2) and will provide a useful cross-check on future damping wing measurements toward high-redshift quasar samples. Perhaps more importantly, GRBs can help extend the quasar measurements past $z \gtrsim 8-9$ to probe earlier phases of the reionization history before bright quasars formed. Measurements of the redshifted 21 cm power spectrum and its evolution may offer more precise measurements of the reionization history (DeBoer et al. 2017) than our GRB forecasts, but these rely on understanding the mapping between the observable, spatial fluctuations in the 21 cm signal and the average ionization history and face a range of systematic challenges (Liu & Shaw 2020). While the GRB constraints also rely on understanding the mapping between an observable (the damping wing) and the average ionization history, they provide an independent constraint from the 21 cm efforts and will be valuable. In the future, it will be interesting to consider joint fits with 21 cm, future cosmic variance limited cosmic microwave background (CMB) EE polarization measurements (which provide mostly an integral constraint on the ionization history; e.g., Heinrich et al. 2017), the patchy kinetic Sunyaev–Zel’dovich effect (Ferraro & Smith 2018), the quasar constraints, Ly α emitter samples (e.g., Mason et al. 2019), and other probes. In the slightly longer term, a sample of as many as ~ 80 $z \geq 6$ GRB afterglows from the Theseus mission (Amati et al. 2018) with follow-up spectroscopy could return constraints as tight as $\sim 5\%$ across much of the EoR and place some limits out to $z \sim 14$. In principle, these measurements could be used to infer the electron-scattering optical depth to CMB photons and thereby sharpen cosmological parameter constraints (e.g., Liu et al. 2016).

We emphasized that an appealing aspect of GRB afterglow probes of reionization is that these constraints may also be combined with escape fraction determinations from the GRB column density distribution and internal or external measurements of the SFRD. In order, however, to build up a sufficient afterglow sample to determine the escape fraction, we recommend a follow-up campaign targeting afterglows at slightly lower redshifts, $z \sim 5$ and 50–100 afterglows. Taken together, the ionization history, escape fraction, and SFRD constraints will provide a powerful test of our census of the sources that reionized the universe.

It will also be interesting to perform follow-up observations to compare the large-scale environment around GRBs with strong and weak damping wings. Surveys of the galaxy distribution and/or line-intensity maps in various emission lines can potentially identify some of the surrounding galaxies responsible for ionizing the intergalactic gas near GRBs that show weak damping wing absorption. It may also be possible to construct maps of the surrounding neutral hydrogen using 21 cm surveys (Beardsley et al. 2015), which would presumably be brighter in the proximity of GRBs showing strong damping wing absorption. The combination of these observations should help in understanding the interplay between the ionizing sources and the neutral gas in the IGM during reionization.

The biggest caveat with our current work is that we assume a simplified three-parameter model for the damping wing

signature toward each GRB afterglow. Although patchy reionization is partly incorporated into this model, we assume a uniform neutral fraction outside of the ionized bubble around the GRB host. This description does not fully account for the generally complicated distribution of neutral and ionized patches toward each afterglow (McQuinn et al. 2008; Mesinger & Furlanetto 2008). While we expect our simplified description to be mostly adequate for forecasting the constraints that may be obtained with future missions, it is insufficient for interpreting upcoming measurements. The model should be least accurate in the late stages of reionization, when the ionized structures become large compared to the effective damping wing length scale. Note that while neglecting scatter in the exterior neutral fraction may cause us to underestimate the error forecasts on this quantity, it also implies that the damping wing feature will be stronger and more detectable toward some lines of sight than in our model; the precise impact of this approximation is unclear and likely sensitive to the reionization model considered.

