



Detection of Cosmological 21 cm Emission with the Canadian Hydrogen Intensity Mapping Experiment

The CHIME Collaboration,

Mandana Amiri¹, Kevin Bandura^{2,3}, Tianyue Chen⁴, Meiling Deng^{5,6,1}, Matt Dobbs^{7,8}, Mateus Fandino^{1,9}, Simon Foreman^{6,5,10}, Mark Halpern¹, Alex S. Hill^{11,5}, Gary Hinshaw¹, Carolin Höfer¹, Joseph Kania¹², T. L. Landecker⁵, Joshua MacEachern¹, Kiyoshi Masui^{4,13}, Juan Mena-Parra^{4,14,15}, Nikola Milutinovic¹, Arash Mirhosseini¹, Laura Newburgh¹⁶, Anna Ordog^{11,5}, Ue-Li Pen^{17,18,14,6}, Tristan Pinsonneault-Marotte¹, Ava Polzin¹⁹, Alex Reda¹⁶, Andre Renard¹⁵, J. Richard Shaw¹, Seth R. Siegel^{7,8,6}, Saurabh Singh^{7,8,20}, Keith Vanderlinde^{14,15}, Haochen Wang^{4,13}, Donald V. Wiebe¹, and Dallas Wulf^{7,8}

¹ Department of Physics and Astronomy, University of British Columbia, Vancouver, BC, Canada; richard@phas.ubc.ca

² Department of Computer Science and Electrical Engineering, West Virginia University, Morgantown, WV, USA

³ Center for Gravitational Waves and Cosmology, West Virginia University, Morgantown, WV, USA

⁴ MIT Kavli Institute for Astrophysics and Space Research, Massachusetts Institute of Technology, Cambridge, MA, USA

⁵ Dominion Radio Astrophysical Observatory, Herzberg Astronomy & Astrophysics Research Centre, National Research Council Canada, Penticton, BC, Canada

⁶ Perimeter Institute for Theoretical Physics, Waterloo, ON, Canada; ssiegel@perimeterinstitute.ca

⁷ Department of Physics, McGill University, Montreal, QC, Canada

⁸ Trottier Space Institute, McGill University, Montreal, QC, Canada

⁹ Department of Physical Sciences, Thompson Rivers University, Kamloops, BC, Canada

¹⁰ Department of Physics, Arizona State University, Tempe, AZ, USA

¹¹ Department of Computer Science, Math, Physics, and Statistics, University of British Columbia-Okanagan, Kelowna, BC, Canada

¹² Department of Physics and Astronomy, West Virginia University, Morgantown, WV, USA

¹³ Department of Physics, Massachusetts Institute of Technology, Cambridge, MA, USA

¹⁴ David A. Dunlap Department of Astronomy & Astrophysics, University of Toronto, Toronto, ON, Canada

¹⁵ Dunlap Institute for Astronomy and Astrophysics, University of Toronto, Toronto, ON, Canada

¹⁶ Department of Physics, Yale University, New Haven, CT, USA

¹⁷ Canadian Institute for Theoretical Astrophysics, Toronto, ON, Canada

¹⁸ Institute of Astronomy and Astrophysics, Academia Sinica, Taipei, Taiwan

¹⁹ Department of Astronomy, Yale University, New Haven, CT, USA

²⁰ Raman Research Institute, Sadashivanagar, Bengaluru, India

Received 2022 February 3; revised 2022 November 8; accepted 2023 January 5; published 2023 April 12

Abstract

We present a detection of 21 cm emission from large-scale structure (LSS) between redshift 0.78 and 1.43 made with the Canadian Hydrogen Intensity Mapping Experiment. Radio observations acquired over 102 nights are used to construct maps that are foreground filtered and stacked on the angular and spectral locations of luminous red galaxies (LRGs), emission-line galaxies (ELGs), and quasars (QSOs) from the eBOSS clustering catalogs. We find decisive evidence for a detection when stacking on all three tracers of LSS, with the logarithm of the Bayes factor equal to 18.9 (LRG), 10.8 (ELG), and 56.3 (QSO). An alternative frequentist interpretation, based on the likelihood ratio test, yields a detection significance of 7.1σ (LRG), 5.7σ (ELG), and 11.1σ (QSO). These are the first 21 cm intensity mapping measurements made with an interferometer. We constrain the effective clustering amplitude of neutral hydrogen (H I), defined as $\mathcal{A}_{\text{H I}} \equiv 10^3 \Omega_{\text{H I}} (b_{\text{H I}} + \langle f\mu^2 \rangle)$, where $\Omega_{\text{H I}}$ is the cosmic abundance of H I, $b_{\text{H I}}$ is the linear bias of H I, and $\langle f\mu^2 \rangle = 0.552$ encodes the effect of redshift-space distortions at linear order. We find $\mathcal{A}_{\text{H I}} = 1.51_{-0.97}^{+3.60}$ for LRGs ($z = 0.84$), $\mathcal{A}_{\text{H I}} = 6.76_{-3.79}^{+9.04}$ for ELGs ($z = 0.96$), and $\mathcal{A}_{\text{H I}} = 1.68_{-0.67}^{+1.10}$ for QSOs ($z = 1.20$), with constraints limited by modeling uncertainties at nonlinear scales. We are also sensitive to bias in the spectroscopic redshifts of each tracer, and we find a nonzero bias $\Delta v = -66 \pm 20 \text{ km s}^{-1}$ for the QSOs. We split the QSO catalog into three redshift bins and have a decisive detection in each, with the upper bin at $z = 1.30$ producing the highest-redshift 21 cm intensity mapping measurement thus far.

Unified Astronomy Thesaurus concepts: Cosmology (343); Large-scale structure of the universe (902); H I line emission (690); Quasars (1319); Emission line galaxies (459)

1. Introduction

Measurements of the large-scale clustering of matter have great potential to improve our understanding of both the early and late universe, probing phenomena ranging from cosmic inflation to

dark energy to galaxy evolution. This large-scale structure (LSS) can be mapped in a variety of ways, including tabulating the locations of luminous objects, using gravitational lensing to relate the distorted appearance of galaxy shapes to mass along the line of sight, identifying the absorption of Ly α photons in the spectra of distant quasars, and isolating so-called secondary anisotropies in maps of the cosmic microwave background (CMB).

Another approach to mapping LSS, 21 cm *intensity mapping*, uses the hyperfine “spin-flip” transition in neutral



Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](#). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

hydrogen (hereafter H I), which has rest wavelength 21.106 cm (rest frequency 1420.406 MHz). The probability of this transition occurring spontaneously in a given hydrogen atom is extremely low, but this is balanced by the large cosmic abundance of H I in such a way that extragalactic 21 cm emission (and/or absorption) is measurable in aggregate. The lack of comparably strong spectral lines at frequencies below 1420 MHz and the optical thinness of the hyperfine transition together imply that we can, if foregrounds can be removed, directly observe a redshift of the 21 cm line. This can then be related to a distance from the observer. Thus, maps of the radio sky at different frequencies contain information about the distribution of H I at different cosmic times, and the spectral and angular fluctuations of these maps can provide us with a 3D picture of this distribution (Battye et al. 2004; Chang et al. 2008; Wyithe & Loeb 2008; Peterson et al. 2009). This idea extends beyond the 21 cm line, and intensity mapping is now being pursued across a wide range of atomic and molecular transitions (Kovetz et al. 2019).

At $z \lesssim 6$, after cosmic reionization has completed, the vast majority of H I is concentrated in the surroundings of galaxies, where it is shielded from ionizing radiation (Villaescusa-Navarro et al. 2018). Thus, a post-reionization 21 cm intensity mapping survey is effectively a coarse-grained galaxy survey, in which galaxies are detected in bulk via their H I content.

The 21 cm brightness temperature fluctuations are therefore highly correlated with galaxy catalogs from other surveys, and this fact has enabled the first detections of LSS using 21 cm intensity mapping. After the initial detection by Pen et al. (2009), which combined existing spectral intensity data from the HIPASS survey with the 6dF galaxy survey, subsequent analyses have used dedicated observations by the Green Bank and Parkes radio telescopes, in concert with galaxy catalogs from the DEEP2, WiggleZ, and 2dF surveys and the extended Baryon Oscillation Spectroscopic Survey (eBOSS; Chang et al. 2010; Masui et al. 2013; Anderson et al. 2018; Tramonte & Ma 2020; Li et al. 2021; Wolz et al. 2022), to detect cross-correlations with signal-to-noise ratios between 4 and 13. Several of these studies have placed constraints on the product $\Omega_{\text{H}\,1} b_{\text{H}\,1} r$, where $\Omega_{\text{H}\,1}$ is the mean H I density as a fraction of the present-day critical density, $b_{\text{H}\,1}$ is the linear bias of H I with respect to matter, and r is a cross-correlation parameter that absorbs uncertainties in the modeling.

In principle, much more powerful measurements of LSS are possible with custom-built telescopes that are optimized for 21 cm observations. This, alongside several other science targets, motivated the design and construction of the Canadian Hydrogen Intensity Mapping Experiment (CHIME).²¹ CHIME is a transit radio interferometer composed of four 20 m \times 100 m cylindrical reflectors, each instrumented with 256 dual-polarized feeds observing at 400–800 MHz. Signals from each feed are processed by an FX correlator and stored for offline cosmological analysis. These signals are also fed to separate back ends devoted to studying fast radio bursts (CHIME/FRB Collaboration et al. 2018) and pulsars (CHIME/Pulsar Collaboration et al. 2021). CHIME Collaboration et al. (2022b) provides an overview of the key features and operational status of the telescope.

²¹ <http://chime-experiment.ca/>

In this paper, we report the first detection of LSS with 21 cm intensity mapping data from CHIME,²² in cross-correlation with galaxies and quasars measured by eBOSS (Dawson et al. 2016). We make use of a stacking approach, which averages sky maps constructed from CHIME observations at the locations of each eBOSS object. The data processing involved in this approach is more straightforward than other cross-correlation methods (e.g., a cross-power spectrum) and involves intermediate data products (such as sky maps) that can be interpreted in terms of features of the telescope and analysis pipeline. These interpretations are vital for examining the performance of our analysis methods, several of which have been custom designed for CHIME.

Using 102 nights of CHIME data, we have achieved significant detections of cross-correlations with eBOSS catalogs of luminous red galaxies (LRGs), emission-line galaxies (ELGs), and quasars (QSOs). We quantify this significance within a Bayesian framework, finding Bayes factors $\mathcal{Z}_1/\mathcal{Z}_0$ (comparing our signal model with a noise-only model) of $\ln(\mathcal{Z}_1/\mathcal{Z}_0) \approx 18.9$ (LRGs), 10.8 (ELGs), and 56.3 (QSOs), each corresponding to decisive evidence on the Jeffreys scale (Jeffreys 1961); an alternative quantification, using a frequentist likelihood ratio test, yields signal-to-noise ratios of 7.1 (LRGs), 5.7 (ELGs), and 11.1 (QSOs).

H I stacking analyses have previously been carried out on interferometric data from the Westerbork Synthesis Radio Telescope (Rhee et al. 2013; Hu et al. 2019, 2020), the Giant Metrewave Radio Telescope (Lah et al. 2007; Kanekar et al. 2016; Rhee et al. 2016, 2018; Bera et al. 2019; Chowdhury et al. 2020), and the Very Large Array (VLA; Chen et al. 2021), as well as on single-antenna data from Parkes (Delhaize et al. 2013; Tramonte et al. 2019; Tramonte & Ma 2020) and the Arecibo Legacy Fast ALFA Survey (Guo et al. 2020). The primary motivation of many of these studies was to improve our understanding of galaxy evolution by probing the reservoirs of H I that serve as fuel for star formation. At $z \gtrsim 0.2$, the 21 cm line is too faint to detect in individual galaxies, but stacking enables a measurement of the average 21 cm flux (and therefore the average H I mass) across all objects in a given catalog, and a sufficiently small beam (possessed by the interferometers above) acts to limit the associated confusion noise. Under certain assumptions about the H I mass–luminosity relation, as well as the completeness and luminosity function of the catalog used for stacking, these measurements can also be used to constrain $\Omega_{\text{H}\,1}(z)$, which controls the overall amplitude of the large-scale 21 cm fluctuations that can be used for cosmology (see Chen et al. 2021 for a recent summary of these constraints).

In contrast to the interferometers mentioned above, CHIME is designed to make 21 cm observations that are intentionally confusion dominated, allowing efficient mapping of the large-scale clustering of 21 cm sources via the corresponding fluctuations in measured 21 cm intensity in broad spatial pixels. Thus, instead of exclusively probing the H I within individual objects in an external catalog, our stacking measurements are broadly sensitive to the nearby structures that are correlated with each object. To infer the value of $\Omega_{\text{H}\,1}(z)$, we must model gravitational and baryonic clustering in addition to the properties of the catalog objects themselves.

²² LSS has previously been detected by cross-correlating CHIME’s first catalog of fast radio bursts with photometric galaxy catalogs (Rafiee-Ravandi et al. 2021).

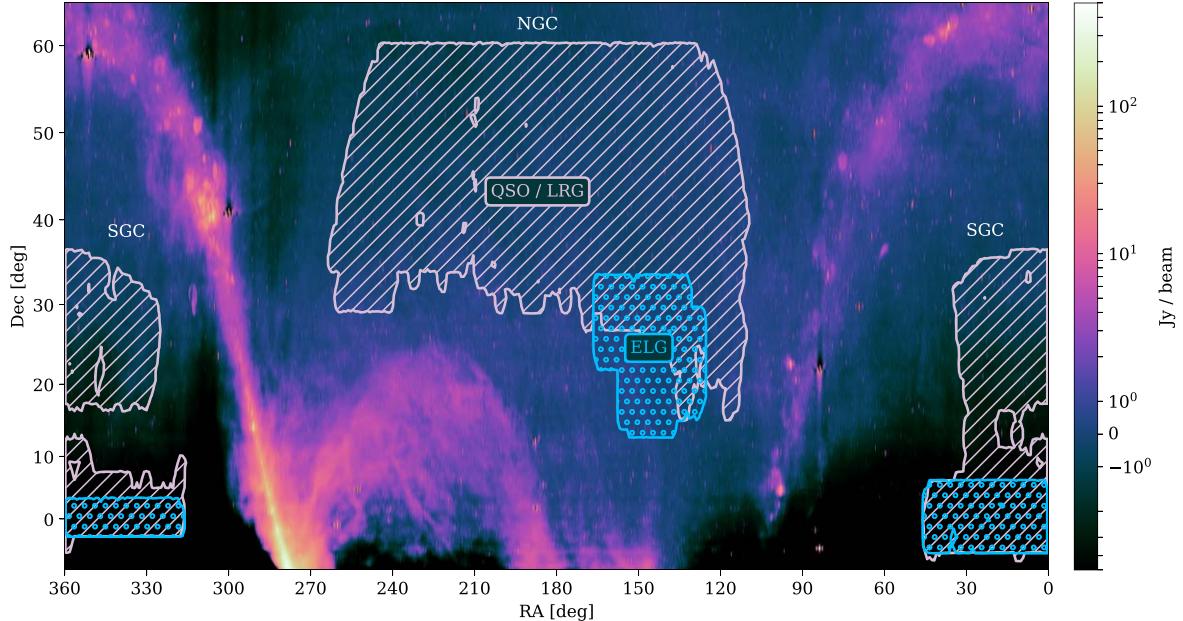


Figure 1. Map of the northern radio sky as measured by CHIME. Shown is the average spectral flux density over the 587.5–800 MHz subband. The hashed regions indicate the spatial footprints of the eBOSS catalogs. The LRG and QSO catalogs share a common footprint indicated by the light-pink hash marks. The footprint of the ELG catalog is indicated by the blue circular hash marks. The eBOSS catalogs are spread across two fields: the NGC and the SGC. We only present results for the NGC field in this work. The color scale is linear between -1 and 1 Jy beam $^{-1}$ and logarithmic otherwise. The map contains negative values because the autocorrelation data have been excluded. The zero-point is defined by setting the median value of a quiet part of the map with R.A. between 135° and 150° equal to zero for each decl. and frequency prior to averaging over the subband.

This modeling is most straightforward at the largest spatial scales, but as part of our analysis, we have needed to apply aggressive filtering that has removed the sensitivity of the data to these well-understood scales. Nevertheless, after marginalizing over the uncertainty associated with modeling of smaller-scale clustering, we are able to constrain an effective H I clustering amplitude \mathcal{A}_{HI} , defined as $\Omega_{\text{HI}}(b_{\text{HI}} + \langle f\mu^2 \rangle)$, where b_{HI} is a linear bias factor that relates large-scale clustering of H I to the clustering of all matter, and $\langle f\mu^2 \rangle$ is an effective quantity involving the linear growth rate, f , and the relative contribution of line of sight and transverse information (described in detail in Section 6.2). With $\langle f\mu^2 \rangle = 0.552$, we obtain $\mathcal{A}_{\text{HI}} = 1.51^{+3.60}_{-0.97}$ (LRGs), $\mathcal{A}_{\text{HI}} = 6.76^{+9.04}_{-3.74}$ (ELGs), and $\mathcal{A}_{\text{HI}} = 1.68^{+1.10}_{-0.67}$ (QSOs), to be compared with fiducial model values of 1.13, 1.21, and 1.37, respectively. While this precision is lower than previous single-antenna measurements, it is significantly more robust in its incorporation of modeling uncertainty: if we were able to fix the values of all small-scale parameters a priori, the precision on \mathcal{A}_{HI} would improve to between 10% and 20% for each sample. We are actively working on improvements to the calibration procedure and data processing that will allow for less aggressive filtering and enable the recovery of larger spatial scales. This is expected to both increase the signal-to-noise ratio of the detection and reduce the modeling uncertainty on the parameter constraints.

This paper (which includes descriptions of several analysis methods that have not previously appeared in the literature) is organized as follows:

1. In Section 2, we describe the CHIME and eBOSS data we use, visualizing the sky coverage in Figure 1 and redshift coverage in Figure 2.
2. In Sections 3 and 4, we describe how CHIME data are processed into stacks at the locations of eBOSS catalog objects, including our procedures for real-time processing

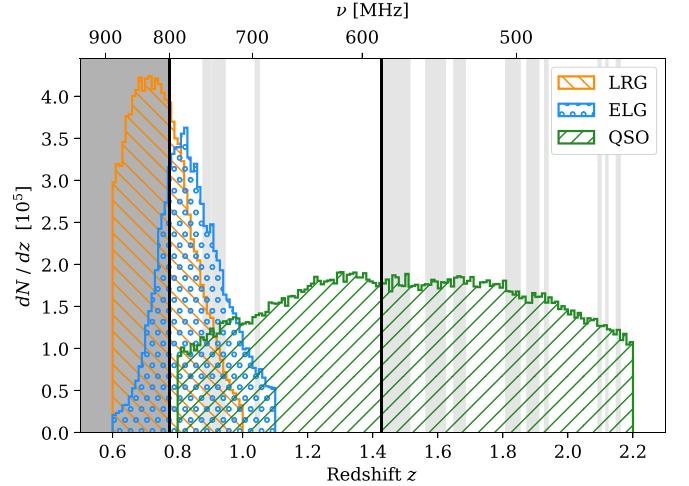


Figure 2. The redshift distribution of the LRG, ELG, and QSO catalogs for the NGC field. The y-axis indicates the number of sources per unit redshift. The upper x-axis indicates the frequency of 21 cm emission from a source at the redshift indicated by the lower x-axis. The dark-gray band denotes the range of frequencies that are outside the CHIME band (400–800 MHz). The light-gray bands denote ranges of frequencies that are inaccessible to CHIME because they are contaminated by a persistent source of RFI. The black solid lines mark the edges of the 587.5–800 MHz subband that will be used in this analysis.

- (Section 3.1), applying additional corrections to individual days of data (Section 3.2), averaging over days (Section 3.3), mapmaking (Section 4.3), beam calibration (Section 4.4), foreground filtering (Section 4.5), masking (Section 4.6), stacking (Section 4.7), and covariance estimation (Section 4.8).
3. In Section 5, we discuss the cosmological scales our analysis probes (Section 5.1), our model for the stacking signal (Section 5.2), our simulation framework

- (Section 5.3), and our simulation-based approach to model fitting (Section 5.4).
4. We begin Section 6 by presenting our stacking measurements and discussing several null tests in Section 6.1. The main results are shown in Figures 19 and 20. We then introduce our model fitting procedure (Section 6.2), visualize the constraints on the parameters of our model and discuss degeneracies (Section 6.3; see Figures 23–25), and quantify the significance of the detected signal (Section 6.4; see Table 3).
 5. In Section 7, we present the results of several validation tests that were performed on the data, related to consistency between the two instrumental polarizations, consistency between jackknives in observing time, beam calibration accuracy, and linearity of the stacking procedure.
 6. In Section 8, we discuss several aspects of the interpretation of these results: confirmation of a systematic bias in the reported QSO redshifts (Section 8.1), uncertainties on our constraints on the HI clustering amplitude A_{HI} (Section 8.2; see Table 8), comparisons of the corresponding Ω_{HI} constraints with previous results from the literature (Section 8.3; see Figure 29), and prospects for constraining the mean HI mass of objects in external catalogs (Section 8.4).
 7. In Section 9, we state our conclusions and discuss the prospects for future 21 cm measurements by CHIME.

We also include four appendices, detailing our Gibbs-sampling-based approach to delay spectrum estimation (Appendix A), the construction of our primary beam model using catalogs of point-source fluxes (Appendix B), the justification for stacking simulated Gaussian 21 cm maps on lognormal mock catalogs (Appendix C), and our construction of simulation-based signal templates used for model fitting (Appendix D).

For computations requiring a cosmological model, we use cosmological parameters from the final Planck data release (specifically, the “TT,TE,EE+lowE+lensing+BAO” parameters from Table 2 of Planck Collaboration et al. 2020).

2. Data

2.1. CHIME

Our analysis uses the CHIME stack data set acquired between 2019 January 1 and November 5. The stack data set is described in CHIME Collaboration et al. (2022b) and consists of the N_{feed}^2 visibilities (with $N_{\text{feed}} = 2048$) after they have been integrated to $\Delta t = 9.9405$ s cadence, calibrated for complex gain variations, and compressed by averaging subsets of redundant baselines. We selected 102 nights from this period to include in the analysis, using criteria that will be described in Section 3.3.1. After masking intervals of poor data quality, these 102 nights contain 521 hr of total integration time on the relevant eBOSS field.

CHIME is sensitive to radio frequencies from 400 to 800 MHz, which corresponds to 21 cm emission from redshifts 2.55 to 0.78. However, frequencies from 400 to 500 MHz suffer from frequent narrowband, transient radio frequency interference (RFI). In addition, approximately 60% of frequencies between 488 and 584 MHz are corrupted by persistent RFI from locally broadcast TV channels. Hence, for this initial analysis we have restricted our attention to the CHIME data acquired in the relatively clean portion of the band between 587.5 and 800 MHz, corresponding to 21 cm emission from

redshifts 1.42 to 0.78. The spectral resolution of the stack data set is $\Delta\nu = 0.390625$ MHz, resulting in 544 frequency channels within this range. We anticipate that the real-time, RFI excision algorithm that was deployed on the CHIME correlator in 2019 mid-October and recent improvements to the offline RFI excision algorithm will enable the inclusion of the lower half of the CHIME band in future analyses.

2.2. eBOSS Catalogs

eBOSS (Dawson et al. 2016), the cosmological survey within the Sloan Digital Sky Survey IV (SDSS-IV; Blanton et al. 2017), was conducted over 4.5 yr using spectrographs previously used for BOSS (Smee et al. 2013), mounted on the Sloan Telescope (Gunn et al. 2006) at the Apache Point Observatory. eBOSS produced four distinct samples of objects, each of which has been used to measure large-scale clustering and place constraints on a variety of cosmological parameters (see Alam et al. 2021 for a summary of these results). In this work, we cross-correlate three of these samples, from SDSS Data Release 16 (Ahumada et al. 2020), with CHIME measurements.

The eBOSS ELG sample (Raichoor et al. 2021) selected targets using imaging from the Dark Energy Camera Legacy Survey (Dey et al. 2019), making special use of emission in the [O II] doublet at $(\lambda 3727, \lambda 3729)$ to obtain efficient and accurate redshift estimates. This resulted in a catalog of 173,736 unique objects over $0.6 < z < 1.1$, spread across two fields: the 550 deg² North Galactic Cap (NGC) and the 620 deg² South Galactic Cap (SGC). These correspond to comoving volumes of 0.63 and $0.71 h^{-3} \text{ Gpc}^3$, respectively. We show both fields, superimposed on a representative CHIME sky map, in Figure 1.

The LRG sample (Ross et al. 2020) is composed of objects from optical imaging taken during previous phases of SDSS (Albareti et al. 2017), along with infrared data from the Wide-field Infrared Survey Explorer satellite (Lang et al. 2016). Selection criteria were designed to target galaxies with $z > 0.6$, with the resulting final catalog containing 174,816 objects over $0.6 < z < 1.0$, distributed between a 2566 deg² NGC field and a 1676 deg² SGC field (shown in Figure 1). The corresponding comoving volumes are 2.2 and $1.4 h^{-3} \text{ Gpc}^3$, respectively.

The QSO sample (Lyke et al. 2020; Ross et al. 2020) is composed of objects observed during previous phases of SDSS and new objects selected from the same imaging data as the LRGs. The QSO catalog used for clustering (as opposed to the QSOs used for Ly α forest studies) contains 343,708 objects over $0.8 < z < 2.2$, covering the same two fields as the LRGs (with comoving volume $12 h^{-3} \text{ Gpc}^3$ for the NGC field and $8.0 h^{-3} \text{ Gpc}^3$ for the SGC field).

Figure 2 shows the redshift distribution of the LRG, ELG, and QSO samples for the NGC field, along with vertical bands indicating redshift ranges that are outside of the CHIME band (dark gray) or excluded owing to persistent RFI (light gray). Persistent RFI obscures 14.2% of the 587.5–800 MHz band, but we mask a larger fraction of the band (not shown in Figure 2) in this analysis, for reasons that will be described in Section 4.2.

The stack of the SGC catalog on the CHIME data is a factor of 3–3.5 times noisier than the stack on the NGC catalog for the same tracer of LSS. There are several reasons for this. First, in the case of the LRG and QSO catalogs there are 50% fewer sources in the SGC field compared to the NGC field. Second, we have less integration time on the SGC field because the range of R.A. occupied by the SGC field transits at CHIME at night in the summer time, whereas the NGC field transits at night in the winter

Table 1
The Redshift Distribution of Each Tracer Used in the Cross-correlation Analysis

Tracer	Frequency Range (MHz)	Source Number		z_{eff}	Redshift Range	
		Total	Nonzero Weight		$z_{\text{min}}-z_{\text{max}}$	$z_{0.16}-z_{0.84}$
LRG	585–800	39706	21615	0.84	0.78–1.00	0.81–0.87
ELG	585–800	63381	31181	0.96	0.78–1.10	0.83–1.03
QSO	585–800	94706	48046	1.20	0.80–1.43	1.00–1.36
QSOb0	700–800	26908	11960	0.97	0.80–1.03	0.85–1.01
QSOb1	650–700	23760	12311	1.12	1.03–1.19	1.07–1.16
QSOb2	585–650	44038	23775	1.30	1.19–1.43	1.23–1.39
QSOb00	745–800	11095	5299	0.84	0.80–0.91	0.82–0.87
QSOb01	700–745	15813	6661	0.99	0.91–1.03	0.96–1.01

Note. Frequency range gives the band that the analysis is limited to, whereas the redshift range gives the spread of source redshifts within that band. The $z_{0.16}-z_{0.84}$ span gives the 16% and 84% weighted percentiles giving an effective range within “ 1σ ” of the effective redshift (the weighted median of the source redshifts). For later analysis we further split the QSO catalog into subbands, denoted by the QSObX and QSObXY tracers.

time, when the nights are longer. Finally, the SGC field is at a lower decl. where the CHIME primary beam response is reduced and where we are forced to use a more aggressive delay filter because of aliasing of foregrounds. For these reasons, we have only a modest detection ($\sim 4\sigma$) of 21 cm emission when stacking on the QSOs in the SGC field, and we do not have a detection for the ELGs and LRGs in the SGC field. In what follows, we present the results for the NGC field only. We note, however, that our measurements in the SGC field are consistent with the amplitude of the 21 cm signal inferred from the catalogs in the NGC field, given the increased noise.

Each object in the eBOSS clustering catalogs includes weight values that account for imaging systematics, close pairs (which can be affected by spectroscopic fiber collisions), and the probability of a catastrophic redshift failure. We found that incorporating these weights into our analysis had a negligible impact on our results. Therefore, we do not employ the eBOSS weights in what follows. The weight given to each object is determined entirely by the sensitivity of the CHIME data at that object’s angular and spectral location.

2.3. Effective Redshift of Tracer Cross-correlations

The effective redshift, z_{eff} , of each catalog, when cross-correlated with CHIME data, is a combination of the redshift distribution of the sources in the catalog, the RFI mask used for the CHIME analysis, and the sensitivity of the CHIME data outside the masked regions. To determine z_{eff} , we first take each catalog, and for every source within it, we extract the inverse variance weights for that source in the processed CHIME data (these weights and how they are propagated through our pipeline will be described in Section 3). We then use these to construct a weighted median of the redshifts of the catalog. Similarly, to define an effective range of each catalog, we take the 16% and 84% weighted percentiles of the redshift distribution (i.e., the 68% equal-tailed interval), which gives a region within “ 1σ ” of the effective redshift. This differs substantially from the minimum–maximum redshift range where the source number density drops at the edges of the redshift distribution, most notably for the low-redshift end of the QSO distribution and the high-redshift end of the LRG distribution. These are all summarized in Table 1.

Some sources have zero weight owing to RFI masking and outlier cuts. In Table 1 we give the total number of sources within the frequency band being analyzed and an effective

source number, defined as the number of sources lying within a voxel with *nonzero* weight. Depending on the frequency range, this is typically $\sim 50\%$ of the total source number.

In Table 1 we also list five additional catalogs that are subdivisions of the QSO catalog, the largest and broadest redshift sample. The three catalogs QSOb0, QSOb1, and QSOb2 divide the redshift span into three roughly equal parts from lowest to highest redshift; the two catalogs QSOb00 and QSOb01 further divide the lowest-redshift catalog into two more catalogs. These additional catalogs will be used in later analysis of the data.

2.4. Coordinate Systems

CHIME is a transit instrument, and as such we are acutely sensitive to the precession of Earth’s polar axis. Historically, the celestial coordinate system has been anchored to the vernal equinox, which makes the coordinate system sensitive to both an unavoidable precession of Earth’s polar axis and an artificial shift in the zero-point of the R.A. coordinate (B1950 and J2000 coordinates are realizations of this anchored at their respective epochs).

The new system outlined in Petit & Luzum (2010) and Kaplan (2005) fixes some of these problems. The fundamental position of sources is given in International Celestial Reference System (ICRS) coordinates, which are fixed and unchanging coordinates that are essentially aligned with J2000 coordinates. Position as seen by an observer on Earth can be given in Celestial Intermediate Reference System (CIRS) coordinates, a frame in which the polar axis shifts with Earth’s precession and the R.A. origin is *minimally rotated*. Unlike previous equinox-based coordinates, CIRS coordinates only contain the minimal shift required to keep the polar alignment. As such, they are much more suited to use in CHIME: over a 5 yr period a typical equinox R.A. position shifts by $4/3$, or around a quarter of a CHIME pixel, whereas a CIRS position changes only by $1/5$, about $1/10$ of a pixel. This means that we are able to trivially align and average data products such as maps over much longer periods.

In this new system Greenwich Apparent Sidereal Time is replaced by Earth rotation angle. Instead of local sidereal time, we use *local Earth rotation angle*, which is equivalent to the current CIRS R.A. of the local meridian.

Throughout this paper the celestial coordinates we use will be CIRS coordinates at the average epoch of the data being analyzed, and any maps presented will be in those coordinates.

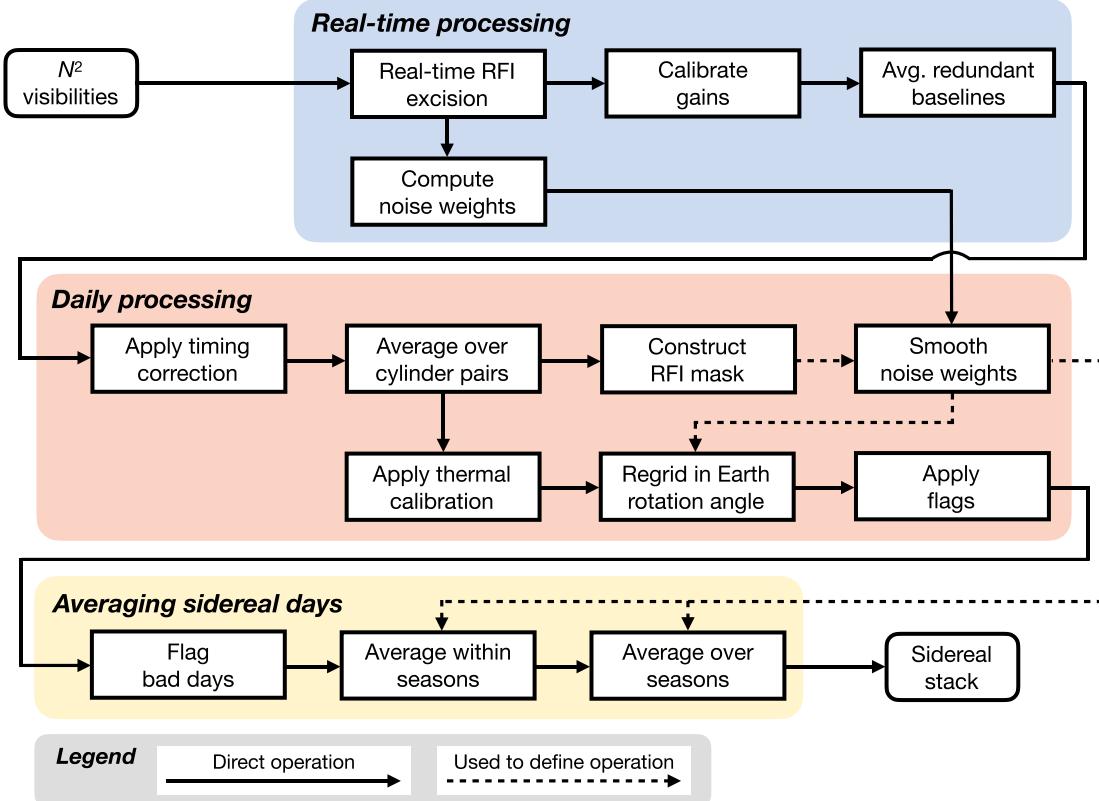


Figure 3. A schematic representation of the data-processing pipeline, starting from visibilities and culminating in the generation of the calibrated average over 102 nights that we refer to as a “sidereal stack.” The real-time pipeline (Section 3.1) performs spectral-kurtosis-based RFI excision (for a subset of the time span used in this work; see Section 3.1.1), gain calibration, and averaging over redundant baselines within each cylinder pair. It also computes fast-cadence noise estimates that are used as weights in several later steps. The daily processing pipeline (Section 3.2) applies a correction for clock drift between different ADCs in the CHIME F-engine, further averages redundant baselines over cylinder pairs, applies an ambient-temperature-dependent gain correction factor, regrids the time axis of each day onto a common grid in local Earth rotation angle, and applies data quality flags. This pipeline also constructs a time–frequency mask that targets longer-timescale RFI and incorporates this mask in a smoothing operation applied to the noise weights. Finally, we average over sidereal days (Section 3.3), first manually identifying and excluding bad days of data before forming several seasonal averages and then averaging these seasons together.

In the absence of better terminology, we will use sidereal day to refer to the interval between transitions of the Earth rotation angle through zero.

3. CHIME Data-processing Pipeline

The CHIME data-processing pipeline can be divided into two parts, the real-time and offline pipelines. The real-time pipeline runs on the CHIME correlator and supporting computing infrastructure. It operates on the digitized voltages measured by the 2048 antenna feeds, and outputs calibrated visibilities at 1024 frequency channels spanning the 400–800 MHz band. These are integrated to roughly 10 s cadence and further compressed by averaging over a subset of the redundant baselines. The offline pipeline runs on Compute Canada’s Cedar cluster. It operates on an archived copy of the visibility data and applies additional RFI masking and calibration, averages over all redundant baselines, interpolates onto a fixed grid in local Earth rotation angle, flags bad data, and averages over sidereal days. These real-time and offline operations, which produce the data product we refer to as a “sidereal stack,” are illustrated in Figure 3.

The CHIME real-time pipeline, offline pipeline, and analysis code used in this work is open source and publicly available. It can be found at <https://github.com/kotekan/>, <https://github.com/radiocosmology/>, and <https://github.com/chime-experiment/>.

3.1. Real-time Processing

We refer the reader to CHIME Collaboration et al. (2022b) for a description of the CHIME correlator, the real-time pipeline, and the archived data products. Below we highlight several aspects of the real-time processing that are relevant for interpreting what follows.

3.1.1. Real-time RFI Excision

CHIME Collaboration et al. (2022b) describe an RFI excision algorithm that runs on the CHIME correlator and is based on the spectral kurtosis statistic calculated at 0.66 ms cadence. This algorithm was deployed for a test period in 2019 June and then turned off until 2019 mid-October. Hence, the majority of the data (82 of 102 nights) used for this analysis did not benefit from fast-cadence RFI excision and rely entirely on the offline, ~ 10 s cadence excision algorithms that will be described in Section 3.2.3. This mixed data set is processed consistently in our analysis, but it simply has a higher rate of flagging in the offline pipeline for the days where real-time excision was not used.

3.1.2. Complex Gain Calibration

The complex gain of each feed is calibrated once per sidereal day by fitting a model to the eigendecomposition of the N_{feed}^2 visibility matrix during the transit of the brightest radio source

that is available at night. The primary calibration source is Cygnus A because it is the brightest radio point source in the sky between 400 and 800 MHz. It is also unresolved by the longest CHIME baselines and has a stable, well-characterized spectral flux density. Cassiopeia A, Taurus A, and Virgo A are used as alternative calibration sources when Cygnus A is transiting during the day. If a source other than Cygnus A was used for calibration, then the resulting gains are corrected for differences in the primary beam pattern of each feed at the location of the calibrator relative to the location of Cygnus A. This “beam ratio” is characterized by averaging the ratio of the gains from the two point sources over many nights. Hence, the complex gain calibration effectively normalizes the primary beam response at each frequency to unity on meridian at the decl. of Cygnus A.

The complex gains are scaled by the flux density of the calibrator source, such that application of the gains converts the visibility data to units of janskys per beam. The flux density of these sources was measured with the Karl G. Jansky VLA in 2014 and 2016 at frequencies ranging from 220 MHz to 48.1 GHz. VLA legacy observations from 1998 also exist at 73.8 MHz for all sources but Cassiopeia A. These measurements are interpolated to the CHIME band using the polynomial expressions provided in Perley & Butler (2017, hereafter P17). The uncertainty on the relative spectral flux density of the calibration sources in the CHIME band is less than 1%. The absolute flux of the P17 scale at these frequencies is determined by measurements of Cygnus A by Baars et al. (1977), which the authors estimate is accurate at 3%–5%.

3.1.3. Compression

The CHIME feeds are located on a regular grid, and as a result the N_{feed}^2 visibilities contain many redundant measurements for each baseline. In order to compress the data, the real-time pipeline performs a weighted average of all redundant baselines formed from feeds on the same pair of cylinders. Correlator inputs that are malfunctioning or otherwise anomalous are identified and flagged in semi-real-time using 10 different tests based on a variety of data products and housekeeping metrics. The weight given to a particular baseline is 0 if either of the inputs that form the baseline is currently flagged and 1 otherwise. This uniform weighting scheme will result in lower sensitivity compared to an inverse variance weighting scheme that accounts for feed-to-feed differences in the noise referred to the sky. We estimate that the magnitude of this degradation in sensitivity is approximately 5%. Note that redundant baselines formed from feeds on different pairs of cylinders are not averaged at this stage. This baseline collation strategy allows for cylinder-dependent corrections and calibrations to be applied offline. Below, we refer to the resulting visibility for baseline \mathbf{b} at frequency ν and time t as $V_{\text{raw}}(\mathbf{b}, \nu, t)$.

3.1.4. Weights

The real-time pipeline estimates the variance of the visibility for each baseline, frequency channel, and ~ 10 s integration by differencing the even and odd 30 ms subintegrations. Since the observed foregrounds do not change significantly on 30 ms timescales, they cancel for this difference, leaving contributions from RFI and intrinsic radiometric noise. This “fast-cadence

estimate” of the variance is propagated through each stage of the real-time and offline pipeline.

In general, there is percent-level agreement between the fast-cadence estimate of the variance and the radiometric estimate calculated from the measured autocorrelation, frequency channel width, and total integration time. Most cases where the two estimates differ correspond to known periods of bad data quality or have a temporal and spectral extent that is characteristic of transient RFI. The fast-cadence estimate is used to construct inverse variance weights that are used to average over sidereal days, average over baselines during mapmaking, and average over sources when stacking on external catalogs. The inverse variance weights are *not* used to average over redundant baselines; instead, we use the uniform weighting scheme described in Section 3.1.3.

3.2. Daily Processing

Here we describe the daily pipeline that applies additional processing to a copy of the archived visibility data. This includes correcting for clock drift, averaging over redundant baselines on different cylinder pairs, identifying and masking RFI, correcting common-mode thermal variations in the amplitude of the gain, interpolating the data onto a common grid in local Earth rotation angle, and finally masking ranges of time with poor data quality. We briefly describe each of these stages. The primary data product output by this processing is the visibility for all unique baselines on each local sidereal day as a function of local Earth rotation angle at $\Delta\nu = 0.390625$ MHz spectral resolution.

3.2.1. Timing Correction

The sampling rate of the analog-to-digital converters (ADCs) that digitize the signal measured by the CHIME feeds is derived from a 10 MHz clock that originates from a GPS-disciplined, oven-controlled crystal oscillator and is distributed to the circuit boards that house the ADCs through a hierarchical network consisting of coaxial cables, power splitters, and amplifiers. Thermal susceptibility of this distribution network results in copies of the clock drifting with respect to one another on timescales set by the different refrigeration cycles of the water chillers used to control the temperature of the electronics. The magnitude of this effect is particularly large between copies of the clock provided to ADCs in different receiver huts, which are temperature controlled by independent chillers.

The CHIME ADCs are housed in eight electronics crates. The thermal drift between the eight copies of the clock that are distributed to the eight crates is measured using a broadband noise source following the procedure described in CHIME Collaboration et al. (2022b). This yields a proxy for the drift, $\delta\tau_{cd}^{\text{clock}}(t)$, between the copies of the clock provided to electronics crate c relative to electronics crate d . The visibility is then corrected as follows:

$$V_{\text{cal},1}(\mathbf{b}, \nu, t) = \exp \{ -2\pi j \nu \langle \delta\tau_{cd}^{\text{clock}}(t) \rangle_{cd \in \mathbf{b}} \} \times V_{\text{raw}}(\mathbf{b}, \nu, t). \quad (1)$$

Here $\langle \delta\tau_{cd}^{\text{clock}}(t) \rangle_{cd \in \mathbf{b}}$ is constructed by averaging the estimates of the relative clock drift between the pairs of crates that digitize the pairs of inputs that form every redundant baseline averaged by the real-time pipeline to obtain $V_{\text{raw}}(\mathbf{b}, \nu, t)$. Applying this correction reduces the standard deviation of the

delay noise on timescales less than 20 minutes from 4.25 to 1.5 ps on average, as inferred from the phase stability of the signal from bright point sources. Note that further improvements have been achieved by applying a more complicated, ADC-dependent correction in real time, but the analysis described in this work uses the simpler, offline correction described above.

3.2.2. Redundant Baseline Collation

The timing correction described in the previous section is the only cylinder-dependent correction that was applied for this analysis. The next stage of the pipeline averages *all* redundant baselines by performing a weighted average over the redundant baselines measured by different cylinder pairs. The weighting scheme used is consistent with the scheme used by the real-time pipeline. Specifically, each cylinder pair is weighted by the number of redundant baselines that were previously averaged by the real-time pipeline. These weights are constructed from the set of correlator input flags that were used by the real-time pipeline at each time sample.

3.2.3. RFI Excision

Narrowband RFI will contaminate the high-delay modes that our analysis relies on to avoid the spectrally smooth foregrounds. Hence, identifying and masking times and frequency channels that are corrupted by RFI is critical to detect the 21 cm signal. The RFI excision occurs in three stages, with each stage generating a single 2D mask in (frequency, time) that is applied to the weight data set for all baselines before proceeding to the next stage. The first stage masks any frequency channel that coincides with a known, persistent source of RFI. There were two sources of persistent RFI in the 587.5–800 MHz band: the mobile LTE bands and the local oscillator (LO) used by the Synthesis Telescope at the Dominion Radio Astrophysical Observatory (DRAO; Landecker et al. 2000). These two sources occupy 14.2% of the band.

The second stage creates a mask by identifying variations in the autocorrelation that have a spectral and temporal extent characteristic of RFI. The average autocorrelation over the 2048 inputs is normalized at each frequency by the median value over the local sidereal day to remove static variations in the bandpass. The median and median absolute deviation (MAD) are then calculated over a 2D moving window in (frequency, time) of size (10 MHz, 7 minutes). Any time and frequency where the autocorrelation deviates from the median by more than 5 times the MAD over the window centered on its location is masked. The window size was calibrated by first identifying RFI events through manual inspection of the autocorrelations acquired on a few typical days and then searching for a window that maximized the fraction of RFI corrupted data that is masked while minimizing the amount of clean sky data that is masked.

The third stage creates a mask by identifying RFI-like variations in the visibility data from the cross-polar, intracylinder baseline with 10 m separation. This baseline has a relatively large number of redundant copies and thus low radiometric noise compared to most other baselines. RFI events are more easily discriminated from the background radio sky in a cross-polar visibility because the RFI is in general polarized, whereas the radio sky is largely unpolarized at the scales at which the 10 m baseline is sensitive. The algorithm for

identifying RFI events is similar to the algorithm applied to the autocorrelations. The median is calculated over a 2D moving window in (frequency, time) of size (4.3 MHz, 1 minute) and subtracted to remove background radio emission from the sky. The MAD is then calculated over a 2D window of size (8.2 MHz, 7.4 minutes). Any time and frequency where the visibility deviates from the median by more than 5 times the MAD over the window centered on its location are masked. The window sizes were chosen using a procedure similar to that described in the previous paragraph.

One common source of transient RFI arises from the reflection of distant broadcast TV channels off meteor ionization trails and aircraft. These appear in known 6 MHz wide bands and last \sim 5 s. A targeted search for these events is performed on the moving median-subtracted, cross-polar visibility by identifying time samples where the majority of frequencies within each TV channel are outliers. The entire TV channel is masked if more than 50% of the frequencies within that TV channel exceed 1.8 times the moving MAD. This results in a false-positive rate equal to the 5 M AD cut used in the standard third-stage excision, assuming a Gaussian noise model.

3.2.4. Thermal Calibration

Common-mode variations in the amplitude of the complex receiver gain are corrected using a linear regression model based on measurements of the outside temperature. Details of the model construction and an evaluation of its performance are provided in CHIME Collaboration et al. (2022b). To briefly summarize, fractional variations in the amplitude of the complex gain inferred from hundreds of bright point-source transits are regressed against the outside temperature as measured by the DRAO weather station at the time of transit. The resulting thermal susceptibility increases with frequency from $0.07\% \text{ K}^{-1}$ at 400 MHz to $0.2\% \text{ K}^{-1}$ at 800 MHz and varies across inputs at the $0.05\% \text{ K}^{-1}$ level. The susceptibility is averaged over the 2048 inputs, and the frequency dependence is fit to a quadratic function. The visibility measured by baseline \mathbf{b} at frequency ν and time t is then corrected as follows:

$$V_{\text{cal},2}(\mathbf{b}, \nu, t) = [1 + \alpha(\nu)(T(t) - T(t_*))]^2 V_{\text{cal},1}(\mathbf{b}, \nu, t), \quad (2)$$

where $\alpha(\nu)$ is the quadratic model for the thermal susceptibility, $T(t)$ is the outside temperature at time t , and $T(t_*)$ is the outside temperature at time t_* at which the complex gain calibration was derived by the real-time pipeline. The quantity t_* is a step function that changes once per sidereal day to the most recent time of transit of the calibrator source. This procedure improves the stability from roughly 0.8% to 0.5% (standard deviation in fractional power units) by correcting the common-mode drift in the amplitude caused by changes in the outside temperature between daily point-source calibrations.

3.2.5. Weight Smoothing

The radiometric noise is not expected to change appreciably on short timescales. In order to reduce the uncertainty on our estimate of the variance of the radiometric noise and also make our estimate less sensitive to transient RFI events, a rolling median filter with a 5-minute window is applied to the time

axis of the inverse variance weight data set. Any time that was masked is ignored when calculating the median and also remains masked after the filtering is applied.

