



COMAP Early Science. IV. Power Spectrum Methodology and Results

Håvard T. Ihle¹, Jowita Borowska¹, Kieran A. Cleary², Hans Kristian Eriksen¹, Marie K. Foss¹, Stuart E. Harper³, Junhan Kim², Jonas G. S. Lunde¹, Liju Philip⁴, Maren Rasmussen¹, Nils-Ole Stutzer¹, Bade D. Uzgil^{5,6}, Duncan J. Watts¹, Ingunn Kathrine Wehus¹, J. Richard Bond^{7,8}, Patrick C. Breyse⁹, Morgan Catha¹⁰, Sarah E. Church¹¹, Dongwoo T. Chung⁷, Clive Dickinson³, Delaney A. Dunne², Todd Gaier⁴, Joshua Ott Gundersen¹², Andrew I. Harris¹³, Richard Hobbs¹⁰, James W. Lamb¹⁰, Charles R. Lawrence⁴, Norman Murray⁷, Anthony C. S. Readhead², Hamsa Padmanabhan¹⁴, Timothy J. Pearson², Thomas J. Rennie³, and David P. Woody¹⁰

(COMAP Collaboration)

¹ Institute of Theoretical Astrophysics, University of Oslo, P.O. Box 1029 Blindern, NO-0315 Oslo, Norway; h.t.ihle@astro.uio.no

² California Institute of Technology, Pasadena, CA 91125, USA

³ Jodrell Bank Centre for Astrophysics, Alan Turing Building, Department of Physics and Astronomy, School of Natural Sciences, The University of Manchester, Oxford Road, Manchester, M13 9PL, UK

⁴ Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA 91109, USA

⁵ California Institute of Technology, 1200 E. California Blvd., Pasadena, CA 91125, USA

⁶ National Radio Astronomy Observatory, Pete V. Domenici Array Science Center, P.O. Box 0, Socorro, NM 87801, USA

⁷ Canadian Institute for Theoretical Astrophysics, University of Toronto, 60 St. George Street, Toronto, ON M5S 3H8, Canada

⁸ Dunlap Institute for Astronomy and Astrophysics, University of Toronto, 50 St. George Street, Toronto, ON M5S 3H4, Canada

⁹ Center for Cosmology and Particle Physics, Department of Physics, New York University, 726 Broadway, New York, NY 10003, USA

¹⁰ Owens Valley Radio Observatory, California Institute of Technology, Big Pine, CA 93513, USA

¹¹ Kavli Institute for Particle Astrophysics and Cosmology & Physics Department, Stanford University, Stanford, CA 94305, USA

¹² Department of Physics, University of Miami, 1320 Campo Sano Avenue, Coral Gables, FL 33146, USA

¹³ Department of Astronomy, University of Maryland, College Park, MD 20742, USA

¹⁴ Département de Physique Théorique, Université de Genève, 24 Quai Ernest-Ansermet, CH-1211 Genève 4, Switzerland

Received 2021 November 19; revised 2022 February 25; accepted 2022 March 25; published 2022 July 13

Abstract

We present the power spectrum methodology used for the first-season COMAP analysis, and assess the quality of the current data set. The main results are derived through the Feed–Feed Pseudo-Cross-Spectrum (FPXS) method, which is a robust estimator with respect to both noise modeling errors and experimental systematics. We use effective transfer functions to take into account the effects of instrumental beam smoothing and various filter operations applied during the low-level data processing. The power spectra estimated in this way have allowed us to identify a systematic error associated with one of our two scanning strategies, believed to be due to residual ground or atmospheric contamination. We omit these data from our analysis and no longer use this scanning technique for observations. We present the power spectra from our first season of observing, and demonstrate that the uncertainties are integrating as expected for uncorrelated noise, with any residual systematics suppressed to a level below the noise. Using the FPXS method, and combining data on scales $k = 0.051\text{--}0.62 \text{ Mpc}^{-1}$, we estimate $P_{\text{CO}}(k) = -2.7 \pm 1.7 \times 10^4 \mu\text{K}^2 \text{ Mpc}^3$, the first direct 3D constraint on the clustering component of the CO(1–0) power spectrum in the literature.

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

1. Introduction

Intensity mapping aims to map out large 3D volumes using bright emission lines as tracers of large-scale matter distribution (Madau et al. 1997; Battye et al. 2004; Peterson et al. 2006; Loeb & Wyithe 2008; Kovetz et al. 2017, 2019). One promising set of lines comprises the rotational transitions of the carbon monoxide (CO) molecule. CO traces cold molecular gas, and is closely linked to star formation (Carilli & Walter 2013).

The CO Mapping Array Project (COMAP) is an intensity mapping experiment targeting CO. This paper, one of a set associated with the first-season COMAP analysis, presents the methodology used to constrain the CO power spectrum with

COMAP data. An overview of the COMAP experiment is presented by Cleary et al. (2022), while the COMAP instrument is described by Lamb et al. (2022).

The low-level COMAP data-processing pipeline is summarized by Foss et al. (2022). This pipeline converts raw uncalibrated observations into 3D maps, using redshifted CO line emission from distant galaxies as a tracer of the cosmic density field. Since the first-season COMAP instrument observes at frequencies between 26 and 34 GHz, and the rotational CO(1–0) transition has a rest frequency of 115 GHz, the current measurements trace galaxy formation at redshifts between $z = 2.4$ and 3.4, during “the epoch of galaxy assembly.” The current limits, forecasts, and modeling at these redshifts is discussed in Chung et al. (2022), while a future phase of COMAP, targeting “the epoch of reionization,” is discussed in Breyse et al. (2022). The use of this instrument for a galactic survey is presented in Rennie et al. (2022).

One common and powerful quantity used to characterize the statistical properties of such 3D cosmic maps is the so-called

power spectrum (or the two-point correlation function), which measures the strength of fluctuations as a function of physical distance (e.g., Lidz et al. 2011; Pullen et al. 2013; Li et al. 2016; Bernal et al. 2019; Chung 2019; Ihle et al. 2019; Uzgil et al. 2019; Gong et al. 2020; Keenan et al. 2022; Yang et al. 2021; Moradinezhad Dizgah et al. 2022). For an isotropic and Gaussian random field, this function quantifies all statistically relevant information in the original data set, but with a far smaller number of data points, and it thus represents a dramatic compression of the full data set. For non-Gaussian fields, additional information can be extracted by the use of other statistics (Breysse et al. 2017, 2019; Ihle et al. 2019; Sato-Polito & Bernal 2022). Even for non-Gaussian fields, however, such as the galactic density field, the power spectrum encapsulates a large fraction of the important information, and it is therefore an efficient tool even for such fields.

However, while compressing hundreds of terabytes of raw data into a handful of power spectrum coefficients certainly makes the interpretation of the data easier in terms of theoretical comparisons, it also makes the final estimates highly sensitive to small systematic effects and instrumental noise. To guide our intuition, we note that current theories predict an intrinsic CO standard deviation per resolution element of no more than a few microkelvin (Breysse et al. 2014; Li et al. 2016; Chung et al. 2022), which is to be compared with a typical system temperature of 44 K for the COMAP instrument; or atmospheric fluctuations of a few kelvin; or sidelobe contributions of a few millikelvin. All such effects must therefore be suppressed by many orders of magnitude in order to establish robust astrophysical constraints.

As described by Lamb et al. (2022) and Foss et al. (2022), the COMAP focal plane consists of 19 different feed horns, arranged in a hexagonal pattern, with about 12' sky separation between the closest feeds. The signal entering each feed horn is sent through its own signal chain, with its own amplifiers and digital high-resolution spectrometers. Each such signal chain is typically referred to as a “feed.” As such, the instrument has many unique features that makes it suited to this process. A few important examples include the highly efficient spectroscopic rejection of common-mode signals, several semi-independent feeds, a configurable scanning strategy, and frequent usage of hardware calibrators. Still, the rejection of systematic errors at the microkelvin level is highly challenging, and the current paper describes several algorithmic methods that can be applied to improve the robustness of the final results.

The rest of the paper is organized as follows. In Section 2, we review various aspects of power spectrum methods and present our adopted COMAP power spectrum estimator, the Feed–Feed Pseudo-Cross-Spectrum (FPXS). Power spectra estimated using data from COMAP’s first observing season are presented in Section 3, and we conclude in Section 4.