We encourage further simulation studies to improve on this treatment. One approach may be to develop a generalized version of the three-parameter model informed by mock damping wing spectra extracted from reionization simulations. In this case, a description might include the extent of the first neutral region, the size of the second ionized patch, and so forth. In this case, the average neutral fraction may no longer even be a parameter in the model. Instead, one considers only the length of each neutral segment and their distance to the source. These could be fit for independently of any particular model, while the simulations could be used in a second, more model-dependent step to calibrate the relationship of these parameters to the volume-averaged neutral fraction. It may also be possible to calculate the average transmission profile around GRB hosts of a given halo mass using a halo model–type (Cooray & Sheth 2002) calculation (see Appendix C of Malloy & Lidz 2015 for work in this direction).

The simulations will also enable mock end-to-end tests of parameter extraction pipelines. Such efforts can test the Fisher matrix approximation adopted in this work, which assumes a quadratic expansion in the logarithm of the likelihood around a fiducial input model. A suite of mock spectra may further help in understanding the sample-to-sample variations in the expected error bars; our Fisher forecasts assume typical samples and do not consider the error on the error bar.

Another important issue concerns the redshift evolution of the DLA column density distribution. Our fiducial model adopts the current empirical distribution of DLA host column densities (Tanvir et al. 2019), but this may be a pessimistic assumption. It might be interesting to use numerical simulations of galaxy formation, capable of modeling the escape fraction of ionizing photons (e.g., Ma et al. 2020), to try and predict the redshift evolution in the DLA column density distribution toward GRB hosts. Interesting in its own right, this could also be used in our forecasts in place of the empirical distribution, which mostly consists of lower-redshift GRBs.

In summary, the prospects appear promising for future GRB missions to extract valuable constraints on the reionization history. There is a good deal of work to be done to better understand the expected damping wing signatures and best prepare for the upcoming missions, but we expect these to play an important role in our understanding of cosmic reionization.

We thank Matt McQuinn for providing the simulation outputs used in computing the ionized skewer size distributions in Figure 3. We thank the Gamow Explorer team for inspiring this work and for useful discussions, especially with Nial Tanvir, Nicholas White, Dieter Hartmann, Joseph Hennawi, Amy Lien, Ruben Salvaterra, and Giancarlo Ghirlanda. A.L. acknowledges support, in part, through NASA ATP grant 80NSSC20K0497. Part of this work was done at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

Appendix A

Here we describe how to account for the distribution in ionized skewer size and the DLA column density PDF in our Fisher matrix calculations. Let us first consider the conditional likelihood of the afterglow transmission data for a sight line with a given bubble size and column density, denoted by $\mathcal{L}(\hat{f}, \mathbf{q}|L, N_{\text{HI}})$, with

$$-\ln \mathcal{L}(\hat{f}, \mathbf{q}|L, N_{\text{HI}}) = \sum_k [\hat{f}(k) - f(k, \mathbf{q})]^2 \times \frac{1}{\text{var}[f(k)]}, \quad (\text{A1})$$

where we have assumed a Gaussian likelihood. Here \hat{f} denotes a measured transmission, while $f(\mathbf{q})$ describes a transmission model, and the sum is over independent spectral pixels. The conditional Fisher matrix, denoted $\tilde{F}_{ij}(L, N_{\text{HI}})$, then follows from the definition of the Fisher matrix as

$$\tilde{F}_{ij}(L, N_{\text{HI}}) = - \left\langle \frac{\partial^2 \ln \mathcal{L}(\hat{f}, \mathbf{q}|L, N_{\text{HI}})}{\partial q_i \partial q_j} \right\rangle, \quad (\text{A2})$$

and this results in Equation (11) in the main body of the text.

It is further useful to define the likelihood of a neutral fraction estimate, denoted \hat{x} here for brevity of notation, for a given realization of bubble size and DLA column density, yet marginalized over the estimates of bubble size and column density. We denote this quantity as $\mathcal{L}(\hat{x}|L, N_{\text{HI}})$. We can compute the width of this distribution around some fiducial volume-averaged neutral fraction by determining the inverse of the conditional Fisher matrix. Specifically,

