HIGH-REDSHIFT ASTROPHYSICS AND COSMOLOGY WITH INTENSITY MAPPING

by

Patrick C. Breysse

A dissertation submitted to Johns Hopkins University in conformity with the requirements for the degree of Doctor of Philosophy.

Baltimore, Maryland

March, 2017

Abstract

Upcoming intensity mapping surveys will provide a powerful probe of astrophysics and cosmology in the distant universe. However, in order to realize the potential of this new technique, we must improve our understanding of the link between intensity maps and physics on all scales. In this thesis, we develop a number of tools for analyzing intensity maps. For most of this work, we illustrate our methods using the example of a CO intensity map at redshift $z \sim 3$, though our results are applicable to many other lines.

We first present the formalism for computing intensity mapping power spectra, and use it to study the detectability of the CO signal. We find that, though the signal amplitude is highly uncertain, a well-designed experiment can detect CO under most assumptions. We then describe a method of simulating intensity maps, and use it to study the problem of foreground line contamination. By masking out the brightest regions of a map, we find that the cosmological information contained in surveys targeting Lyα and CO can be recovered, though astrophysical information is lost. Unfortunately, due to instrumental constraints and high-intensity foreground

ABSTRACT

contamination, this method is less effective for CII surveys. We also demonstrate that the astrophysical information content of foreground lines can be recovered through cross-correlations. By correlating pairs of CO frequency bands, we show that the ^{13}CO isotopologue line can allow unprecedented constraints on molecular gas properties within distant galaxies.

Line emission in intensity maps is highly non-Gaussian, so power spectra alone cannot fully constrain the properties of target populations. To resolve this, we introduce the Voxel Intensity Distribution (VID), which gives the one-point probability distribution of voxel intensities. We demonstrate that the VID provides substantial constraining power beyond what is obtainable from the power spectrum, even in the presence of foreground contamination. We then apply the VID to a hypothetical measurement of the cosmic star formation rate density from a CO survey. We show that, despite model uncertainties, such an observation could place competitive constraints on this crucial quantity while eliminating systematics which hamper existing measurements.

Primary Reader: Marc Kamionkowski

Secondary Reader: Julian Krolik

Acknowledgments

First and formost I need to thank my Ph.D. advisor Marc Kamionkowski for all of his help and support over the past five years. He is a brilliant physicist and a fantastic mentor and I would not be the scientist I am today without his efforts. I am also grateful to my thesis defense committee, Julian Krolik, Toby Marriage, Rigoberto Hernandez, and Eric Switzer for taking the time to provide insightful and constructive comments about this work.

I am indebted to my collaborators over the years, particularly to Ely Kovetz, Mubdi Rahman, and Yacine Ali-Haïmoud, who have passed along their knowledge and experience on topics ranging from astrophysics to numerical methods to job applications. I also need to thank the other members of Marc's group at JHU, including Donghui Jeong, Jens Chluba, Alvise Raccanelli, Simeon Bird, and Illias Cholis for sharing their knowledge over the years.

I am extremely thankful for the companionship of the other graduate students at Hopkins. In particular I need to thank my officemates Alice Cocoros, Dan Pfeffer, Nick Eminizer, and Jon Aguilar for keeping me sane every day, as well as my fellow

group members Julian Muñoz, Tanvi Karwal, and Liang Dai. I also need to thank the myriad members of The Rabble for making Hopkins such a fun place to be. Rachael Alexandroff deserves particular thanks for any number of things, and for keeping me company on the upcoming move to Toronto.

My family has been a constant source of support throughout my life, and in particular as I finish my Ph.D. and move into the academic world. My parents, Pat and Jill Breysse have always had my back, and my father in particular who, as a fellow academic, understands the challenges of an academic career and has been an invaluable source of counsel.

Finally, I could not have accomplished any of this without my wife Heather, who has been my constant partner through college, grad school, and beyond.

Dedication

This thesis is dedicated to my parents, Pat and Jill Breysse

Contents

List of Tables

List of Figures

Chapter 1

Introduction

The past decade has seen an explosion of new data from the distant universe. Observations of the cosmic microwave background (CMB) from space[1,2] and from the ground[3,4,5,6] have provided exquisite images of the structure of the universe in its infancy. Surveys of galaxies, clusters, and quasars[7] have mapped the more nearby universe in unprecedented detail. These observations have led to the dawning of what has come to be known as the "age of precision cosmology", and the formation of a concordance model of cosmic history.

This consensus history takes the form of a model known as ΛCDM, denoting a universe composed primarily of cold dark matter (CDM) and cosmological-constant dark energy (denoted by Λ). Cosmic history in this model begins with a period of inflation, during which space expanded extremely rapidly, leaving behind quantum fluctuations which would become the seeds of large-scale structure in the universe[8].

The early universe consisted of hot plasma, with photons strongly coupled to free electrons driving baryon-acoustic oscillations which left relics still visible today[9]. A few hundred thousand years after the end of inflation (corresponding to redshift $z \sim 1100$), this plasma cooled to the point where electrons and protons combined to form the first atoms, in a period called "recombination". The CMB contains a snapshot of conditions in the universe at this time, with the density fluctuations seeded by inflation seen as hot and cold spots.

After recombination, the universe entered what is known as the dark ages, as small overdensities drew in more baryons and dark matter, which eventually formed bound halos, and the first stars and galaxies. Light from these earliest galaxies ionized the intergalactic medium (IGM) once more, and the dark ages ended as the universe became transparent. These galaxies continued to grow and merge, and to form new stars. The cosmic star formation rate increased continually until a peak at roughly redshift $z \sim 2$, before beginning to decline to the rate we see at the present day.

Our knowledge of the cosmological and astrophysical processes which govern the evolution of the universe has never been greater. However, the cosmic microwave background (CMB) traces only to a two-dimensional surface at $z \sim 1100$, while most galaxy surveys are only sensitive out to $z \sim 1$ or so. At such great distances, even the deepest surveys can only detect the very brightest galaxies, which by definition are likely unrepresentative of the broader population. This leaves a large gap in our understanding of galaxy evolution and the growth of structure, as shown in Figure

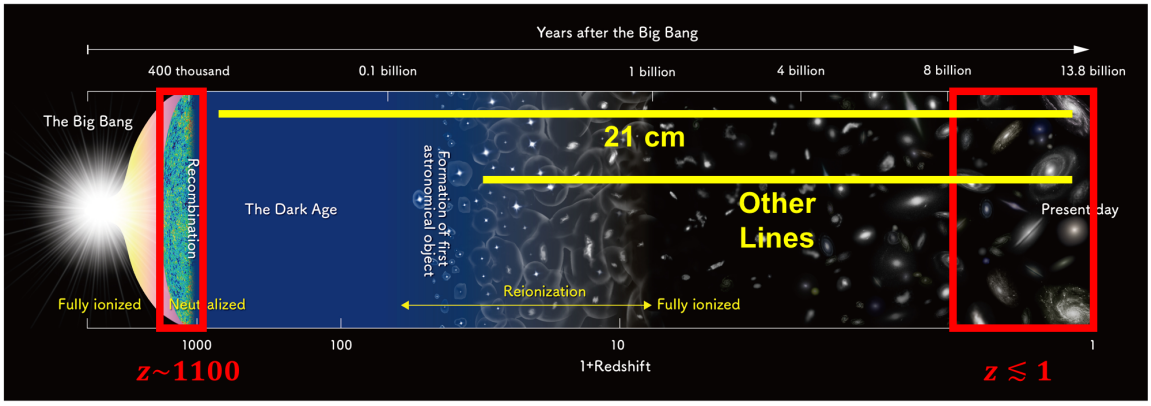

Figure 1.1: Schematic view of the history of the universe. Red frames show the periods observed by the CMB (left) and by galaxy surveys (right). The gap in between includes some of the most important periods in cosmic history, including the birth of the first galaxies during the dark ages, the epoch of reionization, and the growth of galaxies into the forms we see today. Yellow lines show the time periods accessible to intensity mapping surveys targeting the 21 cm line (top) as well as other lines (bottom), including CO, [CII], Lyα, and many others (Original Image Credit: NAOJ).

1.1. The dark ages, reionization, and the peak of star formation are just some of the interesting periods of cosmic history which are not well-sampled by current observations.

Recently, *intensity mapping* has arisen as a powerful probe of this unseen structure. Rather than attempting to study galaxies individually as in typical observations, intensity mapping seeks to probe the intensity fluctuations of a spectral line on spatial scales large compared to any individual source. Such a survey thus makes use of the aggregate emission from all of the galaxies within a target volume, and can make statistical measurements of the entire galaxy population. By studying a single emission line and observing at different frequencies, it is possible to observe these intensity fluctuations in three dimensions, as the observed frequency of a line maps

one-to-one to the emission redshift. This allows the mapping of line emission across cosmic history. Initial intensity mapping efforts focused on the 21 cm transition in neutral hydrogen, but efforts have also begun recently to probe other lines as well, as different lines are sensitive to different astrophysical processes.

The fluctuations observed in an intensity mapping survey depend on the luminosity function and spatial distribution of the source galaxies. The luminosity function depends on the detailed astrophysical conditions within the emitters, such as star formation rates and metallicities. The spatial distribution indirectly traces the underlying dark matter field, the properties of which in turn depend on cosmological parameters. Intensity mapping can thus provide information about a wide variety of cosmological and astrophysical topics. Typically, this information is extracted from a map using its power spectrum, a powerful statistic which has proven valuable for studying both galaxy distributions[10] and the cosmic microwave background[11]. The power spectrum $P(k)$ of a map consisting of brightness temperature T as a function of position $\mathbf{x}$ is defined by

$$\left\langle \widetilde{T}(\mathbf{k})\widetilde{T}^{*}(\mathbf{k}')\right\rangle \equiv (2\pi)^{3}\delta_{D}(\mathbf{k}-\mathbf{k}')P(k) \tag{1.1}$$

where δ_D is the Dirac delta function, the Fourier wavenumber $k = |\mathbf{k}|$, and $\widetilde{T}(\mathbf{k})$ is the Fourier transform of $T(\mathbf{x})$. The delta function enters because the universe is isotropic and homogeneous on large scales, and thus the power spectrum must remain

unchanged under rotations or translations.

The most well-known line used for intensity mapping is the 21 cm spin-flip line from neutral hydrogen (see for example Morales & Wyithe[12] and references therein). Experiments such as the Precision Array for Probing the Epoch of Reionization[13] (PAPER), the Murchison Widefield Array[14] (MWA), the LOw-Frequency ARray[15] (LOFAR), the Hydrogen Epoch of Reionization Array[16] (HERA) and the Square Kilometer Array[17] (SKA) seek to study the epoch of reionization. During this era, large volumes of the IGM remained neutral, and thus can be probed by the 21 cm line. These experiments could provide substantial insight into how ionized bubbles grew and merged, and how they were sourced[18]. Other experiments like the Canadian Hydrogen Intensity Mapping Experiment[19] (CHIME) and the Hydrogen Intensity and Real-time Analysis eXperiment[20] (HIRAX) are attempting to study galaxies around $z \sim 1$. The 21 cm line at these redshifts comes entirely from within individual sources rather than the IGM, and thus can be used to map the spatial distribution of galaxies over large volumes.

There has been a surge of interest recently[21] in conducting surveys in other lines besides 21 cm as a probe of high-redshift astrophysics. The line which we will primarily focus on for this work is the 115 GHz rotational transition of the carbon monoxide (CO) molecule. CO is a powerful tracer of molecular gas, which is a critical part of understanding the interstellar medium (ISM) of galaxies. Molecular gas provides the fuel and composes the environment from which stars form across cosmic time.

Its abundance is directly related to the rate at which stars form[22,23], and its direct measurement provides insight into the efficiency through which this process takes place[24].

The rotational transitions of carbon monoxide molecules have long been some of the most potent tracers of the molecular gas phase of the interstellar medium. H_2, though a much more common molecule, is not easily excited in typical molecular cloud conditions. Reasons for this include the fact that H_2 lacks an electric dipole moment and has a low mass. These properties raise the temperature required to excite the molecule out of the range of all but the most extreme molecular clouds. CO, on the other hand, is easily excited. Combined with millimeter emission wavelengths which fall into a relatively transparent atmospheric window, this has led to a great deal of study of CO and its dynamics[25,26,27,28].

The CO intensity mapping signal was first modeled[29] as a possible CMB foreground, before being recognized as an interesting observable in its own right[30,31,32]. The current experimental efforts in CO intensity mapping are primarily targeted at $2 \lesssim z \lesssim 3$. This is a particularly interesting redshift range as it is near the peak of the cosmic star formation rate, the so-called "Cosmic Noon". The CO Power Spectrum Survey (COPSS), a high-resolution interferometric survey aiming to detect CO fluctuations on relatively small scales, recently published the first tentative detection of the CO signal at these redshifts[33,34]. A second experiment known as the CO Mapping Array Pathfinder (COMAP) is currently under construction[35]. COMAP is a

single-dish experiment which aims for larger scales than COPSS.

Another line which has garnered significant experimental interest is the 158 μm CII fine structure line. This line provides a dominant cooling mechanism of interstellar gas[36], and thus is a useful probe of both the ionized and neutral media in star-forming galaxies[27]. This makes it an intriguing choice for intensity mapping observations[37,38]. CII is the intended target of the Tomographic Intensity Mapping Experiment[39](TIME), which seeks to probe star-forming galaxies during the epoch of reionization. Efforts are also underway to target the 1216 ÅLyman α (Lyα) line at reionization[40]. At these redshifts, the UV Lyα line redshifts into the near-infrared, where it will be observable by the Spectro-Photometer for the History of the Universe, Epoch of Reionization, and Ices Explorer[41] (SPHEREx).

Several other lines have been discussed in the literature as targets for intensity mapping. Lines such as Hα and OII when combined with the lines above could expand our understanding of galaxy emission spectra[42], and will also be accessible to SPHEREx at $z \sim 1 - 4$. Though molecular hydrogen H_2 is not a particularly useful probe of the nearby universe, at $z > 10$ it constitutes a significant coolant of primordial gas[43], and could thus be an interesting future target for intensity mapping[44]. Intensity mapping with He II recombination lines could be used to study Population III stars, which form in the earliest galaxies[45]. These galaxies are quite small and faint, which makes them hard to study using conventional methods.

