



COMAP Early Science. VIII. A Joint Stacking Analysis with eBOSS Quasars

Delaney A. Dunne¹, Kieran A. Cleary¹, Patrick C. Breysse², Dongwoo T. Chung^{3,4}, Håvard T. Ihle⁵, J. Richard Bond³, Hans Kristian Eriksen⁵, Joshua Ott Gundersen⁶, Laura C. Keating⁷, Junhan Kim^{1,8}, Jonas Gahr Sturtzel Lunde⁵, Norman Murray³, Hamsa Padmanabhan⁹, Liju Philip^{10,11}, Nils-Ole Stutzer⁵, Doğa Tolgay^{3,12}, Ingunn Katherine Wehus⁵, Sarah E. Church¹³, Todd Gaier¹⁰, Andrew I. Harris¹⁴, Richard Hobbs¹⁵, James W. Lamb¹⁵, Charles R. Lawrence¹⁰, Anthony C. S. Readhead¹, and David P. Woody¹⁵

¹ California Institute of Technology, 1200 E. California Boulevard, Pasadena, CA 91125, USA

² Center for Cosmology and Particle Physics, Department of Physics, New York University, 726 Broadway, New York, NY, 10003, USA

³ Canadian Institute for Theoretical Astrophysics, University of Toronto, 60 St. George Street, Toronto, ON M5S 3H8, Canada

⁴ Dunlap Institute for Astronomy and Astrophysics, University of Toronto, 50 St. George Street, Toronto, ON M5S 3H4, Canada

⁵ Institute of Theoretical Astrophysics, University of Oslo, P.O. Box 1029 Blindern, N-0315 Oslo, Norway

⁶ Department of Physics, University of Miami, 1320 Campo Sano Avenue, Coral Gables, FL 33146, USA

⁷ Institute for Astronomy, University of Edinburgh, Blackford Hill, Edinburgh, EH9 3HJ, UK

⁸ Department of Physics, Korea Advanced Institute of Science and Technology (KAIST), 291 Daehak-ro, Yuseong-gu, Daejeon 34141, Republic of Korea

⁹ Département de Physique Théorique, Université de Genève, 24 Quai Ernest-Ansermet, CH-1211 Genève 4, Switzerland

¹⁰ Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA 91109, USA

¹¹ Brookhaven National Laboratory, P.O. Box 5000, Upton, NY 11973-5000, USA

¹² Department of Physics, University of Toronto, 60 St. George Street, Toronto, ON, M5S 1A7, Canada

¹³ Kavli Institute for Particle Astrophysics and Cosmology & Physics Department, Stanford University, Stanford, CA 94305, USA

¹⁴ Department of Astronomy, University of Maryland, College Park, MD 20742, USA

¹⁵ Owens Valley Radio Observatory, California Institute of Technology, Big Pine, CA 93513, USA

Received 2023 May 10; revised 2024 February 23; accepted 2024 February 23; published 2024 April 1

Abstract

We present a new upper limit on the cosmic molecular gas density at $z = 2.4\text{--}3.4$ obtained using the first year of observations from the CO Mapping Array Project (COMAP). COMAP data cubes are stacked on the 3D positions of 243 quasars selected from the Extended Baryon Oscillation Spectroscopic Survey (eBOSS) catalog, yielding a 95% upper limit for flux from CO(1–0) line emission of $0.129 \text{ Jy km s}^{-1}$. Depending on the balance of the emission between the quasar host and its environment, this value can be interpreted as an average CO line luminosity L'_{CO} of eBOSS quasars of $\leq 1.26 \times 10^{11} \text{ K km pc}^2 \text{ s}^{-1}$, or an average molecular gas density ρ_{H_2} in regions of the Universe containing a quasar of $\leq 1.52 \times 10^8 M_{\odot} \text{ cMpc}^{-3}$. The L'_{CO} upper limit falls among CO line luminosities obtained from individually targeted quasars in the COMAP redshift range, and the ρ_{H_2} value is comparable to upper limits obtained from other line intensity mapping (LIM) surveys and their joint analyses. Further, we forecast the values obtainable with the COMAP/eBOSS stack after the full 5 yr COMAP Pathfinder survey. We predict that a detection is probable with this method, depending on the CO properties of the quasar sample. Based on the achieved sensitivity, we believe that this technique of stacking LIM data on the positions of traditional galaxy or quasar catalogs is extremely promising, both as a technique for investigating large galaxy catalogs efficiently at high redshift and as a technique for bolstering the sensitivity of LIM experiments, even with a fraction of their total expected survey data.

Unified Astronomy Thesaurus concepts: CO line emission (262); Cosmological evolution (336); High-redshift galaxies (734); Molecular gas (1073); Quasars (1319); Radio astronomy (1338)

1. Introduction

Spectral line intensity mapping (LIM) is an emerging observational technique with the potential to enhance our understanding of the Universe by constraining the global properties of galaxies over cosmic time. LIM surveys do not aim to resolve individual galaxies, but instead measure three-dimensional fluctuations in the integrated emission from many galaxies, allowing for efficient mapping of galaxies across large cosmic volumes (see Kovetz et al. 2019 for a review). Because LIM measures integrated emission, it is sensitive to the faintest galaxies, which are nearly impossible to detect in traditional surveys but nevertheless make up the

bulk of the galaxy population at any given cosmic time. Observationally, the field is still in its early stages, with dedicated LIM instruments only beginning to publish early autocorrelation constraints (e.g., COMAP, Cleary et al. 2022; the Hydrogen Epoch of Reionization Array, HERA, Abdurashidova et al. 2022; and MeerKAT, Paul et al. 2023).

Currently, the results being released from these first LIM surveys are primarily upper limits (for COMAP, Chung et al. 2022) and are not yet sufficient to achieve the measurements described above. This is partially a consequence of the newness of the field—strong autocorrelation detections of a given spectral line are required to measure its cosmological fluctuations, and most LIM surveys have simply not been integrating long enough to achieve the needed sensitivities (for example, COMAP projects a secure autocorrelation detection after five years; Chung et al. 2022). However, cosmic line emission may be detectable even with present-day



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

LIM data sets, if careful signal processing techniques making use of external data sets are applied.

In particular, averaging together intensity-map voxels (3D pixels) that are known to contain galaxies (by comparison with some other survey) in a stacking analysis has the potential to make a detection of the average CO line temperature associated with catalog galaxies, even with an intensity map insufficiently sensitive for other analyses (Silva et al. 2021). Stacking, or coadding, is an established technique for improving sensitivity in traditional, targeted galaxy surveys (e.g., Stanley et al. 2019; Jolly et al. 2021; Lujan Niemeyer et al. 2022; Romano et al. 2022), and has recently been extended to LIM observations as well (Keenan et al. 2022). The efficacy of this technique will depend on factors such as the number of traditional survey objects that fall into the LIM survey footprint and the redshift accuracy in the traditional catalog (Chung et al. 2019; Silva et al. 2021), as well as the chosen weighting scheme (Sinigaglia et al. 2022).

In this paper, we use LIM data from the CO Mapping Array Project (COMAP; Cleary et al. 2022), the first survey to use a purpose-built instrument (the COMAP Pathfinder, described in detail in Lamb et al. 2022), to impose direct constraints on the clustering-scale CO power spectrum at high redshifts. We use COMAP Season 1 data, taken during the first 13 months observing season of the project. COMAP observations cover a frequency range of 26–34 GHz (redshifts of $z = 2.4$ – 3.4) and encompass three $\sim 4 \text{ deg}^2$ fields. The data reduction and map-making processes for these Season 1 data are described in Foss et al. (2022), while power spectrum estimation techniques and constraints are described in Ihle et al. (2022) and Chung et al. (2022). Observations are continuing with the Pathfinder to complete the nominal 5 yr survey, at the end of which a detection of the CO auto-power spectrum is forecast. The COMAP fields were selected to overlap with the HETDEX survey of Ly α emitters (Gebhardt et al. 2021), allowing us to perform stacking and cross-correlation analyses of these two data sets.

As a preliminary step, we investigate in this work the potential for the BOSS/eBOSS quasar sample (Dawson et al. 2016) to provide additional constraints and inform our understanding of the relationship between active galaxies and star formation. The eBOSS catalog was developed with the intention of studying the Ly α forest and therefore covers a redshift range overlapping that of COMAP—the eBOSS DR16 catalog includes spectroscopic observations of $>239,000$ quasars in the range $z > 2.1$. Depending on the assumptions made, a stack of COMAP data on the positions of eBOSS quasars can be treated as a measurement of either the average CO luminosity of the quasars themselves or the average molecular gas density at quasar positions.

In addition to its extensive size, which allows for a high-sensitivity stack, the combination of COMAP with the eBOSS catalog enables the study of the CO properties of a large sample of high-redshift quasars. Quasar feedback, particularly in relation to the molecular gas traced by CO transitions, is poorly understood on a statistical scale. Current studies mainly involve individual objects, using the CO rotational ladder in the host galaxies of active galactic nuclei (AGN) to study the effects of AGN on their surroundings. These are complex measurements requiring extremely CO-bright objects, and thus the pool of available galaxies for study is small and biased, especially at high redshifts. LIM measurements will be able to

Table 1
The Locations of the Three Fields Used in the COMAP Pathfinder Survey

Field	R.A. (J2000)	Decl. (J2000)
Field 1	01 ^h 41 ^m 44 ^s .4	00° 00′ 00″.0
Field 2	11 ^h 20 ^m 00 ^s .0	52° 30′ 00″.0
Field 3	15 ^h 04 ^m 00 ^s .0	55° 00′ 00″.0

provide a more complete picture of the wider population of AGN and their host galaxies (e.g., Breyse & Alexandroff 2019).

In this work, we will investigate the CO properties of the 243 eBOSS quasars in the COMAP fields through stacking, thus providing constraints on the molecular gas content of a large sample of high-redshift quasars for the first time. We will additionally use this analysis to explore the viability of stacking as an LIM technique. As such, we will forecast the sensitivity of the COMAP Pathfinder survey to CO emission from the eBOSS quasars after five years of observing. Section 2 introduces both the COMAP Pathfinder (Section 2.1) and eBOSS surveys (Section 2.2), and Section 3 describes the stacking methodology we use in our analysis. Section 4 describes verification of the stacking analysis. We present results in Section 5, including forecasts of the results of stacking the full Pathfinder survey (Section 5.3). We discuss the implications of these results in Section 6.