$$\begin{aligned} \text{var}[\langle x_{\text{HI}} \rangle | L, N_{\text{HI}}] &= \left[\left\langle \frac{\partial^2 \ln \mathcal{L}(\hat{x}|L, N_{\text{HI}})}{\partial \hat{x}^2} \right\rangle \Big|_{\hat{x}=\langle x_{\text{HI}} \rangle} \right]^{-1} \\ &= \langle x_{\text{HI}} \rangle^2 [\tilde{F}^{-1}]_{q_{\text{xHI}}, q_{\text{xHI}}}, \end{aligned} \quad (\text{A3})$$

where the factor of $\langle x_{\text{HI}} \rangle^2$ in the second equality is needed to move from the variance of q_{xHI} (see Section 3) to that of $\langle x_{\text{HI}} \rangle$ itself.

Next, we can finally incorporate the bubble size and column density distributions. The likelihood of \hat{x} , marginalized over the bubble size and column density PDFs, is

$$\mathcal{L}(\hat{x}) = \int dL dN_{\text{HI}} \frac{dP_L}{dL} \frac{dP_{\text{NH1}}}{dN_{\text{HI}}} \mathcal{L}(\hat{x}|L, N_{\text{HI}}). \quad (\text{A4})$$

The nonconditional Fisher matrix of this one-parameter distribution is given by

$$F_{\langle x_{\text{HI}} \rangle, \langle x_{\text{HI}} \rangle} = - \left\langle \frac{\partial^2 \ln \mathcal{L}(\hat{x})}{\partial \hat{x}^2} \right\rangle \Big|_{\hat{x}=\langle x_{\text{HI}} \rangle}, \quad (\text{A5})$$

which is just the inverse variance of the \hat{x} distribution. We can evaluate this expression and relate the results to the conditional variance in Equation (A3). First, note that

$$\begin{aligned} \frac{\partial \ln \mathcal{L}(\hat{x})}{\partial \hat{x}} &= \int dL dN_{\text{HI}} \frac{dP_L}{dL} \frac{dP_{\text{NH1}}}{dN_{\text{HI}}} \\ &\times \frac{\mathcal{L}(\hat{x}|L, N_{\text{HI}})}{\mathcal{L}(\hat{x})} \frac{\partial \ln \mathcal{L}(\hat{x}|L, N_{\text{HI}})}{\partial \hat{x}}. \end{aligned} \quad (\text{A6})$$

We take the second derivative of Equation (A6), using that $\partial \mathcal{L}(\hat{x}|L, N_{\text{HI}}) / \partial \hat{x} \Big|_{\hat{x}=\langle x_{\text{HI}} \rangle} = \partial \mathcal{L}(\hat{x}) / \partial \hat{x} \Big|_{\hat{x}=\langle x_{\text{HI}} \rangle} = 0$. This gives

$$\begin{aligned} - \left\langle \frac{\partial^2 \ln \mathcal{L}(\hat{x})}{\partial \hat{x}^2} \right\rangle \Big|_{\hat{x}=\langle x_{\text{HI}} \rangle} &= - \int d\hat{x} dL dN_{\text{HI}} \mathcal{L}(\hat{x}) \frac{dP_L}{dL} \frac{dP_{\text{NH1}}}{dN_{\text{HI}}} \\ &\times \frac{\mathcal{L}(\hat{x}|L, N_{\text{HI}})}{\mathcal{L}(\hat{x})} \frac{\partial^2 \ln \mathcal{L}(\hat{x}|L, N_{\text{HI}})}{\partial \hat{x}^2} \Big|_{\hat{x}=\langle x_{\text{HI}} \rangle}, \end{aligned} \quad (\text{A7})$$

where the integration over \hat{x} , weighted by $\mathcal{L}(\hat{x})$, takes care of the ensemble averaging. We can rearrange the order of integration and rewrite this as