In this dissertation, we will explore several topics relevant to extracting useful

astrophysical and cosmological information from intensity mapping surveys. We will focus on lines which are predominantly emitted from within discrete sources such as galaxies, neglecting for now the more well-studied case of diffuse IGM 21 cm emission at reionization. For most of our computations, we will consider intensity mapping of the CO line at $z \sim 3$, but our methods are readily generalizable to many other lines. We will start in Chapter 2 with a discussion of the form of the intensity mapping power spectrum, both in three-dimensional spatial coordinates and two-dimensional angular coordinates.

In Chapter 3 we study the feasibility of an IM survey targeting CO at $z \sim 3$. We consider four models for cosmological CO emission from the literature and estimate the strength of the CO signal predicted by each[46]. In doing so, we demonstrate a simple one-parameter family of models which can account for a wide variety of assumptions about CO emission. When we calculate the power spectra using these models, we find that the difference in amplitude between the largest amplitude signal and the smallest is roughly two orders of magnitude, which clearly illustrates the lack of theoretical understanding of star formation physics at high redshifts. When studying the prospects for an experiment to detect this signal, we explore the effect of the instrumental parameters on the predicted signal-to-noise ratio and demonstrate the importance of carefully considering the values of these parameters when designing a survey. In particular, attempting to survey too large of an area of the sky or choosing a spectrograph with insufficient resolution can significantly decrease the chance of

detecting the CO signal.

In Chapter 4, we explore a means of creating simulated intensity maps using log-normal density fields[47]. We find that by assuming a power spectrum and drawing a realization of a corresponding density field we can simulate maps much more quickly than usual N-body methods. In Chapter 5 we use these simulations to study the problem of foreground line confusion in intensity mapping. In any intensity map, there will be a contribution from lines besides the target emitted by galaxies along the line of sight. For some cases, such as Lyα and CII, these foreground lines are virtually certain to be significantly brighter than the signal. For others, such as CO, the foreground lines are still fainter on average, but is possible to model them in such a way that the total power spectrum is dominated by Poisson noise from foreground sources. We show in this chapter that by blindly masking the brightest voxels, or three-dimensional pixels, in a map, it is in some situations possible to recover the shape of the signal power spectrum. Much of a map's astrophysical information is lost in this process, but cosmological information may be recoverable. For the models and survey parameters we use here, this method is effective for CO and Lyα experiments, but fails for CII.

Though foreground lines are a nuisance in many cases, they do carry useful information of their own about conditions where they are emitted. In Chapter 6, we discuss how one such line, the 110 GHz transition of ^{13}CO, can be used to dramatically improve the knowledge obtained from a CO intensity map[48]. CO maps target

the ^{12}CO isotopologue, but always contain emission from the rarer ^{13}CO species. The ^{12}CO emission will always dominate, but by cross-correlating pairs of frequency channels in an intensity map one can obtain a measurement of the fainter ^{13}CO emission. As we demonstrate, this allows one to study both the relative abundance of the two species, which depends on nucleosynthesis processes, and the column density of molecular clouds in which the CO molecules are contained.

Virtually all work on intensity mapping in the literature has focused on the power spectrum as a means of extracting useful information from a map. In Chapter 7, we introduce an alternative statistic which enables much more detailed study of luminosity functions in intensity maps[49]. The Voxel Intensity Distribution, or VID, is the probability distribution of voxel intensities, and is related to the luminosity function by a technique called $P(D)$ analysis. We show that, even in the presence of line and continuum foregrounds, the VID allows a high-quality measurement of the luminosity function. In Chapter 8, we demonstrate how this luminosity function measurement can be used to constrain the cosmic star formation history (SFH)[50], a subject of significant astrophysical interest. Current measurements of the SFH rely mostly on deep observations of bright galaxies, which are then extrapolated to estimate the contribution from faint galaxies. We show that the VID of a CO intensity map can place competitive constraints on the SFH while directly probing the faint galaxies which are undetectable to conventional methods.

Throughout this work, we assume a fiducial ΛCDM cosmology with $(\Omega_m, \Omega_\Lambda, h, \sigma_8, n_s) =$

$[0.27, 0.73, 0.7, 0.8, 0.96]$. For the convenience of the reader, various supporting details can be found in the Appendices.

Chapter 2

Intensity Mapping Power Spectra

As with many cosmological observables, the primary statistic used when studying an intensity map is the power spectrum. In this chapter, we will demonstrate how to compute the power spectrum of a given population of galaxies. We will study first the three-dimensional case in Section 2.1 followed by the two-dimensional case in Section 2.2.

2.1 3D Power Spectrum

Consider a survey which observes the brightness temperature $T(\mathbf{x})$ in a line with rest frequency $\nu_{\rm em}$ at position $\mathbf{x}$ and redshift z. Assume that all of the observed emission comes from galaxies which are small compared to the angular resolution, and that the line width of any individual source is small compared to the frequency

resolution. In this limit, we can assume a halo model[51] and write the power spectrum (Equation (1.1)) in the form

$$P(k,z) = \overline{T}^2(z)\bar{b}^2(z)P_m(k,z) + P_{\rm shot}(z), \tag{2.1}$$

where $\overline{T}(z)$ is the sky-averaged brightness temperature, $\bar{b}$ is the luminosity weighted galaxy bias, $P_m(k,z)$ is the linear power spectrum of the underlying dark matter distribution, and $P_{\rm shot}(z)$ is the contribution from galaxy shot noise (see, for example,[30,46,52]). We have assumed for now that the bias is scale-independent, though this assumption may not be valid on very small[53] or very large[54] scales., though it should be close to scale-independent on the scales observed most near-future experiments.

If we assume a ΛCDM cosmology, we can compute the dark matter power spectrum from a Boltzmann code such as CAMB[55]. To determine the other quantities of interest, we assume a luminosity function $\Phi(L,z)$ for our source galaxies. The intensity $I(\nu_{\rm obs})$ observed at frequency $\nu_{\rm obs}$ can be determined from the cosmological radiative transfer equation[30]

$$I(\nu_{\rm obs}) = \frac{c}{4\pi}\int_0^\infty \epsilon\left[\nu_{\rm obs}(1+z')\right]\frac{dz'}{H(z')(1+z')^4}, \tag{2.2}$$

where $H(z)$ is the Hubble parameter. The average proper volume emissivity $\epsilon(\nu)$ in

the target line is given by

$$\epsilon(\nu, z) = \delta(\nu - \nu_{\rm em})(1+z)^3 \int_0^\infty L\Phi(L, z)dL. \tag{2.3}$$

At redshift $z = \nu_{\rm em}/\nu_{\rm obs} - 1$ this yields an intensity

$$I(z) = \frac{c}{4\pi\nu_{\rm em}H(z)} \int_0^\infty L\Phi(L, z)dL. \tag{2.4}$$

Converting this to a mean brightness temperature yields

$$\overline{T}(z) = \frac{c^3(1+z)^2}{8\pi\nu_{\rm em}^3 k_B H(z)} \int_0^\infty L\Phi(L, z)dL, \tag{2.5}$$

where k_B is Boltzmann's constant.

The shot noise amplitude $P_{\rm shot}$ gives the power that would be measured in the absence of clustering. It is proportional to the mean square of galaxy luminosities. Written in brightness temperature units, it takes the form

$$P_{\rm shot} = \left[\frac{c^3(1+z)^2}{8\pi\nu_{\rm em}^3 k_B H(z)}\right]^2 \int_0^\infty L^2\Phi(L, z)dL. \tag{2.6}$$

If we assume a functional form $b(L)$ for the bias of a galaxy with luminosity L, we

can write the luminosity-averaged bias as

$$\bar{b} = \frac{\int Lb(L)\Phi(L)dL}{\int L\Phi(L)dL}. \tag{2.7}$$

We can then combine Equations (2.5–2.7) with Equation (2.1) to obtain the power spectrum of our target line.

It is common practice[40,35] to assume the luminosity of a galaxy is some function $L(M)$ of its mass, and to replace $\Phi(L)$ with a mass function dn/dM. In this type of model, Equation (2.5) would become

$$\overline{T}(z) = \frac{c^3(1+z)^2}{8\pi\nu_{\rm em}^3 k_B H(z)} \int_0^\infty L(M)\frac{dn}{dM}dM, \tag{2.8}$$

with Equations (2.6) and (2.7) modified similarly. One can also assume some scatter around this mean $L(M)$ relation in an effort to take into account the effects of galaxy properties besides mass which can alter the luminosity distribution[35].

2.2 Angular Power Spectrum

Though most IM experiments plan to analyze maps in terms of the three-dimensional power spectrum, there are some situations in which it may be preferable to consider the two-dimensional angular power spectrum instead. This is particularly useful when comparing emission from sources at very different redshifts, where the same angular

scale on the sky corresponds to very different spatial scales at the two redshifts.

The exact expression for converting a 3D power spectrum $P(k)$ to a 2D spectrum C_ℓ is

$$C_\ell = \frac{2}{\pi} \int k^2 P(k) \left[\int f(r) j_\ell(kr) dr \right]^2 dk, \tag{2.9}$$

where r is the comoving distance, j_ℓ is the spherical Bessel function of the first kind, and $f(r)$ is the selection function which determines the range of distances we are sensitive to[56]. For our purposes, we will always assume that we are computing the power spectrum of a single frequency band of an intensity map, and for simplicity we will take $f(r)$ to be a top hat with a width given by the frequency bandwidth of our instrument.

Equation (2.9) is somewhat time consuming to evaluate exactly, so we will instead compute angular spectra in two different limits and interpolate between them. In the limit where the width δr of the selection function satisfies $\ell \delta r / r \gg 1$, i.e. when the width of the observed shell is large compared to the scale of fluctuations being considered, we will use the well-known Limber approximation[57,58]. In this limit, we can write

$$C_\ell \approx \int \frac{H(z)}{c} \frac{f^2(z)}{r^2(z)} P\left[k = \ell / r(z)\right] dz. \tag{2.10}$$

In the opposite limit, where $\ell \delta r / r \ll 1$, the spatial scales being probed are large compared to the shell width, so we can approximate $f(r)$ as a delta function. The

angular power spectrum in this limit is then

$$C_\ell \approx \frac{2}{\pi} \int k^2 P(k) j_\ell^2 \left[kr(z)\right] dk. \tag{2.11}$$

Some attempts have been made to further simplify Equation (2.11) by assuming $P(k)$ is a power law[59], but these approximations only work if the slope of the power law is less than 2. Since the slopes of $P_m(k)$ and thus $P(k)$ are roughly 3 on our scales of interest, this approximation is invalid in our case.

2.3 Summary

We have presented functional forms for computing 3D and 2D power spectra for a given intensity map. In the upcoming chapters, we will compute these spectra for various line emission models and demonstrate how they can be constrained by different experiments. Note that the computations in this chapter neglect additional power from correlations between galaxies in the same dark matter halo, i.e. the so-called "one-halo term" commonly seen in halo models[51]. For all of the models we will consider below, this term is subdominant to the clustering and shot-noise terms in Equation (2.1). Because the models used to predict IM power spectra are quite uncertain (See Chapter 3), it is not necessary to compute power spectra to the level of accuracy which would necessitate including the one-halo term.

Chapter 3

CO Intensity Mapping at Moderate Redshifts

In this chapter we will consider the detectability of the CO intensity mapping signal at our fiducial redshift of $z \sim 3$, expanding on previous work by Pullen et al.[31] (denoted as P13 for the remainder of this chapter). Since this is the redshift where the cosmic star formation rate is expected to be highest[60], this is a good place to attempt a first detection of this signal. Once the techniques for measuring and analyzing this signal are demonstrated successfully at this moderate redshift, surveys could be performed at redshifts corresponding to other periods, such as the epoch of reionization[30,61].

This chapter is primarily based on work published in Ref.[46]. In Section 3.1 below we will summarize four models from the literature and use them to estimate the

power spectrum of CO fluctuations. In order to do this, we parameterize the CO luminosity of a halo as a linear function of its mass. In Section 3.2 we will estimate the signal-to-noise ratios that would be obtained with an optimal survey under the assumptions of these models. We also study how this signal-to-noise ratio varies with various aspects of survey design and find that the detectability of the signal depends strongly on survey parameters.

3.1 Modeling CO Emission

Our goal is to get a rough estimate of the range of possible CO power spectra. To that end, we will assume a one-parameter model for the CO luminosity $L_{\rm CO}$ of a halo with mass M of the form

$$\frac{L_{\rm CO}}{L_\odot} = A\frac{M}{M_\odot}. \tag{3.1}$$

We then need a model of CO emission which can allow us to calculate the parameter A. Below, we consider four such models, then compute angular CO power spectra.

3.1.1 Estimating A

The first and simplest model we consider is proposed by Visbal & Loeb[62], hereafter referred to as VL10. They estimate the star-formation rate for a halo of mass M by assuming that a fraction $f_\star = 0.1$ of the baryons in a halo form stars at a constant rate over a time period $t_s \approx 10^8$ years, where $f_{\rm duty}$ is the ratio of this time to the

Hubble time at redshift $z_{\rm CO}$. This gives a star formation rate (SFR) S of

$$S = \frac{f_\star}{t_s}\frac{\Omega_b}{\Omega_m}M. \tag{3.2}$$

With this relation, we can determine the CO luminosity of a halo if we have a relationship between SFR and $L_{\rm CO}$. In this model, this relation is obtained by measuring the ratio of CO luminosity to SFR from M82, which is observed to be $3.7 \times 10^3\ L_\odot/(M_\odot/{\rm yr})$[63]. Combining these scaling relations allows us to set the value of the parameter A from Equation (3.1). The value in this model is

$$A_{\rm VL10} = 6.24 \times 10^{-7}. \tag{3.3}$$

This value can be used with Equations (2.5-2.7) to determine the average brightness temperature. It can easily be seen from this equation that in this model, the amplitude of both the clustering and shot noise components of the power spectrum will be proportional to the square of the parameter A.