We assume a Λ CDM cosmology based broadly on nine-year WMAP (Hinshaw et al. 2013) results throughout, with $\Omega_m = 0.286$, $\Omega_\Lambda = 0.714$, $\Omega_b = 0.047$, and $H_0 = 100 h \text{ km s}^{-1} \text{ Mpc}^{-1}$ with $h = 0.7$. These are the same values used in all previous COMAP works. Unless otherwise indicated, distances are given as proper values.

2. Data

2.1. COMAP

The COMAP Season 1 data consist of 13 months of observations using the COMAP Pathfinder instrument, a single-polarization 19 feed spectrometer array fielded on a 10.4 m telescope located at the Owens Valley Radio Observatory (OVRO) in California (the Pathfinder is described in detail in Lamb et al. 2022). The Pathfinder instrument observes between 26 and 34 GHz, and is therefore sensitive to the CO(1–0) rotational transition ($\nu_0 = 115.27 \text{ GHz}$) emitted in a redshift range of $z = 2.4$ – 3.4 . These observations were toward three $\sim 4 \text{ deg}^2$ cosmological fields (shown in Table 1).

During observations, the telescope is positioned at the leading edge of the field and performs a scanning motion while the cosmological field drifts through. This usually takes 3–10 minutes, after which the telescope is pointed to the leading edge of the field again and the procedure is repeated. This means that over time, we are scanning the fields from several different directions leading to a smoothly varying noise distribution in the final maps, with low noise in the central regions of the field, and gradually increasing levels of noise toward the edges of the fields.

Raw time-ordered data (TOD) from the telescope are processed by applying a series of filters with the goal of removing correlated noise, standing waves, continuum foregrounds, ground contamination and other systematic effects from the TOD. This process and the COMAP scanning strategies are described in detail by Foss et al. (2022). The

output of this pipeline is a set of three calibrated 3D intensity maps, each of angular size $\sim 4'9''\text{--}4'4''$ with 31.25 MHz spectral resolution and a $4'9''\text{--}4'4''$ beam FWHM. At small scales these maps are dominated by uncorrelated Gaussian noise (Foss et al. 2022). The maps are somewhat elongated in the R.A. direction, giving us appreciable coverage in a 3D region with dimensions (in comoving coordinates) roughly $300\text{ Mpc} \times 200\text{ Mpc}$ in directions perpendicular to the line of sight (at redshift $z = 2.9$) and about 1000 Mpc along the line of sight, for a total volume of order $6 \times 10^7\text{ Mpc}^3$ comoving for each field.

As described by Foss et al. (2022), Ihle et al. (2022), and Rennie et al. (2022), the maps are calibrated to the total power entering the telescope. About 72% of this power comes from the main beam, with another roughly 10% of the power in the near sidelobes (within about $30'$). To take this scale dependence into account, Ihle et al. (2022) use a model of the COMAP beam to construct a beam transfer function for the power spectrum. For the purposes of our stacking analysis, since we are mostly interested in the smallest scales in the map, we simply divide the map (and the corresponding uncertainty map) by a single overall factor of 0.72 to get the right calibration at scales corresponding to the main beam. We mask any map voxel (3D pixel) with an integration time below 1000 s (50,000 “hits,” a stricter cut than is taken for the overall COMAP pipeline; Foss et al. 2022).

2.2. eBOSS

The Baryon Oscillation Spectroscopic Survey (BOSS; Dawson et al. 2013) and its extension (eBOSS; Ahumada et al. 2020) together encompass 11 years of observations using the BOSS spectrograph at Apache Point Observatory (Gunn et al. 2006; Smee et al. 2013). We use the large ($\sim 240,000$ object) eBOSS spectroscopic sample of quasi-stellar objects (QSOs) at $z \geq 2.4$, which was intended to enable studies of the Ly α forest at $z \sim 1$. These observations targeted objects that were selected based on Wide-field Infrared Survey–Sloan Digital Sky Survey (SDSS) DR13 imaging data. We use the eBOSS DR16 superset catalog (Ahumada et al. 2020; Lyke et al. 2020) released in 2020—the final iteration of the BOSS/eBOSS catalog. While principally composed of traditional quasars, the superset applied no quasar-selection pipeline and thus contains a small percentage (3.3% in the overlap with COMAP fields) of bright galaxies, broad absorption line quasars, and damped Ly α systems.

Objects in the eBOSS catalog are identified and their redshifts are determined simultaneously, through the fitting of templates spanning galaxies, quasars, and stars to the eBOSS spectra (stepping over redshift; Bolton et al. 2012). Spectra are optical, and at the redshifts of interest the templates to be fit typically include some combination of Ly α , H α and H β , and a range of optical and near-ultraviolet forbidden metal lines. The fit with the minimum reduced χ^2 value is used, provided $\chi_r^2 > 0.01$ and there is no other similarly confident fit determination greatly offset in redshift from the best-fit value. Visual follow-up is used to confirm the least confident identifications and to test the algorithmic determination.¹⁶ Tests of the classification algorithm on blank-sky spectra showed a $<2\%$ false positive rate—the eBOSS redshift




Figure 1. The redshift distribution of the eBOSS objects used for this stacking analysis.

determinations are generally considered highly trustworthy. Based on the pixel size of the detector, the uncertainty in the eBOSS redshifts is at most $\sim 207\text{ km s}^{-1}$, or $|\Delta z| \leq 0.01$ (Bautista et al. 2017; approximately one COMAP channel). We note that this uncertainty does not account for the systematic gas in- or outflows that may be present in quasars; we discuss this additional consideration in Section 4.2.3.

We remove duplicates and any object with an SDSS `NEGATIVE_EMISSION` warning from the eBOSS catalog, and then cut to the COMAP fields. In total, 243 eBOSS objects lie in the COMAP fields after these cuts. The average redshift of these objects is $z = 2.73$, with their redshift distribution as shown in Figure 1. Their spatial distribution is shown in Figure 2, against the COMAP Season 1 sensitivity.

3. Stacking Methods

Broadly, the stack is a 3D coaddition of the COMAP maps on the spatial and spectral positions of the quasars in the eBOSS catalog: we extract a 3D “cubelet” centered around the spatial position and redshifted CO(1–0) frequency of each source, and then combine them together into a single stacked cubelet containing the average luminosity of each source.

3.1. Determining Physical Quantities from the COMAP Data Cubes

The COMAP data cubes are provided in brightness temperature units, $T_{b,i,j,k}$, where i and j index voxels in directions perpendicular to the line of sight (spatial axes) and k indexes voxels parallel to the line of sight (spectral axis). The stacking pipeline begins by converting $T_{b,i,j,k}$ into velocity-integrated line flux, $(S\Delta\nu)_{i,j,k}$:

$$(S\Delta\nu)_{i,j,k} = \frac{2\nu^2 k_B T_{b,i,j,k}}{c^2} \Omega_{\text{vox}} \Delta\nu_k. \quad (1)$$

In this case, $T_{b,i,j,k}$ is the brightness temperature in the voxel, Ω_{vox} is the solid angle of the sky subtended by the voxel (which is constant across the map), and $\Delta\nu_k$ is the channel width

¹⁶ Other redshift values, in particular those based on specific emission lines, are also available. For a full list of the various methods for redshift determination, see: https://data.sdss.org/datamodel/files/BOSS_QSO/DR16Q/DR16Q_Superset_v3.html.




Figure 2. The distribution in space and redshift of the eBOSS objects used for this stacking analysis, plotted on the COMAP Season 1 median map rms for each field (see Foss et al. 2022).

in km s^{-1} in the k th frequency channel,

$$\Delta v_k = \frac{\Delta \nu}{\nu_{\text{obs},k}} c. \quad (2)$$

We follow Solomon et al. (1997) to calculate L'_{CO} in each voxel,

$$L'_{\text{CO},i,j,k} = \frac{c^2}{2k_B} (S\Delta v)_{i,j,k} \frac{D_{L,k}^2}{\nu_{\text{obs},k}^2 (1+z_k)^3}, \quad (3)$$

where $\nu_{\text{obs},k}$ is the observed frequency of the k th frequency channel, z_k is the redshift required to place the 115.27 GHz CO(1–0) line into that channel’s frequency range, and $D_{L,k}$ is the luminosity distance associated with that redshift. We perform this conversion before stacking because several of these values vary across the COMAP data cube.

3.2. Extraction of Cubelets

For each eBOSS quasar, we extract from the converted COMAP L'_{CO} cubes a 3D “cubelet” centered around the spatial position and redshifted CO(1–0) frequency of the quasar. These cubelets have dimensions $42' \times 42' \times 1.28 \text{ GHz}$ and contain CO line luminosity values $L'_{\text{CO},l,m,n,p}$, where l and m are in spatial directions, n is along the spectral axis, and p indexes each cubelet. These are purposefully large compared to the COMAP spatial and spectral resolution, so they can be visually searched for potential large-scale fluctuations. If any cubelet has its central voxel masked, or more than half of the voxels immediately adjacent to the central voxel in all three directions masked, it is excluded from the catalog of cubelets to be stacked.

3.3. Stacking

We then stack the CO emission by directly combining the full 3D cubelets associated with each individual galaxy, voxel-by-voxel. This is done to allow for additional testing of the fidelity of any potential detections. To account for the inhomogeneous noise response across the map (Section 2.1), values are weighted by their rms noise using inverse-variance weighting. Thus, the line luminosity in each voxel of the

stacked cubelet is

$$L'_{\text{CO},l,m,n} = \frac{\sum_p L'_{\text{CO},l,m,n,p} / \sigma_{l,m,n,p}^2}{\sum_p 1 / \sigma_{l,m,n,p}^2}, \quad (4)$$

for cubelet voxel line luminosities $L'_{\text{CO},l,m,n,p}$ with rms noise $\sigma_{l,m,n,p}$. The uncertainty in each voxel of the resulting stacked cubelet is then

$$\sigma_{l,m,n} = \sqrt{\frac{1}{\sum_p 1 / \sigma_{l,m,n,p}^2}}. \quad (5)$$

We additionally attempt to lessen any interactions between asymmetries in the position of catalog galaxies in the data cubes and the inhomogeneous COMAP noise response by rotating each cubelet randomly in its spatial axes by either 0° , 90° , 180° , or 270° before performing the coaddition.