3.2.6. Sidereal Regridding

The next stage of the daily processing pipeline interpolates the visibilities onto a fixed grid in local Earth rotation angle ϕ that ranges from 0° to 360° with 4096 samples, giving a spacing $d_\phi = 5.27$. The interpolation algorithm assumes that the processed visibilities, $V_{\text{cal},2}(\mathbf{b}, \nu, t)$, are sampled from some regularly gridded sky visibility, $V_{\text{grid}}(\mathbf{b}, \nu, \phi)$, and corrupted by both noise and RFI, denoted as $n(\mathbf{b}, \nu, t)$. Since the sky visibility is band limited by its maximum fringe rate and the chosen sample rate is more than twice that, the following relation holds:

$$V_{\text{cal},2}(\mathbf{b}, \nu, t_g) = \sum_h K_{gh} V_{\text{grid}}(\mathbf{b}, \nu, \phi_h) + n(\mathbf{b}, \nu, t_g), \quad (3)$$

where $K_{gh} = \text{sinc}\left(\frac{\Delta\phi_{gh}}{d_\phi}\right)$ is the interpolation kernel with $\Delta\phi_{gh} = \phi(t_g) - \phi_h$, and the summation runs over the regular grid in local Earth rotation angle. The infinite support of the kernel is computationally problematic, so a common approximation involves truncating the kernel by multiplying it with a window function. We use the Lanczos kernel, which is given by

$$K_{gh} = \begin{cases} \text{sinc}\left(\frac{\Delta\phi_{gh}}{d_\phi}\right) \text{sinc}\left(\frac{\Delta\phi_{gh}}{ad_\phi}\right) & |\Delta\phi_{gh}| \leq ad_\phi \\ 0 & |\Delta\phi_{gh}| > ad_\phi. \end{cases} \quad (4)$$

The parameter a controls the kernel width and couples at most $4a + 1$ samples of V_{grid} .

We use a Wiener filter to invert Equation (3) and solve for the regularly gridded sky visibility, given the noisy, RFI contaminated data. Let $\mathbf{v}_{\text{cal},2}$ denote the vector containing the time-ordered visibility for a given frequency and baseline. The regularly gridded visibility \mathbf{v}_{grid} for that frequency and baseline is estimated as

$$\hat{\mathbf{v}}_{\text{grid}} = \mathbf{C} \mathbf{K}^T \mathbf{N}^{-1} \mathbf{v}_{\text{cal},2}, \quad (5)$$

where

$$\mathbf{C} = (\mathbf{S}^{-1} + \mathbf{K}^T \mathbf{N}^{-1} \mathbf{K})^{-1}, \quad (6)$$

$\mathbf{N}^{-1} = (N_{\text{noise}} + N_{\text{RFI}})^{-1}$ is the inverse covariance of the noise and RFI, and \mathbf{S}^{-1} is the inverse covariance of the sky visibility. The noise covariance is assumed to be diagonal and equal to the fast-cadence estimate of the variance described in Section 3.1.4. The RFI covariance is also assumed to be diagonal and equal to infinity for times and frequencies that are missing or have been masked by the procedure described in Section 3.2.3 and equal to zero otherwise. Finally, the sky covariance is assumed to be diagonal and constant as a function of baseline, frequency, and sidereal angle, such that $\mathbf{S}^{-1} = s_{\text{max}}^{-2} \mathbf{I}$, where \mathbf{I} is the identity matrix and $s_{\text{max}} = 1 \times 10^4 \text{ Jy beam}^{-1}$ is chosen to be around the maximum flux observed on the sky. Each frequency and baseline is solved independently. This is made computationally tractable by utilizing the fact that \mathbf{C}^{-1} is a band matrix, which is a consequence of the compact support of the Lanczos kernel.

Choosing the kernel width, a , is a balance between the computational cost of the regridding (which is $O(a^2)$) and the accuracy of the reconstruction. We use $a = 5$ in this work, which has deviations away from the ideal sinc transfer of $\lesssim 10^{-3}$ for the typical range of fringe rates in this analysis.

The covariance of the filtered signal is $\langle \hat{\mathbf{v}}_{\text{grid}} \hat{\mathbf{v}}_{\text{grid}}^\dagger \rangle = \mathbf{C}$ as given by Equation (6). The weight data set that tracks the inverse variance of the noise present in the visibilities is therefore updated to $w(\mathbf{b}, \nu, \phi_h) = 1/C_{hh}(\mathbf{b}, \nu)$. The interpolation scheme introduces ringing in the visibilities at the edge of any large gap of missing or masked data, with the post-interpolation weights at these edges gradually transitioning to s_{max}^{-2} . In order to mask these artifacts, we apply a baseline-dependent threshold to the weight data set, setting it to zero if it is less than 50% of the average weight over all frequencies and sidereal angles. As these samples lie at the edge of large periods of missing data, the relative increase in the amount of data flagged is small.

3.2.7. Daytime, Moon, and Data Flags

Next, we apply a series of flags that exclude certain time ranges from further analysis. The weight data set is set to zero for any time sample that meets one or more of the following criteria: (1) the Sun is above the horizon (52% flagged), (2) the Moon is within 5° of the meridian (3% flagged), (3) the data quality is poor as indicated by a “bad data flag” in our database (36% flagged).

The database that is used for the last item is updated external to the daily pipeline and contains a variety of flag types based on different metrics for data quality. The following data flag types were employed in this analysis:

Rain: Mask any time where the accumulated rainfall during the 30 hr prior was greater than 1 mm. This condition finds intervals where a large number of feeds are likely to be wet. Precipitation at the site causes analog signal corruption in 4%–12% (interquartile range) of the feeds due to water pooling on the focal line (CHIME Collaboration et al. 2022b; 10% flagged).

Jumps: Mask any time where the autocorrelation for five or more feeds has shown a sudden ($\lesssim 30$ minutes), broadband increase of more than 20% in the past 30 hr. This condition, too, is designed to find intervals where a large number of feeds are likely to be wet (29% flagged).

Correlator restart: Mask the interval between a correlator restart and the next daily point-source calibration. The FPGA resynchronization that occurs during a correlator restart introduces a change in the relative phase between feeds digitized by different ADC chips that is nonnegligible with respect to our requirements on phase stability (7.1% flagged).

Acquisition restart: Mask the interval between a restart of the data acquisition software and application of the calibration gains (2.8% flagged).

Bad calibration: Mask any time where the calibration gains were not updated in the past 24 hr. Also mask intervals where poor-quality gains were applied to the visibility data as determined by several metrics that are generated by the real-time pipeline and monitored by the telescope operator (0.4% flagged).

CHIME acquired 245 days of integration time during the 309-day period between 2019 January 1 and November 6. The

64 days of instrument downtime consisted of 55 days of planned hardware maintenance and software upgrades and 9 days of unintended interruptions due to power failures, cooling failures, and other accidental outages. The flags described above exclude 70% of the remaining data from the stacking analysis, with the daytime and rain/jumps flags representing the primary sources of data loss. After applying these flags, the total integration time is 1760 hr, of which 834 hr was spent observing the range of R.A. containing the NGC field. This total is further reduced by the sidereal day flags that will be described in Section 3.3.1.

3.3. Averaging Sidereal Days

After the individual days have been flagged and processed to a common grid in local Earth rotation angle, the days are averaged together to produce a high-sensitivity measurement of the sky.

This process is complicated by the presence of *noise crosstalk*, a bias in the zero-level of a non-autocorrelation visibility. Physically this requires both signal chains being correlated to share a common source of thermal noise. We expect there to be two major contributors to this: one is the leaking of thermal noise generated within the low-noise amplifier on one signal chain that is broadcast by the antenna and received (directly, or by an indirect path) by another antenna; another source is antennas seeing thermal emission from the ground by a direct path or by scattering or diffraction from nearby structures. The correlation of this common source of noise yields a bias in the visibility between the two antennas. A thorough investigation of the source of the crosstalk is beyond the scope of this paper; however, we will discuss the crosstalk briefly in Section 3.3.4.

We observe the crosstalk to be relatively stable in time, varying slowly over the course of 1 day (we quantify this in Section 3.3.4). In practice, this allows the crosstalk removal to be performed by estimating and removing a single time-independent signal from each day for each frequency and baseline. However, as the crosstalk signal is not known a priori and must be measured from the data, it is degenerate with any constant sky signal within the time period being used to estimate it.

As we use only nighttime data spread over a year, there is no single period in common between all days that we can choose as a reference. To account for this, we break the sidereal averaging into two stages: the first operates on data taken from each quarter of the year, and the second combines those into a full stacking of the data.

3.3.1. Sidereal Day Flags

Prior to averaging the sidereal streams, we make further cuts to the data. Any sidereal day with less than 80% of the day remaining after applying the data flags in Section 3.2.7 is rejected, as is any day where less than half the crosstalk reference range is available (see next section).

Finally, each day is manually inspected via a standardized set of visualizations:

1. A delay power spectrum for each baseline generated by averaging over all unmasked time samples (see Appendix A). This presents a holistic summary of all

elements of the data and is particularly powerful for illustrating poor RFI flagging and misbehaving baselines.

2. A sensitivity plot showing the estimated point-source flux sensitivity found by appropriately averaging the fast-cadence estimate of the variance over all baselines at each time and frequency. This is another summary of the whole data set and is a good diagnostic of RFI excision performance.
3. A sky map (see Section 4.3) at two different frequencies and its difference from a day-averaged map. This is not a complete summary, as it does not incorporate information from every frequency, but is very effective at identifying poor calibration.

Each day was inspected by at least two people, and any day flagged by at least one person was removed from further analysis.

After all these cuts are applied, 102 sidereal days remain for averaging. After also applying the flags described in Section 3.2.7, the 102 sidereal days contain 1073 hr of integration time. Of this, 521 hr was spent observing the range of R.A. containing the NGC field.

3.3.2. Sidereal Averaging (Seasonal)

The first stage of sidereal averaging combines data from a single quarter of each calendar year and assigns each “good” day of data into alternating partitions of the data. By splitting into partitions per quarter, we are able to produce two jackknife splits of our data that have approximately the same sensitivity and sidereal coverage; these will be used for consistency tests in Section 7.2.

For each quarter, we pick a single hour-long range in local Earth rotation angle that is observed within the nighttime for the entire quarter and avoids the transits of bright point sources. This time range is used to reference the crosstalk signal for the entire quarter. We illustrate these ranges for each quarter, as well as how the quarters overlap with the eBOSS NGC field, in Figure 4.

Every day we calculate the median over this time range for each visibility and subtract it from the data for that day. Assuming that the crosstalk signal is approximately constant across the day, this procedure will remove that day’s crosstalk contamination and a small amount of the sky signal, which is the same across all days within the quarter. It is important to use consistent estimates of the crosstalk; therefore, if more than 70% of the data within this reference range are missing for a frequency on a given day, the entire frequency will be flagged out for the whole day. This differs from the initial selection discussed in Section 3.3.1, as it is determined from the full frequency-dependent missing data mask for that day, not just the frequency-independent data flags.

After the crosstalk has been removed consistently from all days within the quarter, the days within each partition are averaged together with an inverse variance weighting.

3.3.3. Sidereal Averaging (All)

The second stage of sidereal averaging is to combine the data for all quarters and partitions. As the crosstalk removal uses a different sky reference region for each quarter, a simple averaging would introduce discontinuities at the boundaries. To account for this, we exploit the overlap in local Earth rotation

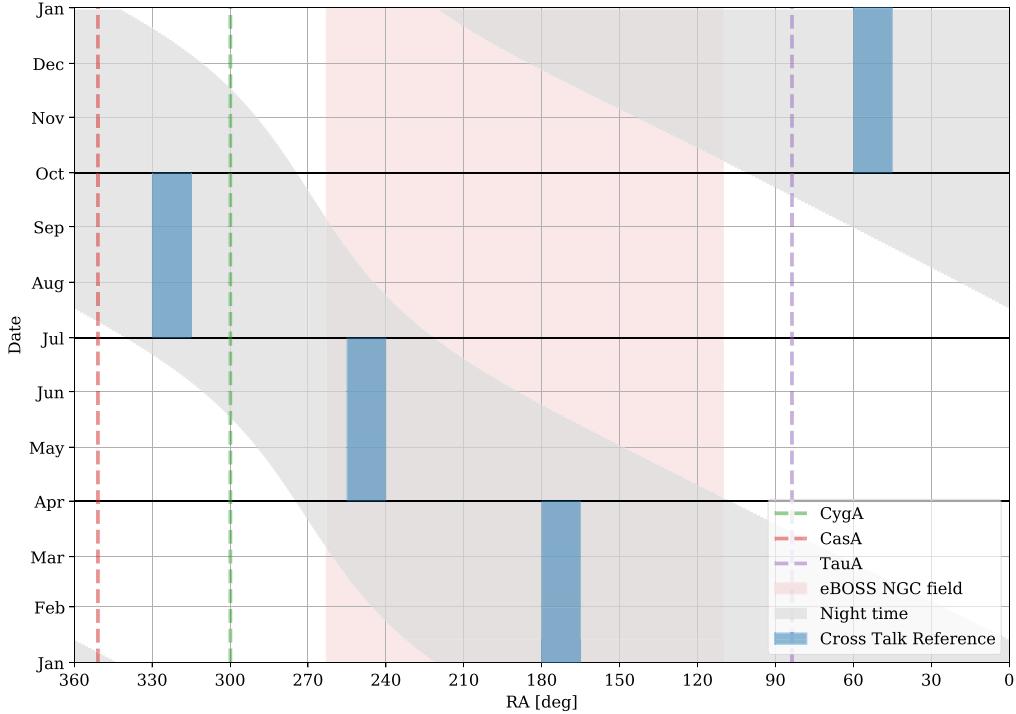


Figure 4. In the first stage of combining individual days, we combine sidereal days within each quarter of the year. The figure above shows the R.A. range of nighttime data for each day within a quarter (gray shaded region). The bulk of the sensitivity to the eBOSS NGC field (pink band) comes from the first two quarters of the year, where the local nighttime better overlaps with the NGC R.A. range. In order to combine the days, we need to consistently reference the mean level of each day to remove crosstalk. For each quarter, we pick a single hour of local Earth rotation angle (blue boxes) for which we compute the median and subtract it from each day’s sidereal stream. These regions are chosen to be within the nighttime and avoid the transit of bright point sources, in order to minimize the bias from gain variations.

angle of the nighttime data for each quarter with its neighbors to solve for the differences and set a common reference.

To do this, we treat our estimate of the regularly gridded visibility $v_{\text{grid},i}$ (where we have dropped the \wedge symbol to simplify notation) for each frequency and baseline within a partition i (out of p total partitions) as being composed of a signal v that we are interested in that is constant for all partitions, a noise n_i , and a residual crosstalk contribution x_i that is different for each partition and also incorporates the bias from the per-partition crosstalk referencing. We write this as

$$v_{\text{grid},i} = v + n_i + x_i. \quad (7)$$

We model the statistics of each component as having zero mean with covariance matrices $\langle vv^\dagger\rangle = S$, $\langle n_i n_i^\dagger\rangle = N_i$, and $\langle x_i x_i^\dagger\rangle = X$. As the crosstalk has little time variation, we model the residuals as a low-rank contribution (with rank k), allowing us to factorize the covariance as $X = UU^\dagger$, where U is a rectangular matrix. Though the crosstalk referencing means that the modes x_i may be very different, we assume that the crosstalk statistics are the same across all partitions, so X does not depend on i . The noise matrix N_i is assumed to be diagonal and includes both the noise expected in the data (Section 3.1.4) and any masking that has been applied (Sections 3.2.3 and 3.2.7), encoded in the standard way of setting the inverse variance to zero for masked samples. Although the averaging over sidereal days has reduced the number of samples that are flagged entirely, ranges of R.A. observed during the daytime for the entire quarter and badly RFI-contaminated frequencies will still be masked.

To solve for the signal, we start by writing a Wiener estimator for s treating both n_i and x_i as a generalized noise

$$\hat{v} = \mathbf{C} \left[\sum_i (N_i + X)^{-1} v_{\text{grid},i} \right], \quad (8)$$

where the covariance matrix \mathbf{C} is defined by

$$\mathbf{C}^{-1} = \mathbf{S}^{-1} + \sum_i (N_i + X)^{-1}. \quad (9)$$

A naive application of this scheme would require tracking and inverting a large matrix for each frequency, but we can simplify it by repeated application of the Woodbury matrix identity.²³ First, we expand the $N_i + X$ term, allowing us to regroup Equation (9) as

$$\mathbf{C}^{-1} = \mathbf{C}_0^{-1} + \mathbf{W}\mathbf{W}^\dagger, \quad (10)$$

where

$$\mathbf{C}_0^{-1} = \mathbf{S}^{-1} + \sum_i N_i^{-1} \quad (11)$$

and \mathbf{W} is a block matrix,

$$\mathbf{W} = (\mathbf{W}_0 \mid \mathbf{W}_1 \mid \dots \mid \mathbf{W}_{p-1}), \quad (12)$$

with one block for each partition, and within each block are k columns for each crosstalk mode and a row for each R.A.

²³ The Woodbury matrix identity allows us to expand the inverse of a low-rank update to a matrix with known inverse. In its most general form it is written as $(\mathbf{A} + \mathbf{UCV})^{-1} = \mathbf{A}^{-1} - \mathbf{A}^{-1}\mathbf{U}(\mathbf{C}^{-1} + \mathbf{V}\mathbf{A}^{-1}\mathbf{U})^{-1}\mathbf{V}\mathbf{A}^{-1}$, with \mathbf{A} and \mathbf{C} square, but potentially different sizes.

sample. The blocks are

$$\mathbf{W}_i = \mathbf{N}_i^{-1} \mathbf{U} (\mathbf{I}_k + \mathbf{U}^\dagger \mathbf{N}_i^{-1} \mathbf{U})^{-1/2}, \quad (13)$$

where \mathbf{I}_k is the identity matrix of size k , and each \mathbf{W}_i can be interpreted as a noise-weighted projection operator onto the crosstalk basis for each partition.

The estimator in Equation (8) can be rewritten as

$$\hat{\mathbf{v}} = \mathbf{C} \left[\sum_i (\mathbf{N}_i^{-1} - \mathbf{W}_i \mathbf{W}_i^\dagger) \mathbf{v}_{\text{grid},i} \right]. \quad (14)$$

A second application of the Woodbury identity, this time to Equation (10), allows us to write \mathbf{C} in a more easily applied form:

$$\mathbf{C} = \mathbf{C}_0 + \mathbf{C}_0 \mathbf{W} (\mathbf{I}_{(k \times p)} - \mathbf{W}^\dagger \mathbf{C}_0 \mathbf{W})^{-1} \mathbf{W}^\dagger \mathbf{C}_0. \quad (15)$$

To produce the final estimate for the stacked signal, we need to generate \mathbf{W}_i for each day and retain it, accumulate $(\mathbf{N}_i^{-1} - \mathbf{W}_i \mathbf{W}_i^\dagger) \mathbf{v}_{\text{grid},i}$ to generate $\mathbf{C}^{-1} \hat{\mathbf{v}}$, and then finally apply the deconvolving matrix \mathbf{C} , which can be done efficiently by evaluating matrix-vector products from right to left in Equation (15) using the accrued \mathbf{W}_i rather than explicit construction of \mathbf{C} . Conceptually this final step uses the noise-weighted overlaps between the different partitions (in the $\mathbf{I} - \mathbf{W}^\dagger \mathbf{C}_0 \mathbf{W}$ term) to solve for a consistent bias and remove it.

In the implementation within our pipeline we model the crosstalk as a single time-independent constant mode per day (i.e., $k = 1$ and $\mathbf{U} \propto \mathbf{1}$, where $\mathbf{1}$ is a column vector filled with ones). We also assume that both X and S are much larger than the instrumental noise N for unmasked samples. This means that the estimator we use does not depend on S at all, nor on the scale of X (but it does depend on the form), and, importantly, means that \mathbf{C}_0 is a diagonal matrix. However, this does produce one singular mode, the sidereal average of each visibility, that must be regularized externally. Finally, rather than using the N for each baseline, we use an average over all baselines, which ensures that the same linear combinations of partitions are used for all baselines at a given frequency. We use these same linear combinations when updating the baseline-dependent weights in the final stack, although we drop the small correction to the weights that comes from removing the crosstalk, which primarily affects the off-diagonal elements of the noise covariance that we do not track in our analysis, for memory reasons.

As the sidereal-time-independent component of the sky is entirely degenerate with a constant noise bias, the mean of each sidereal stream is a singular mode. To regularize this degenerate mode, we add a constant offset to set the median in time of the full sidereal day to zero.

3.3.4. Crosstalk Properties

Throughout this analysis we have made the strong assumption that we can model the crosstalk as a constant mode per day. To test the validity of this, we use the estimate of the true sky \hat{V} produced above and compare it with the data for each individual day $V_{\text{grid},i}$. Assuming that the true sky estimate is faithful, this allows us to estimate the time-dependent crosstalk within that day. Practically, we do this by regressing these inputs against each other in short time intervals for each

baseline and frequency, i.e., we fit a model

$$V_{\text{grid},i}(\mathbf{b}, \nu, \phi) = G(\mathbf{b}, \nu) \hat{V}(\mathbf{b}, \nu, \phi) + X(\mathbf{b}, \nu), \quad (16)$$

where G is an overall scaling factor and X is the crosstalk estimate for the interval. We use 64 intervals over the sidereal day, which gives a time resolution of ~ 22 minutes. We exclude daytime intervals from our fit, as well as periods around bright source and lunar transits.

The accuracy of the inferred crosstalk is limited by the fact that we do not have an independent estimate of the true sky, and so we depend on the averaging over days that produced it to suppress time/R.A.-dependent crosstalk variation, and that assuming that the visibilities of the true sky have zero median does not significantly bias the crosstalk estimate.

Using the extracted crosstalk estimates across each day, we compute three summaries of the time dependence, which we show in Figure 5. The first is the mean crosstalk calculated over the full time range analyzed in this paper; the second is the *interday variation*, which is defined as the variance over days of the mean of each sidereal day; and the final summary is the *intraday variation*, which is the mean over days of the variance within each sidereal day. In all cases we have used a simple outlier cut on the intraday variance to remove anomalous time samples and frequencies, though some residual contributions remain that generate the baseline-independent horizontal banding visible in both the inter- and intraday variation. The following summary statistics we quote are all for instrumental Stokes I and at a frequency (≈ 613.7 MHz) observed to have low levels of interference.

As we might expect, the crosstalk is significantly stronger for baselines within a cylinder, with the rms crosstalk taken over north-south (NS) baselines ~ 30 times larger than between neighboring cylinders and ~ 70 and ~ 130 times larger than the two- and three-cylinder separations. For the intracylinder crosstalk, the signal is strongly concentrated in the shortest baselines, with 90% of the cumulative rms coming from baselines $\lesssim 5.5$ m and 99% from $\lesssim 12.8$ m. A delay-space analysis indicates that most of the crosstalk contributions correspond to specific path lengths. Within a single cylinder, the major contributions appear to be from the indirect paths between two feeds, with one reflection from the cylinder vertex (with smaller contribution from paths with multiple reflections); between cylinders, the major contribution is from the direct geometric path between the feeds, with smaller contributions from this path combined with a straight up-down reflection of the signal at either feed. These observations lead us to believe that the dominant crosstalk mechanism is from rebroadcast amplifier noise, though we cannot rule out a contribution from ground pickup. Regardless, the crosstalk removal and analysis that follows do not depend on the origin.

The importance of the time dependence of the crosstalk differs depending on the cylinder separation, with the fractional contributions being larger at longer cylinder separations, where the typical crosstalk is lower. Overall we find that the rms intraday crosstalk variations are $\approx 3\%$, 5% , 9% , and 15% of the mean level for the zero-to-three-cylinder separations, respectively, and the interday variations $\approx 5\%$, 8% , 10% , and 17% . We believe that the low intraday variation justifies our decision to use a single crosstalk reference region in each day, but the small separation between the inter- and intraday variation means that to make further improvements we will need to account for the time variation across each day.

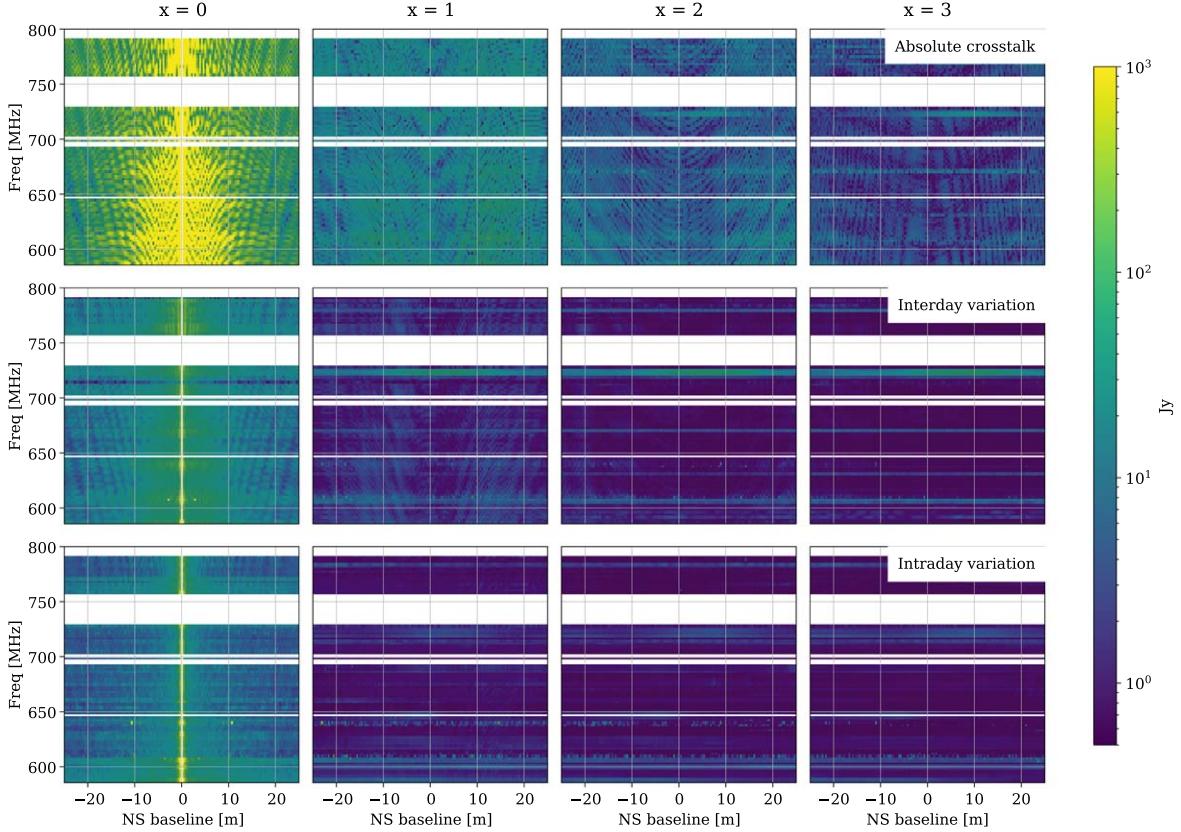


Figure 5. The crosstalk observed in CHIME has complex spatial, spectral, and temporal behavior. In the top row we show the magnitude of the time-averaged crosstalk in instrumental Stokes I as a function of cylinder separation (columns), north-south baseline length (horizontal axis), and frequency (vertical axis). In the bottom two rows we compare the time variability calculated on different scales: the middle row shows the variation observed on timescales longer than a day, and the bottom row shows scales less than a day. In both cases the amount of variation is much lower than the average crosstalk level, with the typical fractional variation $\lesssim 5\%$ for the zero cylinder separation, and slightly higher at larger separations. The intraday variation is particularly susceptible to RFI and instrumental issues, leading to the excess variation seen in horizontal bands.

4. Stacking Pipeline

We have developed a dedicated pipeline to stack the CHIME data on the angular and spectral locations of the sources in a spectroscopic catalog. The pipeline takes as input the sidereal stack that is generated by the CHIME data-processing pipeline as described in Section 3. It subtracts the signal from the four brightest point sources and masks corrupted frequency channels. Next, it constructs a map of the sky at each frequency channel, deconvolving a model for the primary beam pattern in the process. It applies a high-pass filter to the frequency axis of each map pixel to remove foregrounds. It then masks frequency channels and pixels that are outliers. Finally, it stacks the maps on the angular and spectral locations of the sources in a catalog. The entire process is visualized in Figure 6. In what follows, we describe each stage of the pipeline.

4.1. Point-source Subtraction

The CHIME feeds have nonnegligible off-axis response. As a result, the signal from the four brightest point sources—Cygnus A, Cassiopeia A, Taurus A, and Virgo A—is significant whenever these sources are above the horizon. The first stage of the stacking pipeline performs a targeted removal of these four sources. All other foregrounds are removed from the data using a spectral filter that will be described in Section 4.5. Prior to filtering, the spectral structure introduced by the on-axis response of the instrument is calibrated (see Section 4.4) and deconvolved from the data (see

Section 4.3). The effectiveness of this calibration will be quantified in Section 7.3 by comparing the recovered flux density of a set of standard calibrators to measurements made by other telescopes.

The following model is assumed for the contribution of the four brightest point sources to the visibility measured by baseline \mathbf{b} at local Earth rotation angle ϕ and frequency ν :

$$V_{\text{psrc}}(\mathbf{b}, \nu, \phi) = \sum_{s=1}^4 a_s(\nu, \phi) e^{j2\pi\nu\mathbf{b} \cdot \hat{\mathbf{n}}(\theta_s, \phi - \phi_s)/c}, \quad (17)$$

where $a_s(\nu, \phi)$, θ_s , and ϕ_s denote the primary-beam-modulated amplitude, decl., and R.A. of source s , respectively; $\hat{\mathbf{n}}$ is the unit vector pointing toward the source's location; and c is the speed of light. At every frequency and local Earth rotation angle we estimate the set of source amplitudes \mathbf{a} using weighted linear regression:

$$\hat{\mathbf{a}} = (\mathbf{Z}^\dagger \mathbf{N}^{-1} \mathbf{Z})^{-1} \mathbf{Z}^\dagger \mathbf{N}^{-1} \mathbf{v}. \quad (18)$$

Here \mathbf{v} is a vector containing the visibilities for a selection of baselines,

$$Z_{is} = e^{j2\pi\nu\mathbf{b}_i \cdot \hat{\mathbf{n}}(\theta_s, \phi - \phi_s)/c} \quad (19)$$

is the geometric phase factor for baseline i and source s , and \mathbf{N} is the noise covariance. As before, we assume that the noise covariance is diagonal and equal to the propagated fast-cadence estimate of the variance (see Section 3.1.4).

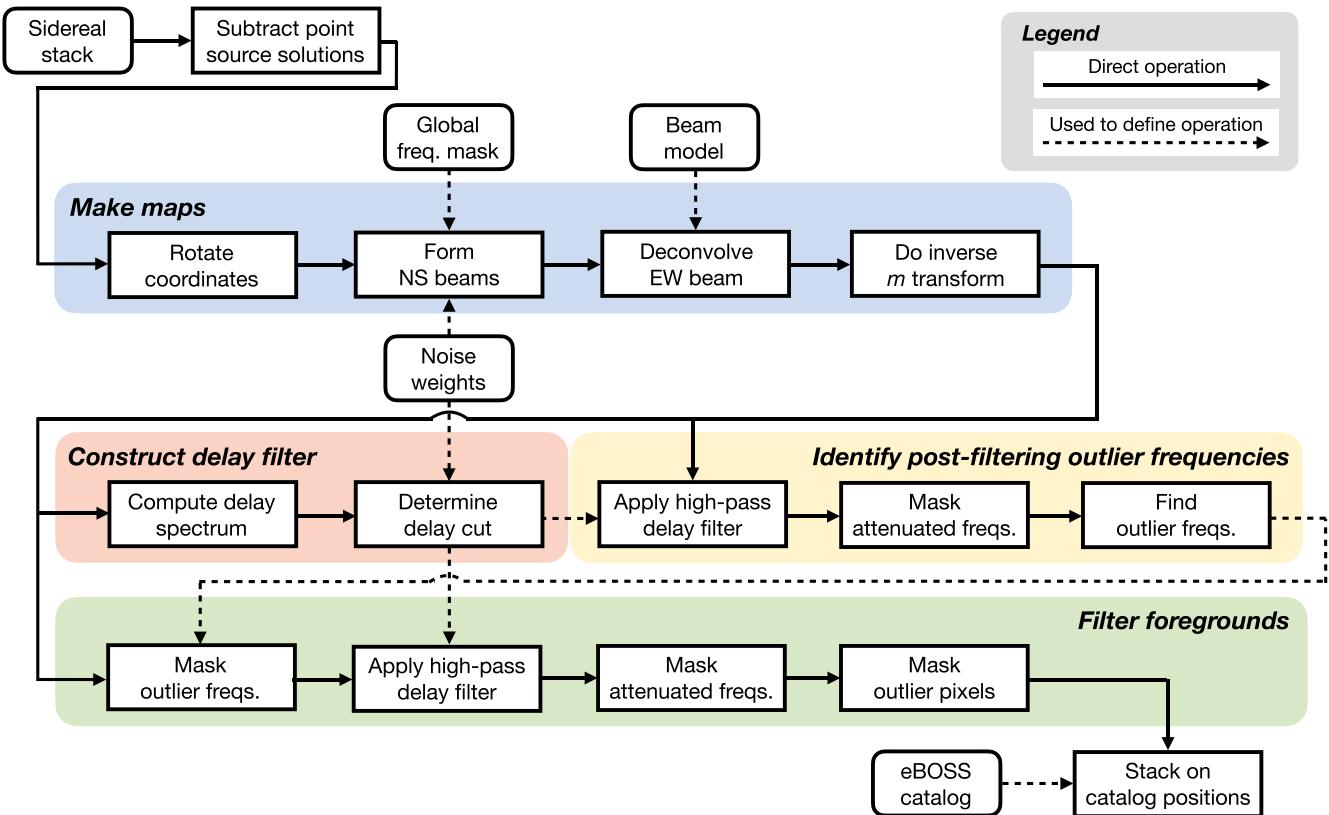


Figure 6. A schematic representation of the stacking pipeline, proceeding from the sidereal stack described in Section 3.3. After subtraction of the four brightest point sources (Section 4.1), sky maps are formed (Section 4.3), accounting for a global frequency mask (Section 4.2) and a model for the primary beam pattern (Section 4.4). Our foreground filtering scheme is designed to reject components of the data with variance far in excess of the expected thermal noise and includes a high-pass delay filter (Section 4.5) and several additional masking operations (Section 4.6), including masking of frequencies that are attenuated by the delay filter. Finally, the filtered maps are stacked at the positions of objects in each eBOSS catalog (Section 4.7).

The amplitude a_s is equal to the spectral flux density of source s modulated by the power beam pattern of the instrument at the source's coordinates and is expected to vary slowly as a function of frequency and hour angle. To improve the signal-to-noise ratio, the best-fit amplitude for each source is smoothed in (ν, ϕ) by iteratively applying a 2D moving average window with size (1.2 MHz, $0^{\circ}44$) and number of iterations (12, 8). The model for the four brightest sources is then computed using Equation (17) and subtracted from the data. Note that only visibilities measured by baselines consisting of feeds on different cylinders are used to solve for the source amplitudes—because contamination from diffuse Galactic emission and noise crosstalk is significantly reduced for these *intercylinder baselines*—but the resulting model is subtracted from all baselines.

4.2. Frequency Mask

The inverse variance weights are multiplied by a global frequency mask that completely excludes certain frequency channels from the stacking analysis. The list below gives the conditions under which a frequency channel is masked and the fraction of the 587.5–800 MHz band that meets each condition.

1. Mask any frequency channel that coincides with a known, persistent source of RFI. There were two sources of persistent RFI in the 587.5–800 MHz band: the mobile LTE bands, and the LO used by the Synthesis Telescope

at DRAO (Landecker et al. 2000). These are shown in Figure 2 (14.2% masked).

2. Mask any frequency channel where the sidereal stack is missing a subset (or all) of the full sidereal day, which prevents a straightforward application of the m -mode transform required for mapmaking. This could be due to a GPU node that was not operational for a significant portion of 2019, as one example (12.5% masked).
3. Mask any frequency channel where the total integration time over the range of R.A. coinciding with the NGC field is less than 75% of the maximum over frequencies. Again, most often this is due to a temporarily nonoperational GPU node (5.9% masked).
4. Mask any frequency channel where manual inspection of the foreground-filtered map in an initial iteration of the analysis revealed residuals that are large relative to the expected radiometric noise and corrupt a significant fraction of the NGC field. This procedure is described in greater detail in Section 4.6 (14.7% masked).

In total, these four conditions mask 47.2% of the 587.5–800 MHz band.

4.3. Mapmaking

The next step in the data processing is to construct a map from the sidereal visibilities. We use a mapmaking technique that is tailored to CHIME, or effectively any transit radio interferometer consisting of cylindrical telescopes oriented in the NS direction with a close-packed array of antennas along

the axis of each cylinder. The technique draws on the work presented in Shaw et al. (2014) and Masui et al. (2017) but is distinct and has not been described elsewhere, so we go into considerable detail in this section. Note that these types of maps are referred to as *deconvolved ringmaps* in CHIME Collaboration et al. (2022b).

4.3.1. Baseline Configuration

To good approximation, the CHIME baselines \mathbf{b}_c are located on a 2D grid that lies in the plane tangent to Earth's surface at (latitude, longitude) $\equiv (\Lambda, \Phi) = (49^\circ 320709, -119^\circ 623677)$. The 2D grid is given by

$$\mathbf{b}_c = x d_x \hat{\mathbf{x}}_c + y d_y \hat{\mathbf{y}}_c, \quad (20)$$

where $\hat{\mathbf{x}}_c$ is the unit vector that is orthogonal to the cylinder, $\hat{\mathbf{y}}_c$ is the unit vector parallel to the cylinder, $d_x = 22.0$ m is the (center-to-center) cylinder spacing, $d_y = 0.3048$ m is the spacing of the feeds along the focal line, and the grid indices are denoted by $x \in [-3, 3]$ and $y \in [-255, 255]$.

The sidereal visibilities are arranged onto this 2D grid. Let $V_{xy}^{pq}(\nu, \phi)$ denote the visibility measured at frequency ν and local Earth rotation angle ϕ by the baseline at the (x, y) grid position. The variables $p, q \in \{X, Y\}$ refer to the polarizations of the two antennas that form the baseline, with the dipole of the X and Y polarizations oriented in the $\hat{\mathbf{x}}_c$ - and $\hat{\mathbf{y}}_c$ -directions, respectively. The analysis presented in this work will only use the co-polar baselines, XX and YY , so that $p = q$, and we drop the redundant index in the notation going forward. Note that it is assumed that the visibilities have conjugate symmetry about the origin, specifically

$$V_{-x, -y}^p \equiv (V_{xy}^p)^*. \quad (21)$$

The CHIME cylinders were aligned with the NS direction by design. However, we have empirically determined that the cylinders are rotated by $\psi = -0^\circ 071$ with respect to true astronomical north using observations of a large number of bright point sources (CHIME Collaboration et al. 2022b). Let $\mathbf{b} = \mathbf{R}(\psi) \mathbf{b}_c$ denote the baselines in a coordinate system where $\hat{\mathbf{x}}$ is aligned with the east–west (EW) direction, $\hat{\mathbf{y}}$ is aligned with the NS direction, and

$$\mathbf{R}(\psi) = \begin{bmatrix} \cos \psi & -\sin \psi \\ \sin \psi & \cos \psi \end{bmatrix} \quad (22)$$

is the rotation matrix that transforms between the cylinder-based coordinate system and the NS based coordinate system.

The measured visibility is the true sky visibility \mathcal{V} corrupted by noise n ,

$$V_{xy}^p(\nu, \phi) = \mathcal{V}_{xy}^p(\nu, \phi) + n_{xy}^p(\nu, \phi). \quad (23)$$

The sky visibility \mathcal{V} is the integral of the spectral flux density, S , of the sky multiplied by the primary beam pattern, A , of the two feeds and a geometric phase factor set by the baseline between the feeds:

$$\begin{aligned} \mathcal{V}_{xy}^p(\nu, \phi) &= \int |A^p(\nu, \theta', \phi - \phi')|^2 e^{j2\pi\nu\mathbf{b} \cdot \hat{\mathbf{n}}(\theta', \phi - \phi')/c} \\ &\quad \times S(\nu, \theta', \phi') \cos \theta' d\theta' d\phi'. \end{aligned} \quad (24)$$

Here c is the speed of light and $\hat{\mathbf{n}}(\theta', \phi - \phi')$ is the unit vector pointing toward decl. θ' and hour angle HA $\equiv \phi - \phi'$ and is

given by

$$\begin{aligned} \hat{\mathbf{n}}(\theta', \phi - \phi') &= -\cos \theta' \sin(\phi - \phi') \hat{\mathbf{x}} \\ &\quad + (\cos \Lambda \sin \theta' - \sin \Lambda \cos \theta' \cos(\phi - \phi')) \hat{\mathbf{y}} \\ &\quad + (\sin \Lambda \sin \theta' + \cos \Lambda \cos \theta' \cos(\phi - \phi')) \hat{\mathbf{z}}, \end{aligned} \quad (25)$$

with Λ denoting the latitude of the telescope. Note that Equation (24) assumes that the primary beam pattern is the same for all feeds of a given polarization. It also assumes that there are no residual complex gain variations.

4.3.2. North–South Beamforming

The CHIME power beam, $|A|^2$, is reasonably compact in the hour-angle direction. The FWHM of the main lobe is $\lesssim 2^\circ 1$ ($2^\circ 5$) for the Y (X) polarization in the 587.5–800 MHz band, and the sidelobes are $\lesssim 1\%$ (CHIME Collaboration et al. 2022b). If we restrict the integral in Equation (24) to the range of hour angles covering the main lobe of the primary beam, then we can expand the geometric phase to first order in the small angles HA and ψ to obtain

$$\begin{aligned} \mathbf{b} \cdot \hat{\mathbf{n}}(\theta', \phi - \phi') &= x d_x [\sin(\theta' - \Lambda) \psi - \cos \theta' (\phi - \phi')] \\ &\quad + y d_y \sin(\theta' - \Lambda). \end{aligned} \quad (26)$$

The geometric phase due to the y component of the baseline is given by the second term in Equation (26). Since this term depends only on decl. and does not depend on hour angle, we can form a linear combination of all visibilities with the same x so that only signal from a specific decl., θ , adds coherently:

$$V_x^p(\nu, \phi, \theta) = \sum_y W_{xy}^p(\nu, \phi) V_{xy}^p(\nu, \phi) e^{-j2\pi\nu y d_y \sin(\theta - \Lambda)/c}, \quad (27)$$

where

$$W_{xy}^p(\nu, \phi) = w_{xy}^p(\nu, \phi) / \sum_{y'} w_{xy'}^p(\nu, \phi) \quad (28)$$

denotes the relative weights, which are normalized to preserve point-source flux. This beamforming operation is repeated for a grid of pointings that span from horizon to horizon and are equally spaced in $\sin(\theta - \Lambda)$. This operation can be done efficiently with a fast Fourier transform (FFT), but, in practice, we evaluate the expression directly to ensure that the grid of pointings is the same for all frequencies.

We will refer to $V_x^p(\nu, \phi, \theta)$ as the hybrid beamformed visibility, since the NS component of the baseline has been beamformed to a specific decl., but there is still fringing associated with the EW component of the baseline. Combining Equations (24)–(27), we obtain the following theoretical expression for the hybrid beamformed visibilities:

$$\begin{aligned} \mathcal{V}_x^p(\nu, \phi, \theta) &= \int b_{\text{synth}}^{\theta, p}(\nu, \theta, \theta', \phi) B_x^p(\nu, \theta', \phi - \phi') \\ &\quad \times S(\nu, \theta', \phi') \cos \theta' d\theta' d\phi', \end{aligned} \quad (29)$$

where

$$\begin{aligned} B_x^p(\nu, \theta', \phi - \phi') &= |A^p(\nu, \theta', \phi - \phi')|^2 e^{j2\pi\nu x d_x [\sin(\theta' - \Lambda) \psi - \cos \theta' (\phi - \phi')]/c} \\ &= |A^p(\nu, \theta', \phi - \phi')|^2 e^{j2\pi\nu x d_x [\sin(\theta' - \Lambda) \psi - \cos \theta' (\phi - \phi')]/c} \end{aligned} \quad (30)$$

will be referred to as the beam transfer function and

$$b_{\text{synth},x}^{\hat{\theta},p}(\nu, \theta, \theta', \phi) = \sum_y W_{xy}^p(\nu, \phi) e^{j2\pi\nu y d_y [\sin(\theta' - \Lambda) - \sin(\theta - \Lambda)]/c} \quad (31)$$

is the synthesized beam in the $\hat{\theta}$ -direction. The top panel of Figure 7 shows an example of $b_{\text{synth}}^{\hat{\theta}}$ for the weighting scheme used in this analysis.

The absolute weights in Equation (28) are set to the inverse variance of the corresponding visibility, i.e.,

$$w_{xy}^p(\nu, \phi) = [\text{Var}(V_{xy}^p(\nu, \phi))]^{-1}, \quad (32)$$

which will maximize the point-source sensitivity since the amplitude of a true point source is the same for all baselines. We describe how the variance of the visibilities is estimated in Section 3.1.4. The inverse variance weights scale approximately as the number of redundant baselines that are averaged together by the real-time pipeline to produce $V_{xy}^p(\nu, \phi)$, which scales with the NS baseline distance as $(256 - |y|)$. As a result, the inverse variance weights produce an approximately triangular window function in y . This yields a synthesized beam $b_{\text{synth}}^{\hat{\theta}}$ that has an FWHM ranging from $0^{\circ}.35$ at 585 MHz to $0^{\circ}.25$ at 800 MHz and sidelobes that range from 0.05 to 10^{-4} of the peak. Note that, instead of inverse variance weights, we could set the weights to any window function that further suppresses the sidelobes at the expense of point-source sensitivity.

In principle, the synthesized beam $b_{\text{synth}}^{\hat{\theta}}$ depends on both the EW baseline distance x and the local Earth rotation angle ϕ , because the inverse variance weights change with these parameters. However, the weights that are used in this analysis yield a synthesized beam that is quite stable with ϕ and similar across x . Indeed, the standard deviation of the synthesized beam over ϕ is at most 0.1% (relative to the peak) over all polarizations, frequencies, and decl., and the standard deviation over x is at most 2%. In order to simplify the derivation that follows, we will drop the dependence of the synthesized beam on both ϕ and x . This assumption can be enforced directly—while maintaining roughly the same sensitivity—by explicitly using the triangular window function, or, in other words, by setting $w_{xy}(\nu, \phi) = 256 - |y|$. Doing so, we find no appreciable change in either the signal or noise in the stacks on the eBOSS catalogs.

The regularly gridded baselines do not Nyquist sample the visibility of the sky in the \hat{y} -direction for frequencies $\nu \geq \frac{c}{2d_y} \approx 492$ MHz, which includes all frequencies considered in this analysis. As a result, the hybrid beamformed visibilities will suffer from aliasing. In this derivation, the effects of aliasing are encoded in the synthesized beam $b_{\text{synth}}^{\hat{\theta}}$. Let $\beta(\nu) \equiv \frac{c}{\nu d_y} - 1$. If $\sin(\theta - \Lambda) < -\beta(\nu)$ or $\sin(\theta - \Lambda) > \beta(\nu)$, then the synthesized beam will have two main lobes, one centered on the desired decl. $\theta' = \theta$ and a duplicate centered on the frequency-dependent aliased decl. $\theta'_{\text{alias}}(\nu)$, given by

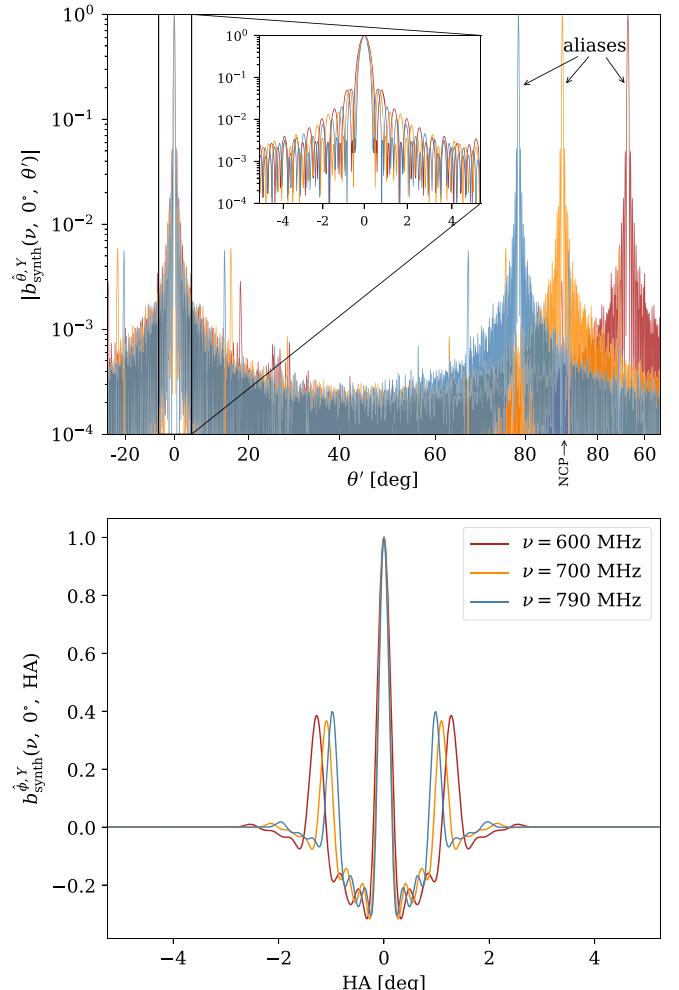


Figure 7. The synthesized beam (i.e., the point-spread function of the map) for the Y polarization array at decl. $\theta = 0^{\circ}$. Each color corresponds to a different frequency as described in the legend in the bottom panel. The top panel shows the synthesized beam in the $\hat{\theta}$ -direction (see Equation (31)). The x -axis is uniformly spaced in the sine of the zenith angle and spans from horizon to horizon, with the region to the right of the NCP annotation corresponding to the antipodal transit at hour angle = 180° . The synthesized beam has sensitivity to both the beamformed decl. at $\theta' = 0^{\circ}$ and—due to aliasing—a second frequency-dependent decl. at $\theta'_{\text{alias}}(\nu)$ (see Equation (33)). The inset panel zooms in on $\pm 5^{\circ}$ from the beamformed decl. The bottom panel shows the synthesized beam in the $\hat{\phi}$ -direction (see Equation (47)). The exclusion of intracylinder baselines in the mapmaking procedure results in negative shoulders on either side of the main lobe. Aliasing in the $\hat{\phi}$ -direction results in two *grating lobes* with amplitudes that are 40% of the amplitude of the main lobe. As the regularization parameter $\eta \rightarrow 0$ (see Equation (45)), the primary beam is perfectly deconvolved (assuming an accurate primary beam model) and the grating lobes disappear. For this analysis we have chosen a relatively large value of η , which results in better point-source sensitivity but larger grating lobes.

the equation

$$\begin{aligned} & \sin(\theta'_{\text{alias}}(\nu) - \Lambda) \\ &= \begin{cases} \sin(\theta - \Lambda) + \frac{c}{\nu d_y} & \text{if } -1 \leq \sin(\theta - \Lambda) \leq -\beta(\nu) \\ \sin(\theta - \Lambda) - \frac{c}{\nu d_y} & \text{if } \beta(\nu) \leq \sin(\theta - \Lambda) \leq 1. \end{cases} \end{aligned} \quad (33)$$

Hence, the hybrid beamformed visibility $V_x^p(\nu, \phi, \theta)$ will contain equal contributions from the sky (modulated by the

beam transfer function) at θ' and $\theta'_{\text{alias}}(\nu)$. This is illustrated in the top panel of Figure 7. At the upper edge of the band, β (800 MHz) = 0.23, which implies that there is a stripe of the sky centered on zenith (specifically $\theta \in [36.0^\circ, 62.6^\circ]$) that is free from aliases at all CHIME frequencies. Outside of this stripe, the aliased sky is heavily attenuated in the intercylinder baselines by utilizing the fact that it will fringe at a different rate than the true sky. This is discussed further below.