$$\begin{aligned} - \left\langle \frac{\partial^2 \ln \mathcal{L}(\hat{x})}{\partial \hat{x}^2} \right\rangle \Big|_{\hat{x}=\langle x_{\text{HI}} \rangle} &= - \int dL dN_{\text{HI}} \frac{dP_L}{dL} \frac{dP_{\text{NH1}}}{dN_{\text{HI}}} \\ &\times \int d\hat{x} \mathcal{L}(\hat{x}|L, N_{\text{HI}}) \frac{\partial^2 \ln \mathcal{L}(\hat{x}|L, N_{\text{HI}})}{\partial \hat{x}^2} \Big|_{\hat{x}=\langle x_{\text{HI}} \rangle} \\ &= \int dL dN_{\text{HI}} \frac{dP_L}{dL} \frac{dP_{\text{NH1}}}{dN_{\text{HI}}} \frac{1}{\text{var}[\langle x_{\text{HI}} \rangle | L, N_{\text{HI}}]}. \end{aligned} \quad (\text{A8})$$

Although the notation is a little cumbersome, we arrive at a simple result. The nonconditional Fisher matrix, $F_{\langle x_{\text{HI}} \rangle, \langle x_{\text{HI}} \rangle}$ —i.e., the inverse variance of the neutral fraction estimate—is a weighted average of the conditional neutral fraction inverse variance, where the weighting is over the bubble size and DLA column density PDFs. That is, our final expression is

$$F_{\langle x_{\text{HI}} \rangle, \langle x_{\text{HI}} \rangle} = \int dL dN_{\text{HI}} \frac{dP_L}{dL} \frac{dP_{\text{NH1}}}{dN_{\text{HI}}} \times \frac{1}{\text{var}[\langle x_{\text{HI}} \rangle | L, N_{\text{HI}}]}. \quad (\text{A9})$$

In summary, our procedure is to first compute the conditional Fisher matrices of Equations (11) and (A2) for many different values of L, N_{HI} , spanning the full range of their PDFs. We then invert each conditional Fisher matrix, using Equation (A3) to determine the inverse variance of the neutral fraction estimate for that L, N_{HI} . Finally, we average together the results using Equation (A9) to weight by the PDFs of L, N_{HI} . The final error bar follows simply from this equation as

$$\sigma_{\langle x_{\text{HI}} \rangle} = [N_{\text{GRB}}(z_s) F_{\langle x_{\text{HI}} \rangle, \langle x_{\text{HI}} \rangle}]^{-1/2}, \quad (\text{A10})$$

where the $N_{\text{GRB}}(z_s)$ factor accounts for the number of independent GRB sight lines in the redshift bin.

Appendix B

Here we show how sensitive our results are to variations around some of the assumptions in our modeling. First, we consider how the results depend on our treatment of the lower limit of integration in our damping wing computations (Equation (1)). Recall that throughout the text, we took each GRB to reside at the center of its redshift bin (at redshift z_s) and computed the damping wing using the average neutral fraction at bin center, truncating the integral at the edge of each redshift bin, $z_s - 0.5$. This neglects redshift evolution across the bin and, more importantly, the presence of neutral gas at redshifts below the lower-redshift end of the bin in question. While this estimate is imperfect, it simplifies the calculations and is a reasonable approximation given that the damping wing is dominated by relatively nearby gas in the IGM.

As one simple test of the impact of this approximation, we simply adjust the lower limit of the redshift integration from $z_{\text{end}} - 0.5$ to $z_{\text{end}} = 5.5$ while keeping $\langle x_{\text{HI}} \rangle$ fixed at its bin center value. This latter redshift value coincides with the edge of the lowest-redshift bin in our analysis and is close to the end of reionization in the model. This case deliberately exaggerates the impact of neutral gas in the IGM below the edge of each redshift bin, especially in the highest-redshift bins (while it gives identical results for the lowest-redshift bin). We then follow the procedure in Section 4 to compute the error bars on the ionization history for this alternate choice of z_{end} . The results of this test are shown in Figure 15; the error bars tighten slightly in this case, especially in the highest-redshift bin. However, the difference is relatively small, and the case with $z_{\text{end}} = 5.5$ deliberately exaggerates the effect of the neutral gas below the bin edge. Hence, our results are not strongly impacted by the choice of z_{end} .

