The deconvolved distribution estimator: enhancing reionization-era CO line-intensity mapping analyses with a cross-correlation analogue for one-point statistics

Dongwoo T. Chung [1,2]★ Ishika Bangari,[3] Patrick C. Breysse [4] Håvard T. Ihle,[5] J. Richard Bond,[1]
Delaney A. Dunne,[6] Hamsa Padmanabhan [7] Liju Philip,[6] Thomas J. Rennie[8]
and Marco P. Viero [6] (COMAP Collaboration)

[1]*Canadian Institute for Theoretical Astrophysics, University of Toronto, 60 St. George Street, Toronto, ON M5S 3H8, Canada*
[2]*Dunlap Institute for Astronomy and Astrophysics, University of Toronto, 50 St. George Street, Toronto, ON M5S 3H4, Canada*
[3]*David A. Dunlap Department of Astronomy and Astrophysics, University of Toronto, 50 St. George Street, Toronto, ON M5S 3H4, Canada*
[4]*Center for Cosmology and Particle Physics, Department of Physics, New York University, 726 Broadway, New York, NY 10003, USA*
[5]*Institute of Theoretical Astrophysics, University of Oslo, PO Box 1029 Blindern, N-0315 Oslo, Norway*
[6]*California Institute of Technology, 1200 E. California Blvd., Pasadena, CA 91125, USA*
[7]*Département de Physique Théorique, Université de Genève, 24 Quai Ernest-Ansermet, CH-1211 Genève 4, Switzerland*
[8]*Jodrell Bank Centre for Astrophysics, Alan Turing Building, Department of Physics and Astronomy, The University of Manchester, Oxford Road, Manchester M13 9PL, UK*

ABSTRACT
We present the deconvolved distribution estimator (DDE), an extension of the voxel intensity distribution (VID), in the context of future observations proposed as part of the CO Mapping Array Project (COMAP). The DDE exploits the fact that the observed VID is a convolution of correlated signal intensity distributions and uncorrelated noise or interloper intensity distributions. By deconvolving the individual VID of two observables away from their joint VID in a Fourier-space operation, the DDE suppresses sensitivity to interloper emission while maintaining sensitivity to correlated components. The DDE thus improves upon the VID by reducing the relative influence of uncorrelated noise and interloper biases, which is useful in the context of COMAP observations that observe different rotational transitions of CO from the same comoving volume in different observing frequency bands. Fisher forecasts suggest that the theoretical sensitivity in the DDE allows significant improvements in constraining power compared to either the cross power spectrum or the individual VID data, and matches the constraining power of the combination of all other one- and two-point summary statistics. Future work should further investigate the covariance and model-dependent behaviour of this novel one-point cross-correlation statistic.

Key words: methods: statistical – diffuse radiation – large-scale structure of Universe.

1 INTRODUCTION

Line-intensity mapping (LIM; see reviews by Kovetz et al. 2017, 2019; Bernal & Kovetz 2022) is a nascent paradigm for surveying the cosmic web at early cosmic epochs. Instead of tracing the large-scale structure (LSS) of the Universe with individual resolved galaxies, LIM proposes to use unresolved emission in atomic or molecular spectral lines across large cosmological volumes. The resulting measurements of clustering of and shot noise in the cosmological line emission should allow statistical inferences of properties of the entire population of line emitters, such as the luminosity function at both faint and bright ends, the bias with which line emission traces the underlying LSS, and so on.

While LIM thus works around challenges that conventional surveys face in target resolution and selection bias, that does not render LIM an *easier* endeavour. Mitigation of systematics and

contaminants is key to being able to leverage the rich statistical promise of LIM. In this respect, a LIM measurement of one spectral line on its own gives at best an incomplete and tenuous picture of cosmic evolution. Cross-correlations involving LIM experiments and other tracers of LSS will strongly reject disjoint sources of error while enabling multitracer astrophysics on cosmological scales to completely probe the environmental and topological factors that drive events like cosmic reionization and galaxy assembly (see e.g. Sun et al. 2019). Indeed, some of the earliest halo models used to forecast prospects for higher-frequency LIM experiments intended cross-correlation of carbon monoxide (CO) and ionized carbon ([C II]) signals with 21 cm observations or other LSS tracers as a central science case (Gong et al. 2011; Lidz et al. 2011; Gong et al. 2012; Pullen et al. 2013). Naturally cross-correlations remain a key consideration in LIM survey forecasting, design, and analysis (see e.g.: Breysse & Alexandroff 2019; Keenan, Keating & Marrone 2022; Pullen et al. 2022).

The bulk of cross-correlation forecasts focus on two-point statistics, examining the possibility of detecting cross power spectra.

But two-point statistics only provide complete information about the correlation of fields if the fields are Gaussian, whereas LIM signals end up being significantly non-Gaussian at small scales due to the non-linear nature of the process of structure formation and the resulting spatial distribution of line emitters. On the other hand, one-point statistics like the voxel intensity distribution (VID) – i.e. the histogram of voxel intensities or temperatures across the LIM data cube – will be sensitive to small-scale, non-Gaussian information. As previous studies have shown, one-point statistics like the VID thus critically complement two-point statistics like power spectra by breaking degeneracies inherent to the latter and thus improving the constraining power of LIM data, as shown explicitly in previous studies (Breysse et al. 2017; Ihle et al. 2019; Breysse 2022).

An extension of the VID involves conditioning voxel intensities on external data, e.g. obtaining the VID only across voxels in the LIM data cube that overlap with a galaxy detection in an external data set, and comparing this to the VID obtained across voxels in the LIM data cube without an overlapping galaxy. Comparing different conditional VID data rejects biases disjoint between the data sets just as cross-correlation two-point statistics do, but still retaining non-Gaussian information. Breysse, Anderson & Berger (2019) confirmed this in an explicit simulation study of the conditional VID between 21 cm data and counts-in-cells in simulations (or more precisely, of the ratio of Fourier-transformed conditional VID data, representing an operation that deconvolves foreground biases uncorrelated with the conditioning galaxy distribution). The EXperiment for Cryogenic Large-Aperture Intensity Mapping (EXCLAIM) also recently presented conditional VID forecasts, with [C II] intensities conditioned on galaxy counts-in-cells (Pullen et al. 2022).

However, certain contexts require an extension of the VID that deals with the joint probability density function (PDF) of correlated voxel intensity components, rather than discrete conditional distributions. One such context is found in the CO Mapping Array Project (COMAP; Cleary et al. 2022), a dedicated single-dish LIM experiment that targets rotational transitions of CO. A key aim of COMAP in later phases is cross-correlation of $z \sim 7$ CO(2–1) emission observed in the Ka band (around 30 GHz) with $z \sim 7$ CO(1–0) emission observed in the Ku band (around 15 GHz). In such a study of molecular gas in the late Epoch of Reionization (EoR), we wish to leverage a robust joint-PDF extension of the VID capable of dealing with two continuous line-intensity fields, which would strongly complement the cross power spectrum detections forecast for future COMAP phases.

In this paper, we describe the deconvolved distribution estimator (DDE; Breysse, Chung & Ihle 2022b). Derived from the individual and joint PDF of correlated signal intensities, the DDE is designed to be robust to independent noise or other disjoint sources of bias in the individual VID or joint PDF, as its calculation deconvolves uncorrelated biasing distributions away from the joint PDF. We demonstrate the potential utility of the DDE in the context of the COMAP, and specifically the COMAP Expanded Reionization Array (COMAP-ERA) concept proposed by Breysse et al. (2022a). In doing so, we will answer the following questions:

(i) Is the DDE robust to contaminants like noise and interloper emission, as is the analytic expectation?
(ii) How much could the DDE fundamentally improve constraints on $z \sim 7$ CO emission in simulated COMAP-ERA observations?

We structure the paper as follows. In Section 2 we motivate and define the DDE in more mathematical detail. Section 3 then defines the COMAP-ERA simulations in which we propose to examine the

potential of the DDE. We consider the results of these simulations in Section 4 before concluding in Section 5.

Unless otherwise stated, we assume base-10 logarithms, and a ΛCDM cosmology with parameters $\Omega_m = 0.286$, $\Omega_\Lambda = 0.714$, $\Omega_b = 0.047$, $H_0 = 100h \, \mathrm{km \, s^{-1} \, Mpc^{-1}}$ with $h = 0.7$, $\sigma_8 = 0.82$, and $n_s = 0.96$, to maintain consistency with previous simulations used by Ihle et al. (2019). Distances carry an implicit h^{-1} dependence throughout, which propagates through masses (all based on virial halo masses, proportional to h^{-1}) and volume densities ($\propto h^3$).

2 DDE MOTIVATION

In this section, we discuss the fundamentals of the DDE but refer the reader to Breysse et al. (2022b) for further theoretical grounding.

While the LIM signal traces large-scale structure and therefore has two-point correlations in comoving space as described by the power spectrum (or higher point correlations – see e.g. Beane & Lidz 2018), we may model the one-point statistics of the signal approximately as a random variable following some probability distribution $\mathcal{P}_S(T)$. The variance of this signal is entirely independent of the variance of the noise in the LIM survey, which follows its own probability distribution $\mathcal{P}_N(T)$. The noise is Gaussian for an ideal spectroscopic cube from a radiometer survey; the signal is strongly non-Gaussian and asymmetric.

In the case of COMAP reionization observations, the Ka-band and Ku-band measurements of temperatures T_{Ka} and T_{Ku} will contain correlated signals due to CO(2–1) and CO(1–0) emission from $z \sim 7$, respectively. However, the Ka-band will also contain a lower redshift interloper component in the form of CO(1–0) emission from $z \sim 3$. This component follows its own probability distribution $\mathcal{P}_I(T_{\mathrm{Ka}})$, which is independent of the signals[1] and of the radiometer noise.