The second model we consider is Model A from P13. The CO luminosity function is calculated similarly to the VL10 model, but instead of using the M82 normalization, they use a set of empirical scaling relations described by Carilli[64]. They first relate CO luminosity to FIR luminosity, then FIR luminosity to SFR, and then SFR to halo mass in a similar manner to VL10. This gives a luminosity function which is

still linear in mass, but with a different normalization

$$A_{\rm P13A} = 2 \times 10^{-6}. \tag{3.4}$$

Model B from P13 uses a slightly different method to calculate the CO brightness temperature. Instead of trying to calculate a luminosity function, this model assumes that the SFR follows the Schechter function[65]:

$$\Phi(S)dS = \phi_\star \left(\frac{S}{S_\star}\right) \exp\left(-\frac{S}{S_\star}\right) \frac{dS}{S_\star}, \tag{3.5}$$

where $\phi_\star$ is a characteristic density and $S_\star$ is a characteristic star formation rate. Integrating this function gives the cosmic SFR density, which can then be combined with the same SFR-CO luminosity scaling relation used in model P13A to get an estimate of the CO volume emissivity. This emissivity can then be entered into Equation (2.2) to get the CO brightness temperature. Though the calculation of $\overline{T}_{\rm CO}$ in this model is somewhat more involved than the one described in Equations (2.2)-(2.5), we can get a reasonable estimate of the brightness temperature in model P13B simply by adjusting the value of A. P13 state that $\overline{T}_{\rm CO}$ in model B is roughly a factor of 4.8 higher than in model A at $z \sim 3$, so the brightness temperature in model P13B can be approximated by choosing

$$A_{\rm P13B} = 9.6 \times 10^{-6}. \tag{3.6}$$

It should be noted that a full computation of the power spectrum using this model should predict a ratio of shot noise to clustering power which is different than the P13A and VL10 models. However, this effect will be small compared to the differences in overall amplitudes between our models. We will therefore continue for now using the approximation that the amount of power in this model on all scales depends only on the value of A. This assumption will be relaxed in subsequent chapters.

The final model we look at here is proposed by Righi et al.[29], which we will refer to as R08. Instead of just assuming that some fraction of halos are forming stars at any given time, the R08 model assumes that star forming episodes happen following major merger events. They estimate that the mass $M_\star$ of stars formed when two halos of mass M_1 and M_2 merge into a halo of mass M is

$$M_\star = 4\frac{\Omega_b}{\Omega_m} f_\star \frac{M_1 M_2}{M}. \tag{3.7}$$

From there, one can calculate the merger rates for halos of a given mass and integrate over all possible masses to determine the total star formation rate. The authors then use the same M82 normalization from the VL10 model to calculate the CO luminosity. As with model P13B, the full calculation of $\overline{T}_{\rm CO}$ in R08 is more complicated than what we have shown here thus far, but we can obtain a good approximation of their result by choosing the correct value for A. In this case the necessary value is approximately

twice the one used for model P13A

$$A_{\rm R08} = 4.0 \times 10^{-6}. \tag{3.8}$$

Note that we have again assumed that the shot noise to clustering ratio for this model is the same as models P13 and VL10.

3.1.2 CO Power Spectra

With values for A in hand, we can compute power spectra for these models. The mean brightness temperature, shot noise amplitude, and bias take the forms

$$\overline{T}(z) = A\frac{c^3(1+z)^2}{8\pi\nu_{\rm em}^3 k_B H(z)} f_{\rm duty}(z) \int_0^\infty M\frac{dn}{dM}dM, \tag{3.9}$$

$$P_{\rm shot} = A^2\left[\frac{c^3(1+z)^2}{8\pi\nu_{\rm em}^3 k_B H(z)}\right]^2 f_{\rm duty}(z) \int_0^\infty L^2\Phi(L,z)dL, \tag{3.10}$$

and

$$b(z) = \frac{\int_{M_{\rm min}}^\infty M\frac{dn}{dM}b(M,z)dM}{\int_{M_{\rm min}}^\infty M\frac{dn}{dM}dM}, \tag{3.11}$$

where we have assumed that only halos with masses greater than $M_{\rm min}$ emit CO. We choose a fiducial value for $M_{\rm min}$ of 10^9 $M_\odot$. For the mass-dependent bias, $b(M,z)$, we use

$$b(M,z) = 1 + \frac{\nu(M,z) - 1}{\delta_c}, \tag{3.12}$$

where $\delta_c = 1.69$, $\nu(M,z) = \delta_c/\sigma(M,z)$, and $\sigma(M,z)$ is the RMS density fluctuation in a spherical region containing mass M[66,67].

It is clear that there is a large amount of theoretical uncertainty regarding the amplitude of the expected CO signal. This justifies our decision to simplify our estimation of $\overline{T}_{\rm CO}$ in models P13B and R08, since the differences between our approximations and the full calculations will be considerably smaller than the differences between the models. For the remainder of this paper we will consider only model P13A at $z = 3$ unless stated otherwise. However, the reader should bear in mind the broad range of theoretical possibilities when following the rest of our results.

For comparison to P13 we will consider intensity maps in terms of their angular power spectra. Figure 3.1 shows the clustering (solid lines) and shot noise (dashed lines) terms of the angular power spectrum for our fiducial model P13A for instruments with different bandwidths. These power spectra are calculated at $z = 3$. As the bandwidth increases, the amplitude of the signal falls off sharply. This makes it very difficult to find this signal in WMAP or Planck data since those instruments had very wide frequency bands. This is likely why the attempt in P13 to find the CO signal in WMAP data was unsuccessful. Also note that at small bandwidths, the clustering term approaches a maximum amplitude. This is due to the use of Equation (2.11) at low ℓ's. If only the Limber approximation was used, the clustering term would continue to increase beyond the limit seen here. Figure 3.2 shows this effect for the two narrowest bandwidths from Figure 3.1.

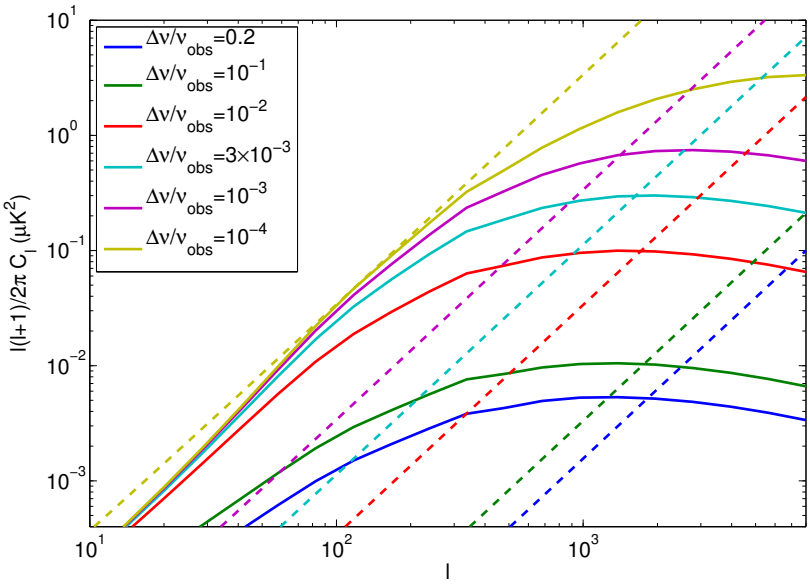

Figure 3.1: Clustering term of the angular CO power spectrum at $z_0 = 3$ for instruments with different frequency bandwidths. The solid lines show the contribution from the clustering term and the dashed lines show the contributions from shot noise. The clustering term loses its dependence on frequency bandwidth when the width of the spatial shell being observed becomes smaller than the size of the features being probed at a given ℓ value.

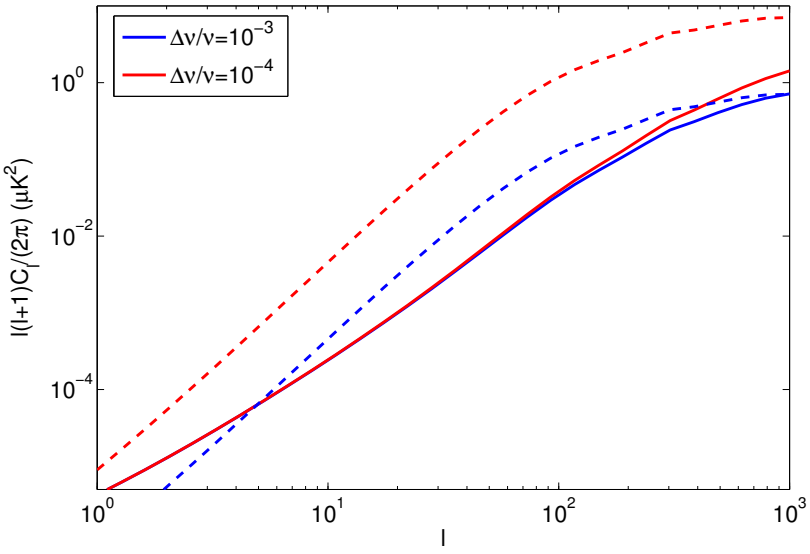

Figure 3.2: Comparison of full power spectrum (solid) with Limber approximation (dashed) for two narrow frequency bandwidths. The Limber approximation fails in this case because the width of the redshift shell being probed is small compared to the spatial size of the brightness fluctuations.

3.2 Dependence on Survey Design

We now introduce a sample instrument to measure the CO signal and estimate the signal-to-noise ratio it would produce. We consider a spectrograph with a 1 GHz total bandwidth targeted at $z = 3$. This bandwidth is split into 35 channels, giving a spectral resolution $R = 1000$. We assume an observation time $t_{\rm obs} = 1$ yr and choose the detector sensitivity s and beam size $\theta_{\rm fwhm}$ to be 800 $\mu{\rm K}\sqrt{\rm s}$ and 10 arcminutes respectively. These values are comparable to those for experiments currently under consideration[68,31]. For a survey covering a solid angle Ω_s, the instrumental noise can be modeled[69] as a random field on the sky with a power spectrum

$$C_\ell^n = \frac{s^2 \Omega_s}{t_{\rm obs} B_\ell^2}, \tag{3.13}$$

where B_ℓ is the beam profile, typically approximated as a Gaussian,

$$B_\ell = e^{-\theta_{\rm fwhm}^2 \ell(\ell+1)/(16 \ln 2)}. \tag{3.14}$$

The angular power spectrum can be measured in each channel of this spectrograph, after which the signals can be stacked to increase the overall SNR. The SNR for such an instrument with N_{ch} channels is given by[70]

$${\rm SNR}^2 = N_{ch} \sum_\ell \frac{C_\ell^2}{\sigma_\ell^2}. \tag{3.15}$$

where

$$\sigma_\ell = \sqrt{\frac{8\pi}{\Omega_s(2\ell+1)}} C_\ell^n. \tag{3.16}$$

Combining Equations (3.15) and (3.16) with Equation (3.13) yields

$$SNR^2 = \frac{N_{ch}}{8\pi}\left(\frac{A}{A_{\rm P13}}\right)^4 \frac{t_{\rm obs}^2}{s^4\Omega_s}\sum_\ell (2\ell+1)\left(C_\ell^{\rm P13A}\right)^2 e^{\theta_{\rm fwhm}^2 \ell(\ell+1)/(8\ln 2)}, \tag{3.17}$$

where $C_\ell^{\rm P13A}$ is the angular power spectrum in the P13A model.

This expression is only valid if one is testing against the null hypothesis, i.e. if one is only interested in seeing if this signal exists at all. It is useful for a first detection attempt, but when trying to obtain useful cosmological information from a signal it is necessary to include an extra cosmic variance term in Equation (3.16) to account for the limited number of modes available in the survey. The signal variance in this case is given by

$$\sigma_\ell = \sqrt{\frac{8\pi}{\Omega_s(2\ell+1)}}\left(C_\ell + C_\ell^n\right). \tag{3.18}$$

Under the null hypothesis, Equation (3.17) clearly shows that a smaller, higher resolution survey will always give a higher SNR. If cosmic variance is included, surveys which are too small will yield smaller SNR's because they include fewer modes. These behaviors can clearly be seen in Figure 3.3, which shows how SNR depends on survey area and beam size, and Figure 3.4, which shows how SNR depends on area for our fiducial 10 arcmin beam. We choose our fiducial value of the survey area to

be $\Omega_s = 4$ deg^2 because this is the value which maximizes the SNR with cosmic variance included. This is a much smaller area than what was chosen for a similar spectrograph suggested in P13, but is comparable to the areas suggested by more recent proposals[35].

Given the wide variation in the signals predicted by the four models we discussed above, we predict a wide range of possible SNR. Figure 3.5 shows the SNR as a function of the parameter A, with the values for the four models discussed above marked by dashed red lines. The curve for the null hypothesis is simply a power law since $SNR \propto A^2$ when cosmic variance is neglected. Possible values of SNR range from 3.2 (VL10) to 760 (P13B) under the null hypothesis and from 2.8 (VL10) to 37 (P13B) including cosmic variance. All of the models except the most pessimistic have a good chance to detect the signal.

The predicted SNR is also sensitive to other theoretical parameters in the models such as $M_{\rm min}$ and t_*. These parameters are poorly constrained, so there is some additional uncertainty beyond that shown in Figure 3.5. For example, as seen in Figure 1 of P13, if $M_{\rm min}$ is increased to 10^{10} $M_\odot$ then the average brightness temperature falls off by a factor of ~ 2. This would in turn decrease the null hypothesis SNR by a factor of 4. Changing the value of t_* would change $f_{\rm duty}$ and have a similar effect on the predicted SNR.