The first sidelobe of the synthesized beam in the $\hat{\theta}$ -direction has an amplitude that is 5% of the amplitude of the main lobe, the next sidelobe is 1%–2%, and beyond roughly 2° separation all sidelobes are below 0.5%. We therefore assume that a hybrid beamformed visibility is dominated by the sky at a narrow range of decl. centered on θ . This assumption will start to break down at right ascensions that coincide with bright foregrounds. This problem is mitigated to a certain extent by subtracting the four brightest point sources directly from the sidereal visibilities, as explained in Section 4.1, and using only intercylinder baselines that resolve out the bright, diffuse Galactic emission, which will be explained below.

4.3.3. Primary Beam Deconvolution

Since the beam transfer function does not change appreciably on scales less than the FWHM of the synthesized beam, it can be brought outside of the integral over θ' in Equation (29), resulting in the following equation:

$$\begin{aligned} \mathcal{V}_x^p(\nu, \phi, \theta) &= \int B_x^p(\nu, \theta, \phi - \phi') \cos \theta \, d\phi' \\ &\times b_{\text{synth}}^{\hat{\theta}, p}(\nu, \theta, \theta') S(\nu, \theta', \phi') \, d\theta'. \end{aligned} \quad (34)$$

Given a model for the primary beam A , the beam transfer function is computed using Equation (30) and then deconvolved from the hybrid beamformed visibilities at each decl. to recover the flux density of the sky S convolved with $b_{\text{synth}}^{\hat{\theta}}$. The construction of the primary beam model will be described in Section 4.4. The FFT of the hybrid beamformed visibility is taken along the ϕ -axis,

$$\tilde{V}_{xm}^p(\nu, \theta) = \frac{1}{N_\phi} \sum_n V_x^p(\nu, \phi_n, \theta) e^{-jm\phi_n}. \quad (35)$$

Here the sum runs over the $N_\phi = 4096$ samples on the grid in local Earth rotation angle and the m -modes range over $[-2047, 2048]$. We will refer to this operation as the m -mode transform going forward. The same operation is performed on the beam transfer function. Figure 8 provides an example of the m -mode transform of both the hybrid beamformed visibility and the beam transfer function.

The beam transfer function is then deconvolved from the data in m -space using a Tikhonov regularization scheme

$$\tilde{M}_m^p(\nu, \theta) = \frac{\sum_x W_x^p(\nu) \tilde{B}_{xm}^p(\nu, \theta) \tilde{V}_{xm}^p(\nu, \theta)}{\eta + \sum_x W_x^p(\nu) |\tilde{B}_{xm}^p(\nu, \theta)|^2}, \quad (36)$$

where

$$W_x^p(\nu) = w_x^p(\nu) / \sum_x w_x^p(\nu) \quad (37)$$

denotes the relative weight given to each EW baseline and η is a regularization parameter. The different EW baselines measure

a largely disjoint set of m -modes, with each baseline primarily sensitive to the range of m values centered on

$$m_{\text{center},x}(\nu, \theta) = -2\pi\nu \cos(\theta) x d_x / c, \quad (38)$$

with width

$$m_{\text{width}}(\nu, \theta) = 2\pi\nu \cos(\theta) w / c, \quad (39)$$

where $w = 20$ m is the width of the cylinder. However, there is some mild overlap that is dependent on the aperture illumination and accounted for by the m -mode transform of the primary beam pattern. Equation (36) first performs a weighted average of the measurements made by the different EW baselines and then deconvolves the primary beam by effectively dividing by the corresponding weighted average of the m -mode transform of the beam transfer function. The regularization parameter η is the assumed inverse signal-to-noise ratio. It defines which m -modes are signal dominated, and hence should be divided by the beam, and conversely which m -modes are noise dominated, and hence should not be amplified further by dividing by the beam.

We set

$$w_x^p(\nu) = \begin{cases} 0 & x = 0 \\ [\sigma_x^p(\nu)]^{-2} & |x| > 0, \end{cases} \quad (40)$$

where

$$[\sigma_x^p(\nu)]^2 = \frac{1}{N_\phi^2} \sum_n [\sigma_x^p(\nu, \phi_n)]^2 \quad (41)$$

is the variance of the noise in the m -mode transform of the hybrid beamformed visibility and

$$[\sigma_x^p(\nu, \phi)]^2 = \left(\sum_y w_{xy}^p(\nu, \phi) \right)^{-1} \quad (42)$$

is the variance of the noise in the hybrid beamformed visibility. This weighting scheme masks all intracylinder baselines and propagates the inverse variance weights through the beamforming and m -mode transform for intercylinder baselines. The redundancy of the array results in $W_x^p(\nu) \approx [0.0, 0.5, 0.333, 0.166]$ for $|x| = [0, 1, 2, 3]$, corresponding to the intracylinder autocorrelation that is removed, the three-fold redundancy in the one-cylinder separation, the twofold redundancy in the two-cylinder separation, and the single appearance of the three-cylinder separation.

The intracylinder baselines are masked for this analysis because they contain two sources of contamination that are significantly reduced in the intercylinder baselines: (1) large-scale diffuse Galactic emission, and (2) noise crosstalk (see Section 3.3). Since the noise crosstalk changes slowly with time, it contaminates only low m , which is where the signal from the sky resides for intracylinder baselines. Note that the signal from the sky at decl. near the north celestial pole (NCP) will also appear at low m , even for intercylinder baselines. However, the maximum decl. of sources in the eBOSS NGC field is 60° , which is far enough from the NCP that the crosstalk contamination in the intercylinder measurements is negligible.

The beam transfer function of the intercylinder baselines is largely insensitive to the range of m -modes occupied by the aliased sky for the decl. considered in this analysis. This can be

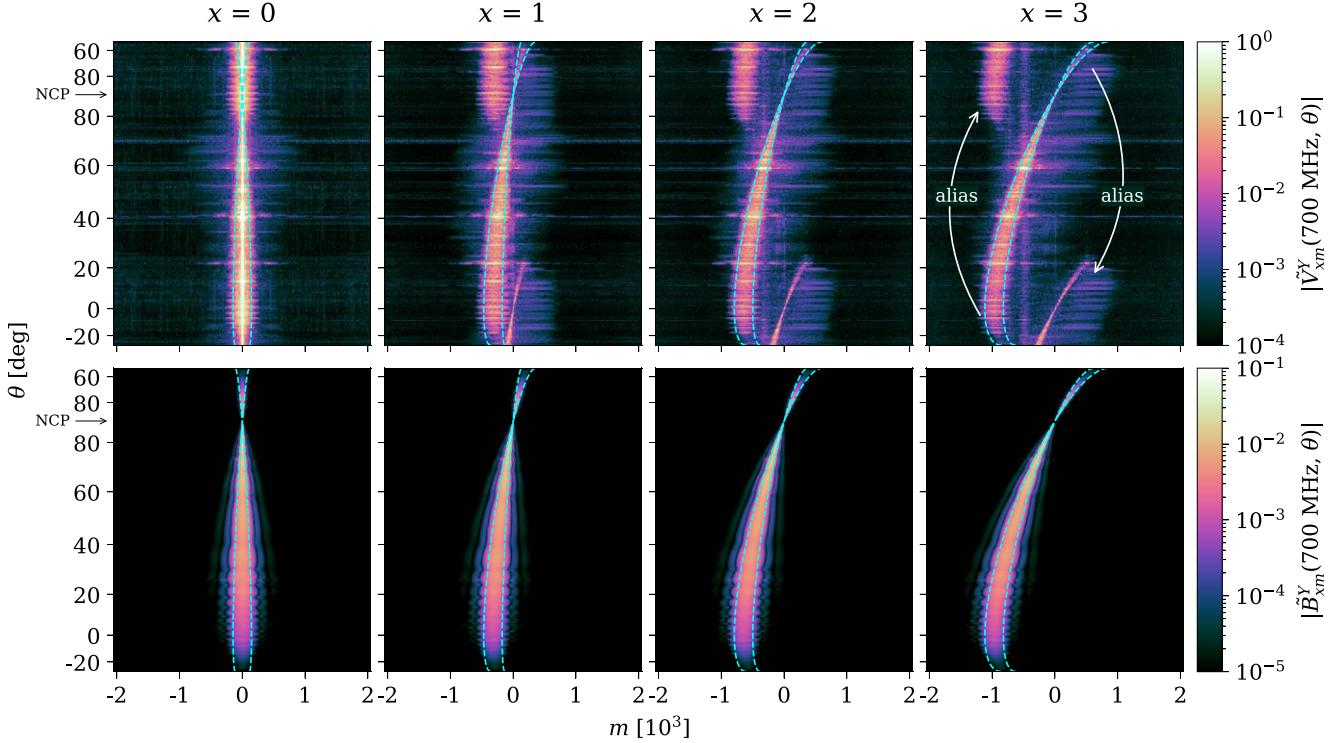


Figure 8. The m -mode transform of the hybrid beamformed visibilities (top row) and beam transfer function (bottom row) for the Y polarization array at frequency $\nu = 700.78125$ MHz. Each column shows a different EW baseline separation. The dashed cyan lines are given by $m_{\text{center},x}(\theta, \nu) \pm \frac{1}{2}m_{\text{width}}(\theta, \nu)$ (see Equations (38) and (39)) and enclose the range of m where that baseline has the bulk of its sensitivity to the sky. By convention, the positive intercylinder baselines ($x > 0$) measure a negative fringe rate (negative m) for the sky south of the NCP and a positive fringe rate (positive m) for the sky north of the NCP. In the top row, the aliased sky is annotated in the $x = 3$ column and is also clearly visible in the $x = 2$ and $x = 1$ column, but it overlaps entirely with true sky for the intracylinder baselines at $x = 0$. In general, the aliased sky and true sky are well separated in m -space for the intercylinder baselines. The bright features at approximately 20° , 40° , and 60° decl. correspond to residual signal from Taurus A, Cygnus A, and Cassiopeia A, respectively. The leakage from Cassiopeia A to other decl. is clearly visible in the $x = 3$ panel. In the bottom row, there is ringing outside of the dashed cyan lines because our model for the primary beam pattern has been truncated so that it only includes the main lobe (see Appendix B).

shown in a rough way using Equations (33), (38), and (39). The NGC field spans decl. from 13° to 60° , and over this range there is zero overlap between $m_{\text{center},x}(\nu, \theta) \pm \frac{1}{2}m_{\text{width}}(\nu, \theta)$ and $m_{\text{center},x}(\nu, \theta_{\text{alias}}(\nu)) \pm \frac{1}{2}m_{\text{width}}(\nu, \theta_{\text{alias}}(\nu))$ for all ν and for all $x > 0$. However, examining Figure 8, it is clear that for $\theta \lesssim 20^\circ$ our actual beam model does have some sensitivity to the aliased sky for the $x = 1$ baseline. Therefore, although the deconvolution procedure will heavily attenuate the aliased sky, it is still expected to introduce some nonnegligible contamination.

The regularization parameter is set to $\eta = 10^{-4}$. This value was chosen by constructing a map for several different values of η between 10^{-6} and 10^{-3} and choosing the value that maximizes the point-source sensitivity. Note that smaller values of the regularization parameter result in better deconvolution of the primary beam in the $\hat{\phi}$ -direction, but also higher noise, and were thus disfavored for the analysis presented in this work.

Finally, the deconvolved map is obtained by taking the inverse m -mode transform

$$M^p(\nu, \theta, \phi) = \sum_m \tilde{M}_m^p(\nu, \theta) e^{jm\phi}. \quad (43)$$

4.3.4. Map Normalization

In order to determine the correct normalization for the map, we consider a radio sky that contains a single point source with

unit flux density at decl. θ and local Earth rotation angle ϕ' . The m -mode transform of the hybrid beamformed visibilities at that decl. is given by

$$\tilde{V}_{xm}^p(\nu, \theta) = \tilde{B}_{xm}^p(\nu, \theta) e^{-jm\phi'}. \quad (44)$$

The source profile along the ϕ -axis of the resulting map is therefore

$$\begin{aligned} a^p(\nu, \theta, \phi - \phi') \\ = \sum_m \frac{\sum_x W_x^p(\nu) |\tilde{B}_{xm}^p(\nu, \theta)|^2}{\eta + \sum_x W_x^p(\nu) |\tilde{B}_{xm}^p(\nu, \theta)|^2} e^{jm(\phi - \phi')}, \end{aligned} \quad (45)$$

and the peak flux density of the source is $a^p(\nu, \theta, 0)$, which in general is not equal to unity. Therefore, in order to preserve the point-source flux through the mapmaking process, the map is normalized as

$$M^p(\nu, \theta, \phi) \rightarrow \frac{M^p(\nu, \theta, \phi)}{a^p(\nu, \theta, 0)} \quad (46)$$

and the synthesized beam in the $\hat{\phi}$ -direction is given by

$$b_{\text{synth}}^{\hat{\phi}, p}(\nu, \theta, \phi - \phi') = \frac{a^p(\nu, \theta, \phi - \phi')}{a^p(\nu, \theta, 0)}. \quad (47)$$

The bottom panel of Figure 7 shows an example of $b_{\text{synth}}^{\hat{\phi}}$ for the weighting scheme, regularization parameter, and beam model employed in this analysis.

The resulting map is modeled as

$$M^p(\nu, \theta, \phi) = \mathcal{M}^p(\nu, \theta, \phi) + n^p(\nu, \theta, \phi). \quad (48)$$

Here \mathcal{M} is related to the flux density of the sky through the relation

$$\mathcal{M}^p(\nu, \theta, \phi) = \int b_{\text{synth}}^{\hat{\theta}, p}(\nu, \theta, \theta') b_{\text{synth}}^{\hat{\phi}, p}(\nu, \theta, \phi - \phi') \times S(\nu, \theta', \phi') \cos \theta d\theta' d\phi', \quad (49)$$

where the synthesized beams in the $\hat{\theta}$ - and $\hat{\phi}$ -directions can be calculated directly from Equations (31) and (47), respectively. The quantity $n^p(\nu, \theta, \phi)$ represents the noise in the map.

4.3.5. Variance Estimation

The variance of the noise in the map is estimated as

$$[\sigma_{\text{map}}^p(\nu, \theta, \phi)]^2 = F^p(\nu, \phi) [\sigma_{\text{map}}^p(\nu, \theta)]^2, \quad (50)$$

where $[\sigma_{\text{map}}^p(\nu, \theta)]^2$ is obtained by propagating the variance given by Equation (41) through the mapmaking (Equation (36)), inverse m -mode transform (Equation (43)), and normalization (Equation (46)) procedure. The integration time in the sidereal stack is nonuniform, primarily due to seasonal changes in the length of the day and the likelihood of rainfall. As a result, the variance of the noise depends on the local Earth rotation angle. This dependence is lost when propagating the variance through the forward and inverse m -mode transform. The factor $F^p(\nu, \phi)$ approximately recovers this dependence and is given by the weighted average over EW baselines of the fractional change in the variance of the noise in the hybrid beamformed visibilities, i.e.,

$$F^p(\nu, \phi) = \sum_x W_x^p(\nu) \frac{[\sigma_x^p(\nu, \phi)]^2}{\frac{1}{N_\phi} \sum_n [\sigma_x^p(\nu, \phi_n)]^2}, \quad (51)$$

where $\sigma_x^p(\nu, \phi)$ is given by Equation (42).

This procedure for propagating the variance through the mapmaking has been validated as follows. We generate visibilities that have been randomly drawn from a circularly symmetric, complex Gaussian distribution with mean 0 and variance equal to the expected variance of the radiometric noise,

$$V_{xy}^p(\nu, \phi) \sim \mathcal{N}\left(0, \frac{1}{2} \sigma_{xy}^p(\nu, \phi)^2\right) + j \mathcal{N}\left(0, \frac{1}{2} \sigma_{xy}^p(\nu, \phi)^2\right), \quad (52)$$

where $\mathcal{N}(\mu, \sigma^2)$ denotes a Gaussian distribution with mean μ and variance σ^2 . Estimation of the expected radiometric variance $\sigma_{xy}^p(\nu, \phi)^2 = w_{xy}^p(\nu, \phi)^{-1}$ is described in Section 3.1.4. The mapmaking procedure is then applied to this *Gaussian noise realization* in an identical manner as to the data. The sample variance of the map pixels is calculated in a 2D rolling window in (θ, ϕ) and compared to the estimate given by Equation (50). In general, we find good agreement ($\lesssim 5\%$) between the two. This technique of processing a Gaussian noise realization using the same pipeline that is applied to the data will be used in other comparisons below.

4.4. Beam Calibration

Our primary beam model is obtained by deconvolving a model for the radio sky that consists only of extragalactic point sources from the visibilities measured with baselines that have a large EW component. The long baselines resolve out the diffuse Galactic emission, making a point-source-only sky model a reasonable description of the data. There are several high-resolution, large-area sky surveys that can be interpolated to the 400–800 MHz CHIME band to construct this sky model. At lower frequencies we rely on the VLA Low-frequency Sky Survey (Cohen et al. 2007) at 74 MHz and the Westerbork Northern Sky Survey (Rengelink et al. 1997) at 326 MHz. At higher frequencies we rely on the NRAO VLA Sky Survey (Condon et al. 1998) at 1400 MHz and the Green Bank survey (Gregory et al. 1996) at 4850 MHz. The method used to deconvolve the sky model from the data is similar to the method used to construct a map, which was described in the previous section. Whereas the mapmaker deconvolves a model for the primary beam from the hybrid beamformed visibilities to estimate the intensity of the sky, the beam calibration deconvolves a model for the sky intensity from the same hybrid beamformed visibilities in order to estimate the primary beam. Appendix B describes this method in detail.

The resulting power beam, $|A^Y(\nu, \theta, \phi)|^2$, for the Y polarization array is shown in the top row of Figure 9. We briefly describe the main features of the CHIME primary beam pattern, referring the reader to CHIME Collaboration et al. (2022b) for a more in-depth discussion. The large (50%) ripples that are evident in the frequency and decl. direction are the result of multipath interference. Radiation from the sky can be absorbed and then reradiated by feeds or reflected off the ground plane. It then reflects off the cylinder and interferes with the primary path from the sky. The period of the ripple is ~ 30 MHz and is set by the 5 m focal length of the CHIME cylinders. Harmonics at 60 and 90 MHz, which arise from multiple reflections off the focal line and cylinder, are also significant, although they may not be distinguishable by eye in Figure 9. The narrowing of the beam in the hour-angle direction as one moves toward higher frequencies is due to diffraction through the 20 m aperture. The apparent widening of the beam in the hour-angle direction as one approaches the NCP is simply due to the fact that a point at decl. θ travels $\cos \theta$ degrees on the sky for every degree in hour angle that elapses. The beam is normalized to 1.0 on meridian at the decl. of Cygnus A ($40^\circ 73392$) at each frequency in order to match how the data are normalized during complex gain calibration. This imprints the interference pattern at the decl. of Cygnus A onto all other decl. The power beam for the X polarization array exhibits the same general features but is slightly wider in both the hour-angle and decl. direction and also has a lower response at zenith because the dipole illuminates the cylinder less efficiently.

In order to characterize the effect that the ripples in the beam have on our final stacking result, we repeat our analysis with a “control” beam that has the same large-scale properties as the default beam model, but without the small-scale structure in the frequency and decl. direction. The control beam is a modified version of the analytical beam model proposed in Shaw et al. (2015, hereafter S15) for cylindrical telescopes. To briefly summarize the S15 model, the beam pattern of the antenna (henceforth “base” beam) is assumed to be that of a horizontal dipole mounted a distance $\lambda/4$ over a conducting ground plane.

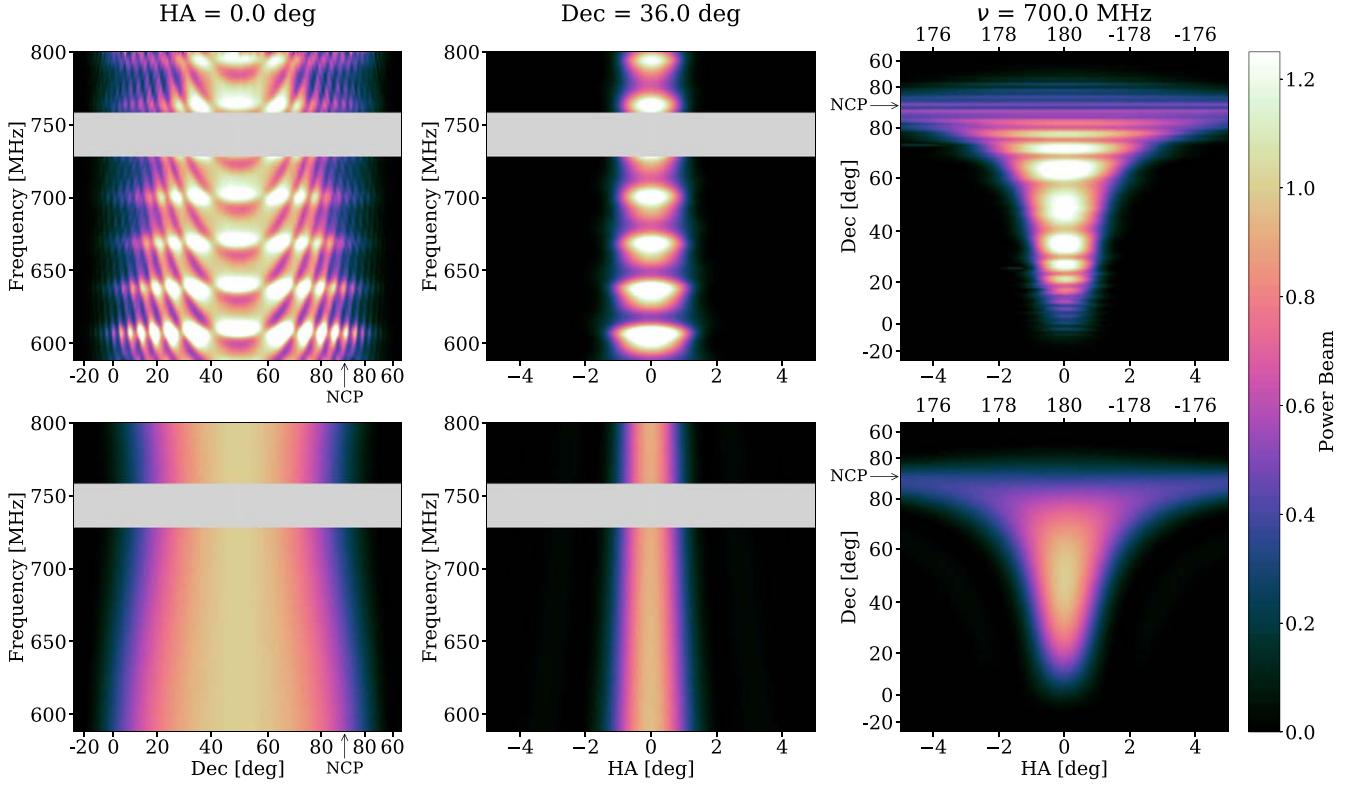


Figure 9. 2D slices through the 3D primary beam models. We show the power beam for the Y polarization array. The top row is the default beam model, obtained by deconvolving a model for the radio emission from extragalactic point sources from the visibilities measured with long EW baselines. The bottom row is the control beam model, which has similar global properties to the default, but without the small-scale spectral structure. Left: beam model as a function of decl. and frequency on the meridian (hour angle = 0°). The decl. axis is uniformly spaced in the sine of the zenith angle. The region to the right of the NCP annotation corresponds to the antipodal transit at hour angle = 180° . Middle: beam model as a function of hour angle and frequency at a decl. of 36.0° . Right: beam model as a function of hour angle and decl. at 700 MHz. The region above the NCP annotation corresponds to the antipodal transit at hour angles given by the upper x -axis. The beam has been normalized to 1.0 on meridian at the decl. of Cygnus A ($40^\circ 73392$) at each frequency in order to match how the data are normalized by the calibration procedure. The gray band denotes frequencies where we do not have a valid model for the beam owing to persistent RFI in mobile LTE bands.

The response in the EW direction is the result of solving the Fraunhofer diffraction problem for a dipole illuminating an aperture with width equal to the 20 m cylinder width. The response in the NS direction is simply the reflected amplitude of the base beam. The primary beam of the telescope is then the outer product of these two 1D functions.

In this work, the S15 model for the base beam is modified to more accurately describe existing measurements of the CHIME primary beam. The assumption that the base beam for the X polarization is the base beam for Y polarization rotated by 90° is abandoned. The FWHM of the base beam in the EW direction is assumed to be polarization dependent but frequency independent and is obtained by performing a fit to holographic observations of several bright sources made in conjunction with the John A. Galt 26 m telescope (see CHIME Collaboration et al. 2022b for a description of these measurements). The FWHM of the base beam in the NS direction is assumed to be polarization and frequency dependent and is obtained by fitting a flattened Gaussian to the meridian profile of the default beam at each frequency and then fitting the resulting FWHM as a function of frequency to a third-order polynomial in order to smooth over the small-scale ripples while retaining large-scale variations observed in the width of the meridian beam with frequency. The rest of the procedure is unchanged: the beam model is given by the outer product of an EW response obtained by solving the Fraunhofer diffraction problem and a NS response obtained from the reflected base beam amplitude.

The resulting beam model is shown in the bottom row of Figure 9.

4.5. Foreground Filtering

The deconvolved map described in Section 4.3 is dominated by emission from extragalactic point sources, which is expected to be a factor of $\sim 10^3$ – 10^5 brighter than the 21 cm signal of interest (Santos et al. 2005). This foreground contamination can be separated from the 21 cm signal on the basis of spectral scale; the foregrounds are expected to be spectrally smooth, whereas the 21 cm signal varies rapidly with frequency (Shaver et al. 1999; Oh & Mack 2003; Liu & Tegmark 2011). For each pixel in the map, we apply a high-pass filter along the frequency axis to suppress the foregrounds while retaining some fraction of the 21 cm signal.

Designing an adequate filter is complicated by the fact that—as discussed in Section 4.2—47.2% of the band has been masked in order to remove RFI-like features and other narrowband, instrumental artifacts. The DAYENU technique (Ewall-Wice et al. 2021) is used to construct a linear filter for the irregularly sampled map spectra that achieves the required suppression at large spectral scales. In what follows, we briefly summarize this technique.

Let τ denote the delay, which is the Fourier transform dual to frequency ν . The following simple model is assumed for the

covariance of the map as a function of τ :

$$\tilde{C}(\tau, \tau') = \begin{cases} \frac{\epsilon^{-1}}{2\tau_{\text{cut}}} \delta^D(\tau, \tau') & |\tau| \leq \tau_{\text{cut}} \\ \Delta\nu \delta^D(\tau, \tau') & |\tau| > \tau_{\text{cut}}, \end{cases} \quad (53)$$

where the region below τ_{cut} is the region of delay space contaminated by bright foregrounds, ϵ is a small number that corresponds to the ratio of the radiometric noise to foreground variance, $\Delta\nu = 0.390625$ MHz is the width of the frequency channel, and $\delta^D(\tau, \tau')$ is the Dirac delta function. This model results in the following analytical formula for the covariance between frequency channel ν_m and ν_n :

$$C_{mn} = \epsilon^{-1} \text{sinc}[2\pi\tau_{\text{cut}}(\nu_m - \nu_n)] + \delta_{mn}. \quad (54)$$

where $\text{sinc}(x) \equiv \sin(x)/x$ and δ_{mn} is the Kronecker delta.

To construct the filter, the delay cut $\tau_{\text{cut}}^p(\theta)$ and stop-band rejection ϵ are specified. Note that we allow the delay cut to vary as a function of polarization and decl. Equation (54) is then evaluated for each pair of frequency channels in the 587.5–800 MHz band. Rows and columns of the covariance matrix that correspond to masked frequencies are zeroed and the Moore–Penrose pseudo-inverse is calculated:

$$\mathbf{R}^p(\theta) = [\mathbf{m}^T \mathbf{C}^p(\theta) \mathbf{m}]^+, \quad (55)$$

where \mathbf{m} is a vector that is 1 for valid frequencies and 0 for masked frequencies. The filter is then applied to each map pixel independently,

$$M_{\text{hp}f}^p(\nu_m, \theta, \phi) = \sum_n R_{mn}^p(\theta) M^p(\nu_n, \theta, \phi). \quad (56)$$

The weights are also propagated through the filtering operation according to

$$w_{\text{hp}f}^p(\nu_m, \theta, \phi) = \left[\sum_n (R_{mn}^p(\theta))^2 (\sigma_{\text{map}}^p(\nu_n, \theta, \phi))^2 \right]^{-1}, \quad (57)$$

where $\sigma_{\text{map}}^p(\nu, \theta, \phi)$ is given by Equation (50).

In order to find an appropriate delay cut, the delay power spectrum of the map is estimated as

$$P(\tau, \theta) = \text{Var}_\phi \{ \tilde{M}^p(\tau, \theta, \phi) \}, \quad (58)$$

where \tilde{M} denotes the Fourier transform of the map along the frequency axis. Direct calculation of \tilde{M} through the FFT will result in a point-spread function in delay space that has large sidelobes owing to the band-limited and irregularly spaced nature of our map spectra. This will leak power from the bright foreground to higher delays, thus biasing our determination of τ_{cut} . To address this, the delay power spectrum is estimated using a Gibbs sampling method, which is described in detail in Appendix A.

Figure 10 shows the delay power spectrum of the map as a function of decl. for each polarization. Note that the variance was calculated over $\phi \in [110^\circ, 263^\circ]$, which corresponds to the range of R.A. covered by the eBOSS NGC field. The delay power spectrum is normalized by the delay power spectrum of the expected radiometric noise. This is obtained by applying the mapmaking and delay power spectrum estimation to a Gaussian noise realization randomly drawn according to Equation (52).

At high delays, the measured spectrum is in a good agreement with our expectation for the noise, and at low

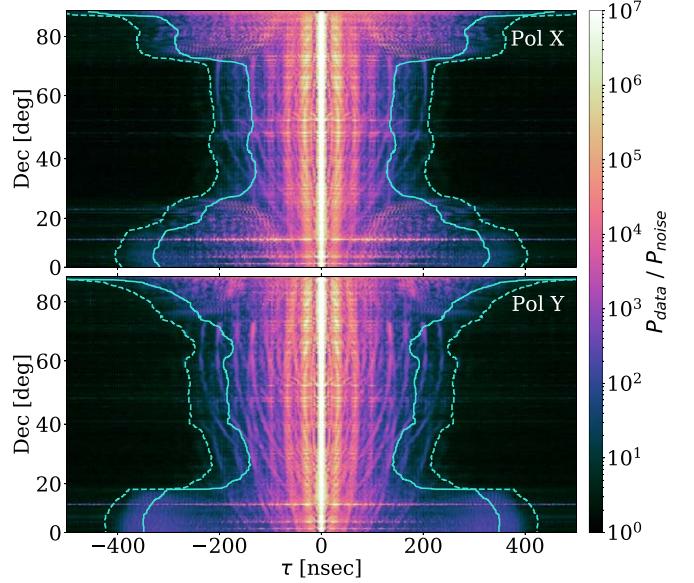


Figure 10. Delay power spectrum of the deconvolved map for the X (top) and Y (bottom) polarizations, normalized by our expectation for the radiometric noise. The delay power spectrum is obtained by computing the variance of the delay spectrum of the map over $\phi \in [110^\circ, 263^\circ]$ (see Appendix A for details). The dashed line is obtained by finding the minimum delay where $P_{\text{data}}/P_{\text{noise}} < 3$ at each decl. The solid line is obtained by subtracting 75 ns from the dashed line and corresponds to the delay cut that is used in this analysis. Note that the x -axis has been restricted in this figure: we are sensitive out to 1250 ns, but beyond 500 ns the measured spectrum matches our expectation for radiometric noise.

delays we are dominated by foreground emission. Ideally all foreground power would be contained within the bright peak centered on 0 ns. However, the ripples in the primary beam are imprinted on the foregrounds, leaking power to higher delays. The three additional peaks observed at integer multiples of ~ 30 ns correspond to interference of the primary path through the telescope with secondary paths that have undergone one, two, and three additional reflections off the focal line and cylinder. The amplitude of these peaks has been reduced by deconvolving the model for the primary beam; however, they are still significant compared to our expectation for the noise. We are actively working on improving the accuracy of our beam model and implementing a deconvolution procedure that better addresses the off-axis response (see S15 for one example) to further reduce the amplitude of these peaks. The “U”-shaped tracks in the delay power spectrum correspond to the brightest point sources moving through the far sidelobes. In this case, there is a delay associated with the EW component of the baseline that is not corrected by the mapmaking procedure because it assumes that the instrument has no sensitivity outside the main lobe of the primary beam. For each bright source, there are three “U”-shaped tracks corresponding to the three intercylinder, EW baseline separations that extend out to progressively higher delays. Finally, at large zenith angle, or low and high decl., the foreground power extends out to higher delays owing to aliasing of the sky in the baselines with one-cylinder EW separation, as explained in Section 4.3.

The stop-band rejection is set to $\epsilon = 10^{-12}$, which is much smaller than the inverse of the dynamic range of the delay power spectrum ($\sim 5 \times 10^{-9}$). This ensures that the brightest foreground features near 0 ns delay are attenuated to well below the radiometric noise level.

Our initial delay cut is defined as the minimum delay where the measured power spectrum is less than 3 times the power spectrum of the Gaussian noise realization. This is indicated by the dashed cyan line in Figure 10 and results in an aggressive filter that yields a map that is dominated by radiometric noise. However, the dominant contamination at delays just below our aggressive cut originates from a few bright sources in the far sidelobes, which are easily masked. In an attempt to maximize signal-to-noise ratio, we examine four different delay cuts that correspond to the aggressive cut minus [25 ns, 50 ns, 75 ns, 100 ns]. For each cut, a foreground filter is constructed and applied to both the data and a simulation of the 21 cm signal. Next, the regions around known bright point sources are masked. The foreground-filtered data and signal are then pushed through the rest of the analysis pipeline, which is described in the sections that follow. As the delay cut is reduced, the relative increase in the noise is compared to the relative reduction in the amplitude of the simulated 21 cm signal. The aggressive cut minus 75 ns results in the maximum signal-to-noise ratio of the four values tested. This is indicated by the solid blue line in Figure 10 and will be used as the delay cut $\tau_{\text{cut}}^p(\theta)$ for the rest of the analysis.

4.6. Additional Masking

The foreground filter heavily attenuates the signal at frequencies near the edges of the 587.5–800 MHz band and at frequencies neighboring large spans of masked frequencies. These heavily attenuated frequencies would be improperly upweighted when stacking on external catalogs because the pipeline accounts for the fact that the noise has been attenuated but does not account for the fact that the signal has also been attenuated. To address this, at each polarization and decl. the median value of the nonzero diagonal elements of the filter is calculated. Any frequency where the diagonal element of the filter is less than 20% of the median is masked. This removes approximately 4.2% of the 587.5–800 MHz band.

The simulations described in Section 5.3 predict that the rms of the radiometric noise is more than an order of magnitude larger than the rms of the 21 cm signal in the foreground-filtered, deconvolved map. The distribution of map pixel values is largely set by the radiometric noise, and the 21 cm signal is a small perturbation that is only evident after averaging over a large number of sources. The map does contain residual foregrounds, RFI, and instrumental artifacts that are large compared to the propagated fast-cadence estimate for the noise. However, this excess noise is for the most part restricted to specific frequency bins or localized to regions on the sky. The subset of the data that exceeds our expectation for the noise is masked using the following procedure.

The foreground-filtered, deconvolved map is standardized by dividing the value of each map pixel by the standard deviation from the fast-cadence estimate. These standardized maps are examined manually at each frequency channel. Any channel that contains residuals that both are large compared to the expected noise and corrupt a significant portion of the NGC field are masked. Note that the delay filter couples frequency channels, so a channel may show significant residuals due to the filter leaking some narrowband artifact from an adjacent channel. This can be disentangled for the most part by identifying artifacts with a common spatial profile across frequencies and then masking the channel where that artifact has the largest magnitude. It could also be

automated through an iterative procedure of masking and foreground filtering. In the end, 14.7% of the 587.5–800 MHz band is discarded in this way. We note that newer versions of the pipeline with improved RFI excision have reduced this fraction to roughly 5%, and most of these frequency channels are believed to be recoverable in future analyses by making additional improvements to the RFI excision algorithm and by further vetting the time ranges that are included in the sidereal stack. The frequency mask generated through this procedure is applied to the unfiltered map, and the foreground filter is reapplied.

The total fraction of the 587.5–800 MHz band that remains after removing persistent RFI bands, frequencies that do not have complete sidereal coverage, frequencies that have low integration time for the NGC field, frequencies that show excess noise, and frequencies near the edges of the large gaps of missing data is 48.6%. Finally, any map pixel whose absolute value is greater than 6σ is masked, where σ is again obtained from the fast-cadence noise estimate. This removes 1.4% of the remaining map pixels within the NGC field. The choice of a 6σ threshold was informed by signal injection simulations that are described in Section 7.4. The threshold is large enough that the resulting bias in the amplitude of the 21 cm signal is small compared to the statistical uncertainty. Figure 11 shows the map at several stages of the pipeline processing for a typical frequency channel; the third and fourth panels depict the application of the 6σ mask.

Figure 12 shows in black the standard deviation of the map pixels within the NGC field as a function of frequency after all masking has been applied. The noise for a single frequency channel is on average $0.6 \text{ mJy beam}^{-1}$, with an increase to $1.0 \text{ mJy beam}^{-1}$ in the upper 50 MHz of the band. The increased noise at high frequencies is driven by a reduction in the primary beam response on meridian when averaged over the decl. spanned by the NGC field, a mild increase in the system temperature, and frequent flagging of the highest frequencies ($\gtrsim 794 \text{ MHz}$) by the threshold applied in the sidereal regridding stage of the pipeline (see Section 3.2.6). This last item will be corrected in future revisions of the pipeline. The noise in the map is on average 50% greater than the expected radiometric noise, which is shown in red. To generate the expected radiometric noise, visibilities are drawn randomly according to Equation (52) and then propagated through the mapmaking and foreground filtering procedure. For comparison, we also show in blue the standard deviation of the map pixels in a jackknife of even and odd days. The procedure for constructing this jackknife will be described in Section 7.2. The noise in the jackknife is in better agreement with the expected radiometric noise, except in a few 6 MHz wide bands where there is still unmasked, transient RFI. This is due to the fact that the residual foregrounds are largely due to instrument chromaticity that is the same from day to day and thus cancels in the jackknife. Note that a slightly different frequency mask was used for the jackknife because the frequency channels at the upper edge of the band do not have full coverage of the sidereal day in the even or odd split. This results in the jackknife noise dropping below the expected radiometric noise because the large filter attenuation at the upper edge of the band is pushed to lower frequencies.

Using a more aggressive mask (3σ) removes 6.6% of the map pixels within the NGC field and brings the measured noise to within 22% of the expected radiometric noise on average at

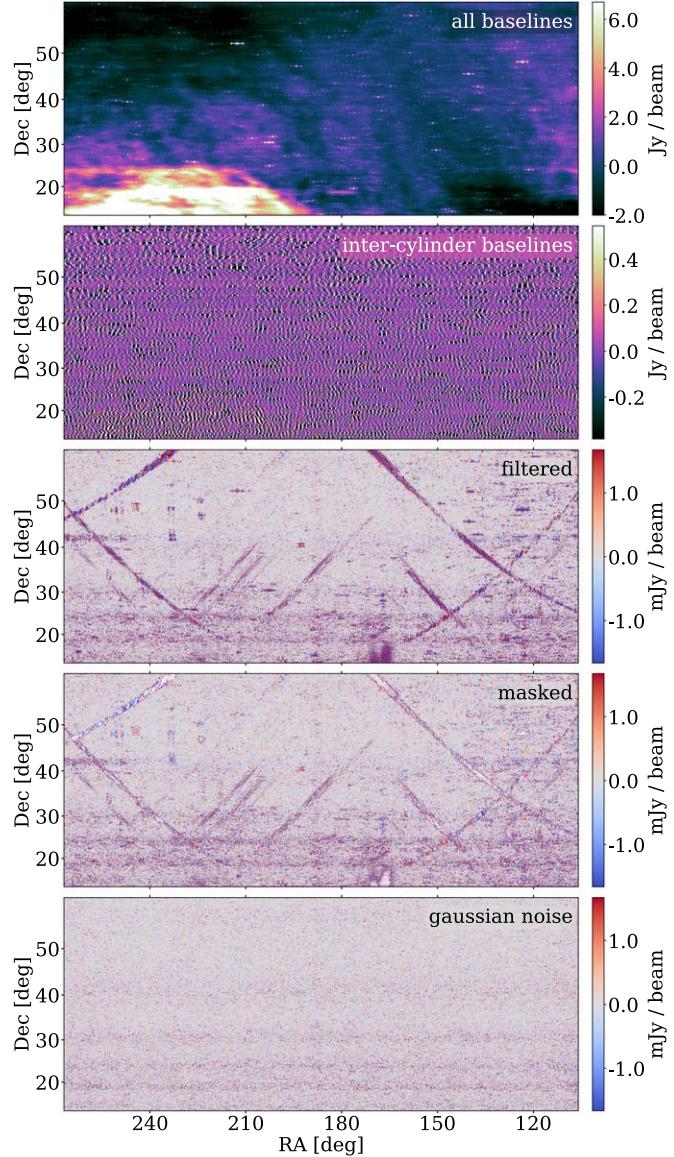


Figure 11. The deconvolved map at 700.78125 MHz at several stages of the processing. The range of R.A. and decl. matches that of the eBOSS NGC field. The top panel shows the map constructed from all Y polarization baselines (excluding autocorrelations). The second panel shows the map constructed from only the intercylinder baselines, which resolve out the diffuse Galactic emission, leaving primarily emission from extragalactic point sources. The range on the color scale has been compressed by a factor of ~ 15 . The third panel shows the intercylinder map after applying the delay filter. The range on the color scale has been further compressed by a factor of ~ 300 . Residuals associated with very bright sources in the far sidelobes, bright sources in the main lobe, and instrumental artifacts are evident. However, they are localized, and the fourth panel shows the result of masking any pixel that is more than 6 times the standard deviation of the expected radiometric noise (1.3% of the pixels). This can be compared to the bottom panel, which shows a realization of the radiometric noise generated according to Equation (52). The horizontal features in the bottom two panels are due to the ripples in the primary beam pattern, which are imprinted on the noise during deconvolution.

the expense of introducing a significant nonlinearity into the analysis pipeline. Using a more aggressive mask (3σ) and more aggressive delay filter with a cutoff that is 75 ns larger (indicated by the dashed cyan line in Figure 10) brings the measured noise to within 12% of the expected radiometric noise on average. However, this results in a significant

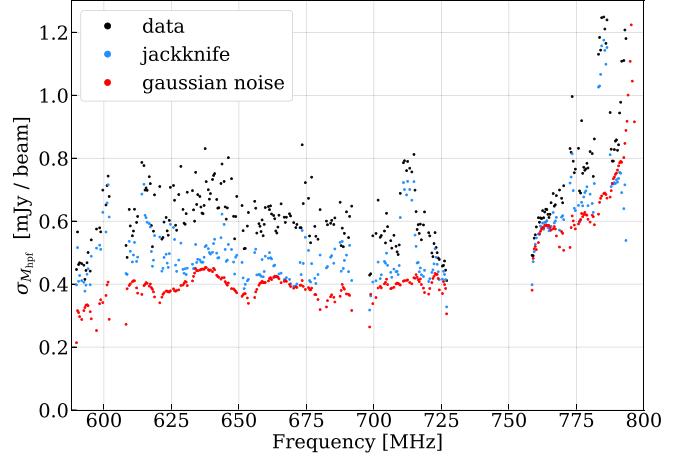


Figure 12. Standard deviation of pixels within the NGC field in the deconvolved, foreground-filtered map. Black denotes the measured standard deviation. Red denotes the expected standard deviation due to radiometric noise, which is based on the fast-cadence estimate of the variance. Blue denotes the measured standard deviation in a jackknife of even and odd days (see Section 7.2). The measured noise in a single frequency channel is on average $0.6 \text{ mJy beam}^{-1}$, with an increase to $1.0 \text{ mJy beam}^{-1}$ in the upper 50 MHz of the band. This is on average 50% greater than the expected radiometric noise, due to residual foregrounds and RFI. The residual foregrounds are largely the same from day to day and therefore cancel in the even-odd jackknife, but the RFI does not. As a result, the jackknife is in better agreement with the radiometric noise outside of certain 6 MHz wide bands that still suffer from unmasked, transient RFI.

reduction in the amplitude of the stacked 21 cm signal in simulations, and a better signal-to-noise ratio is anticipated using the less aggressive delay cut. Note that with either mask or delay cut the excess noise from residual foreground, RFI, and instrumental artifacts is comparable to or less than the radiometric noise for the 102-day sidereal stack, assuming that they add in quadrature.

4.7. Stacking

For each source in a given eBOSS catalog, a spectral cube centered on the source’s location is extracted from the deconvolved, foreground-filtered map. First, the R.A. and decl. of the source are converted from ICRS to CIRS coordinates to account for the precession and nutation of Earth’s polar axis. The redshift of the source is converted to the frequency of the redshifted 21 cm emission,

$$\nu_{21\text{cm}} = \frac{1420.406 \text{ MHz}}{1 + z}. \quad (59)$$

The map pixel and frequency channel closest to these coordinates are found, and ± 50 pixels (channels) are extracted in the angular (frequency) directions. This results in a spectral cube that spans $\pm 3^\circ$ in R.A./decl. and ± 20 MHz in frequency.

The stacked signal is given by the weighted average of the spectral cubes over all sources in the catalog:

$$\begin{aligned} d^p(\Delta\nu, \Delta\theta, \Delta\phi) &= \sum_s W_{\text{hp}}^p(\nu_s + \Delta\nu, \theta_s + \Delta\theta, \phi_s + \Delta\phi) \\ &\times M_{\text{hp}}^p(\nu_s + \Delta\nu, \theta_s + \Delta\theta, \phi_s + \Delta\phi), \end{aligned} \quad (60)$$

where $(\nu_s, \theta_s, \phi_s)$ denote the frequency channel and map pixel closest to the coordinates of source s and

$$\begin{aligned} W_{\text{hp}}^p(\nu_s + \Delta\nu, \theta_s + \Delta\theta, \phi_s + \Delta\phi) \\ = \frac{w_{\text{hp}}^p(\nu_s + \Delta\nu, \theta_s + \Delta\theta, \phi_s + \Delta\phi)}{\sum_s w_{\text{hp}}^p(\nu_s + \Delta\nu, \theta_s + \Delta\theta, \phi_s + \Delta\phi)} \end{aligned} \quad (61)$$

denotes the relative weight given to source s , with the absolute weight w_{hp}^p given by Equation (57).

Note that we make no attempt to interpolate the spectral cubes onto a common grid relative to the coordinates of the source. Instead, we take a forward modeling approach where the stacking procedure is applied to simulations in order to characterize how the pixelization alters a stack of the 21 cm signal. This will result in a small degradation in signal-to-noise ratio because we are not stacking on the true peak, but given the pixelization used, we estimate this to be only $\sim 3\%$ for the NGC field.

For simplicity, all model fitting and parameter estimation uses only the central pixel of the stack as a function of frequency offset. Going forward we will use $d^p(\Delta\nu) \equiv d^p(\Delta\nu, 0, 0)$ to describe the stack of the pixels closest to the coordinates of the sources.