The probability distribution of the sum of independent random variables is the convolution of the probability distributions of the individual variables. Therefore the observed brightness temperature, being the sum of signal and noise, will follow a distribution obtained by convolving the signal and noise distributions:

$$\mathcal{P}(T) = \mathcal{P}_S * \mathcal{P}_N(T). \tag{1}$$

In cases like Ka-band COMAP observations, the interloper distribution also enters the convolution:

$$\mathcal{P}(T_{\mathrm{Ka}}) = \mathcal{P}_S * \mathcal{P}_{N,\mathrm{Ka}} * \mathcal{P}_I(T_{\mathrm{Ka}}). \tag{2}$$

We provide a graphical representation in Fig. 1. Note that in the latter case, we can treat the interloper as an additional non-Gaussian component of 'noise', meaning we treat $\mathcal{P}_{N,\mathrm{Ka}} * \mathcal{P}_I$ as a total noise distribution $\mathcal{P}_N$.

Now consider the joint probability distribution between independent observations of temperatures T_1 and T_2 (T_{Ka} and T_{Ku} in the case of reionization-epoch COMAP). If the noise is independent between observations and follows distributions $\mathcal{P}_{N1}(T_1)$ and $\mathcal{P}_{N2}(T_2)$, common knowledge holds that

$$\mathcal{P}_N(T_1, T_2) = \mathcal{P}_{N1}(T_1)\mathcal{P}_{N2}(T_2). \tag{3}$$

On the other hand, if the signals in each observation are perfectly correlated (which is close to what simulations suggest for the low-J CO lines – see Yang et al. 2022), and follow effectively the same probability distribution $\mathcal{P}_S$ for both T_1 and T_2 (up to any difference

---

[1] We expect the clustering of the emission to be largely decorrelated between the interloper and the signals, given the large interval between $z \sim 3$ and $z \sim 7$.
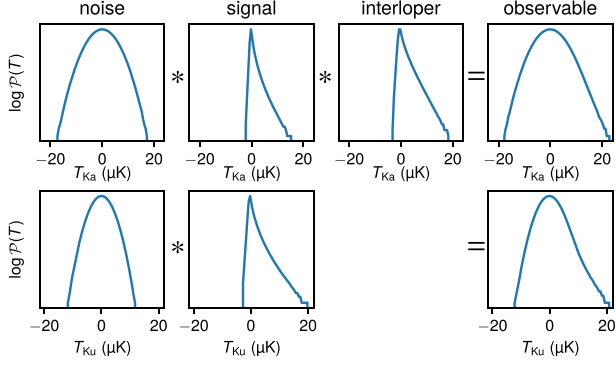
Figure 1. Illustration for Ka- and Ku-band observations (upper and lower rows, respectively) of 1D distributions of noise, (mean-subtracted) signal, and (for Ka-band only) interloper voxel intensities being convolved into the observed probability distribution.
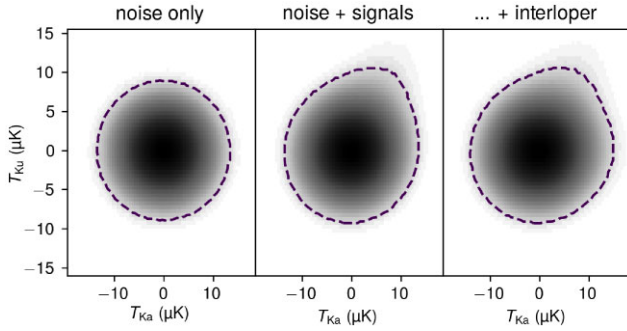

Figure 2. Bivariate probability distributions for voxel intensity given perfect white noise in both Ka- and Ku-bands (left-hand panel), white noise plus correlated signals (middle), and with the additional component of a Ka-band interloper line (right-hand panel). We also draw $\sim 3.5\sigma$ contours (dashed) to aid visualization of how signal and interloper intensity distributions convolve with the noise-only distribution.

in normalization, which can be removed with rescaling), then the joint distribution would be the product of the signal distribution and a delta function:

$$\mathcal{P}_S(T_1, T_2) = \mathcal{P}_S(T_1)\delta(T_2 - T_1). \tag{4}$$

As with the univariate distributions, the joint distribution of the total observed T_1 and T_2 follows the convolution of the signal and noise distributions:

$$\mathcal{P}(T_1, T_2) = \mathcal{P}_S * \mathcal{P}_N(T_1, T_2). \tag{5}$$

Without working explicitly through the maths, we can illustrate graphically in Fig. 2 that the joint distribution of the total signal-plus-noise T_1 and T_2 in this case has a clear correlation between higher T_1 and higher T_2, a result of the signal distribution that modifies the noise distribution (which we also show) through convolution. In the case of the reionization-epoch COMAP observation, the interloper line also modifies the distribution through convolution, but purely in the T_{Ka} dimension and introducing no correlation with T_{Ku}. This makes sense because the interloper joint distribution would be entirely a function of T_{Ka} and thus completely separable (and absorbable into $\mathcal{P}_{N,Ka}(T_{Ka})$ as before). However, it remains the case that this interloper convolution interferes with our main correlated signal. So we want to devise a statistic that is based on the joint $\mathcal{P}(T_1, T_2)$ but sensitive only to $\mathcal{P}_S$.

Now consider the Fourier transforms $\widetilde{\mathcal{P}}$ of each of the univariate and joint distributions considered here, and in particular the description of the Fourier transform of the joint distributions in terms of the Fourier transforms of the univariate distributions. We use $\widetilde{T}_1$ and $\widetilde{T}_2$ to denote the Fourier duals of the temperature variables T_1 and T_2.

The Fourier transform of a convolution of functions is the product of the Fourier transforms of the individual functions, univariate or otherwise:

$$\widetilde{\mathcal{P}}\left(\widetilde{T}_1\right) = \widetilde{\mathcal{P}}_S\left(\widetilde{T}_1\right) \widetilde{\mathcal{P}}_{N1}\left(\widetilde{T}_1\right) \tag{6}$$

$$\widetilde{\mathcal{P}}\left(\widetilde{T}_2\right) = \widetilde{\mathcal{P}}_S\left(\widetilde{T}_2\right) \widetilde{\mathcal{P}}_{N2}\left(\widetilde{T}_2\right) \tag{7}$$

$$\widetilde{\mathcal{P}}\left(\widetilde{T}_1, \widetilde{T}_2\right) = \widetilde{\mathcal{P}}_S\left(\widetilde{T}_1, \widetilde{T}_2\right) \widetilde{\mathcal{P}}_N\left(\widetilde{T}_1, \widetilde{T}_2\right) \tag{8}$$

For separable joint distributions like the noise distribution described in equation (3), the Fourier transform is simply the product of the Fourier transforms of the separated distributions:

$$\widetilde{\mathcal{P}}_N\left(\widetilde{T}_1, \widetilde{T}_2\right) = \widetilde{\mathcal{P}}_{N1}\left(\widetilde{T}_1\right) \widetilde{\mathcal{P}}_{N2}\left(\widetilde{T}_2\right). \tag{9}$$

In the specific case of the perfectly correlated joint signal described in equation (4), working through the Fourier transform results in

$$\widetilde{\mathcal{P}}_S\left(\widetilde{T}_1, \widetilde{T}_2\right) = \widetilde{\mathcal{P}}_S\left(\widetilde{T}_1 + \widetilde{T}_2\right). \tag{10}$$

In analogue to normalized time-series cross-correlation or covariance (e.g. Pearson's bivariate correlation coefficient), we define the DDE here as a normalized measure of correlation between variables T_1 and T_2 via their Fourier duals $\widetilde{T}_1$ and $\widetilde{T}_2$:

$$\mathcal{D} \equiv \frac{\widetilde{\mathcal{P}}\left(\widetilde{T}_1, \widetilde{T}_2\right)}{\widetilde{\mathcal{P}}\left(\widetilde{T}_1\right) \widetilde{\mathcal{P}}\left(\widetilde{T}_2\right)} - 1. \tag{11}$$

We can then ask what $\mathcal{D}$ is for T_1 and T_2 when including both signal and noise (with any separable interlopers presumed as fully described by one of the univariate Fourier-transformed noise distributions) as described above. Substituting equations (6)–(10) into equation (11), we find

$$\mathcal{D} = \frac{\widetilde{\mathcal{P}}_S\left(\widetilde{T}_1, \widetilde{T}_2\right) \widetilde{\mathcal{P}}_N\left(\widetilde{T}_1, \widetilde{T}_2\right)}{\widetilde{\mathcal{P}}_S\left(\widetilde{T}_1\right) \widetilde{\mathcal{P}}_{N1}\left(\widetilde{T}_1\right) \widetilde{\mathcal{P}}_S\left(\widetilde{T}_2\right) \widetilde{\mathcal{P}}_{N2}\left(\widetilde{T}_2\right)} - 1 \tag{12}$$

$$= \frac{\widetilde{\mathcal{P}}_S\left(\widetilde{T}_1 + \widetilde{T}_2\right) \widetilde{\mathcal{P}}_{N1}\left(\widetilde{T}_1\right) \widetilde{\mathcal{P}}_{N2}\left(\widetilde{T}_2\right)}{\widetilde{\mathcal{P}}_S\left(\widetilde{T}_1\right) \widetilde{\mathcal{P}}_{N1}\left(\widetilde{T}_1\right) \widetilde{\mathcal{P}}_S\left(\widetilde{T}_2\right) \widetilde{\mathcal{P}}_{N2}\left(\widetilde{T}_2\right)} - 1 \tag{13}$$

$$= \frac{\widetilde{\mathcal{P}}_S\left(\widetilde{T}_1 + \widetilde{T}_2\right)}{\widetilde{\mathcal{P}}_S\left(\widetilde{T}_1\right) \widetilde{\mathcal{P}}_S\left(\widetilde{T}_2\right)} - 1. \tag{14}$$

The upshot is that at least in this idealized case, $\mathcal{D}$ depends solely on the one-point statistics of the signal common to both observations, provided that any noise (including interlopers) is independent between the observations. Had there been no common signal – i.e. had we had entirely unrelated signal distributions $\mathcal{P}_S(T_1)$ and $\mathcal{P}_S(T_2)$ in the two measurements – we would have obtained $\mathcal{D} = 0$, so any deviation from $\mathcal{D} = 0$ indicates the presence of a correlated signal in both distributions. (An anticorrelated signal would similarly result in deviation from $\mathcal{D} = 0$, as we could have defined the joint distribution in equation (4) with $\delta(T_1 + T_2)$ and ended up with $\widetilde{P}_S(\widetilde{T}_1 - \widetilde{T}_2)$ as the joint signal distribution in equation (10).)