As shown in Figure 3.3 above, the parameters of an intensity mapping survey must be carefully chosen to maximize the chance to detect the signal. Here we explore the

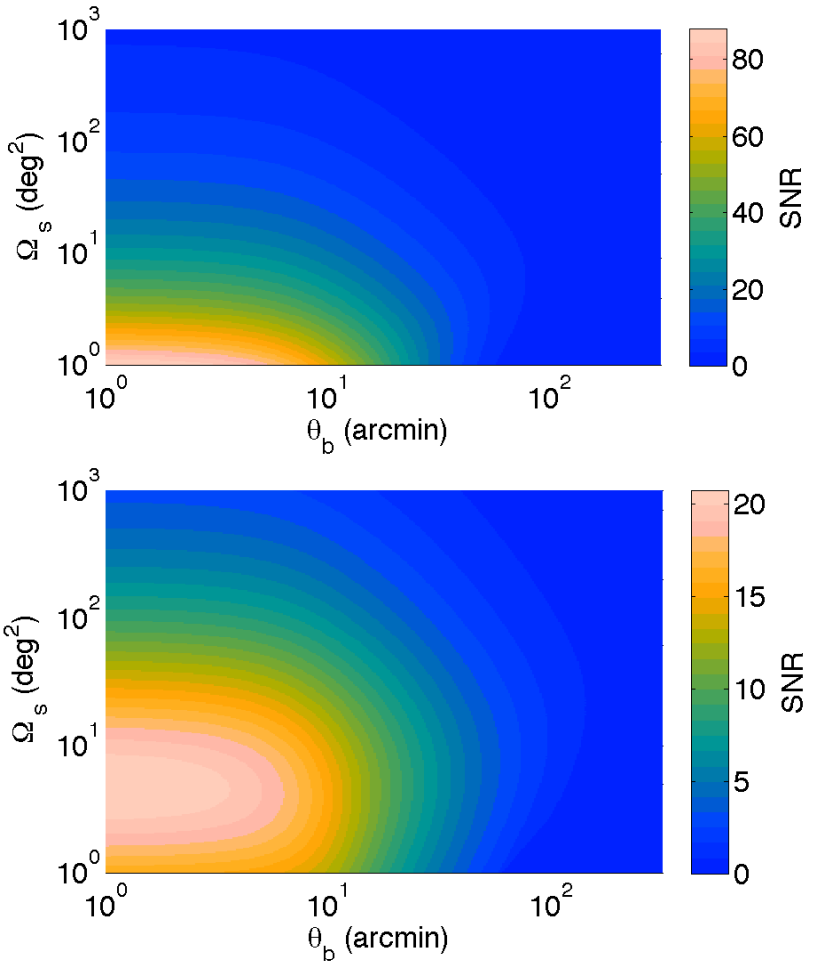

Figure 3.3: SNR as a function of beam size and survey area for our hypothetical spectrograph in the null hypothesis (top panel) and including cosmic variance (bottom panel). A smaller, higher-resolution survey will always improve the chances of a simple detection, but surveys which are too small lose cosmological information because they include fewer modes.

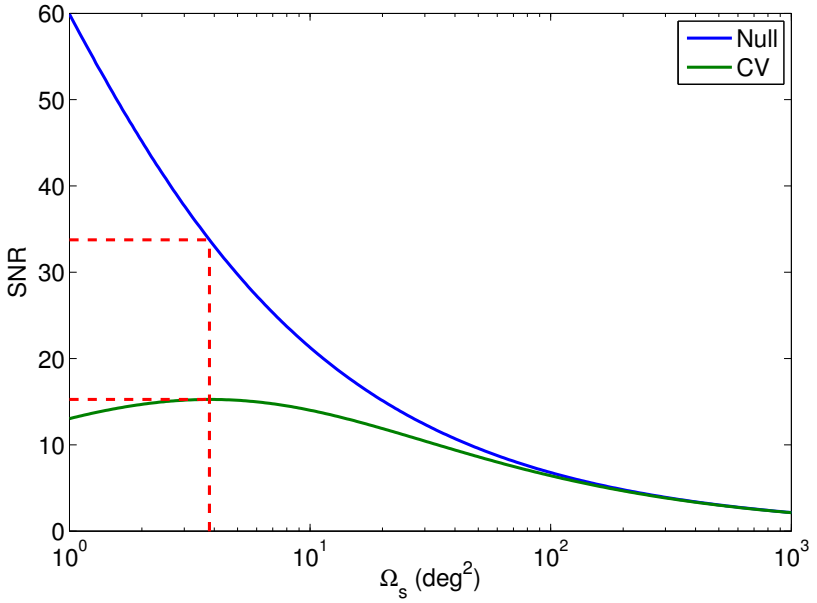

Figure 3.4: SNR as a function of survey area with a beam size of 10 arcmin with and without cosmic variance. The same behaviors seen in Figure 3.3 are visible here as well.

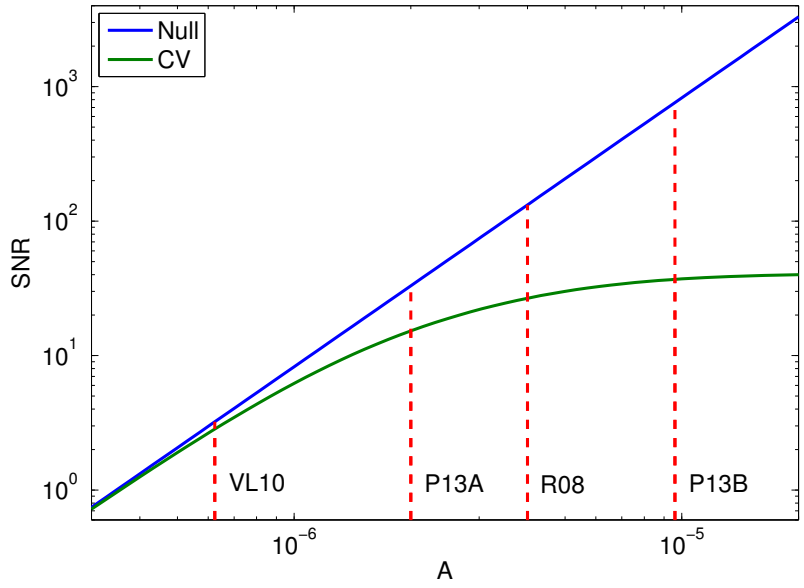

Figure 3.5: Signal-to-noise ratio as a function of parameter A for our hypothetical spectrograph with and without cosmic variance. Values for the four models discussed above are marked.

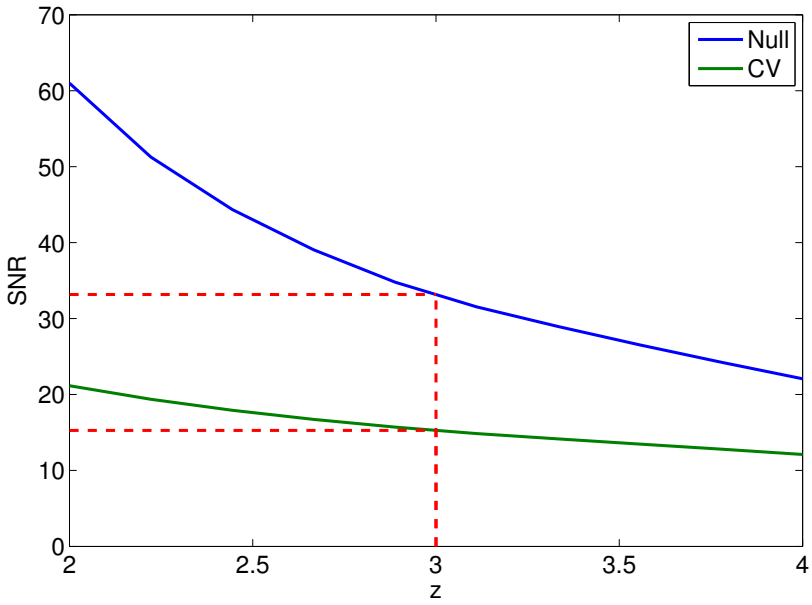

Figure 3.6: Signal-to-noise ratio as a function of central redshift for our hypothetical spectrograph with and without the null hypothesis, with the value at $z = 3$ marked. The total bandwidth of the spectrograph is held constant at 1 GHz while the central frequency is varied.

dependence of SNR on some of the other survey parameters. The first possibility we consider is altering the redshift targeted by the survey to see if $z = 3$ is actually the best redshift to target. Figure 3.6 shows the SNR as a function of the central redshift of the survey for our optimal spectrograph. It is clear that SNR can be increased by targeting lower redshifts. However, when cosmic variance is included the changes are relatively minor, so the target redshift can be altered somewhat without significantly affecting the SNR.

Figure 3.7 shows the effect of varying the number of frequency channels in the spectrographs. The total 1 GHz bandwidth is held constant, so if N_{ch} is increased the width of an individual channel is decreased. The shapes of the curves in Figure 3.7

are due to several factors. For small values of N_{ch}, increasing the number of channels increases the amplitude of the clustering power spectrum as seen in Figure 3.1. As N_{ch} gets larger though, this effect lesses as the clustering power spectrum approaches a constant value. In addition, the telescope sensitivity s is proportional to[68] $\Delta\nu^{-1/2}$, so the amplitude of the noise power spectrum increases for spectrographs with smaller channels, causing the decrease in SNR seen in the null hypothesis term and the slowed increase seen in the cosmic variance term. Finally, for very small channels the shot noise power spectrum overtakes the clustering spectrum causing the signal amplitude to increase again, slowing the decrease in SNR in the null hypothesis term. It can be seen from Figure 3.7 that a lower resolution spectrograph produces a higher null hypothesis SNR. However, this increase is minor and with it comes a decrease in cosmic variance SNR. Thus it may be preferable to use a spectrograph with N_{ch} near our chosen value. Note that we have assumed arbitrarily-narrow line widths. At some point, an experiment with high enough frequency resolution will resolve the structure of emission lines, smoothing out small-scale power along the line of sight.

Another obvious way to improve the SNR of a survey is simply to increase the total observation time. How much SNR can be gained by observing for longer periods is less obvious. Longer observing times decrease the amplitude of the noise power spectrum, but eventually the C_ℓ cosmic variance term in Equation (3.18) starts to dominate over the C_ℓ^n term. At this point, it is more useful to survey a larger area of sky rather than spend additional time on an already deeply studied patch. This effect is illustrated in

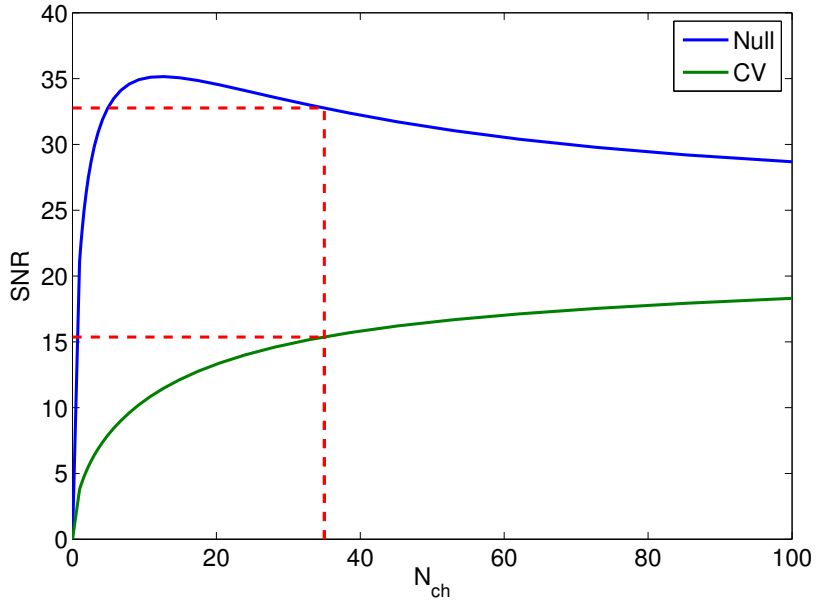

Figure 3.7: Signal to noise as a function of number of spectrograph channels with and without cosmic variance. The fiducial value $N_{ch} = 35$ is marked. The SNR shown here is obtained by stacking the signals from each of the N_{ch} channels.

Figure 3.8, which shows the SNR as a function of survey area and observing time for our two spectrographs, assuming the values calculated above are for a 1 year survey. For longer observations, the maximum SNR is obtained by surveying a larger area of the sky.

3.3 Summary

We have presented a study based on several models of CO emission with the goal of designing an optimal survey aimed at detecting it. We briefly discussed four models which estimated the intensity of CO emission using slightly different methods and we found that the large theoretical uncertainties in the calculation lead to a broad

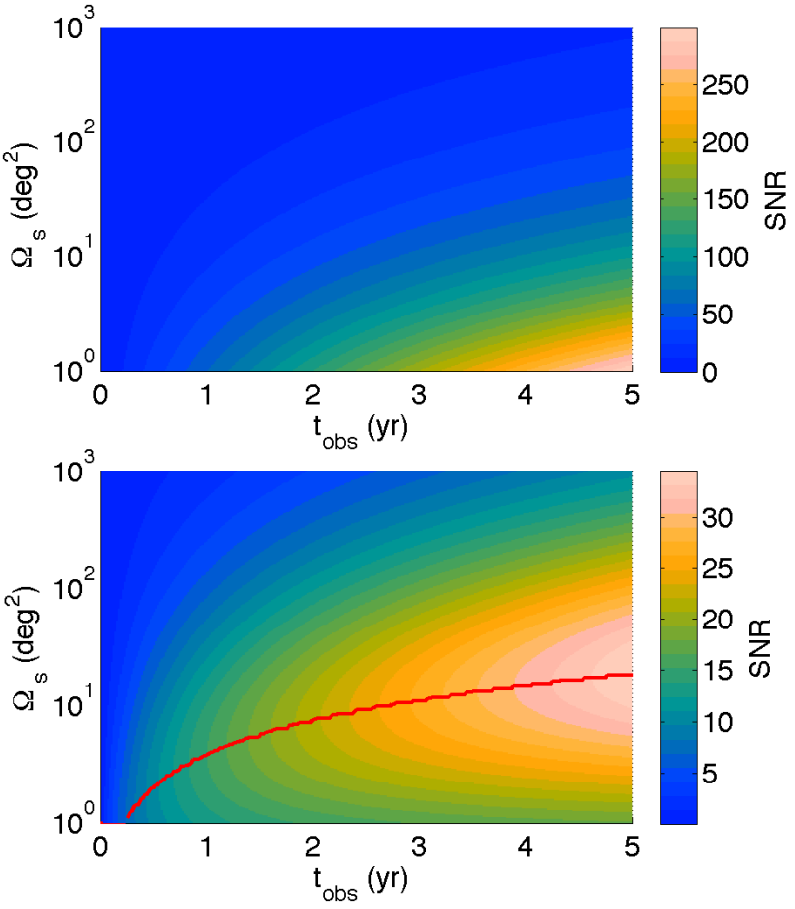

Figure 3.8: Signal to noise as a function of survey time and area for the null hypothesis (top panel) and including cosmic variance (bottom panel). The red line in the bottom panel shows the optimal survey area for a given observing time assuming a 10 arcminute beam.

range of possible values. When calculating signal-to-noise ratios for a representative of these models, model P13A, we found that the optimal survey to detect this signal is one which deeply surveys a relatively small portion of the sky. We found that the exact target redshift is not especially important, and that an instrument with a higher spectral resolution can gain a slight increase in SNR. The instruments we describe here are able to attain a reasonable SNR given model P13A, and they could provide a much stronger detection considering either model R08 or P13B. Model VL10 is much more pessimistic, but since it is the most simplistic of the models (relying only on one galaxy for normalization), it appears less likely than the others.