A second choice relates to the range of wavelengths included in our Fisher matrix calculations (Equation (11)). In the computations in the text, we include the wavelength range between Ly α at the source redshift, at the observed wavelength $\lambda_{\text{obs}} = \lambda_{\alpha}(1 + z_s)$, and out to 1500 Å redward of this wavelength (independent of z_s). In Figure 16, we consider variations around this choice. Specifically, for the case of our fiducial GRB redshift distribution, we compute the marginalized fractional error on the neutral fraction in each redshift bin, including a range of different choices for the long-wavelength cutoff. In each bin, the results converge, provided $\gtrsim 500$ Å redward of Ly α at the source redshift are included (as might also be expected based on Figures 4 and 6), justifying our fiducial choice of 1500 Å.

Finally, the calculations in the text assume sky background-limited measurements, which are appropriate for the case of ground-based follow-up spectroscopic observations. As discussed in Section 4.2, space-based observations with the JWST will instead be in the source-dominated regime. At a fixed S/N, the transmission variance in the source-dominated case will be slightly smaller in absorbed regions than in the sky-dominated regime (see Equation (12)). This slightly improves the ability to measure the transmission profile close to Ly α at the source redshift, which helps in determining the bubble size and column density parameters. This then leads to less strong parameter degeneracies than in the sky-dominated case of Figures 7–9 and slight improvements on the marginalized neutral fraction error forecasts. Figure 17 compares the

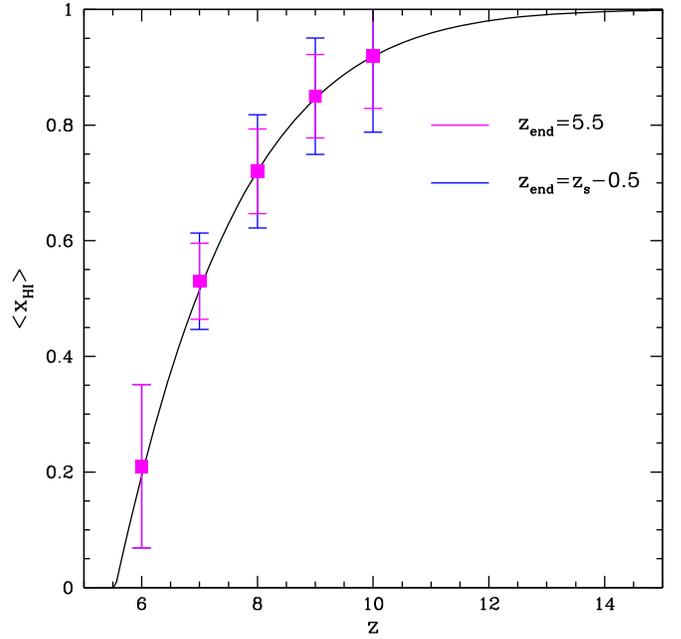


Figure 15. Error bar forecasts for different choices of z_{end} . The blue error bars show our fiducial model assumption for the lower-redshift limit on the damping wing integral in Equation (1). In this case, $z_{\text{end}} = z_s - 0.5$; that is, we adopt the average neutral fraction at bin center and integrate down to the lower edge of the bin. The magenta points instead adopt the average neutral fraction at bin center but integrate all the way down to $z_{\text{end}} = 5.5$. We adopt the fiducial GRB redshift distribution with $20 z \geq 6$ GRBs. Note that the two cases coincide for the $z_s = 6$ bin. The results are not strongly sensitive to the treatment of z_{end} .