4.8. Noise Covariance Estimation

The probability distribution of the noise in the stack must be characterized in order to derive accurate uncertainties on the inferred model parameters. As discussed in Section 4.6, residual foregrounds and RFI are expected to be subdominant but significant contributors to the noise, and both are likely correlated between frequencies. More generally, the foreground filter couples all frequency channels, ensuring a nonzero correlation between frequency offsets in the stack. These factors are not accounted for in the propagated fast-cadence noise estimate, which only includes the radiometric contribution to the noise and does not account for the correlation between frequency channels. To develop an accurate noise model, we stack the data on a large number of random mock catalogs and examine the distribution of values.

Each eBOSS clustering catalog has a corresponding “random” catalog that approximates the 3D selection function of the clustering catalog and is more than 40 times as dense (Ross et al. 2020; Raichoor et al. 2021). We randomly sample the random catalog without replacement to generate a mock catalog that has the same number of sources as the true catalog. The deconvolved, foreground-filtered map is then stacked on the mock catalog following the same procedure described in Section 4.7. This process is repeated $N_{\text{mock}} = 10,000$ times.

The noise covariance of the stacked data is estimated using the sample covariance of the mocks

$$\hat{\Sigma} = \frac{1}{N_{\text{mock}} - 1} \sum_{m=1}^{N_{\text{mock}}} (\mathbf{d}_m - \hat{\mu})^T (\mathbf{d}_m - \hat{\mu}), \quad (62)$$

where

$$\mathbf{d}_m = [d_m^p(\Delta\nu_i, 0, 0); 0 \leq i \leq 100, p \in \{X, Y\}] \quad (63)$$

is a vector containing the stacked signal at the central pixel as a function of frequency offset for both polarizations for the m th

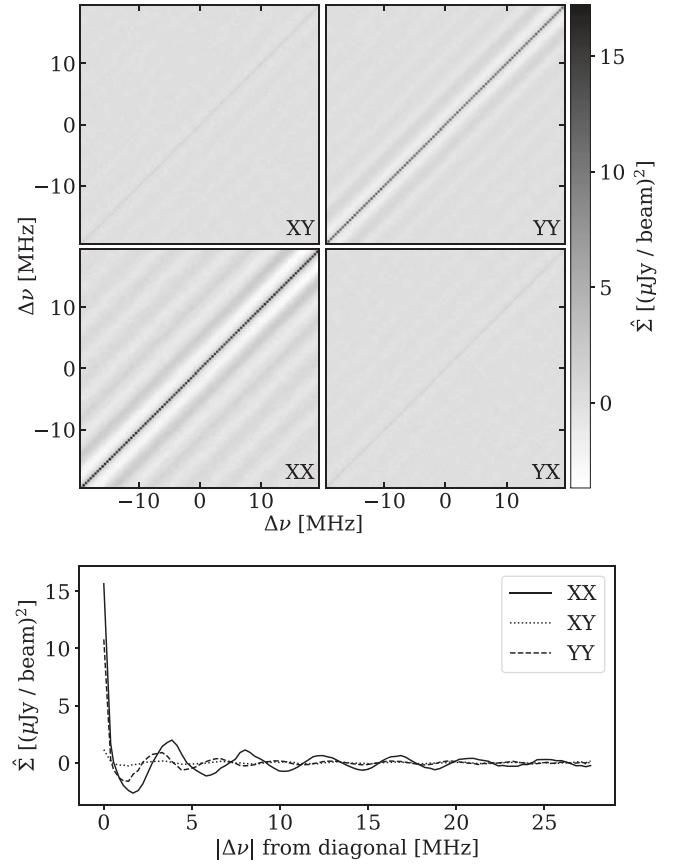


Figure 13. Estimated noise covariance of the NGC QSO stack, obtained by computing the sample covariance of stacks on 10,000 random mock catalogs. Each subpanel in the top panel shows the covariance between frequency offsets for a different pair of polarizations. The bottom panel shows the average value of the covariance as a function of distance from the central diagonal of each subpanel, with the different polarization pairs denoted using different line styles as indicated in the legend. The noise in the two polarizations is largely independent. However, there is nonnegligible correlation in the noise between frequency offsets within a polarization. The sinc-like dependence on frequency offset is introduced by the foreground filter (see Equation (54)). The period of the ripple is different for the two polarizations because a different delay cut was used on average.

mock catalog and

$$\hat{\mu} = \frac{1}{N_{\text{mock}}} \sum_{m=1}^{N_{\text{mock}}} \mathbf{d}_m \quad (64)$$

is the sample mean of the mocks. An example of the sample covariance is shown in Figure 13.

We find that the sample mean for a given frequency offset and polarization is nonzero at a level larger than expected given the standard error. The rms of the sample mean over all frequency channels and polarizations is $\sigma_{\hat{\mu}} = 0.7\text{--}1.9 \mu\text{Jy beam}^{-1}$ depending on the tracer, which is roughly 20% of the sample standard deviation over mock catalogs and a factor of 20 times larger than the standard error. This sample mean over mocks is subtracted from the stack on the true catalog to ensure a consistent noise model.

We find that the distribution of values observed in the mocks is consistent with a multivariate Gaussian whose covariance and mean are set to the sample variance and mean as calculated above.

4.9. Future Improvements

In this section we summarize improvements to the data processing that are expected in the short term. Real-time RFI excision was deployed on the CHIME correlator in 2019 October, and its performance is summarized in CHIME Collaboration et al. (2022b). In addition, we have developed more effective methods for offline RFI excision that can be applied retroactively to all past data. We have also improved the tests that are used to identify anomalous data in the individual sidereal days. As a result, we expect to recover most of the 14.7% of the 585–800 MHz band that was masked during manual inspection of the foreground-filtered maps. We also expect that these changes will mitigate artifacts responsible for the excess noise observed at high delays and enable an analysis of the 400–585 MHz band.

Our analysis must include intracylinder baselines in order to recover the large angular scales relevant for measuring the imprint of the baryonic acoustic oscillation (BAO) on the power spectrum of 21 cm emission. Preliminary work suggests that we can include the majority of intracylinder baselines in our analysis if we also apply a high-pass filter along the time axis to remove the additive noise crosstalk. This also removes the largest angular modes of the sky but retains those angular modes that probe the BAO.

Similarly, our analysis must use a less aggressive delay filter in order to recover large line-of-sight scales. Currently we are limited at low delays by sources in the far sidelobes and uncalibrated spectral structure in the primary beam. The development of a full-sky model for the primary beam and the deconvolution of the off-axis response has been a major focus of the CHIME collaboration over the past several years. Some of the progress on the full-sky beam model has been reported in CHIME Collaboration et al. (2022a, 2022b). An algorithm for performing the deconvolution is described in Shaw et al. (2014) and S15.

Although not currently limiting our results, the stability of the complex receiver gains is expected to improve in future analyses. In 2020 July an ADC-dependent correction for the clock drift was deployed on the CHIME correlator, which is expected to improve the delay noise on short timescales ($\lesssim 20$ minutes) by roughly a factor of 2. In addition, techniques are being developed to perform an offline correction for thermal expansion of the focal line, which is the dominant source of phase instability on longer timescales (CHIME Collaboration et al. 2022b).

Finally, the results presented in this work employed 521 hr of total integration time. The CHIME archive now contains almost 20 times this amount after accounting for the flags described in Sections 3.2.7 and 3.3.1. Inclusion of these data will significantly reduce the radiometric noise and certain types of systematic errors, for example, those due to RFI.

5. Signal Modeling and Simulations

Interpreting our stacking measurements requires that we are able to predict the cosmological signal within them and that we understand the performance of our analysis pipeline, including any signal loss that has occurred. In this section we discuss the framework to address these: a parameterized model of the cosmological signal, a simulation pipeline producing synthetic time streams and source catalogs, and a scheme for using these

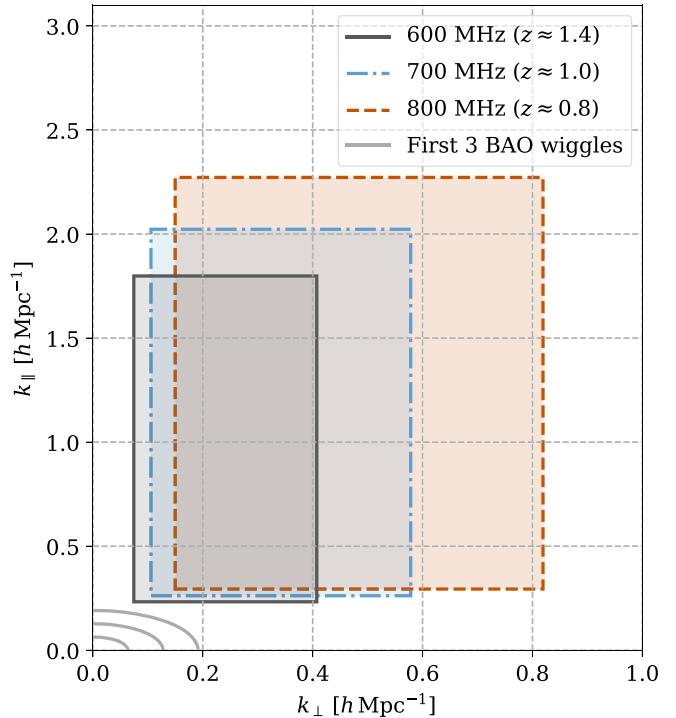


Figure 14. Approximate physical scales probed by the stacking measurements, as comoving wavenumbers along the line of sight (k_{\parallel}) or transverse to it (k_{\perp}). We evaluate the range of scales at three observing frequencies that span the relevant portion of the CHIME band. The accessible values of k_{\parallel} are determined by the CHIME frequency channel width and delay filtering prescription, while the ranges of k_{\perp} arise from the synthesized beamwidth and choice to exclude intracylinder baselines; see main text for details. The maxima of the first three BAO wiggles in the matter power spectrum are shown by the gray lines, making it apparent that the measurements in this work are insensitive to BAO scales and instead mainly probe the nonlinear regime of structure formation.

simulations to predict the stack signal from the parameters of our model.

5.1. Cosmological Scales Being Probed

To set the stage for the modeling approach described later in this section, in Figure 14 we show the approximate range of physical scales probed by our stacking measurements, represented as comoving wavenumbers k_{\parallel} (along the line of sight) and k_{\perp} (transverse to the line of sight). This range depends on observing frequency due to the relationship between frequency and radial distance and also due to chromaticity of CHIME’s beam response, so we show results at three frequencies within the portion of the band used in our analysis.

The foreground filter described in Section 4.5 acts roughly as a high-pass filter in k_{\parallel} , with the minimum accessible k_{\parallel} determined by the delay cut τ_{cut} ; for Figure 14, we use $\tau_{\text{cut}} = 200$ ns, reflective of the typical delay cut within the decl. covered by the eBOSS catalogs. The sensitivity at high k_{\parallel} is attenuated by the finite width of CHIME’s frequency channels, which we approximate as top hats with width 390.625 kHz.

Similarly, the sensitivity at high k_{\perp} is determined by the profile of the synthesized beam associated with the maps described in Section 4.3. For Figure 14, we use the simplified 1D beamforming result from Masui et al. (2017) to obtain NS and EW synthesized beam profiles based on CHIME’s feed layout and the analytical (“control”) primary beam model

discussed in Section 4.4, take the geometric mean of the FWHMs in each direction, and translate this into a comoving wavenumber at each plotted frequency. Finally, since intracylinder baselines are excluded from our analysis, we are not sensitive to any angular scales that are only probed by pure NS baselines; these scales are determined by the EW primary beam profile, and we translate the EW FWHM into a minimum accessible k_{\perp} .

Note that a more thorough treatment of the scales being probed is possible, in which the stacking measurements can be related to an integral of the galaxy–HI cross-power spectrum multiplied by a transfer function $W(k_{\parallel}, k_{\perp})$ that precisely encodes the sensitivity of our analysis to a given Fourier mode. Such a treatment is currently under development and will be presented in a forthcoming publication (CHIME Collaboration 2023, in preparation), but preliminary results are in good agreement with the estimates in Figure 14.

In this figure, we also show the maxima of the first three BAO wiggles in the matter power spectrum, located at multiples of $k_{\text{BAO}} = 2\pi/r_{\text{drag}} \approx 0.064 h^{-1} \text{Mpc}$. It is clear that our delay filter and exclusion of intracylinder baselines have effectively filtered out any sensitivity to BAO scales from our stacking measurements. The scales that remain are beyond the reach of analytical perturbative methods for LSS statistics in Fourier space (e.g., d’Amico et al. 2020; Ivanov et al. 2020; Chen et al. 2022); while these scales have some overlap with those accessible to hybrid simulation–perturbation theory methods (e.g., Kokron et al. 2021), the majority of our signal-to-noise ratio lies at even smaller scales, implying that we cannot immediately apply those methods in our present analysis.

Halo-based models for HI (e.g., Padmanabhan 2021) and galaxy clustering can in principle describe the full range of scales shown in Figure 14. However, we have found that a simpler model, which makes efficient use of our simulation framework described in Section 5.3, is fully capable of describing the observed signal while allowing for marginalization over hard-to-predict properties of nonlinear clustering. We describe this model and its application to our measurements in the following subsections.

5.2. Signal Model

Cosmological modeling of the distribution of galaxies²⁴ and HI typically begins with the matter overdensity $\delta_m(\mathbf{x}, z) \equiv [\rho_m(\mathbf{x}, z) - \bar{\rho}_m(z)]/\bar{\rho}_m(z)$, where an overbar denotes a spatial average. In our modeling we assume that galaxies and HI are each linearly biased tracers of the total matter density. The overdensity corresponding to galaxy or HI number density, δ_g or δ_{HI} , can then be written in Fourier space as

$$\delta_{\alpha}(\mathbf{k}; z) = [b_{\alpha}(z) + f(z)\mu^2]\tilde{D}_{\alpha}^{\text{FoG}}(k\mu, z)\delta_m(\mathbf{k}, z) + \epsilon_{\alpha}(z), \quad (65)$$

with $\alpha \in [g, \text{HI}]$. In Equation (65), b_{α} is the bias factor (assumed to be scale independent), and the $f(z)\mu^2$ term encodes the effect of redshift-space distortions at linear order (Kaiser 1987), with f as the logarithmic growth rate and $\mu \equiv k_{\parallel}/k$. We aim to capture the key nonlinear contributions to the two-

point statistics of the fields: we include real-space nonlinear clustering in δ_m itself; the impact of small-scale velocities on redshift-space observations (“Finger of God”; Jackson 1972) is modeled with the damping function $\tilde{D}_{\alpha}^{\text{FoG}}$; and finally, we include a term ϵ_{α} in Equation (65), which is uncorrelated with δ_m and represents the contribution of shot noise to δ_{α} .

In our analysis we will only require the two-point statistics of the correlated fields. These are captured entirely by the power spectrum of two fields:

$$\begin{aligned} P_{\alpha\beta}(\mathbf{k}; z_{\alpha}, z_{\beta}) &= [b_{\alpha}(z_{\alpha}) + f(z_{\alpha})\mu^2][b_{\beta}(z_{\beta}) + f(z_{\beta})\mu^2] \\ &\times \tilde{D}_{\alpha}^{\text{FoG}}(k\mu, z_{\alpha})\tilde{D}_{\beta}^{\text{FoG}}(k\mu, z_{\beta})P_m(\mathbf{k}; z_{\alpha}, z_{\beta}) \\ &+ P_{\alpha\beta}^{\text{shot}}(z_{\alpha}, z_{\beta}). \end{aligned} \quad (66)$$

The ingredients required to complete our model are functions for the nonlinear matter power spectrum P_m , the linear bias b_{α} , the Finger-of-God function $\tilde{D}_{\alpha}^{\text{FoG}}$, and the shot noise $P_{\alpha\beta}^{\text{shot}}$. We discuss our fiducial choices for these ingredients in the following sections.

5.2.1. Matter Power Spectrum

As input to our simulations, we use the halo model prediction for the nonlinear matter power spectrum from Mead et al. (2021), as implemented in the CAMB code (Lewis et al. 2000). We have also considered the Halofit fitting functions from Smith et al. (2003) and Takahashi et al. (2012) and have found that these different choices affect the final stacking amplitude in the simulations by at most $\sim 3\%$, with little change in the shape. Thus, the uncertainty arising from the specific choice of nonlinear matter power spectrum is far subdominant to the uncertainty inherent in our assumption of linear, scale-independent bias in Equation (65).

5.2.2. Linear Bias

We assume the following for the linear bias of each eBOSS sample:

$$b_{\text{ELG}}(z) = 1.5 + 0.7(z - 0.85), \quad (67)$$

$$b_{\text{LRG}}(z) = 2.03 + 0.86(z - 0.4) + 0.13(z - 0.4)^2, \quad (68)$$

$$b_{\text{QSO}}(z) = 2.38 + 1.4(z - 1.55) + 0.28(z - 1.55)^2. \quad (69)$$

The ELG bias uses the redshift evolution of the linear bias predicted by the simulations of Merson et al. (2019), normalized such that Equation (67) evaluates to the bias measurement from de Mattia et al. (2021) at the mean redshift of the eBOSS ELG sample. The LRG bias is based on Zhai et al. (2017), who fit a halo occupation distribution model to small-scale clustering of a combined BOSS+eBOSS LRG sample and computed the linear bias from this model. Specifically, Equation (68) is the result of a quadratic fit to the best-performing bias model from Figure 12 of Zhai et al. (2017). The QSO bias is taken from the fitting function in Laurent et al. (2017), based on measurements of the eBOSS QSO correlation function in four redshift bins.

For the linear bias of HI, we follow Cosmic Visions 21 cm Collaboration et al. (2018) in smoothly interpolating between measurements from the IllustrisTNG simulations (Villaescusa-Navarro et al. 2018) at $z < 2$ and the analytical model from

²⁴ For brevity, we refer to ELGs, LRGs, and QSOs as “galaxies” in this section.

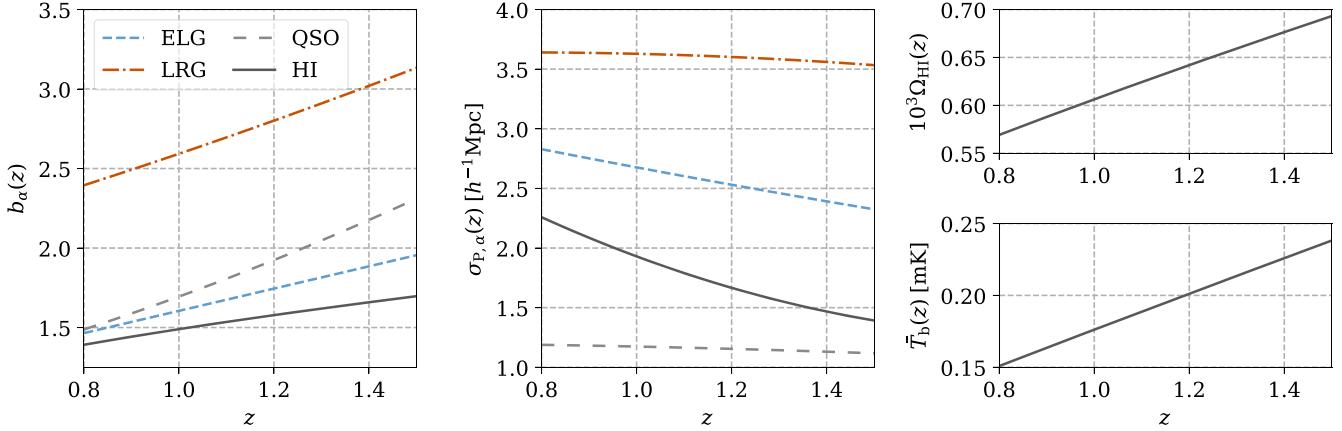


Figure 15. Fiducial models for various redshift-dependent quantities used in our simulated sky maps. See Sections 5.2.2–5.2.4 for discussions of how each model was chosen. Left: linear bias of each eBOSS sample and H I. Middle: Finger-of-God damping scale, with same line styles as in the left panel. Top right: H I density, as a fraction of the critical density at $z = 0$. Bottom right: mean 21 cm brightness temperature.

Castorina & Villaescusa-Navarro (2017) at $z > 2$.²⁵ We show our bias models for H I and each eBOSS sample in the left panel of Figure 15.

5.2.3. Finger-of-God Models

We model the Finger-of-God damping in Fourier space as a Lorentzian:

$$\tilde{D}_\alpha^{\text{FoG}}(k\mu, z) = \frac{1}{1 + k^2 \mu^2 \sigma_{P,\alpha}(z)^2 / 2}, \quad (70)$$

where the damping scale $\sigma_{P,\alpha}$ can approximately be associated with the pairwise velocity dispersion of galaxies or H I emitters on nonlinear scales. The (constant-redshift) Fourier conjugate of this function is an exponential in comoving distance (e.g., Scoccimarro 2004),

$$D_\alpha^{\text{FoG}}(x_\parallel, z) = \frac{e^{-|x_\parallel| \sqrt{2} / \sigma_{P,\alpha}(z)}}{\sqrt{2} \sigma_{P,\alpha}(z)}, \quad (71)$$

and we implement the Finger-of-God effect by convolving our simulated maps with this kernel along the line-of-sight axis. This is equivalent to multiplying the 3D auto-power spectrum of tracer α by $\tilde{D}_\alpha^{\text{FoG}}(k\mu, z)^2$ and multiplying the cross-power spectrum of H I and tracer α by $\tilde{D}_{\text{HI}}^{\text{FoG}}(k\mu, z) \times \tilde{D}_\alpha^{\text{FoG}}(k\mu, z)$.

For each eBOSS sample, Fourier-space clustering measurements have been analyzed using Finger-of-God models similar to what we describe above. For ELGs and LRGs, de Mattia et al. (2021) and Gil-Marín et al. (2020) use a squared Lorentzian function multiplied into the 3D galaxy power spectrum, finding best-fit values of $\sigma_{P,\text{ELG}} = 2.79 h^{-1} \text{ Mpc}$ at $z_{\text{eff}} = 0.85$ and $\sigma_{P,\text{LRG}} = 3.64 h^{-1} \text{ Mpc}$ at $z_{\text{eff}} = 0.7$ (where we quote the average of separate fits to the NGC and SGC fields). For QSOs, Zarrouk et al. (2018) use a Gaussian Finger-of-God model and perform fits that isolate the contribution to this model from small-scale velocities (as opposed to QSO redshift errors, which have a similar effect on the observed clustering). Taking the average of their best-fit σ_P values for the “3-multipole” and “3-wedge” analyses yields $\sigma_{P,\text{QSO}} =$

$1.3 h^{-1} \text{ Mpc}$ at $z_{\text{eff}} = 1.48$. We find that this is roughly equivalent to Lorentzian damping with $\sigma_{P,\text{QSO}} = 1.12 h^{-1} \text{ Mpc}$.

We use these values to fix the amplitude of our fiducial $\sigma_P(z)$ models for each sample. We compute the redshift dependence from a simple model in which $\sigma_P(z)$ scales like a weighted average of the velocity dispersion $\sigma_v^2(M, z)$ of a dark matter halo of mass M , weighted by the halo mass function dn/dM and the mean satellite occupation in a mass- M halo:

$$\sigma_{P,\alpha}(z) \propto \frac{1+z}{H(z)} \left[\frac{\int dM \frac{dn}{dM} N_{\text{sat},\alpha}(M) \sigma_v^2(M, z)}{\int dM \frac{dn}{dM} N_{\text{sat},\alpha}(M)} \right]^{1/2}. \quad (72)$$

To evaluate Equation (72), we use the halo mass function from Tinker et al. (2008) and the eBOSS halo occupation distribution models from Alam et al. (2020). The final results for $\sigma_{P,\alpha}(z)$, incorporating the amplitude constraints described above, are well fit by quadratic functions of redshift, which we present below:

$$\frac{\sigma_{P,\text{ELG}}(z)}{h^{-1} \text{ Mpc}} = 2.79 - 0.77(z - 0.85) + 0.083(z - 0.85)^2, \quad (73)$$

$$\frac{\sigma_{P,\text{LRG}}(z)}{h^{-1} \text{ Mpc}} = 3.64 + 0.019(z - 0.7) - 0.19(z - 0.7)^2, \quad (74)$$

$$\frac{\sigma_{P,\text{QSO}}(z)}{h^{-1} \text{ Mpc}} = 1.12 - 0.14(z - 1.48) - 0.058(z - 1.48)^2. \quad (75)$$

For H I, we choose the damping scale based on simulations from Sarkar & Bharadwaj (2019), who attempt to account for the motion of H I within galaxies in addition to the contribution from the velocity dispersion within dark matter halos. They assume that the Finger-of-God damping of the 3D H I power spectrum are given by a Lorentzian, and they fit a $\sigma_{P,\text{HI}}(z)$ relation to their simulations. We use these results, multiplied by a factor of $2^{-1/2}$ to translate to the damping given by a squared Lorentzian (as implied by our Equation (70)). The adopted $\sigma_{P,\text{HI}}(z)$ model is well fit by a quadratic function of redshift,

²⁵ This bias model has been implemented in the PUMANoise code, available from <https://github.com/slosar/PUMANoise>.

given by

$$\frac{\sigma_{P,\text{HI}}(z)}{h^{-1} \text{Mpc}} = 1.93 - 1.48(z - 1) + 0.81(z - 1)^2. \quad (76)$$

Over the redshift range of interest, this $\sigma_{P,\text{HI}}(z)$ model is within 20% of the values obtained in Villaescusa-Navarro et al. (2018) from fits of a squared Lorentzian to measurements from the IllustrisTNG simulations.

We plot our models for eBOSS and H I damping scales in the middle panel of Figure 15.

5.2.4. 21 cm Brightness Temperature

We convert simulated maps of δ_{HI} into brightness temperature fluctuations by multiplying by the mean 21 cm brightness temperature $\bar{T}_b(z)$. Recall that, after the end of reionization, the spin temperature T_s is high compared to both the background CMB temperature and $T_* = h_{\text{Pl}}\nu_{21}/k_{\text{B}}$. In this limit, the 21 cm brightness temperature can be written as (e.g., Bull et al. 2015)

$$T_b(\mathbf{x}, z) = \frac{3\hbar c^3 A_{10}}{16k_{\text{B}}\nu_{21}^2} \frac{(1+z)^2}{H(z)} n_{\text{HI}}(\mathbf{x}, z), \quad (77)$$

where n_{HI} is the comoving H I number density and A_{10} is the Einstein coefficient for spontaneous emission in the 21 cm line. Using $n_{\text{HI}}(\mathbf{x}, z) = \bar{n}_{\text{HI}}(z)[1 + \delta_{\text{HI}}(\mathbf{x}, z)]$, we can write

$$\begin{aligned} \delta T_b(\mathbf{x}, z) &\equiv T_b(\mathbf{x}, z) - \bar{T}_b(z) \\ &= \bar{T}_b(z)\delta_{\text{HI}}(\mathbf{x}, z), \end{aligned} \quad (78)$$

which justifies our method of converting maps of δ_{HI} into T_b . Using $\bar{n}_{\text{HI}}(z) = \Omega_{\text{HI}}(z)\rho_{c,0}/(m_{\text{p}} + m_{\text{e}})$, where $\rho_{c,0} = 3H_0^2/8\pi G$ is the critical density today,²⁶ we can write \bar{T}_b as

$$\begin{aligned} \bar{T}_b(z) &= \left[\frac{9\hbar c^3 A_{10} H_{100}}{128\pi k_{\text{B}}\nu_{21}^2 G(m_{\text{p}} + m_{\text{e}})} \right] \\ &\times h \frac{H_0}{H(z)} \Omega_{\text{HI}}(z) (1+z)^2, \end{aligned} \quad (79)$$

where $H_{100} = 100 \text{ km s}^{-1} \text{ Mpc}^{-1}$ and $h = H_0/100$. The pre-factor in square brackets is independent of cosmology, consisting only of fundamental constants and A_{10} . Using $A_{10} = 2.8843e - 15 \text{ s}^{-1}$ (Gould 1994), Equation (79) can be written more compactly as²⁷

$$\bar{T}_b(z) \approx 191.06 \left[h \frac{H_0}{H(z)} \Omega_{\text{HI}}(z) (1+z)^2 \right] \text{ mK}. \quad (80)$$

For $\Omega_{\text{HI}}(z)$, we use the fitting function from Crighton et al. (2015), which was determined from a compilation of Ω_{HI} estimates over $0 < z < 5$:

$$\Omega_{\text{HI}}(z) = 4 \times 10^{-4} (1+z)^{0.6}. \quad (81)$$

²⁶ Following the common convention in the 21 cm intensity mapping literature, we have defined $\Omega_{\text{HI}}(z)$ in terms of the comoving H I number density at redshift z and the critical density at $z = 0$.

²⁷ Other versions of Equation (80) in the literature have prefactors that vary significantly from 180 to 190 mK, most of which is accounted for by using values of A_{10} from older calculations. The value quoted in the main text is taken from a recent review of atomic transition properties (Wiese & Fuhr 2009), which takes its A_{10} value for hydrogen from Gould (1994).

We plot Equations (80) and (81) in the right panels of Figure 15.

5.2.5. Shot Noise

The cross-correlation between maps of H I and the distribution of galaxies in a given sample will be sensitive to the H I content of the galaxies. Specifically, the 3D cross-power spectrum of $\delta T_b(\mathbf{x}, z)$ and $\delta_g(\mathbf{x}, z)$ contains a cross shot-noise contribution of the form (e.g., Wolz et al. 2017)

$$P_{Tg}^{\text{shot}}(z_{\text{HI}}, z_g) = C_{\text{HI}}(z_{\text{HI}}) \langle M_{\text{HI}} \rangle_g, \quad (82)$$

where $\langle M_{\text{HI}} \rangle_g$ is the mean H I mass per galaxy in the sample and

$$C_{\text{HI}}(z) = \frac{3\hbar c^3 A_{10}}{16k_{\text{B}}\nu_{21}^2(m_{\text{p}} + m_{\text{e}})} \frac{(1+z)^2}{H(z)}. \quad (83)$$

In principle, $\langle M_{\text{HI}} \rangle_g$ depends on redshift, but for simplicity we consider a single value that is averaged over the entire sample. We also write the shot-noise contribution as being constant for all z_g , but note that there is expected to be a gradual, scale-dependent decorrelation as $|z_{\text{HI}} - z_g|$ increases, due to relative displacements of sources between different time slices.

5.2.6. Model Parameters

To produce a parameterized model of the 21 cm signal, we use the ingredients presented in Sections 5.2.1–5.2.5 as a basis and introduce a finite number of parameters that will scale their magnitude but not their redshift dependence. In total, our model contains seven parameters that are used to model the contributions to the cross-power spectrum:

Ω_{HI} : One of the key quantities controlling the stack signal is the total amount of neutral hydrogen in the universe. Although this quantity is expected to be redshift dependent, in this paper we use the model given in Equation (81) as a baseline and use a single redshift-independent parameter Ω_{HI} to scale the fiducial model about an effective redshift z_{eff} , which gives

$$\Omega_{\text{HI}}(z) = \Omega_{\text{HI}} \left[\frac{\Omega_{\text{HI}}^{\text{fid}}(z)}{\Omega_{\text{HI}}^{\text{fid}}(z_{\text{eff}})} \right]. \quad (84)$$

b_{HI}, b_g : To control the bias of the 21 cm field and galaxy density fields, which are again expected to be redshift dependent, we scale the models given in Section 5.2.2, giving

$$b_{\text{HI}}(z) = b_{\text{HI}} \left[\frac{b_{\text{HI}}^{\text{fid}}(z)}{b_{\text{HI}}^{\text{fid}}(z_{\text{eff}})} \right] \quad (85)$$

for the 21 cm field and the equivalent definition for the galaxy density,

$$b_g(z) = b_g \left[\frac{b_g^{\text{fid}}(z)}{b_g^{\text{fid}}(z_{\text{eff}})} \right]. \quad (86)$$

M_{10} : The strength of the shot-noise contribution is governed by the mass of neutral hydrogen typically associated with a tracer galaxy $\langle M_{\text{HI}} \rangle_g$. We control this quantity with the

parameter M_{10} defined by

$$\langle M_{\text{HI}} \rangle_g = M_{10} \times 10^{10} M_{\odot}. \quad (87)$$

α_{NL} : The shape of the high- k real-space cross-power spectrum is uncertain because of nonlinear gravitational evolution and baryonic effects. We let this shape vary using a linear mode that interpolates from a linear to a nonlinear power spectrum

$$P(k) = \alpha_{\text{NL}} P_{\text{NL}}(k) + (1 - \alpha_{\text{NL}}) P_{\text{L}}(k). \quad (88)$$

For $P_{\text{NL}}(k)$ we use the model described in Section 5.2.1, and for $P_{\text{L}}(k)$ we use a power spectrum with the same parameters but with the Halofit corrections turned off. This parameter is valid for $\alpha_{\text{NL}} > 0$, where values above 1 correspond to increasing the power contributed by nonlinear evolution. Although this parameter is not physically motivated, we expect it to capture the effects of nonlinearities at the level that can be measured in this work.

$\alpha_{\text{FoG,HI}}$, $\alpha_{\text{FoG,g}}$: To account for uncertainties in the Finger-of-God smoothing, we allow redshift-independent scaling of both the 21 cm and tracer velocity dispersion σ_p :

$$\sigma_{P,\text{HI}}(z) = \alpha_{\text{FoG,HI}} \sigma_{P,\text{HI}}^{\text{fid}}(z), \quad (89)$$

$$\sigma_{P,g}(z) = \alpha_{\text{FoG,g}} \sigma_{P,g}^{\text{fid}}(z). \quad (90)$$

Put together, these give a model for the cross-power spectrum of the 21 cm emission and the galactic tracer, controlled by the parameters given above. Written out fully, this gives

$$\begin{aligned} & P_{Tg}(k, \mu; z_1, z_2) \\ &= \Omega_{\text{HI}} \left[\frac{T_b^{\text{fid}}(z_1)}{\Omega_{\text{HI}}^{\text{fid}}(z_{\text{eff}})} \right] \frac{D^+(z_1)}{D^+(z_{\text{fid}})} \frac{D^+(z_2)}{D^+(z_{\text{fid}})} \\ & \times \left(b_{\text{HI}} \left[\frac{b_{\text{HI}}^{\text{fid}}(z_1)}{b_{\text{HI}}^{\text{fid}}(z_{\text{eff}})} \right] + f(z_1) \mu^2 \right) \\ & \times \left(b_g \left[\frac{b_g^{\text{fid}}(z_2)}{b_g^{\text{fid}}(z_{\text{eff}})} \right] + f(z_2) \mu^2 \right) \\ & \times (P_{\text{L}}(k, z_{\text{fid}}) + \alpha_{\text{NL}} [P_{\text{NL}}(k, z_{\text{fid}}) - P_{\text{L}}(k, z_{\text{fid}})]) \\ & \times D_{\text{HI}}^{\text{FoG}}(\alpha_{\text{FoG,HI}} k \mu; z_1) D_g^{\text{FoG}}(\alpha_{\text{FoG,g}} k \mu; z_2) \\ & + M_{10} [C_{\text{HI}}(z_1) 10^{10} M_{\odot}], \end{aligned} \quad (91)$$

where we have highlighted the individual parameters in bold. Note that we evaluate the matter power spectrum at a fiducial redshift z_{fid} and apply the linear growth factor $D^+(z)$ to scale it to other redshifts.

We also require one more parameter to describe an apparent frequency or redshift offset between the HI and galaxies. This will be needed to account for systematic redshift errors in the eBOSS catalogs (see Section 8.1). As it is an observational effect, we do not include it in the cross-power spectrum description (where it would manifest itself as a phase rotation).

$\Delta\nu_0$: This parameter shifts the stack signal away from being centered at zero frequency lag. Positive values of $\Delta\nu_0$ move the peak of the signal to higher frequencies and thus to lower redshifts.

5.3. Simulations

We make extensive use of simulations in this work, for interpreting our stacking measurements in terms of physical

models, determining the signal transfer function, and quantifying the linearity of our analysis pipeline via injection of simulated signals into the data. In this section, we describe our simulation methodology for generating sky maps of 21 cm emission and galaxy density (Section 5.3.1), propagating these through to mock galaxy catalogs (Section 5.3.2) and CHIME time streams (Section 5.3.3), and finally performing the stacking procedure (Section 5.3.4). The associated steps are schematically shown in Figure 16. Note that we do not attempt to simulate foregrounds, instead relying on several data-based tests to assess the contribution of residual foregrounds to our sky maps and cross-correlation measurements.

5.3.1. Map Generation

Each simulation produces a pair of correlated δ_g and δ_{HI} maps of the sky generated as follows. The input real-space matter power spectrum (Section 5.2.1), evaluated at $z = 1$, is transformed to a 3D correlation function using the `hankl` Python package (Karamanis & Beutler 2021) via the `FFTlog` method. We additionally employ Richardson extrapolation to repeated computations with increasingly fine k sampling in order to reduce numerical errors. We then transform this to a multifrequency angular power spectrum $C_{\ell}(\nu, \nu')$ and perform further frequency integrals over top hats with width 0.390625 MHz in order to mimic the effect of CHIME’s frequency channelization.

We form a set of N_{ν} HEALPix maps (Gorski et al. 2005) from a Gaussian realization of this angular (matter) power spectrum, use the linear growth factor for our fiducial cosmology to scale each map to the redshift corresponding to its frequency, and multiply by the bias $b_{\alpha}(z)$ (Section 5.2.2). In tandem, we generate the same number of maps of the gravitational potential ϕ , to which we apply a finite-difference second derivative in the radial direction and appropriate prefactors to generate a velocity field that is added to the biased matter to include linear redshift-space distortions. These maps are then convolved with a frequency kernel designed to reproduce the desired form of Finger-of-God damping in Fourier space (Section 5.2.3).

Finally, the maps corresponding to δ_{HI} are multiplied by the mean 21 cm brightness temperature $\bar{T}_b(z)$ (Section 5.2.4), while a lognormal transform is applied to the δ_g maps, to ensure that $\delta_g \geq -1$ everywhere; this allows $1 + \delta_g$ to be used to construct a probability density function from which to draw mock catalogs (see Section 5.3.2). Note that we do not apply a lognormal transform to the T_b maps: when Gaussian temperature maps are stacked on mock catalogs generated from lognormal δ_g maps, the two-point statistics are equivalent to the case where both sets of maps are Gaussian (see Appendix C for details).

Our baseline simulations set the shot-noise contribution ϵ_X to zero, but we require the ability to add shot noise to ascertain its impact on the stacking signal. We incorporate this into our simulations by adding correlated realizations of white noise to each pair of δ_{HI} and δ_g maps, such that their cross-power spectrum will contain the contribution from Equation (82) (the auto-power spectra of these maps are never used). Specifically, for each map voxel, we draw a random number from a Gaussian with $\sigma = [C_{\text{HI}}(z) \langle M_{\text{HI}} \rangle_g / (\bar{T}_b(z) V_{\text{vox}})]^{1/2}$, where V_{vox} is the voxel volume, and add this value to the same voxel in the δ_{HI} and δ_g maps.²⁸

²⁸ This method of adding correlated shot noise adds unphysical contributions to the auto-power of the δ_{HI} and δ_g maps, which will also affect the variance of the cross-power between them, but this effect is completely negligible for our purposes.

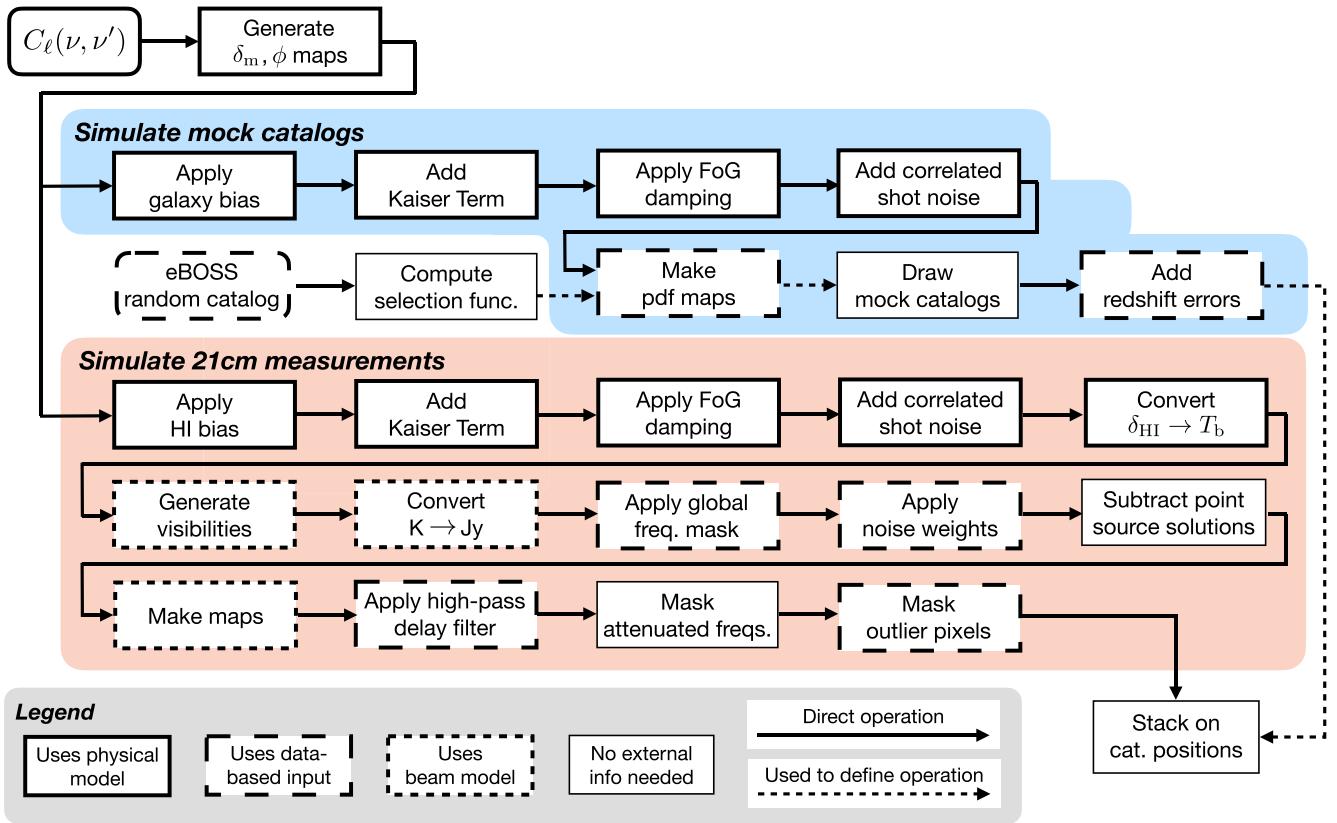


Figure 16. A schematic representation of the simulation pipeline. Starting from the multifrequency angular power spectrum $C_\ell(\nu, \nu')$ corresponding to an input matter power spectrum, we generate correlated full-sky maps of the matter overdensity δ_m and gravitational potential ϕ at redshifts corresponding to each CHIME frequency channel and transform these into maps of LRG/ELG/QSO overdensity and 21 cm brightness temperature using the models described in Section 5.2 and the procedures in Section 5.3.1. Mock LRG/ELG/QSO catalogs are constructed from the corresponding maps (Section 5.3.2), while mock CHIME observations are formed from the 21 cm maps (Section 5.3.3), and these observations are then processed with the same stacking pipeline as the data (Section 5.3.4). In this diagram, boxes with long dashed outlines are defined using inputs from eBOSS or CHIME observations, rather than simulations (e.g., the delay cuts in the delay filter are those from Figure 10).

5.3.2. Mock Catalogs

For each galaxy sample we consider, we create mock catalogs of N objects for each pair of simulated δ_{HI} and δ_g maps. To do so, we select the pixel indices and frequency channels from a probability density function given by

$$\mathcal{P}(\mathbf{x}) \propto S(\mathbf{x})[1 + \delta_g(\mathbf{x})], \quad (92)$$

where $S(\mathbf{x})$ is a sample-specific selection function. Once a voxel is selected, galaxies are assigned positions within it according to uniform random distributions and further displaced by simulated redshift errors as described below.

We obtain approximate galaxy selection functions from the public random catalogs associated with each eBOSS sample. In detail, for each sample, we build a histogram of object positions with 32 redshift bins from $0.8 < z < 2.5$ and a HEALPix angular pixelization with $N_{\text{side}} = 16$ (roughly 3.7° resolution). We then form a rank-7 approximation to this distribution by performing a singular value decomposition of the histogram (represented as a $N_z \times N_{\text{pix}}$ matrix). Finally, we upsample this to the HEALPix resolution of the input maps and apply Gaussian smoothing in the angular direction (with width equal to the original pixel size) to apodize any sharp boundaries. Using this as the selection function for generating mocks ensures that we reproduce the large-scale footprint and modulations of each galaxy sample without introducing smaller-scale features of the catalogs into our simulations.

We generate random redshift errors using a separate scheme for each sample, based on estimates of redshift error distributions (represented as line-of-sight velocities) published by the eBOSS team. For LRGs, Ross et al. (2020) examined pairs of observations of the same target and found that the distribution of redshift differences was well fit by a Gaussian with $\sigma = 91.8 \text{ km s}^{-1}$, corresponding to a redshift uncertainty of $\sigma = 65.6 \text{ km s}^{-1}$ per object. For ELGs, Raichoor et al. (2021) quote three redshift error percentiles based on repeated observations; we find that these values are well fit by a Tukey lambda distribution with $\lambda = -0.4$ and $\sigma = 11.88 \text{ km s}^{-1}$.

For QSOs, Lyke et al. (2020) find that, over the entire QSO catalog, the distribution of redshift differences between repeated observations is well fit by a double Gaussian. This implies that the single-observation redshift errors are also described by a double Gaussian, with $\sigma_1 = 150 \text{ km s}^{-1}$, $\sigma_2 = 1000 \text{ km s}^{-1}$, and 18% of objects having errors drawn from the wider Gaussian.²⁹ Though we use this model for our primary analysis, there is evidence that it does not completely capture the distribution of QSO redshift errors. We discuss the discrepancies and the effect on our analysis in Section 8.1.

²⁹ These double-Gaussian parameters are quoted in Lyke et al. (2020) as corresponding to the distribution of *redshift differences between repeated observations* shown in their Figure 4, but in our own comparison we found that the quoted widths of the two Gaussians correspond to the distribution of *single-object redshift errors* implied by this figure.

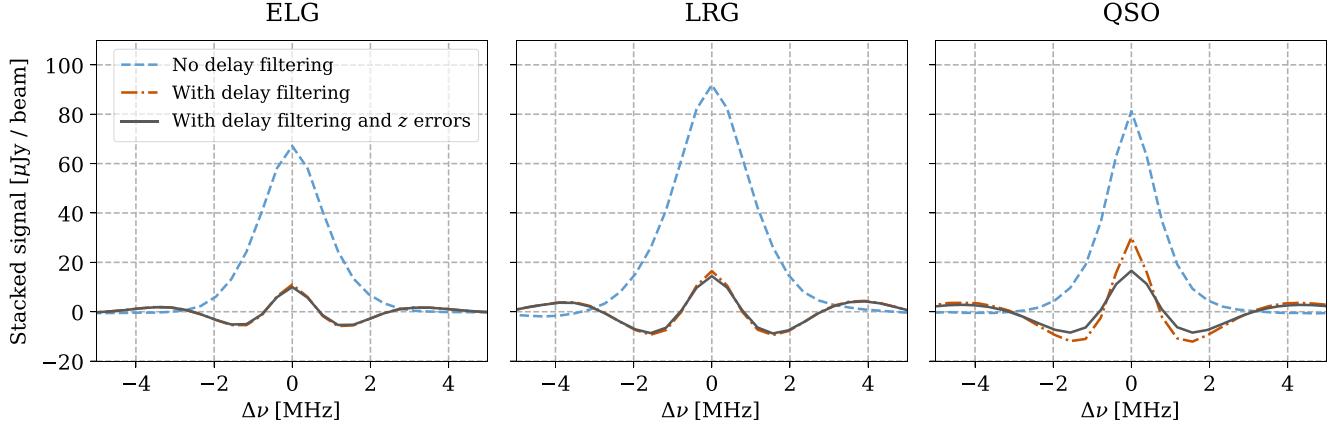


Figure 17. Results of stacking simulated observations containing only 21 cm signal on mock ELG, LRG, and QSO catalogs correlated with the input signal, generated according to the procedure in Section 5.3. The three panels correspond to simulations that use the selection functions and redshift error distributions of the three eBOSS samples we consider. The stacking amplitude in the absence of delay filtering (blue dashed lines) is heavily suppressed by the delay filter (red dotted-dashed lines) and further suppressed by the inclusion of random redshift errors in the catalogs (black solid lines).

We do not attempt to simulate catastrophic redshift errors, which the above references estimate to occur in less than 1% of the LRG and ELG samples and as much as 2% of the QSO sample. The effect of these errors on our stacking measurements is a simple suppression of the overall amplitude, by an amount equal to the catastrophic error fraction.

5.3.3. Time Streams

We make use of the m -mode formalism (S15) to translate simulated 21 cm maps into visibilities. In this formalism, the spherical harmonic coefficients $a_{\ell m}^P(\nu)$ of sky maps for Stokes parameter $P \in \{I, Q, U, V\}$ are related to the sidereal-time Fourier transform of the visibility time stream, $\tilde{V}_{xy,m}^P(\nu)$, via multiplication by a beam transfer matrix $B_{xy;\ell m}^{P,P}(\nu)$:

$$\tilde{V}_{xy,m}^P(\nu) = \sum_P \sum_\ell B_{xy;\ell m}^{P,P}(\nu) a_{\ell m}^P(\nu). \quad (93)$$

After performing this multiplication, we convert the result to a visibility time stream by inverse Fourier transforming in m , applying zero-padding to the Fourier transform such that the corresponding time resolution matches that of the observed sidereal stacks.