It is important to note that all of our calculations in this chapter have taken into account only instrumental noise and cosmic variance. We have not included any estimates of the impact of foregrounds on the signal-to-noise ratios above. Since we are looking at line emission, foregrounds with continuous spectra should be fairly easy to remove[68]. However, it is possible for other lines besides the CO line we want to study to be redshifted into the same frequency range. This line confusion could be mitigated by cross-correlating the CO signal with another map of the same area, either another intensity map in a different frequency or a more traditional map of galaxies or quasars. We will discuss the effect of this line confusion in Chapter 5.

The results of this work suggest that it is feasible to design a survey to detect the CO auto power spectrum if foregrounds are not a major concern. However, we have shown that the parameters of such a survey should be considered carefully in order

to maximize the chance of detection.

Chapter 4

Simulating Intensity Maps

In this chapter we will describe a method for creating simulated intensity maps. Past efforts[68,35] to simulate intensity maps used galaxy distributions generated from N-body dark matter simulations. However, N-body simulations are numerically intensive, and it is therefore difficult to generate a large sample of simulated maps. We will therefore employ an alternative method, in which we assume the matter power spectrum and luminosity function of a galaxy distribution is known *a priori*. This allows us to quickly simulate large volumes. We use matter power spectra generated using CAMB[55]. These simulations will then be used in Chapter 5 to demonstrate methods of foreground line cleaning, and in Chapter 7 to test computations of one-point statistics of intensity maps.

This chapter and the next are primarily based on work published in Ref.[47]. We first describe in Section 4.1 a method for simulating maps in which galaxy positions

are uncorrelated. The power spectra of these maps will only contain Poisson noise, so we will refer to these as "shot-noise" simulations. In Section 4.2 we will then show how to add realistic galaxy clustering to these simulations. We will illustrate these methods using a hypothetical CO survey, modeled using Model A of Pullen et al.[31], though simulations can be generated for many different lines.

4.1 Shot Noise Simulations

When simulating a map, we first define a grid of voxels, with the solid angle of each voxel defined by the angular resolution of a hypothetical instrument, the voxel depth by the spectral resolution, and the total area defined by the proposed survey area. For our example simulations, we use a fiducial CO survey similar to that described in Section 3.2 observing at 30 GHz (corresponding to CO at $z = 3$) covering 550 sq. deg. with an angular resolution of 10 arcmin and $\delta\nu/\nu = 10^{-3}$. Since each voxel in this simulated map is large compared to a galaxy, we can assume that the line emitters are essentially point sources. The expected mean number of sources in a given voxel is

$$\langle N \rangle (z) = \overline{n} V_{\rm pix} = V_{\rm vox}(z) \int_0^\infty \Phi(L) dL = V_{\rm vox} f_{\rm duty} \int L(M) \frac{dn}{dM} dM, \qquad (4.1)$$

where $\overline{n}$ is the total halo number density and $V_{\rm vox}$ is the comoving volume of a single voxel.

If clustering is neglected, galaxies will be Poisson distributed on the sky. Thus

we can draw the number of sources in each voxel from a Poisson distribution with mean $\langle N \rangle$. Assuming the line luminosities of individual sources are uncorrelated, we can then randomly assign a luminosity to each source by drawing from a probability distribution

$$\mathcal{P}(L) = \frac{1}{\overline{n}} \Phi(L). \tag{4.2}$$

If we model luminosities through some mass-luminosity relation $L(M)$, we instead draw a mass from the distribution

$$\mathcal{P}(M) = \frac{1}{\overline{n}} \frac{dn}{dM}, \tag{4.3}$$

then calculate the line luminosity from $L(M)$ for each galaxy. Once we have luminosities for each galaxy in a voxel, we can compute a brightness temperature using

$$T = \frac{X_{LT}}{V_{\rm vox}} \sum_{i=1}^{N} L_i, \tag{4.4}$$

where L_i are the luminosities of the N galaxies contained within $V_{\rm vox}$, and we have defined

$$X_{\rm LT} \equiv \frac{c^3 (1+z)^2}{8\pi k_B \nu^2 H(z)}, \tag{4.5}$$

with the goal of keeping our notation compact (Compare Equations (2.5) and (4.4)).

Once we have a finished map, we then need to compute its power spectrum. Since we are simulating relatively small regions of the sky, we can perform this calculation

in the flat sky approximation[71]. For simplicity, we assume the survey area is a square located near the equator of whatever sky coordinates are in use.

Consider a map with $N_{\rm pix}$ voxels at positions $\mathbf{x} = (x_i, y_i)$, where x_i and y_i run from 1 to $\sqrt{N_{\rm pix}}$. The intensity at each voxel $T(\mathbf{x})$ can be decomposed into Fourier modes $a_{\mathbf{k}}$ through the discrete Fourier transform,

$$a_{\mathbf{k}} = \frac{1}{N_{\rm pix}} \sum_{\mathbf{x}} T(\mathbf{x}) e^{2\pi i \mathbf{x}\cdot\mathbf{k}/N_{\rm pix}}. \tag{4.6}$$

The angular power spectrum of the map is then

$$C_{\ell=2\pi k/\sqrt{\Omega_s}} = \Omega_s C_k, \tag{4.7}$$

where Ω_s is the total solid angle of the survey and

$$C_k = \left\langle |a_{\mathbf{k}}|^2 \right\rangle. \tag{4.8}$$

The average in Equation (4.8) is taken over all k modes in the interval $k - 1/2 \leq |\mathbf{k}| < k + 1/2$. Note that in this calculation $\mathbf{x}$ and $\mathbf{k}$ are in units of voxels and (voxels)$^{-1}$ respectively.

The left hand column of Figure 4.1 shows a map of CO emission in Model P13B at $z = 3$ from Section 3.1 simulated using the above technique along with its power spectrum. The power spectrum has the expected scale-independent form of a shot-noise

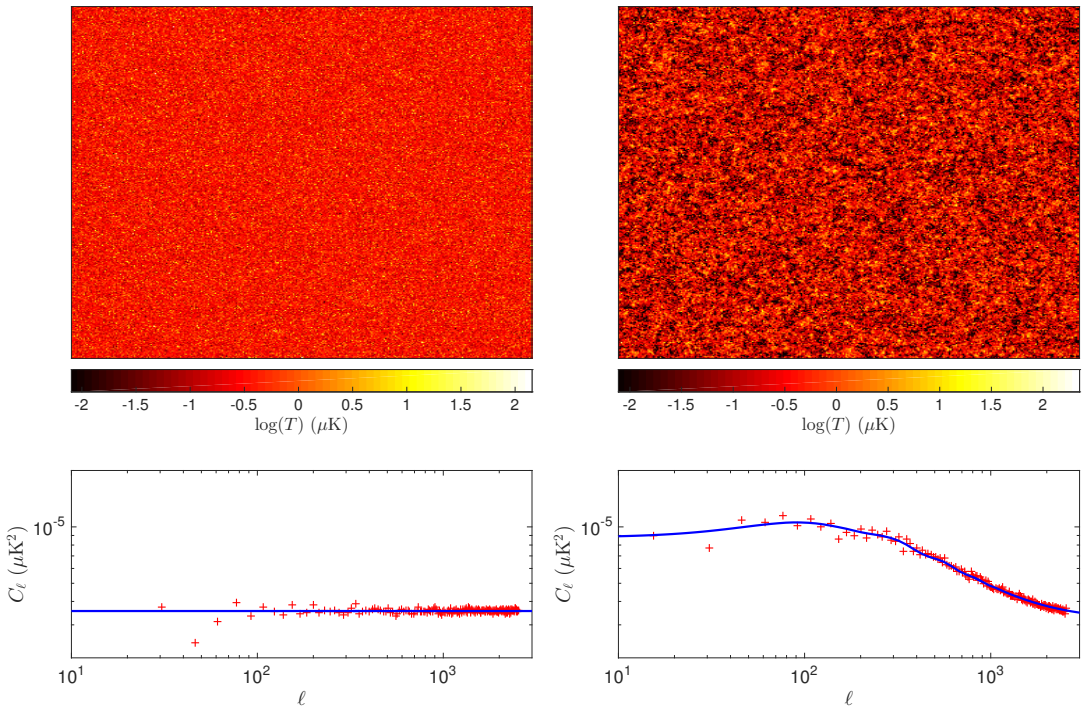

Figure 4.1: (Top left) Simulated shot-noise CO intensity map from $z = 3$ using Model P13B from Section 3.1. (Bottom left) Power spectrum of simulated shot-noise map (red points) overplotted with the predicted shot noise power spectrum. (Top right) Simulated CO intensity map from $z = 3$ with clustering included. (Bottom right) Power spectrum of simulated clustering map overplotted with the predicted power spectrum including both clustering and shot noise components.

dominated sample. The solid line is the predicted shot noise amplitude calculated using the method described in Chapter 2 for this model, included to show the consistency between our model and our simulations. The scatter of the points around this line is sample variance error due to the finite size of our simulation. There is no instrumental noise included in the simulated map.

4.2 Adding Clustering

The simulations above assumed that galaxies are randomly distributed on the sky, but in reality the large scale structure of the universe imposes a clustering pattern on the galaxy distribution. Since we have predictions for the angular power spectra of large-scale clustering, we can generate galaxy density fluctuations by drawing random fields with the desired statistics. This allows us to create galaxy populations without simulating their entire history.

We want to simulate a galaxy density field $\delta(\mathbf{x}) = [n(\mathbf{x}) - \bar{n}]/\bar{n}$ which includes clustering. The δ map should have zero mean and its power spectrum should be the halo power spectrum. In a pixel located at $\mathbf{x}$, instead of drawing the number of galaxies from a Poisson distribution with mean $\langle N \rangle$, we use $\langle N \rangle [\delta(\mathbf{x}) + 1]$. We draw luminosities in the same manner as before. This will give a map with a power spectrum that contains both clustering and shot noise components. Note that we have neglected here the luminosity-dependence of galaxy bias, where brighter galaxies are expected to be more strongly clustered. See Chapter 7 for a more detailed discussion of the effect of luminosity-dependent bias on our calculations.

It only remains then to generate a density field with the desired halo power spectrum. If the field is Gaussian, the process is fairly straightforward. In a full-sky Gaussian random field, each spherical harmonic coefficient $a_{\ell m}$ is drawn from a Gaussian distribution with variance $\sqrt{C_\ell}$. In the flat sky approximation, we draw each Fourier mode with magnitude in the range $k - 1/2 \leq |\mathbf{k}| < k + 1/2$ randomly from

a Gaussian with variance $\sqrt{C_k}$. This map can then be converted back to spatial coordinates using the inverse discrete Fourier transform. One important thing to note is that the resulting map must be real, while the Fourier modes are complex. This means that we must impose the condition that $a_{-\mathbf{k}} = a^*_{\mathbf{k}}$.

Unfortunately, at the redshifts we are considering, the galaxy distribution is highly non-Gaussian. Attempting to impose a Gaussian-distributed $\delta(\mathbf{x})$ produces pixels where $\langle N \rangle [\delta(\mathbf{x}) + 1]$ is negative, which is obviously unphysical. A better approximation for the galaxy distribution would be to use a log-normal distribution[72]. This distribution has the important property that it is zero for negative densities and it provides a reasonably good fit to the observed galaxy distribution. Our procedure for generating a log-normal random field uses the fact that that a log-normal map $\delta_{LN}(\mathbf{x})$ can be generated from a Gaussian map $\delta_G(\mathbf{x})$ with variance σ^2_G using

$$\delta_{LN}(\mathbf{x}) = e^{\delta_G(\mathbf{x}) - \sigma^2_G/2} - 1. \tag{4.9}$$

There exists a convenient relation between the correlation functions $\xi(r)$ of two maps related in this way:

$$\xi_G(r) = \ln\left[1 + \xi_{LN}(r)\right]. \tag{4.10}$$

This allows us to generate log-normal random fields $\delta_{LN}(\mathbf{x})$ with the desired characteristics.

We start with the galaxy power spectrum $C_\ell^{\rm gal}$ and convert it to a flat-sky ap-

proximation $C_k^{\rm gal}$. From this, we can calculate the correlation function we want our log-normal field to have using

$$\xi_{LN}(r) = \frac{1}{2\pi}\int_0^\infty kC_k^{\rm gal}J_0(kr)dk, \tag{4.11}$$

where J_0 is the Bessel Function of the first kind. We then calculate ξ_G using Equation (4.10) and convert this to a new angular power spectrum. We can then draw a Gaussian random field δ_G with this power spectrum using the procedure outlined above, then convert it to a log-normal random field using Equation (4.9).

The result of this process is a map like the one shown in the top right panel of Figure 4.1, with a power spectrum as shown in the bottom right panel. It is obvious when comparing these plots to those in the left hand column that the second map is much more strongly clustered and that the power spectrum contains a distinct scale-dependent component. The blue curve in this plot shows the predicted CO power spectrum from our example model. This predicted spectrum is in good agreement with the simulated data, which again shows the consistency between our theoretical models and our simulations.