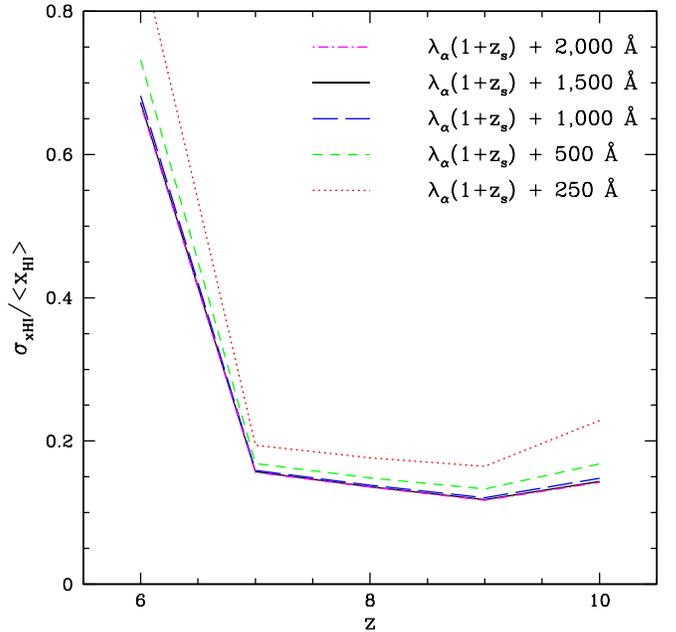


Figure 16. Fractional error bars vs. redshift for different choices of the maximum wavelength entering into the Fisher matrix computations of Equation (11). Again, the figure shows variations for our fiducial GRB redshift distribution. The text adopts $\lambda_{\text{max}} = \lambda_{\alpha}(1 + z_s) + 1500$ Å, but the results are insensitive to this choice.

sky-dominated case with alternate scenarios where the sky flux is 0.1 and 0.01 times the source flux. The improvements are relatively small, so we expect rather similar constraints at a fixed S/N in each of the sky- and source-dominated cases.

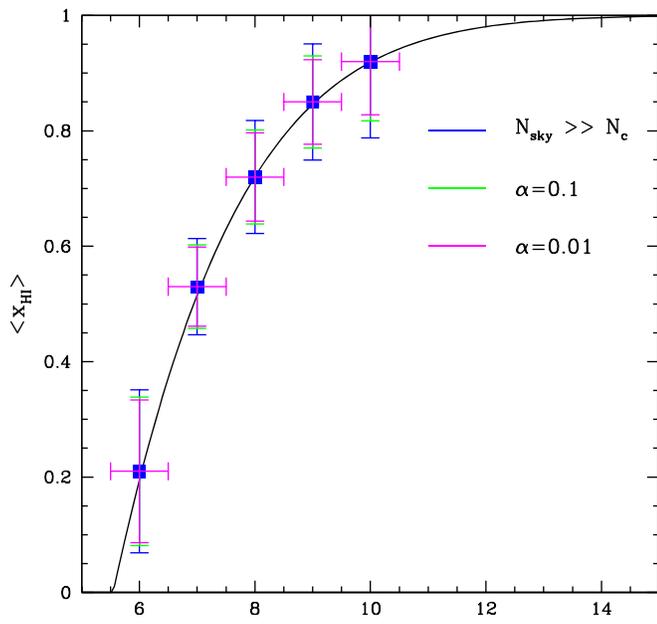


Figure 17. Error bar forecasts for different sky background to source count ratios. In each case, the S/N at the continuum is fixed at $S/N = 20$, but the blue error bars take the sky background–dominated limit adopted in the main text, while the green error bars assume the sky background is 0.1 times the source flux density, and the magenta errors give the case where the sky background is 100 times smaller than the source. In each case, we assume the fiducial GRB redshift distribution. The results are not sensitive to the assumptions here, although the errors are slightly smaller in the source-dominated case.