In practice, we must deal with discrete histograms, and discrete fast Fourier transforms rather than continuous Fourier transforms (with corresponding normalization factors necessary as part of the computational implementation depending on convention). However, the basic idea is sound and we proceed to apply it to simulated COMAP observations.

3 DEMONSTRATION: COMAP SIMULATION

We need to define three key ingredients in order to simulate COMAP survey volumes: a model for line emission (Section 3.1), parameters for the COMAP-ERA survey (Section 3.2), and large numbers of approximate cosmological simulations (Section 3.3) to which we can apply the first two ingredients.

3.1 Fiducial CO model

Of the CO models considered by Breysse et al. (2022a), we use the adaptation of the model of Li et al. (2016) by Keating et al. (2020) at all redshifts. The basic idea behind the model of Li et al. (2016) is to connect halo mass M_h and redshift z to an average star-formation rate based on the empirical model of Behroozi, Wechsler & Conroy (2013a,b), which is assumed to be related linearly to IR luminosity by a conversion factor of $10^{10} \, L_\odot \, (M_\odot \, yr^{-1})^{-1}$. Empirical power-law fits to data then relate IR luminosity to $CO(J \to J - 1)$ luminosity:

$$\log \frac{L_{IR}}{L_\odot} = \alpha_J \left(\log \frac{L_{CO,J}}{L_\odot} + 4.31 \right) + \beta_J. \quad (15)$$

At this step, Keating et al. (2020) diverge from the fiducial model of Li et al. (2016) out of a need (which we share) to model multiple rotational transitions of CO. The Li et al. (2016)–Keating et al. (2020) model uses the work of Kamenetzky et al. (2016), which obtains $\alpha_1 = 1.27$, $\beta_1 = -1.0$, $\alpha_2 = 1.11$, and $\beta_2 = -0.6$ based on a local sample of galaxies. The model also includes random log-normal scatter to mimic variations in galaxy properties, both around the average star-formation rate for fixed halo mass and redshift, and around the average CO luminosity for fixed star-formation rate. For each property there is an independent scatter introduced with a standard deviation of 0.3 in units of dex, but we calculate the star-formation rate for each halo only once and use it to inform both CO line luminosities. This results in imperfectly but non-negligibly correlated CO(1–0) and CO(2–1) line luminosities, although perhaps less correlated than semi-analytic models of CO emission suggest (see e.g. Yang et al. 2022).

We will want to constrain α_1, β_1, α_2, and β_2 at $z \sim 7$. In addition, most summary statistics end up being chiefly sensitive to the luminosity-weighted average CO temperature–bias product $\langle Tb \rangle_J$ for each CO line, defined as

$$\langle Tb \rangle_J = \frac{c^3(1+z)}{8\pi k_B \nu_{rest,J}^3 H(z)} \int dM_h \, \frac{dn}{dM_h} \, L_{CO,J}(M_h, z) \, b(M_h). \quad (16)$$

This calculation requires: a model relation $b(M_h)$ for the linear halo bias, the scaling with which the number density contrast of haloes of virial mass M_h traces the underlying matter density contrast; a halo mass function dn/dM_h to describe halo number densities; the rest frequency of the $CO(J \to J - 1)$ line $\nu_{rest,J} \approx 115.27 \cdot J$ GHz; the Hubble parameter $H(z)$; and standard physical constants, namely the speed of light c and the Boltzmann constant k_B. In this context, we will use the $b(M_h)$ model of Tinker et al. (2010), and a corrected form of the Tinker et al. (2008) halo mass function given by Behroozi et al. (2013a) (which we will also use below in Section 3.3). Given this, the fiducial model results in $\langle Tb \rangle_1 = 3.37 \, \mu K$ and $\langle Tb \rangle_2 = 2.51 \, \mu K$.

We will refer to inferences both of $\{\alpha_J, \beta_J\}$ and of $\langle Tb \rangle_J$, with the latter calculable from the former by obtaining credibility intervals for $\langle Tb \rangle_J$ implied by the $L_{CO,J}(M_h)$ relations found from the credibility intervals for α_J and β_J. In addition, we will not assume that the IR–CO relation is the same at $z \sim 3$ as it is at $z \sim 7$, and will allow $\alpha_{1,z\sim3}$ and $\beta_{1,z\sim3}$ to be inferred separately from α_1 and β_1 (with the latter pair taken to be referring to $z \sim 7$ parameters).

Table 1. Key parameters for COMAP-ERA observations either taken from Breysse et al. (2022a) or newly assumed for this work.

Parameter	Values for:	
	Ku band	Ka band
Frequency coverage (GHz)	13–17	26–34
Nominal system temperature T_{sys} (K)	20	44
Spectrometer count per dish N_{spd}	38	19
Science channelization $\delta\nu$ (MHz)	7.8125	15.625
Dish-hours per field $N_{dish}t_{obs}$ (h)	57 000	110 000
Noise per voxel σ_N (μK)	2.55	3.85

3.2 COMAP parameters

We base our CO survey parameters on the COMAP-ERA parameters considered by Breysse et al. (2022a). The survey is assumed to span three fields of $2° \times 2°$ each, covering frequencies of 13–17 GHz in the Ku band and 26–34 GHz in the Ka band. These volumes are discretized as a grid of $N_{pix} \times N_{ch} = (30 \times 30) \times 512$ voxels, each spanning four arcminutes in both angular dimensions to match the beam size (assumed to be 4.5 arcmin across the Ka band, and 3.9 arcmin across the Ku band), and 7.8125 or 15.625 MHz respectively in the Ku or Ka band.

Based on the nominal system temperature T_{sys} of each receiver, the spectrometer count[2] per dish N_{spd}, the science channel bandwidth $\delta\nu$, the pixel count N_{pix} per field, and the number of dish-hours per field $N_{dish}t_{obs}$, we can calculate the radiometer noise per voxel as

$$\sigma_N = \frac{T_{sys}}{\sqrt{\delta\nu N_{spd} N_{dish}t_{obs}/N_{pix}}}. \quad (17)$$

We show the resulting values for both Ku- and Ka-band observations in Table 1 alongside a summary of other key parameters.

3.3 Peak-patch simulations

As was the case for the VID at the time of writing of Ihle et al. (2019), the covariance matrix of our observables has aspects not well-understood in a purely analytic formalism, and we will use a large number of simulations to estimate the covariance matrix numerically. We use the peak-patch method (Stein, Alvarez & Bond 2019) to obtain large numbers of dark matter simulations and halo catalogues from independent sets of initial conditions.

Ihle et al. (2019) used 161 independent peak-patch light-cone simulations spanning a comoving volume of $L_{box}^3 = (1140 \, Mpc)^3$ with a resolution of $N_{cells} = 4096^3$. As the light cone extent was equivalent to $9.6° \times 9.6°$ in transverse dimensions and 26–34 GHz in CO(1–0) observing frequency, Ihle et al. (2019) were able to split each light cone into many subfields to simulate thousands of semi-independent COMAP observations at $z \sim 3$. We will still be able to use these same simulations to generate realizations of the $z \sim 3$ CO(1–0) interloper emission in the Ka-band mock data cube. However, we require additional peak-patch simulations at high redshift, and with slightly better mass resolution. Given the filtering scales used in the peak-patch method to find matched density peaks, the minimum resolvable halo mass is given by

---

[2]Breysse et al. (2022a) use the term 'feeds'; we attempt to disambiguate further. Each Ka-band dish will host 19 single-polarization feeds and thus 19 spectrometers; each Ku-band dish will host 19 dual-polarization feeds and thus 38 spectrometers.

$$M_{\text{h,min-res}} = \frac{4}{3}\pi\,\Omega_m\rho_c\left(\frac{2L_{\text{box}}}{N_{\text{cells}}^{1/3}}\right)^3 \tag{18}$$

$$= 9.3\times10^{12}\Omega_m h^2\,\frac{(L_{\text{box}}/\text{Mpc})^3}{N_{\text{cells}}}\,\text{M}_\odot \tag{19}$$

$$= 1.3\times10^{12}\,\frac{(L_{\text{box}}/\text{Mpc})^3}{N_{\text{cells}}}\,\text{M}_\odot. \tag{20}$$

For the peak-patch simulations used by Ihle et al. (2019) this is $2.8\times10^{10}\,\text{M}_\odot$, which is within 10 per cent of the minimum halo mass quoted by Ihle et al. (2019) of $2.5\times10^{10}\,\text{M}_\odot$ (noting that Ihle et al. (2019) correct the halo masses obtained from the peak-patch method via abundance matching to the Tinker et al. (2008) HMF). While this is sufficient for the $z\sim3$ CO signal, the $z\sim7$ CO signal will have greater contribution from the lower mass halo population and we want to make sure the statistics for that population are as correct as possible down to at least several times lower mass. At the same time, we still want to obtain sufficiently large volumes to be divisible into many semi-independent $z\sim7$ COMAP observations. We therefore choose a box size of $L_{\text{box}} = 960\,\text{Mpc}$ and $N_{\text{cells}} = 5640^3$, so that $M_{\text{h,min-res}} = 6.4\times10^9\,\text{M}_\odot$.