4.3 Summary

We have demonstrated here a method for quickly creating simulations of intensity mapping signals. The simulated maps give the intensity of a line in a given frequency

bin as a function of position. Currently, we can only produce simulations of two-dimensional volumes, but we can approximate true intensity maps by stacking a number of these 2D simulations equal to the number of frequency channels in a survey. This is only a rough approximation of a true survey, as it does not take into account correlations between bins, but the approximation is good enough for us to begin investigating a number if interesting aspects of intensity mapping surveys, such as the problem of line foregrounds described in Chapter 5.

Chapter 5

Masking Line Foregrounds in Intensity Maps

One major difficulty in intensity mapping surveys is the problem of foregrounds. Every intensity mapping survey will have to deal with many types of foreground emission. Foregrounds with continuum frequency spectra such as dust or synchrotron are problematic, but the removal of these foregrounds is a well studied problem, especially for 21 cm surveys[73,74]. Because continuum emission is spectrally smooth, the contamination is mostly confined to Fourier modes with a small line-of-sight component, and can be removed by simply subtracting these modes. Line foregrounds are a more difficult problem. If a spectral line other than the target is redshifted into the same observing band, it is not easy to tell the two lines apart. 21 cm surveys are not expected to suffer from this issue since there are so few lines at such low

frequencies[75]. However surveys in other lines will require a better understanding of possible line foregrounds. In this chapter, we will discuss potential confusing lines in surveys targeting CO, Lyα, and CII, and study a method for removing them.

One method which can be used to remove these foreground lines is to cross correlate an intensity map with another map in a different frequency, or with some other tracer of large-scale structure[62,76]. Though each map will have its own foreground lines, they will come from different redshifts and thus be uncorrelated. The signals from the two maps will be correlated, as they come from the same redshift, leaving behind a cross spectrum which only depends on the two target populations. However, this method has two issues: The first is simply that it requires a more complicated (and costly) observation, since it requires observation of two signals. Secondly, it is difficult to reconstruct the auto power spectrum of the target line from the cross spectrum of the two maps (See work by Gong et al.[77], hereafter G14) due to added parameter degeneracies in the cross spectrum.

Another possible method for foreground separation is to examine the anisotropy in the combined power spectrum, as foreground lines will behave differently from the signal along the line of sight[78,79]. This is only possible at high signal-to-noise ratios though, which may not be attainable for early-generation surveys. A third technique would be to locate bright foreground emitters in galaxy surveys and mask the map voxels which contain them[39,37]. This again requires a significant amount of data beyond the survey itself.

We therefore seek a foreground-cleaning method which will work without any additional data and at modest signal-to-noise ratios. Though "foreground" lines could technically come from lower or higher redshifts than the target line, only lower redshift lines are likely to pose a problem, as the signal from an intensity mapping survey typically grows weaker with redshift. The fact that these lines come from lower redshifts offers a potential way to remove them. At lower redshifts, galaxy masses tend to be larger, so we expect there to be more very bright sources of a foreground line than a target line. This means that the brightest voxels in a survey will tend to be foreground galaxies, and the foreground contamination can be at least partially removed by masking out the brightest voxels in a survey. This technique tends to bias the target power spectrum because some signal is masked along with the foregrounds[68]. However, as we will show below, much of the *cosmological* information in the power spectrum is preserved after masking, even though most of the *astrophysical* information is lost. G14 also explored this technique for Lyα, though they do not appear to show any biasing of the signal power spectrum.

Using the simulation routines laid out in Chapter 4, we will demonstrate below that voxel masking is an effective technique for removing foregrounds when the signal is considerably brighter on average than the foregrounds, or when the voxel size of the survey is small enough that individual foreground sources can be isolated effectively. We find that CO intensity maps meet the first criterion, and Lyα intensity maps presumably meet the second. CII maps on the other hand, likely will not meet either

criterion, and therefore cannot be easily cleaned using simple voxel masking.

It is important to note that we found in Chapter 3 that the amplitude of a CO intensity mapping power spectrum is extremely uncertain, and that different models yield very different results. We have every reason to believe that the modeling of each and every signal and foreground line we discuss here is similarly uncertain. Therefore it should be noted that the details of the luminosity function modeling described in this paper should be taken with a grain of salt and the exact amplitudes of the power spectra discussed below could vary significantly from the values we use. Nevertheless, our intention is to explore the effects of voxel masking on contaminated intensity maps, and our general conclusions should hold for many different models.

This chapter is organized as follows. Section 5.1 describes how we model power spectra for the three signal lines we consider and their foregrounds, and Section 5.2 shows simulated maps for these models. Section 5.3 explains our method of removing foreground lines from these maps through voxel masking, and Section 5.4 shows the results of this masking. These results and pertaining issues are then discussed in Section 5.5, and we summarize in Section 5.6.

5.1 Signal and Foreground Models

Because our analysis here involves comparing intensity maps at different redshifts, we will again define our maps in this chapter in angular coordinates and study them in

terms of their angular power spectra C_ℓ. Real intensity mapping surveys will survey the sky in several frequency bands, which corresponds to several different redshift slices. The intensities in these slices will be correlated due to the existence of line-of-sight Fourier modes. However, we can obtain a reasonable approximation of a true survey by treating each frequency band as an independent map, which can then be stacked with the others to improve signal to noise (see Chapter 3). If we collapsed our whole 3D volume down to 2D, then we would lose a large amount of the information in our map. Discussing the angular power spectrum of individual slices, however, only sacrifices the information present in the line-of-sight modes[31]. All of the maps discussed in this chapter should be thought of as a single slice of a full 3D survey.

Since intensity mapping experiments are likely to have fairly narrow frequency bands, it is reasonable to assume that the quantity $\overline{T}$ does not change significantly over the width of a single band. This means that we can write the angular power spectrum as

$$C_\ell(z) = \overline{T}^2(z) C_\ell^{\rm gal}(z) + C^{\rm shot}(z), \tag{5.1}$$

where we take z to be the central redshift of the band. The galaxy angular power spectrum is calculated from $P_{\rm gal}$ using

$$C_\ell = \frac{2}{\pi} \int k^2 P_{\rm gal}(k) \left[\int f(r) j_\ell(kr) dr \right]^2 dk, \tag{5.2}$$

where r is the comoving distance, $j_\ell(kr)$ is the spherical Bessel function and $f(r)$ is the

selection function which is determined by the width of a single frequency band. This integral is computationally difficult to evaluate, so we again compute it approximately in low- and high-ℓ limits using the approximations from Section 2.2. The shot noise spectrum $C^{\rm shot}$ is scale independent, so the Limber approximation (Equation 2.10) can be used on all scales.

The values of $\overline{T}$ and $P_{\rm shot}$ for a given line depend sensitively on the exact shape of the luminosity function of that line. This luminosity function in turn depends on the conditions within the emitting galaxies which are highly uncertain, especially for high redshift sources. We can attempt to predict $L(M)$ or $\Phi(L)$ for different lines using various empirical observations. However, we found in Chapter 3 that different models assumed for CO emission yielded power spectra which spanned roughly two orders of magnitude in amplitude. Therefore the models we summarize below should not be interpreted as precise predictions of the power spectra for these lines. Rather, they are intended as a means to gain understanding of how the *shapes* of the luminosity functions of foreground lines can affect how strongly they contaminate a target line, separately from the overall amplitude ambiguity.

5.1.1 CO

When modeling the emission of CO and its foreground lines, we assume that only halos with masses above some cutoff mass $M_{\rm min}$ (assumed here to be $10^9\ M_\odot$) can emit the line in question, and we assume that only a fraction $f_{\rm duty}$ equal to the timescale

of star formation over the age of the universe of halos are emitting the line at any given time. For generality, we assume that the luminosity of a halo is a power law in its mass

$$\frac{L}{L_\odot} = A\left(\frac{M}{M_\odot}\right)^b, \tag{5.3}$$

where A and b are free parameters. This is a simple generalization of the model used in Chapter 3, where we assumed that $b = 1$. We again calibrate this relationship using a series of empirical scalings, starting with a relation between the line luminosity and FIR luminosity of the form

$$\frac{L_{\rm FIR}}{L_\odot} = C_{\rm FIR}\left(\frac{L'_{\rm line}}{\rm K\ km\ s^{-1}\ pc^2}\right)^{X_{\rm FIR}}, \tag{5.4}$$

where $C_{\rm FIR}$ and $X_{\rm FIR}$ are constants set through observations. In the above relation, the line luminosity is given in units commonly used for spectral line observations. The conversion to solar luminosities is given by[64]

$$\frac{L_{\rm line}}{L_\odot} = 3\times10^{-11}\left(\frac{\nu_{em}}{1\ \rm GHz}\right)^3\frac{L'_{\rm line}}{\rm K\ km\ s^{-1}\ pc^2}. \tag{5.5}$$

These relations, in combination with the star formation rate-FIR luminosity and halo mass-star formation rate relations given by Pullen et al.[31], we get the following

expressions for A and b:

$$A = 3 \times 10^{-11} \left(\frac{6.5 \times 10^{-8}}{C_{\rm FIR}} \right)^{1/X_{\rm FIR}} \left(\frac{\nu_{em}}{1 \text{ GHz}} \right)^3, \tag{5.6}$$

$$b = \frac{5}{3} X_{\rm FIR}^{-1}. \tag{5.7}$$

For CO, we again use an observed[80] CO-FIR relation with $C_{\rm FIR} = 1.4 \times 10^{-5}$ and $X_{\rm FIR} = 5/3$. This gives the values $A = 1.9 \times 10^{-6}$ and $b = 1$ (Compare to Chapter 3).

When considering possible foreground lines, we consider only lines with lower rest frame frequencies than our target lines. This is because, as shown in G14, projecting two power spectra from different redshifts boosts the amplitude of the lower redshift spectrum relative to the higher redshift one. Because we are working here with angular power spectra instead of three dimensional power spectra, this projection effect is included naturally in our calculations. For a CO survey targeted at $z = 3$, we are therefore concerned with lines that have emission frequencies between the 115 GHz CO rest frame frequency and the 28.8 GHz observing frequency.

The CO(1-0) line is expected to be considerably brighter on average than any other line in this range, i.e. $\overline{T}_{\rm CO}$ should be much greater than $\overline{T}$ for any foreground line. If we assumed that $L(M)$ were linear in M for all lines this would mean that no foreground line could dominate over the CO signal. However, the addition of the power law dependence on M means that the CO map could still be contaminated by

Table 5.1: Various parameters for the target CO(1-0) line as well as several possible foreground lines. Parameters include the emission frequency ν, the parameters of the FIR correlation $C_{\rm FIR}$ and $X_{\rm FIR}$, the $L(M)$ parameters A and b, and the observable parameters $\overline{T}$ and $C^{\rm shot}$.

Line	ν (GHz)	$C_{\rm FIR}$	$X_{\rm FIR}$	A	b	$\overline{T}$ (μK)	$C^{\rm shot}$ (μK^2)
CO[31]	115	1.35×10^{-5}	1.67	1.9×10^{-6}	1	0.60	3.5×10^{-6}
HCN[81]	88	794	1.0	1.7×10^{-15}	1.67	0.097	6.3×10^{-5}
CN[82]	113	1.6×10^{4}	0.89	6.9×10^{-18}	1.87	0.034	6.3×10^{-6}
CS[83]	49	2.1×10^{4}	1.0	1.1×10^{-17}	1.67	0.0083	1.5×10^{-5}
HCO$^+$[84]	89	158	1.11	7.4×10^{-14}	1.5	0.035	2.9×10^{-6}

the brightest sources for a foreground line which has $b > 1$. The shot noise term in particular is very sensitive to the value of b. For more general luminosity functions than the one we use here, this would mean that a line is potentially a problem if it falls off more slowly at high luminosities than the target line.

We consider foreground lines emitted by four molecules: HCN, HCO$^+$, CN, and CS. All of these molecules have higher critical densities than CO, and thus tend to trace denser regions of galaxies. For each molecule, we use an empirical correlation with FIR to estimate A and b. The results and sources for each line are given in Table 5.1, along with those of CO for reference. Note that for the CN and CS lines the FIR relations used were measured for higher order transitions (3-2 and 7-6, respectively). Thus we have made the assumption that the intensities of these lines are independent of which transition is being considered. The situation in reality is likely not so simple, but this approximation will suffice for our purposes.

Figure 5.1 power spectra for each of these lines compared with that of CO. Because all of the foreground lines have $b > b_{\rm CO}$, they are all dominated by shot noise. A

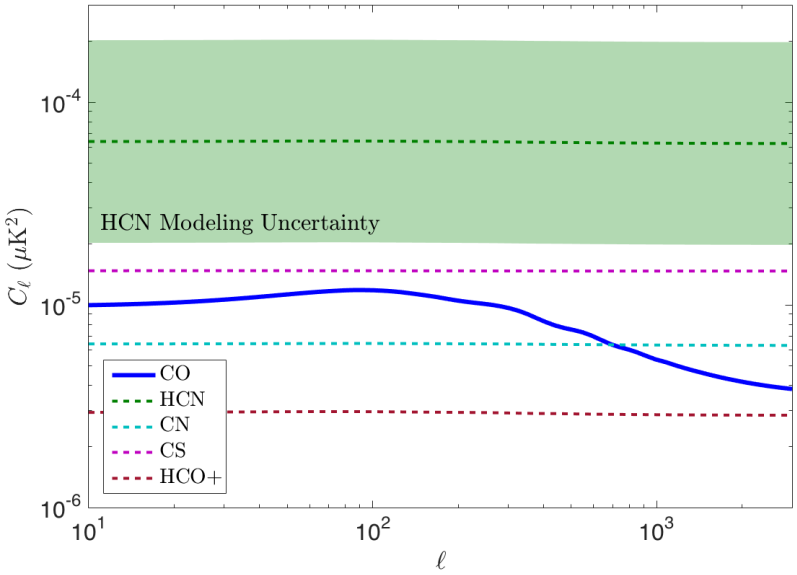

Figure 5.1: Power spectra for lines described in Table 5.1. The solid blue line is the target CO power spectrum, and the dashed lines show the spectra for the foreground lines. Note that all of the foregrounds are shot noise dominated. The shaded region shows a rough estimate of the theoretical uncertainty in these power spectra.

possible qualitative reason for this is the fact that all of these lines tend to trace denser gas than CO, and thus their luminosities may be more sensitive to the environments in their host galaxies.