3.4. Obtaining Luminosity and Density Limits

Once a full 3D stack cubelet is obtained, we calculate the average line luminosity of the included quasars in two steps. First, we calculate the line luminosity in each frequency channel, $L'_{\text{CO},n}$, by summing over the central three spatial pixels (spaxels) of the stacked cubelet (the region indicated with a black box in Figures 3 and 4),

$$L'_{\text{CO},n} = \sum_{l,m=-1}^{+1} L'_{\text{CO},l,m,n} \quad (6)$$

which leaves us with a 1D frequency spectrum. This 3×3 region corresponds to a $6'$ square aperture, roughly 1.5 times the size of the COMAP main beam, meaning that most signal from any catalog object located anywhere in the central spaxel will not be spread outside this aperture by the beam. Thus, we avoid excluding any potential signal from the spectrum, while keeping the aperture as small as possible to mitigate the effects of beam dilution and maximize our signal-to-noise ratio (S/N).

Finally, to determine the average CO luminosity of the stacked catalog objects L'_{CO} , we sum the central seven




Figure 3. A diagram displaying the methodology used for this stacking analysis. From the full 3D COMAP data cube (left), smaller 3D cubelets are cut out centered on the position of each eBOSS object (center). These cubelets are then averaged into a single 3D stack (top right), which is used to determine the average COMAP spectrum of the eBOSS objects (bottom right).




Figure 4. A noiseless simulated stacking detection, normalized to the expected integrated luminosity. (a) The central three frequency channels coadded to provide a view of the spatial distribution of signal. The black box in the center indicates the smaller aperture used to calculate single stack values. (b) The same spatial representation of the stack, smoothed with a $4'$ Gaussian kernel to approximate the COMAP main beam. (c) The frequency spectrum of the stack. The value in each spectral channel is determined by summing over the spatial aperture indicated in panel (a). The three frequency channels making up the stack aperture are highlighted in gray. While the simulation is noiseless, neighboring halos may show up in some objects' cutouts, appearing as small fluctuations in the stacked spectra.

frequency channels of the stacked spectrum,

$$L'_{\text{CO}} = \sum_{n=-3}^{+3} L'_{\text{CO},n}. \quad (7)$$

The seven-channel width corresponds to 218.75 MHz, or 2065 km s^{-1} at the average redshift of the galaxy catalog objects. This chosen spectral width allows for a velocity offset between the CO emission and the optical lines used to

determine the eBOSS redshifts, which can be quite significant (see Section 4.2.3). Instrumental aliasing between adjacent COMAP science channels is negligible (Lamb et al. 2022), so broadening considerations in the spectral axes are almost entirely astrophysical.

We use the average line luminosity to determine an average molecular gas density in the stack. We use the Bolatto et al. (2013) CO-to-H₂ conversion factor of $3.6 M_{\odot} (\text{K km s}^{-1} \text{pc}^2)^{-1}$ (calculated for the Milky Way) to determine the molecular gas mass associated with our stacked line luminosity, and then convert this mass into a molecular gas density ρ_{H_2} by dividing by the *comoving* volume of the stack aperture.

This is the simplest possible version of a stack, ignoring other available information such as the significance of the catalog objects’ detections with eBOSS, or the optical luminosity of the catalog objects, both of which could also be used for weighting (see, for example, Sinigaglia et al. 2022). Functionally, this means assuming a δ function for the luminosity function. While we do plan to explore refinements to this methodology in the future, we assume for now that the noise in the Season 1 COMAP data cube is much greater than any variations across the eBOSS catalog.

4. Data Verification

4.1. Bootstrapped Error Analysis

While we believe the noise in our maps to be Gaussian distributed (“white” noise: see Section 2.1, also Foss et al. 2022) over the entire map, this stacking analysis introduces a new sampling strategy, and it is possible that the white-noise assumption does not hold under this new sampling. We therefore perform a spatial “bootstrap” test on each stack to confirm the noise distribution in our maps and to determine the uncertainty in each of the stacked values.

This is done by binning the actual eBOSS catalog by COMAP field and by redshift (into three redshift bins), and then generating an artificial catalog with the same overall redshift and field distribution as the real catalog but with random 3D positions. Variations in catalog number density between fields, as well as instrumental effects such as the changing main beam with frequency and potential systematic errors correlating to on-sky position, will thus appear both in the bootstrapped stacks and the real stack. The LIM map is then stacked on the artificial catalog. This is repeated for 10,000 different random artificial catalog iterations for each real stack run to fully characterize the stack noise response.

4.2. Constant-luminosity Simulation Tests

4.2.1. Simulation Methodology

As an end-to-end test of the stacking pipeline, we create mock COMAP observations of simulated cosmic CO signal, generate mock galaxy catalogs associated with these cubes, and perform the stacking analysis on these simulated data. Mock CO emission is generated by painting CO luminosities onto a dark matter halo catalog from peak-patch *N*-body simulations (Stein et al. 2019). Because we are looking to check the efficacy of the stacking pipeline rather than make realistic cosmological predictions, we paint each halo with the same luminosity regardless of its mass. We artificially broaden the CO emission line in each halo by an effective velocity v_{eff} of 319 km s^{-1} (see Section 4.2.3 below), again to more easily test

the effects of the stacking pipeline. To generate mock COMAP data cubes, we additionally beam-smooth the emission in the spatial axes by convolving the mock cubes with a 2D Gaussian kernel with a $4'$ FWHM to approximate the COMAP main beam.

The artificial CO emission cube generated from these simulations is then converted into a time stream, scaled by a factor of 1000 to simulate an (unrealistically) strong detection, and added into real time-ordered data from the first COMAP observing season (N.O. Stutzer et al., 2024 in preparation). The positions of the simulated halos should have no correlation with any existing real cosmic structure, so this is an excellent way to simulate observations with the actual noise structure that COMAP observes. This mock time stream is passed through the COMAP pipeline, so any effects of the several pipeline filters (Foss et al. 2022) will show up in the mock data cube.

In addition to the mock COMAP data cubes, we generate an artificial galaxy catalog by randomly selecting 1000 dark-matter halos from the most massive 30% of halos in the peak-patch catalog (dark-matter masses between $9 \times 10^{13} M_{\odot}$ and $4 \times 10^{12} M_{\odot}$), and taking these halos to be the objects emitting brightly enough in some other galaxy tracer to be detected in a traditional galaxy survey. Only the 3D positions of the halos are required for the stacking analysis, so we do not calculate any other mock parameters for the galaxy survey.

4.2.2. Spatial Attenuation due to COMAP Beam

We then run the stacking pipeline on 10 different realizations of the mock data cube and galaxy catalog. This is done principally to confirm that the COMAP pipeline is not attenuating the stacking signal in any unexpected ways, and to determine the size of the stacking aperture discussed in Section 3.4. An example of a resulting simulated stack, shown in Figure 4, confirms that the $6' \times 6'$ spatial aperture we selected indeed encloses the majority of the stacked flux, without incorporating much empty space. The line luminosity output by this stack is 88.5% of the input value, meaning the input signal is attenuated by a factor of 1.13. This attenuation is due to the COMAP main beam spreading signal outside of the stack’s 3×3 spatial aperture (Figure 4). Instead of broadening the aperture further and incorporating more empty space into the stack, which would affect our sensitivity, we multiply the values output by the stack by 1.13 to correct for this attenuation.

4.2.3. Spectral Attenuation due to Offsets between Optical and CO Redshifts

Although the cataloged eBOSS redshifts are considered very precise, high-redshift quasars often contain large systematic outflows/inflows that cause velocity offsets between optical emission lines (such as those used by eBOSS) and the CO emission, which traces the cold gas of the galaxy itself (Banerji et al. 2017; Herrera-Camus et al. 2020; Bischetti et al. 2021). This is an important consideration, as even small redshift uncertainties can cause significant attenuation in the signal of LIM/galaxy catalog cross-correlation analyses (Chung et al. 2019, 2021). The shot component of the cross-power spectrum (i.e., the cross power at small scales, the most direct analog to a stacking analysis) is where this effect is at its worst, although COMAP is most sensitive to larger spatial scales than those

where attenuation due to redshift uncertainty is truly catastrophic.

To quantify systemic offsets between the cataloged redshift for our eBOSS sample (z_{opt}) and the redshift of the CO emission (z_{CO}), we assemble a sample of quasars that have (clean) spectroscopic optical redshifts in the COMAP redshift range from SDSS and have been individually detected in a molecular gas tracer line. These come from three studies:

1. Eight objects from Hill et al. (2019), who targeted CO(3–2) in 13 QSOs selected from the Keck Baryonic Structure Survey (KBSS) for additional cosmic web study with SCUBA-2. CO measurements used NOEMA (the Northern Extended Millimeter Array; Chenu et al. 2016). The z_{opt} values for these objects are specifically from eBOSS.
2. Two objects from the sample of nine QSO MUSEM quasars targeted for CO(6–5) and CO(7–6) observations with the SEPIA180 (Belitsky et al. 2018) receiver on APEX in Muñoz-Elgueta et al. (2022). These objects also have optical redshifts from eBOSS.
3. Two quasars from Riechers et al. (2011), who targeted CO(1–0) in five quasar host galaxies using both the Extended Very Large Array (EVLA), and the Green Bank Telescope. Optical redshifts are from SDSS for these quasars.

In each case, we exclude objects with no available eBOSS/SDSS redshift or that fall far outside the COMAP redshift range. We also exclude objects with SDSS `NEGATIVE_EMISSION` warnings, for consistency with our own generated eBOSS catalog. Over this entire compiled catalog, the quasars have an average molecular gas emission-line FWHM of 319 km s^{-1} .

We plot z_{opt} against z_{CO} for each of these quasars in Figure 5. We find that z_{opt} is systematically smaller than z_{CO} with a mean value of $(z_{\text{CO}} - z_{\text{opt}})/(1 + z_{\text{opt}}) = 0.00397 \pm 0.00408$. This offset does not appear to be correlated with either redshift or extrapolated CO(1–0) line luminosity. At the average redshift of our eBOSS quasar catalog, it corresponds to a negative frequency offset of 122 MHz, or 1185 km s^{-1} (~ 4 COMAP spectral channels). This is in good agreement with observed bulk velocity offsets in quasars, which are typically $\sim 1000\text{--}2000 \text{ km s}^{-1}$ (Tytler & Fan 1992; Orellana et al. 2011).

In addition to exploring corrections in the bulk velocity offset between quasar gas phases, we also investigate the scatter in the relation. The standard deviation in $(z_{\text{CO}} - z_{\text{opt}})/(1 + z_{\text{opt}})$ in the individually detected quasar sample is $\sigma_{\delta z} = 0.00408$ from this effect alone, which is large enough to move a portion of the signal from quasars in the stack outside of our defined seven-channel frequency aperture. We note that the variety of CO transitions used in the molecular gas measurements of the plotted quasars could create a larger scatter than a pure CO(1–0) analysis, as they trace slightly different phases of (dense) gas (e.g., Muñoz-Elgueta et al. 2022).