2. Methods

We begin our discussion with an overview of the fundamental algorithms used for COMAP power spectrum estimation. For other recent examples of the use of power spectrum analysis on intensity mapping data, see, e.g., Mertens et al. (2020), Keating et al. (2020), and Keenan et al. (2022).

2.1. Autospectrum Analysis

Let m_{ijk} denote a 3D map, and let us call each resolution element in this map a voxel. i , j , and k are then the voxel indices. We define the power spectrum $P(\mathbf{k})$ of this map to be the variance of its Fourier components, $f_{\mathbf{k}}$:

$$P(\mathbf{k}) = \frac{V_{\text{vox}}}{N_{\text{vox}}} \langle |f_{\mathbf{k}}|^2 \rangle, \quad (1)$$

where \mathbf{k} is the wavevector of a given Fourier mode, V_{vox} is the volume of each voxel, and N_{vox} is the total number of voxels.

If we assume that the map is statistically isotropic, then the power spectrum will only be a function of the magnitude of the wavevector, $P(\mathbf{k}) = P(k)$. In observational cosmology, we often want to distinguish the angular directions (denoted by the x and y coordinates) from the line-of-sight (LOS) direction (denoted by the z coordinate). This is because the map typically has different properties in different directions; for example, due to instrumental beam effects or redshift space distortions (Hamilton 1998; Chung 2019). It is therefore often useful to define the power spectrum in terms of parallel (LOS) modes, $k_{\parallel} \equiv |k_z|$, and perpendicular (angular) modes, $k_{\perp} \equiv \sqrt{k_x^2 + k_y^2}$. We can estimate the power spectrum in a given set of \mathbf{k} -bins, $\{\mathbf{k}_i\}$, from a given map as

$$P(\mathbf{k}_i) \approx \frac{V_{\text{vox}}}{N_{\text{vox}} N_{\text{modes}}} \sum_{j=1}^{N_{\text{modes}}} |f_{\mathbf{k}_j}|^2 \equiv P_{\mathbf{k}_i}, \quad (2)$$

where $P_{\mathbf{k}_i}$ is the estimated power spectrum in bin number i and $N_{\text{modes},i}$ is the number of Fourier components with wavenumber $\mathbf{k}_j \approx \mathbf{k}_i$ (i.e., in the bin corresponding to wavenumber \mathbf{k}_i).

Assuming that foreground and systematic contributions have already been removed to negligible levels through preprocessing, the power spectrum of a cleaned line intensity map is typically modeled as a sum of a signal and noise component (assumed to be statistically independent):

$$P(\mathbf{k}) = P_{\text{signal}}(\mathbf{k}) + P_{\text{noise}}(\mathbf{k}). \quad (3)$$

If we are able to estimate the noise power spectrum through independent means—for example, by using a noise model or simulations—we can extract the signal power spectrum simply by subtracting the estimated noise,

$$P_{\text{signal}}(\mathbf{k}_i) \approx P_{\mathbf{k}_i} - P_{\text{noise}}^{\text{est}}(\mathbf{k}_i), \quad (4)$$

where $P_{\text{noise}}^{\text{est}}(\mathbf{k}_i)$ is the estimated noise power spectrum in bin number i .

If the map consists of uniformly distributed white noise, then the noise power spectrum is independent of \mathbf{k} and given by

$$P_{\text{noise}} = V_{\text{vox}} \sigma_T^2, \quad (5)$$

where σ_T is the white noise standard deviation in each voxel (in units of kelvin). In our case, this magnitude of the white noise level is determined by the radiometer equation,

$$\sigma_T = \frac{T_{\text{sys}}}{\sqrt{\delta_\nu \tau}}, \quad (6)$$

where T_{sys} is the system temperature of the detector; δ_ν is the frequency resolution of each voxel; and τ is the total time each pixel is observed.

In addition to this instrumental noise contribution, there is an intrinsic uncertainty when estimating the signal power

spectrum from a map, called sample variance, which arises from the limited number of Fourier modes in the map. Together, these contributions give us the uncertainty of the power spectrum:

$$\begin{aligned} \sigma_P &\equiv \sqrt{\langle (P_{k_i} - P(k_i))^2 \rangle} \\ &\approx \underbrace{\frac{P_{\text{noise}}(k_i)}{\sqrt{N_{\text{modes}}}}}_{\text{Thermal noise}} + \underbrace{\frac{P_{\text{signal}}(k_i)}{\sqrt{N_{\text{modes}}}}}_{\text{Sample variance}}, \end{aligned} \quad (7)$$

where N_{modes} is the number of Fourier modes in bin number i , and the last approximation is exact when the Fourier modes are assumed to be independent Gaussian variables.

If the power spectrum is noise dominated, we can reduce this intrinsic uncertainty in two ways. First, we can observe the same area of sky for a longer time, thus decreasing the noise power spectrum contribution to the uncertainty. Alternatively, we can cover a larger sky area, and thus increase the number of measured Fourier modes. As long as we are noise dominated, a simple analysis suggests that observing a small area for a long time is more efficient for making a first detection than spreading the observations over a larger area. In a realistic situation, however, there are several other factors that must be taken into account, including the choice of angular resolution and scanning strategy constraints, and these will typically limit how small a field it is possible to observe.

Another source of uncertainty in estimating the signal is the accuracy of the estimated noise power spectrum model. If this model is biased or uncertain, then the associated residuals will propagate directly into the estimate of $P_{\text{signal}}(k_i)$.

2.2. Pseudospectrum Analysis

As described above, there are several challenges with an autospectrum analysis, as will be discussed both in this and the following sections. First of all, if the noise in the map is not uniform, which it generally is not, the noise power spectrum will be dominated by the parts of the map with the highest noise levels. In order to address this, it is necessary to devise a method that puts more weight on the parts of the map with low noise, and less weight on the parts of the map with high noise.

The standard method of accounting for this is through inverse noise-variance weights. That is, we weigh the map, m , by the noise level map, σ_m (the map given by the expected standard deviation of the white noise in each voxel), before we compute the power spectrum,

$$\tilde{P}_{k_i} = \frac{V_{\text{vox}}}{N_{\text{vox}} N_{\text{modes}}} \sum_{j=1}^{N_{\text{modes}}} |\tilde{f}_{k_j}|^2, \quad (8)$$

where \tilde{P} denotes the pseudospectrum and \tilde{f} are the Fourier components of the noise-weighted map,

$$\tilde{m} \equiv w m \quad (9)$$

and

$$w \equiv \mathcal{N} \frac{1}{\sigma_m^2}. \quad (10)$$

\mathcal{N} is a single overall normalization constant (which we will get back to), and σ_m is, as usual, the noise level map.

On a general note, the term ‘‘pseudospectrum’’ typically refers to a power spectrum estimator that is computed from a

biased estimator of the true sky map, and as such is itself biased; see Hivon et al. (2002). This may be contrasted with more conventional power spectrum estimators that aim to estimate the power spectrum of the true sky signal. The statistical information content of the pseudospectrum and the unbiased power spectrum is identical, and the main difference between the two classes of estimators concerns their ease of interpretation; while the unbiased power spectrum may be directly compared with theoretical models and other literature results, the pseudospectrum is experiment dependent, and typically requires simulations for proper statistical interpretation.

In our setting, we use the pseudospectrum to take into account both masked voxels (by setting $\sigma_m \rightarrow \infty$ for voxels that are excluded from further analysis) and varying noise levels across the map. Both these operations lead to *mode mixing*, i.e., different-signal Fourier modes are mixed together, and the estimated signal pseudospectrum is therefore a distorted version of the true signal power spectrum. However, since we know exactly how the signal map has been distorted, we can, at least in principle, calculate the exact mode-mixing matrix that is needed to reconstruct the mode mixing and obtain an unbiased signal spectrum from the pseudospectrum (Hivon et al. 2002). How feasible this is for a specific case depends on the details of the map dimensions and computational resources. For more details on mode mixing, see Appendix D.