We carry out separate versions of this procedure with beam transfer matrices corresponding to the default or control beam models from Section 4.4. We compute these matrices using `driftscan` (Shaw et al. 2020a), with several performance optimizations: precision truncation using the `bitshuffle` library (Masui et al. 2015), omitting frequencies that fall outside of the mask described in Section 4.2, and only computing the $P=I$ components (since the 21 cm signal is unpolarized).

Up to this point, the simulated data are in temperature units. To transform into spectral flux density units, we first compute the beam solid angle for the assumed beam model:

$$\Omega^P(\nu) = \int d\theta \cos \theta \int d\phi |A^P(\nu, \theta, \phi)|^2. \quad (94)$$

We then multiply the visibilities by the standard Rayleigh-Jeans conversion factor and the beam solid angle, normalized by the power beam evaluated at $\text{HA} = 0'$ and a reference decl.

θ_{ref} :

$$V_{xy}^P(\nu, \phi) |_{\text{Jy}} = \frac{2 \times 10^{26} k_B \nu^2}{c^2} \frac{\Omega^P(\nu)}{|A^P(\nu, \theta_{\text{ref}}, 0)|^2} V_{xy}^P(\nu, \phi) |_{\text{K}}. \quad (95)$$

With this normalization, a visibility corresponding to a point source that transits at $\theta = \theta_{\text{ref}}$ has an amplitude equal to the flux of the source. For consistency with CHIME’s beam and complex gain calibration, we set θ_{ref} to the decl. of Cygnus A.

From here, the simulated visibilities are processed in the same way as the real data: a global frequency mask and noise weights described in Sections 4.2 and 3.1.4 are applied, the contributions of the four brightest point sources are inferred and subtracted, beam-deconvolved maps are constructed as in Section 4.3, delay filtering is applied with the decl.-dependent delay cuts from Section 4.5, and the masking operations in Section 4.6 are applied. Just as we simulate visibilities for each of the default and control beam models, we also perform two versions of the mapmaking step, assuming either beam model; thus, we obtain four simulated data sets corresponding to each pair of assumed and deconvolved beam, and we compare the results in Section 7.3 in order to estimate the systematic uncertainty arising from our choice of beam model.

5.3.4. Mock Source Stacking

Finally, we stack the simulated observations on the associated mock catalogs, following the procedure in Section 4.7. Figure 17 shows stacking results corresponding to simulations of each eBOSS sample, for a single LSS realization but averaged over 100 mock catalogs of 400,000 objects each, in order to suppress shot noise associated with the catalog size. In the absence of delay filtering, the stacking amplitude inferred from these simulations for QSOs is greater than for ELGs and less than for LRGs; the former follows from ELGs having lower bias and higher Finger-of-God suppression than QSOs, while the latter is due to the higher bias of LRGs than QSOs, which wins over the more severe Finger-of-God effect for LRGs (see Figure 15).

The delay filter significantly suppresses the signal level, reducing the zero-lag amplitude by around 80% for ELGs and LRGs and 63% for QSOs. We attribute the lower suppression for QSOs to their milder Finger-of-God suppression at small

scales: the delay filter removes sensitivity to the largest scales (see Section 5.1), and the remaining smaller-scale contribution is larger for QSOs than for the other tracers owing to a smaller amount of suppression. Finally, redshift errors in the simulated catalogs reduce the zero-lag amplitude by no more than 10% for ELGs and LRGs but by 40% for QSOs, thanks to the much wider distribution of QSO redshift errors discussed above.

5.4. Template Calculation

To interpret our results, we need to be able to calculate the expected signal from stacking on a given catalog for an underlying set of parameters θ . We call this quantity the *template*, denoted by $s(\Delta\nu; \theta)$. Though the template is entirely determined by the cross-power spectrum in Equation (91), propagating this through the instrumental transfer function and our analysis procedure is challenging to do both efficiently and accurately, and so it will be left to a follow-up paper (CHIME Collaboration 2023, in preparation).

In this work, we instead use our simulation capability to calculate the templates. In brief, we generate LSS realizations corresponding to several modes, each of which is defined by a specific combination of model parameters; average over a sufficiently large number of random mock catalogs to estimate the stack signal for each mode; and calculate the full template for arbitrary parameter values by making linear combinations of the template modes and applying an effective treatment for the Finger of God. Overall the errors in this approach are $\lesssim 1\%$. We describe this approach in detail in Appendix D.

In Figure 18, we display the change in the H I-tracer cross-power spectrum (left panels) corresponding to variations of each of our eight model parameters, along with the corresponding change in the predicted stack signal (right panels). Access to the full k range shown in the left panels would allow nondegenerate constraints on several of these parameters, due to their different impacts on the cross-power spectrum. However, our filtering choices imply that the stack signal is only sensitive to nonlinear scales ($k \gtrsim 0.3 h^{-1} \text{Mpc}$ or so, as shown in Figure 14), and as a result, we are left with significant parameter degeneracies, which can be inferred from the similar variations in the right panels of Figure 18.

6. Results

6.1. Stacking Measurements

The top left panel of Figure 19 shows the result of stacking the deconvolved, foreground-filtered maps on the 3D positions in the eBOSS NGC QSO catalog. It is shown as a function of R.A. offset and decl. offset at 0 MHz frequency offset, averaged over the two polarizations, in other words, $d(0, \Delta\theta, \Delta\phi)$ in the notation of Section 4.7. Also shown in the top row is our best-fit model for the 21 cm emission based on the simulations described in the preceding section and the residuals obtained by subtracting the best-fit model from the data. The residuals can be compared to the three panels in the second row, which correspond to three different techniques for estimating the noise present in the stack. The left panel is the result of applying the stacking procedure to a Gaussian noise realization generated according to Equation (52). The middle panel is the result of applying the stacking procedure to a jackknife of even and odd days (see Section 7.2). Finally, the right panel is the result of stacking the data on a random mock catalog.

The noise in the residuals is consistent with that observed in the random mock catalog. Both are in excess of the noise in the even–odd jackknife, owing to the fact that residual foregrounds are highly correlated between even and odd days and therefore cancel in the jackknife. The noise observed in the even–odd jackknife is in excess of that observed in the Gaussian noise realization owing to unflagged RFI and variations in the foregrounds from day to day caused by instrument instability.

The third row of Figure 19 shows 1D slices of both data and the best-fit model. The negative shoulders in the R.A. direction that are observed in both the data and model are caused by the exclusion of intracylinder baselines from our analysis. The grating lobes in the R.A. direction, which are shown in Figure 7, have largely averaged away in the stack because their location varies with frequency and decl. It is important to note that the angular information displayed in Figure 19 was not used to constrain the model. For simplicity, the model is only fit to the central pixel of the stack as a function of frequency. A full 3D fit could further improve the signal-to-noise ratio and help break the degeneracy between the amplitude A_{HI} and the Finger-of-God damping, but we leave that for a future analysis.

For all three tracers, the spatial extent of the signal is consistent with the synthesized beam computed directly from Equations (47) to (31) and averaged over sources, indicating that the 21 cm signal is unresolved. Figure 20 shows the central pixel of the stack as a function of frequency, i.e., $d(\Delta\nu, 0, 0)$, for the three tracers in black. The-dark gray and light-gray contours indicate the central 68% and 95% of values, respectively, observed when stacking the maps on 10,000 random mock catalogs as outlined in Section 4.8. The red line indicates our best-fit model for the signal. Note that although the two polarizations are fit jointly, to simplify the figure we show only their weighted average, with the weights set to the inverse variance as measured by the random mock catalogs. Also note that the polarization- and frequency-dependent mean value of the noise has been characterized using the random mock catalogs and subtracted from both the stack on the true catalog and the stack on the mock catalogs that are shown in the figure.

The best-fit model shown in both Figures 19 and 20 consists of fixing all nonlinear parameters at their fiducial values and allowing the parameters governing the large-scale clustering of H I to vary. This model has been described in Section 5.2.6. The bottom row shows the result of subtracting the best-fit model from the data and compares to the same gray mock catalog contours shown in the top row. For all tracers, the residuals are consistent with our noise model based on the random mock catalogs. This is also true for all QSO redshift bins, which are not shown.

6.2. Model Fitting

We assume that the noise in the stacked source data is described by a Gaussian and that the signal is described by the model given in Section 5.4. This means that the likelihood function $\mathcal{L}(\mathbf{d}|\theta)$ of observing the stacked signal \mathbf{d} given a template $s(\theta)$ with model parameters $\theta = [\Delta\nu_0, \Omega_{\text{HI}}, b_{\text{HI}}, b_g, M_{10}, \alpha_{\text{NL}}, \alpha_{\text{FoG,HI}}, \alpha_{\text{FoG,g}}]$ is described by a multivariate Gaussian

$$\mathcal{L}(\theta) = \mathcal{P}(\mathbf{d}|\theta) \quad (96)$$

$$= \frac{1}{|2\pi\Sigma|^{1/2}} \exp\left(-\frac{1}{2}\chi^2\right), \quad (97)$$

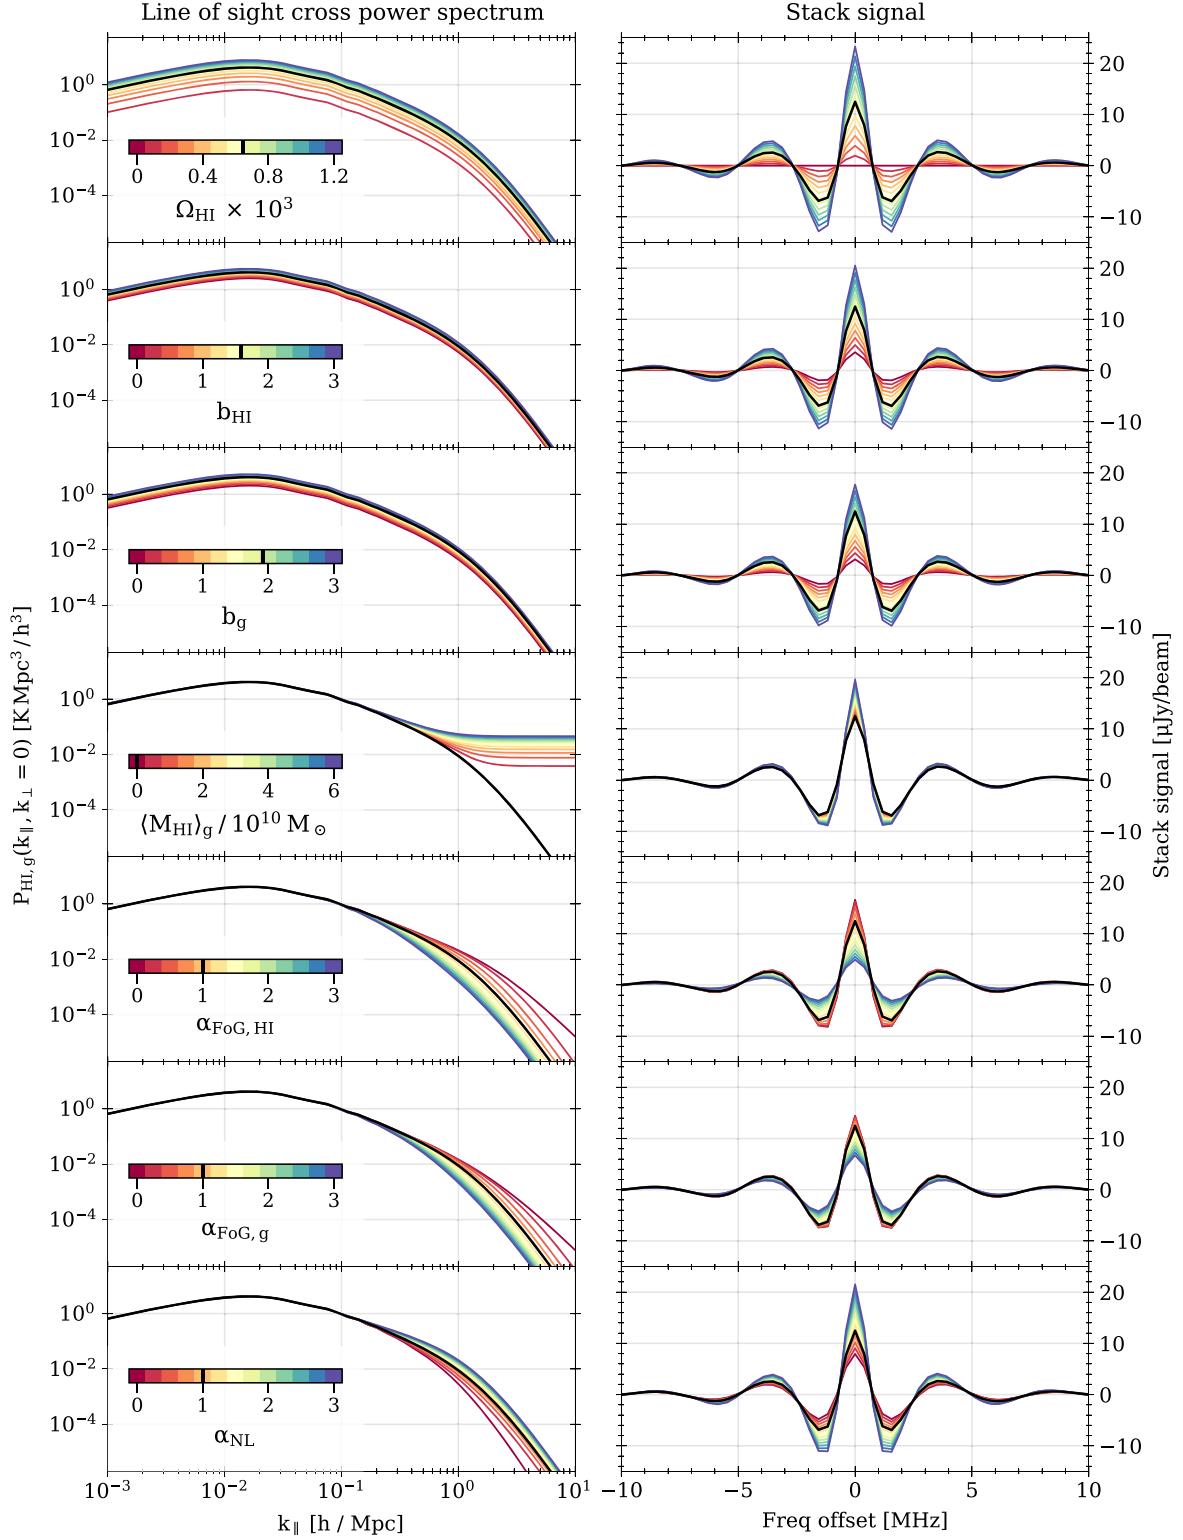


Figure 18. The simulated stack signal after processing through the CHIME pipeline. Each row shows the effect of varying a parameter on the theoretical 21 cm QSO cross-power spectrum in the left panel and the expected signal observed by CHIME in the right panel. The variation for each parameter is chosen to be over a range consistent with our prior uncertainties. The fiducial model used within our modeling is indicated by the thick black line for each panel, and the location within range of each parameter is indicated by the black line within the color bar.

with

$$\chi^2 = (\mathbf{d} - s(\theta) - \mu)^T \Sigma^{-1} (\mathbf{d} - s(\theta) - \mu), \quad (98)$$

where $s(\theta)$ is the model for the 21 cm signal and μ and Σ^{-1} are the mean and inverse covariance of the noise, respectively, which are estimated using the sample mean and covariance of the mock catalogs as outlined in Section 4.8.

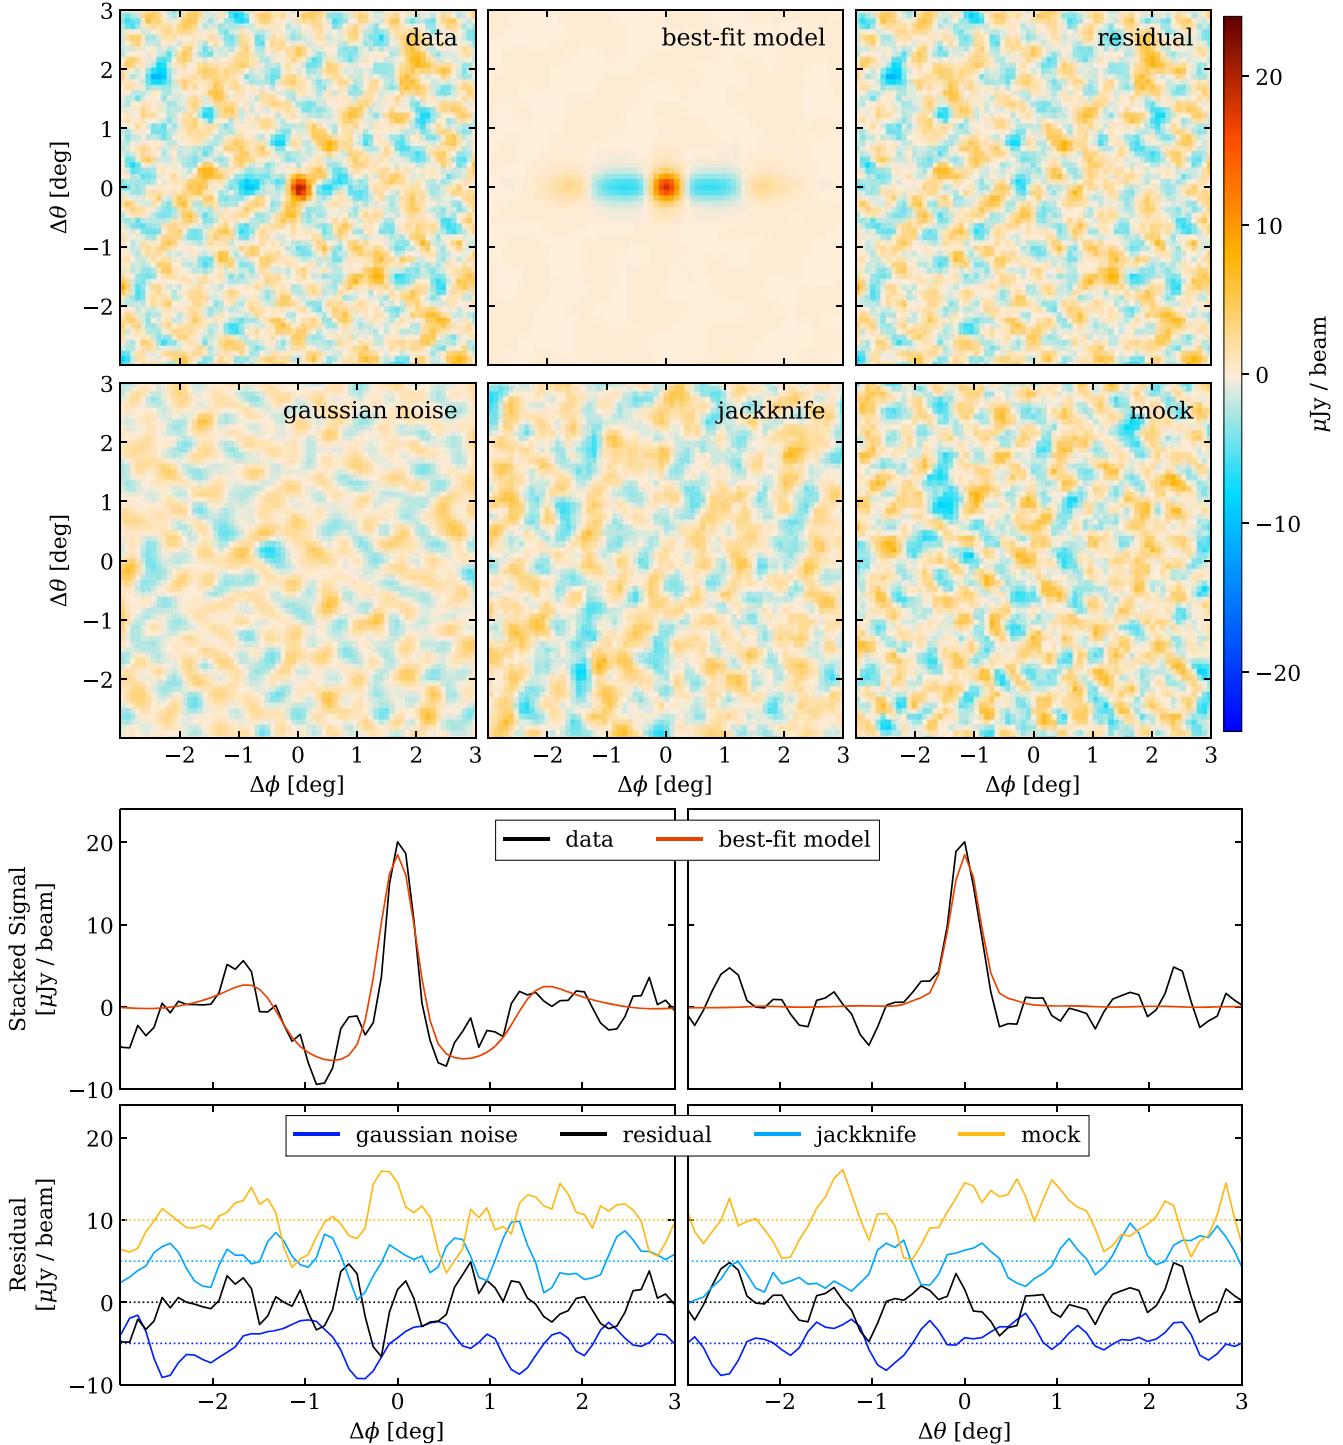


Figure 19. The stacked signal at $\Delta\nu = 0$ MHz as a function of R.A. offset ($\Delta\phi$) and decl. offset ($\Delta\theta$) for the QSO catalog. The top row shows, from left to right, the data, best-fit model, and residual. The second row shows, from left to right, the result of stacking the QSO catalog on a Gaussian noise realization, stacking the QSO catalog on a jackknife of even and odd days, and stacking a random mock catalog on the data. The third row shows a slice of the data in black and the best-fit model in red at $\Delta\theta = 0^\circ$ on the left and $\Delta\phi = 0^\circ$ on the right. The bottom row shows, for these same slices, the residuals in black compared to the Gaussian noise realization in dark blue, the jackknife in light blue, and the random mock catalog in orange. Note that, to facilitate the comparison, the slices in the bottom row have been offset by an amount indicated by the dotted line of the same color.

We employ a Markov Chain Monte Carlo (MCMC) to sample from the joint posterior distribution,

$$\mathcal{P}(\boldsymbol{\theta}|\mathbf{d}) = \frac{1}{\mathcal{Z}} \mathcal{L}(\boldsymbol{\theta}) \pi(\boldsymbol{\theta}), \quad (99)$$

where $\pi(\boldsymbol{\theta})$ is the prior probability distribution over the model parameters and \mathcal{Z} is the normalization constant such that the posterior integrates to unity.

We use noninformative priors for most parameters, ascribing equal prior probability over large ranges. For the nonlinear

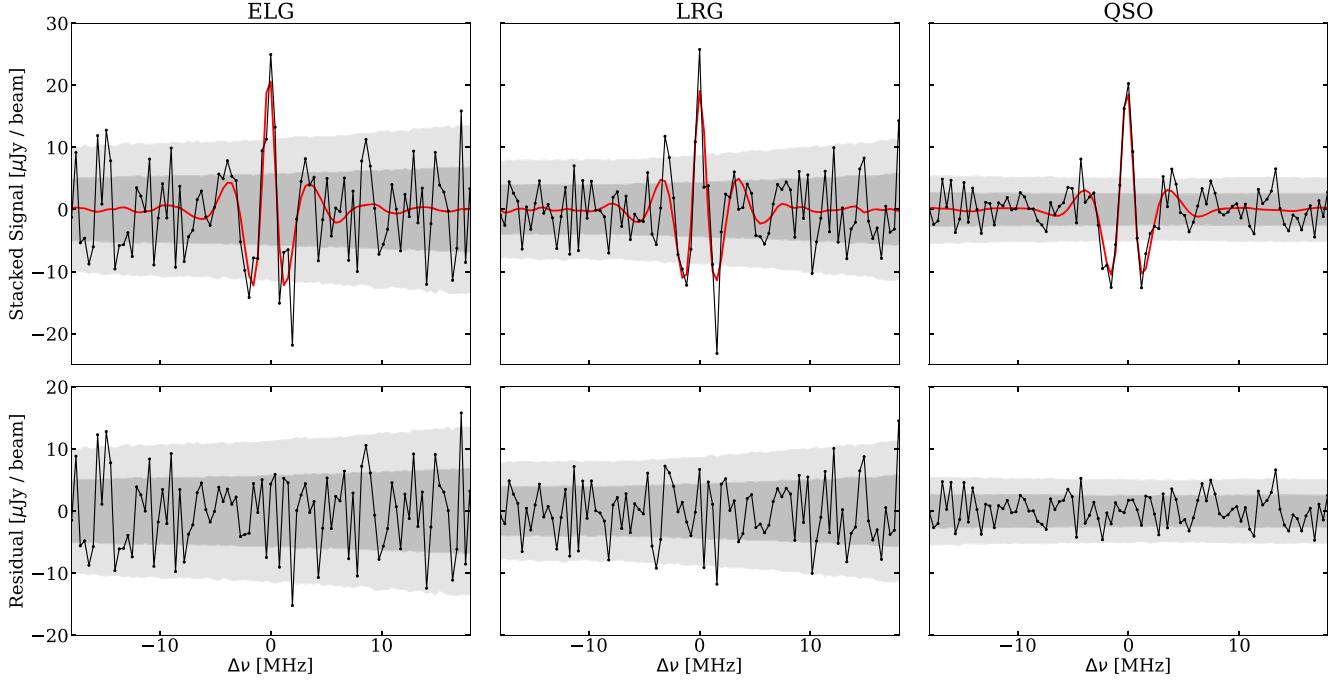


Figure 20. Top: the stacked signal as a function of frequency offset for the ELG, LRG, and QSO catalogs. The data are shown in black, and the best-fit model is shown in red. Bottom: the residuals obtained by subtracting the best-fit model from the data. For both the top and bottom rows, the dark-gray and light-gray bands indicate the central 68% and 95% of values, respectively observed when applying the same stacking procedure to 10,000 mock catalogs.

parameters we choose to do this even where there is some external information either from simulations or more strongly from analysis of the eBOSS data themselves (e.g., on the Finger-of-God scale; see Section 5.2.3), as it is difficult to combine the different prescriptions for modeling the nonlinear scales. These analyses guide our choice of fiducial model, but we allow a wide range of variation around them when trying to fit the data.

The one exception to this is for the galactic bias b_g . As it is a large-scale parameter, it is less susceptible to systematic differences in the modeling, and we instead use a prior informed by modeling of the eBOSS tracers. For the QSOs, our fiducial model is that from Laurent et al. (2017), and to get an uncertainty on this, we fit a shift in the amplitude to the two lowest-redshift bins in their analysis (which overlap with that of this paper), which gives an uncertainty of 3% about the fiducial model. For the LRGs we translate the overall results of Zhai et al. (2017) of $b = 2.30 \pm 0.03$ into a 1.3% uncertainty on the amplitude of the bias model used here. Finally, for the ELGs we symmetrize the measurements of b_1 from Tamone et al. (2020) to give an uncertainty of 10% for the ELG linear bias.

For $\Omega_{\text{H}\,1}$, which gives an overall normalization to the signal, we use a prior symmetric about zero, despite the fact that physically $\Omega_{\text{H}\,1} \geq 0$. This is to ensure that our priors do not give an artificial bias toward positive signal and give a more robust estimation of the detection significance. However, we do enforce that $b_{\text{H}\,1} \geq 0$ to exclude an unphysical mode of high probability with both $\Omega_{\text{H}\,1} < 0$ and $b_{\text{H}\,1} < 0$.

We summarize our choice of priors in Table 2.

The affine-invariant ensemble sampler from the `emcee` package (Foreman-Mackey et al. 2013) is used to sample from the joint posterior distribution. We run 32 samplers initialized from random locations within the region defined by Table 2. To analyze the chains, we need to reduce them down to an

Table 2
The Prior Placed on Each Parameter during Our Analysis

Parameter	Type	Description
<i>Standard parameters</i>		
$\Omega_{\text{H}\,1}$	Uniform	Range: -10^{-2} to 10^{-2}
$b_{\text{H}\,1}$	Uniform	Range: 0 to 10
b_g	Gaussian	Mean: $\bar{b}_g = b_g^{\text{fid}}(z_{\text{eff}})$ standard deviation: QSOs 3%, LRGs 1.3%, ELGs 10%
$\Delta\nu_0$	Uniform	Range: -0.8 MHz to 0.8 MHz
<i>Nonlinear parameters</i>		
M_{10}	Uniform	Range: 0 to 20; Fixed: 0
α_{NL}	Uniform	Range: 0 to 5; Fixed: 1
$\alpha_{\text{FoG},\text{H}\,1}$	Uniform	Range: 0 to 5; Fixed: 1
$\alpha_{\text{FoG},g}$	Uniform	Range: 0 to 5; Fixed: 1

Note. There are two classes of parameters in our analysis, *standard* parameters that capture the large-scale quantities we hope to constrain, and *nuisance* parameters that model the signal on small, *nonlinear* scales. In our analysis the latter group of parameters will either be marginalized over or be fixed to their fiducial values in order to assess the contribution of modeling uncertainties to our constraints.

independent set of samples, but we must first determine how correlated the samples actually are. To do this, we first define the sample autocorrelation $\hat{\rho}_i^s(\tau)$ of a parameter θ_i within the s th chain as

$$\hat{\rho}_i^s[\tau] = \frac{1}{N - \tau} \sum_{t=1}^{N-\tau} (\theta_i^s[t] - \bar{\theta}_i^s)(\theta_i^s[t + \tau] - \bar{\theta}_i^s), \quad (100)$$

where $\bar{\theta}_i^s$ is the sample mean for parameter θ_i , and then using this, we define the autocorrelation length for a single parameter and chain as

$$\zeta_i^s = \frac{1}{\hat{\rho}_i^s[0]} \sum_{\tau} \hat{\rho}_i^s[\tau]. \quad (101)$$

For a single aggregate autocorrelation length ζ summarizing the chain convergence we take the mean over the different samplers and then use the longest autocorrelation across the different model parameters. That is,

$$\zeta = \max_i (\text{mean}_s \zeta_i^s). \quad (102)$$

This gives the number of MCMC steps over which the chains remain correlated. Only samples from each chain separated by ζ samples can be considered independent. The first $10 \times \zeta$ samples in each chain are discarded as burn-in. The chains are then thinned by ζ and concatenated. The parameter space is high dimensional and has complex degeneracies, which means that the correlation lengths are large, $\zeta \gtrsim 500$ in the full parameter space. We also make extensive use of the GetDist package (Lewis 2019) for analyzing the MCMC chains.

6.3. Parameter Constraints

In Figure 21 we show the constraint on the default model parameters for the QSO catalog. We show constraints for both a model where all parameters are allowed to vary and a model where the nonlinear parameters are fixed to their fiducial values ($M_{10} = 0$, $\alpha_{NL} = 1$, $\alpha_{FoG,H1} = 1$, and $\alpha_{FoG,g} = 1$). The constraints show that certain parameter combinations are highly degenerate, most notably $\Omega_{H1} - b_{H1}$, but also correlations with the nonlinear parameters $\alpha_{FoG,H1}$ and $\alpha_{FoG,g}$. As these degeneracies limit our ability to make a cosmological interpretation of our results, it is worth attempting to understand them.

The most severe degeneracy in our model is between Ω_{H1} and b_{H1} , and it is clearly apparent in both the full and fixed models. The origin of this can be seen in Equation (91), which, simplified slightly down to the linear terms, has

$$P_{H1,g}(k, \mu) \propto \Omega_{H1}(b_{H1} + f\mu^2)(b_g + f\mu^2)P(k), \quad (103)$$

which contains a multiplicative $\Omega_{H1}b_{H1}$ term, responsible for the curved degeneracy seen in the $\Omega_{H1} - b_{H1}$ panel of Figure 21. Previous 21 cm cross-correlation analyses (Masui et al. 2013; Switzer et al. 2013; Wolz et al. 2022) gave constraints directly on the combination $\Omega_{H1}b_{H1}r$, where r is a scale-independent cross-correlation parameter that absorbs modeling uncertainties on nonlinear scales; however, this is not sufficient for the analysis here. Although transforming our constraints to be in terms of $\Omega_{H1}b_{H1}$ removes the curved degeneracy,³⁰ we find that a linear degeneracy against Ω_{H1} remains. This can be understood straightforwardly as the effect of the Kaiser redshift-space distortions. As CHIME has higher resolution in the frequency direction versus the angular direction, and we have removed low- k_{\parallel} modes by foreground filtering, the sensitivity in this analysis is biased toward wavenumbers with higher μ (which is illustrated in Figure 14). As both b_{H1} and f are of order unity, the contribution of the Kaiser term is important and cannot be neglected.

³⁰ In fact, we actually sample within a transformed basis by replacing the parameter Ω_{H1} with $\Omega_{H1}b_{H1}$. This substantially improves convergence, as the remaining linear degeneracy is easily navigated by the affine-invariant sampler, where the original curved degeneracy was not. To do this, we need to carefully adjust the prior applied in the sampler to ensure that the prior on Ω_{H1} remains uniform.

To account for this, we transform to a plane of $(\Omega_{H1}b_{H1}) - \Omega_{H1}$ and determine a linear combination of these parameters that minimizes their variance. For a single LRG, ELG, or QSO sample g , the solution for an exactly linear degeneracy can be found by using the MCMC samples to construct the covariance matrix between $\Omega_{H1}b_{H1}$ and Ω_{H1} , which we write as $\mathbf{C}_{\Omega b, g}$, and then finding the eigenvector with minimal eigenvalue, which gives the linear combination we are searching for. We will use this combination as our primary amplitude parameter

$$A_{H1} \equiv 10^3 \Omega_{H1}(b_{H1} + \langle f\mu^2 \rangle), \quad (104)$$

where we make the interpretation that the coefficient $\langle f\mu^2 \rangle$ is the sensitivity-weighted average $f\mu^2$ that this CHIME analysis is probing. We perform this optimization on the chains with fixed nonlinear parameters, as this gives a cleaner separation from other degenerate parameters.

The $\langle f\mu^2 \rangle$ coefficient preferred by each tracer differs slightly from ~ 0.45 (QSOb00) to ~ 0.62 (QSOb2), which we would expect, as both f and CHIME's sensitivity change with redshift. As we would like to be able to compare our measurements between tracers, we would instead like a single effective $\langle f\mu^2 \rangle$. To do this, we minimize the covariance $\mathbf{C}_{\Omega b, \text{all}}$, defined by

$$\mathbf{C}_{\Omega b, \text{all}}^{-1} = \sum_g \mathbf{C}_{\Omega b, g}^{-1}, \quad (105)$$

where we sum over the tracers QSOb0, QSOb1, QSOb2, LRG, and ELG (we exclude the other QSO tracers to avoid double-counting the data). The form of $\mathbf{C}_{\Omega b, \text{all}}$ is motivated by considering each tracer to be a different measurement in the $(\Omega_{H1}b_{H1}) - \Omega_{H1}$ plane: if each distribution were Gaussian and all were consistent, the covariance on the combined distribution would be given by $\mathbf{C}_{\Omega b, \text{all}}$. After this procedure, we derive an effective $\langle f\mu^2 \rangle \approx 0.552$, which we fix for the rest of this analysis. The overall loss of constraining power from fixing a single value is small, with a drop of $\sim 7\%$ for the worst affected tracer (full QSO catalog).

The second degeneracy we focus on is between the Finger-of-God parameters. If we examine the cross-power spectrum given by Equation (91) and expand the Finger-of-God damping factors defined in Equation (70) assuming $k_{\parallel}\sigma_P \gg 1$, we find that

$$P_{H1,g}(k, \mu) \propto D_{H1}^{\text{FoG}}(\alpha_{FoG,H1}\mu k)D_g^{\text{FoG}}(\alpha_{FoG,g}\mu k) \quad (106)$$

$$\approx \frac{4}{k_{\parallel}^4 \sigma_{P,H1}^2 \sigma_{P,g}^2} (\alpha_{FoG,H1} \alpha_{FoG,g})^{-2}. \quad (107)$$

For most of the region of CHIME's k -space sensitivity (see Figure 14) we are close to this regime, and so we expect there to be an approximate degeneracy of the form $\alpha_{FoG,H1}\alpha_{FoG,g}$, which can be seen in the $\alpha_{FoG,H1} - \alpha_{FoG,g}$ panel of Figure 21. This motivates us to transform to two new parameters

$$\alpha_{FoG,+} = (\alpha_{FoG,H1} \alpha_{FoG,g})^{1/2} \quad (108)$$

$$\alpha_{FoG,-} = \log(\alpha_{FoG,H1}/\alpha_{FoG,g}), \quad (109)$$

where in the large k_{\parallel} limit $\alpha_{FoG,+}$ controls the amount of Finger-of-God damping and $\alpha_{FoG,-}$ does not affect the cross-power spectrum. The logarithm in the definition of $\alpha_{FoG,-}$ is to limit the effect of small $\alpha_{FoG,g}$ values generating extremely large values for this parameter.

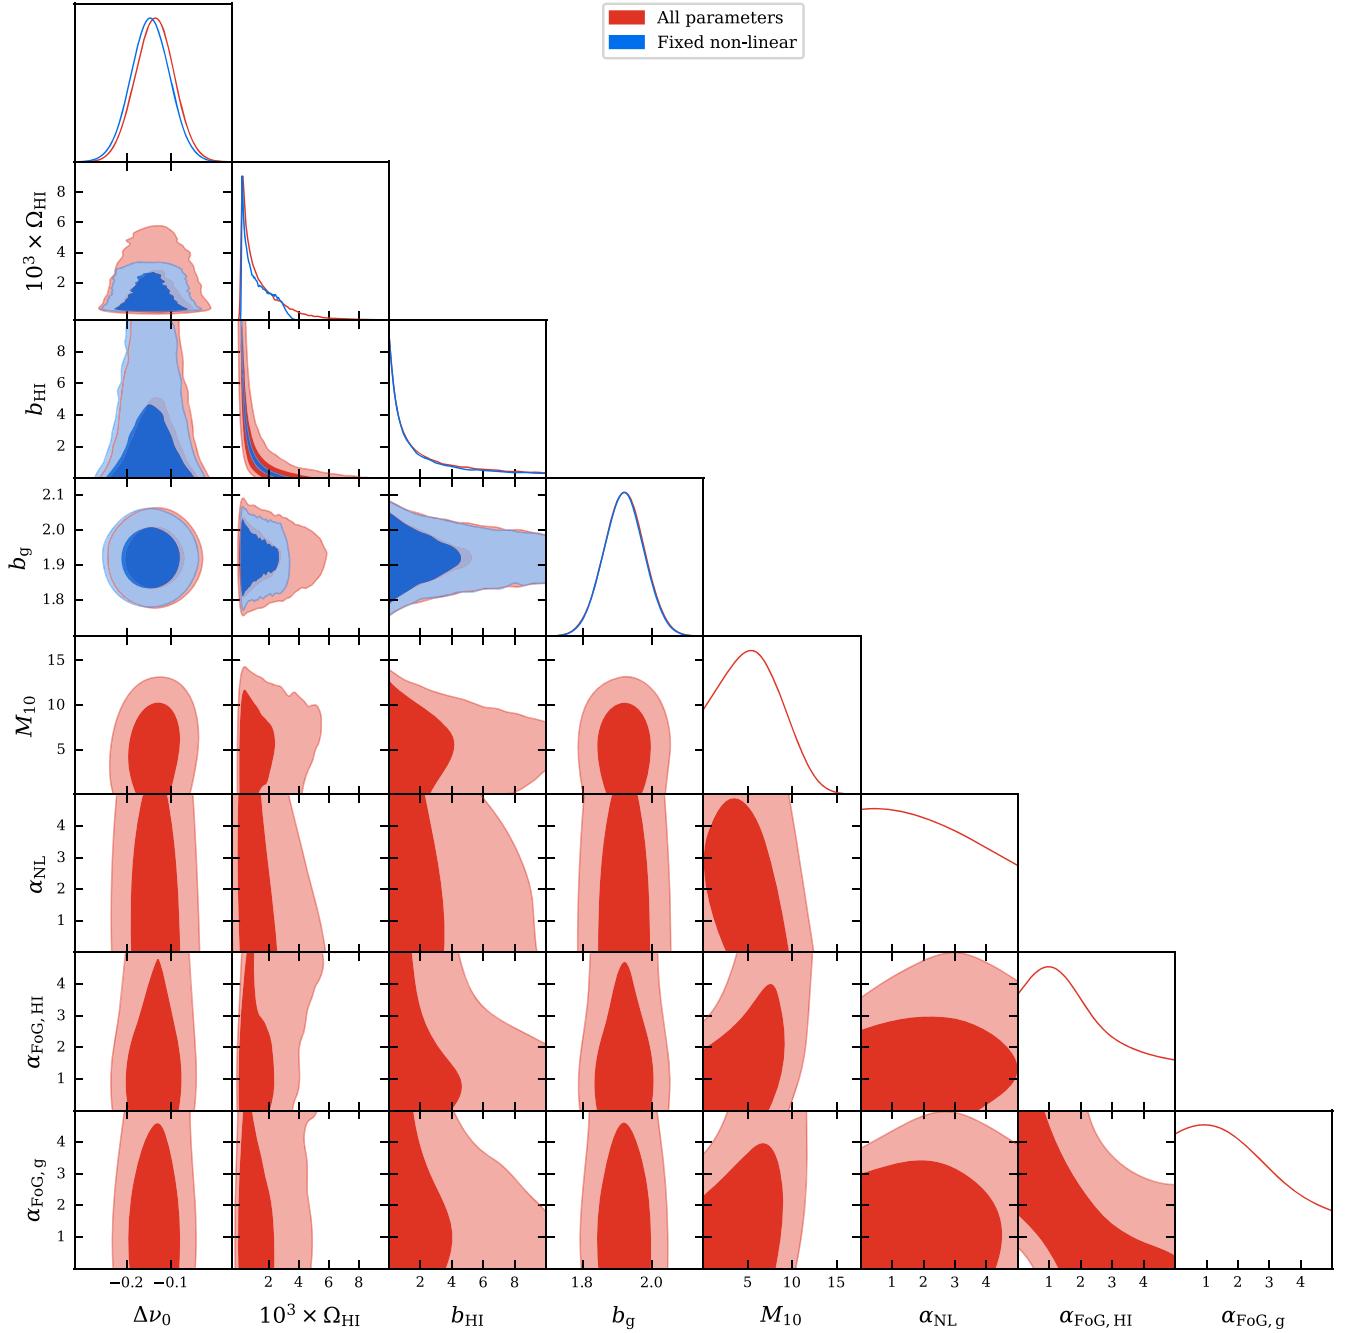


Figure 21. The constraints on the model derived from the cross-correlation of CHIME and the full eBOSS QSO sample. The red contours show the constraints on the full parameter set (described in Section 5.2.6), whereas the blue contours show the constraints if we fix the nonlinear parameters to their fiducial values and only allow $\Delta\nu_0$, Ω_{HI} , b_{HI} , and b_g to vary. The blue contours are not shown at all for the nonlinear parameters (M_{10} , α_{NL} , $\alpha_{\text{FoG,HI}}$, and $\alpha_{\text{FoG,g}}$), as they are not being varied, as will be the case for future constraint plots. There are significant degeneracies between parameters, notably $\Omega_{\text{HI}} - b_{\text{HI}}$, but also within the nonlinear parameters such as $\alpha_{\text{FoG,HI}} - \alpha_{\text{FoG,g}}$.

In Figure 22 we show these new parameters and how they are correlated with the parameters they are derived from. The new amplitude-like parameter \mathcal{A}_{HI} clearly flattens the degeneracy, capturing all the information in Ω_{HI} and b_{HI} . Similarly, the parameter $\alpha_{\text{FoG},+}$ correlates with the amplitude parameter, whereas the orthogonal combination $\alpha_{\text{FoG},-}$ does not, although there is interesting behavior observed at low $\alpha_{\text{FoG},+}$ where we are even further from the regime where we can make the high- k_{\parallel} expansion used in Equation (107).

One of the key remaining degeneracies is that between the overall amplitude, \mathcal{A}_{HI} , and the combined Finger-of-God

strength, $\alpha_{\text{FoG},+}$. This can be understood physically: on the scales that CHIME observes, the Finger-of-God damping reduces the stacked signal amplitude, and so an increase in $\alpha_{\text{FoG},+}$ must be compensated by an increase in the underlying 21 cm signal amplitude to remain consistent with the measurements.

In Figures 23–25 we show the constraints for the QSO, ELG, and LRG tracers stacked over the full 585–800 MHz band for the amplitude parameter, \mathcal{A}_{HI} , the frequency offset, $\Delta\nu_0$, the shot noise, M_{10} , and the two nonlinear nuisance parameters, $\alpha_{\text{FoG},+}$ and α_{NL} . In all cases we find an excellent

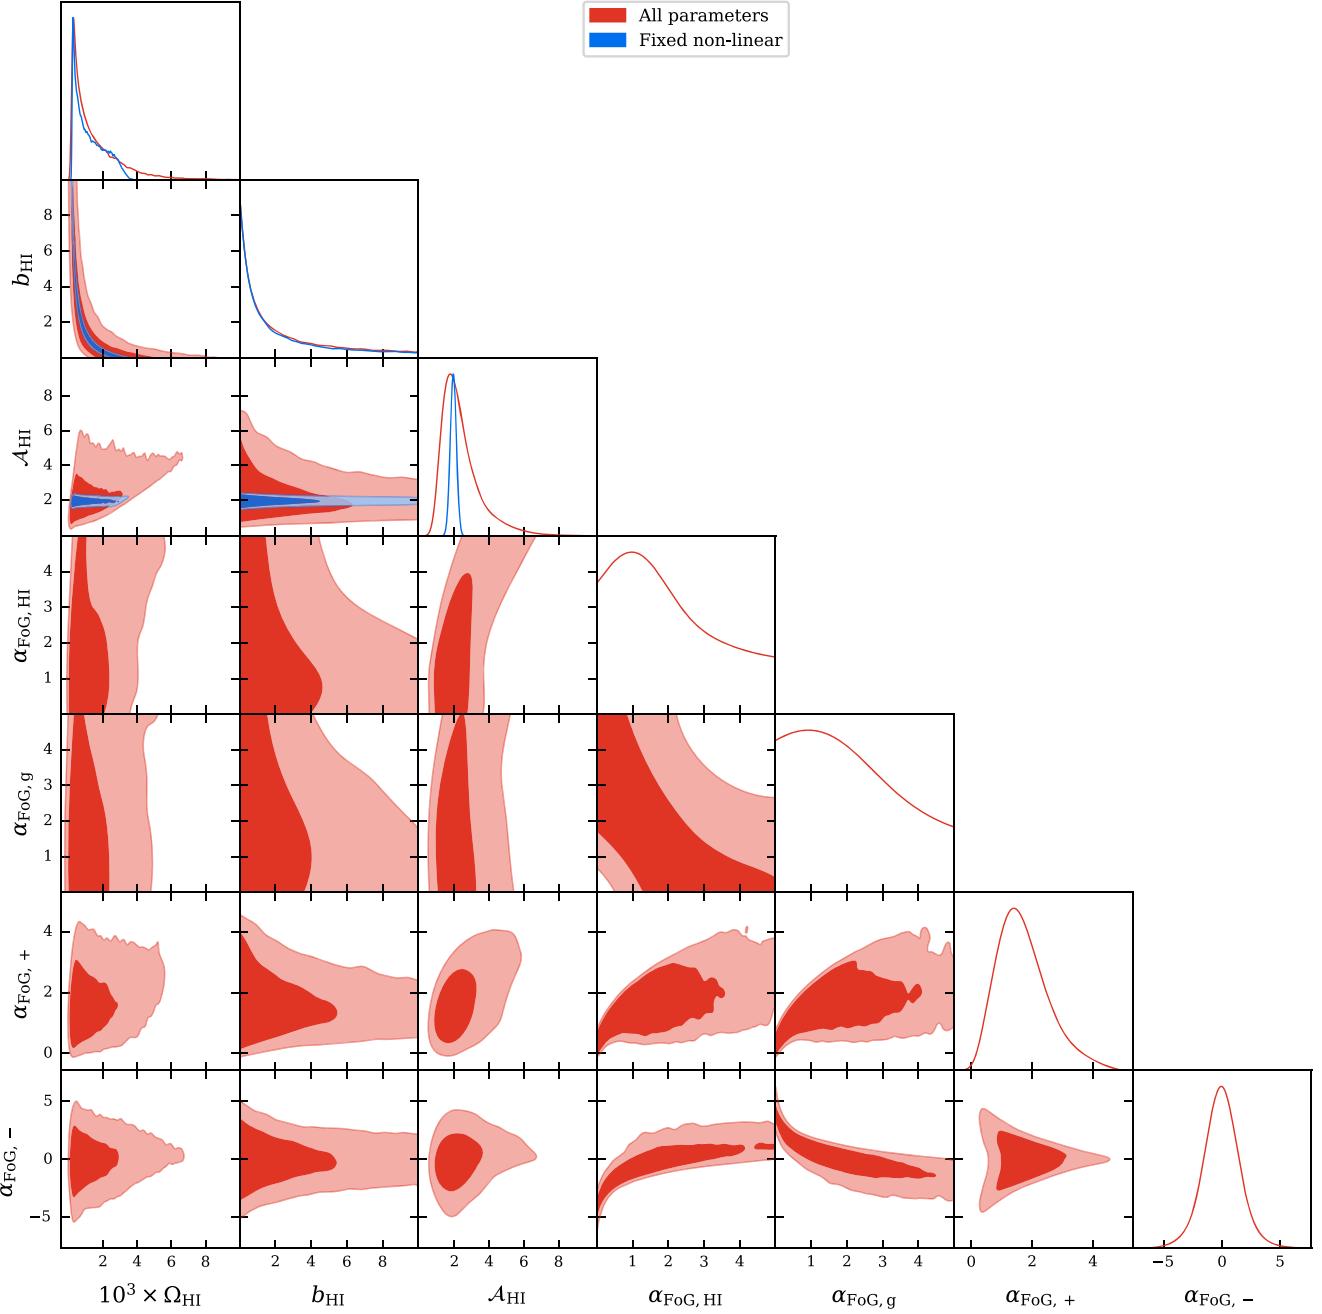


Figure 22. The parameter constraints corresponding to the QSO catalog, showing the derived parameters, A_{HI} , $\alpha_{FoG,+}$, and $\alpha_{FoG,-}$, and their correlations with the parameters they are derived from (Ω_{HI} , b_{HI} , $\alpha_{FoG,H}$, and $\alpha_{FoG,g}$). This figure illustrates that these new parameters are less degenerate than the original parameters, with A_{HI} decorrelating the Ω_{HI} – b_{HI} plane for both the fixed nonlinear models and full models, and with $\alpha_{FoG,+}$ – $\alpha_{FoG,-}$ effectively decorrelating the $\alpha_{FoG,H}$ – $\alpha_{FoG,g}$ plane for the full model, although the structure of the $\alpha_{FoG,+}$ – $\alpha_{FoG,-}$ posterior surface remains complex. After this, the key behavior in the constraints can be captured by just two quantities, A_{HI} and $\alpha_{FoG,+}$, with a simple positive correlation.

goodness of fit with χ^2_{\min} being close to the 202 degrees of freedom we expect (this comes from the 101 frequency offsets in the stacked data for each of the X and Y polarizations).