The original $z\sim3$ simulations used the SciNet General Purpose Cluster (GPC; Loken et al. 2010), before its decommissioning in 2018 April. Each run used 900 s of compute time on 2048 Intel Xeon E5540 cores (or 256 nodes), using $\simeq2.4$ TB of RAM (roughly half the available RAM, as most GPC nodes had 16 GB RAM). Our new $z\sim7$ simulations use the successor to the GPC, the Niagara cluster (Ponce et al. 2019), whose nodes use a mix of Intel Xeon 6148 and 6248 CPU cores at the time of this work. These simulations also use ≈900 s runtime per realization with 1880 cores (or 47 nodes), but with a peak memory footprint of $\simeq7.5$ TB (around 90 per cent of the available RAM) owing to the higher resolution. Just as the 161 $z\sim3$ realizations from Ihle et al. (2019) took only $\sim82\,000$ CPU hours, we are able to generate 270 realizations after only $\approx127\,000$ core-hours. We estimate (as did Ihle et al. 2019; Stein et al. 2019) that this is three orders of magnitude faster than an equivalent N-body simulation.

The resulting halo catalogues span $z = 5.8{-}7.9$ and $6°\times6°$. As with the $z\sim3$ halo catalogues, we correct peak-patch halo masses through abundance matching to the HMF of Tinker et al. (2008), but in this case with high-redshift corrections from Appendix G of Behroozi, Wechsler & Conroy (2013b). We use these corrections only in mass-correcting our new simulations as they principally apply at $z > 3$.

By dividing the $z\sim3$ simulations into 16 sub-patches each spanning $2°\times2°$, and the $z\sim7$ simulations into 9 sub-patches each spanning the same, we obtain 2430 semi-independent $2°\times2°$ sky realizations, which we believe should be sufficient for a reasonably high-quality estimation of the covariance in the context of this early investigatory work.

We use `limlam_mocker`[3] to assign halo luminosities and create simulated brightness temperature cubes for Ku-band CO(1–0) and Ka-band CO(2–1), as well as the Ka-band CO(1–0) interloper. We also include the line broadening model described by Chung et al. (2021). We also apply a high-pass transfer function to mimic the effect of the COMAP pipeline filtering out large-scale angular and frequency modes as described by Foss et al. (2022) and Ihle et al. (2022). However, while we adapt the form suggested by Chung et al. (2022), we adjust the exponents slightly to reflect that we project the data cube onto comoving space at $z\sim7$, rather than at $z\sim3$ as for
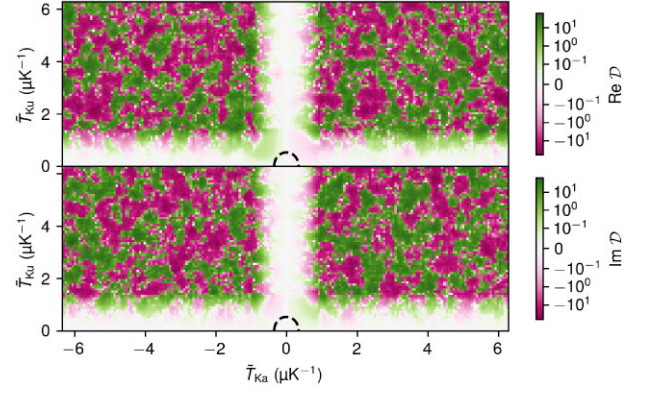
[3]https://github.com/georgestein/limlam_mocker

Figure 3. Illustration of the DDE real and imaginary parts (upper and lower sub-panels, respectively) when all signal, noise, and interloper cubes are accounted for, using the median across all realizations. The dashed ellipse indicates the low-variance cut in both $\widetilde{T}$ dimensions described in the main text as a function of the noise distribution.

COMAP Pathfinder analyses:

$$\mathcal{T}_{\text{hp}}(k_\perp, k_\parallel) = \frac{1}{\left(1 + e^{5-138\text{Mpc}\cdot k_\perp}\right)\left(1 + e^{5-144\text{Mpc}\cdot k_\parallel}\right)}, \tag{21}$$

where $k_\parallel$ is the line-of-sight component of the wavevector $\boldsymbol{k}$, and $k_\perp$ is the magnitude of the transverse part of $\boldsymbol{k}$ (so $k_\perp^2 + k_\parallel^2 = k^2$).

We calculate the following summary statistics for each data cube:

(i) the Ku-band auto power spectrum, with 101 k-bins of width $\Delta k = 0.02$, centred at values ranging from 0.01 to $2.09\,\text{Mpc}^{-1}$;

(ii) the Ka-band auto power spectrum, with the same k-bins used for the Ku-band;

(iii) the cross power spectrum between the Ku- and Ka-band data cubes, again using the same k-bins;

(iv) the Ku-band 'auto' VID, using a histogram across $0.5\,\mu\text{K}$-wide bins spanning $T_{\text{Ka}}\in(-36, 36)\,\mu\text{K}$;

(v) the Ka-band 'auto' VID, using a histogram across $0.5\,\mu\text{K}$-wide bins spanning $T_{\text{Ka}}\in(-48, 48)\,\mu\text{K}$;

(vi) and the DDE, based on fast Fourier transforms of the individual and joint VID with the same binning of T_{Ku} and T_{Ka} as used for the individual VID calculations.

4 RESULTS

4.1 Visual inspection

We show the typical DDE for our simulated observations in Fig. 3, and show slices of the DDE with different components of the simulation included or excluded in Fig. 4. Since the noise distribution is the most visible component, we suggest that the most usable modes of the DDE (i.e. those least affected by noise) will be larger scale modes than the Fourier scale set by the instrumental noise. We set the cutoff to correspond to when the half-wavelength of the mode equals the full-width at half-maximum of the noise distribution:

$$\widetilde{T}_{\text{Ka}} < \frac{\pi}{2.355\sigma_{N,\text{Ka}}}; \tag{22}$$

$$\widetilde{T}_{\text{Ku}} < \frac{\pi}{2.355\sigma_{N,\text{Ku}}}. \tag{23}$$

The values at $\widetilde{T} = 0$ are also unusable as they are always equal to $0 + 0i$, and we only choose the modes with $\widetilde{T}_{\text{Ku}} > 0$ as the other half of the Fourier modes in the DDE are completely dependent. For our $0.5\,\mu\text{K}$-wide bins spanning $T_{\text{Ka}}\in(-48, 48)\,\mu\text{K}$ and $T_{\text{Ku}}\in(-36,$
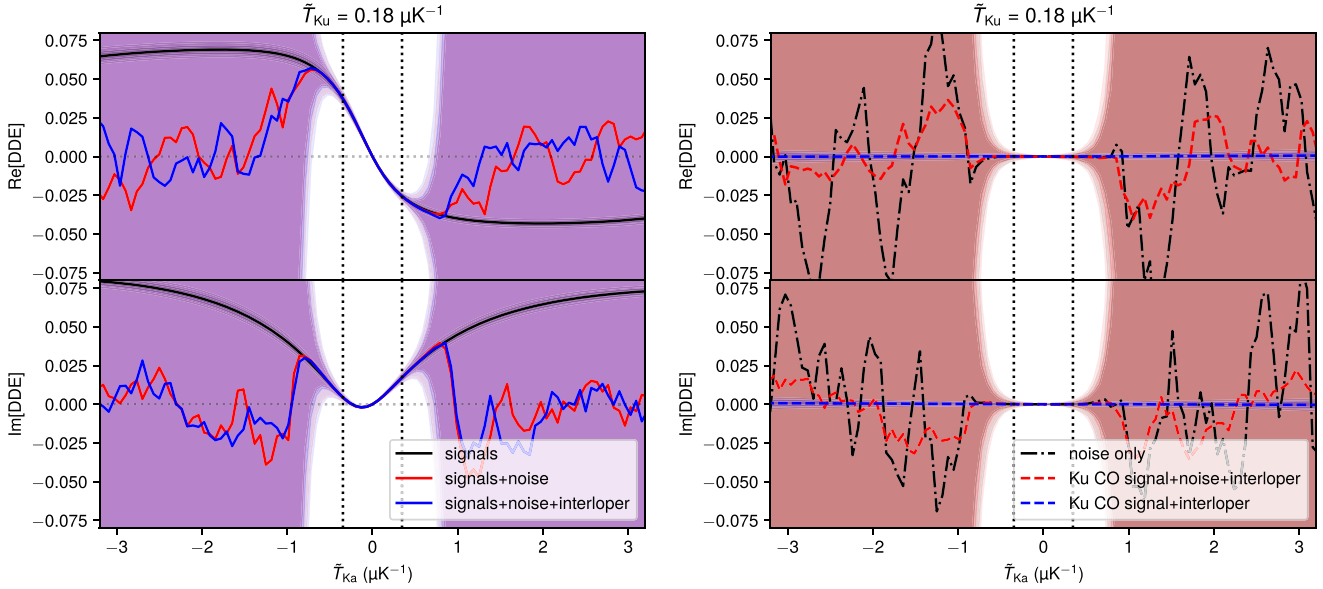
Figure 4. Slices of the DDE and low-variance cut (now shown as dotted vertical lines), again showing real and imaginary parts separately (upper and lower sub-panels, respectively, within each panel). Across the left-hand and right-hand panels, we show how the DDE changes if certain simulation components are omitted, in particular whether the Ka-band observation contains the CO(2–1) signal correlated with the Ku-band CO(1–0) signal (left-hand panel) or whether the Ka-band observation contains no such correlated signal (right-hand panel). Dark (light) shaded areas around each curve indicate 68 per cent (95 per cent) sample intervals, with the presence of noise significantly increasing interval widths.