Although mode mixing does complicate the physical interpretation of the pseudospectrum, there are several ways of dealing with this without having to calculate and invert the full mode-mixing matrix. First of all, if the analysis involves comparisons with signal simulations, then one may simply apply the same weight matrix to each simulation, making the observed and simulated power spectra statistically compatible. Second, if the level of mode mixing is modest, then the pseudospectrum may be an adequate estimator for the signal power spectrum for a given application, especially on smaller scales. This typically holds particularly well for noise-dominated applications, for which a single power estimate covering a large range in k is desired; in that case, the mode mixing often has a minimal effect on the estimates, and the pseudospectrum is often a perfectly valid estimate in its own right. The accuracy of this approximation must be assessed for each use case.

In cases for which the pseudospectrum is intended to be used as a direct estimator, it is necessary to choose a value for the normalization factor \mathcal{N} in Equation (10). Establishing the *formally correct* value for this normalization is not entirely well defined, as you are essentially trying to approximate the effect of an entire matrix with a single number (see Appendix D for more details). However, we can make a simple and fairly reasonable choice as follows:

$$\mathcal{N} = \frac{1}{\sqrt{\langle \frac{1}{\sigma_m^4} \rangle}}, \quad (11)$$

where $\langle \rangle$ denotes the average over the whole map. To make the results easier to interpret, we therefore apply this normalization to all the results shown in this paper. For analyses that employ the full mode-mixing matrix, or in which the pseudospectrum is compared directly to simulations, this normalization is completely irrelevant.

To roughly estimate the expected level of mode mixing, we calculate the ratio of the pseudospectrum and the autospectrum




Figure 1. Ratio of the signal pseudospectrum to the signal autospectrum, based on 10 signal realizations.

for 10 signal realization maps. Figure 1 shows the mean and standard deviation of the mode mixing in each of the main power spectrum bins. Overall, we see that at the scales where we have most of our sensitivity, the effect of mode mixing is fairly modest, typically in the 5%–30% range. Thus, even at face value, the pseudospectrum does provide a reasonable order-of-magnitude estimate of the true power spectrum, even if it may not be appropriate for precision analysis. We also note that these results suggest that, if anything, an upper limit obtained by interpreting the pseudospectrum at face value will be a slightly weaker (i.e., more conservative) upper limit than we would get by accounting for the mode mixing.

We leave it for future work to estimate the mode-mixing matrix and undo the mode-mixing bias in the pseudospectra. For the rest of this paper, we will interpret the pseudospectra at face value.

2.3. Cross-spectrum Analysis

A general challenge when using either the auto- or pseudospectrum is that highly accurate estimates of the noise contribution are required to estimate the signal power spectrum. In many cases, this can be very challenging, and any systematic error will directly bias the final signal estimate.

One way to avoid this complication is to use the so-called cross-spectrum, $C(\mathbf{k})$. While the power spectrum quantifies the variance of the Fourier components of a single map, the cross-spectrum quantifies the covariance between the Fourier modes of two different maps:

$$C(\mathbf{k}) = \frac{V_{\text{vox}}}{N_{\text{vox}}} \langle \text{Re}\{f_{1\mathbf{k}}^* f_{2\mathbf{k}}\} \rangle$$

$$\approx \frac{V_{\text{vox}}}{N_{\text{vox}} N_{\text{modes}}} \sum_{j=1}^{N_{\text{modes}}} \text{Re}\{f_{1\mathbf{k}_j}^* f_{2\mathbf{k}_j}\} \equiv C_{\mathbf{k}_j}. \quad (12)$$

Here, $\text{Re}\{\}$ denotes the real part of a complex number, and f_1 and f_2 are the Fourier components of two maps, m_1 and m_2 .

Clearly, if m_1 and m_2 are identical, then the cross-spectrum is equivalent to the autospectrum. The advantage of the cross-spectrum, however, is that if the maps m_1 and m_2 are made from different data, the noise contributions are independent, and they do not contribute to the mean of the cross-spectrum, only to its variance. Therefore, it is not necessary to estimate

and subtract the noise power spectrum to obtain an unbiased signal estimate, but rather

$$\langle C_{\mathbf{k}_i} \rangle = P_{\text{signal}}(\mathbf{k}_i). \quad (13)$$

Of course, a proper noise estimate is still necessary for uncertainty estimation, but the requirements on this are typically far less stringent than for the estimator mean.

Although the cross-spectrum significantly reduces the precision needed when estimating the noise power spectrum, we do pay a price in the form of somewhat lower intrinsic sensitivity. For instance, when splitting the data into two independent parts, and cross-correlating these, we do lose a factor of at least $\sqrt{2}$ from the fact that we do not exploit the autocorrelations within each data set separately. Fortunately, this problem can be remedied by splitting the data into more independent maps, and averaging the cross-spectra of all possible combinations. A lower limit on the cross-spectrum sensitivity is given by

$$\sigma_C^{N_{\text{split}}} \geq \sqrt{\frac{1}{1 - 1/N_{\text{split}}}} \sigma_P, \quad (14)$$

where N_{split} is the number of different map splits and σ_P is the optimal sensitivity of the autospectrum derived from the full data set.

The cross-spectrum has some other very important advantages with respect to the autospectrum, as well. As discussed in the introduction, one of the major challenges for an experiment like COMAP, in which we have to integrate down the noise by several orders of magnitude in order to measure a small signal, are systematic errors. However, since the cross-spectrum may only be biased by structures common to the two maps, one can try to ensure that any known systematic effects contribute to the two maps independently. In that case, the systematic effects will not bias the signal estimate. Combining this insight with splitting the data into multiple parts allows us to design a power spectrum statistic that is far more robust to systematics than the autospectrum.

We define a pseudo-cross-spectrum in an analogous manner to the pseudo-auto-spectrum. The only subtlety is that we make sure to apply the same weight map, w , for both maps. Explicitly, we adopt the following weight map,

$$w_{1,2} \propto \frac{1}{\sigma_{m_1} \sigma_{m_2}}, \quad (15)$$

for both m_1 and m_2 when calculating the pseudo-cross-spectrum, $\tilde{C}_{\mathbf{k}_i}$.

2.4. The Feed–Feed Pseudo-Cross-Spectrum Method

The idea of the FPXS method is to combine all the insights from the preceding sections to construct a single statistic for the CO signal, which has a high intrinsic sensitivity, uses proper noise weighting, and is robust against instrumental and other systematic errors. In that respect, we first note that the COMAP focal plane consists of 19 feeds, each with its own amplifiers and detectors. Furthermore, many systematic errors are particular to each feed, due to different passbands, amplifiers, cables, beams, etc. We may therefore split the data according to the feeds (i.e., make one map per feed), and then compute the cross-spectra of all the different feed combinations, without ever correlating two maps from the same feed.

Second, we also note that one of the most troublesome systematic errors for COMAP is ground pickup. This is mainly because the ground contamination correlates with the pointing, and it therefore does not average down in the same way as any systematic error that is random in the time domain (and hence independent in different observations). We can make ourselves as robust as possible against any residual ground signal in our map by also splitting the data by the elevation of the observations, so that we never take the cross-spectrum of two different data sets taken at the same elevation.¹⁵

With these considerations in mind, we define the following procedure for calculating the FPXS:

1. We split the data into disjoint sets, sorted according to elevation. For simplicity, we assume for now that we split the data into two sets, A and B , where A contains all the observations taken at elevations below the median elevation, and B contains all the observations from the higher elevations. We can easily generalize this to a case where we split the data into more than two sets.
2. For each set, A and B , we generate maps for each of the 19 feeds. We denote the different maps according to data set and feed, such that A_{13} indicates the map that combines all data from data set A for feed number 13.
3. We then calculate the pseudo-cross-spectrum, $\tilde{C}_{k_i}^{ij}$, for all the different map combinations of A_i and B_j where $i \neq j$.
4. Next, we compute the average pseudo-cross-spectrum, \tilde{C}_{k_i} , by noise weighting all the different cross-spectra:

$$\tilde{C}_{k_i}^{\text{FPXS}} = \left(\sum_{i \neq j} \frac{1}{\sigma_{\tilde{C}_{k_i}^{ij}}^2} \right)^{-1} \sum_{i \neq j} \frac{\tilde{C}_{k_i}^{ij}}{\sigma_{\tilde{C}_{k_i}^{ij}}^2}. \quad (16)$$

Here, $\sigma_{\tilde{C}_{k_i}^{ij}}$ is the uncertainty (standard deviation) in k -bin number i of the pseudo-cross-spectrum of the maps A_i and B_j , and the sum is over all combinations of i and j except the cases where $i = j$. Under the naive assumption that all cross-spectra are independent, the uncertainty of the combined cross-spectrum is given by

$$\sigma_{\tilde{C}_{k_i}^{\text{FPXS}}} = \left(\sum_{i \neq j} \frac{1}{\sigma_{\tilde{C}_{k_i}^{ij}}^2} \right)^{-1/2}. \quad (17)$$

The data can of course be split in other ways, to make ourselves less susceptible to other systematic effects, but we have found that using the feeds and elevation splits yields good results for the current data set.