We perform a separate analysis to determine signal attenuation due to velocity offsets in quasars between the optical lines measured by eBOSS and the CO(1–0) line. This is addressed in detail in Appendix A. We find that after this effect is accounted for, the input signal is attenuated by an additional factor of 1.58. As above, we multiply the values output by the stack by this factor to correct for this source of signal




Figure 5. The optically determined redshift values compared with CO-based redshifts for a selection of quasars individually detected in CO in the COMAP redshift range. The color map is the extrapolated CO(1–0) luminosity of each object, and the line where $z_{\text{opt}} = z_{\text{CO}}$ is shown for reference. In the bottom panel we indicate the mean offset of the sample with a dotted line and the 1σ scatter in that scatter with the shaded region.

attenuation. Combined with the spatial attenuation discussed above, this is an overall attenuation factor of 1.79.

5. Results

5.1. CO(1–0) Stack Results

The final COMAP/eBOSS stack is shown in Figure 6. We find, at our current sensitivity level, no detection of CO emission associated with the 243 eBOSS quasars in the COMAP fields. The results of the COMAP/eBOSS bootstrap test (discussed in Section 4.1) are shown in Figure 7. As expected, the maps remain consistent with Gaussian noise centered at $T_b = 0 \text{ K}$ even in this stacking analysis. As this is a more robust calculation of the distribution of stack values expected from random noise, we use the 95% confidence region calculated from a Gaussian fit to these bootstrapped L'_{CO} values as the uncertainty in our stack.

We set a 95% upper limit on the CO frequency-integrated line flux of $S\Delta\nu \leq 0.129 \text{ Jy km s}^{-1}$, corresponding to a limit on the CO line luminosity of $L'_{\text{CO}} \leq 1.26 \times 10^{11} \text{ K km pc}^2 \text{ s}^{-1}$ and a limit on the average molecular gas density in regions containing a quasar of $\rho_{\text{H}_2} \leq 1.52 \times 10^8 \text{ M}_\odot \text{ Mpc}^{-3}$.

Additionally, as the quasar velocity offset discussed in Section 4.2.3 is large enough to move any flux from our stack outside our chosen stack aperture, we correct for it by applying a 122 MHz frequency offset to our stack. We again find no significant CO emission, with a new upper limit of $L'_{\text{CO}} \leq 8.69 \times 10^{10} \text{ K km pc}^2 \text{ s}^{-1}$ (Figures 10 and 11). Accounting for this frequency offset thus does not meaningfully change our results at our current level of sensitivity.




Figure 6. The result of stacking the COMAP data cubes on the 3D positions of eBOSS quasars. (a) The central seven frequency channels coadded to provide a view of the spatial distribution of signal. The black box in the center indicates the smaller aperture used to calculate single stack values. (b) The same spatial representation of the stack, smoothed with a $4'$ Gaussian kernel to approximate the COMAP main beam. (c) The frequency spectrum of the stack. The value in each spectral channel is determined by summing over the spatial aperture indicated in panel (a). The seven frequency channels making up the 3D stack aperture are highlighted in gray. All distance scales are given as proper values.

5.2. Foreground Lines

A major concern for LIM experiments is contamination of signal at the redshift of interest by other (primarily foreground) spectral lines, which may also be redshifted into the frequency range of the experiment by coincidence (Cheng et al. 2016; Lidz & Taylor 2016). CO LIM experiments suffer less from this type of contamination than LIM experiments targeting other spectral lines because the (primarily molecular) lines that fall into our frequency range are expected to be faint compared to CO(1–0); (Cleary et al. 2022). Joint analyses such as stacking, which require emission in two different spectral lines to be detected from the target sources, are also expected to be fairly robust against foreground contamination (Chung et al. 2019). Because joint analyses are targeted toward specific emission lines, however, they present a unique opportunity for testing theoretical predictions that there will be minimal foreground emission present. Additionally, non-CO molecular emission lines are also interesting tracers of molecular gas in their own rights, and worth investigating.

Thus, we perform the stacking analysis on four other molecular emission lines that will likely fall into the COMAP frequency range, in addition to the $^{12}\text{CO}(1-0)$ spectral line for which COMAP was designed. These lines are the ^{13}CO isotope’s $J = 1 \rightarrow 0$ rotational transition, the HCN $J = 1 \rightarrow 0$ transition, the CS $J = 2 \rightarrow 1$ transition, and the CN^- radical’s $J = 1 \rightarrow 0$ rotational transition. These lines are discussed in more detail in Appendix C or in Breyse et al. (2015). We make




Figure 7. A histogram of the L'_{CO} values returned from 10,000 stacks performed on random 3D positions in the COMAP data cubes, with the same spectral and field distribution of objects as the real eBOSS catalog. A Gaussian fit to the histogram is shown in orange. The 66% confidence interval of the fitted Gaussian (i.e., 1 standard deviation about the fitted mean) is shown in dark gray, and the mean value is shown as a dotted gray line (see Section 3.4).

no detection of any foreground line. The resulting upper limits are shown in Table 2, and representations of the stacked cubelets are shown in Appendix C.

Table 2
Upper Limits Obtained by Stacking the Data Cube on the CO(1–0) Line and Expected Foreground Emission Lines

Spectral Line	No. of Objects	ν_{rest} (GHz)	Redshift Range	$S\Delta\nu$ 95% Upper Limit (Jy km s ^{−1})	L' 95% Upper Limit (K km pc ² s ^{−1})
¹² CO(1–0)	243	115.27	2.4–3.4	≤0.129	≤1.26 × 10 ¹¹
¹² CO(1–0) with offset ^a	244	115.15	2.4–3.4	≤0.270	≤8.69 × 10 ¹⁰
HCN(1–0)	595	88.63	1.6–2.4	≤0.196	≤6.21 × 10 ¹⁰
CS(2–1)	479	91.98	1.7–2.5	≤0.285	≤1.18 × 10 ¹¹
¹³ CO(1–0)	311	110.20	2.2–3.2	≤0.236	≤1.14 × 10 ¹¹
CN [−] (1–0)	269	113.49	2.3–3.4	≤0.316	≤1.29 × 10 ¹¹

Note.

^a Shifted by the mean velocity offset of z_{eBOSS} from z_{CO} for the individually detected quasars in Section 4.2.3 (see Figure 5).

5.3. Forecasts

As we are working with only the first year of data from COMAP’s planned multiyear Pathfinder survey, we also project the stack sensitivity after five years of observing. The eBOSS survey has finished, so we only forecast improvements to the stack sensitivity arising from deeper COMAP data, instead of other possible stack improvements such as including more catalog objects or refining their redshift or positional accuracy. Under this assumption, the stack will improve by an identical factor to the sensitivity of the COMAP data cubes themselves.

Detailed projections were made by Foss et al. (2022) to forecast improvements to COMAP’s sensitivity. Primarily, this improvement will be due to the obvious increase in total integration time by the end of the survey, but the Pathfinder’s first season of operation was partially used for commissioning and refining the survey design, so we also expect significant improvements to the overall percentage of usable data. The effect of these improvements on COMAP’s power spectrum sensitivity specifically are broken down in Table 2 and collated in Equation (26) of Foss et al. (2022). Almost all of them apply to the stacking analysis, but some are power spectrum specific:

1. The Foss et al. (2022) forecasts are made by projecting the sensitivity of a single COMAP field to the sensitivity of all three fields. As we are already using the full COMAP on-sky area in the stack analysis, we do not include the factors of $\sqrt{3}$ corresponding to this projection.
2. Improvements to the filtering and map-making stages of the COMAP pipeline should result in these stages removing roughly 10% less actual signal from the maps (this is the transfer function efficiency, E_{TF}), but will mostly act on the large angular scales to which the stack is not sensitive. We therefore do not multiply by E_{TF} when forecasting for the stacks.
3. The process of calculating the power spectrum from the completed intensity map involves generating many jackknifed cross-spectra and only keeping those that pass quality tests. The retention efficiency of these steps, parameterized by Foss et al. (2022) as $E_{\chi^2_{\text{PK}}}$ and E_{split} , will likely improve in future COMAP seasons as the earlier pipeline stages become more robust. Since these data cuts use statistics calculated from already generated intensity maps, they only affect the power spectrum itself and we do not include them here.

All other efficiency improvements apply at the map level and thus will apply to the stack. The forecast stacking sensitivity for the full COMAP Pathfinder survey then becomes

$$\sigma_{\text{stack}}^{5\text{yr}} = \frac{\sigma_{\text{stack}}^{\text{S1}}}{\sqrt{D_{\text{map}}^{5\text{yr}}}}, \quad (8)$$

where

$$D_{\text{map}}^{5\text{yr}} \equiv \frac{T^{\text{S1}} E_{\text{tot}}^{\text{S1}} + (5 \cdot 365 - T^{\text{S1}}) E_{\text{tot}}^{\text{proj}}}{T^{\text{S1}} E_{\text{tot}}^{\text{S1}}}. \quad (9)$$

T^{S1} is the calendar duration of Season 1, 440 days. Both instances of E_{tot} indicate total observing efficiency, which combines bare on-sky observing efficiency E_{obs} with the fraction of data retained after all pipeline steps E_{data} . For the exact breakdown of these quantities, see Foss et al. (2022). $E_{\text{tot}}^{\text{S1}}$ is the actual Season 1 efficiency, and $E_{\text{tot}}^{\text{proj}}$ is the projected value after five years. Adjusting the values calculated by Foss et al. (2022) to exclude the power spectrum only factors discussed above, these are 9.15% and 33.2%, respectively. We thus project that the sensitivity of the COMAP/eBOSS stack will improve by a factor of $D_{\text{map}}^{5\text{yr}} = 25.2$, corresponding to a CO(1–0) flux sensitivity $\sigma_{\text{stack}}^{5\text{yr}}$ of 0.026 Jy km s^{−1} at the 95% level.