36) μK, this in principle leaves 32 usable DDE bins,[4] each with a real and imaginary part.

Looking at the increase in DDE variance introduced by noise in Fig. 5, it is qualitatively true that the noise variance is much lower inside the $\widetilde{T}$ cutoff than outside it. It is possible that the cut is somewhat conservative and DDE data at slightly larger values of $\widetilde{T}$ could be usable, but considering only the bins within the cutoff will be sufficient for the exploratory purposes of this work.

Of interest is whether what we expect based on the continuous distribution formalism of Section 2 holds in these numerical simulations. With our fiducial binning the outcome is as good as can be expected, at least by eye. But overall outlook here is mixed due to some numerical instability. In a preliminary version of this analysis that used coarser 2 μK bins, introducing noise did in fact affect the DDE by dampening its deviation from $0 + 0i$, contrary to our expectation. On the other hand, once noise was introduced in addition to the correlated signals, introducing the Ka-band interloper emission did not appear to further affect the DDE even in this case. Furthermore, in that case also, when there was no correlated signal in the Ku-band simulated observation, taking the DDE between the Ka- and Ku-band cubes resulted in a value of $0 + 0i$ throughout, as expected.

All this suggests some robustness of the DDE (or rather the fact of its deviation from $0 + 0i$) modulo choice of binning, and with sufficiently fine histogramming[5] the robustness matches
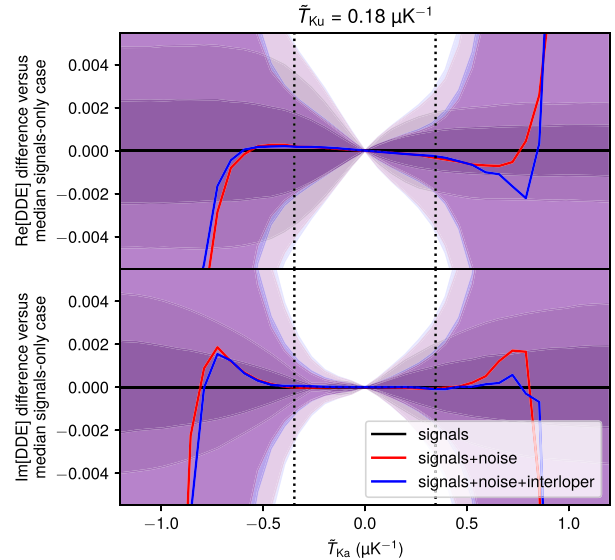

Figure 5. Same as the left-hand panel of Fig. 4, but now showing residuals relative to the median obtained from realizations with only the correlated signals. Again, dark (light) shaded areas around each curve indicate 68 per cent (95 per cent) sample intervals.

theoretical expectations much better. Furthermore, even before we consider astrophysical interpretation, the level of variance in Fig. 4 suggests a very high-significance detection of correlated distributions is achievable with COMAP-ERA.

---

[4]By contrast, using the present binning scheme, only two DDE bins would be usable with data from the COMAP-EoR experiment preceding COMAP-ERA, also conceptualized by Breysse et al. (2022a). However, an actual COMAP-EoR analysis would tailor the VID binning appropriately to the wider noise distribution, so would have more than just two DDE bins available.

[5]Note in particular that both the noise per voxel σ_N and the CO $\langle T_b \rangle_J$ are around 3 μK, so the histogramming likely needs to oversample this typical

fluctuation magnitude by a factor of several, which is the case with 0.5 μK bins but not 2 μK bins.

The important aspect for parameter inference, however, is the covariance (which we will estimate numerically) and the interplay of the covariance with the partial derivatives of the DDE with respect to each of the high-redshift CO parameters α_J and β_J as well as the nuisance low-redshift CO parameters $\alpha_{1, z \sim 3}$ and $\beta_{1, z \sim 3}$. We estimate this numerically from a smaller set of 96 simulations, for which we repeat calculations of all summary statistics while shifting one parameter away from its fiducial value. α_J and $\alpha_{1, z \sim 3}$ are allowed to shift ± 0.02 away from their fiducial values; β_J and $\beta_{1, z \sim 3}$ shift a larger ± 0.2 away as the signals are not as sensitive to the same absolute change in those parameters.

We show the resulting central difference quotient estimates in Fig. 6. We note with interest that the imaginary part of the DDE reacts in a visibly qualitatively different way to shifts in α_2 compared to shifts in α_1; the same applies to β_J. That is, changing CO(2–1) parameters alters the shapes of the individual and joint VID in a different way to changing CO(1–0) parameters. Different temperature-space offsets introduced in the distribution shapes would correspond to different phases introduced in the DDE and thus the different partial derivatives for the real and imaginary parts of the DDE that we see in Fig. 6 against α_2 and β_2 versus against α_1 and β_1. This will become important in a moment in considering the astrophysical constraints possible with this measure, as we will discuss in Section 4.2.

Note also that we find non-zero partial derivatives of the DDE against nuisance parameters. As the range of values indicated in Fig. 6 demonstrates, the DDE is still significantly less sensitive to $\alpha_{1, z \sim 3}$ and $\beta_{1, z \sim 3}$ compared to α_J and β_J, and it is likely that some if not much of what we find is numerical noise. None the less, given the consistency of some qualitative trends between the partial derivatives with respect to $\alpha_{1, z \sim 3}$ and $\beta_{1, z \sim 3}$, some real trends may well exist. The effect of discrete VID binning cannot be discounted here, given our earlier observation that coarser binning reduced robustness of the DDE against noise. It is similarly possible that even finer binning than considered here would reduce any of these apparently real trends related to interloper parameters, which we leave for future, more extensive work to examine in greater detail. Either way, in this work we will end up forecasting non-zero constraining power on these $z \sim 3$ CO parameters, although the DDE should constrain α_J and β_J somewhat more tightly.

4.2 Fisher analysis

Our simulations allow us to numerically estimate the covariance matrix given the fiducial parameters and the full suite of 2430 simulations described in Section 3.3. The previously discussed smaller set of 96 simulations at various points in the local neighbourhood is sufficient to then estimate derivatives of observables with respect to the $z \sim 7$ α_J and β_J as well as $\alpha_{1, z \sim 3}$ and $\beta_{1, z \sim 3}$. This in turn allows us to run a Fisher analysis to estimate the constraining power on the $z \sim 7$ α_J and β_J, with $\alpha_{1, z \sim 3}$ and $\beta_{1, z \sim 3}$ still allowed to vary as nuisance parameters to allow for the possibility of redshift evolution (although the fiducial values are the same as α_1 and β_1 at $z \sim 7$).

Note that of the 64 real variables associated with the usable DDE bins (32 real parts and 32 imaginary parts), only 29 are usable after masking variables until no pair of variables has a correlation coefficient above 0.98. The masked DDE then has an acceptably stable numerically estimated covariance matrix.[6] We graphically

show the correlation coefficient matrix (which has the same structure as the covariance matrix, except with variances for each observable normalized away) in Appendix A.

The numerically estimated covariance matrix is calculated relative to one simulated $2° \times 2°$ patch, so we divide it by three to reflect the fact that the full COMAP-ERA survey is of three independent fields. Once we have the full COMAP-ERA covariance matrix $\mathbf{C}_{ab}$ for a vector of combined or individual observables $\boldsymbol{O}_a$, the Fisher matrix in the basis of parameters $\{\lambda_i\} = \{\alpha_1, \alpha_2, \beta_1, \beta_2, \alpha_{1, z \sim 3}, \beta_{1, z \sim 3}\}$ is

$$\mathbf{F}_{ij} = \sum_{a,b} \frac{\partial \boldsymbol{O}_a}{\partial \lambda_i} \mathbf{C}_{ab}^{-1} \frac{\partial \boldsymbol{O}_b}{\partial \lambda_j}. \tag{24}$$

Inverting this will yield the parameter covariance matrix, allowing us to calculate expected parameter constraints under the assumption of Gaussian covariances throughout. As some non-Gaussianity is present in our signals and variances, and as we numerically estimate all covariances and derivatives informing the Fisher forecast, we only trust the resulting parameter constraints on a qualitative level. In other words, the exact shapes and widths of the posterior distributions should not be considered robust predictions, but relative strengths of constraints between different sets of observables should be credible.

4.2.1 Cross-statistics only

We show expected resulting constraints on α_J and β_J, as well as on the nuisance parameters $\alpha_{1, z \sim 3}$ and $\beta_{1, z \sim 3}$, from the Ka–Ku cross-spectrum $P_{\times}(k)$ and the DDE in the upper panels of Fig. 7. The expected achievable constraining power of the DDE is extremely high, far beyond that expected for the cross power spectrum. The fundamental sensitivity of the DDE appears to allow for very strong parameter constraints by itself, without use of individual VID data or power spectra. This may be a peculiar aspect of the reionization-epoch CO signal or even our model of it, and may not necessarily hold in other LIM contexts.

We also note that the DDE does actually appear sensitive to $\alpha_{1, z \sim 3}$ and $\beta_{1, z \sim 3}$. However, the sensitivity to those parameters is two to three times poorer than to any of the corresponding $z \sim 7$ α_J and β_J parameters. We tabulate the Fisher forecast errors for these parameters in Appendix B, not only for $P_{\times}(k)$ and/or the DDE but for all statistics considered in this work.