2.5. White Noise Simulations

Until now, we have not discussed how to estimate the noise power spectrum and the corresponding noise uncertainty of the power spectrum. In general, estimating the noise power spectrum precisely is very difficult, since one needs to take into account not only the intrinsic white noise level of the data, but also the effect of the different filtering procedures in the low-level data analysis, as well as any correlated noise contribution.

¹⁵ The ground contamination also depends on azimuth, but since most of the problematic ground contamination happens at the highest or lowest elevations, it is most natural to divide the data according to elevation.

Since we use a cross-spectrum method, however, the noise spectrum is only used to estimate the uncertainties of the power spectrum, not its mean level, and the requirements on the absolute noise spectrum are therefore somewhat relaxed. Explicitly, if we make an error of a few percent in our noise estimate, we will not bias the estimated signal spectrum, only misestimate the error bars by a few percent. While clearly not ideal, this is usually not critical, considering all the other simplifying assumptions introduced in the analysis. On the other hand, if we had adopted an autospectrum method, an error of a few percent on the noise power spectrum could easily have rendered our signal estimate unusable, even in the case of very high intrinsic sensitivity.

For this reason, we therefore adopt a simple approach to noise power spectrum estimation: we assume that the noise in the maps is uncorrelated white noise, and generate noise simulations, m_i , by drawing random samples in each voxel from a Gaussian distribution with zero mean and a standard deviation given by the value of the noise level map, σ_m . We then estimate error bars by generating a large number of noise simulation maps, calculating the power spectrum from each, and finally taking the standard deviation in each k -bin of interest. This gives us uncertainties on the noise contribution to each power spectrum bin, but neglects the intrinsic uncertainty in the signal power spectrum itself. However, as we are still completely noise dominated, this intrinsic uncertainty of the signal spectrum should be negligible.

2.6. Transfer Functions

Until now, we have assumed that the sky maps produced by the low-level analysis pipeline are unbiased. For multiple reasons, this is not the case. First of all, the instrument does not have infinite resolution, and the instrumental beam will therefore smooth out the signal on small angular scales. The same effect occurs due to the finite spectral resolution of the instrument in the frequency dimension. Second, the various filters and mapmaking procedures in the analysis pipeline generally remove some of the signal, mostly on larger angular and spectral scales. In the following, we take these effects into account through so-called transfer functions. These are functions in the k_{\parallel} - k_{\perp} plane that quantify the fraction of the signal power that is retained in each k -bin, and allow us to establish unbiased estimates of the power spectrum from biased sky maps.

In general, a transfer function, $T(\mathbf{k})$, is defined through the following relation:

$$\langle P_k \rangle = T(\mathbf{k})P_{\text{signal}}(\mathbf{k}) + P_{\text{noise}}(\mathbf{k}), \quad (18)$$

where P_k is the power spectrum calculated from the final map and $P_{\text{signal}}(\mathbf{k})$ is the actual physical signal power spectrum. We decompose the full transfer function into different parts, and derive each separately. We then multiply the transfer functions together to get the full transfer function.

In writing down Equation (18) with a transfer function, $T(\mathbf{k})$, that is not a function of the signal, we have implicitly assumed that the effects we are accounting for using the transfer function are linear, that they do not depend on the properties of the signal, and that they can therefore be estimated using any signal model. While this is a good approximation in many cases (e.g., the beam effect is purely linear and most of the low-level filters are linear, assuming the same scanning pattern), it is only an approximation. However, even any residual theoretical dependence

on the input signal will typically be small in practice, since the signal power spectrum in any reasonable model is very smooth and has no sharp features. The most important effects to get right when estimating the transfer function are the noise distribution and scanning pattern. This is even more important because we are working with pseudospectra, where the noise level will affect the weighting and the mode mixing. That is why, even though it is costly to produce simulated data, we use about 63 hr of simulated data (thus ensuring a realistic scanning pattern as well as noise distribution) when we estimate the pipeline transfer function. Since we are using the pseudospectrum anyway, and not accounting for the mode mixing, we are already accepting errors of the order of 10%, which puts less stringent constraints on the precision of the rest of our procedures.

2.6.1. Instrumental Beam Transfer Function

Due to the finite resolution of the instrument, we cannot measure the cosmological signal on the smallest angular scales. In order to take this effect into account, we introduce a beam transfer function. For now, we assume the beam to be both achromatic (i.e., constant in frequency) and azimuthally symmetric. We construct an azimuthally symmetric beam model by averaging the full 2D (azimuth-elevation) beam model (Lamb et al. 2022) and inserting an exponential cutoff at around $30'$.

The ambient load calibration discussed in Foss et al. (2022) measures all of the power entering the feed horns, including the power that comes from the ground and all of the sky above the horizon. However, any power on scales larger than the modes we are sensitive to is essentially lost. To get a proper, scale-dependent, calibration, our beam model is normalized using observations of Jupiter and TauA, which show that 72% of the power is in the central $6/4$ of the beam (Rennie et al. 2022).

In addition, by including the beam model out to about $30'$, we take into account the roughly 10% of extra power that is retained at larger angular scales. We could include the beam model further out, but we are already hitting diminishing returns, so not much more would be gained.

Our (unnormalized) beam model can be seen in the top panel of Figure 2. The corresponding beam transfer function is estimated using signal-only simulations. That is, we generate a large number of 3D signal realizations and convolve our azimuthally symmetric beam model with the angular dimensions of the map. We then calculate the power spectrum of each of the signal realization maps, with and without beam smoothing. The estimated transfer function is given as the ratio of the average of these:

$$T^{\text{beam}}(\mathbf{k}) \approx \frac{\langle P_{\mathbf{k}}^{\text{signal,beam}} \rangle}{\langle P_{\mathbf{k}}^{\text{signal}} \rangle}, \quad (19)$$

where $P_{\mathbf{k}}^{\text{signal,beam}}$ is the power spectrum calculated from a beam-smoothed signal realization map and $P_{\mathbf{k}}^{\text{signal}}$ is the power spectrum calculated from the non-smoothed one.

Figure 2 shows the beam transfer function derived using 100 signal simulations. We see that the beam smoothing suppresses the power on small angular scales, corresponding to the main beam FWHM of about $4.5'$. We also see that although we lose sensitivity from our main beam efficiency, we retain some of this power on larger scales by making use of a beam model up to around $30'$.

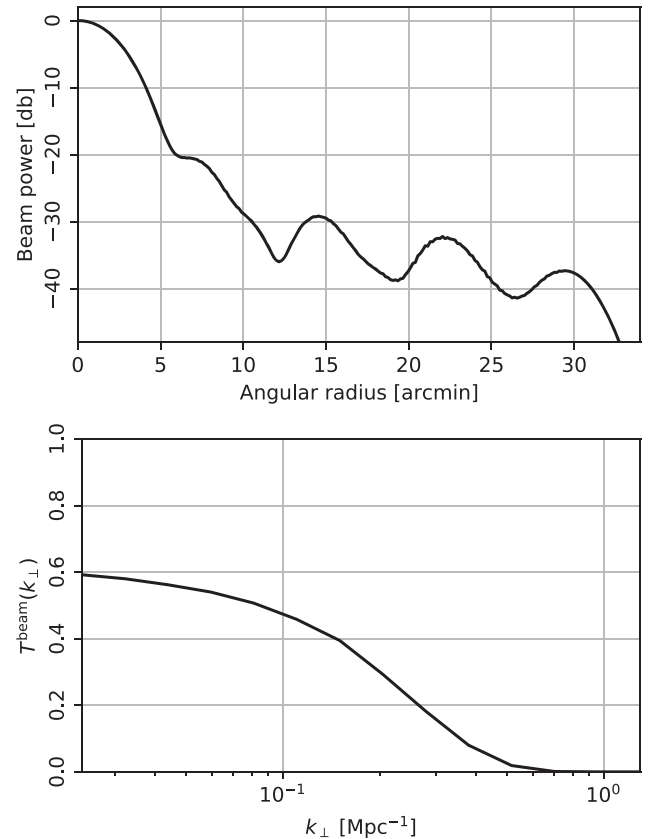


Figure 2. Top: radially symmetric instrumental beam model. Bottom: the resulting beam transfer function, after taking into account the main beam efficiency.