While we have chosen in this work to focus specifically on the CO properties of quasars, we note that more general high-redshift spectroscopic catalogs can be used for stacking if the goal is to extract the CO properties of the Universe as a whole. In particular, we have already made projections (Chung et al. 2019, 2022) for the cross-correlation of COMAP data with the the Hobby–Eberly Telescope Dark Energy EXperiment (HETDEX; Gebhardt et al. 2021) catalog of Ly α emitting galaxies. We plan in future to perform a stacking analysis on the same catalog. To first order, the stack sensitivity increases with the number of objects N in the spectroscopic catalog by a factor of $\sqrt{N}$, so we project significant sensitivity increases with this larger catalog. We note, however, that different galaxy tracers are subject to different biases, and any anticorrelation between the brightness of the tracers used and the CO brightness of a given object will affect this prediction; we plan to quantify this effect in future works.

6. Discussion

In general, how we interpret the results of LIM/galaxy survey stacking analyses depends on the relative contribution

of the CO emission associated with the catalog galaxy compared to that of other sources within the COMAP beam. We can bracket the range of possible interpretations by considering the two extreme cases: (i) the galaxies traced by our chosen survey dominate the brightness temperature of their surroundings (making the stack simply a measurement of the average CO luminosity, L'_{CO} , of the survey objects), and (ii) the cataloged galaxies are unspectacular themselves, and instead tend to fall into overdense regions of the large-scale structure (making the stack an measurement of the average CO luminosity of bright regions of the Universe, which can be converted into their molecular gas density, ρ_{H_2}).

Due to the dearth of CO observations of quasars at cosmic noon, it is not yet clear which of these cases is closer to the truth; we discuss the available evidence below. Neither assumption, however, invalidates the other in the context of the upper limits we present here. Quasar companions will not be distinguishable from the objects themselves at COMAP resolutions, meaning any nearby objects will cause a source confusion effect pushing the measured L'_{CO} value upwards. Conversely, abnormal amounts of CO emission from the quasars themselves will drive the local ρ_{H_2} upwards. In both cases, therefore, the stack values should likely be treated as upper limits even once a confident detection is made at this level of analysis.

6.1. CO in Quasars

The response of the molecular gas in quasar host galaxies to the extremely active SMBH at their centers is likely extremely complicated, especially at cosmic noon. The abundances of both quasars and gas-rich, submillimeter-bright galaxies (SMGs) seem to peak during this epoch, suggesting a link between the two (Simpson et al. 2012). Indeed, both theoretical and observational arguments have suggested that the SMBH activity driving quasars is triggered from normal SMGs by mergers rich in cold and cool gas (Sanders et al. 1988; Hopkins et al. 2008; Brusa et al. 2018; Herrera-Camus et al. 2020; Bischetti et al. 2021), suggesting that, at least in their initial phases, quasars should be extremely rich in cold gas.

Whether that cold gas survives the subsequent barrage of heat and pressure from its central SMBH, however, is still uncertain. SMG studies of molecular gas have shown that the presence of an AGN does not seem to significantly affect the gas properties of its host (Bothwell et al. 2013; Banerji et al. 2017), and observed gas masses in quasars themselves follow typical SMG values closely (Hill et al. 2019), but other observations have shown that CO spectral-line energy distributions (SLEDs) in quasar hosts tend toward higher excitation states than normal (e.g., Brusa et al. 2018). At least for the purposes of COMAP, which measures the lowest CO excitation state, this would read as damped CO emission in quasar hosts. Additionally, many separate quasar studies have observed large-scale outflows of all gas phases from the host galaxy, although a fraction of this gas may remain in the molecular phase and thus be observable with COMAP (Bothwell et al. 2013; García-Burillo et al. 2014; Feruglio et al. 2015; Brusa et al. 2018).

In the limiting case that much of the CO luminosity contributing to the LIM maps at the quasar positions comes from the quasars themselves, then the upper limit on L'_{CO} determined from our COMAP/eBOSS stack becomes a direct upper limit on the average CO line luminosity of the 243

eBOSS quasars included in the stack. Under this analysis, the L'_{CO} value can only ever be an upper limit, because the cosmic volume covered by the stack is large ($9 \text{ Mpc} \times 9 \text{ Mpc} \times 2065 \text{ km s}^{-1}$, in proper units) compared to an individual quasar, so it is unlikely that the quasars are the only sources contributing to the CO luminosity of the region. Because we are reporting an upper limit, however, this remains valid.

The measured average L'_{CO} value is shown in Figure 8, plotted against other measures of quasar line luminosity in a similar redshift range for comparison. These include the sample of 13 QSOs from KBSS, detected in CO(3–2) with NOEMA by Hill et al. (2019); the sample of nine objects from the QSO MUSEUM, surveyed with APEX in CO(6–5) and CO(7–6) by Muñoz-Elgueta et al. (2022); and the sample of five QSOs detected in CO(1–0) by Riechers et al. (2011) using the Extended Very Large Array and the Green Bank Telescope.

The Hill et al. (2019) CO(3–2) measurements are converted into CO(1–0) line luminosities using the conversion factor $L'_{\text{CO}(3-2)}/L'_{\text{CO}(1-0)} = 0.97 \pm 0.19$ from Carilli & Walter (2013). In several cases, a single KBSS object was found to be associated with two different CO sources slightly offset on-sky. For each of these objects, we sum both CO sources together to obtain the plotted L'_{CO} value, as both sources would fall well into the COMAP beam. Muñoz-Elgueta et al. (2022) converted their higher- J CO measurements into molecular gas masses directly by fitting to CO SLEDs, incorporating also the [C I] luminosity. In order to determine CO(1–0) line luminosities associated with the objects, we extrapolate their calculated M_{H_2} values to L'_{CO} values using the Milky Way α_{CO} conversion factor of $3.6 M_{\odot} (\text{K km s}^{-1} \text{ pc}^{-2})^{-1}$ (Bolatto et al. 2013), for consistency with our own previous analyses. These objects are mostly nondetections, and the SLED fitting returned a range of M_{H_2} values, so these ranges are what we plot in Figure 8. Riechers et al. (2011) report CO(1–0) line luminosities directly.

Figure 8 shows that the COMAP/eBOSS stacking upper limit is already comparable with the line luminosities of (the brightest) individually surveyed quasars in its redshift range, even using only the first season of COMAP data. In some cases, the COMAP value is actually a stricter limit, particularly for the APEX objects (Muñoz-Elgueta et al. 2022). This illustrates a powerful potential application for LIM—directly detecting CO in individual objects becomes prohibitively expensive very quickly at these redshifts. A statistically significant survey would require integration times that are not practical on high-demand community instruments. LIM-based stacking analyses, able to efficiently survey large samples of galaxies, could thus act as an important complement to these individual detection-based surveys by placing constraints on the ensemble emission from many galaxies. The minimum spatial scales accessible from an LIM instrument are large compared to galaxy sizes, meaning such analyses would have to remain upper limits even if a detection were made, but these still have the potential to be extremely enlightening.

Under this interpretation, our forecasted COMAP Pathfinder sensitivity to CO luminosity from a stack on the same 243 eBOSS quasars after five years of observing falls below the mean luminosity of the individually detected quasars in Figure 8, meaning it has the potential to be genuinely constraining—we will likely be able to tell with COMAP Y5 if these three samples are representative of the eBOSS quasar population as a whole. If we are able to constrain the L'_{CO} of eBOSS quasars to this degree, it will provide significant insight




Figure 8. The COMAP/eBOSS 95% upper limit on L'_{CO} (the black bar), plotted against quasars detected individually in CO from Riechers et al. (2011), Hill et al. (2019), and Muñoz-Elgueta et al. (2022). Because Muñoz-Elgueta et al. (2022) report a range of molecular gas mass values M_{H_2} from SED fitting of multiple lines, we show the L'_{CO} values determined from those values as a range using vertical bars. Upper limits are shown as inverted triangles. The extent of the COMAP/eBOSS bar indicates the redshift range to which COMAP is sensitive, and the marker is located at the mean redshift of the eBOSS objects on which we stack. A forecast for the stack sensitivity at the end of the full COMAP Pathfinder Survey is plotted in gray (see Section 5.3).

into the complex feedback cycles present in quasars, and thus into star formation processes as a whole at cosmic noon.

6.2. Cosmic Molecular Gas Density

The other limiting case—that quasars trace large-scale structure but do not themselves dominate the local CO emission—is perhaps better supported, as it follows from hierarchical structure formation. SMBH mass and galaxy mass have been shown several times to be correlated, meaning quasars will preferentially be found in massive galaxies that are likely to be at the center of large dark-matter halos (e.g., Gebhardt et al. 2000; Hopkins et al. 2008). This argument is additionally supported by clustering measurements of both BOSS and SDSS Stripe 82 quasars (White et al. 2012; Timlin et al. 2018), as well as observations of the individual objects shown in Figure 8: a quarter of the Hill et al. (2019) objects, for example, have nearby CO-bright companions (also Banerji et al. 2017, 2018; Bischetti et al. 2021; Decarli et al. 2021). The argument that quasars require major mergers to ignite also necessitates at least one nearby massive galaxy in their recent past.

In this limit, where the quasar host itself is contributing negligibly to a local overdensity of CO emission, we can use the stack’s upper limit on T_b to calculate the average molecular gas density ρ_{H_2} in bright regions at the COMAP redshifts. As the stack is only sensitive to a single, ~ 9 (comoving) Mpc spatial scale (the size of our chosen aperture), this value is most similar to a random “shot” power component of the ρ_{H_2} determined from a power spectrum analysis. We plot this value in Figure 9, alongside several other measurements of the same value made using a variety of techniques, including COMAP’s

own auto-spectrum based early science constraint (Chung et al. 2022). These other measurements include:

1. COPSS (Keating et al. 2016), an LIM survey targeting the same CO(1–0) transition as COMAP. COPSS used the Sunyaev–Zel’dovich Array, and is subject to interferometric insensitivity to large scales. The ρ_{H_2} value from COPSS is thus based mainly on the shot-noise component of the CO power spectrum, converted to a molecular gas density using the Milky Way α_{CO} value of $3.6 M_{\odot} (\text{K km s}^{-1} \text{pc}^{-2})^{-1}$ (the same value we adopt in this work). Two other COPSS-determined values (Keenan et al. 2022) are plotted, both determined using joint analyses combining COPSS data with (after correcting for main beam weighting) 145 spectroscopically confirmed cataloged galaxies. Keenan et al. (2022) cross-correlated the COPSS data with a data cube obtained by gridding the catalog objects in 3D space (yielding one limit on ρ_{H_2}) and fit the resulting cross-power spectrum jointly with both the auto-power spectrum of COPSS alone and the galaxy–galaxy power spectrum of their galaxy catalog (yielding a second ρ_{H_2} limit).
2. ALMACAL (Klitsch et al. 2019) used archival Atacama Large Millimeter/submillimeter Array (ALMA) calibrators as background sources to search blindly for CO absorption lines. As these calibrator sources must be observed frequently, ALMACAL boasts >1500 hr of observations on a community instrument over a wide enough area to not be limited by cosmic variance. However, these calibrators do not extend past $z \sim 2$.
3. The Millimeter-wave Intensity Mapping Experiment (mmIME) used a combination of archival and targeted




Figure 9. The upper limit on the average molecular gas density determined from the COMAP/eBOSS stack (black line), compared to other ρ_{H_2} measurements made using different survey techniques. The upper limit from the COMAP Season 1 power spectrum analysis is shown in purple. Also shown is the COPSS LIM survey’s measured value (Keating et al. 2016), as well as upper limits from COPSS joint analyses with galaxy surveys (Keenan et al. 2022). Shaded boxes are measurements made by large-scale traditional surveys COLDz (Riechers et al. 2019), ASPECS (González-López et al. 2019), and PHIBBS2 (Lenić et al. 2020). We show the forecast 2σ sensitivity level of the COMAP survey after five years stacked on eBOSS as a gray upper limit.