Using random draws from the parameter confidence ellipses of Fig. 7, we can translate expected constraints on α_J and β_J into constraints on $\langle Tb \rangle_J$, shown in the lower panels of Fig. 7. With the Ka–Ku cross-spectrum alone, the degeneracy between α_J or between β_J represents an inability to constrain $\langle Tb \rangle_J$ independently, as the cross-spectrum is proportional to their product $\langle Tb \rangle_1 \langle Tb \rangle_2$. The cross-spectrum constraint on this product is $\langle Tb \rangle_1 \langle Tb \rangle_2 = 8.6^{+1.1}_{-0.6}$ (68 per cent interval).

However, as previously noted the DDE is somehow sensitive to CO(1–0) and CO(2–1) in different ways to the cross-spectrum, thus allowing for constraints of $\langle Tb \rangle_1$ and $\langle Tb \rangle_2$ separately. This is of extreme interest as it would allow us to simultaneously probe excitation through line ratios and molecular gas content at high redshift. The DDE constraint on the product of the two temperature–

---

[6]Despite high correlations between the remaining DDE bins, the masked DDE covariance matrix has a condition number of $\sim 10^7$, which is acceptable
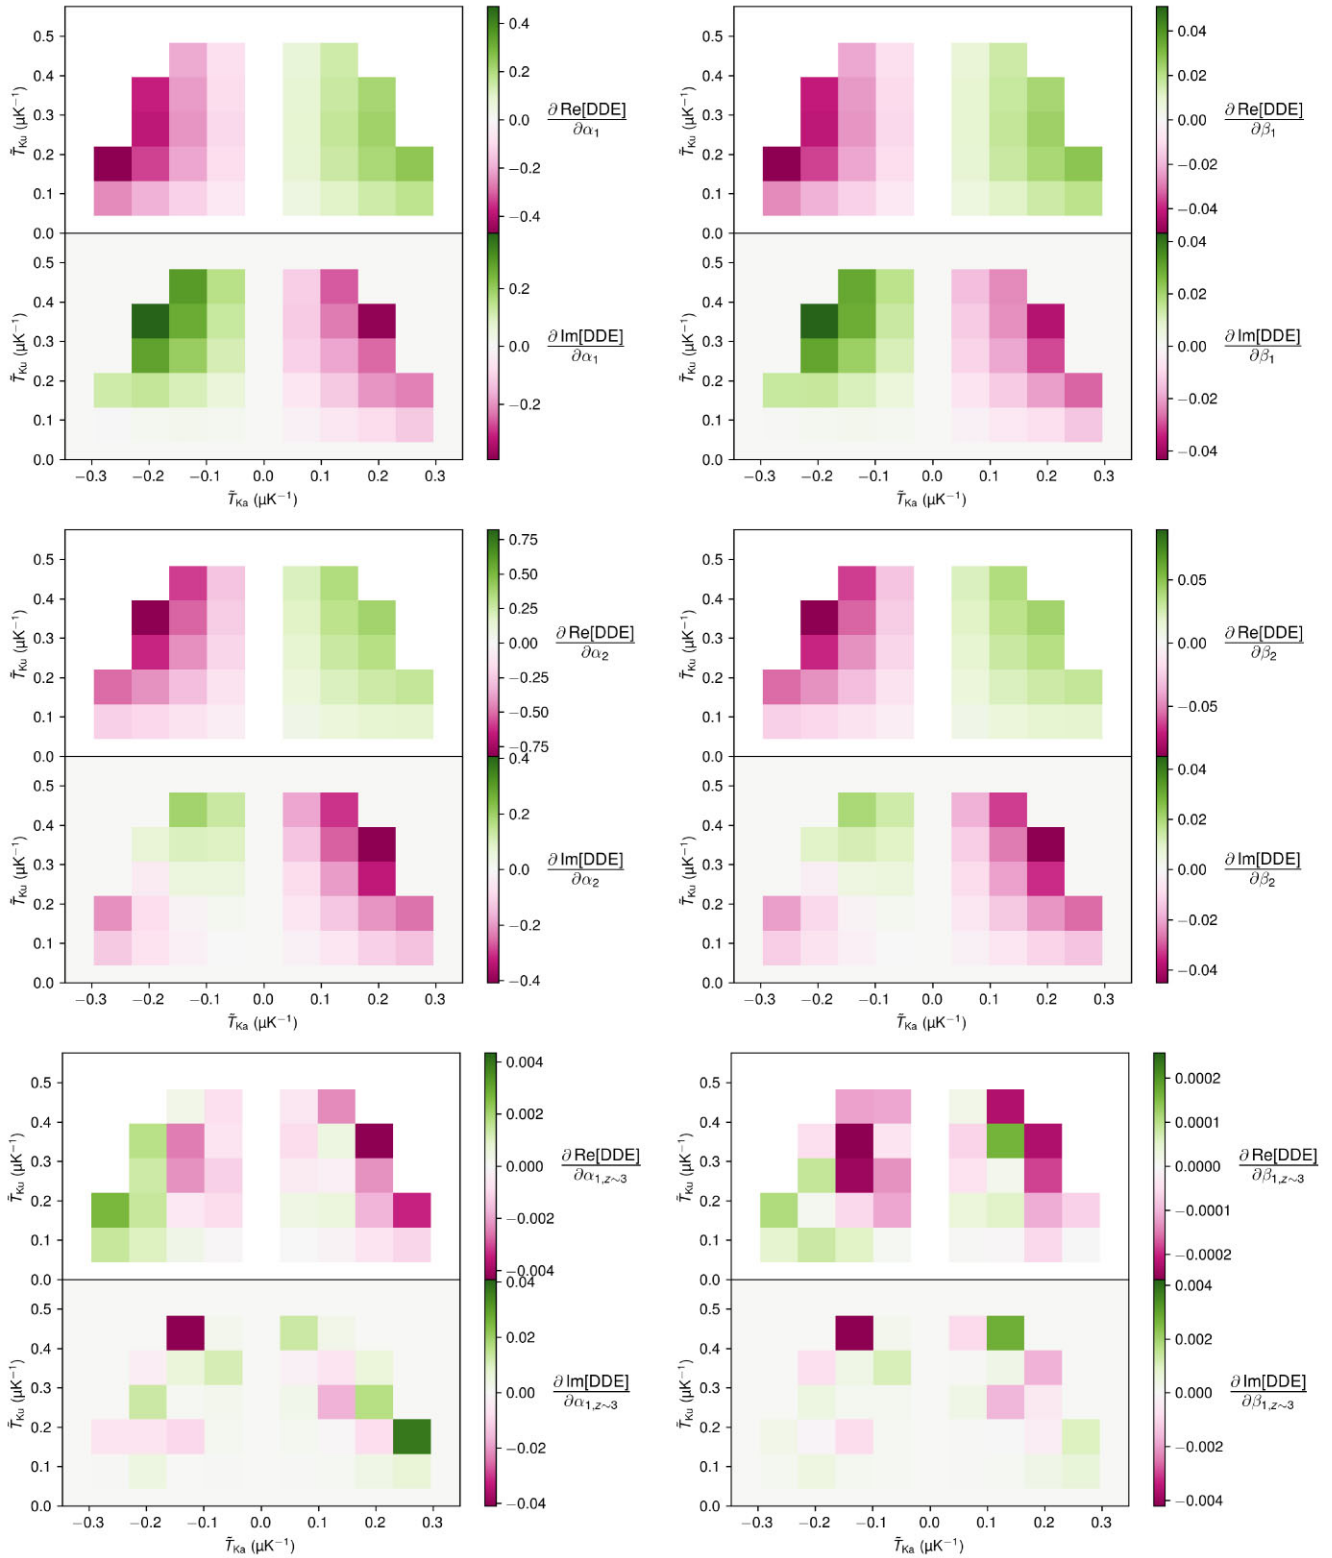
for 64-bit float arithmetic. When combining the masked DDE with other variables, we rescale all DDE values by a factor of 10^6 to decrease the range of variances and improve the condition of the covariance matrix to be well beyond what is required given 64-bit float precision.