2.6.2. Frequency Resolution Transfer Function

Our current analysis, for simplicity and computational efficiency, uses fairly wide bins in frequency, of 31.25 MHz. This can be compared to intrinsic CO line widths of order 30 MHz (Chung et al. 2021), which will give the smallest scales present in the CO signal we are trying to observe. Once we are ready to claim a detection, we will increase the frequency resolution by at least a factor of 2 to get slightly more sensitivity to the small-scale CO signal, but for now this is not a high priority. We will, however, take into account the bias induced by the current bin size. Often such effects are taken into account by applying an analytic pixel window function, but this is not sufficient here, since the presence of structure on scales smaller than our bins means that some of this power can be aliased into our power spectrum. As the effect depends on the small-scale structure of the signal, there is no model-independent way to take this effect into account, and we will have to use simulations.

We estimate the frequency binning transfer function, $T^{\text{freq}}(\mathbf{k})$, by comparing power spectra of the simulated signal (using the default model in Chung et al. 2022) on a high-resolution frequency grid to power spectra of the simulated signal on our current frequency grid, both binned in our current $k_{\parallel} \times k_{\perp}$ bins. The transfer function derived using 50 such signal simulations is shown in Figure 3. We see a decrease in power toward smaller-LOS scales, but with an increase in the final bin, which we believe is the effect of aliasing the smaller-scale structure into this bin.




Figure 3. Frequency binning transfer function.

2.6.3. Pipeline Transfer Function

Each step of the analysis pipeline—including low-level filtering, calibration, and mapmaking—affects how much of the true sky signal is present in the final maps and power spectra. We estimate the transfer function of these operations by processing the sum of the raw data and a known signal-only time-ordered simulation through the analysis pipeline, following the exact same procedure as for the raw data alone. The pipeline transfer function may then be estimated as

$$T^{\text{pipeline}}(\mathbf{k}) \approx \left\langle \frac{P_{\mathbf{k}}^{\text{full}} - P_{\mathbf{k}}^{\text{noise}}}{P_{\mathbf{k}}^{\text{signal}}} \right\rangle, \quad (20)$$

where $P_{\mathbf{k}}^{\text{full}}$ is the power spectrum calculated from the maps derived from the raw data with added signal, $P_{\mathbf{k}}^{\text{noise}}$ is the power spectrum derived from the same data but without the added signal, and $P_{\mathbf{k}}^{\text{signal}}$ is the power spectrum derived from the raw signal simulation that was added to the raw data.

In Figure 4, the 2D binned pipeline transfer function for the Coudé Echelle Spectrometer (CES) data is shown. The transfer function peaks at intermediate k s, with efficiencies of ~ 0.8 – 0.85 around the peak region. We see that we lose the largest scales both in the angular and LOS directions. This is due to the various filters applied to the time-ordered data to remove correlated noise and systematics, in addition to the effects of the scanning strategy. For more details, see Foss et al. (2022).

2.6.4. Unbiased Signal Estimate

Figure 5 shows the full transfer function, combining all the effects discussed above. Correcting the FPXS with the above transfer function, we can establish an unbiased estimate of the




Figure 4. Pipeline transfer function for the cylindrically averaged power spectrum for constant-elevation scans. This transfer function is based on a single signal realization and roughly 3 hr of data.

signal pseudospectrum,

$$\tilde{P}_{\text{signal}}(\mathbf{k}) \approx \tilde{C}_{\mathbf{k}} \equiv \frac{\tilde{C}_{\mathbf{k}}^{\text{FPXS}}}{\tilde{T}_{\mathbf{k}}^{\text{full}}}, \quad (21)$$

where $\tilde{P}_{\text{signal}}(\mathbf{k})$ is the signal pseudospectrum and $\tilde{T}_{\mathbf{k}}^{\text{full}} = \tilde{T}_{\mathbf{k}}^{\text{beam}} \tilde{T}_{\mathbf{k}}^{\text{freq}} \tilde{T}_{\mathbf{k}}^{\text{pipeline}}$ is the full estimated transfer function for the pseudospectrum. The uncertainty of this signal estimate is given by

$$\sigma_{\tilde{P}_{\text{signal}}(\mathbf{k})} = \sigma_{\tilde{C}_{\mathbf{k}}} \equiv \frac{\sigma_{\tilde{C}_{\mathbf{k}}}^{\text{FPXS}}}{\tilde{T}_{\mathbf{k}}^{\text{full}}}. \quad (22)$$

2.6.5. Spherical Averaging

Due to the transfer function, different \mathbf{k} -modes corresponding to the same $k = |\mathbf{k}|$ bin have very different sensitivities. In order to get the best result, we need to take this into account when we calculate the spherically averaged power spectrum. As before, we use inverse noise-variance weighting to achieve this, giving us the following estimate for the unbiased spherically averaged pseudo-cross-spectrum:

$$\tilde{C}_{k_i} \equiv \frac{1}{\sum_{|\mathbf{k}| \in k_i} w_{\mathbf{k}} |\mathbf{k}| \in k_i} \sum_{|\mathbf{k}| \in k_i} w_{\mathbf{k}} \tilde{C}_{\mathbf{k}}, \quad (23)$$

where $w_{\mathbf{k}} \equiv 1/\sigma_{\tilde{C}_{\mathbf{k}}}^2$ and k_i denotes the i th k -bin.

For simplicity, we only calculate the spherical average of the cross-spectra that have already been cylindrically averaged and binned. This means that we use the bin centers of the $k_{\parallel} \times k_{\perp}$ bins to represent all the modes in the bin, which means that a few \mathbf{k} -modes get shifted back or forth by one bin in the spherically averaged cross-spectrum. Since there are no sharp




Figure 5. Full transfer function.

features in the signal power spectrum, this bias is modest, and not very important for the first-season analysis.

3. Power Spectrum Results

As described in Foss et al. (2022), after the COMAP time-ordered data have been filtered and calibrated, and the bad observations have been removed, the cleaned data set is compressed into a set of 3D maps. We make separate maps for the Lissajous scans and the CES scans, since these tend to have different systematics and statistical properties.

3.1. FPXS Results

We estimate separate cross-spectra for the Lissajous and CES data, for each of the three CO fields that we have observed (Foss et al. 2022). Since we found clear excess power in the Lissajous spectra, we do not include them in the main results, and we will here focus on the CES data. The power spectrum results for the Lissajous data are presented in Appendix A.

We split the data in two parts, according to the elevation of the observations, and use the FPXS method on these two sets of feed maps in order to minimize systematics. We also calculate a χ^2 statistic for each of the 16×15 different feed–feed cross-spectra,¹⁶ $\tilde{C}_{k_i}^{ij}$.

Based on these χ^2 statistics, denoted $\chi_{C(k)}^2$, we decided to reject all the spectra involving feed 8, since they showed very clear excesses in almost all spectra. This reduced the amount of data by 12.5% for all fields. We also saw clear structure in several of the spectra involving the low-elevation data from feeds 16 and 17 in the Field 1 results. This led us to remove all

¹⁶ As discussed in Foss et al. (2022), all data from feeds 4, 6, and 7 are rejected at an earlier stage of data selection. This leaves the data from 16 of the full 19 feeds.




Figure 6. Spherically averaged mean pseudo-cross-spectra for CES observations of Field 1 (blue), Field 2 (orange), and Field 3 (green). These spectra were generated from all the accepted data using the FPXS cross-spectrum statistic. In addition, the full transfer function has been applied, to debias the signal estimate. The data points from the different fields are offset slightly in k from their actual values to make them easier to distinguish.

spectra involving these feeds from the low-elevation data set for Field 1, thereby increasing the data loss to 24.2% for this field. In addition to the spectra that were removed by hand, we also rejected all spectra with more than a 5σ excess in $\chi_{C(k)}^2$ before we calculated the FPXS mean spectrum.