ALMA and Atacama Compact Array (ACA) observations at frequencies tracing multiple CO transitions (Keating et al. 2020). As with COPSS, these constraints are primarily constraints on the spectral shot power.

4. Finally, the shaded boxes in Figure 9 correspond to ρ_{H_2} measurements made using traditional (i.e., resolved and targeted) galaxy surveys. These include ASPECS (González-López et al. 2019; Uzgil et al. 2019) and PHIBBS2 (Lenić et al. 2020), each tracing multiple CO rotational transitions, as well as COLDz (Riechers et al. 2019), which traces exclusively CO(1–0) at $z \sim 3$. While each of these traditional surveys provide ρ_{H_2} values that seem like much better constraints than any of the LIM-based numbers, they cover much smaller survey volumes and are thus extremely susceptible to cosmic variance effects. Additionally, as traditional surveys, they are insensitive to any emission below their detection limits, and so will miss any contribution to the total gas density from fainter objects, likely underestimating the total gas density.

Clearly, these ρ_{H_2} measurements were made using a diverse range of different techniques, so some of the scatter in Figure 9 can be attributed simply to the biases associated with each strategy. There is also likely additional scatter present due to variations in the CO-to- H_2 ratio, both astrophysically and in the conversion factor α_{CO} chosen for each analysis, as well as the relative intensity of the different CO line transitions, which will also change based on environment.

In our case, while we are indeed subject to the assumptions associated with our choice of α_{CO} and the bias of CO emission as a galaxy tracer, we are also only investigating the regions of space associated with eBOSS quasars, so the stacking analysis necessarily introduces a secondary bias factor. As discussed above, regions of this size surrounding quasars are likely to be more dense than the Universe as a whole, biasing the stacked ρ_{H_2} value upwards because of clustering. The upper limit we report here is therefore likely an overestimate—an analysis that

fully explored clustering effects would be more constraining. Because we report an upper limit we find this to be acceptable. We will explore the effects of clustering more thoroughly in future work.

Due to this additional bias, the ρ_{H_2} value from the eBOSS stack is more similar to a cross-correlation analysis, such as the early Wilkinson Microwave Anisotropy Probe cross-correlations with SDSS and BOSS quasars (Pullen & Hirata 2013; Pullen et al. 2013). Unlike these cross-correlation analyses, the stack’s ρ_{H_2} value is only sensitive to a single spatial scale (the scale of the 3D stack aperture, ~ 9 Mpc). We note that, because of our definition of a voxel in the COMAP data cubes, this scale does vary across the cube, and the ρ_{H_2} values being stacked are thus probing variable volumes. We explore this more thoroughly in Appendix B.

Even with these caveats, however, our predicted value is comparable to the values returned from the different COPSS joint analyses, reinforcing the viability of stacking as an LIM tool introduced by Keenan et al. (2022). More promising is our forecast sensitivity for the COMAP Y5/eBOSS stack, which actually falls into the regions reported by traditional galaxy surveys. As these galaxy survey values are likely underestimates of the actual cosmic H_2 density, our stacking analysis should detect emission from even very pessimistic models of cosmic CO, provided our forecasts hold true. As we already predict that our autocorrelation analysis will allow us to discriminate between CO models by COMAP Y5 (Chung et al. 2022), stacking will provide another valuable tool in characterizing the amount and properties of molecular gas emission at cosmic noon.

Additionally, quasar catalogs are far from the only spectroscopic information available on astrophysical objects at our redshifts. As discussed in Section 5.3, increasing the number of objects in our galaxy catalog will increase the sensitivity of our stack, but introducing new objects detected using different tracers will also provide new opportunities for analysis. It will be very interesting to compare, for example, the CO properties

of the Ly α emitting galaxies in the HETDEX survey (Gebhardt et al. 2021) to the CO properties of the eBOSS quasars. Galaxy evolution is an inherently multitracer problem, so there will be significant analysis benefits to collating as many different surveys (whether traditional resolved surveys or LIM experiments) of the same region of space as possible. As the LIM field progresses, there will be more and more opportunities to do so.

7. Summary and Conclusions

By stacking CO line intensity maps from COMAP’s first year of science operations on the 3D positions of quasars from the SDSS eBOSS survey, we have obtained a new upper limit on the cosmic abundance of molecular hydrogen gas. We describe the methodology behind the stack in detail, including the COMAP rms-based weighting scheme we use when coadding and the bootstrapped error analysis technique we use to quantify uncertainties. We use the stack to search additionally for interloper emission from foreground galaxies in four molecular lines, finding no evidence of a detection in any case. We can interpret our stacked upper limit as a constraint on the CO emission from either the cataloged quasars themselves, or the regions of the Universe surrounding those quasars. Likely, the most realistic interpretation is some combination of both cases. Additionally, we forecast stacking results for the COMAP Pathfinder survey after five years of observing.

In the limit that any potential signal in the stack is dominated by CO emission from the eBOSS quasars themselves, we treat our measurement as the average CO luminosity of these objects. We compare this average luminosity to resolved CO measurements of quasars in the COMAP redshift range and find that our upper limit already probes CO luminosities fainter than those of the brightest objects observed to date. The quasar studies to which we compare typically required hours of observations using large community facilities to detect CO in only a handful of objects. Determining the CO properties of a large sample of objects such as our subset of the eBOSS catalog would thus be prohibitively expensive through traditional means. LIM stacks, therefore, could potentially be extremely useful as a tool for studying large samples of high-redshift galaxies and quasars.

Conversely, in the limit that the eBOSS quasars do not significantly contribute to the integrated CO emission of their surroundings, the stacked flux measurement can be converted to a measurement of the cosmic molecular gas density. While we compare this value directly to other ρ_{H_2} measurements from various sources, its interpretation is somewhat more complicated, as the molecular gas density in the regions traced by the stack depends heavily on the bias of the cataloged quasars toward large-scale structure. Additionally, the stack probes only a single spatial scale, unlike the more conventional power spectrum based LIM measurements. These differences make stacking an excellent complement to other LIM analyses.

We propose, therefore, that stacking analyses with existing galaxy catalogs are a promising addition to the LIM analysis toolbox, especially when using LIM to approach as complex and multitracer a problem as galaxy evolution. To take full advantage of their potential benefits, stacking analyses should be performed on catalogs using as many different galaxy tracers as possible to probe this phase space more fully. We aim to investigate stacking as a galaxy analysis tool more fully in

future works, including by stacking on the extensive HETDEX catalog of Ly α emitters (Gebhardt et al. 2021).

Acknowledgments

This material is based upon work supported by the National Science Foundation under grant Nos. 1518282, 1517108, 1517598, 1517288, 1910999, and 2206834; by the Keck Institute for Space Studies under “The First Billion Years: A Technical Development Program for Spectral Line Observations;” and by a seed grant from the Kavli Institute for Particle Astrophysics and Cosmology.

D.D. acknowledges support from NSF Award 2206834. K. C. acknowledges support from NSF Awards 1518282, 1910999, and 2206834. P.C.B. is supported by the James Arthur Postdoctoral Fellowship. D.C. is supported by a CITA/Dunlap Institute postdoctoral fellowship. The Dunlap Institute is funded through an endowment established by the David Dunlap family and the University of Toronto. H.I., H.K.E., and I.W. acknowledge support from the Research Council of Norway through grant 251328. J.G. acknowledges support from the Keck Institute for Space Science, NSF AST-1517108, and University of Miami and thanks Hugh Medrano for assistance with cryostat design. J.K. is supported by a Robert A. Millikan Fellowship from Caltech. H.P.’s research is supported by the Swiss National Science Foundation via Ambizione grant PZ00P2_179934.

Parts of the work were carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

This research made use of the SDSS-IV eBOSS survey. 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.sdss4.org.

SDSS-IV is managed by the Astrophysical Research Consortium for the Participating Institutions of the SDSS Collaboration including the Brazilian Participation Group, the Carnegie Institution for Science, Carnegie Mellon University, Center for Astrophysics—Harvard & Smithsonian, the Chilean Participation Group, the French Participation Group, Instituto de Astrofísica de Canarias, The Johns Hopkins University, Kavli Institute for the Physics and Mathematics of the Universe (IPMU)/University of Tokyo, the Korean Participation Group, Lawrence Berkeley National Laboratory, Leibniz Institut für Astrophysik Potsdam (AIP), Max-Planck-Institut für Astronomie (MPIA Heidelberg), Max-Planck-Institut für Astrophysik (MPA Garching), Max-Planck-Institut für Extraterrestrische Physik (MPE), National Astronomical Observatories of China, New Mexico State University, New York University, University of Notre Dame, Observatório Nacional/MCTI, The Ohio State University, Pennsylvania State University, Shanghai Astronomical Observatory, United Kingdom Participation Group, Universidad Nacional Autónoma de México, University of Arizona, University of Colorado Boulder, University of Oxford, University of Portsmouth, University of Utah, University of Virginia, University of Washington, University of Wisconsin, Vanderbilt University, and Yale University.

This research made use of NASA’s Astrophysics Data System Bibliographic Services. For the purpose of open access,

the authors have applied a Creative Commons Attribution (CC BY) licence to any Author Accepted Manuscript version arising from this submission.

Finally, we would like to thank the anonymous referee whose comments and suggestions improved this manuscript.