ORCID iDs

Adam Lidz  <https://orcid.org/0000-0002-3950-9598>

Lluís Mas-Ribas  <https://orcid.org/0000-0003-4584-8841>

Guochao Sun  <https://orcid.org/0000-0003-4070-497X>

References

- Aghanim, N., Akrami, Y., Ashdown, M., et al. 2020, *A&A*, 641, A6
 Amati, L., O’Brien, P., Götz, D., et al. 2018, *AdSpR*, 62, 191
 Amati, L., O’Brien, P. T., Götz, D., Bozzo, E., & Santangelo, A. 2021, *Proc. SPIE*, 11444, 114442J
 Barkana, R., & Loeb, A. 2001, *PhR*, 349, 125
 Barkana, R., & Loeb, A. 2004, *ApJ*, 601, 64
 Beardsley, A. P., Morales, M. F., Lidz, A., Malloy, M., & Sutter, P. M. 2015, *ApJ*, 800, 128
 Blain, A. W., & Natarajan, P. 2000, *MNRAS*, 312, L35
 Bouwens, R. J., Smit, R., Labbé, I., et al. 2016, *ApJ*, 831, 176
 Chen, H.-W., Prochaska, J. X., & Gnedin, N. Y. 2007, *ApJL*, 667, L125
 Chornock, R., Berger, E., Fox, D. B., et al. 2013, *ApJ*, 774, 26
 Cooray, A., & Sheth, R. K. 2002, *PhR*, 372, 1
 Cucchiara, A., Levan, A. J., Fox, D. B., et al. 2011, *ApJ*, 736, 7
 Davies, F. B., Hennawi, J. F., Bañados, E., et al. 2018a, *ApJ*, 864, 143
 Davies, F. B., Hennawi, J. F., Bañados, E., et al. 2018b, *ApJ*, 864, 142
 DeBoer, D. R., Parsons, A. R., Aguirre, J. E., et al. 2017, *PASP*, 129, 045001
 Dvorkin, C., & Smith, K. M. 2009, *PhRvD*, 79, 043003
 Fan, X., Barth, A., Banados, E., et al. 2019, *BAAS*, 51, 121
 Fan, X., Strauss, M. A., Becker, R. H., et al. 2006, *AJ*, 132, 117