For all tracers, the amplitude constraints are significantly weakened by marginalizing over the nonlinear parameters compared to the case of fixed nonlinear parameters. However, the posteriors are non-Gaussian and highly skewed such that, despite the large credible interval, the probability that $A_{HI} \leq 0$ is negligible. Even though the nonlinear parameters are degenerate with the amplitude, the amplitude must be nonzero for a signal to be seen.

Interpretation of these constraints is complicated by a volume factor pushing the constraints toward larger Finger-of-God smoothing effects. The originally uniform prior on $\pi(\alpha_{FoG,H}, \alpha_{FoG,g})$ transforms to a $\pi(\alpha_{FoG,+}, \alpha_{FoG,-}) \propto \alpha_{FoG,+}$. As the stack signal $\propto A_{HI}/\alpha_{FoG,+}^2$, our broad noninformative priors give an unintentional upward pressure on A_{HI} , as there is more volume at higher FoG damping levels. This can be resolved by placing a flat prior on $\alpha_{FoG,+}$, but as it is not a physical parameter, this is difficult to justify. Future analysis will need to have data that can break this degeneracy internally, or use better modeling that allows for the prior bounds on $\alpha_{FoG,H}$ and $\alpha_{FoG,g}$ to be reduced.

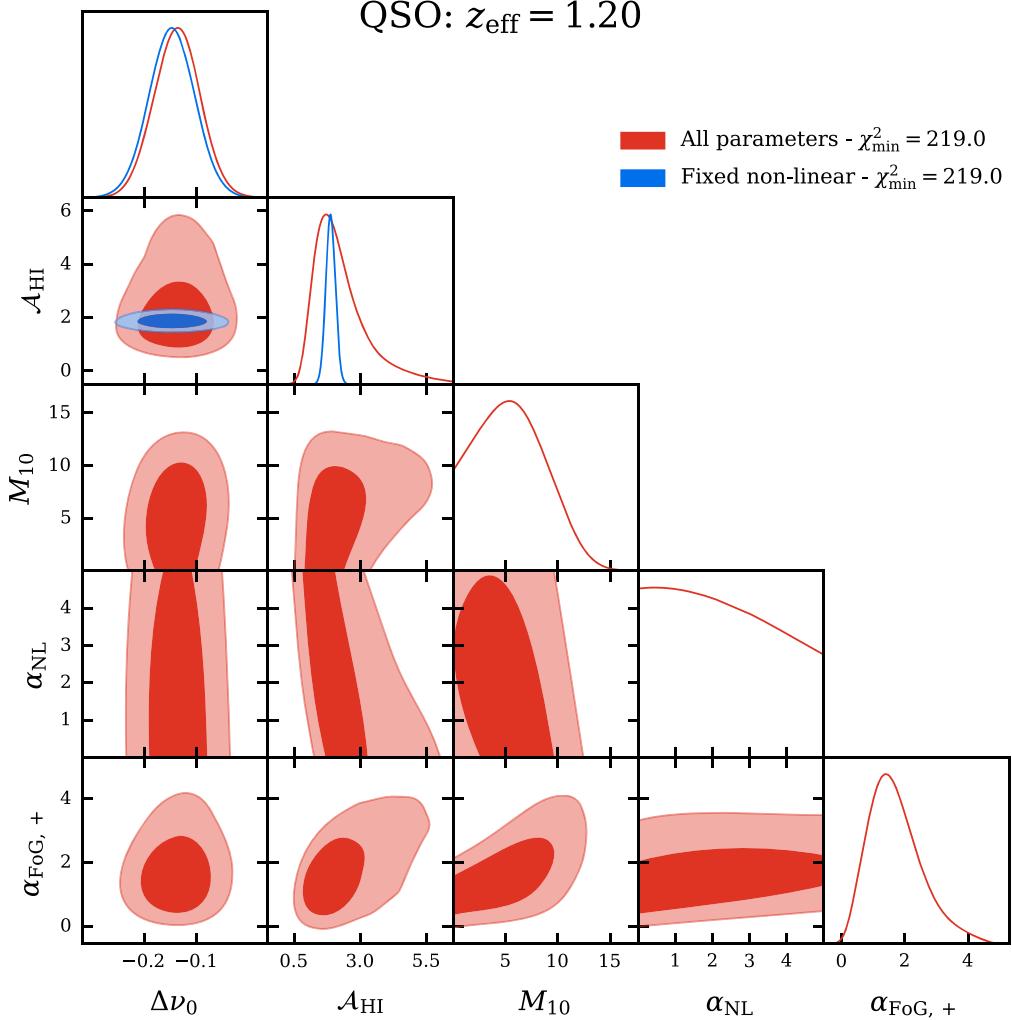


Figure 23. The constraints from stacking the QSO catalog over the full frequency band chosen for this analysis. We reduce the original set of eight parameters down to five: the amplitude-like A_{HI} , the frequency offset $\Delta\nu_0$, the correlated shot noise M_{10} , and the two relevant nonlinear nuisance parameters $\alpha_{FoG,+}$ and α_{NL} . The fits with all five parameters free (red contours) or with the M_{10} , $\alpha_{FoG,+}$, and α_{NL} fixed to their fiducial values (blue contours) both result in an excellent goodness of fit, with $\chi^2_{\min} \approx 219$ for 202 degrees of freedom. We discuss the physical interpretation of these constraints in Sections 8.1, 8.3, and 8.4.

6.4. Detection Significance

Assessing the significance of the detection is difficult for two reasons. First, the posterior distributions of the full set of parameters are highly non-Gaussian, which means that a naive “mean over standard deviation” figure does not accurately represent the significance of a parameter being nonzero. Second, there is not a single amplitude-like parameter that we can use to assess significance. Although we are primarily interested in Ω_{HI} or A_{HI} , whose posterior distributions include the projected degeneracies with α_{NL} and $\alpha_{FoG,x}$, they are also somewhat degenerate with M_{10} , and this should be captured, as any measurement of M_{10} should also count toward a detection.

One way of describing the detection significance is by way of a Bayesian model comparison. In this case, we seek to compare two explanations of the data, one in which the signal is represented by the full signal model given above (\mathcal{M}_1), and a null model where the signal is exactly zero and the data are entirely noise (\mathcal{M}_0). To compare these, we need to calculate the marginal likelihood, or *Bayesian evidence*, \mathcal{Z} , which is the normalization constant for the posterior distribution shown in

Equation (99):

$$\mathcal{Z} = \mathcal{P}(\mathbf{d}|\mathcal{M}) \quad (110)$$

$$= \int \mathcal{P}(\mathbf{d}|\boldsymbol{\theta}, \mathcal{M}) \mathcal{P}(\boldsymbol{\theta}|\mathcal{M}) d^n\theta \quad (111)$$

$$= \int \mathcal{L}(\boldsymbol{\theta}) \pi(\boldsymbol{\theta}) d^n\theta. \quad (112)$$

The evidence allows us to compare the relative probability of two models given the observed data

$$\frac{\mathcal{P}(\mathcal{M}_1|\mathbf{d})}{\mathcal{P}(\mathcal{M}_0|\mathbf{d})} = \frac{\mathcal{P}(\mathbf{d}|\mathcal{M}_1)}{\mathcal{P}(\mathbf{d}|\mathcal{M}_0)} \times \frac{\mathcal{P}(\mathcal{M}_1)}{\mathcal{P}(\mathcal{M}_0)} \quad (113)$$

$$= \frac{\mathcal{Z}_1}{\mathcal{Z}_0} \times \frac{\mathcal{P}(\mathcal{M}_1)}{\mathcal{P}(\mathcal{M}_0)}, \quad (114)$$

where the $\mathcal{P}(\mathcal{M})$ terms give the prior probabilities of the models. We assume that the model prior probabilities are equal from this point on, and we focus solely on the Bayesian evidence ratio $\mathcal{Z}_1/\mathcal{Z}_0$ (often termed the Bayes factor).

Calculating the evidence directly is challenging, as the region of high likelihood is typically much smaller than the

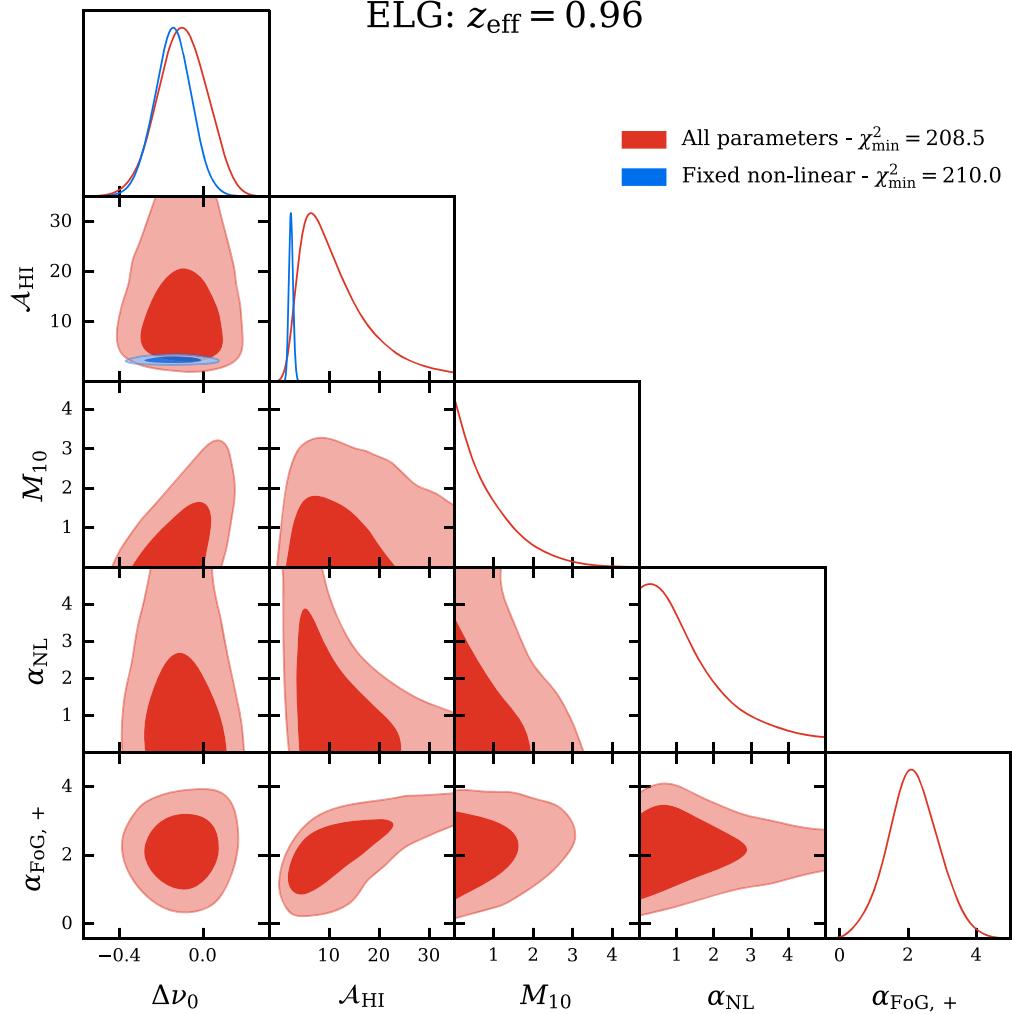


Figure 24. Parameter constraints from stacking the ELG catalog, in the same format as Figure 23.

prior volume, and so estimates tend to be dominated by sample noise. The standard techniques for evidence calculation are variants on nested sampling (Skilling 2006), but here we instead use the simpler process of thermodynamic integration (Gelman & Meng 1998), as we do not need the extra efficiency of nested-sampling-based techniques. To do this, we introduce the quantity

$$\mathcal{Z}(\lambda) = \int \mathcal{L}(\theta)^\lambda \pi(\theta) d^n\theta. \quad (115)$$

Noting that $\mathcal{Z}(0) = 1$ and $\mathcal{Z}(1) = \mathcal{Z}$, we can write the quantity we want to calculate as

$$\ln \mathcal{Z} = \int_0^1 \frac{\partial \ln \mathcal{Z}(\lambda)}{\partial \lambda} d\lambda. \quad (116)$$

This transformation is useful because we can write the integrand as

$$\frac{\partial \ln \mathcal{Z}(\lambda)}{\partial \lambda} = \frac{1}{\mathcal{Z}(\lambda)} \int \ln \mathcal{L}(\theta) \mathcal{L}(\theta)^\lambda \pi(\theta) d^n\theta \quad (117)$$

$$= \langle \ln \mathcal{L}(\theta) \rangle_\lambda, \quad (118)$$

where the $\langle \dots \rangle_\lambda$ denotes an expectation evaluated against a posterior with the likelihood raised to the power λ . This gives us a straightforward way of calculating $\ln \mathcal{Z}$: first, on a discrete

grid in λ , we use a standard MCMC sampler to draw from the unnormalized distribution $\mathcal{L}(\theta)^\lambda \pi(\theta)$, and then we estimate $\langle \ln \mathcal{L}(\theta) \rangle_\lambda$ from these samples; second, we numerically integrate over these estimates to calculate $\ln \mathcal{Z}$.

We calculate the evidence for the signal model, \mathcal{M}_1 , by multiple sampling runs (as described in Section 6.2) generated at different λ . As the bulk of the variation in the integrand is around $\lambda \sim 0$, we use the common choice of a grid regularly spaced in $\lambda^{1/5}$ (Calderhead & Girolami 2009), and as the integrands are smooth and well behaved, we find that a Romberg integration over 33 samples achieves sufficient accuracy. For the evidence calculation, we use shorter chains per λ step than for the parameter estimation, with only 15,000 samples per chain. After removing the initial samples for burn-in and thinning to the independent samples, this leaves ~ 700 samples for each λ step. To estimate the error on each evidence calculation, we bootstrap resample the set of points at each λ step, integrate over the resampled sets, and estimate the sample variance over bootstrap sets. This gives a typical error in $\ln \mathcal{Z}$ of ~ 0.1 . In contrast, the null signal model, \mathcal{M}_0 , is a zero-parameter model, and so its evidence is simply the likelihood of the data evaluated at zero signal. That is,

$$\ln \mathcal{Z}_0 = -\frac{1}{2} [\ln |2\pi\Sigma| + \chi_0^2], \quad (119)$$

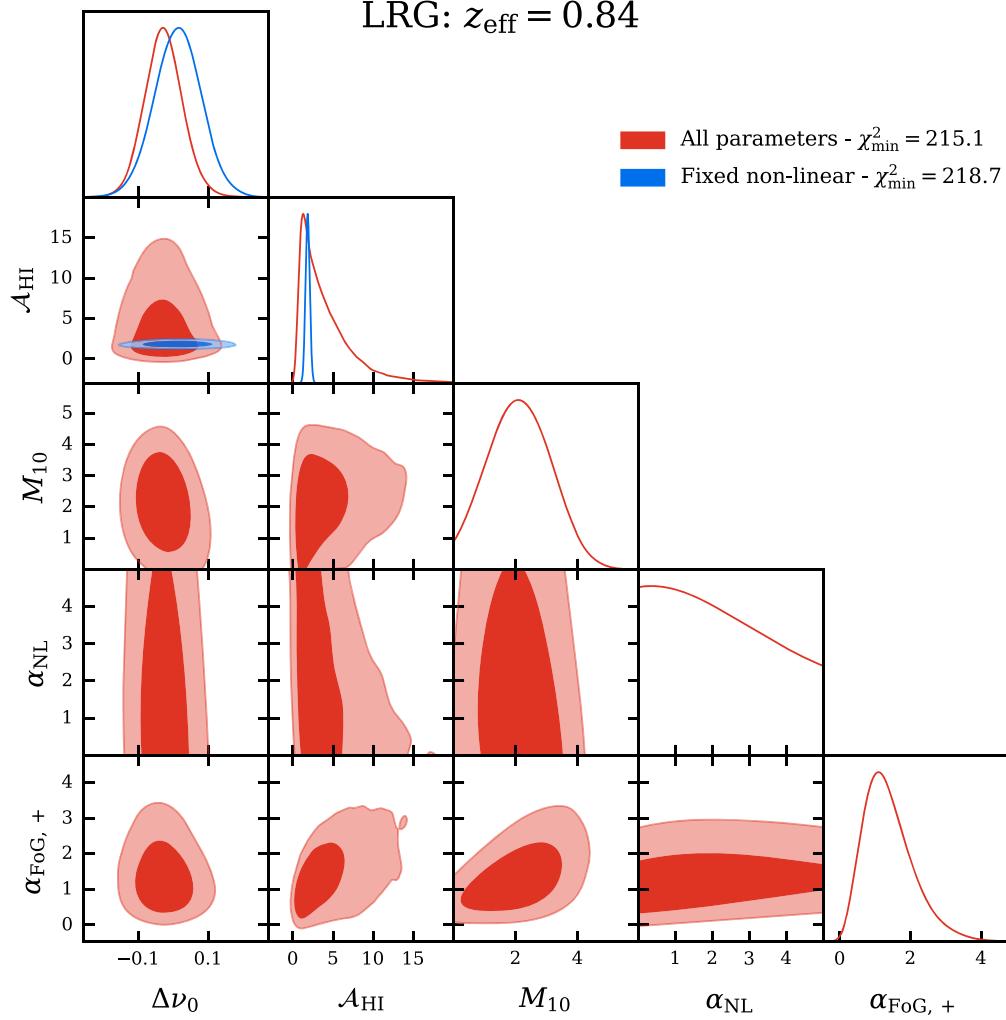


Figure 25. Parameter constraints from stacking the LRG catalog, in the same format as Figure 23.

with

$$\chi_0^2 = (\mathbf{d} - \boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1} (\mathbf{d} - \boldsymbol{\mu}), \quad (120)$$

so that we do not need any MCMC scheme to calculate it.

With $\ln \mathcal{Z}_0$ and $\ln \mathcal{Z}_1$, we have both of the ingredients required to give the Bayes factor. To enable a comparison with other significance estimates, we can turn the evidence ratio into an effective “number of sigma.” Assuming that the only two models that could explain the data are \mathcal{M}_0 and \mathcal{M}_1 and giving them equal prior probabilities, $\mathcal{P}(\mathcal{M}_0) = \mathcal{P}(\mathcal{M}_1) = 1/2$, we can write the probability of the null model as

$$\mathcal{P}(\mathcal{M}_0|\mathbf{d}) = \frac{1}{1 + \mathcal{Z}_1/\mathcal{Z}_0}. \quad (121)$$

We turn this into an effective number of σ , $N_{\mathcal{Z}}$, via

$$N_{\mathcal{Z}} = \Phi^{-1}(1 - \mathcal{P}(\mathcal{M}_0|\mathbf{d})), \quad (122)$$

where $\Phi^{-1}(x)$ is the inverse cumulative distribution function of the standard normal distribution.

An alternative, frequentist method of estimating the detection significance is to use a likelihood ratio test. First, we compute the ratio of the maximum likelihood values between a

model with no signal and one with the full signal model

$$\lambda = 2 \ln \left(\frac{\mathcal{L}(\hat{\boldsymbol{\theta}}_{\text{ML}})}{\mathcal{L}_0} \right) \quad (123)$$

$$= \Delta \chi^2, \quad (124)$$

with $\Delta \chi^2 = \chi_0^2 - \chi_{\text{min}}^2$. This quantity is asymptotically χ^2 distributed with degrees of freedom equal to the effective number of model parameters. As our model has several notable degeneracies, the effective number of model parameters will be less than the total number of parameters. We use the Bayesian model dimensionality (Handley & Lemos 2019)

$$d_M = 2[\langle (\ln \mathcal{L})^2 \rangle - \langle \ln \mathcal{L} \rangle^2] \quad (125)$$

as an estimate of the number of parameters, where the expectation $\langle \dots \rangle$ is taken over the posterior. Taking an average of this over the set of tracers, we find $d_M \sim 4.4$, and so we use 4 as the effective number of parameters. Using this, we can ascribe a detection significance via the probability for a $\chi^2_{\nu=4}$ distribution to exceed λ . We again turn this into an effective number of sigma using the inverse cumulative distribution

Table 3
The Detection Significance for Each Tracer

Tracer	Bayesian			Likelihood Ratio		Amplitude N_A
	$\ln(\mathcal{Z}_1/\mathcal{Z}_0)$	$\ln(\mathcal{Z}_1/\mathcal{Z}_2)$	N_Z	$\Delta\chi^2$	N_{LR}	
LRG	18.9	-1.5	5.7	60.3	7.1	7.8
ELG	10.8	-2.4	4.1	40.8	5.7	6.4
QSO	56.3	-2.2	10.3	133.5	11.1	11.6
QSOb0	23.9	-2.3	6.5	66.2	7.5	8.1
QSOb1	19.6	0.8	5.8	53.2	6.6	7.3
QSOb2	16.9	-0.9	5.3	50.0	6.4	7.1
QSOb00	7.6	1.5	3.3	27.8	4.5	5.3
QSOb01	14.6	-1.6	4.9	46.3	6.1	6.8

Note. We calculate the detection significance by three different methods, using the Bayesian evidence, a likelihood ratio test, and the amplitude constraints on the best-fit model. The log Bayes factors, $\ln(\mathcal{Z}_1/\mathcal{Z}_0)$, all exceed the highest threshold $\gtrsim 4.6$ for *decisive* evidence according to the scale of Jeffreys (1961), and the $\Delta\chi^2$ values all have p -values $\lesssim 10^{-5}$. We also convert each to an effective number of sigmas for comparison, giving N_Z , N_{LR} , and N_A , respectively. The results are similar; however, N_Z suffers from the choice of wide priors, and the amplitude ratio N_A is overly optimistic, as it does not account for the previous fitting of the template. We also give the evidence ratio $\ln(\mathcal{Z}_1/\mathcal{Z}_2)$ comparing the full signal compared to fixing the nonlinear parameters. For most catalogs there is no evidence in favor of the full model ($\ln \mathcal{Z} < 0$), that is, the improved fit is not sufficient to support the additional free parameters. Only for the QSOb00 tracer is there moderate evidence to support the full model.

function of a standard normal distribution:

$$N_{LR} = \Phi^{-1} \left(1 - \int_{\lambda}^{\infty} \chi_4^2(x) dx \right). \quad (126)$$

As a final estimate of the detection significance, we take the best-fit (minimum χ^2) template as a fixed single template and then fit that directly to the data with a varying amplitude, A . As the likelihood is Gaussian, the distribution of A can be computed exactly and is Gaussian with mean of 1 and variance $(\Delta\chi^2)^{-1}$. This directly gives the number of sigma of detection, $N_A = \sqrt{\Delta\chi^2}$. We expect this quantity to overestimate the detection significance, as the template has already been adjusted to fit the data.

Table 3 shows the detection significance for each tracer calculated by each method given above. The Bayes factors, $\ln(\mathcal{Z}_1/\mathcal{Z}_0)$, are $\gtrsim 4.6$ for all tracers, which corresponds to *decisive* evidence for a cross-correlation detection according to the interpretations of Jeffreys (1961) and Kass & Raftery (1995). The numbers of sigmas for each method are reasonably close, with the Bayesian-evidence-based number N_Z the lowest of the three and the amplitude parameter the highest. The common criticism of evidence calculations is that they are dependent on the prior widths, and, as is the case here, a choice that is intended to be noninformative for the purpose of parameter estimation can significantly lower the evidence compared to a less conservative choice of prior. In our case, parameters like M_{10} could be significantly narrower without influencing the parameter estimation, which would boost the Bayes factor. Although we do not attempt it here, one resolution to this for nested comparisons (of which this is one), advocated by Gordon & Trotta (2007), is to optimize the prior widths centered on the value implied by the nested model to maximize the Bayes factor.

We also calculate the evidence for the signal model where we fix the nonlinear parameters, which we call \mathcal{Z}_2 , and give the log Bayes factor relative to the full signal model, $\ln(\mathcal{Z}_1/\mathcal{Z}_2)$, in Table 3. In most cases $\ln(\mathcal{Z}_1/\mathcal{Z}_2)$ is negative, that is, there is not sufficient improvement in the fits to justify the expanded model from statistical arguments alone, and in the remaining cases the evidence is marginal.

Our rationale for varying the nonlinear parameters is to explore what our data tell us about the large-scale H α distribution while including the genuine uncertainties in the modeling. With that in mind, we do not take this as an indication that we should fix these nonlinear parameters, but as one that they are not meaningfully constrained, as they allow the model to overfit the data.

7. Validation

In this section we describe several consistency tests that were performed on the analysis and inform the systematic errors that are placed on the result. These tests consist of evaluating whether the measurements made by the two polarizations are consistent, evaluating whether the signal is the same from day to day, estimating the uncertainty on the amplitude of the signal due to beam calibration errors, and evaluating the linearity of the analysis pipeline.

7.1. Consistency between Polarizations

The following procedure is used to determine whether measurements made with the X and Y baselines are consistent given our model for the noise and 21 cm signal. The two polarizations are jointly fit to a restricted model and an unrestricted model. For the restricted model, both polarizations are described by the same set of parameters, θ , as outlined in Section 6.2. The version of the model that holds the nonlinear parameters fixed at their fiducial values is employed for this exercise, since the version that allowed them to vary did not yield a significantly better fit to the data for any tracer or QSO redshift bin. For the unrestricted model, the polarizations are described by a different set of parameters, θ_X and θ_Y . The maximum likelihood estimate of the parameters is obtained for each model using the L-BFGS-B optimization algorithm. The following test statistic is then calculated:

$$\Delta\chi^2 = \chi^2(\hat{\theta}_{\text{res}}) - \chi^2(\hat{\theta}_{\text{unres}}), \quad (127)$$

where χ^2 is given by Equation (98), $\hat{\theta}_{\text{res}} \equiv \hat{\theta}$ denotes the maximum likelihood parameter estimates for the restricted model, and $\hat{\theta}_{\text{unres}} \equiv [\hat{\theta}_X, \hat{\theta}_Y]$ denotes the maximum likelihood parameter estimates for the unrestricted model. The χ^2 values and the test statistic are quoted in Table 4 for all tracers and all QSO redshift bins.

The test statistic will follow a χ^2 distribution with $\Delta\text{DOF} = \text{DOF}_{\text{res}} - \text{DOF}_{\text{unres}}$ degrees of freedom under the null hypothesis that the two polarizations are described by the same model. Naively we expect ΔDOF to be equal to the number of model parameters, since the unrestricted model has twice the number of parameters as the restricted model. However, the model parameters are highly degenerate, so that using the number of parameters would likely overestimate ΔDOF and bias the test toward accepting the null hypothesis.

To avoid this, the distribution of the test statistic under the null hypothesis is empirically measured using the random mock catalogs. We generate 10,000 realizations of our data by adding

Table 4
Model-dependent Test for Consistency between Polarizations

Tracer	χ^2		$\Delta\chi^2$	ΔDOF	PTE
	Restricted	Unrestricted			
LRG	218.6	214.2	4.4	2.3	0.12
ELG	210.1	209.3	0.7	2.3	0.77
QSO	219.0	213.9	5.0	2.3	0.10
QSOb0	210.5	208.3	2.2	2.3	0.37
QSOb1	202.1	200.8	1.2	2.3	0.62
QSOb2	220.4	214.3	6.1	2.3	0.05
QSOb00	185.5	184.8	0.6	2.2	0.78
QSOb01	235.8	233.4	2.3	2.2	0.35

Note. For each tracer and redshift bin, we report the minimum χ^2 obtained when fitting a model in which the two polarizations are described by the same set of parameters (restricted) and a different set of parameters (unrestricted). The distribution of the difference, $\Delta\chi^2$, under the null hypothesis that the polarizations are described by the same set of parameters is calibrated using random mock catalogs and approximately follows a theoretical χ^2 distribution with the quoted ΔDOF degrees of freedom. The PTE provides the fraction of random mock catalogs that exceed the value observed in the data.

the best-fit, restricted model and a stack on a random mock catalog. We then fit each realization to the restricted and unrestricted models and calculate the test statistic. The probability to observe a value of the test statistic in excess of that observed in the data is then determined from the empirical cumulative distribution function. The results are presented in the last column of Table 4. For all tracers and QSO redshift bins, the null hypothesis that the two polarizations are described by the same set of model parameters is accepted with the probability to exceed (PTE) > 0.05 . We also note that the empirical distributions are reasonably well described by a χ^2 distribution with $\Delta\text{DOF} \approx 2.3$ degrees of freedom.

7.2. Consistency between Even and Odd Days

The 102 sidereal days that were used to construct the sidereal stack are split into two subsets by chronologically ordering the days that went into each seasonal stack and then separating the even days into one set and the odd days into the other set (see Section 3.3.2). The two sets have size 50 and 52 sidereal days and a mean date that differs by 53 hr. Each set is averaged using the procedure outlined in Section 3.3. This yields two estimates of the visibilities that are then differenced according to

$$\Delta V = \frac{1}{c} (V_{\text{even}} - V_{\text{odd}}), \quad (128)$$

with

$$c = \begin{cases} \frac{w_{\text{even}} + w_{\text{odd}}}{\sqrt{w_{\text{even}} w_{\text{odd}}}} & \text{if } (w_{\text{even}} > 0) \wedge (w_{\text{odd}} > 0); \\ 0 & \text{otherwise.} \end{cases} \quad (129)$$

Here $V \equiv V_{xy}^p(\nu, \phi)$ and $w \equiv w_{xy}^p(\nu, \phi)$ denote the visibilities and corresponding weights. The quantity c is a scale factor that will set the variance of the radiometric noise in the difference equal to that in the weighted average. In the limit that the even and odd splits have equal radiometric noise, and hence equal weight, then $c = 2$. In reality, the two splits have slightly different weights such that $c = 2.0065 \pm 0.018$ over the

baselines, frequency, and right ascensions examined. The processing described in Section 3.3 through Section 4.7 is applied to the differenced visibility, with the caveat that we use the global frequency mask, delay cut, and primary beam model that were previously derived from the weighted average of the full set of days.

The cosmological 21 cm signal is constant as a function of sidereal day and is expected to cancel in the difference. The radiometric noise, on the other hand, will be independent in the two subsets and therefore will remain in the difference. Transient RFI is also expected to be independent in the two subsets and remain in the difference. Residual foregrounds caused by spectral leakage due to a chromatic instrument transfer function will be the same from day to day and hence cancel in the difference. Residual foregrounds due to seasonal changes in the instrument transfer function will also cancel. On the other hand, residual foregrounds due to changes in the instrument transfer function from day to day will remain.

Since a significant portion of the noise in the stack is due to residual foregrounds that will be mitigated by the differencing procedure, the covariance matrix of the even–odd difference is expected to change relative to the covariance matrix of the full data set. We recalibrate the covariance matrix with mock catalogs as outlined in Section 4.8. We find better agreement between the even–odd difference covariance and the expected radiometric noise, suggesting that the majority of $\sim 50\%$ excess noise in the full set is primarily due to foregrounds that are static from one day to the next.

Under the null hypothesis that the observed signal is the same on even and odd days, stacking the even–odd difference on the true catalog should be statistically equivalent to stacking on a random mock catalog. The distribution of the χ^2 test statistic for the random mock catalogs is well described by a theoretical χ^2 distribution with 202 degrees of freedom. Table 5 quotes the χ^2 value of the stack on each tracer and QSO redshift bin, as well as the fraction of the 10,000 random mock catalogs that have a χ^2 test statistic in excess of that observed for the true catalog. We find that all tracers and redshift bins have a PTE greater than 0.05, except for the LRG catalog, which has a PTE of 0.025, and the subset of the QSO catalog with a redshift between 0.91 and 1.03 (QSOb01), which has a PTE of 0.023. If the large χ^2 values are due to differences in the observed signal on even and odd days, then we would expect to see a copy of the signal in the jackknife. We recompute the test statistic using only frequencies $|\Delta\nu| < 5$ MHz, where the magnitude of the signal is largest. We find that the PTE for the LRG catalog increases to 0.12, the QSOb01 catalog decreases to 0.009, and all other tracers and QSO redshift bins have a value greater than 0.05. This leads us to conclude that the large χ^2 observed when stacking the jackknife on the LRG catalog originates from a rare noise fluctuation rather than differences in the signal on even and odd days. However, the large χ^2 for the QSOb01 catalog warrants additional investigation.

To explore this further, we perform a model-dependent analysis of the even and odd days that is similar to the analysis used to check for consistency between polarizations, which was described in Section 7.1. Each split is processed independently through the pipeline and then stacked on the true catalog and the random mock catalogs for both splits to ensure that the covariance matrix captures correlated noise between them. The two splits are

Table 5
Model-independent Test for Consistency between Even and Odd Days

Tracer	$ \Delta\nu \leq 20$ MHz (202 DOF)		$ \Delta\nu \leq 5$ MHz (50 DOF)	
	χ^2	PTE	χ^2	PTE
LRG	242.4	0.025	62.0	0.12
ELG	199.1	0.54	47.2	0.58
QSO	202.0	0.49	59.2	0.17
QSOb0	233.5	0.064	66.3	0.062
QSOb1	177.2	0.90	35.1	0.95
QSOb2	190.5	0.71	57.9	0.21
QSOb00	207.5	0.38	55.9	0.26
QSOb01	243.5	0.023	76.3	0.009

Note. For each tracer and redshift bin, we report the χ^2 test statistic when stacking the catalog on a jackknife of even and odd days. Under the null hypothesis that the 21 cm signal is the same on even and odd days, this will follow a χ^2 distribution with the stated degrees of freedom. The PTE provides the fraction of random mock catalogs that exceed the value observed when stacking on the true catalog.

jointly fit to both a restricted and unrestricted model. The restricted model describes both splits with the same set of parameters. The unrestricted model describes each split with a different set of parameters. We employ the version of our model where the nonlinear parameters are held fixed at their fiducial values (see Table 2). We compute $\Delta\chi^2$ as given by Equation (127) and calibrate its distribution under the null hypothesis using the random mock catalogs. The results are presented in Table 6.

As anticipated, the LRG catalog passes the test (PTE = 0.71) and the QSOb01 catalog fails the test (PTE = 0.02). The QSOb0 catalog, which contains all QSOs with a redshift between 0.80 and 1.03 and is a superset of QSOb01, also fails the test (PTE = 0.005). The discrepancy appears primarily in the amplitude of the signal, with the even-day split exhibiting an approximately 50% larger amplitude than the odd-day split for these two redshift bins. We perform an MCMC fit of both the restricted and unrestricted models and use the posterior distributions of the amplitude parameter \mathcal{A}_{H_1} to characterize the fractional error in \mathcal{A}_{H_1} implied by this discrepancy. This is defined as half the difference in \mathcal{A}_{H_1} between the even and odd splits as measured by the unrestricted model fit divided by the most likely value from the restricted model fit, and it is quoted in the last column of Table 6.

The discrepancy is suggestive of a 50% difference in our calibration between even and odd days at frequencies between 700 and 745 MHz. However, we have ruled out an error in the relative calibration of this magnitude by examining the spectra of 34 bright point sources in the NGC field extracted from the maps prior to foreground filtering. We find that the difference in spectra between the even and odd days is at most 1% over all sources and frequencies.

In order to account for the observed discrepancy, we will assume an additional 25% systematic error on the \mathcal{A}_{H_1} constraint for the QSOb0 and QSOb01 catalogs.

7.3. Beam Calibration Errors

In order to estimate the uncertainty on the default beam model described in Section 4.4, it is compared to independent

measurements of the beam from observations of the Sun and holographic observations of bright point sources made in conjunction with the John A. Galt 26 m telescope (CHIME Collaboration et al. 2022b). Based on these comparisons, we estimate that within the main lobe the beam model is accurate to $\lesssim 5\%$ relative to the beam on meridian at the decl. of Cygnus A. Currently our beam calibration technique is unable to constrain the sidelobes of the beam (for details see Appendix B). The solar and holographic data both suggest that the sidelobes are $\lesssim 1\%$ at hour angles $\lesssim 30^\circ$ and $\lesssim 0.1\%$ at hour angles $\gtrsim 30^\circ$. It is estimated that approximately 10% of the beam solid angle lies outside the region that we are able to measure with the default beam model.

The solar beam measurements are described in CHIME Collaboration et al. (2022a) and span $-23.5^\circ \leq \theta \leq 23.5^\circ$, which corresponds to the range of apparent decl. that the Sun travels between winter and summer solstice. The rms difference between the solar and default beam model is 4% (relative to the beam on meridian at the decl. of Cygnus A) in the region $|\text{HA}| \lesssim 3^\circ$, $|\theta| < 23.5^\circ$, $587.5 \text{ MHz} < \nu < 800 \text{ MHz}$. However, the inferred amplitude of the 21 cm signal is primarily sensitive to the fractional error in the beam on meridian when averaged over the large range of decl. and frequencies covered by the eBOSS catalogs. In order to estimate the systematic error on the 21 cm amplitude due to beam uncertainties, the fractional difference between the default and solar beam model on meridian at the decl. and 21 cm frequency of each source in each catalog is extracted and then averaged using the same weights that are used in the stacking procedure described in Section 4.7. Only 3% of the QSOs and LRGs in the NGC field are at decl. that overlap with the solar data. However, 42% of the ELGs in the NGC field lie at decl. where there are two independent measurements of the beam, and the average fractional difference between these two measurements is 6%.

We have also compared the flux density of the brightest radio sources in the deconvolved map to their expected flux density in order to obtain an additional estimate of the systematic uncertainty on the 21 cm amplitude. The expected flux densities are obtained by interpolating recent measurements made by the VLA to frequencies in the CHIME band (P17). There are 14 sources in total used for this purpose, with an average decl. separation of 5° . These sources do not provide a dense sampling of the decl. axis like the solar data, but they do cover the full range of decl. spanned by the eBOSS catalogs. All 14 sources have interpolated flux densities that are accurate at the subpercent level and are greater than 10 Jy at 600 MHz. The rms of the fractional error in the flux density of these sources in the deconvolved map is 5.0%, 6.4%, and 7.4% for the range of frequencies and decl. spanned by QSO, LRG, and ELG catalogs in the NGC field, respectively. Taking instead the weighted average of the fractional error in the flux at the decl. and 21 cm frequencies nearest to the sources in each catalog yields 0.6%, 2%, and 0.5% for the QSO, LRG, and ELG catalogs in the NGC field. Note that this is an end-to-end test of our ability to recover the true flux of point sources and is sensitive to a variety of potential sources of systematic error, including beam calibration errors, but also complex gain errors and regridding artifacts.

As a final check, we simulate observations of the fiducial 21 cm signal using both the default beam and the control beam. For each of these simulations, we construct a map by deconvolving both the default beam and control beam and

Table 6
Model-dependent Test for Consistency between Even and Odd Days

Tracer	χ^2		$\Delta\chi^2$	ΔDOF	PTE	$\frac{\Delta A_{\text{HI}}}{2A_{\text{HI}}}$
	Restricted	Unrestricted				
LRG	484.0	483.0	1.0	2.4	0.71	$-0.11^{+0.16}_{-0.14}$
ELG	414.7	412.7	2.1	2.3	0.42	$-0.15^{+0.22}_{-0.11}$
QSO	425.2	422.4	2.8	2.6	0.34	$-0.06^{+0.11}_{-0.10}$
QSOb0	447.9	436.7	11.2	2.4	0.005	$0.26^{+0.12}_{-0.12}$
QSOb1	370.6	368.3	2.4	2.5	0.41	$-0.12^{+0.14}_{-0.12}$
QSOb2	452.7	446.4	6.3	2.4	0.052	$-0.16^{+0.16}_{-0.13}$
QSOb00	378.7	378.1	0.6	2.3	0.81	$0.11^{+0.20}_{-0.17}$
QSOb01	468.6	460.1	8.5	2.4	0.020	$0.25^{+0.13}_{-0.11}$

Note. For each tracer and redshift bin, we report the minimum χ^2 obtained when fitting a model in which the even and odd splits are described by the same set of parameters (restricted) and a different set of parameters (unrestricted). The distribution of the difference, $\Delta\chi^2$, under the null hypothesis that the splits are described by the same set of parameters is calibrated using random mock catalogs and approximately follows a theoretical χ^2 distribution with the quoted ΔDOF degrees of freedom. The PTE provides the fraction of random mock catalogs that exceed the value observed in the data. The last column provides the fractional error on the amplitude parameter, A_{HI} , inferred from this comparison.

then stack the map on simulated catalogs. We then examine the fractional difference in the amplitude of the stacked signal for the four different pairs of (simulation beam model, deconvolution beam model) relative to the (default, default) pair that was used in the actual analysis. For all four pairs and all three tracers the observed difference is less than 6%. This provides an estimate of the systematic error due to uncertainty in the interference pattern that modulates the beam. Note that this is a very conservative estimate because the uncertainty on the interference pattern is roughly a factor of 10 less than than the amplitude of the interference pattern itself.

Based on the solar comparison, bright point-source comparisons, and simulations of different beam models, a conservative 8% systematic error on the amplitude of the 21 cm signal will be assumed for all fields.

7.4. Linearity

Many of the elements in our analysis pipeline, such as the delay filtering, are explicitly linear, meaning that they operate independently on the 21 cm signal and foregrounds present in the data. To characterize the linearity of the entire analysis, we inject simulated 21 cm signal into the data, process the signal+data combination in the same way as the data, and stack the results on mock eBOSS catalogs that are correlated with the simulated signal. We also separately perform the stacking on mock catalogs using the data without injected signal and using mock observations containing only the injected signal. In a perfectly linear analysis, the difference of the signal+data and data stacks will be equal to the signal-only stacks, while nonlinearities will cause a violation of this equality. This method has previously been used to characterize signal loss in 21 cm analyses that rely on strongly nonlinear foreground filtering techniques (e.g., Masui et al. 2013; Paciga et al. 2013).

In detail, we generate correlated 21 cm and galaxy number density sky maps and propagate them through to simulated time streams and mock LRG, ELG, or QSO catalogs following the procedures in Section 5.3. The signal-only time stream is added to the sidereal stack derived from the data prior to subtraction of the brightest point sources (i.e., in the first box in Figure 6), and this combined time stream is passed through the same analysis pipeline as the data, culminating in the beam-

deconvolved, filtered, masked map being stacked on the mock catalogs. Prior to delay filtering, the signal-only map has $\text{rms} \sim 0.3 \text{ mJy beam}^{-1}$, compared to $\sim 3 \text{ Jy beam}^{-1}$ for the data map; therefore, the addition of signal to the data map has negligible effect on the determination of the elevation-dependent delay cut (Section 4.5), or on which frequencies are identified as outliers (Section 4.6), so these aspects of the analysis are not regenerated for the signal+data combination.

However, the final masking step—which masks map pixels whose absolute value exceeds a chosen threshold, based on the estimated map noise level—is explicitly nonlinear, so we recompute this mask to determine the impact of this nonlinearity on the recovered signal. We find that this impact is significant, which can be explained as follows. The distribution of pixel values in the signal-only map is symmetric about zero, but this distribution is skewed positive if one only considers pixels containing an object in a given mock catalog, since these objects are more likely to occupy pixels corresponding to matter overdensities, which are also correlated with 21 cm emission. Thus, if a given pixel (containing a catalog object) in the data map is positive and just below the mask threshold, it is more likely to be perturbed above the threshold by the addition of the signal-only map; similarly, a given negative pixel that is just beyond the threshold is more likely to be perturbed within the threshold.

The net effect is that the signal injection increases the number of negative near-threshold unmasked pixels and decreases the number of such positive pixels, resulting in an artificial attenuation of the overall stacking amplitude. A lower threshold will result in a greater number of affected pixels and more severe attenuation, while a higher threshold will mitigate this, but at the expense of decreasing the signal-to-noise ratio owing to a greater number of anomalous pixels being included in the stack.

Figure 26 quantifies these two effects. For each threshold in the figure, we compute the difference of stacks on signal+data and data-only maps, form a “prediction” given by a stack on the corresponding signal-only map, and fit the overall amplitude of the prediction to the stack difference, using the data stack covariance matrix described in Section 4.8. This amplitude indicates the amount of attenuation (shown as solid

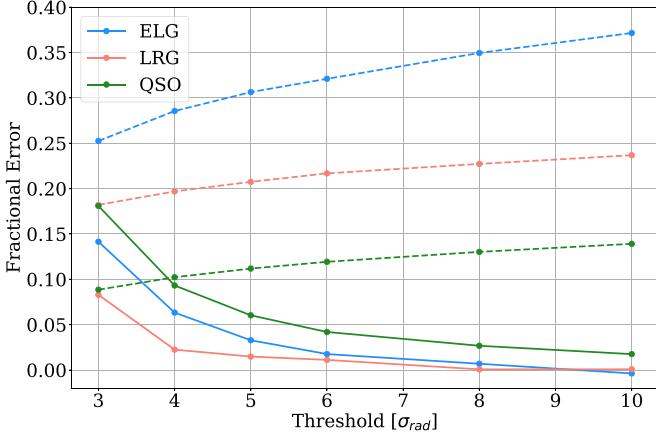


Figure 26. Comparison of the statistical uncertainty on the amplitude of the stacked signal (dashed line) to the bias in the amplitude caused by application of the outlier mask (solid lines) as a function of the threshold used to generate the mask. The threshold is defined in units of the standard deviation of the radiometric noise. The different colors correspond to different tracers of LSS. These measurements were made using the signal injection technique described in the text, wherein we stack on the sum of the data and the fiducial simulation for the 21 cm signal. Increasing the threshold reduces the bias in the recovered 21 cm amplitude but also increases the statistical uncertainty due to inclusion of residual foregrounds. A threshold of six times the expected radiometric noise was chosen for this analysis, which results in a bias in the amplitude that is <4% and small relative to the statistical error for all tracers.

lines in Figure 26) induced by the outlier mask, while the fractional uncertainty on this amplitude (shown as dashed lines) indicates the effect of the mask threshold on the statistical significance of the stacking measurement.

Based on these results, we choose a mask threshold of 6σ , where σ is the estimated radiometric noise in the maps (see Section 3.1.4). For the fiducial 21 cm model assumed in our simulations, this results in signal attenuation of less than 4% for each eBOSS tracer, which is at least a factor of three smaller than the statistical uncertainty. Note that in our actual fits to data, presented in Section 6, the stacking amplitudes are factors of (1.9, 1.5, 1.4), for the (ELG, LRG, QSO) stacks, greater than in our simulations. We have rerun the test in Figure 26, modifying the amplitude of the injected signal accordingly, and have verified that the attenuation level is unchanged, while the fractional uncertainty decreases by the quoted factors. Even with this change, the attenuation is still less than half of the uncertainty for each tracer, which we deem to be acceptable for this analysis.

8. Discussion

8.1. Quasar Redshift Errors

As illustrated in Figure 23, there is a statistically significant frequency offset in the QSO stacks of $\Delta\nu_0 \approx -0.2$ MHz, equal to roughly half the width of a CHIME frequency channel. As this is only seen in the QSO stacks and not within the overlapping LRG and ELG measurements, it is difficult to explain this as an instrumental issue within CHIME. Instead, we interpret this as being a systematic bias in the eBOSS QSO redshifts, stemming from the difficulty of determining a redshift from the complex processes producing a QSO spectrum (see Lyke et al. 2020). QSO emission lines such as C IV are frequently blueshifted from the host galaxy redshift by dynamical and radiative processes within the QSO's accretion

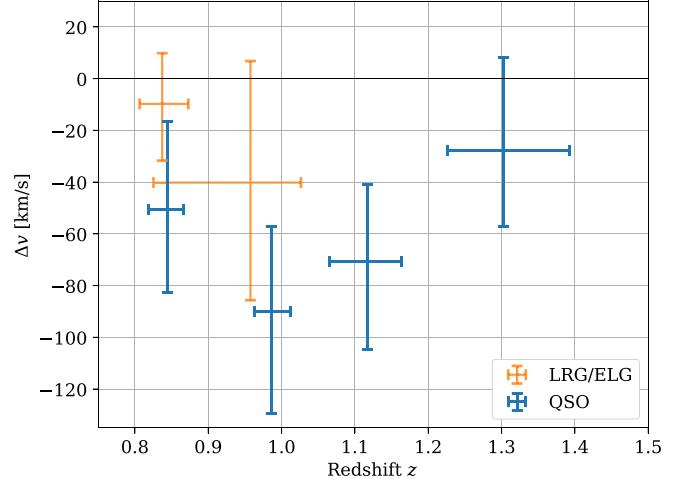


Figure 27. The derived redshift bias for each tracer given as a velocity shift $\Delta v = c\Delta z/(1+z)$. We derived this from the frequency offset $\Delta\nu_0$ measured from each tracer assuming that the source is a systematic bias in the eBOSS catalog redshifts. For both the ELG and LRG catalogs (orange points) there is no discernible bias, but the QSO catalogs (blue points; from left to right, QSOb00, QSOb01, QSOb1, and QSOb2) have a significant bias. This is in agreement with the bias of the z_{PCA} redshift estimates shown in Figure 3 of Lyke et al. (2020).

disk and outflowing winds (Richards et al. 2011; Shen et al. 2016).