Software: Astropy (Astropy Collaboration et al. 2013, 2018), Matplotlib (Hunter 2007).

Appendix A

Signal Attenuation due to Finite Spectral Aperture Size

As discussed in Section 4.2.3, quasars have inherently large velocity offsets between different redshift tracers due to inflowing and outflowing of different phases of the gas in the galaxy. In the stack, this manifests as an offset of the centroid of the stacked emission in the spectral axis and a broadening of this emission. This will serve to reduce the S/N, as emission is spread across spectral channels, and also attenuate the stack luminosity, as these velocity offsets can be wide enough to move signal out of the seven-channel aperture used to calculate the stack luminosity. This effect is compounded by other sources of spectral broadening in the stack, such as the inherent FWHM of the CO line. We defer the full details of the effects of spectral broadening on a stacking analysis to a future work, but we here explore the specific case of eBOSS quasars.

We use the previous studies of molecular gas in individual eBOSS quasars shown in Figure 5 to determine (1) an average CO FWHM for these objects of 319 km s^{-1} , (2) an average offset between the molecular redshift and the eBOSS redshift of $(z_{\text{CO}} - z_{\text{opt}})/(1 + z_{\text{opt}}) = 0.00397$, and (3) a scatter in that offset of 0.00408. We note that these correction factors are illustrative values, limited by the small number of eBOSS

quasars detected in CO. We would require many more CO measurements of individual quasars to determine a more confident calibration, and these values are not available at the moment. Additionally, we note that the assembled literature measurements use a variety of tracers that may trace slightly different phases of molecular gas, and thus this scatter may be enhanced compared to a single-tracer study such as COMAP. We deal with the average offset by simply shifting the stack centroid in the spectral axis by this amount. This remains a nondetection (as discussed in Section 4.2.3); the resulting stack is shown in Figures 10 and 11.

To address the scatter in the velocity offset between tracers (functionally, this introduces a redshift uncertainty in the stack), we use the constant-luminosity simulations discussed in Section 4.2 as a framework. We artificially set the line width of each CO emitter to be 319 km s^{-1} , as measured in Section 4.2.3. We test the attenuation from various values of scatter in $(z_{\text{CO}} - z_{\text{opt}})/(1 + z_{\text{opt}})$ by offsetting the spectral centroid of each halo individually by a randomly selected amount pulled from a normal distribution with standard deviation set to the $(z_{\text{CO}} - z_{\text{opt}})/(1 + z_{\text{opt}})$ scatter. This is done before stacking. We then stack the simulated data cube as normal. We perform 10 different iterations of this analysis for each value of $(z_{\text{CO}} - z_{\text{opt}})/(1 + z_{\text{opt}})$ we test, using the same underlying dark-matter halo distribution as a baseline and varying only the halo velocities. The resulting spectra for the scatter value we measure from individual eBOSS quasars are shown in Figure 12. In Figure 13, we show how this scatter attenuates signal compared to an unbroadened stack. Using the values determined from individual eBOSS quasars, we find that




Figure 10. The 3D stack performed on the eBOSS catalog using the $^{12}\text{CO}(1-0)$ emission wavelength offset by the frequency offset found in individually detected quasars (see Section 4.2.3). See Figure 6 for panel descriptions.




Figure 11. The bootstrapped uncertainty in L'_{CO} for the offset $^{12}\text{CO}(1-0)$ stack (see Figure 7).




Figure 12. Spectra of simulated stacks, offset to account for the velocity difference between eBOSS-determined redshifts and CO-determined redshifts ($(z_{\text{CO}} - z_{\text{opt}})/(1 + z_{\text{opt}}) = 0.00408$). Ten different iterations are shown. The fluctuations in the fringes of the spectrum are due to neighboring objects falling into some cutouts.




Figure 13. The percentage of the input luminosity returned by a constant-luminosity simulated stack as a function of an induced scatter between z_{opt} and z_{CO} in the individual halos making up the stack. The average scatter calculated from the independently detected quasars presented in Section 4.2.3, taken to be the average of our eBOSS quasar stack, is indicated with the vertical black dashed line. We repeat the test ten times—the average values are shown as large purple dots.

the returned signal is 63.1% of the actual value, corresponding to an attenuation factor of 1.58.

Appendix B Cosmic Volume Variation across the Stack

One constraint with our current stacking methodology is that the physical volume of each cubelet is redshift dependent—the COMAP data cubes stretch from redshift 2.4 to redshift 3.4, so the $6' \times 6' \times 218.75$ MHz aperture we sum to determine quantities such as L'_{CO} and ρ_{H_2} varies in size from $3.0 \text{ Mpc} \times 3.0 \text{ Mpc} \times 2519 \text{ km s}^{-1}$ at low redshifts to $2.7 \text{ Mpc} \times 2.7 \text{ Mpc} \times 1929 \text{ km s}^{-1}$ at high redshifts. We investigate here how this affects the stack. Ideally, this would be addressed by rebinning the data at the map level to create voxels that are constant physical sizes, and we plan to implement this in future work.

As in Appendix A, we investigate this effect using the constant-luminosity simulations from Section 4.2, again artificially imposing a line width of 319 km s^{-1} for each CO emitter. We generate three different 1000 halo simulations in which emitters are only located in a specific redshift range: a “low- z ” simulation with halos $2.39 < z < 2.42$, a “mid- z ” simulation with halos $2.82 < z < 2.86$, and a “high- z ” simulation with halos $3.89 < z < 3.43$. Each of these redshift ranges corresponds to 12 spectral channels in a COMAP data cube, meaning that the density of sources is the same in each simulated cube. We create three different iterations of each cube at each redshift.

In the spatial axes, the size of the voxel is varying across the stack at a similar rate to the size of the main beam, which is also defined in on-sky coordinates (Lamb et al. 2022). At the COMAP spatial resolution, nearly any astrophysical object is a point source, so the beam defines the spatial distribution of the emission. The distribution of the emission between map voxels is therefore extremely similar across the entire redshift range—we are stacking together point sources with a similar measured distribution, even if their physical distributions vary.

In the spectral axis, any line broadening is from astrophysical sources and not instrumental sources, as our spectral aliasing is near zero. At low redshifts, this is 1.15 times the width of a map channel, whereas at high redshifts it is 0.9 times the width of a channel. An identical CO emitter will thus (unphysically) appear more strongly peaked at high redshifts than at low redshifts (Figure 14). However, the same amount of emission is still present (spread between more spectral channels) and our seven-channel aperture is more than enough to encompass this spread, so the returned L'_{CO} value remains the same.

While the line luminosity is unaffected by the change in cosmic volume across the stack, the resulting molecular gas density, ρ_{H_2} , does vary across the stack because the volume being averaged over varies across the stack. This is because, in our simulations, all of the input luminosity is coming from a single source but is being spread over variable volumes across the stack due to the main beam. Volumes vary by roughly 37% across the full range of the stack—the same percentage difference we observe in the ρ_{H_2} values output from these simulations.




Figure 14. A comparison of constant-luminosity, constant line width simulations at three different redshifts spanning the width of the COMAP data cubes. The spatial distribution of (stacked) emission at each redshift is shown in the top three panels, and the bottom panel shows the spectrum of the stacked cubelets. While the high-redshift stacks are more peaked than the low-redshift stacks, each stack returns the same CO line luminosity. The low-level fringes in the spectra are background emission from objects other than the one being stacked appearing in the spectrum—asymmetries appear due to the low-redshift and high-redshift cutoffs.

Appendix C Foreground Stacks

As discussed in Section 5.2, there are several other molecular spectral lines whose emission may fall into the COMAP frequency range. These spectral lines could potentially contaminate the expected CO(1–0) LIM signal. Conversely, these spectral lines are important tracers of galaxy properties in their own rights, and any signal would be interesting and informative. We present stacks on each of these foregrounds lines below.

C.1. HCN(1–0)

HCN(1–0) is a tracer of dense ($n(\text{H}_2) \geq 3 \times 10^4 \text{ cm}^{-3}$) molecular gas in galaxies, associated primarily with the star-forming cores of giant molecular clouds (Gao & Solomon 2004). HCN luminosity correlates strongly with CO luminosity, although in a given galaxy L'_{HCN} is roughly an order of magnitude less than L'_{CO} . HCN has a rest frequency of 88.63 GHz, and 595 eBOSS quasars with this rest frequency fall into the COMAP volume. The HCN stack and its corresponding bootstrapped uncertainty determination are shown in Figures 15 and 16.

C.2. CS(2–1)

CS is interesting as a dense gas tracer independent of HCN/CN[−] chemistry, although it is quite faint (Wang et al.

2011; Zhang et al. 2014, e.g.). With a rest frequency of 91.98 GHz, there are 479 eBOSS objects included in this stack. The CS stack and its corresponding bootstrapped uncertainty determination are shown in Figures 17 and 18.

C.3. ¹³CO(1–0)

Molecule ¹³CO, as a rarer and slightly heavier isotopologue of ¹²CO, has a similar rest frequency (110.20 GHz) and traces a similar phase of gas, although with a higher critical density. As an isotopologue, ¹³CO is of particular interest in determining the CO-H₂ conversion factor in galaxies (e.g., Israel 2020). A total of 311 eBOSS objects are included in this stack. The ¹³CO stack and its corresponding bootstrapped uncertainty determination are shown in Figures 19 and 20.

C.4. CN[−](1–0)

CN[−](1–0) traces a very similar phase of gas to HCN (i.e., dense molecular cores; Section C.1), but has a critical density that is lower by roughly a factor of 5, and likely additionally prefers gas layers affected by stellar UV radiation (Riechers et al. 2007). The CN[−] rest frequency of 113.49 GHz is very close to CO(1–0), and 269 eBOSS quasars fall into the COMAP volume at this frequency. The CN[−] stack and its corresponding bootstrapped uncertainty determination are shown in Figures 21 and 22.




Figure 15. The 3D stack performed on the eBOSS catalog using the HCN(1-0) emission wavelength. See Figure 6 for panel descriptions.




Figure 16. The bootstrapped uncertainty in L'_{CO} for the HCN(1-0) stack (see Figure 7).




Figure 17. The 3D stack performed on the eBOSS catalog using the CS(2–1) emission wavelength. See Figure 6 for panel descriptions.




Figure 18. The bootstrapped uncertainty in L'_{CO} for the CS(2–1) stack (see Figure 7).




Figure 19. The 3D stack performed on the eBOSS catalog using the $^{13}\text{CO}(1-0)$ emission wavelength. See Figure 6 for panel descriptions.