- Faucher-Giguère, C.-A. 2020, *MNRAS*, 493, 1614
 Faucher-Giguère, C. A., Lidz, A., Hernquist, L., & Zaldarriaga, M. 2008, *ApJ*, 688, 85
 Ferraro, S., & Smith, K. M. 2018, *PhRvD*, 98, 123519
 Gehrels, N., Chincarini, G., Giommi, P., et al. 2004, *ApJ*, 611, 1005
 Ghirlanda, G., Salvaterra, R., Ghisellini, G., et al. 2015, *MNRAS*, 448, 2514
 Gunn, J. E., & Peterson, B. A. 1965, *ApJ*, 142, 1633
 Haislip, J. B., Nysewander, M. C., Reichart, D. E., et al. 2006, *Natur*, 440, 181
 Hartoog, O. E., Malesani, D., Fynbo, J. P. U., et al. 2015, *A&A*, 580, A139
 Heinrich, C. H., Miranda, V., & Hu, W. 2017, *PhRvD*, 95, 023513
 Keating, L. C., Weinberger, L. H., Kulkarni, G., et al. 2020, *MNRAS*, 491, 1736
 Kistler, M. D., Yuksel, H., Beacom, J. F., & Stanek, K. Z. 2008, *ApJL*, 673, L119
 Kovetz, E. D., Viero, M. P., Lidz, A., et al. 2017, arXiv:1709.09066
 Krühler, T., Malesani, D., Fynbo, J. P. U., et al. 2015, *A&A*, 581, A125
 Kulkarni, G., Keating, L. C., Haehnelt, M. G., et al. 2019, *MNRAS*, 485, L24
 Lamb, D. Q., & Reichart, D. E. 2000, *ApJ*, 536, 1
 Lee, H.-W. 2013, *ApJ*, 772, 123
 Lidz, A. 2016, *Understanding the Epoch of Cosmic Reionization* ed. A. Mesinger (Cham: Springer International), 23
 Lidz, A., McQuinn, M., & Zaldarriaga, M. 2007, *ApJ*, 670, 39
 Liu, A., Pritchard, J. R., Allison, R., et al. 2016, *PhRvD*, 93, 043013
 Liu, A., & Shaw, J. R. 2020, *PASP*, 132, 062001
 Loeb, A., & Furlanetto, S. R. 2013, *The First Galaxies in the Universe* (Princeton, NJ: Princeton Univ. Press)
 Ma, X., Quataert, E., Wetzel, A., et al. 2020, *MNRAS*, 498, 2001
 Madau, P., & Dickinson, M. 2014, *ARA&A*, 52, 415
 Madau, P., Haardt, F., & Rees, M. J. 1999, *ApJ*, 514, 648
 Malloy, M., & Lidz, A. 2015, *ApJ*, 799, 179
 Mason, C. A., Fontana, A., Treu, T., et al. 2019, *MNRAS*, 485, 3947
 Mason, C. A., Treu, T., Dijkstra, M., et al. 2018, *ApJ*, 856, 2
 Mas-Ribas, L., Miralda-Escudé, J., Pérez-Ràfols, I., et al. 2017, *ApJ*, 846, 4
 McGreer, I., Mesinger, A., & D’Odorico, V. 2015, *MNRAS*, 447, 499
 McQuinn, M., Hernquist, L., Zaldarriaga, M., & Dutta, S. 2007a, *MNRAS*, 381, 75
 McQuinn, M., Lidz, A., Zahn, O., Dutta, S., Hernquist, L., & Zaldarriaga, M. 2007b, *MNRAS*, 377, 1043
 McQuinn, M., Lidz, A., Zaldarriaga, M., Hernquist, L., & Dutta, S. 2008, *MNRAS*, 388, 1101
 McQuinn, M., Oh, S. P., & Faucher-Giguère, C.-A. 2011, *ApJ*, 743, 82
 Melandri, A., Bernardini, M. G., D’Avanzo, P., et al. 2015, *A&A*, 581, A86
 Mesinger, A. 2010, *MNRAS*, 407, 1328
 Mesinger, A., & Furlanetto, S. 2008, *MNRAS*, 385, 1348
 Miralda-Escudé, J. 1998, *ApJ*, 501, 15
 Naidu, R. P., Tacchella, S., Mason, C. A., et al. 2020, *ApJ*, 892, 109
 Robertson, B. E., & Ellis, R. S. 2012, *ApJ*, 744, 95
 Salvaterra, R., Della Valle, M., Campana, S., et al. 2009, *Natur*, 461, 1258
 Savaglio, S. 2006, *NJPh*, 8, 195
 Shapiro, P. R., & Giroux, M. L. 1987, *ApJL*, 321, L107
 Sheth, R. K., Mo, H. J., & Tormen, G. 2001, *MNRAS*, 323, 1
 Stratta, G., Ciolfi, R., Amati, L., et al. 2018, *AdSpR*, 62, 662
 Sun, G., Chang, T.-C., Uzgil, B. D., et al. 2021, *ApJ*, 915, 33
 Sun, G., & Furlanetto, S. R. 2016, *MNRAS*, 460, 417
 Tanvir, N. R., Fox, D. B., Levan, A. J., et al. 2009, *Natur*, 461, 1254
 Tanvir, N. R., Fynbo, J. P. U., de Ugarte Postigo, A., et al. 2019, *MNRAS*, 483, 5380
 Tanvir, N. R., Laskar, T., Levan, A. J., et al. 2018, *ApJ*, 865, 107
 Totani, T., Kawai, N., Kosugi, G., et al. 2006, *PASJ*, 58, 485
 Vergani, S. D., Salvaterra, R., Japelj, J., et al. 2015, *A&A*, 581, A102
 White, N. E. 2020, arXiv:2003.01592
 Woosley, S. E. 1993, *ApJ*, 405, 273
 Yang, J., Wang, F., Fan, X., et al. 2020, *ApJ*, 904, 26
 Yonetoku, D., Mihara, T., Sawano, T., et al. 2014, *Proc. SPIE*, 9144, 91442S