Similar to Lyke et al. (2020), we express the redshift error as a velocity that can be connected to our measured frequency offset

$$\Delta v = c \frac{\Delta z}{1+z} \quad (130)$$

$$= c(1+z) \frac{\Delta\nu_0}{\nu_{21}}. \quad (131)$$

In Figure 27, we show the inferred velocity bias for the QSOb00, QSOb01, QSOb1, and QSOb2 stacks, which give nonoverlapping measurements in redshift. Overall, we measure $\Delta v \sim -66$ km s⁻¹ at $\sim 3.3\sigma$, with individual bins ranging from 0.8σ (QSOb2) to 2.5σ (QSOb1). Our analysis does not account for the Doppler shift from Earth's motion around the solar system barycenter; however, while the shift on any individual source may be up to ~ 30 km s⁻¹, on average this effect is small. Taking a weighted mean of the Doppler shift toward each source for each night of observation, we find an average Doppler correction of -3.1 km s⁻¹.

Overall the results in Figure 27 are consistent with those of Lyke et al. (2020, Figure 3), who estimated the systematic bias in the z_{PCA} redshift estimates (which we used for stacking) as compared to redshifts of QSO host galaxies measured using stellar absorption lines. We anticipate that future QSO cross-correlation analyses with higher source numbers and improved processing of the CHIME data will be able to provide useful measurements of this bias across a broad range of redshifts.

The eBOSS QSO redshifts are significantly noisier than those of LRG and ELG samples owing to the broader emission lines, with significant long tails of poor redshift estimates (Lyke et al. 2020). This has a noticeable effect on the stack signal (recall Figure 17), as the convolutional effect of the redshift errors broadens and suppresses the peak of the stack signal.

Uncertainties in the QSO redshift error model can therefore give sizable changes in the constraints on the signal amplitude.

In our primary analysis we use the “double-Gaussian” model of Lyke et al. (2020, Equation (A2)) to describe the QSO redshift uncertainties. However, the model as presented does not seem to match their measurements of the redshift errors in ways that are significant for our analysis. There are two clear differences: first, the fraction of observations that have errors drawn from the wider Gaussian component appears to be smaller in the data (as shown in Figure 4 of Lyke et al. 2020) than the $\sim 18\%$ quoted in the model; second, there appears to be a significant reduction in the errors at low redshift compared to the rest of the sample (as shown in Figure 9 of Lyke et al. 2020), which is expected from the presence of O III and H β in the wavelength range of the spectrograph at redshifts $z \lesssim 1$ (Étienne Burtin, private communication).

To assess the importance of this, we modify the z_{PCA} redshift error model to capture these effects. This change is intended to give a plausible alternative consistent with the data presented in Lyke et al. (2020), though we do not claim that it is more realistic. Producing an improved model would require repeating the analysis of Lyke et al. (2020) and is beyond the scope of this paper. Our model is a straightforward modification of the published “double Gaussian” where we allow the coefficients to be redshift dependent. The redshift error on a single observation of a QSO, as given by a velocity error δv , is drawn from a redshift probability distribution

$$\mathcal{P}(\delta v|z) = \frac{1}{1 + f^{-1}(z)} \times [\mathcal{G}(\delta v, \sigma_1(z)^2) + f^{-1}(z) \mathcal{G}(\delta v, \sigma_2(z)^2)], \quad (132)$$

where \mathcal{G} is the standard normalized Gaussian with

$$\mathcal{G}(x, \sigma^2) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{x^2}{2\sigma^2}}, \quad (133)$$

and $\sigma_1(z)$ and $f(z)$ are smooth step functions centered at $z = 1.0$ with σ_1 transitioning from 90 to 150 km s $^{-1}$,

$$\sigma_1(z) = \left[90 + 30 \left(1 + \tanh \left(\frac{z - 1.0}{0.05} \right) \right) \right] \text{ km s}^{-1}, \quad (134)$$

and $f^{-1}(z)$ changing from zero at low redshift (that is, no errors in the broad distribution) to ~ 0.03 (a value that approximately matches the data of Lyke et al. 2020, Figure 4),

$$f^{-1}(z) = \frac{1}{35} \left(1 + \tanh \left(\frac{z - 1.0}{0.05} \right) \right). \quad (135)$$

It is necessary to change both σ_1 and f , as no single change is able to reproduce the observed low-redshift uncertainties. However, we do leave σ_2 unchanged with a redshift-independent $\sigma_2(z) = 1000$ km s $^{-1}$.

In Figure 28 we show the change in our constraints that occur if we switch to this modified model. The top panel compares the redshift dependence of our new model to the Lyke et al. (2020) model and the measured uncertainties in their Figure 9. The bottom panel shows the change in the inferred amplitude, \mathcal{A}_{HI} , between the two models. We have fixed the nonlinear parameters in these constraints, which gives an indication of the statistical error on our constraints. At all redshifts the difference between the published model and our alternative is larger than the statistical uncertainty on \mathcal{A}_{HI} and

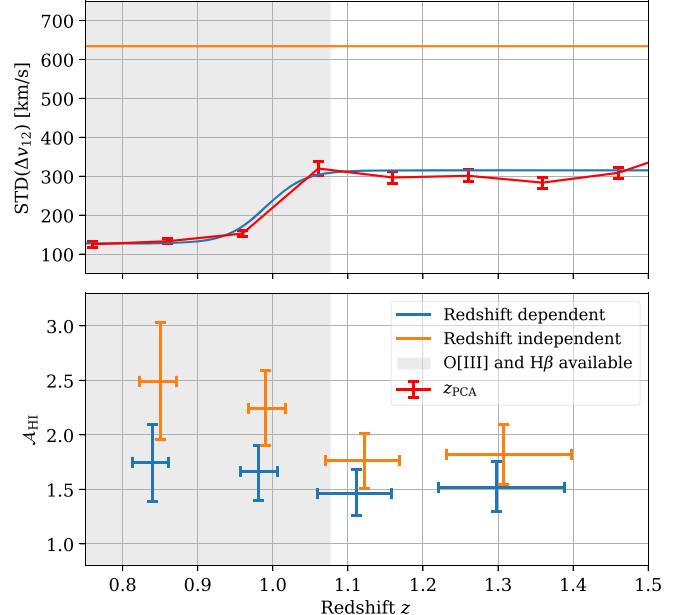


Figure 28. The QSO amplitude constraints are strongly dependent on the model for the QSO redshift errors. In the top panel we show the statistical error on the QSO z_{PCA} redshift estimates given by differences between repeated measurements of the same source. These estimates improve at $z \lesssim 1.1$ (gray shaded region) owing to the availability of the [O III] and H β lines. The measured distribution (red line) is taken from Lyke et al. (2020, Figure 9), the orange line shows the standard deviation of the published redshift-independent error model (Lyke et al. 2020 Equation (A1)), and the blue line is a redshift-dependent model described in the text that gives a plausible fit to the z_{PCA} errors. In the bottom panel we show the amplitude constraints (with fixed nonlinear parameters) for assuming each of these redshift error models. The shifts are significant and are redshift dependent. This suggests that the modeling of the QSO redshift errors is a larger source of uncertainty than the statistical error in our measurements.

suggests that redshift error distribution is a significant source of systematic uncertainty in our analysis. Future analyses will need to resolve the questions in this modeling to make precision constraints on \mathcal{A}_{HI} .

We use the differences observed in Figure 28 to estimate a systematic uncertainty from the redshift error modeling of $|\mathcal{A}_{\text{HI}}^{\text{alt}} - \mathcal{A}_{\text{HI}}^{\text{norm}}|/\sqrt{2}$, where the $\sqrt{2}$ comes from an argument that the models considered are samples from some distribution of plausible models.

8.2. \mathcal{A}_{HI} Constraints and Sources of Error

We are interested in learning about the amplitude of fluctuations in the HI distribution, which is probed most effectively by the parameter \mathcal{A}_{HI} (see Equation (104)) in our analysis. In Section 6.3 we discuss the constraints on \mathcal{A}_{HI} in the case where we allow the full set of parameters to vary and where we pin the nonlinear parameters to their fiducial values.

The uncertainty in the case with fixed nonlinear parameters is dominated by the statistical uncertainty in the data and from the prior on the galactic bias. We assume that this statistical contribution is the same in the case where we allow the nonlinear parameters to vary, with the weaker constraints coming from modeling uncertainties. In this case, and assuming that the modeling errors are multiplicative within the degenerate regions of parameter space, we can roughly separate the uncertainty in the full parameter constraints into statistical and modeling contributions.

Table 7
Sources of Uncertainty

Tracer	Fractional Errors (%)								
	Statistical	Modeling	Flux	Template	Consistency	Beam	Linearity	Redshift Errors	Total
LRG	14	150	4	1	0	8	1	0	151
ELG	18	93	4	0	0	8	2	0	95
QSO	10	49	4	0	0	8	4	14	52
QSOb0	13	73	4	0	25	8	2	24	82
QSOb1	14	76	4	0	0	8	4	13	79
QSOb2	15	73	4	0	0	8	5	13	76
QSOb00	21	191	4	0	0	8	1	25	194
QSOb01	15	76	4	0	25	8	2	21	85

Note. In this table we quantify the sources of error in our measurement. From left to right the sources are as follows: *Statistical*, inferred from the constraints with fixed nonlinear parameters; *Modeling* is the symmetrized error from the constraints varying all parameters, after removing the statistical contribution; *Flux* is from uncertainty in the absolute flux scale (Section 3.1); *Template* is from errors in the template calculation (Section 5.4 and Appendix D); *Consistency* gives systematic errors inferred from issues observed in data validation (Section 7.2); *Beam* lists the uncertainties from an imperfect beam model (Section 7.3); *Linearity* gives a systematic error to incorporate the effect of signal loss during our analysis that is not fully captured by our template calculation (Section 7.4); and *Redshift Errors* adds a systematic error to account for the difference in inferred amplitudes across plausible alternatives to the QSO redshift error model (Section 8.1). The final column, *Total*, combines the extra sources of systematic error with those from the full parameter constraints to give an estimate of the symmetrized fractional error.

There are many potential additional sources of systematic errors in our analysis that have been discussed beyond the modeling uncertainty. These are listed, along with the statistical and modeling uncertainty breakdown, in Table 7. This error budget is dominated by the modeling uncertainties; however, both the systematic error added to cover unexpected validation failures (labeled “Consistency” in Table 7 and discussed in Section 7.2) and the error from uncertainties in the QSO redshift error model (Section 8.1) are larger than the statistical error and thus could be the limiting sources if the modeling of nonlinear scales can be improved.

In Table 8, we summarize the constraints on \mathcal{A}_{HI} for all the tracer catalogs. We show the case both with the nonlinear parameters fixed and when the full set is allowed to vary, illustrating again the substantial increase in the uncertainties from these parameters. We also give a final case including the total error budget from all the systematic contributions above (we have assumed that they are all multiplicative effects). As the modeling uncertainties are dominant, including these extra sources of error gives only marginal increases to the total uncertainty. The most severely affected catalogs are QSOb0 and QSOb01, due to the systematic error contributions from both issues in the QSO redshift error model (which is worse at low redshifts) and from the consistency test failures.

8.3. Ω_{HI} Comparisons

To be able to compare our results directly to measurements of Ω_{HI} from other probes, we need to be able to break the degeneracy between Ω_{HI} and b_{HI} . Although our measurements are unable to do this internally and there are no external measurements of b_{HI} , we can use simulations as a guide.

As an indicator of the uncertainty on the bias, we use the bias measured at $z=1$ from various simulations. Villaescusa-Navarro et al. (2018) use the IllustrisTNG hydrodynamic simulation and find that $b_{\text{HI}}(z=1) \approx 1.49$, and Ando et al. (2019) use another hydrodynamic simulation, the Osaka simulation, to find that $b_{\text{HI}}(z=1) \approx 1.26$ (from their b_0 measurements). Another approach uses semianalytic prescriptions on top of dark-matter-only simulations, such as Spinelli

Table 8
Parameter Constraints for Each Tracer

Tracer	z_{eff}	\mathcal{A}_{HI}			
		Fiducial	Fixed NL	Full NL	
LRG	0.84	1.13	$1.82^{+0.26}_{-0.25}$	$1.51^{+3.60}_{-0.96}$	$1.51^{+3.60}_{-0.97}$
ELG	0.96	1.21	$2.35^{+0.43}_{-0.42}$	$6.76^{+9.01}_{-3.74}$	$6.76^{+9.04}_{-3.79}$
QSO	1.20	1.37	$1.86^{+0.18}_{-0.17}$	$1.68^{+1.06}_{-0.60}$	$1.68^{+1.10}_{-0.67}$
QSOb0	0.97	1.22	$2.27^{+0.31}_{-0.28}$	$2.04^{+2.09}_{-0.94}$	$2.04^{+2.21}_{-1.19}$
QSOb1	1.12	1.31	$1.75^{+0.25}_{-0.25}$	$2.89^{+3.13}_{-1.36}$	$2.89^{+3.17}_{-1.44}$
QSOb2	1.30	1.43	$1.81^{+0.27}_{-0.28}$	$1.63^{+1.55}_{-0.86}$	$1.63^{+1.57}_{-0.90}$
QSOb00	0.84	1.14	$2.49^{+0.52}_{-0.54}$	$1.49^{+4.06}_{-1.65}$	$1.49^{+4.08}_{-1.69}$
QSOb01	0.99	1.23	$2.23^{+0.35}_{-0.34}$	$3.23^{+3.47}_{-1.56}$	$3.23^{+3.64}_{-1.91}$

Note. After reparameterization to avoid degeneracies, the physically interesting parameter is the 21 cm amplitude \mathcal{A}_{HI} . We show the highest posterior density 68% credible intervals for these parameters for both a prior with the nonlinear parameters *fixed* and the *full* parameter space. Comparing the \mathcal{A}_{HI} constraints for the cases of fixed and varying nonlinear parameters, we can see that there is a substantial increase in the uncertainty from modeling the small-scale structure. We also show estimates for \mathcal{A}_{HI} including the effects of the systematic errors listed in Table 7. As the modeling errors are large, the additional uncertainty from this is small.

et al. (2020), who find $b_{\text{HI}}(z=1) \approx 1.22$ or 1.31 (depending on whether the Millennium I or II simulation is used), or Wang et al. (2021), who use an empirically calibrated star formation model to find $b_{\text{HI}}(z=1) \approx 1.27$. Collectively these prescriptions have a mean of ≈ 1.3 and a standard deviation of ≈ 0.1 . With this in mind, we place a simulation-derived Gaussian prior on the bias with a conservative width of 20%, i.e., $\sigma_{b_{\text{HI}}}/b_{\text{HI}}^{\text{fid}} = 0.2$.

We reweight the MCMC chains from our analysis to apply the updated prior on b_{HI} and marginalize over all the other parameters to derive constraints on Ω_{HI} . We give our measurements as the highest posterior density credible interval about the mode of the distribution. In Figure 29 we show the measurements for the LRG and ELG samples, as well as the

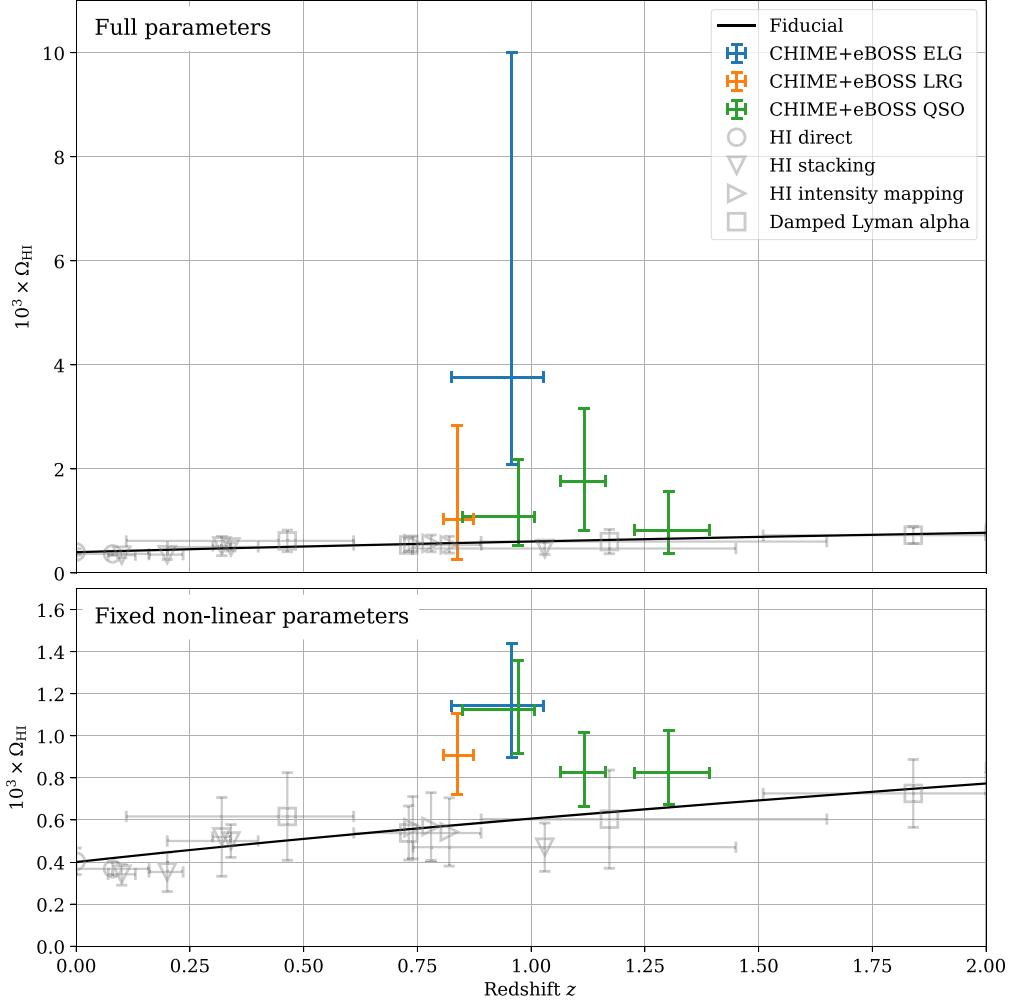


Figure 29. Constraints on Ω_{HI} from this analysis compared to other experiments. In the top panel we show the constraints in this work when varying the full set of modeling parameters, and in the bottom panel we fix the nonlinear parameters, which considerably reduced the uncertainties at the expense of hidden systematic errors. We have selected a representative sample of measurements using independent data sets to place in this figure. The data sets are of four types: at the lowest redshift there are direct 21 cm observations, such as those from ALFALFA (Jones et al. 2018) and the Arecibo Ultra Deep Survey (Xi et al. 2020); at intermediate redshifts source stacking of individual galaxies such as Rhee et al. (2013), who use Westerbork data and low-redshift galaxies observed with CFHT-MOS, and three studies combining GMRT radio data with different optical catalogs, VVDS optical data taken at VIMOS (Rhee et al. 2018), DEEP2 and DEEP3 at low redshift (Bera et al. 2019), and DEEP2 at high redshifts (Chowdhury et al. 2020); also at intermediate redshifts are H I intensity mapping cross-correlations like Wolz et al. (2022), who cross-correlate GBT intensity mapping data against eBOSS and WiggleZ catalogs; at the highest redshifts, measurements are from surveys of damped Ly α systems such as using Hubble Space Telescope ACS and GALEX data (Rao et al. 2017) and using ESO UVES (Zafar et al. 2013).

QSOs split across three redshift bins, compared to measurements from other experiments.

There are four main methods for measuring Ω_{HI} that we include in Figure 29 for comparison:

Direct H I surveys: At the lowest redshifts, blind surveys of the 21 cm line can measure the H I mass function directly, which can be integrated to obtain estimates of Ω_{HI} .

H I stacking: At intermediate redshifts, it is difficult to detect individual galaxies in their 21 cm emission; to get around this, high-resolution radio data can be stacked on the positions of galaxies found in optical catalogs to get an estimate of the average amount of H I per galaxy in the sample. This can then be combined with an optical luminosity function for the sample and corrected for completeness to give an estimate of Ω_{HI} .

H I intensity mapping: Another method is to cross-correlate H I intensity mapping data with optical catalogs. These are distinct from the H I stacking measurements described above

in that they do not resolve the emission in the (average of) individual galaxies, but instead are sensitive to the correlated H I mass in the vicinity of the galaxy. Our results are an example of this technique.

Damped Ly α : At the highest redshifts damped Ly α systems are detected in optical and UV QSO spectra, and the distribution of their observed column densities can be integrated to find Ω_{HI} .

For all measurements, we convert into the Planck 2018 cosmology used in this paper, using $H(z)^2 = H_0^2 [\Omega_m(1+z)^3 + 1 - \Omega_m]$, with $\Omega_m = 0.30964$. In each case, the measurements are effectively a fluxlike quantity, which is multiplied by an area to give an H I mass, divided by a survey volume to give a density, and then divided by the critical density to give Ω_{HI} , though some of these steps are implicit (this is still true for the damped Ly α analysis, though the “area” in the mass is canceled with the one implicit in the volume). If we approximate the observations as coming from a narrow

band in redshift, then the cosmology dependence is

$$\Omega_{\mathrm{H}\,1}(z) \propto \frac{1}{H_0^2} \frac{dz}{dr}, \quad (136)$$

where r is the comoving distance to redshift z . For the Wolz et al. (2022) intensity mapping points we convert their $\Omega_{\mathrm{H}\,1} b_{\mathrm{H}\,1} r$ measurements into $\Omega_{\mathrm{H}\,1}$ constraints with the fiducial bias model we use in this paper to allow a consistent comparison.

In Figure 29 for the constraints both when varying the nonlinear parameters (top panel) and when fixing them (bottom panel) our results are in broad agreement with other $\Omega_{\mathrm{H}\,1}$ constraints. As we would expect, the uncertainties are much larger when allowing the nonlinear parameters to vary, though the distributions are non-Gaussian and the probability of $\Omega_{\mathrm{H}\,1} \leq 0$ is still negligible. We note that while we expect all points to be pushed toward higher values of $\Omega_{\mathrm{H}\,1}$ by the prior volume effect in the FoG parameters discussed in Section 6.3, the CHIME+eBOSS ELG point is noticeably discrepant when the nonlinear parameters are varied. We believe that this is a chance fluctuation where the region further along the $\alpha_{\mathrm{H}\,1} - \alpha_{\mathrm{FoG},+}$ degeneracy is preferred and leads to a $\sim 2\sigma$ shift from the $\Omega_{\mathrm{H}\,1}$ values preferred by the other tracers. This can also be seen clearly in Figure 24, where the preferred range of $\alpha_{\mathrm{FoG},+}$ values is higher than in the QSO and LRG cases (Figures 23 and 25, respectively). When fixing the nonlinear parameters, the ELG constraints are much more consistent with both the other CHIME tracers and the external data sets.

As the constraints with fixed nonlinear parameters do not include the full modeling uncertainties, they show the internal consistency and significance of our measurements but are not good indicators of the plausible range of $\Omega_{\mathrm{H}\,1}$ determined from our data. In all cases we are showing constraints derived with the fiducial eBOSS QSO error model. As discussed in Section 8.1, we believe that this model may bias the QSO constraints (particularly the lowest-redshift bin) to higher values of $\Omega_{\mathrm{H}\,1}$.

8.4. Atomic Hydrogen Content of Galaxies and Quasars

As mentioned previously, our stacking analysis not only probes the correlated clustering of eBOSS catalog objects and H I but also is sensitive to the H I associated with the objects themselves, which sets the value of our M_{10} parameter for each sample (recall that M_{10} is defined as the mean H I mass per catalog object, in units of $10^{10} M_{\odot}$). Figures 23–25 show that the posteriors for M_{10} peak at nonzero values for the QSO and LRG stacks, while for the ELG stack the posterior peaks at $M_{10} = 0$. In each case, however, the model where M_{10} and the other nonlinear parameters ($\alpha_{\mathrm{FoG},+}$ and α_{NL}) are allowed to vary is not strongly preferred over the case where these parameters are fixed to their fiducial values (see Section 6.4); thus, we cannot interpret the posteriors of M_{10} as providing definitive information about the H I content of the objects in each catalog.

Nevertheless, the finite width of these posteriors indicates that future analyses may hold the promise of interesting constraints. In particular, for the ELG stack, the highest posterior density 68% credible interval is $M_{10} < 1.04$. This is consistent with the simulations of Wolz et al. (2022), which were based on the DARK SAGE semianalytical galaxy evolution model (Stevens et al. 2016) and predicted a shot-noise contribution to the H I–ELG cross-power spectrum equivalent to $M_{10} \approx 0.8$ (as inferred from their Figure 12). It is also consistent with the analysis of Chowdhury et al. (2020), who

stacked GMRT 21 cm observations on star-forming galaxies from the DEEP2 survey and found $M_{10} = 1.19 \pm 0.26$ at an effective redshift $z_{\mathrm{eff}} = 1.03$. A cross-correlation analysis with greater power to break the parameter degeneracies in our model would likely improve the constraint on M_{10} to a level where it could fruitfully be compared with these other values.

Empirical information on the H I content of LRGs at $z \sim 1$ is scarce: direct stacking analogous to Chowdhury et al. (2020) has only been carried out at lower redshifts for such red galaxies (e.g., Rhee et al. 2018). Thus, constraints on M_{10} for LRGs (and QSOs) would provide valuable information about the evolution and environments of these objects. On the other hand, inclusion of an external prior on M_{10} , obtainable from, for example, stacking GMRT observations on a subset of objects from each eBOSS catalog, would help to break the degeneracies in our model (or other, more detailed models of H I–galaxy cross-correlations), and we see this as a promising avenue for future investigation.

9. Conclusions

In this paper, we have presented the first detection of cosmological 21 cm emission with the CHIME telescope. This detection is the result of constructing sky maps from CHIME data, filtering and cleaning these maps in various ways, and performing a cross-correlation analysis with catalogs of LRG, ELG, and QSO positions from the eBOSS survey. We have described several aspects of CHIME data processing that have not previously appeared in the literature: these include our procedures for combining multiple sidereal days of observations (Section 3.3), forming beam-deconvolved sky maps from measured visibilities (Section 4.3), measuring delay power spectra using Gibbs sampling (Appendix A), and inferring the primary beam pattern based on external measurements of many radio point sources (Appendix B).

We have filtered bright foregrounds out of the measurements with a high-pass delay filter using the approach of Ewall-Wice et al. (2021), with a decl.-dependent delay cutoff that selects the regime where the fluctuations in the data are close to the expected noise level. This filtering has the effect of removing any sensitivity to linear cosmological scales related to BAOs, such that the signal-to-noise ratio is concentrated at nonlinear scales ($0.3 h \mathrm{Mpc}^{-1} \lesssim k \lesssim \mathrm{few} h \mathrm{Mpc}^{-1}$; see Figure 14).

We perform the cross-correlation by separately stacking CHIME sky maps at the angular and spectral locations of the objects in eBOSS catalogs of ELGs, LRGs, and QSOs. In each case, the spatial extent of the signal is consistent with an unresolved point source (Figure 19), so we present our main results as 1D stacking profiles as a function of frequency offset from the locations of the catalog objects (Figure 20). We achieve significant detections for each catalog, as indicated by Bayes factors Z_1/Z_0 of $\ln(Z_1/Z_0) \approx 18.8$ (LRGs), 10.8 (ELGs), and 56.3 (QSOs), computed by comparing our signal model with a noise-only model; alternatively, a frequentist likelihood ratio test gives signal-to-noise ratios of 7.1 (LRGs), 5.7 (ELGs), and 11.1 (QSOs).

We interpret these measurements using a simulation-based framework (Sections 5.3 and 5.4), within a model that considers H I and galaxies to be linearly biased tracers of the underlying matter distribution, including the leading effects of redshift-space distortions and a correlated shot-noise contribution related to the mean H I mass of the objects in each catalog (Section 5.2). We are able to constrain an effective H I

clustering amplitude $\mathcal{A}_{\text{HI}} \equiv 10^3 \Omega_{\text{HI}} (b_{\text{HI}} + \langle f\mu^2 \rangle)$, where Ω_{HI} is the cosmic abundance of HI, b_{HI} is the linear bias of HI, and $\langle f\mu^2 \rangle$ (equal to 0.552 in this analysis) is an average over the linear growth rate f and an angular factor μ^2 related to the line-of-sight components of the Fourier modes probed in the stacks (Section 6.2). We constrain this amplitude separately for each eBOSS catalog, marginalizing over parameters controlling the scale dependence of nonlinear clustering, obtaining $\mathcal{A}_{\text{HI}} = 1.51^{+3.60}_{-0.97}$ (LRGs), $\mathcal{A}_{\text{HI}} = 6.76^{+9.04}_{-3.79}$ (ELGs), and $\mathcal{A}_{\text{HI}} = 1.68^{+1.10}_{-0.67}$ (QSOs) (see Table 8). Previous cross-correlations between GBT 21 cm maps and galaxy catalogs have measured $\Omega_{\text{HI}} b_{\text{HI}} r$ (where r is a phenomenological cross-correlation parameter) with 15%–25% precision (Chang et al. 2010; Masui et al. 2013; Wolz et al. 2022); our constraints on \mathcal{A}_{HI} are weaker than this, but only due to our more detailed modeling of small-scale clustering, which requires marginalization over several parameters.

We also constrain an overall frequency offset $\Delta\nu$ of the stacking profile. This offset is consistent with zero for ELGs and LRGs, while for QSOs we find $\Delta\nu \approx -0.2$ MHz. We interpret this as a systematic bias in the measured redshifts of the QSOs, corresponding to $\Delta\nu \approx -66$ km s $^{-1}$ in velocity units. As discussed in Section 8.1, this is consistent with what was found by the eBOSS team in Lyke et al. (2020).

Our results point to several interesting directions for future investigation. Our present analysis only considered CHIME frequencies above 585 MHz, corresponding to redshifts less than 1.42, but the eBOSS QSO catalog contains a significant number of QSOs at higher redshift (see Figure 2), and it would be worthwhile to repeat the stacking procedure using these objects, after additional effort to remove transient RFI in CHIME data at the relevant frequencies. In addition, similar future analyses have the potential to constrain the mean HI mass per catalog object. This would provide opportunities for coordination with stacking analyses from higher-resolution interferometers like GMRT (e.g., Chowdhury et al. 2020), which could help to disentangle the contributions from LSS and correlated shot noise and also provide new information about the evolution and properties of LRG, ELG, and QSO samples. In parallel, future cross-correlation analyses could be used to obtain more detailed information about systematic errors in spectroscopic redshifts obtained from optical instruments.

More broadly, many of the methods developed for this analysis are not specific to CHIME but could also be applied to other low-redshift interferometric 21 cm surveys, such as CHORD (Vanderlinde et al. 2019), Tianlai (Li et al. 2020; Wu et al. 2021), HIRAX (Crichton et al. 2022), uGMRT (Chakraborty et al. 2021), and the Ooty Wide Field Array (Subrahmanya et al. 2017), as well as higher-redshift surveys like HERA (DeBoer et al. 2017) and potential future projects (Cosmic Visions 21 cm Collaboration et al. 2018).

Finally, we note that this paper has made use of only a small fraction of the total amount of data collected by CHIME in the past 4 yr. Future improvements in data processing will be focused not only on enabling much more detailed cross-correlation measurements but also on the ultimate goal of measuring BAOs in the auto-power spectrum of 21 cm emission, providing important clues as to the nature of dark energy and the properties of the low-redshift universe.

We thank Étienne Burtin for useful discussions.

We thank the Dominion Radio Astrophysical Observatory, operated by the National Research Council Canada, for

gracious hospitality and expertise. The DRAO is situated on the traditional, ancestral, and unceded territory of the Syilx Okanagan people. We are fortunate to live and work on these lands.

CHIME is funded by grants from the Canada Foundation for Innovation (CFI) 2012 Leading Edge Fund (Project 31170) and the CFI 2015 Innovation Fund (Project 33213) and by contributions from the provinces of British Columbia, Québec, and Ontario. Long-term data storage and computational support for analysis is provided by WestGrid,²⁹ SciNet,³⁰ and Compute Canada,²⁹ and we thank their staff for flexibility and technical expertise that has been essential to this work, particularly Martin Siegert, Lixin Liu, and Lance Couture.

Additional support was provided by the University of British Columbia, McGill University, and the University of Toronto. CHIME also benefits from NSERC Discovery Grants to several researchers and funding from the Canadian Institute for Advanced Research (CIFAR) and from the Dunlap Institute for Astronomy and Astrophysics at the University of Toronto, which is funded through an endowment established by the David Dunlap family. This material is partly based on work supported by the NSF through grants 2008031, 2006911, and 2006548 and by the Perimeter Institute for Theoretical Physics, which in turn is supported by the Government of Canada through Industry Canada and by the Province of Ontario through the Ministry of Research and Innovation.

We thank the Sloan Digital Sky Survey and eBOSS collaborations for publicly releasing the LRG, ELG, and QSO catalogs and supporting mock catalogs used in this work. Funding for the Sloan Digital Sky Survey IV has been provided by the Alfred P. Sloan Foundation, the U.S. Department of Energy Office of Science, and the Participating Institutions. SDSS-IV acknowledges support and resources from the Center for High Performance Computing at the University of Utah. The SDSS website is www.sdss.org.

Software: bitshuffle (Masui et al. 2015), CAMB (Lewis et al. 2000), caput (Shaw et al. 2020c), ch_pipeline (Shaw et al. 2020b), cora (Shaw et al. 2020d), Cython (Behnel et al. 2011), draco (Shaw et al. 2020e), driftscan (Shaw et al. 2020a), emcee (Foreman-Mackey et al. 2013), GetDist (Lewis 2019), hankl (Karamanis & Beutler 2021), h5py (Collette et al. 2021), HDF5 (<http://portal.hdfgroup.org/display/HDF5/HDF5>), HEALPix (Gorski et al. 2005), healpy (Zonca et al. 2019), Matplotlib (Hunter 2007), mpi4py (Dalcin & Fang 2021), networkx (Hagberg et al. 2008), NumPy (Harris et al. 2020), OpenMPI (Gabriel et al. 2004), pandas (McKinney 2010; The pandas development team 2020), peewee (<https://github.com/coleifer/peewee>), SciPy (Virtanen et al. 2020), Skyfield (Rhodes 2019),

Appendix A

Delay Power Spectrum Estimation via Gibbs Sampling

Delay power spectra³¹ measure the power at different time lags observed within a frequency spectrum and are an

²⁹ <https://www.alliancecan.ca>

³⁰ <https://www.scinethpc.ca/>

³¹ For clarity, we will use *delay spectrum* to refer only to the direct Fourier transform of a frequency *intensity* or *flux* spectrum. The *delay power spectrum* will refer only to the variance of this quantity. Though the intensity and flux are both second-order statistics of the electric field and thus are power-like quantities in a physical sense, we do not think that this is ambiguous anywhere in this text.

extremely powerful tool for investigating instrumental effects and the frequency structure of radio emission from the sky (see Figure 10 for an example). Superficially estimating a delay power spectrum involves taking a Fourier transform of a frequency spectrum and estimating the resulting power in the time domain. However, in the presence of interference that causes certain frequencies to be masked out and a large dynamic range between the power at different delays, significant care must be taken to avoid mixing of power between different delays. There are several existing strategies for dealing with this, such as using a CLEAN-like algorithm in delay space (Parsons et al. 2014) and least-squares spectral analysis (LSSA; see Vaníček 1969; Trott et al. 2016).

To understand the challenges involved, consider a noisy observation f of a frequency spectrum with length N_f . This is related to an underlying delay spectrum d by

$$f = Fd + n, \quad (A1)$$

with noise n and where the delay spectrum is assumed to be drawn from the input delay power spectrum $D[\tau_a]$:

$$D[\tau_a] = \langle |d_a|^2 \rangle. \quad (A2)$$

The noise is described by covariance N , and for the moment we treat the noise as being uniform except for entirely missing frequencies that we give infinite noise. We write this as $N^{-1} = \sigma^{-2} M$, where M is a diagonal masking matrix with ones for included frequencies and zeros for missing frequencies. The matrix F , with $F_{ab} = e^{-2\pi j \tau_a \nu_b} / N_f^{1/2}$, is unitary and performs a discrete Fourier transform from the time (delay) domain to the conjugate frequency domain.

A first attempt to estimate the delay power spectrum might start by simply applying the mask M to the observed frequency spectrum and performing an inverse Fourier transform,

$$\hat{d}_{\text{inv}} = F^* M f. \quad (A3)$$

With this estimate of the delay spectrum, we can then infer the delay power spectrum $D[\tau_a]$ by using a variance over N_{obs} observations, indexed by i :

$$\hat{D}[\tau_a] = \frac{1}{N_{\text{obs}}} \sum_{i=1}^{N_{\text{obs}}} |\hat{d}_a^i|^2, \quad (A4)$$

with \hat{d}_a^i being any estimator for d_a such as \hat{d}_{inv} defined above (later we will introduce additional estimators). However, this procedure generates significant leakage between delay channels, with a delay spread function given by the matrix $F^* M F$. In the case of random masking of N_{masked} single frequencies, it can be shown that this gives leakage at the level of $\sim N_{\text{masked}} \sum_a D[\tau_a] / N_f^2$ uniformly across delays, and we are not able to see any structure in the delay power spectrum below this level.

To improve this, we could modify the delay spectrum estimate by deconvolving the delay spread function by its pseudo-inverse, or equivalently use a maximum likelihood estimator³²

$$\hat{d}_{\text{ml}} = (F^* N^{-1} F)^+ F^* N^{-1} f. \quad (A5)$$

³² We note that this is similar to LSSA, though LSSA considers more general cases, such as irregular sampling, and typically restricts the range of delays being solved for to minimize correlations and leakage.

However, as F is unitary, the pseudo-inverse $(F^* N^{-1} F)^+$ is equal to $F^* (N^{-1})^+ F$, and as the noise matrix is diagonal with zeros where samples are masked, $(N^{-1})^+ N^{-1} = M$; together, these imply that $\hat{d}_{\text{ml}} = \hat{d}_{\text{inv}}$. In words, the maximum likelihood estimator is exactly equivalent to inverse Fourier transforming the masked frequency spectra.

Another option is using a Wiener filter instead of a maximum-likelihood-type filter:

$$\hat{d}_w = (D^{-1} + F^* N^{-1} F)^- F^* N^{-1} f, \quad (A6)$$

where $D_{ab} = D[\tau_a] \delta_{ab}$ is the covariance matrix of the delay spectrum signal. By providing information about the distribution of power at various delays, the filter can distinguish delays related to true signal in the masked frequency spectra, resulting in delay spectra with significantly lower leakage and hence cleaner power spectra. However, constructing this requires that we already know the delay power spectrum $D[\tau_a]$, which is the quantity that we are trying to estimate. A close enough guess may minimize the leakage enough to produce accurate delay power spectrum estimates, but there is no knowing in advance whether this is the case.

A resolution to this is to jointly solve for both the delay spectrum and the delay power spectrum, a problem that is tractable by Gibbs sampling (Geman & Geman 1984), an MCMC technique for drawing samples from a joint distribution where the conditional distributions are easily sampled. In particular, we draw inspiration from techniques used for power spectrum estimation of the CMB (e.g., Eriksen et al. 2004; Wandelt et al. 2004).

We want to infer both the delay spectrum d and the delay power spectrum (equivalent to the diagonal matrix D) by drawing samples from the joint probability distribution $\mathcal{P}(d, D|f)$. Gibbs sampling allows us to do that by alternately drawing from the conditional distributions $\mathcal{P}(d|D, f)$ and $\mathcal{P}(D|d, f)$; the ensuing set of samples will eventually converge to the joint distribution, and we can take the mean over D samples as an estimate of the delay power spectrum. We now describe how to sample from each conditional distribution.

Starting with $\mathcal{P}(d|D, f)$, we can use Bayes's theorem to write

$$\mathcal{P}(d|D, f) \propto \mathcal{P}(f|d, D) \mathcal{P}(d|D). \quad (A7)$$

The first term on the right-hand side is the likelihood function for the frequency spectrum, which for Gaussian noise can be written as $\mathcal{P}(f|d, D) = \mathcal{G}_C(f - Fd, N)$, where \mathcal{G}_C is a circularly symmetric complex Gaussian distribution:

$$\mathcal{G}_C(z, C) = \frac{1}{|\pi C|} e^{-z^* C^{-1} z}. \quad (A8)$$

We will also model the conditional prior distribution for the delay spectrum as Gaussian, with $\mathcal{P}(d|D) = \mathcal{G}_C(d, D)$. Combining these together and grouping the terms in d , we find that the conditional distribution is

$$\mathcal{P}(d|D, f) = \mathcal{G}_C(d - \hat{d}_w, C), \quad (A9)$$

where $C^{-1} = D^{-1} + F^* N^{-1} F$. Thus, the mean of the conditional distribution is just the Wiener filter of Equation (A6), with the standard covariance. Although drawing from this can be done by solving for the mean, followed by inversion and factorization of C^{-1} to add a random fluctuation, it is more

efficiently done by constructing

$$\mathbf{C}^{-1} \mathbf{d} = \mathbf{F}^\dagger \mathbf{N}^{-1} \mathbf{f} + \mathbf{D}^{-1/2} \mathbf{w}_1 + \mathbf{F}^\dagger \mathbf{N}^{-1/2} \mathbf{w}_2, \quad (\text{A10})$$

where \mathbf{w}_1 and \mathbf{w}_2 are standard Gaussian random samples, and then solving for \mathbf{d} (Jewell et al. 2004).

The conditional distribution for the delay power spectrum is more straightforward. We wish to calculate the conditional distribution $\mathcal{P}(\mathbf{D}|\mathbf{d}, \mathbf{f})$, which is independent of \mathbf{f} , as all the information about \mathbf{D} is contained within \mathbf{d} . Using a flat prior on the elements of \mathbf{D} and the prior $\mathcal{P}(\mathbf{d}|\mathbf{D})$, we find that $\mathcal{P}(\mathbf{D}|\mathbf{d}, \mathbf{f}) \propto \mathcal{G}_C(\mathbf{d}, \mathbf{D})$. Assuming that \mathbf{D} is diagonal, we can rewrite this in terms of the sample variance estimates $\hat{D}[\tau_a]$ for each delay τ_a , which are sufficient statistics for the diagonal elements of \mathbf{D} itself, $D[\tau_a]$. The sample variance $\hat{D}[\tau_a]$ has a χ^2 distribution,

$$N_{\text{obs}} \frac{\hat{D}[\tau_a]}{D[\tau_a]} \sim \chi^2(N_{\text{obs}}), \quad (\text{A11})$$

and so we can draw samples from $\mathcal{P}(\mathbf{D}|\mathbf{d}) \propto \mathcal{P}(\hat{D}[\tau_a]|D[\tau_a])$ by drawing a standard χ^2 deviate for each delay $x_a \sim \chi^2(N_{\text{obs}})$ and setting the new sample for $D[\tau_a]$ to be $N_{\text{obs}} \hat{D}[\tau_a]/x_a$.

Our practical implementation of this algorithm is as follows:

1. Pick a set of data whose delay spectra are expected to be similar enough that we can average over them. For computing delay spectra from visibilities, this might consist of all R.A. samples for individual baselines (after stacking over redundant copies). For the map delay spectrum described in Section 4.5, we choose the set of

R.A. samples at each polarization and decl. As above, we use N_{obs} to denote the size of this set.

2. Apply an apodization window to each frequency spectrum, if desired. A Nuttall window is used in Section 4.5.
3. Choose an initial guess $D_0[\tau_a]$ for the delay power spectrum. We use a white spectrum with amplitude 10 Jy beam $^{-1}$ in this work.
4. Loop over the following steps until convergence has been achieved:
 - (a) For each element i of the set of N_{obs} spectra, draw the n th delay spectrum sample $\mathbf{d}_n^i \leftarrow \mathcal{P}(\mathbf{d}|\mathbf{f}, \mathbf{D}_n)$ using Equation (A10). Note that each \mathbf{d}_n^i is a delay spectrum with N_f elements.
 - (b) Draw the $(n+1)$ th delay power spectrum sample $D_{n+1}[\tau_a] \leftarrow \mathcal{P}(D[\tau_a]|\{\mathbf{d}_n^i\}_{i=1}^{N_{\text{obs}}})$, using the N_{obs} delay spectra drawn at step n to compute $\hat{D}[\tau_a]$ in Equation (A11).
5. Take the average of the converged samples, after removing burn-in and performing any necessary thinning. In Section 4.5, we halt after 100 samples and take the median over the final 50 samples as an estimate of the delay power spectrum.

In summary, the Gibbs sampling approach is a statistically well-motivated technique that iteratively deconvolves the delay spectra, uses them to update a delay power spectrum, and uses this to improve the next deconvolution round.

In Figure 30 we apply the various estimators discussed above to a synthetic data set with high dynamic range in delay

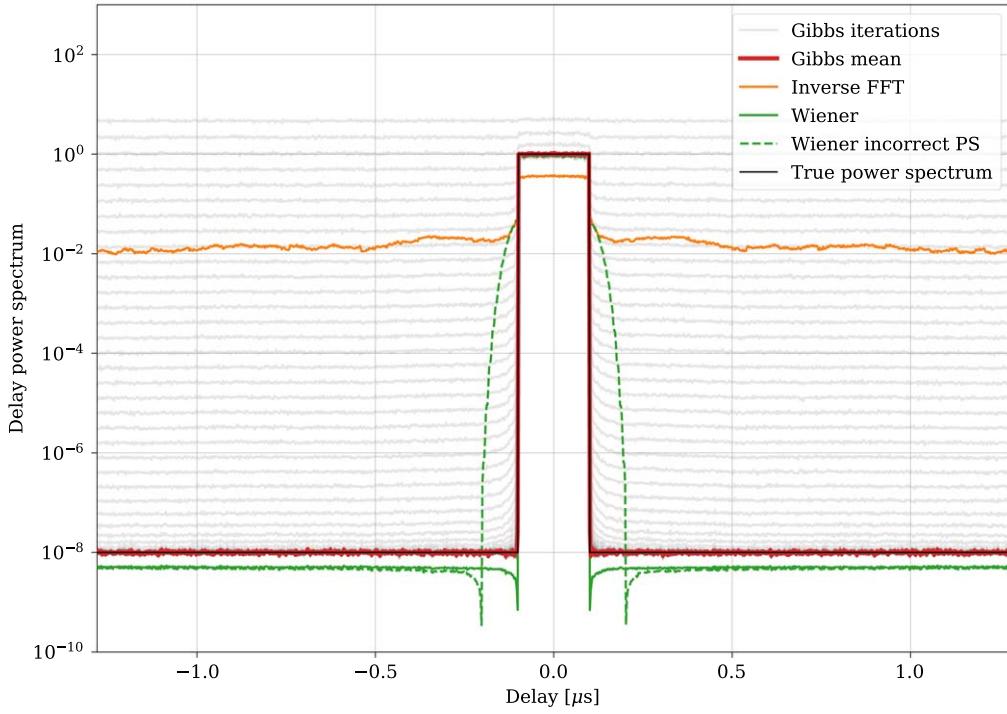


Figure 30. To test the performance of the delay power spectrum estimation techniques discussed in Appendix A, we generate a set of random delay spectra with a true delay power spectrum (black line) consisting of very high power at low delays and a plateau of low power outside this region. We Fourier transform these into frequency spectra and apply the CHIME RFI mask used in this analysis (Section 3.2.3). In orange we show the direct inverse estimate, which has significant leakage at the $\sim 10^{-2}$ level. The Wiener filter estimate (green solid) produces a much closer estimate, correcting most of the leakage effects, at the expense of needing a good starting guess. If generated with a poor initial delay power spectrum (high power over twice the range of delays as the true power spectrum), the estimate is significantly worse (green dashed). The Gibbs sampler produces an estimate that is much closer to the true power spectrum (red). Like any MCMC scheme, attention must be paid to the convergence of the chain. We used 100 samples and derived our estimate from the median of the last half. The full chain is shown in gray and can be seen converging from a poor starting guess of a flat delay power spectrum over ~ 30 samples.

space and a realistic frequency mask. We clearly see that the Gibbs-sampling-based estimator is able to accurately recover the input spectrum, while the naive inverse Fourier and Wiener estimators show various degrees of discrepancy.