Of the lines we consider, the HCN power spectrum clearly dominates. Though HCN is only $\sim 4\%$ as bright as CO on average, it has a higher value of b, which means that it produces a small number of bright sources which contribute a large amount of shot noise. This is how a line with relatively small average intensity can still be a problematic foreground. For the remainder of this paper, we will consider only the HCN foreground line since in our modeling the other lines are subdominant and should behave similarly under masking.

The shaded region in Figure 5.1 shows a very crude estimate of the theoretical

uncertainty in the HCN power spectrum, roughly an order of magnitude. In Chapter 3, we found that the amplitude of the CO spectrum could also vary by roughly an order of magnitude in either direction. For convenience we only show this shading for HCN, but all of the other spectra are similarly uncertain, if not more. Given the vast amount of uncertainty both in the modeling of these spectra as well as in the empirical measurements used in our model, we do not attempt to make a more accurate estimate of the error bars on these spectra.

5.1.2 Lyman α

For the Lyman α line we consider a hypothetical survey targeted at $z = 7$, and we follow the modeling of G14 and references therein, which we summarize briefly here. Since the Lyα line will be observed in the infrared rather than the radio, we use intensity units here instead of brightness temperature. Recombination and collision processes within galaxies give a Lyα luminosity

$$\frac{L_{\mathrm{Ly}\alpha}(M,z)}{L_\odot} = 5.1\times10^8 f_{\mathrm{Ly}\alpha}(z)\left[1-f_{\mathrm{esc}}^{\mathrm{ion}}(M)\right]\frac{SFR(M,z)}{M_\odot\ \mathrm{yr}^{-1}}, \tag{5.8}$$

where $f_{\mathrm{Ly}\alpha} = 3.34\times10^{-3}\times(1+z)^{2.57}$ is the fraction of Lyα photons not absorbed by dust, and $f_{\mathrm{esc}}^{\mathrm{ion}} = \exp\left(-5.18\times10^{-3}\times M^{0.244}\right)$ is the escape fraction of ionizing photons. The star formation rate estimated from simulated galaxy populations[85] is

parameterized by

$$SFR(M, z = 7) = 1.6\times10^{-26} \left(\frac{M}{M_\odot}\right)^{2.59} \left(1 + \frac{M}{M_1}\right)^{-0.62} \left(1 + \frac{M}{M_2}\right)^{0.4} \left(1 + \frac{M}{M_3}\right)^{-2.25}. \tag{5.9}$$

We take $M_{\rm min} = 10^8\ M_\odot$ and $M_{\rm max} = 10^{13}\ M_\odot$. Note that for simplicity we neglect Lyα emission from the IGM. The IGM contribution is small compared to the halo emission, and it is not as easily simulated using our methods. We calculate the bias for Lyα following G14.

For Lyα at $z = 7$, three foreground lines are considered: Hα coming from $z \sim 0.5$, OIII from $z \sim 0.9$, and OII from $z \sim 1.6$. For these lines, the luminosity function is assumed to be a Schechter function

$$\Phi(L)dL = \phi_* \left(\frac{L}{L_*}\right)^\alpha \exp\left(-\frac{L}{L_*}\right) \frac{dL}{L_*}, \tag{5.10}$$

where $\Phi(L)$ is the comoving number density of halos with luminosities between L and $L + dL$, and ϕ_*, L_*, and α are parameters which can be determined from observations[86]. With the above relations, we can calculate angular power spectra for Lyα and the three foreground lines. Following G14, we assume that the halo bias for these lines is proportional to the halo mass. Figure 5.2 shows the calculated power spectra for Lyα and the three foreground lines along with the total foreground spectrum. Note that the foreground lines dominate entirely over the signal, and unlike CO, the foregrounds have a significant clustering component.

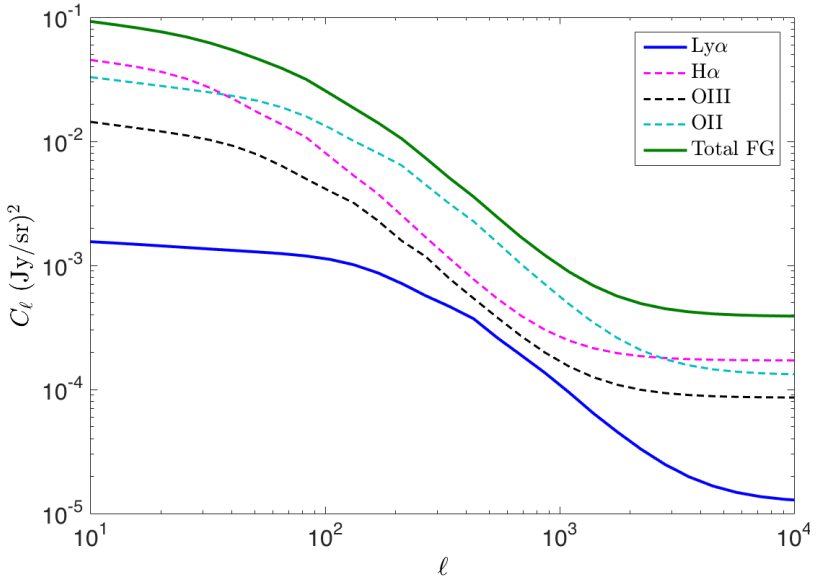

Figure 5.2: Theoretical power spectra for Lyα and its three foreground lines

5.1.3 CII

We base our simulations of CII intensity maps on the modeling done by Silva et al.[38], hereafter S15. We consider a possible survey targeted at $z = 7$. For the CII luminosity we use model $\mathbf{m_2}$ from Table 1 of S15, where the luminosity is given by

$$\log 10 \left(\frac{L_{\rm CII}}{L_\odot} \right) = \log 10 \left(\frac{\psi(M,z)}{M_\odot} \right) + 6.9647, \tag{5.11}$$

where

$$\psi(M,z) = M_0 \left(\frac{M}{M_a} \right)^{a_{\rm CII}} \left(1 + \frac{M}{M_b} \right)^{b_{\rm CII}}. \tag{5.12}$$

The parameters in the formula for ψ at $z = 7$ are $M_0 = 6.6 \times 10^{-6}\ M_\odot$, $M_a = 10^8\ M_\odot$, $M_b = 1.6 \times 10^{11}\ M_\odot$, $a_{\rm CII} = 2.25$, and $b_{\rm CII} = -2.3$.

The primary foregrounds for a CII survey come from CO transitions. Since such

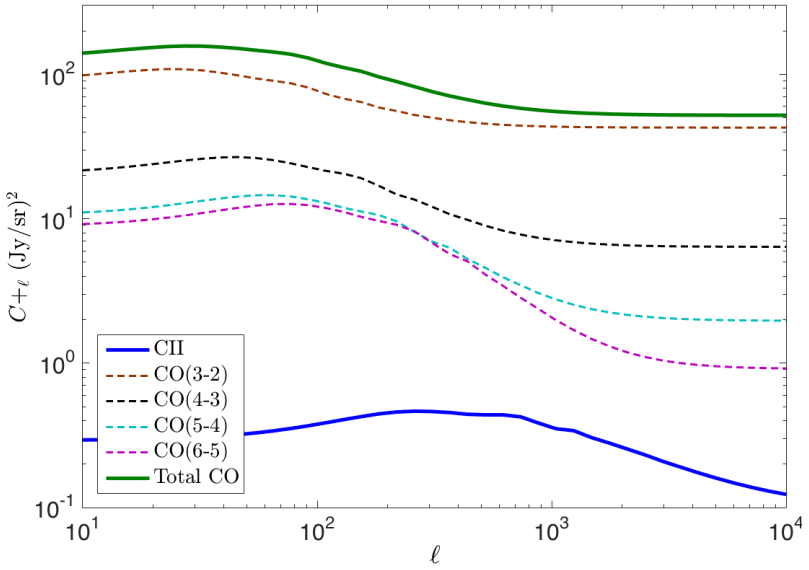

Figure 5.3: Theoretical power spectra for CII and its four foreground CO lines.

a survey would observe at roughly 240 GHz, we need not worry about the CO(1-0) line at 115 GHz or the CO(2-0) line at 230 GHz. The higher order transitions are potentially problematic though. We estimate the intensities of the higher order CO lines using our formulae above along with the line ratios observed for submillimeter galaxies[27]. This table gives the ratio $L'/L'_{CO(1-0)}$ for the transitions from CO(2-1) to CO(5-4). Following S15, we consider the CO(6-5) line as well, and assume that it has the same luminosity ratio as the CO(5-4) transition. We also calculate the bias by assuming that the halo bias is proportional to the halo mass. Figure 5.3 shows the calculated power spectra for these CO lines as well as the spectrum of the target CII line. As with Lyα, the foreground lines are strongly clustered and considerably brighter than the target line.

5.2 Simulated Maps

With the simulation method from Chapter 4, we can now simulate maps for any line for which we can define either a function $L(M)$ or a luminosity function $\Phi(L)$. Since we are working in angular space, it is also very easy to combine maps for different lines. A map containing two lines is simply the sum of the maps for the individual lines, as the signal and foreground emitters should be separated by large enough distances as to be uncorrelated. Using Equation 5.1, we can obtain best fit power spectra for these simulated maps through least-squares fitting assuming the same $C_\ell^{\rm gal}$ we used to generate the map with $\overline{T}$ and $C^{\rm shot}$ as free parameters.

If we simulate a map for CO at $z = 3$ and a foreground HCN line at $z = 2$, we get power spectra like the ones shown in Figure 5.4. The best fit power spectra are plotted as solid lines. As expected, the effect of foreground contamination is to significantly increase the shot noise of the map. The amplitudes of the power spectra in these simulations depend on exactly how many halos are drawn and exactly what masses are assigned to them. This means that the result can vary somewhat from the theoretical predictions. The HCN power spectrum is particularly sensitive to this, since it depends so much on the highest mass halos. The shaded region in Figure 5.4 shows the 95% confidence range of the HCN power spectrum, calculated by comparing the results of 500 simulated HCN maps.

Figure 5.5 shows simulated maps for CO at $z = 3$ and Lyα and CII at $z = 7$ along with their foregrounds in order to illustrate the qualitative differences between them.

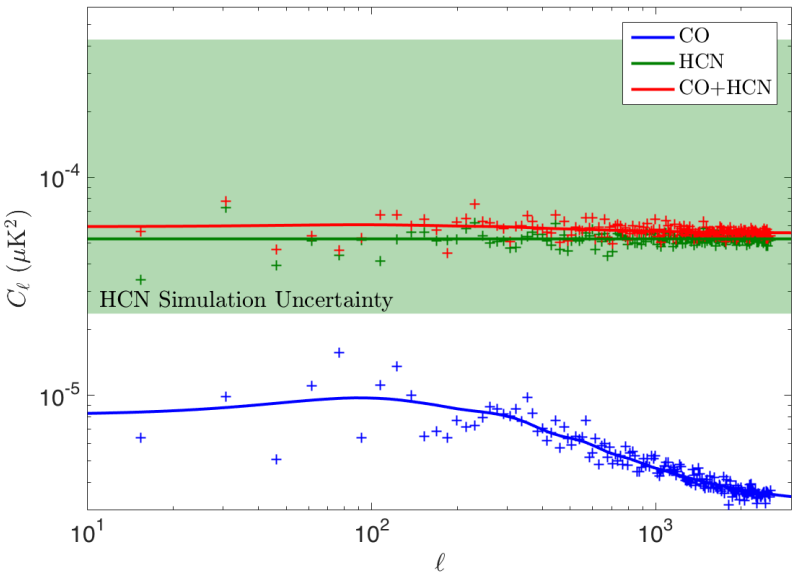

Figure 5.4: Power spectra of simulated maps of CO at $z = 3$ (blue), HCN at $z = 2$ (green), and the sum of the two (red). Solid curves are best fit power spectra allowing $\langle T \rangle$ and C_ℓ^S to vary. As expected, the HCN foreground contributes a significant amount of extra shot noise at high ℓ. The shaded region corresponds to the 95% confidence interval for the simulated HCN maps. Variations from simulation to simulation can cause the power spectrum amplitude to change significantly.

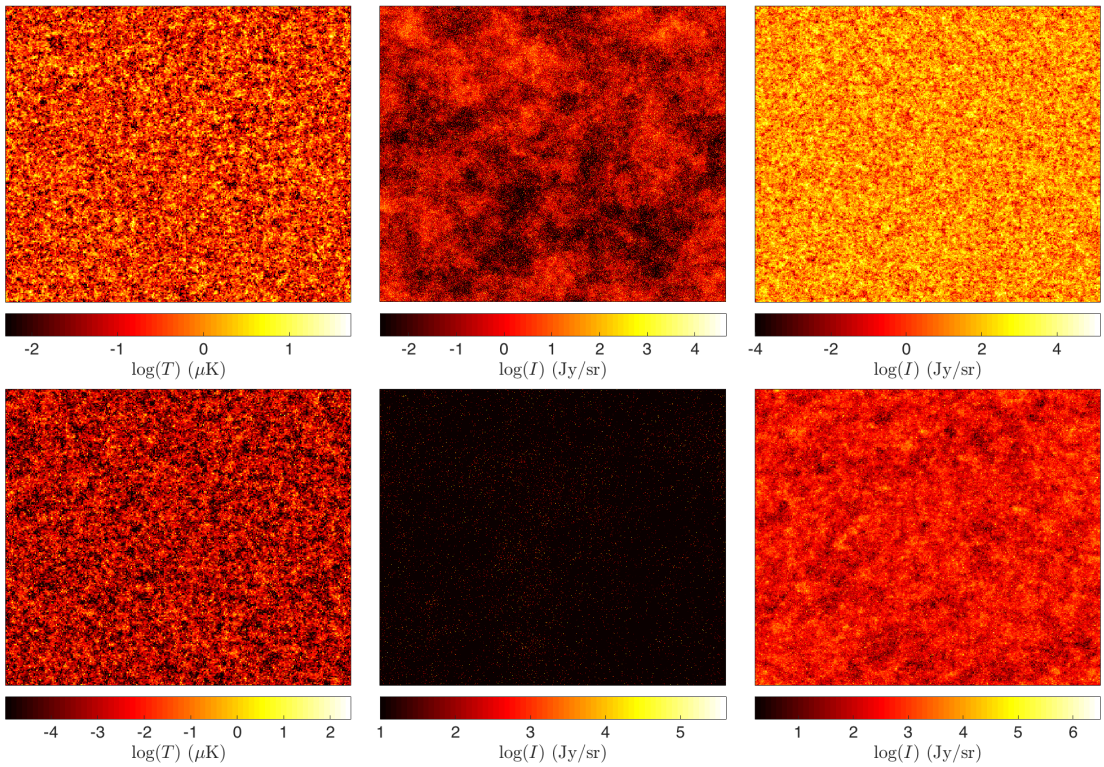

Figure 5.5: Simulated maps of CO at $z = 3$, Lyα at $z = 7$, and CII at $z = 7$ (top row, left to right) along with their foregrounds (bottom row). The CO simulations cover 550 deg^2 with 10 arcmin resolution, the Lyα simulations cover 1 deg^2 with 0.1 arcmin resolution, and the CII simulations cover 100 deg^2 with 3.2 arcmin resolution.