In the automatic 5σ cut, the fractions of the remaining spectra that were removed for the CES data were given by 1/182 for Field 1, 159/210 for Field 2, and 65/210 for Field 3. As discussed in Foss et al. (2022), the fact that such large fractions of data were removed at this stage (especially for Fields 2 and 3) suggests that large improvements in sensitivity can be achieved in the future, if we can identify the data affected by systematic errors at an earlier stage of the pipeline.

The resulting spherically averaged pseudo-cross-spectra are shown in Figure 6. We see that the results for the CES data appear to be largely flat, with fluctuations that are consistent with our white noise estimate. This demonstrates that we are in fact averaging down the noise, as expected for uncorrelated noise, and that the various potential systematic errors are suppressed to a level below the noise. At this point in the COMAP survey, this is a key outcome, given that our fiducial theoretical model predicts a signal on the order of $kP_{CO} \sim 10^3 \mu\text{K}^2 \text{Mpc}^2$ at our target redshift (Chung et al. 2022), well below the noise level shown here. This signal estimate is highly uncertain, however, and as discussed in Chung et al. (2022) these data already rule out some of the most optimistic models.

Combining these data points into a single measurement of the average CO power spectrum over the range $k = 0.051 - 0.62 \text{ Mpc}^{-1}$, we get

$$P_{CO}(k) = -2.7 \pm 1.7 \times 10^4 \mu\text{K}^2 \text{Mpc}^3. \quad (24)$$

This estimate is based on the pseudospectrum, and, as discussed in Section 2.2, it is a somewhat biased estimate of the signal, but should be a conservative estimate if used as an upper limit, as we do in Chung et al. (2022). This is the first direct 3D constraint on the clustering component of the CO(1–0) power spectrum in the literature.




Figure 7. Cylindrically averaged mean pseudo-cross-spectra for CES observations (top row). The second row shows the spectra divided by the corresponding white noise uncertainties.

Figure 7 shows the corresponding cylindrically averaged power spectrum in $k_{\parallel} \times k_{\perp}$ space. We see that the noise blows up very quickly as we move away from the region where the transfer function peaks. This illustrates the importance of taking the transfer function into account during the spherical averaging if you want maximum sensitivity in the 1D spectrum. In the region where we have appreciable sensitivity, the spectra look consistent with white noise, as they do for the 1D spectra. The bottom row shows the spectra divided by their corresponding white noise uncertainties, to better illustrate what happens at all scales.

3.2. Null Tests

Given that the current data appear to be largely consistent with white noise, the importance of null tests is less critical than it would be if a potential detection had been made. Still, null tests are a useful consistency check, and they may be useful in identifying and highlighting specific systematic errors, as well as potentially provide hints regarding the nature of the Lissajous excesses.

In order to get a sensitive set of null tests, we can calculate cross-spectra between maps of our different fields. In these null tests, any systematic related to standing waves or from residual large-scale ground contamination could still show up, while the signal should not contribute at all. Moreover, as long as we center the fields appropriately, they are roughly as sensitive as our original spectra, while other null tests are typically less sensitive. We therefore perform the same kind of FPXS power spectrum estimation as for the main results, but use the high-elevation data from one field and the low-elevation data from another field. In this way, we obtain two null tests per field pair, one where the first field uses the low-elevation data while the

second field uses the high-elevation data (denoted A), and another pair (denoted B) where the first field uses the high-elevation data and the other the low-elevation data. This gives a total of 6 null tests for each scanning method.

Figure 8 shows the results from these calculations, and we see that the null spectra are consistent with white noise expectations for all the CES data. For the Lissajous data, however, we see that most of the null tests show large excesses in power, consistent with our interpretation of the main Lissajous data containing systematics.

Table 1 shows χ^2 statistics calculated from each of the null tests in Figure 8, as well as the single-field results from Figures 6 and 10, by combining the eight data points of each spectrum. Specifically, we calculate the “probability to exceed,” which is defined as 1 minus the cumulative distribution function of the χ^2 distribution with the given number of degrees of freedom (here, 8). If the data are given by white noise, these statistics should be evenly distributed between 0% and 100%, while if we have excess power present, then the values should tend to be small. The results in the table support and quantify the statements that we made above, that the CES data looks consistent with white noise and the Lissajous data has clear power excesses present.

Although it is hard to interpret these values precisely, the fact that we see a clear excess in most of the Lissajous spectra made by combining maps from different fields suggests that whatever the systematic effect that gives rise to this excess, it needs to be common to all fields.

4. Conclusion

In this paper, we have introduced FPXS as a robust method for estimating the CO signal power spectrum from 3D intensity




Figure 8. FPXS spectra of maps from different fields. Here, A denotes cross-spectra of the low-elevation map from the first field and the high-elevation map from the second field, while B denotes the opposite combination. The data points from the different spectra are offset slightly in k from their actual values to make them easier to distinguish.

Table 1
 χ^2 Statistics from Science Results (Left) and Null Tests Using All the Different Field Combinations (Right)

Fields	χ^2 , Probability to Exceed									
	All	Field 1	Field 2	Field 3	1 \times 2, A	1 \times 2, B	1 \times 3, A	1 \times 3, B	2 \times 3, A	2 \times 3, B
CES	33%	17%	30%	52%	9%	73%	52%	69%	5%	28%
Lissajous	0.02%	0.1%	3%	72%	3%	3%	0.5%	58%	0.3%	0.3%

Note. These values were calculated by combining the data points shown in Figures 6, 8, and 10 into a single χ^2 value for each spectrum.

maps produced by the COMAP data analysis pipeline. We have discussed how to account for signal loss due to both filtering and beam smoothing, and we have estimated their magnitudes for the first-year COMAP observations with simulations. Computing the FPXS from the actual COMAP data, we find that the current data set is consistent with white noise for constant-elevation scan data, and the uncertainties average down with time, as expected for ideal data. Equivalently, these results suggest that all systematic errors are lower than the white noise level in our main sensitivity range.

In contrast, the FPXS results from the Lissajous scan data show clear signs of systematic errors. Further modeling and analysis work is required before these data can be used for astrophysical analysis.

Null tests largely seem to be consistent with our main results, with all the CES null tests being consistent with white noise, while most of the Lissajous null tests show clear excesses, supporting our assumption that the excesses seen in the main Lissajous spectra are the result of systematics.

Future analysis will involve explicit estimation of the mode-mixing matrix (see Appendix D), to undo the mode-mixing effect and present unbiased power spectrum estimates. We can also increase our sensitivity (by up to a factor of $\sqrt{2}$) by splitting the data into more than the current low- and high-elevation sets.

This material is based upon work supported by the National Science Foundation, under grant Nos. 1517108, 1517288, 1517598, 1518282, and 1910999, and by the Keck Institute for Space Studies, under “The First Billion Years: A Technical Development Program for Spectral Line Observations.” Parts of the work were carried out

at the Jet Propulsion Laboratory (JPL), California Institute of Technology, under a contract with the National Aeronautics and Space Administration, and funded through the internal Research and Technology Development program.

D.T.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. C.D. acknowledges support from an STFC Consolidated Grant (ST/P000649/1).

J.B., H.K.E., M.K.F., H.T.I., J.G.S.L., M.R., N.O.S., D.W., and I.K.W. acknowledge support from the Research Council of Norway through grants 251328 and 274990, and from the European Research Council (ERC) under the Horizon 2020 Research and Innovation Program (grant agreement No. 819478, COSMOGLOBE). J.G. acknowledges support from the University of Miami and is grateful to Hugh Medrano for assistance with cryostat design. S.H. acknowledges support from an STFC Consolidated Grant (ST/P000649/1). J.K. is supported by a Robert A. Millikan Fellowship from Caltech. At JPL, we are grateful to Mary Soria for assembly work on the amplifier modules and to Jose Velasco, Ezra Long, and Jim Bowen for the use of their amplifier test facilities. H.P. acknowledges support from the Swiss National Science Foundation through Ambizione Grant PZ00P2_179934. P.C.B. is supported by the James Arthur Postdoctoral Fellowship. We thank Isu Ravi for her contributions to the warm electronics and antenna-driven characterization. The scientific color maps `roma` and `tokyo` (Crameri 2021) are used in this study to prevent visual distortion of the data and exclusion of readers with color-vision deficiencies (Crameri et al. 2020).