Appendix B Estimating the Primary Beam by Deconvolving a Model for the Point-source Sky

In this appendix, we describe the algorithm that is used to directly reconstruct the average primary beam pattern of the CHIME antennas. We provide a schematic representation of this algorithm in Figure 31. First, a model for the radio emission from extragalactic point sources is constructed from measurements made by other telescopes. The *specfind v2* table (Vollmer et al. 2010) in the Vizier database is queried for flux measurements of all known sources between decl. -40° and 85° . For each source, all available measurements of the flux are fit to a power-law with frequency

$$s(\nu) = a \left(\frac{\nu}{600 \text{ MHz}} \right)^\gamma, \quad (\text{B1})$$

where the amplitude a and exponent γ are allowed to float. The fit is done by performing a weighted linear regression of the logarithm of the flux to the logarithm of the frequency. The uncertainties provided in the *specfind v2* table are used to construct inverse variance weights. These uncertainties are 20% of the measured flux (Vollmer et al. 2005), and the power-law model is in general a good fit given these large uncertainties. Only sources with $s(600 \text{ MHz}) > 15 \text{ mJy}$ that have at least one measurement on either side of the CHIME band are included in the sky model. There are 97,941 sources in total that meet these criteria. All of the sources have at least three flux measurements, with six flux measurements on average.

Our model for the visibility measured by baseline \mathbf{b} at frequency ν and local Earth rotation angle ϕ is then given by

$$S(\mathbf{b}, \nu, \phi) = \sum_i s_i(\nu) e^{j2\pi\nu\mathbf{b}\cdot\hat{\mathbf{n}}(\theta_i, \phi - \phi_i)/c} \delta(\phi, \phi_i), \quad (\text{B2})$$

where $s_i(\nu)$ is a power-law model for the flux of the i th source; $\hat{\mathbf{n}}'(\theta_i, \phi - \phi_i)$ is the unit vector pointing in the direction of the i th source and is given by Equation (25), with ϕ_i and θ_i denoting the source's R.A. and decl. in CIRS coordinates, respectively; and

$$\delta(\phi, \phi_i) = \begin{cases} 1 & |\phi - \phi_i| < \frac{1}{2}\Delta\phi, \\ 0 & \text{otherwise} \end{cases}, \quad (\text{B3})$$

with $\Delta\phi = 0^\circ.0879$ denoting the sample spacing of the data in local sidereal angle. The sum in Equation (B2) runs over all sources.

The following identical operations are then performed on the sidereal visibilities V and sky model S . We first arrange the baselines onto a 2D grid and then beamform in the $\hat{\mathbf{y}}$ -direction using Equation (27). The weights used in the beamformer are given by

$$w_{xy}^p(\nu, \phi) = W_\alpha \left(\frac{1}{2} \left[1 + \frac{\nu y}{\nu_{\min} y_{\max}} \right] \right), \quad (\text{B4})$$

where W_α denotes the Dolph–Chebyshev window, $\alpha = 60 \text{ dB}$ is the peak-to-sidelobe ratio, $\nu_{\min} = 587.5 \text{ MHz}$ is the minimum frequency examined, and $y_{\max} = 255$ corresponds to the maximum baseline distance in the $\hat{\mathbf{y}}$ -direction.

The window function in Equation (B4) will result in a frequency-independent synthesized beam in the $\hat{\theta}$ -direction that has an $\text{FWHM} = 0^\circ.385$ and sidelobes that are $\lesssim 10^{-3}$ of the peak amplitude. The Dolph–Chebyshev window minimizes the main lobe width for a given number of baselines and equiripple peak-to-sidelobe ratio. It will degrade the point-source sensitivity relative to the inverse variance weighting scheme discussed in Section 4.3; however, the loss of sensitivity is not problematic for beam calibration because it relies on a foreground signal that is $\gtrsim 500$ times brighter than the noise. The low equiripple sidelobes help to ensure that each formed beam is sensitive to the primary beam at a narrow range of decl.

The argument of the window function is scaled with frequency so that the synthesized beam in the $\hat{\theta}$ -direction is frequency independent. Essentially the resolution at every frequency is degraded to the resolution at the lowest frequency. This ensures that all frequencies are sensitive to the same decl., so that any errors in our sky model are not further modulated by a frequency-dependent synthesized beam pattern.

Next, we multiply the hybrid beamformed visibilities by a cosine-tapered window that is unity for $125^\circ < \phi < 255^\circ$ and transitions to zero over a span of 15° . This restricts our attention to a relatively quiet portion of the radio sky, avoiding sharp features in the Galactic emission that are present in the data but not in our model, and also avoiding regions of the sky contaminated by Cygnus A and Cassiopeia A in the sidelobes, which this technique is unable to account for properly. This range of ϕ also coincides with the range covered by the eBOSS NGC field. The m -mode transform is then taken.

The model for the primary beam is obtained by cross-correlating the sky model and the visibilities in m -mode space,

$$\tilde{B}_{xm}^p(\nu, \theta) = \frac{\tilde{S}_{xm}^*(\nu, \theta) \tilde{V}_{xm}^p(\nu, \theta)}{|\tilde{S}_{xm}(\nu, \theta)|^2 + \sigma_x^p(\nu)^2}, \quad (\text{B5})$$

where \tilde{S} , \tilde{V} , and \tilde{B} denote the m -mode transform of the hybrid beamformed visibilities for the sky model, data, and beam transfer function, respectively, and σ is an estimate of the noise in \tilde{V} .

The m -mode transform of the beam transfer function, \tilde{B} , is multiplied by a cosine-tapered mask to remove any m -modes that cannot originate from the sky near meridian. This mask is unity for $m_{\text{center},x} \pm 0.75m_{\text{width}}$ and then smoothly transitions to zero by $m_{\text{center},x} \pm m_{\text{width}}$ (see Equations (38) and (39)). The inverse m -mode transform is then calculated to obtain our estimate of the beam transfer function B for each EW baseline separation x . The beam transfer function is then “fringes-topped,” or in other words, it is multiplied by the complex conjugate of the exponential term in Equation (30), to recover $|A^p(\nu, \theta, \phi)|^2$, which we will refer to as the power beam. Note that this method yields four distinct estimates of the power beam, one for each EW baseline separation.

We find that the resulting estimate of the power beam exhibits small-scale variations along the decl. axis that are highly correlated as a function of frequency and hour angle. We suspect that these variations are due to errors in the flux of the sources in the sky model, and we remove them as follows. At each decl., the logarithm of the power beam at $\phi = 0^\circ$ is fit to a

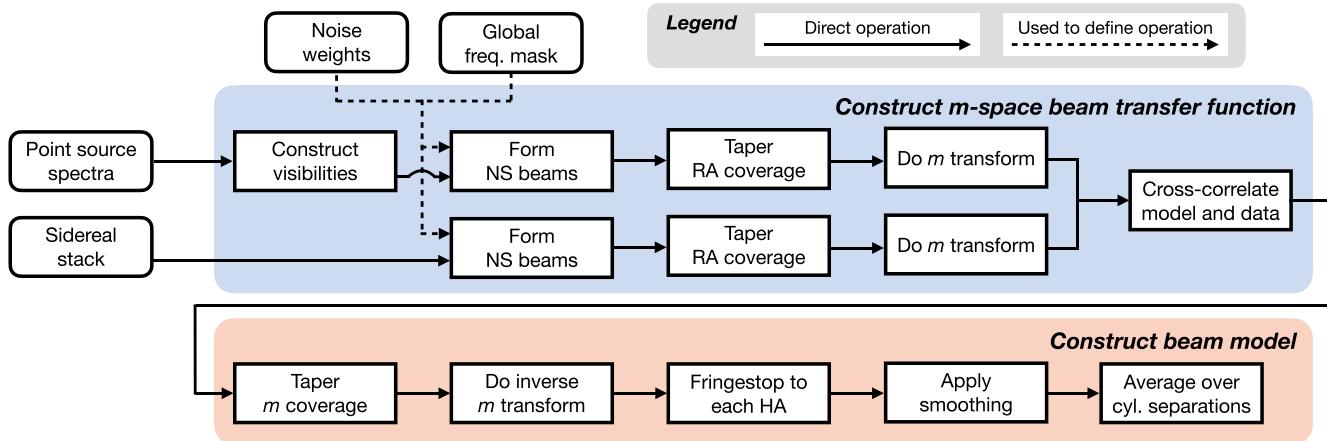


Figure 31. A schematic representation of the construction of the primary beam model used in our stacking analysis. Externally measured spectra of 97,941 radio point sources are propagated into mock visibilities, which, after several transformations, are cross-correlated with CHIME observations to construct a beam transfer function that assumes that the sky is solely composed of these sources. Further transformations are applied to minimize sensitivity to this assumption and remove artifacts.

fourth-order polynomial in frequency. This logarithmic polynomial model is then high-pass filtered along the $\hat{\theta}$ -direction so that only variations on scales $\lesssim 3^\circ$ are preserved. The power beam at each decl. is then divided by the exponential of the high-pass-filtered, logarithmic polynomial model.

The uncertainty in the power beam is estimated at each frequency and decl. by examining the variance at large hour angle ($0.087 \leq |\cos \theta \sin \phi| \leq 0.42$). This uncertainty varies significantly as a function of decl. based on the brightness of the sources at that decl. We apply a 2D Savitzky–Golay filter in (ν, θ) space to low-pass-filter the beam model. For each (ν, θ, ϕ) a fourth-order Chebyshev polynomial in both ν and θ is fit to a small window centered on that location. The best-fit polynomial model is evaluated at that location to obtain the low-pass-filtered version of the beam model. The variance in the beam model at large hour angle is used to estimate the weights in the fit and properly account for the decl.-dependent uncertainties. The size of the window changes between three distinct values based on the decl. and frequency in order to retain features in the beam at progressively smaller scales as one moves to lower decl. In addition to smoothing the beam, the low-pass filter interpolates the beam to the majority of the frequencies that have been masked because of missing data or RFI.

Even after applying the 2D smoothing operation, there are still sharp features in the beam along the frequency axis that we believe originate from unflagged RFI present in the sidereal visibilities. These sharp features will leak foreground power to small-spectral scales when the beam model is deconvolved from the data. To address this, at each (θ, ϕ) we apply an eighth-order low-pass Butterworth filter along the frequency axis. The cutoff used for the low-pass filter is decl. dependent in order to retain what we suspect are actual features of the beam. The cutoff ranges from 125 to 200 ns.

The final estimate of the power beam is obtained from a weighted average of the estimate from baselines with a 44 and 66 m EW component. The baselines with a 0 m EW component are contaminated by diffuse Galactic emission, which is not present in our sky model, and also coupled noise that varies slowly as a function of Earth rotation angle and thus appears at low m values that overlap with the range of m at which the meridian sky fringes. The baselines with 22 m EW components

are also contaminated by coupled noise, albeit to a lesser extent.

This technique is currently unable to measure the primary beam accurately at hour angles greater than $\approx 2^\circ$, where the first-order approximation for the geometric phase given in Equation (26) begins to break down and an additional term that depends on the NS baseline distance, decl., and hour angle becomes relevant. The phase due to this term will be equal to the first-order phase at a new “effective” decl. given by

$$\theta'_{\text{eff}} = \arcsin(\cos \Lambda \sin \theta' - \sin \Lambda \cos \theta' \cos(\phi - \phi')) + \Lambda. \quad (\text{B6})$$

As a result, bright sources will exhibit a “U”-shaped track in the hybrid beamformed visibilities as they move out of the meridian beam centered on their true decl. θ' and into meridian beams centered on more northern decl. at θ'_{eff} . The recovered primary beam will be attenuated at large hour angles by a factor $b_{\text{synth}}^{\hat{\theta}, p}(\nu, \theta, \theta'_{\text{eff}}, \phi) / b_{\text{synth}}^{\hat{\theta}, p}(\nu, \theta, \theta', \phi)$. In the main lobe of the primary beam the attenuation is less than 6% for the polarizations, frequencies, and decl. considered in this work, but in the sidelobes it quickly becomes significant. We are actively exploring extensions to this algorithm that are capable of recovering the sidelobes as well.

Appendix C Stacking on Lognormal Galaxy Density Realizations

In Section 5.3.1, we made the following statement: if simulated galaxy catalogs are drawn from lognormal realizations of the galaxy density δ_g and correlated Gaussian-distributed HI maps are stacked on the resulting galaxy positions, the measured stacking signal is the same as it would be if the galaxy catalogs were drawn from Gaussian realizations of the galaxy density. In this appendix, we justify this statement. We will make use of the following Gaussian integrals: if δ , α are n -component vectors and C is a symmetric, positive-definite $n \times n$ matrix, we can write

$$\int d^n \delta e^{-\frac{1}{2} \delta^T C^{-1} \delta + \alpha \cdot \delta} = \sqrt{(2\pi)^n \det C} e^{\frac{1}{2} \alpha^T C \alpha}, \quad (\text{C1})$$

$$\int d^n\delta \delta_i e^{-\frac{1}{2}\delta^T C^{-1}\delta + \alpha \cdot \delta} = \sqrt{(2\pi)^n \det C} e^{\frac{1}{2}\alpha^T C \alpha} \sum_j C_{ij} \alpha_j. \quad (C2)$$

Consider an idealized version of the stacking analysis, in which we average the HI overdensity δ_{HI} at a 3D separation \mathbf{r} from the location of each of N galaxies in a catalog:

$$S(\mathbf{r}) = \frac{1}{N} \sum_{i=1}^N \delta_{\text{HI}}(\mathbf{x}_i + \mathbf{r}). \quad (C3)$$

Suppose that δ_{HI} has Gaussian statistics, while the galaxy overdensity δ_g , from which the galaxy positions are drawn, is lognormal, related to a Gaussian field δ_G by

$$1 + \delta_g(\mathbf{x}) = e^{\delta_G(\mathbf{x}) + \mu(\mathbf{x})}. \quad (C4)$$

In addition, let $C_{ab}(\mathbf{x}, \mathbf{x}') \equiv \langle \delta_a(\mathbf{x}) \delta_b(\mathbf{x}') \rangle$, where $a, b \in \{\text{G, HI}\}$. Following the standard procedure for lognormal fields, we fix $\mu(\mathbf{x})$ in Equation (C4) such that $\langle 1 + \delta_g(\mathbf{x}) \rangle = 1$. We can compute the relevant ensemble average by defining δ to be δ_g evaluated at a finite number of points and writing

$$\begin{aligned} \langle 1 + \delta_g(\mathbf{x}) \rangle &= \int d^n\delta [1 + \delta_g(\mathbf{x})] \mathcal{P}(\delta) \\ &= \int d^n\delta e^{\delta_G(\mathbf{x}) + \mu(\mathbf{x})} \frac{1}{\sqrt{(2\pi)^n \det C_{\text{GG}}}} e^{-\frac{1}{2}\delta^T C_{\text{GG}}^{-1}\delta} \\ &= e^{\mu(\mathbf{x}) + \frac{1}{2}C_{\text{GG}}(\mathbf{x}, \mathbf{x})}, \end{aligned} \quad (C5)$$

where in the second equality we substituted Equation (C4) and the probability distribution function for a Gaussian random field, and in the final equality we used Equation (C1). Setting this to unity implies that $\mu(\mathbf{x}) = -(1/2)C_{\text{GG}}(\mathbf{x}, \mathbf{x})$.

We wish to show that stacked HI overdensity in Equation (C3) approaches the same result whether the galaxy positions are drawn from the lognormal field in Equation (C4) or from the Gaussian field δ_G itself. To do so, we first consider an ensemble average over galaxy positions in the catalog, keeping the underlying fields (δ_g and δ_{HI}) fixed:

$$\begin{aligned} \langle S(\mathbf{r}) \rangle_{\text{cat}} &= \frac{1}{N} \sum_{i=1}^N \langle \delta_{\text{HI}}(\mathbf{x}_i + \mathbf{r}) \rangle_{\text{cat}} \\ &= \frac{1}{N} \sum_{i=1}^N \int_V d^3\mathbf{x}_i \delta_{\text{HI}}(\mathbf{x}_i + \mathbf{r}) \mathcal{P}(\mathbf{x}_i | \delta_g), \end{aligned} \quad (C6)$$

with the integral evaluated over the survey volume V . The probability distribution function of galaxy positions, $\mathcal{P}(\mathbf{x} | \delta_g)$, is given by $\mathcal{P}(\mathbf{x} | \delta_g) = V^{-1}[1 + \delta_g(\mathbf{x})]$. By combining Equations (C4)–(C6) and the expressions for $\mathcal{P}(\mathbf{x} | \delta_g)$ and $\mu(\mathbf{x})$, we obtain

$$\langle S(\mathbf{r}) \rangle_{\text{cat}} = \frac{1}{V} \int_V d^3\mathbf{x} \delta_{\text{HI}}(\mathbf{x} + \mathbf{r}) e^{\delta_G(\mathbf{x}) - \frac{1}{2}C_{\text{GG}}(\mathbf{x}, \mathbf{x})}. \quad (C7)$$

We now take the ensemble average of Equation (C7) over the density fields, as well as the catalog positions. Defining δ to contain both δ_{HI} and δ_G and letting C represent the joint

covariance, we have

$$\begin{aligned} \langle \delta_{\text{HI}}(\mathbf{x} + \mathbf{r}) e^{\delta_G(\mathbf{x})} \rangle_{\text{fields}} &= \int d^n\delta \delta_{\text{HI}}(\mathbf{x} + \mathbf{r}) e^{\delta_G(\mathbf{x})} \frac{1}{\sqrt{(2\pi)^n \det C}} e^{-\frac{1}{2}\delta^T C^{-1}\delta} \\ &= e^{\frac{1}{2}C_{\text{GG}}(\mathbf{x}, \mathbf{x})} C_{\text{HI}, \text{G}}(\mathbf{x} + \mathbf{r}, \mathbf{x}), \end{aligned} \quad (C8)$$

where we used Equation (C2) in the second equality. Combining this with Equation (C7) and assuming that the fields have translation-invariant statistics, we arrive at

$$\langle S(\mathbf{r}) \rangle_{\text{cat, fields}} = C_{\text{HI}, \text{G}}(\mathbf{r}, \mathbf{0}), \quad (C9)$$

which is also what we would obtain if the galaxy positions were drawn directly from δ_G itself.

Appendix D Template Calculation

In this appendix, we discuss the challenge of calculating the signal templates for arbitrary parameter combinations and the approach we take in this work. Other than the frequency bias parameter $\Delta\nu_0$, the parameters described in Section 5.2.6 affect the properties of the underlying LSS, or the 21 cm or tracer density fields. This suggests that one way of calculating the template is to produce a realization of the 21 cm field and a correlated tracer catalog given a set of parameters and then simulate a CHIME time stream from the 21 cm field using a model for the instrumental transfer function, repeat the analysis procedure done to the actual data (flagging, filtering, and mapmaking), and finally stack the output on the mock catalog. By repeating this procedure and averaging the results, we can estimate the expected signal.

Unfortunately, a full Monte Carlo of this procedure is challenging, as even a single iteration requires around 900 core-hours of compute time (dominated by the time stream generation from input sky maps). Instead, we utilize the ergodic principle. We have an overlapping volume of $\gtrsim 10 h^{-3} \text{ Gpc}^3$ (covered by the eBOSS QSO sample), but the stacking is probing scales $\sim 10 h^{-1} \text{ Mpc}$. This gives many quasi-independent regions of that size within the volume, and so on those scales we expect the volume average to approach the ensemble average, or equivalently, that averaging over independent mock source catalogs drawn from a single LSS realization should give the same as averaging over completely independent LSS realizations. Though this naive picture will break down on larger scales where the cosmic variance contribution is significant, we find that the cosmic variation in the stack signal is small: around 0.8% for the LRG sample, 0.2% for the ELGs, and 0.3% for QSOs (estimated by comparing the zero-lag amplitude for stacks drawn from distinct LSS realizations). This is no more than the variation between single catalogs drawn from the same LSS realization ($\sim 0.7\%$ for LRGs and QSOs, $\sim 1.2\%$ for ELGs), although we average over a sufficiently large number of catalogs to reduce this contribution to well below the cosmic variance level.

While this gives us a tractable method of computing the template for a given set of parameters, it still requires a costly time stream simulation for each set of values. To avoid this, we note that as both the process of observation and analysis are linear (other than data-derived RFI and bright pixel masking), if we can isolate individual terms in the cross-power spectrum

description, they map to distinct contributions to the stacked signal.

For the moment we will fix the Finger-of-God parameters $\alpha_{\text{FoG,H}_1}$ and $\alpha_{\text{FoG,g}}$, as well as the nonlinear power spectrum parameter α_{NL} , and to make the notation more compact, we will define scaling parameters about the fiducial model: $\alpha_{\Omega} = \Omega_{\text{H}_1}/\Omega_{\text{H}_1}^{\text{fid}}(z_{\text{eff}})$, $\alpha_{\text{H}_1} = b_{\text{H}_1}/b_{\text{H}_1}^{\text{fid}}(z_{\text{eff}})$, and $\alpha_g = b_g/b_g^{\text{fid}}(z_{\text{eff}})$. With this we can rewrite Equation (91) as

$$P_{\text{H}_1,g}(\alpha_{\Omega}, \alpha_{\text{H}_1}, \alpha_g, M_{10}) = \alpha_{\Omega}(\alpha_{\text{H}_1}\alpha_g P_{hg} + \alpha_{\text{H}_1} P_{hv} + \alpha_g P_{vg} + P_{vv}) + M_{10} P_{\text{sn}}, \quad (\text{D1})$$

where we have left the k , μ , and z dependence implicit. The power spectrum terms on the right-hand side are

$$P_{hg} = T_b^{\text{fid}}(z_1)b_{\text{H}_1}^{\text{fid}}(z_1)b_g^{\text{fid}}(z_2)P \quad (\text{D2})$$

$$P_{hv} = T_b^{\text{fid}}(z_1)b_{\text{H}_1}^{\text{fid}}(z_1)f(z_2)\mu^2P \quad (\text{D3})$$

$$P_{vg} = T_b^{\text{fid}}(z_1)b_g^{\text{fid}}(z_2)f(z_1)\mu^2P \quad (\text{D4})$$

$$P_{vv} = T_b^{\text{fid}}(z_1)f(z_1)f(z_2)\mu^4P \quad (\text{D5})$$

$$P_{\text{sn}} = C_{\text{H}_1}(z_1)10^{10} M_{\odot}, \quad (\text{D6})$$

with

$$P = [P_{\text{L}}(k) + \alpha_{\text{NL}}(P_{\text{NL}}(k) - P_{\text{L}}(k))] \times D_{\text{H}_1}^{\text{FoG}}(\alpha_{\text{FoG,H}_1} k\mu; z_1) D_g^{\text{FoG}}(\alpha_{\text{FoG,g}} k\mu; z_2). \quad (\text{D7})$$

The linearity of the simulation and analysis procedure means that the stack signal should be separable into distinct terms like Equation (D1). If we write the template generated by the given parameters as $s(\alpha_{\Omega}, \alpha_{\text{H}_1}, \alpha_g, M_{10})$, where we make the dependence on $\Delta\nu$ implicit, we find

$$s(\alpha_{\Omega}, \alpha_{\text{H}_1}, \alpha_g, M_{10}) = \alpha_{\Omega}(\alpha_{\text{H}_1}\alpha_g s_{hg} + \alpha_{\text{H}_1} s_{hv} + \alpha_g s_{vg} + s_{vv}) + M_{10} s_{\text{sn}}, \quad (\text{D8})$$

where each s_{xy} term is the stack signal corresponding to the cross-power spectrum term P_{xy} as defined above. However, as we cannot directly propagate a cross-power spectrum into a stack signal, we must determine these terms indirectly. This can be done by running simulations through with specific α parameters that generate known linear combinations of the s_{xy} . By choosing these simulated parameters judiciously, we can easily invert these combinations to generate the individual s_{xy} terms. One such choice is

$$s_{hg} = s(1, 1, 1, 0) - s(1, 1, 0, 0) - s(1, 0, 1, 0) + s(1, 0, 0, 0) \quad (\text{D9})$$

$$s_{hv} = s(1, 1, 0, 0) - s(1, 0, 0, 0) \quad (\text{D10})$$

$$s_{vg} = s(1, 0, 1, 0) - s(1, 0, 0, 0) \quad (\text{D11})$$

$$s_{vv} = s(1, 0, 0, 0) \quad (\text{D12})$$

$$s_{\text{sn}} = s(0, 0, 0, 1). \quad (\text{D13})$$

Each of the five unique combinations of parameters passed to $s(\alpha_{\Omega}, \alpha_{\text{H}_1}\alpha_g, M_{10})$ in the equation above requires a separate simulation to determine, but after that, we can use these modes and Equation (D8) to determine the stacked template for any combination of parameters.

This scheme allows us to exactly treat the effects of the three linear parameters and the shot-noise contribution on the template. Incorporating the effect of the nonlinear power

spectrum shape is straightforward, as the parameterization used for α_{NL} means that the output stack signal is a simple linear mixing of two terms, as it is in the cross-power spectrum given by Equation (D7). This means that four more simulations can be used to generate templates at any α_{NL} . However, the template is a nonlinear function of the Finger-of-God parameters and so cannot be exactly generated in a finite number of modes.

To account for this, we start by noting that the effect of the Finger-of-God treatment we use (see Section 5.2.3), at constant time and in comoving distance, is a convolution of the underlying fields along the line of sight. In a narrow enough interval in redshift, such that we could ignore evolutionary effects and the constant frequency spacing of our measurement maps to a constant separation in comoving distance, this effect commutes with the stacking, and we could apply it via convolving a post-simulation s_{xy} template to the desired $\alpha_{\text{FoG,x}}$ value, rather than needing to incorporate it into the simulations directly. However, as our sources are located over wide redshift intervals, evolution of the cosmological fields and the pairwise velocity dispersion σ_P cannot be neglected, and in addition the RFI masking, redshift-dependent source number density, and sensitivity further break the stationarity of the radial axis.

However, even if there is no exact mapping from the Finger-of-God effects into a convolution on the template modes, we can still attempt to find an effective one that is accurate in the vicinity of the fiducial Finger-of-God parameters $\alpha_{\text{FoG,x}} = 1$. To do this, we use a transfer function in delay, τ , the Fourier conjugate of the frequency separation $\Delta\nu$, of the form

$$D_x^{\text{eff}}(\tau, \alpha_{\text{FoG,x}}) = \frac{2 + \tau^2\sigma_{\text{eff}}^2}{2 + \alpha_{\text{FoG,x}}\tau^2\sigma_{\text{eff}}^2}, \quad (\text{D14})$$

that will be applied to the templates with the fiducial Finger-of-God strength. To motivate this choice, we note that within a short redshift interval

$$\tau \approx -\frac{1}{2\pi\nu_{21}} \frac{c}{H(z)}(1+z)^2 k_{\parallel}, \quad (\text{D15})$$

and so if we set

$$\sigma_{\text{eff}} = 2\pi\nu_0 \frac{H(z)}{c(1+z)^2} \sigma_P \quad (\text{D16})$$

and then apply the transfer function above to the template, the numerator would effectively undo the Lorentzian Finger-of-God model with the fiducial $\alpha_{\text{FoG,x}} = 1$ and the denominator would reapply it with the desired $\alpha_{\text{FoG,x}}$. Thus, we would have transformed the template mode from the fiducial to the desired $\alpha_{\text{FoG,x}}$ parameter. To take into account the wide redshift range and nonstationarity, we estimate an effective smoothing width σ_{eff} by finding the value that minimizes the template error at a higher $\alpha_{\text{FoG,x}} = 1.2$ compared to an exact simulation at the same value. This effective convolution approach is applicable over a wide range of values of $\alpha_{\text{FoG,x}}$, with errors at $\alpha_{\text{FoG,x}} = 0$ or 3 of $\lesssim 0.5\%$ and much smaller around the pivot $\alpha_{\text{FoG,x}} = 1$. Computationally this requires an additional eight simulations, one for each of the four $(\alpha_{\text{H}_1}, \alpha_g)$ combinations with a perturbed value of $\alpha_{\text{FoG,H}_1} = 1.2$, and an additional four with $\alpha_{\text{FoG,g}} = 1.2$.

The final effect we need to apply is the frequency shift $\Delta\nu_0$. This is performed in Fourier space by phase rotating the delay transform of the template.

ORCID iDs

Mandana Amiri [ID](https://orcid.org/0000-0001-6523-9029) <https://orcid.org/0000-0001-6523-9029>
 Kevin Bandura [ID](https://orcid.org/0000-0003-3772-2798) <https://orcid.org/0000-0003-3772-2798>
 Tianyue Chen [ID](https://orcid.org/0000-0003-0173-6274) <https://orcid.org/0000-0003-0173-6274>
 Meiling Deng [ID](https://orcid.org/0000-0001-8123-7322) <https://orcid.org/0000-0001-8123-7322>
 Matt Dobbs [ID](https://orcid.org/0000-0001-7166-6422) <https://orcid.org/0000-0001-7166-6422>
 Mateus Fandino [ID](https://orcid.org/0000-0002-6899-1176) <https://orcid.org/0000-0002-6899-1176>
 Simon Foreman [ID](https://orcid.org/0000-0002-0190-2271) <https://orcid.org/0000-0002-0190-2271>
 Mark Halpern [ID](https://orcid.org/0000-0002-1760-0868) <https://orcid.org/0000-0002-1760-0868>
 Alex S. Hill [ID](https://orcid.org/0000-0001-7301-5666) <https://orcid.org/0000-0001-7301-5666>
 Gary Hinshaw [ID](https://orcid.org/0000-0002-4241-8320) <https://orcid.org/0000-0002-4241-8320>
 Carolin Höfer [ID](https://orcid.org/0000-0003-4887-8114) <https://orcid.org/0000-0003-4887-8114>
 Joseph Kania [ID](https://orcid.org/0000-0002-3354-3859) <https://orcid.org/0000-0002-3354-3859>
 T. L. Landecker [ID](https://orcid.org/0000-0003-1455-2546) <https://orcid.org/0000-0003-1455-2546>
 Joshua MacEachern [ID](https://orcid.org/0000-0001-8064-6116) <https://orcid.org/0000-0001-8064-6116>
 Kiyoshi Masui [ID](https://orcid.org/0000-0002-4279-6946) <https://orcid.org/0000-0002-4279-6946>
 Juan Mena-Parra [ID](https://orcid.org/0000-0002-0772-9326) <https://orcid.org/0000-0002-0772-9326>
 Nikola Milutinovic [ID](https://orcid.org/0000-0001-8292-0051) <https://orcid.org/0000-0001-8292-0051>
 Arash Mirhosseini [ID](https://orcid.org/0000-0002-2626-5985) <https://orcid.org/0000-0002-2626-5985>
 Laura Newburgh [ID](https://orcid.org/0000-0002-7333-5552) <https://orcid.org/0000-0002-7333-5552>
 Anna Ordog [ID](https://orcid.org/0000-0002-2465-8937) <https://orcid.org/0000-0002-2465-8937>
 Ue-Li Pen [ID](https://orcid.org/0000-0003-2155-9578) <https://orcid.org/0000-0003-2155-9578>
 Tristan Pinsonneault-Marotte [ID](https://orcid.org/0000-0002-9516-3245) <https://orcid.org/0000-0002-9516-3245>
 Ava Polzin [ID](https://orcid.org/0000-0002-5283-933X) <https://orcid.org/0000-0002-5283-933X>
 Alex Reda [ID](https://orcid.org/0000-0001-6967-7253) <https://orcid.org/0000-0001-6967-7253>
 Andre Renard [ID](https://orcid.org/0000-0003-3463-7918) <https://orcid.org/0000-0003-3463-7918>
 J. Richard Shaw [ID](https://orcid.org/0000-0002-4543-4588) <https://orcid.org/0000-0002-4543-4588>
 Seth R. Siegel [ID](https://orcid.org/0000-0003-2631-6217) <https://orcid.org/0000-0003-2631-6217>
 Saurabh Singh [ID](https://orcid.org/0000-0001-7755-902X) <https://orcid.org/0000-0001-7755-902X>
 Keith Vanderlinde [ID](https://orcid.org/0000-0003-4535-9378) <https://orcid.org/0000-0003-4535-9378>
 Haochen Wang [ID](https://orcid.org/0000-0002-1491-3738) <https://orcid.org/0000-0002-1491-3738>
 Donald V. Wiebe [ID](https://orcid.org/0000-0002-6669-3159) <https://orcid.org/0000-0002-6669-3159>
 Dallas Wulf [ID](https://orcid.org/0000-0001-7314-9496) <https://orcid.org/0000-0001-7314-9496>

References

- Ahumada, R., Prieto, C. A., Almeida, A., et al. 2020, [ApJS](#), **249**, 3
 Alam, S., Aubert, M., Avila, S., et al. 2021, [PhRvD](#), **103**, 083533
 Alam, S., Peacock, J. A., Kraljic, K., Ross, A. J., & Comparat, J. 2020, [MNRAS](#), **497**, 581
 Albareti, F. D., Prieto, C. A., Almeida, A., et al. 2017, [ApJS](#), **233**, 25
 Anderson, C. J., Luciwi, N. J., Li, Y. C., et al. 2018, [MNRAS](#), **476**, 3382
 Ando, R., Nishizawa, A. J., Hasegawa, K., Shimizu, I., & Nagamine, K. 2019, [MNRAS](#), **484**, 5389
 Baars, J. W. M., Genzel, R., Pauliny-Toth, I. I. K., & Witzel, A. 1977, [A&A](#), **61**, 99
 Battye, R. A., Davies, R. D., & Weller, J. 2004, [MNRAS](#), **355**, 1339
 Behnel, S., Bradshaw, R., Citro, C., et al. 2011, [CSE](#), **13**, 31
 Bera, A., Kanekar, N., Chengalur, J. N., & Bagla, J. S. 2019, [ApJL](#), **882**, L7
 Blanton, M. R., Bershadsky, M. A., Abolfathi, B., et al. 2017, [AJ](#), **154**, 28
 Bull, P., Ferreira, P. G., Patel, P., & Santos, M. G. 2015, [ApJ](#), **803**, 21
 Calderhead, B., & Girolami, M. 2009, [Comput. Stat. Data Anal.](#), **53**, 4028
 Castorina, E., & Villaescusa-Navarro, F. 2017, [MNRAS](#), **471**, 1788
 Chakraborty, A., Datta, A., Roy, N., et al. 2021, [ApJL](#), **907**, L7
 Chang, T.-C., Pen, U.-L., Bandura, K., & Peterson, J. B. 2010, [Natur](#), **466**, 463
 Chang, T.-C., Pen, U.-L., Peterson, J. B., & McDonald, P. 2008, [PhRvL](#), **100**, 091303
 Chen, Q., Meyer, M., Popping, A., et al. 2021, [MNRAS](#), **508**, 2758
 Chen, S.-F., Vlah, Z., & White, M. 2022, [JCAP](#), **2022**, 008
 CHIME/FRB Collaboration, Amiri, M., Bandura, K., et al. 2018, [ApJ](#), **863**, 48
 CHIME Collaboration, Amiri, M., Bandura, K., et al. 2022a, [ApJ](#), **932**, 100
 CHIME Collaboration, Amiri, M., Bandura, K., et al. 2022b, [ApJS](#), **261**, 29
 CHIME/Pulsar Collaboration, Amiri, M., Bandura, K. M., et al. 2021, [ApJS](#), **255**, 5
 Chowdhury, A., Kanekar, N., Chengalur, J. N., Sethi, S., & Dwarakanath, K. S. 2020, [Natur](#), **586**, 369
 Cohen, A. S., Lane, W. M., Cotton, W. D., et al. 2007, [AJ](#), **134**, 1245
 Collette, A., Kluyver, T., Caswell, T. A., et al. 2021, [h5py/h5py](#): v3.5.0, Zenodo, doi:[10.5281/zenodo.5585380](https://doi.org/10.5281/zenodo.5585380)
 Condon, J. J., Cotton, W. D., Greisen, E. W., et al. 1998, [AJ](#), **115**, 1693
 Cosmic Visions 21 cm Collaboration, Ansari, R., Arena, E. J., et al. 2018, arXiv:[1810.09572](https://arxiv.org/abs/1810.09572)
 Crichton, D., Aich, M., Amara, A., et al. 2022, [JATIS](#), **8**, 011019
 Crighton, N. H. M., Murphy, M. T., Prochaska, J. X., et al. 2015, [MNRAS](#), **452**, 217
 Dalcin, L., & Fang, Y.-L. L. 2021, [CSE](#), **23**, 47
 d'Amico, G., Gleyzes, J., Kokron, N., et al. 2020, [JCAP](#), **2020**, 005
 Dawson, K. S., Kneib, J.-P., Percival, W. J., et al. 2016, [AJ](#), **151**, 44
 de Mattia, A., Ruhmann-Kleider, V., Raichoor, A., et al. 2021, [MNRAS](#), **501**, 5616
 DeBoer, D. R., Parsons, A. R., Aguirre, J. E., et al. 2017, [PASP](#), **129**, 045001
 Delhaize, J., Meyer, M. J., Staveley-Smith, L., & Boyle, B. J. 2013, [MNRAS](#), **433**, 1398
 Dey, A., Schlegel, D. J., Lang, D., et al. 2019, [AJ](#), **157**, 168
 Eriksen, H. K., O'Dwyer, I. J., Jewell, J. B., et al. 2004, [ApJS](#), **155**, 227
 Ewall-Wice, A., Kern, N., Dillon, J. S., et al. 2021, [MNRAS](#), **500**, 5195
 Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, [PASP](#), **125**, 306
 Gabriel, E., Fagg, G. E., Bosilca, G., et al. 2004, in [Recent Advances in Parallel Virtual Machine and Message Passing Interface](#), ed. D. Kranzlmüller, P. Kacsuk, & J. Dongarra (Berlin: Springer), 97
 Gelman, A., & Meng, X.-L. 1998, [StaSc](#), **13**, 163
 Geman, S., & Geman, D. 1984, [ITPAM](#), PAMI-6, 721
 Gil-Marín, H., Bautista, J. E., Paviot, R., et al. 2020, [MNRAS](#), **498**, 2492
 Gordon, C., & Trott, R. 2007, [MNRAS](#), **382**, 1859
 Gorski, K. M., Hivon, E., Banday, A. J., et al. 2005, [ApJ](#), **622**, 759
 Gould, R. J. 1994, [ApJ](#), **423**, 522
 Gregory, P. C., Scott, W. K., Douglas, K., & Condon, J. J. 1996, [ApJS](#), **103**, 427
 Gunn, J. E., Siegmund, W. A., Mannery, E. J., et al. 2006, [AJ](#), **131**, 2332
 Guo, H., Jones, M. G., Haynes, M. P., & Fu, J. 2020, [ApJ](#), **894**, 92
 Hagberg, A. A., Schult, D. A., & Swart, P. J. 2008, in [Proc. 7th Python in Science Conf.](#), ed. G. Varoquaux, T. Vaught, & J. Millman, 11
 Handley, W., & Lemos, P. 2019, [PhRvD](#), **100**, 023512
 Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, [Natur](#), **585**, 357
 Hu, W., Catinella, B., Cortese, L., et al. 2020, [MNRAS](#), **493**, 1587
 Hu, W., Hoppmann, L., Staveley-Smith, L., et al. 2019, [MNRAS](#), **489**, 1619
 Hunter, J. D. 2007, [CSE](#), **9**, 90
 Ivanov, M. M., Simonović, M., & Zaldarriaga, M. 2020, [JCAP](#), **2020**, 042
 Jackson, J. C. 1972, [MNRAS](#), **156**, 1P
 Jeffreys, H. 1961, [Theory of Probability](#) (3rd ed.; Oxford: Oxford Univ. Press)
 Jewell, J., Levin, S., & Anderson, C. H. 2004, [ApJ](#), **609**, 1
 Jones, M. G., Haynes, M. P., Giovanelli, R., & Moorman, C. 2018, [MNRAS](#), **477**, 2
 Kaiser, N. 1987, [MNRAS](#), **227**, 1
 Kanekar, N., Sethi, S., & Dwarakanath, K. S. 2016, [ApJL](#), **818**, L28
 Kaplan, G. H. 2005, [USNOC](#), **179**
 Karamanis, M., & Beutler, F. 2021, arXiv:[2106.06331](https://arxiv.org/abs/2106.06331)
 Kass, R. E., & Raftery, A. E. 1995, [J. Am. Stat. Assoc.](#), **90**, 773
 Kokron, N., DeRose, J., Chen, S.-F., White, M., & Wechsler, R. H. 2021, [MNRAS](#), **505**, 1422
 Kovetz, E., Breysse, P. C., Lidz, A., et al. 2019, [BAAS](#), **51**, 101
 Lah, P., Chengalur, J. N., Briggs, F. H., et al. 2007, [MNRAS](#), **376**, 1357
 Landecker, T. L., Dewdney, P. E., Burgess, T. A., et al. 2000, [A&AS](#), **145**, 509
 Lang, D., Hogg, D. W., & Schlegel, D. J. 2016, [AJ](#), **151**, 36
 Laurent, P., Eftekhari, S., Le Goff, J.-M., et al. 2017, [JCAP](#), **2017**, 017
 Lewis, A. 2019, arXiv:[1910.13970](https://arxiv.org/abs/1910.13970)
 Lewis, A., Challinor, A., & Lasenby, A. 2000, [ApJ](#), **538**, 473
 Li, J., Zuo, S., Wu, F., et al. 2020, [SCPMA](#), **63**, 129862
 Li, L.-C., Staveley-Smith, L., & Rhee, J. 2021, [RAA](#), **21**, 030
 Liu, A., & Tegmark, M. 2011, [PhRvD](#), **83**, 103006
 Lyke, B. W., Higley, A. N., McLane, J. N., et al. 2020, [ApJS](#), **250**, 8
 Masui, K., Amiri, M., Connor, L., et al. 2015, [A&C](#), **12**, 181
 Masui, K. W., Shaw, J. R., Ng, C., et al. 2017, arXiv:[1710.08591](https://arxiv.org/abs/1710.08591)
 Masui, K. W., Switzer, E. R., Banavar, N., et al. 2013, [ApJL](#), **763**, L20
 McKinney, W. 2010, in [Proc. 9th Python in Science Conf.](#), ed. S. van der Walt & J. Millman, 56, doi:[10.25080/Majora-92bf1922-00a](https://doi.org/10.25080/Majora-92bf1922-00a)

- Mead, A. J., Brieden, S., Tröster, T., & Heymans, C. 2021, *MNRAS*, **502**, 1401
- Merson, A., Smith, A., Benson, A., Wang, Y., & Baugh, C. 2019, *MNRAS*, **486**, 5737
- Oh, S. P., & Mack, K. J. 2003, *MNRAS*, **346**, 871
- Paciga, G., Albert, J. G., Bandura, K., et al. 2013, *MNRAS*, **433**, 639
- Padmanabhan, H. 2021, *IJMPD*, **30**, 2130009
- The pandas development team 2020, pandas-dev/pandas: Pandas v1.2.0 Zenodo, doi:10.5281/zenodo.7344967
- Parsons, A. R., Liu, A., Aguirre, J. E., et al. 2014, *ApJ*, **788**, 106
- Pen, U.-L., Staveley-Smith, L., Peterson, J. B., & Chang, T.-C. 2009, *MNRAS*, **394**, L6
- Perley, R. A., & Butler, B. J. 2017, *ApJS*, **230**, 7
- Peterson, J. B., Aleksan, R., Ansari, R., et al. 2009, arXiv:0902.3091
- Petit, G., & Luzum, B. 2010, *ITN*, **36**, 1
- Planck Collaboration, Aghanim, N., Akrami, Y., et al. 2020, *A&A*, **641**, A6
- Rafiei-Ravandi, M., Smith, K. M., Li, D., et al. 2021, *ApJ*, **922**, 42
- Raichoor, A., de Mattia, A., Ross, A. J., et al. 2021, *MNRAS*, **500**, 3254
- Rao, S. M., Turnshek, D. A., Sardane, G. M., & Monier, E. M. 2017, *MNRAS*, **471**, 3428
- Rengelink, R. B., Tang, Y., de Bruyn, A. G., et al. 1997, *A&AS*, **124**, 259
- Rhee, J., Lah, P., Briggs, F. H., et al. 2018, *MNRAS*, **473**, 1879
- Rhee, J., Lah, P., Chengalur, J. N., Briggs, F. H., & Colless, M. 2016, *MNRAS*, **460**, 2675
- Rhee, J., Zwaan, M. A., Briggs, F. H., et al. 2013, *MNRAS*, **435**, 2693
- Rhodes, B. 2019, Skyfield: High Precision Research-grade Positions for Planets and Earth Satellites Generator, Astrophysics Source Code Library, ascl:1907.024
- Richards, G. T., Kruczak, N. E., Gallagher, S. C., et al. 2011, *AJ*, **141**, 167
- Ross, A. J., Bautista, J., Tojeiro, R., et al. 2020, *MNRAS*, **498**, 2354
- Santos, M. G., Cooray, A., & Knox, L. 2005, *ApJ*, **625**, 575
- Sarkar, D., & Bharadwaj, S. 2019, *MNRAS*, **487**, 5666
- Scoccimarro, R. 2004, *PhRvD*, **70**, 083007
- Shaver, P. A., Windhorst, R. A., Madau, P., & de Bruyn, A. G. 1999, *A&A*, **345**, 380
- Shaw, J. R., Foreman, S., Nitsche, R., et al. 2020a, radiocosmology/driftscan: 20.5.0, v20.5.0, Zenodo, doi:10.5281/zenodo.5949193
- Shaw, J. R., Masui, K., Hincks, A. D., et al. 2020b, chime-experiment/ch_pipeline: 20.10.0, v20.10.0, Zenodo, doi:10.5281/zenodo.5846379
- Shaw, J. R., Masui, K., Nitsche, R., et al. 2020c, radiocosmology/caput: 20.10.0, v20.10.0, Zenodo, doi:10.5281/zenodo.5846375
- Shaw, J. R., Nitsche, R., Foreman, S., & Kefala, A. 2020d, radiocosmology/cora: 20.5.0, v20.5.0, Zenodo, doi:10.5281/zenodo.5949195
- Shaw, J. R., Sigurdson, K., Pen, U.-L., Stebbins, A., & Sitwell, M. 2014, *ApJ*, **781**, 57
- Shaw, J. R., Sigurdson, K., Sitwell, M., Stebbins, A., & Pen, U.-L. 2015, *PhRvD*, **91**, 083514
- Shaw, R., Nitsche, R., Siegel, S. R., et al. 2020e, radiocosmology/draco: 20.10.0, v20.10.0, Zenodo, doi:10.5281/zenodo.5828373
- Shen, Y., Brandt, W. N., Richards, G. T., et al. 2016, *ApJ*, **831**, 7
- Skilling, J. 2006, *BayAn*, **1**, 833
- Smee, S. A., Gunn, J. E., Uomoto, A., et al. 2013, *AJ*, **146**, 32
- Smith, R. E., Peacock, J. A., Jenkins, A., et al. 2003, *MNRAS*, **341**, 1311
- Spinelli, M., Zoldan, A., De Lucia, G., Xie, L., & Viel, M. 2020, *MNRAS*, **493**, 5434
- Stevens, A. R. H., Croton, D. J., & Mutch, S. J. 2016, *MNRAS*, **461**, 859
- Subrahmanya, C. R., Manoharan, P. K., & Chengalur, J. N. 2017, *JApA*, **38**, 10
- Switzer, E. R., Masui, K. W., Bandura, K., et al. 2013, *MNRAS*, **434**, L46
- Takahashi, R., Sato, M., Nishimichi, T., Taruya, A., & Oguri, M. 2012, *ApJ*, **761**, 152
- Tamone, A., Raichoor, A., Zhao, C., et al. 2020, *MNRAS*, **499**, 5527
- Tinker, J., Kravtsov, A. V., Klypin, A., et al. 2008, *ApJ*, **688**, 709
- Tramonte, D., & Ma, Y.-Z. 2020, *MNRAS*, **498**, 5916
- Tramonte, D., Ma, Y.-Z., Li, Y.-C., & Staveley-Smith, L. 2019, *MNRAS*, **489**, 385
- Trott, C. M., Pindor, B., Procopio, P., et al. 2016, *ApJ*, **818**, 139
- Vanderlinde, K., Liu, A., Gaensler, B., et al. 2019, in Canadian Long Range Plan for Astronomy and Astrophysics White Papers, LRP2020, ed. P. Barmby, M. Dobbs, B. Gaensler et al. (Victoria: Canadian Astronomical Society), 28
- Vaníček, P. 1969, *Ap&SS*, **4**, 387
- Villaescusa-Navarro, F., Genel, S., Castorina, E., et al. 2018, *ApJ*, **866**, 135
- Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *NatMe*, **17**, 261
- Vollmer, B., Davoust, E., Dubois, P., et al. 2005, *A&A*, **436**, 757
- Vollmer, B., Gassmann, B., Derrière, S., et al. 2010, *A&A*, **511**, A53
- Wandelt, B. D., Larson, D. L., & Lakshminarayanan, A. 2004, *PhRvD*, **70**, 083511
- Wang, Z., Chen, Y., Mao, Y., et al. 2021, *ApJ*, **907**, 4
- Wiese, W. L., & Fuhr, J. R. 2009, *JPCRD*, **38**, 565
- Wolz, L., Blake, C., & Wyithe, J. S. B. 2017, *MNRAS*, **470**, 3220
- Wolz, L., Pourtsidou, A., Masui, K. W., et al. 2022, *MNRAS*, **510**, 3495
- Wu, F., Li, J., Zuo, S., et al. 2021, *MNRAS*, **506**, 3455
- Wyithe, J. S. B., & Loeb, A. 2008, *MNRAS*, **383**, 606
- Xi, H., Staveley-Smith, L., For, B.-Q., et al. 2020, *MNRAS*, **501**, 4550
- Zafar, T., Péroux, C., Popping, A., et al. 2013, *A&A*, **556**, A141
- Zarrouk, P., Burtin, E., Gil-Marín, H., et al. 2018, *MNRAS*, **477**, 1639
- Zhai, Z., Tinker, J. L., Hahn, C., et al. 2017, *ApJ*, **848**, 76
- Zonca, A., Singer, L., Lenz, D., et al. 2019, *JOSS*, **4**, 1298