Figure 20. The bootstrapped uncertainty in L'_{CO} for the $^{13}\text{CO}(1-0)$ stack (see Figure 7).




Figure 21. The 3D stack performed on the eBOSS catalog using the $\text{CN}^-(1-0)$ emission wavelength. See Figure 6 for panel descriptions.




Figure 22. The bootstrapped uncertainty in L'_{CO} for the $\text{CN}^-(1-0)$ stack (see Figure 7).

ORCID iDs

Delaney A. Dunne <https://orcid.org/0000-0002-5223-8315>
 Kieran A. Cleary <https://orcid.org/0000-0002-8214-8265>
 Patrick C. Breyse <https://orcid.org/0000-0001-8382-5275>
 Dongwoo T. Chung <https://orcid.org/0000-0003-2618-6504>
 Håvard T. Ihle <https://orcid.org/0000-0003-3420-7766>
 J. Richard Bond <https://orcid.org/0000-0003-2358-9949>
 Hans Kristian Eriksen <https://orcid.org/0000-0003-2332-5281>
 Joshua Ott Gundersen <https://orcid.org/0000-0002-7524-4355>

Laura C. Keating <https://orcid.org/0000-0001-5211-1958>
 Junhan Kim <https://orcid.org/0000-0002-4274-9373>
 Norman Murray <https://orcid.org/0000-0002-8659-3729>
 Hamsa Padmanabhan <https://orcid.org/0000-0002-8800-5740>
 Liju Philip <https://orcid.org/0000-0001-7612-2379>
 Nils-Ole Stutzer <https://orcid.org/0000-0001-5301-1377>
 Doğa Tolgay <https://orcid.org/0000-0002-3155-946X>
 Ingunn Katherine Wehus <https://orcid.org/0000-0003-3821-7275>
 Andrew I. Harris <https://orcid.org/0000-0001-6159-9174>

James W. Lamb  <https://orcid.org/0000-0002-5959-1285>
 Anthony C. S. Readhead  <https://orcid.org/0000-0001-9152-961X>

References

- Abdurashidova, Z., Aguirre, J. E., Alexander, P., et al. 2022, *ApJ*, **925**, 221
 Ahumada, R., Prieto, C. A., Almeida, A., et al. 2020, *ApJS*, **249**, 3
 Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, *AJ*, **156**, 123
 Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, *A&A*, **558**, A33
 Banerji, M., Carilli, C. L., Jones, G., et al. 2017, *MNRAS*, **465**, 4390
 Banerji, M., Jones, G. C., Wagg, J., et al. 2018, *MNRAS*, **479**, 1154
 Bautista, J. E., Busca, N. G., Guy, J., et al. 2017, *A&A*, **603**, A12
 Belitsky, V., Lapkin, I., Fredrixon, M., et al. 2018, *A&A*, **612**, A23
 Bischetti, M., Feruglio, C., Piconcelli, E., et al. 2021, *A&A*, **645**, A33
 Bolatto, A. D., Wolfire, M., & Leroy, A. K. 2013, *ARA&A*, **51**, 207
 Bolton, A. S., Schlegel, D. J., Aubourg, É., et al. 2012, *AJ*, **144**, 144
 Bothwell, M. S., Smail, I., Chapman, S. C., et al. 2013, *MNRAS*, **429**, 3047
 Breyse, P. C., & Alexandroff, R. M. 2019, *MNRAS*, **490**, 260
 Breyse, P. C., Kovetz, E. D., & Kamionkowski, M. 2015, *MNRAS*, **452**, 3408
 Brusa, M., Cresci, G., Daddi, E., et al. 2018, *A&A*, **612**, A29
 Carilli, C., & Walter, F. 2013, *ARA&A*, **51**, 105
 Cheng, Y.-T., Chang, T.-C., Bock, J., Bradford, C. M., & Cooray, A. 2016, *ApJ*, **832**, 165
 Chenu, J.-Y., Navarrini, A., Bortolotti, Y., et al. 2016, *ITTST*, **6**, 223
 Chung, D. T., Breyse, P. C., Cleary, K. A., et al. 2022, *ApJ*, **933**, 186
 Chung, D. T., Breyse, P. C., Ihle, H. T., et al. 2021, *ApJ*, **923**, 188
 Chung, D. T., Viero, M. P., Church, S. E., et al. 2019, *ApJ*, **872**, 186
 Cleary, K. A., Borowska, J., Breyse, P. C., et al. 2022, *ApJ*, **933**, 182
 Dawson, K. S., Kneib, J.-P., Percival, W. J., et al. 2016, *ApJ*, **151**, 44
 Dawson, K. S., Schlegel, D. J., Ahn, C. P., et al. 2013, *AJ*, **145**, 10
 Decarli, R., Arrigoni-Battaia, F., Hennawi, J. F., et al. 2021, *A&A*, **645**, L3
 Feruglio, C., Fiore, F., Carniani, S., et al. 2015, *A&A*, **583**, A99
 Foss, M. K., Ihle, H. T., Borowska, J., et al. 2022, *ApJ*, **933**, 184
 Gao, Y., & Solomon, P. M. 2004, *ApJ*, **606**, 271
 García-Burillo, S., Combes, F., Usero, A., et al. 2014, *A&A*, **567**, A125
 Gebhardt, K., Bender, R., Bower, G., et al. 2000, *ApJL*, **539**, L13
 Gebhardt, K., Mentuch Cooper, E., Ciardullo, R., et al. 2021, *ApJ*, **923**, 217
 González-López, J., Decarli, R., Pavesi, R., et al. 2019, *ApJ*, **882**, 139
 Gunn, J. E., Siegmund, W. A., Mannery, E. J., et al. 2006, *AJ*, **131**, 2332
 Herrera-Camus, R., Janssen, A., Sturm, E., et al. 2020, *A&A*, **635**, A47
 Hill, R., Chapman, S. C., Scott, D., et al. 2019, *MNRAS*, **485**, 753
 Hinshaw, G., Larson, D., Komatsu, E., et al. 2013, *ApJS*, **208**, 19
 Hopkins, P. F., Hernquist, L., Cox, T. J., & Kereš, D. 2008, *ApJS*, **175**, 356
 Hunter, J. D. 2007, *CSE*, **9**, 90
 Ihle, H. T., Borowska, J., Cleary, K. A., et al. 2022, *ApJ*, **933**, 185
 Israel, F. P. 2020, *A&A*, **635**, A131
 Jolly, J.-B., Knudsen, K., Laporte, N., et al. 2021, *A&A*, **652**, A128
 Keating, G. K., Marrone, D. P., Bower, G. C., et al. 2016, *ApJ*, **830**, 34
 Keating, G. K., Marrone, D. P., Bower, G. C., & Keenan, R. P. 2020, *ApJ*, **901**, 141
 Keenan, R. P., Keating, G. K., & Marrone, D. P. 2022, *ApJ*, **927**, 161
 Klitsch, A., Péroux, C., Zwaan, M. A., et al. 2019, *MNRAS*, **490**, 1220
 Kovetz, E., Breyse, P. C., Lidz, A., et al. 2019, *BAAS*, **51**, 101
 Lamb, J. W., Cleary, K. A., Woody, D. P., et al. 2022, *ApJ*, **933**, 183
 Lenkić, L., Bolatto, A. D., Förster Schreiber, N. M., et al. 2020, *AJ*, **159**, 190
 Lidz, A., & Taylor, J. 2016, *ApJ*, **825**, 143
 Lujan Niemeyer, M., Komatsu, E., Byrohl, C., et al. 2022, *ApJ*, **929**, 90
 Lyke, B. W., Higley, A. N., McLane, J. N., et al. 2020, *ApJS*, **250**, 8
 Muñoz-Elgueta, N., Arrigoni Battaia, F., Kauffmann, G., et al. 2022, *MNRAS*, **511**, 1462
 Orellana, G., Nagar, N. M., Isaak, K. G., et al. 2011, *A&A*, **531**, A128
 Paul, S., Santos, M. G., Chen, Z., & Wolz, L. 2023, arXiv:2301.11943
 Pullen, A. R., Chang, T.-C., Doré, O., & Lidz, A. 2013, *ApJ*, **768**, 15
 Pullen, A. R., & Hirata, C. M. 2013, *PASP*, **125**, 705
 Rennie, T. J., Harper, S. E., Dickinson, C., et al. 2022, *ApJ*, **933**, 187
 Riechers, D. A., Carilli, C. L., Maddalena, R. J., et al. 2011, *ApJL*, **739**, L32
 Riechers, D. A., Pavesi, R., Sharon, C. E., et al. 2019, *ApJ*, **872**, 7
 Riechers, D. A., Walter, F., Cox, P., et al. 2007, *ApJ*, **666**, 778
 Romano, M., Morselli, L., Cassata, P., et al. 2022, *A&A*, **660**, A14
 Sanders, D. B., Soifer, B. T., Elias, J. H., Neugebauer, G., & Matthews, K. 1988, *ApJL*, **328**, L35
 Silva, M. B., Baumschlager, B., Cleary, K. A., et al. 2021, arXiv:2111.05354
 Simpson, J. M., Smail, I., Swinbank, A. M., et al. 2012, *MNRAS*, **426**, 3201
 Sinigaglia, F., Elson, E., Rodighiero, G., & Vaccari, M. 2022, *MNRAS*, **514**, 4205
 Smee, S. A., Gunn, J. E., Uomoto, A., et al. 2013, *AJ*, **146**, 32
 Solomon, P. M., Downes, D., Radford, S. J. E., & Barrett, J. W. 1997, *ApJ*, **478**, 144
 Stanley, F., Jolly, J. B., König, S., & Knudsen, K. K. 2019, *A&A*, **631**, A78
 Stein, G., Alvarez, M. A., & Bond, J. R. 2019, *MNRAS*, **483**, 2236
 Timlin, J. D., Ross, N. P., Richards, G. T., et al. 2018, *ApJ*, **859**, 20
 Tytler, D., & Fan, X.-M. 1992, *ApJS*, **79**, 1
 Uzgil, B. D., Carilli, C., Lidz, A., et al. 2019, *ApJ*, **887**, 37
 Wang, J., Zhang, Z., & Shi, Y. 2011, *MNRAS*, **416**, L21
 White, M., Myers, A. D., Ross, N. P., et al. 2012, *MNRAS*, **424**, 933
 Zhang, Z.-Y., Gao, Y., Henkel, C., et al. 2014, *ApJL*, **784**, L31