Finally, we want to thank the anonymous referee, whose comments and suggestions have helped to improve and clarify this manuscript.

Software: Matplotlib (Hunter 2007); Astropy, a community-developed core Python package for astronomy (Astropy Collaboration et al. 2013).

Appendix A Lissajous Power Spectrum Results

When looking at the Lissajous data, we found clear excess power in the spectra. For this reason, we do not use any Lissajous data for our final science results, but include the Lissajous power spectra here for completeness. We use the exact same approach for the Lissajous data as we did for the CES data. We derive a separate pipeline transfer function for the Lissajous data (see Figure 9), since the properties of these scans are a bit different from the CES scans. As noted in Foss et al. (2022), we find that the transfer function for the Lissajous data preserves a bit more large-scale angular structure than the CES one, but it does not make much qualitative difference in terms of the complete transfer function. In the automatic 5σ cut based on the $\chi^2_{C(k)}$ statistics, the fractions of remaining spectra that were removed for the Lissajous data were given by 132/182 for Field 1, 92/210 for Field 2, and 109/210 for Field 3.

The Lissajous FPXS results for the spherically averaged power spectrum are shown in Figure 10. In contrast to the CES data, these data do not appear equally well behaved. Here, we see clear signs of excess power on large scales in both Field 1 and Field 2. These excesses suggest that large-scale systematic errors are still present for the Lissajous scans, and may be caused by either residual atmospheric variations or ground pickup from the far sidelobes, for instance.

These residuals are even more prominent when considering the 2D $k_{\parallel} \times k_{\perp}$ power spectrum, as shown in Figure 11. Here, we see some clear regions exhibiting systematic power excess. This is seen most clearly in the second row of the figure, which shows the power spectrum divided by the expected white noise fluctuations, thus corresponding to power measured in units of standard deviation. In particular, for Field 3, we see a bright region on the largest angular scales, and on scales between $k_{\parallel} \sim 0.03\text{--}0.1 \text{ Mpc}^{-1}$ in the frequency direction. We also see fairly bright regions at around $k_{\perp} = 0.2 \text{ Mpc}^{-1}$ and between $k_{\parallel} \sim 0.06\text{--}0.6 \text{ Mpc}^{-1}$ in the Field 1 data, which is right in the middle of our most sensitive region.




Figure 9. Pipeline transfer function for the cylindrically averaged power spectrum for the Lissajous scans. This transfer function is based on a single signal realization and roughly 3 hr of data.




Figure 10. Spherically averaged mean pseudo-cross-spectra for the Lissajous observations of Field 1 (blue), Field 2 (orange), and Field 3 (green). These spectra were generated from all the accepted data using the FPXS cross-spectrum statistic. In addition, the full transfer function has been applied, to debias the signal estimate. The data points from the different fields are offset slightly in k from their actual values to make them easier to distinguish.




Figure 11. Cylindrically averaged mean pseudo-cross-spectra for the Lissajous observations (top row). The second row shows the spectra divided by the corresponding white noise uncertainties.

Appendix B Fourier Conventions

In this Appendix, we present the conventions for the discrete Fourier transformations used in this paper. All the conventions are consistent with the default conventions in NumPy’s (Harris et al. 2020) FFT library. The forward transformation is given by

$$f_l = \sum_{m=0}^{n-1} x_m \exp\left(-2\pi i \frac{ml}{n}\right), \quad l = 0, \dots, n-1,$$

where x_m are the discrete values of the function in real space and f_l are the Fourier coefficients. The inverse transformation is then given by

$$x_m = \frac{1}{n} \sum_{l=0}^{n-1} f_l \exp\left(2\pi i \frac{ml}{n}\right).$$

We define the physical wavenumber as follows:

$$\begin{aligned} k &\equiv \frac{2\pi j}{\Delta x n}, \quad j \in \left\{-\frac{n}{2}, \dots, -1, 0, 1, \dots, \frac{n}{2}\right\} \\ &= 2\pi \cdot \text{np.fft.fftfreq}(n, \Delta x). \end{aligned}$$

Appendix C Definition of Cosmological Map Grid

Since Fourier transforms require a rectangular grid, we assume that the 3D temperature maps can be approximated by a rectangular grid in comoving cosmological parameters. We assume that all the voxels have the same shape and size as the middle voxel at redshift $z_{\text{mid}} \approx 2.9$.

The comoving length corresponding to an angular separation $\delta\theta$, for a given redshift z , is given by

$$\delta l_{\perp} = r(z)\delta\theta = \delta\theta \int_0^z \frac{cdz'}{H(z')}, \quad (\text{C1})$$

where $r(z)$ is the comoving distance traveled by light emitted from redshift z to us.

The comoving radial distance corresponding to a small redshift interval $\delta z = z_1 - z_2 = \nu_0/\nu_1^{\text{obs}} - \nu_0/\nu_2^{\text{obs}} \approx (1+z)^2 \delta\nu^{\text{obs}}/\nu_0$, where $z_1 > z_2$, is given by

$$\delta l_{\parallel} = \int_{z_2}^{z_1} \frac{cdz}{H(z)} \approx \frac{c\delta z}{H(z)} \approx \frac{c}{H(z)} \frac{(1+z)^2 \delta\nu^{\text{obs}}}{\nu_0}, \quad (\text{C2})$$

where $\nu_0 \approx 115.27$ is the emission frequency of the CO $1 \rightarrow 0$ line we are studying and $\delta\nu^{\text{obs}} = 31.25$ MHz is the resolution of our frequency bins.

Given a pixel width of $2'$, we then get the following voxel dimensions:

$$\delta l_{\perp} = 3.63 \text{ Mpc}, \quad (\text{C3})$$

$$\delta l_{\parallel} = 4.26 \text{ Mpc}. \quad (\text{C4})$$

Appendix D Mode Mixing and the Master Algorithm

In order to understand the mode-mixing effect, let us consider in more detail the Fourier transform of a weighted map:¹⁷

¹⁷ We work in 2D here to save some indices; the generalization to 3D is straightforward.

$$\tilde{f}_{k_1 k_2} = \sum_{m_1=0}^{n-1} \sum_{m_2=0}^{n-1} x_{m_1 m_2} W_{m_1 m_2} \times \exp\left(-2\pi i \frac{m_1 k_1 + m_2 k_2}{n}\right). \quad (\text{D1})$$

Here, $x_{m_1 m_2}$ is the map, $W_{m_1 m_2}$ is the weight map, and $\tilde{f}_{k_1 k_2}$ is the Fourier transform of the weighted map. We can insert the expression for the inverse Fourier transform of x and W ,

$$\begin{aligned} \tilde{f}_{k_1 k_2} &= \frac{1}{n^4} \sum_{k'_1=0}^{n-1} \sum_{k'_2=0}^{n-1} f_{k'_1 k'_2} \sum_{k''_1=0}^{n-1} \sum_{k''_2=0}^{n-1} f_{k''_1 k''_2}^W \\ &\times \sum_{m_1=0}^{n-1} \sum_{m_2=0}^{n-1} \exp\left(-2\pi i \frac{m_1(k'_1 + k''_1 - k_1) + m_2(k'_2 + k''_2 - k_2)}{n}\right), \end{aligned} \quad (\text{D2})$$

where f and f^W are the Fourier transforms of x and W , respectively. Working through the algebra, we get

$$\begin{aligned} \tilde{f}_{k_1 k_2} &= \frac{1}{n^2} \sum_{k'_1=0}^{n-1} \sum_{k'_2=0}^{n-1} f_{k'_1 k'_2} \\ &\sum_{k''_1=0}^{n-1} \sum_{k''_2=0}^{n-1} f_{k''_1 k''_2}^W \delta_{k'_1(k_1 - k'_1) \% n, k'_2(k_2 - k'_2) \% n} \\ \tilde{f}_{k_1 k_2} &= \frac{1}{n^2} \sum_{k'_1=0}^{n-1} \sum_{k'_2=0}^{n-1} f_{k'_1 k'_2} f_{(k_1 - k'_1) \% n, (k_2 - k'_2) \% n}^W \\ \tilde{f}_{k_1 k_2} &= \sum_{k'_1=0}^{n-1} \sum_{k'_2=0}^{n-1} f_{k'_1 k'_2} \underbrace{\frac{1}{n^2} f_{(k_1 - k'_1) \% n, (k_2 - k'_2) \% n}^W}_{\equiv K_{k_1, k_2, k'_1, k'_2}}, \end{aligned} \quad (\text{D3})$$

where $\%$ denotes the modulo operation and we have defined the mode-mixing amplitude $K_{k, k'}$.