Figure 6. Partial derivatives of the real and imaginary parts of the DDE with respect to model parameters, as indicated by the label next to the colour bars. Compare in particular the partial derivatives of the imaginary part of the DDE for negative $\widetilde{T}_{Ka}$ against α_1 or β_1 versus those against α_2 or β_2.
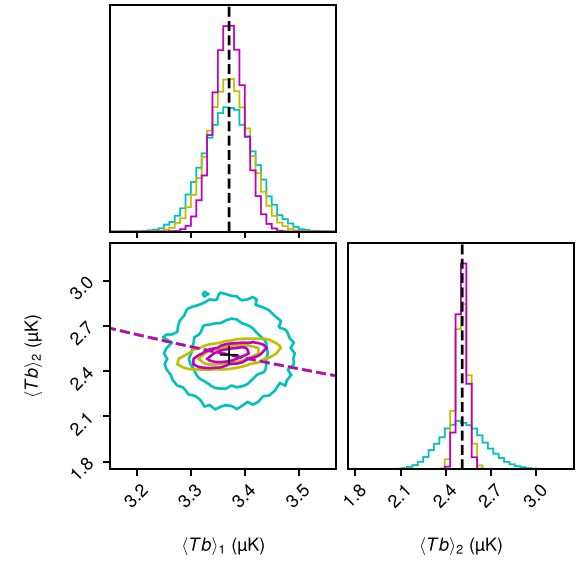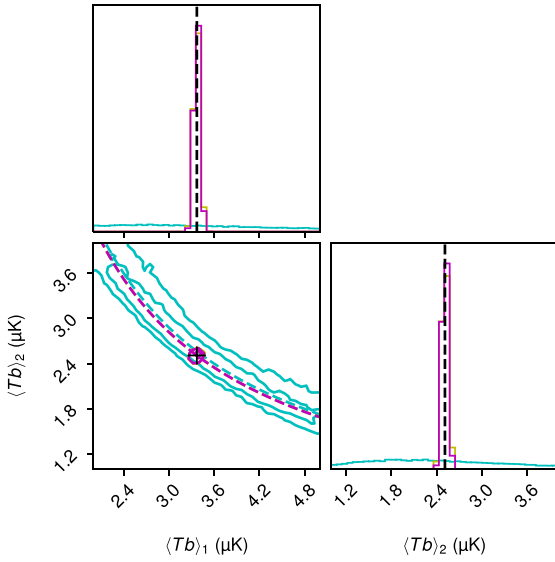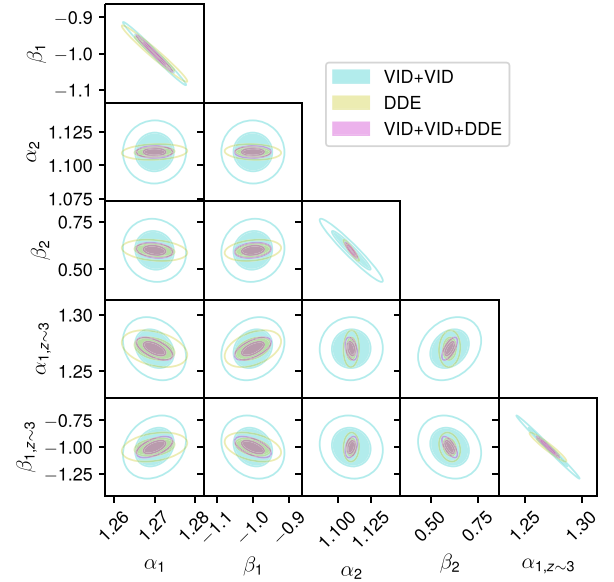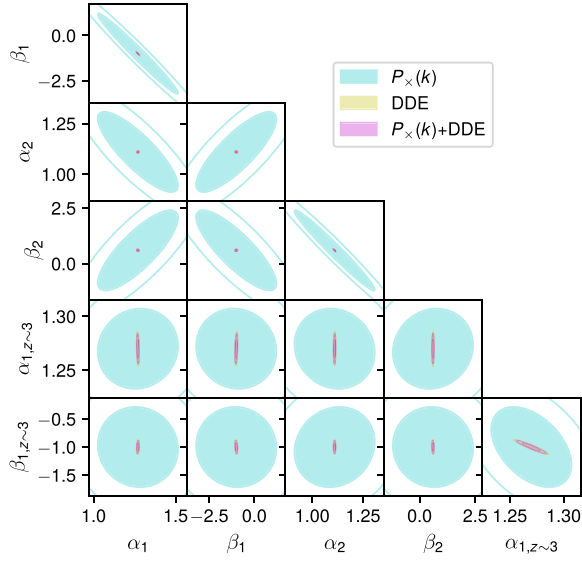
Figure 7. Corner plots of CO model parameter constraints (upper portion of plot) and constraints on $\langle Tb \rangle_J$ (lower portion of plot) based on the cross power spectrum alone (cyan), the masked DDE alone (yellow), and the combination of the two (magenta). (Here, the last is effectively the same as the DDE-only case, obscuring the DDE-only contours as a result. Figs 8 and 9 show the DDE-derived constraints in greater detail.) The filled (unfilled) ellipses or inner (outer) contours in each triangle plot indicate 68 per cent (95 per cent) credibility intervals. The dashed lines in the $\langle Tb \rangle_1$–$\langle Tb \rangle_2$ contour plot mark out $\langle Tb \rangle_1 \langle Tb \rangle_2 = 8.5\ \mu K^2$, which is consistent with constraints from all summary statistics.

bias products is also still tighter than with the cross spectrum, at $\langle Tb \rangle_1 \langle Tb \rangle_2 = 8.5 \pm 0.2$ (68 per cent interval).

4.2.2 One-point statistics only

We repeat the exercise but replace the cross-spectrum with the VID from both bands (excising the first 50 and last 30 bins so that we only consider bins with significant variance). The constraints shown in Fig. 8 suggest that the DDE significantly collapses the parameter space volume surrounding the $z \sim 7$ CO(2–1) parameters α_2 and

Figure 8. Same as Fig. 7, but now comparing the DDE against the combination of the individual Ku- and Ka-band VID data (possibly also combined with the DDE) instead of the cross power spectrum.

β_2 (as well as their nuisance $z \sim 3$ CO(1–0) counterparts). This is sensible as the Ka-band VID has significant interloper emission and is thus less sensitive to the $z \sim 7$ signal than either the Ku-band VID (which has no such interloper to contend with) or the DDE (which largely rejects the interloper emission by design).

The $\langle Tb \rangle_1$–$\langle Tb \rangle_2$ constraints shown in Fig. 8 reflect this, as both the VID and the DDE constrain $\langle Tb \rangle_1$ and $\langle Tb \rangle_2$ with little degeneracy but the DDE significantly narrows the plausible range of $\langle Tb \rangle_2$. This in turn also happens to significantly narrow the plausible range of $\langle Tb \rangle_1 \langle Tb \rangle_2$, from $\langle Tb \rangle_1 \langle Tb \rangle_2 = 8.5 \pm 0.5$ (68 per cent interval) for the auto VID data alone to the previously mentioned $\langle Tb \rangle_1 \langle Tb \rangle_2 = 8.5 \pm 0.2$ with the DDE.

Thus, in principle, the fundamental sensitivity of COMAP-ERA would allow strong constraints on reionization-epoch CO from one-point statistics alone. This does assume the data cubes are sufficiently clean and free of systematics and foregrounds, but again this would
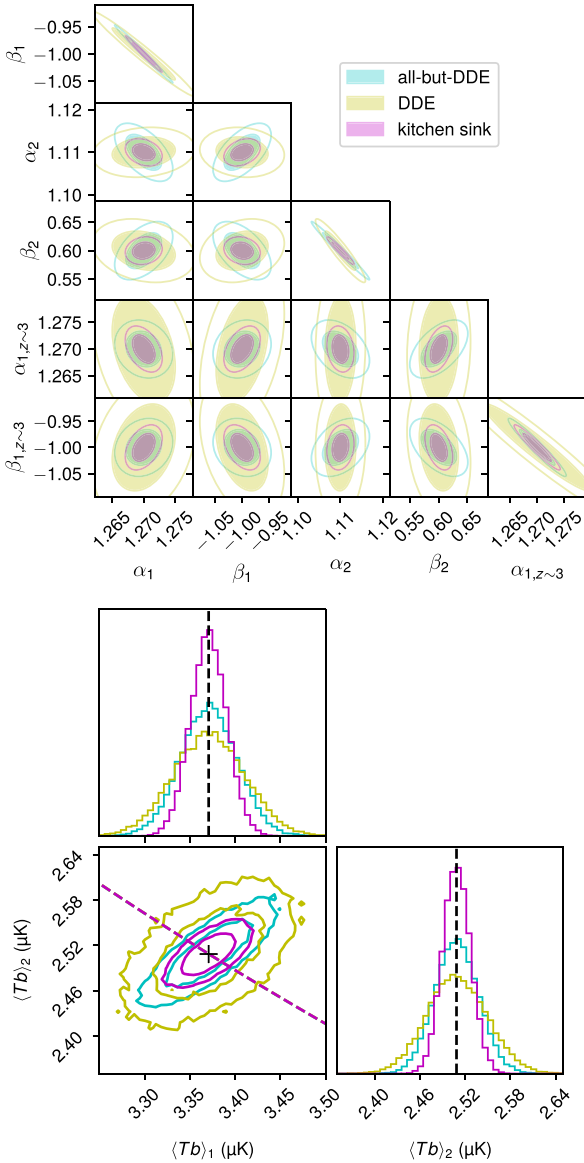
Figure 9. Same as Fig. 7 or Fig. 8, but now comparing the DDE against not just the cross power spectrum or just one-point statistics, but the 'all-but-DDE' combination of all summary statistics from COMAP-ERA data other than the DDE (possibly also combined with the DDE in a 'kitchen sink' scenario).

affect the 'auto' VID far more than the DDE, which should again be significantly more robust (even if not perfectly insensitive) to such factors by design.

4.2.3 The kitchen sink

We now consider throwing everything we have at the problem – all auto- and cross-spectra as well as both VIDs – and examine in Fig. 9 how the DDE might *still* improve constraints.

The combination of every observable other than the DDE finally matches the DDE-based constraint of $\langle Tb\rangle_1\langle Tb\rangle_2 = 8.5 \pm 0.2$ (68 per cent interval). With the DDE and all-but-DDE constraints on similar levels, the DDE contributes non-negligibly to the 'kitchen-sink' constraint using all observables, which is $\langle Tb\rangle_1\langle Tb\rangle_2 = 8.5 \pm 0.1$ (68 per cent interval).

Neither $\{\alpha_J, \beta_J\}$ nor $\langle Tb\rangle_J$ necessarily provide an intuitive picture of physical conditions at reionization, so we may also consider how the parameter posteriors translate to other quantities. The example we will consider briefly here is the ratio of the average CO(2–1) and CO(1–0) line brightness temperatures:

$$\frac{\langle T\rangle_2}{\langle T\rangle_1} = \frac{\nu_{\mathrm{rest},1}^3}{\nu_{\mathrm{rest},2}^3}\frac{\int \mathrm{d}M_h\,(\mathrm{d}n/\mathrm{d}M_h)\,L_{\mathrm{CO},2}(M_h,z)}{\int \mathrm{d}M_h\,(\mathrm{d}n/\mathrm{d}M_h)\,L_{\mathrm{CO},1}(M_h,z)}, \quad (25)$$

with $z = 6.68$ at the midpoints of the COMAP observing frequency bands. The ratio of these temperatures is an approximate global measure of excitation of the CO rotational transitions overall and thus characterizes the temperatures and densities of the interstellar medium hosting the CO gas.