Adopting vector notation, we may now write the pseudospectrum as follows:

$$\tilde{P}(\mathbf{k}) = \frac{V_{\text{vox}}}{N_{\text{vox}}} \langle \tilde{f}_{\mathbf{k}} \tilde{f}_{\mathbf{k}}^* \rangle \quad (\text{D4})$$

$$= \frac{V_{\text{vox}}}{N_{\text{vox}}} \frac{1}{n^{2D}} \sum_{\mathbf{k}'} \sum_{\mathbf{k}''} \langle f_{\mathbf{k}'} f_{\mathbf{k}''}^* \rangle K_{\mathbf{k}, \mathbf{k}'} K_{\mathbf{k}, \mathbf{k}''}^* \quad (\text{D5})$$

$$= \frac{1}{n^{2D}} \sum_{\mathbf{k}'} \sum_{\mathbf{k}''} P(\mathbf{k}') \delta_{\mathbf{k}', \mathbf{k}''} K_{\mathbf{k}, \mathbf{k}'} K_{\mathbf{k}, \mathbf{k}''}^* \quad (\text{D6})$$












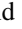

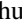
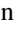
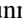
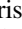
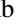



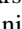
$$= \sum_{\mathbf{k}'} P(\mathbf{k}') \frac{1}{n^{2D}} \frac{|K_{\mathbf{k}, \mathbf{k}'}|^2}{M_{\mathbf{k}, \mathbf{k}'}} \quad (\text{D7})$$

where D is the number of dimensions of the map and we have defined the mode-mixing matrix $M_{\mathbf{k}, \mathbf{k}'}$. We see that the autospectrum and the pseudospectrum are related by a linear transformation, so all the information in one is also there in the other.

Within the cosmic microwave background field, accounting for mode mixing by explicitly calculating and inverting $M_{\mathbf{k}, \mathbf{k}'}$ is often referred to as the MASTER algorithm (Hivon et al. 2002; Leung et al. 2022). Doing this requires that we calculate the mode mixing between each Fourier mode and all the other

Fourier modes, so for 3D maps this scales poorly with the map dimensions. On the other hand, the algorithm parallelizes trivially, and the matrix must only be computed once for a given weight map, after which the same operation may be applied efficiently to any number of simulations. Whether this is feasible depends on the details of the individual use case. Some methods exist in the literature to approximate this procedure in a faster way; see, e.g., Louis et al. (2020).

ORCID iDs

Håvard T. Ihle  <https://orcid.org/0000-0003-3420-7766>
 Kieran A. Cleary  <https://orcid.org/0000-0002-8214-8265>
 Hans Kristian Eriksen  <https://orcid.org/0000-0003-2332-5281>
 Marie K. Foss  <https://orcid.org/0000-0001-8896-3159>
 Stuart E. Harper  <https://orcid.org/0000-0001-7911-5553>
 Junhan Kim  <https://orcid.org/0000-0002-4274-9373>
 Liju Philip  <https://orcid.org/0000-0001-7612-2379>
 Nils-Ole Stutzer  <https://orcid.org/0000-0001-5301-1377>
 Bade D. Uzgil  <https://orcid.org/0000-0001-8526-3464>
 Duncan J. Watts  <https://orcid.org/0000-0002-5437-6121>
 Ingunn Kathrine Wehus  <https://orcid.org/0000-0003-3821-7275>
 J. Richard Bond  <https://orcid.org/0000-0003-2358-9949>
 Patrick C. Breyse  <https://orcid.org/0000-0001-8382-5275>
 Dongwoo T. Chung  <https://orcid.org/0000-0003-2618-6504>
 Clive Dickinson  <https://orcid.org/0000-0002-0045-442X>
 Delaney A. Dunne  <https://orcid.org/0000-0002-5223-8315>
 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>
 Hamsa Padmanabhan  <https://orcid.org/0000-0002-8800-5740>
 Timothy J. Pearson  <https://orcid.org/0000-0001-5213-6231>
 Thomas J. Rennie  <https://orcid.org/0000-0002-1667-3897>

References

- Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, *A&A*, **558**, A33
- Battye, R. A., Davies, R. D., & Weller, J. 2004, *MNRAS*, **355**, 1339
- Bernal, J. L., Breyse, P. C., Gil-Marín, H., & Kovetz, E. D. 2019, *PhRvD*, **100**, 123522
- Breyse, P. C., Anderson, C. J., & Berger, P. 2019, *PhRvL*, **123**, 231105
- Breyse, P. C., Chung, D. T., Cleary, K. A., et al. 2022, *ApJ*, **933**, 188
- Breyse, P. C., Kovetz, E. D., Behroozi, P. S., Dai, L., & Kamionkowski, M. 2017, *MNRAS*, **467**, 2996
- Breyse, P. C., Kovetz, E. D., & Kamionkowski, M. 2014, *MNRAS*, **443**, 3506
- Carilli, C. L., & Walter, F. 2013, *ARA&A*, **51**, 105
- Chung, D. T. 2019, *ApJ*, **881**, 149
- 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
- Cleary, K. A., Borowska, J., Breyse, P. C., et al. 2022, *ApJ*, **933**, 182
- Cramer, F. 2021, Scientific colour maps, v7.0.0, Zenodo, doi:10.5281/zenodo.1243862
- Cramer, F., Shephard, G. E., & Heron, P. J. 2020, *NatCo*, **11**, 5444
- Foss, M. K., Ihle, H. T., Borowska, J., et al. 2022, *ApJ*, **933**, 184
- Gong, Y., Chen, X., & Cooray, A. 2020, *ApJ*, **894**, 152
- Hamilton, A. J. S. 1998, in *The Evolving Universe*, ed. D. Hamilton (Dordrecht: Kluwer), 185
- Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, *Natur*, **585**, 357
- Hivon, E., Górski, K. M., Netterfield, C. B., et al. 2002, *ApJ*, **567**, 2
- Hunter, J. D. 2007, *CSE*, **9**, 90
- Ihle, H. T., Chung, D., Stein, G., et al. 2019, *ApJ*, **871**, 75
- 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
- Kovetz, E., Breysse, P. C., Lidz, A., et al. 2019, *BAAS*, 51, 101
- Kovetz, E. D., Viero, M. P., Lidz, A., et al. 2017, [arXiv:1709.09066](#)
- Lamb, J. W., Cleary, K. A., Woody, D. P., et al. 2022, [ApJ](#), 933, 183
- Leung, J. S. Y., Hartley, J., Nagy, J. M., et al. 2022, [ApJ](#), 928, 109
- Li, T. Y., Wechsler, R. H., Devaraj, K., & Church, S. E. 2016, [ApJ](#), 817, 169
- Lidz, A., Furlanetto, S. R., Oh, S. P., et al. 2011, [ApJ](#), 741, 70
- Loeb, A., & Wyithe, J. S. B. 2008, [PhRvL](#), 100, 161301
- Louis, T., Naess, S., Garrido, X., & Challinor, A. 2020, [PhRvD](#), 102, 123538
- Madau, P., Meiksin, A., & Rees, M. J. 1997, [ApJ](#), 475, 429
- Mertens, F. G., Mevius, M., Koopmans, L. V. E., et al. 2020, [MNRAS](#), 493, 1662
- Moradinezhad Dizgah, A., Nikakhtar, F., Keating, G. K., & Castorina, E. 2022, [JCAP](#), 2022, 026
- Peterson, J. B., Bandura, K., & Pen, U. L. 2006, [arXiv:astro-ph/0606104](#)
- Pullen, A. R., Chang, T.-C., Doré, O., & Lidz, A. 2013, [ApJ](#), 768, 15
- Rennie, T. J., Harper, S. E., Dickinson, C., et al. 2022, [ApJ](#), 933, 187
- Sato-Polito, G., & Bernal, J. L. 2022, [arXiv:2202.02330](#)
- Uzgil, B. D., Carilli, C., Lidz, A., et al. 2019, [ApJ](#), 887, 37
- Yang, S., Somerville, R. S., Pullen, A. R., et al. 2021, [ApJ](#), 911, 132