The value of this ratio recovered with the DDE is 0.712 ± 0.010, and while the 'kitchen-sink' constraint of 0.712 ± 0.004 is significantly better, the addition of the DDE does very marginally narrow it from the all-but-DDE constraint of 0.712 ± 0.005. Of course, improvements from the DDE would be more significant if some subset of auto power spectra or individual VID data were considered unreliable.

5 CONCLUSIONS

We have now answered the questions set out by the Introduction of this paper:

(i) *Is the DDE robust to contaminants like noise and interloper emission, as is the analytic expectation?* For appropriately fine choices of binning, we are able to define a region of Fourier-dual temperature space unaffected by noise where the DDE is an unbiased measurement of the correlation between the shapes of the distributions of two correlated variables. While we do not demonstrate perfect insensitivity to interloper emission, our analysis does show relative insensitivity of the DDE to interloper CO(1–0) parameters compared to the equivalent $z \sim 7$ CO parameters.

(ii) *How much could the DDE fundamentally improve constraints on $z \sim 7$ CO emission in simulated COMAP-ERA observations?* The DDE potentially significantly improves constraints on the $z \sim 7$ CO line model, with (for example) constraints on $\langle Tb\rangle_1\langle Tb\rangle_2$ becoming 3–5 times tighter than either the cross power spectrum alone or 'auto' one-point statistics alone. Unlike the cross power spectrum, the DDE actually appears able to separately constrain $\langle Tb\rangle_1$ and $\langle Tb\rangle_2$. Even in comparison to the combination of all other available summary statistics, the DDE contributes additional constraining power to $z \sim 7$ CO model parameters.

Our numerical investigations, while preliminary, make it apparent that the non-Gaussianity of the signals imprint themselves in the DDE just as much as in the VID. Just as the the VID breaks the degeneracy inherent in the auto power spectrum between the mean CO temperature and luminosity-averaged tracer bias (Breysse 2022), the DDE breaks the degeneracy inherent in the cross power spectrum between the intensities of the two lines as we have demonstrated numerically. Future work should provide a non-numerical explanation for this degeneracy breaking.

The present investigation also suggests the DDE can potentially enhance science output from cross-correlations, particularly in scenarios where auto spectra and/or VID data are untrusted. Real-world application of the DDE will still need better understanding of covariance, dependence on signal and noise modelling, and so on. However, at minimum, the clearly different degeneracies from the cross power spectrum and other observables show that the DDE merits further investigation on these fronts.

ACKNOWLEDGEMENTS

DATA AVAILABILITY

The simulated data underlying this article will be shared on reasonable request to the corresponding author.

REFERENCES

Astropy Collaboration 2013, A&A, 558, A33
Astropy Collaboration 2018, AJ, 156, 123
Beane A., Lidz A., 2018, ApJ, 867, 26
Behroozi P. S., Wechsler R. H., Conroy C., 2013a, ApJ, 762, L31
Behroozi P. S., Wechsler R. H., Conroy C., 2013b, ApJ, 770, 57
Bernal J. L., Kovetz E. D., 2022, A&AR, 30, 5
Breysse P. C., 2022, preprint (arXiv:2209.01223)
Breysse P. C., Alexandroff R. M., 2019, MNRAS, 490, 260
Breysse P. C., Kovetz E. D., Behroozi P. S., Dai L., Kamionkowski M., 2017, MNRAS, 467, 2996
Breysse P. C., Anderson C. J., Berger P., 2019, Phys. Rev. Lett., 123, 231105
Breysse P. C. et al., 2022a, ApJ, 933, 188
Breysse P. C., Chung D. T., Ihle H. T., 2022b, preprint (arXiv:2210.14902)
Chung D. T. et al., 2021, ApJ, 923, 188
Chung D. T. et al., 2022, ApJ, 933, 186
Cleary K. A. et al., 2022, ApJ, 933, 182
Foss M. K. et al., 2022, ApJ, 933, 184
Gong Y., Cooray A., Silva M. B., Santos M. G., Lubin P., 2011, ApJ, 728, L46
Gong Y., Cooray A., Silva M., Santos M. G., Bock J., Bradford C. M., Zemcov M., 2012, ApJ, 745, 49
Ihle H. T. et al., 2019, ApJ, 871, 75
Ihle H. T. et al., 2022, ApJ, 933, 185
Kamenetzky J., Rangwala N., Glenn J., Maloney P. R., Conley A., 2016, ApJ, 829, 93
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. D. et al., 2017, preprint (arXiv:1709.09066)
Kovetz E. et al., 2019, Bull. Am. Astron. Soc., 51, 101
Li T. Y., Wechsler R. H., Devaraj K., Church S. E., 2016, ApJ, 817, 169
Lidz A., Furlanetto S. R., Oh S. P., Aguirre J., Chang T.-C., Doré O., Pritchard J. R., 2011, ApJ, 741, 70
Loken C. et al., 2010, J. Phys. Conf. Ser., 256, 012026
Ponce M. et al., 2019, Proceedings of the Practice and Experience in Advanced Research Computing on Rise of the Machines (Learning), PEARC '19. Association for Computing Machinery, New York, NY, USA. Available at: https://doi.org/10.1145/3332186.3332195
Pullen A. R., Chang T.-C., Doré O., Lidz A., 2013, ApJ, 768, 15
Pullen A. R. et al., 2022, preprint (arXiv:2209.02497)
Stein G., Alvarez M. A., Bond J. R., 2019, MNRAS, 483, 2236
Sun G., Hensley B. S., Chang T.-C., Doré O., Serra P., 2019, ApJ, 887, 142
Tinker J., Kravtsov A. V., Klypin A., Abazajian K., Warren M., Yepes G., Gottlöber S., Holz D. E., 2008, ApJ, 688, 709
Tinker J. L., Robertson B. E., Kravtsov A. V., Klypin A., Warren M. S., Yepes G., Gottlöber S., 2010, ApJ, 724, 878
Yang S., Popping G., Somerville R. S., Pullen A. R., Breysse P. C., Maniyar A. S., 2022, ApJ, 929, 140

APPENDIX A: CORRELATION COEFFICIENT MATRIX BETWEEN ALL OBSERVABLES

Fig. A1 shows the correlation coefficient matrix between all observables calculated in this work, across all 2430 simulations of COMAP survey fields. Aside from strong correlations between the different bins of the masked DDE (which do not affect the stability
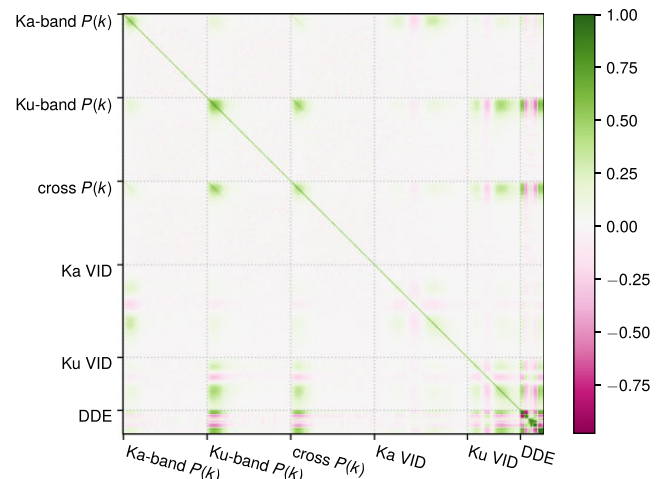

Figure A1. The correlation coefficient matrix between all observables across the 2430 simulated COMAP survey fields used in this work. We calculate and bin all observables as summarized at the end of Section 3 except for the DDE, which needs to be masked as explained in Section 4.2.

of the Fisher forecasts as explained in Section 4.2), we also find correlations between the individual VID bins and the DDE. However, correlations between the Ka-band VID and the DDE are weaker, presumably due to the interloper signal present in the Ka-band data.

Similar points apply for the low-k power spectrum bins, which have correlations with VID bins likely introduced by the high-pass transfer function applied to the simulated data cubes.

APPENDIX B: TABULATED FISHER FORECAST PARAMETER ERRORS

We tabulate the marginalized Fisher forecast errors for each $z \sim 7$ CO–IR power-law parameter in Table B1, based on different sets of summary statistics. For the CO(2–1) parameters α_2 and β_2, the DDE has constraining power equivalent to the combination of every other simulated observable combined.

Table B1. Marginalized Fisher forecast parameter errors for each summary statistic combination considered in Section 4.2.

| Statistics | Marginalized Fisher forecast 1σ error for: | | | | | |
	α_1	β_1	α_2	β_2	$\alpha_{1, z \sim 3}$	$\beta_{1, z \sim 3}$
DDE	0.0033	0.0313	0.0024	0.0225	0.0066	0.0545
$P_\times(k)$	0.16	1.46	0.13	1.19	0.02	0.47
$P_\times(k)$ + DDE	0.0030	0.0282	0.0022	0.0204	0.0056	0.0467
VID + VID	0.0032	0.0350	0.0095	0.0669	0.0113	0.1184
VID+VID + DDE	0.0018	0.0203	0.0018	0.0163	0.0041	0.0396
all-but-DDE	0.0019	0.0215	0.0028	0.0206	0.0022	0.0224
kitchen-sink	0.0012	0.0135	0.0013	0.0098	0.0017	0.0174

This paper has been typeset from a TEX/LATEX file prepared by the author.