The CO simulations cover 550 deg^2 with 10 arcmin resolution and $\Delta\nu/\nu_{\rm obs} = 10^{-3}$ (based on our results from Chapter 3), the Lyα simulations cover 1 deg^2 with 0.1 arcmin resolution and $\Delta\nu/\nu_{\rm obs} = 1/40$ (based on G14), and the CII simulations cover 100 deg^2 with 3.2 arcmin resolution and $\Delta\nu/\nu_{\rm obs} = 1.7 \times 10^{-3}$ (based on S15).

The HCN foreground map is considerably fainter on average than the CO map, but though the brightest sources are difficult to see on this image the colorbar shows that the HCN map extends to higher intensities than the CO map. The Lyα maps cover a much smaller region of the sky than the others, so the clustering features in

this image appear much larger. In addition, each voxel in the Lyα foreground map covers a rather small volume of space, so most of the voxels in this map are dark. Though the Lyα foregrounds are brighter than Lyα on average, most of the intensity in this map comes from a small number of voxels. The foreground map for CII is much brighter than the signal, and the voxels in the CII map are large enough that there is strong emission in most voxels.

Though the models we use in this work are fairly simple, this simulation method is straightforward to generalize to different, more complex models. Any power spectrum and luminosity function can be input to get a simulated map. One important effect that we have not taken into account here is the correlation between maps made at different frequencies due to line-of-sight Fourier modes. Since the target lines we consider here are widely separated from the foreground lines in frequency space this effect should not be significant for the problems discussed here. However, it will need to be taken into account when attempting to accurately simulate the three-dimensional data taken by realistic intensity mapping experiments.

5.3 Voxel Masking

As noted above, the majority of the contamination in CO surveys comes from a few bright foreground emitters which add a large amount of shot noise to a survey. This means that the foreground effect could be mitigated if we simply mask out

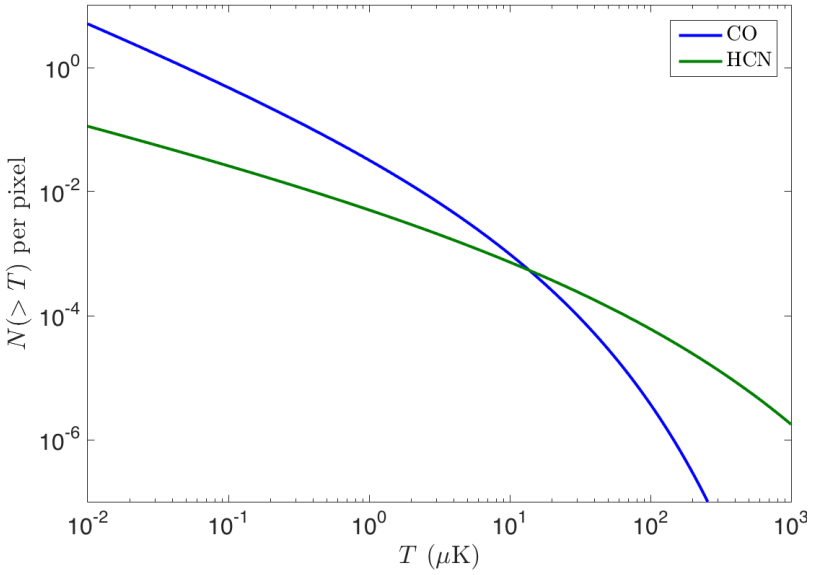

Figure 5.6: Number of sources per voxel which contribute brightness temperature greater than T. At low intensities, the target CO line dominates as expected. However, at high intensities the HCN emitters begin to dominate because the HCN luminosity function does not fall off as quickly as that of CO.

the brightest voxels in our survey. G14 noted that a similar argument holds for other lines, though perhaps not to the same extent. Figure 5.6 shows the number of sources/voxel in a CO survey which produce brightness temperatures above a given value. At brightness temperatures above a few hundred K, there are more HCN emitters than CO emitters. This is primarily due to the fact that $b_{\rm HCN} > b_{\rm CO}$, with some additional contribution from the fact that there are more high mass halos at $z \sim 2$ than at $z \sim 3$.

This means that if a given voxel has a very high intensity, it is likely that the extra flux is coming from a foreground HCN emitter. Thus if we mask all voxels brighter than a given value we will remove on average much more foreground emitters than target emitters. In addition, the foreground sources we mask are exactly the bright

sources which produce the worst contamination. Therefore we expect that we can clean the foregrounds out of a map by masking out the brightest voxels.

The effects of voxel masking on the power spectra of these maps are difficult to understand analytically. We can make a very rough estimate if we assume that we can perfectly isolate and remove all sources which contribute an intensity greater than some cutoff. The ratio of the average brightness temperature of a line before and after applying this cutoff is

$$\frac{\overline{T}_{\rm masked}}{\overline{T}_{\rm unmasked}} = \frac{\int_{M_{\rm min}}^{M_{\rm cut}} L(M) \frac{dn}{dM} dM}{\int_{M_{\rm min}}^{M_{\rm max}} L(M) \frac{dn}{dM} dM}, \tag{5.13}$$

If we choose to remove $\sim 1\%$ of the voxels from a map, $M_{\rm cut}$ is the mass of a halo such that 1% of halos have $M > M_{\rm cut}$. For CO, we have $M_{\rm cut} = 6.2 \times 10^{10}\ M_\odot$ and for HCN we have $9.5 \times 10^{10}\ M_\odot$. The ratio of the shot noise before and after masking is similar to Equation (5.13) with $L(M)$ replaced with $L^2(M)$.

This analytical procedure predicts that the ratio of the masked to unmasked CO intensity will be ~ 0.7 and that of the masked to unmasked CO shot noise will be ~ 0.06. The ratio of masked to unmasked HCN shot noise is predicted to be $\sim 2 \times 10^{-5}$. Therefore, although masking voxels causes the CO power spectrum to decrease, we expect the HCN spectrum to decrease far more, leaving behind a signal dominated power spectrum. However, the procedure described here neglects a few key effects which would be present in a real map. If a map voxel contains a bright

source, masking it will also remove all of the fainter sources present in the same voxel. In addition, some voxels will be masked because they contain a large number of faint sources rather than a single bright one. In order to fully account for these effects, we need to rely on our simulated maps, as we do below.

In order to use this procedure one must have some means of determine how many voxels should be masked. For CO, and any other hypothetical line where all of the foregrounds are shot-noise dominated, the answer is simple. We only need to mask until the masked power spectrum shows clustering behavior out to the desired angular scale. Even if the CO foregrounds are faint enough that CO dominates entirely, this masking will reveal the CO clustering behavior on scales which are normally obscured by shot noise. For lines with clustered foregrounds, the answer is less clear. If we have a reasonable estimate of the luminosity functions of the signal and foregrounds, we can use a plot like Figure 5.6 to predict a cutoff intensity (this is the method used in G14).

5.4 Masking Results

Here we show the results of applying the masking procedure described above to CO, Lyα, and CII intensity maps.

5.4.1 CO

Figure 5.7 shows the effect of voxel masking on the power spectra of three simulated maps: one with just CO, one with HCN, and the sum of the two. The dots and dashed curves show the power spectra and best fit curves for the three maps with no masking. The foreground power spectrum entirely dominates on all scales. The pluses and solid curves show the spectra of the maps after all voxels brighter than 6 μK are masked. This value corresponds to masking roughly 1% of the voxels in the map with both signal and foregrounds. After masking, the foreground power spectrum has dropped dramatically, and the CO power spectrum has fallen by a much smaller amount. But most importantly, the red total power spectrum is very similar to the CO power spectrum. The power spectrum of the map with foregrounds is nearly identical to that of a map without them. There remains some small amount of shot noise contamination, but this could be removed by choosing a lower cutoff value for masking. Thus it appears that foreground contamination in CO intensity maps can indeed be mitigated by masking bright voxels.

It is worth mentioning that the amount by which the CO power spectrum drops differs somewhat from the simple prediction made in Section 4. The clustering term in our simulation decreases roughly 40% less than the simple calculation predicts, and the drop in shot noise is ~ 3 times less than predicted. As mentioned above, this happens because the calculation in Section 5.3 does not take into account the fact that there are multiple sources in each voxel, which will either be masked along with

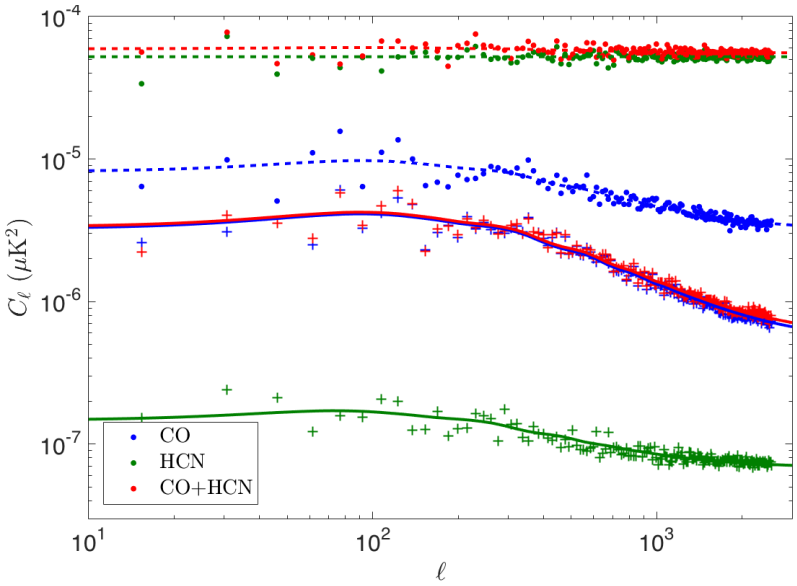

Figure 5.7: Power spectra and best fit curves for simulated maps of CO at redshift 3 (blue), HCN at redshift 2 (green), and the sum of the two (red). Dashed curves/dots show the spectra before masking, solid curves/pluses show the results of masking all voxels above 6 μK ($\sim$ 1% of the voxels in the total map). After masking, the maps with and without foregrounds have very similar power spectra.

a single bright source or add together to mimic a bright source. This discrepancy makes it difficult to use Equation (5.13) to predict the unmasked spectrum from a masked one.

5.4.2 Lyman α

To facilitate comparison between our results and those given in G14, we choose to mask 3% of the voxels in our Lyα simulations. Figure 5.8 shows the effectiveness of this masking. The dashed lines show the power spectra of the signal, foregrounds, and total map before masking, and the solid lines show the spectra after masking. For the sake of visibility, we have plotted only the total foreground spectra rather than

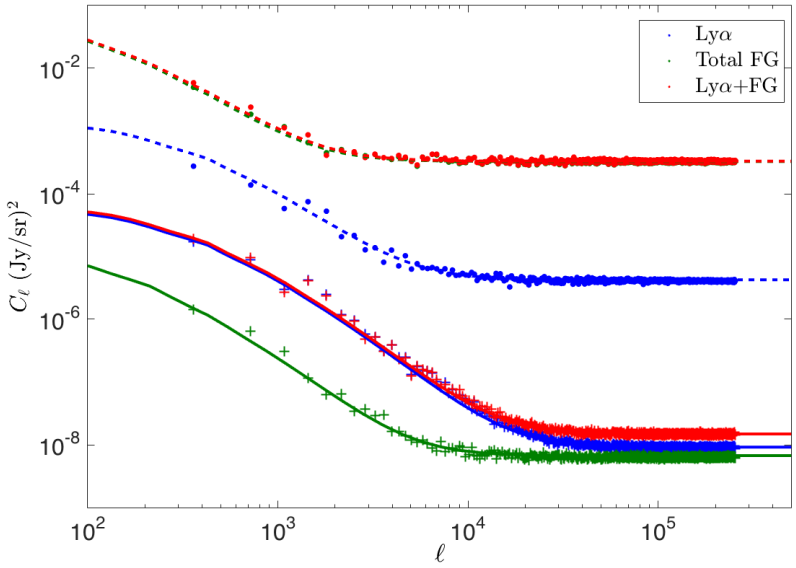

Figure 5.8: Power spectra and best fit curves for simulated maps of Lyα at $z = 7$ (blue), the sum of the three foreground lines (green), and all four lines combined (red). Dashed curves/dots show the spectra before masking, solid curves/pluses show the results of masking all voxels above ~ 100 Jy/sr (3% of the voxels in the total map). As with CO, the masking removes a large amount of the foreground contamination.

those of the three individual lines. After masking, the foreground has dropped below the signal and the map which includes both signal and foregrounds gives a spectrum very similar to the map without foregrounds. Thus it appears that for the model and telescope resolution simulated here it is possible to remove most of the foreground contribution to the power spectrum.

5.4.3 CII

Figure 5.9 shows the power spectra of the CII simulations before and after masking 1% of the voxels. Note that unlike in our CO and Lyα simulations, the foreground lines dominate over the signal even after masking. Though we have plotted the results

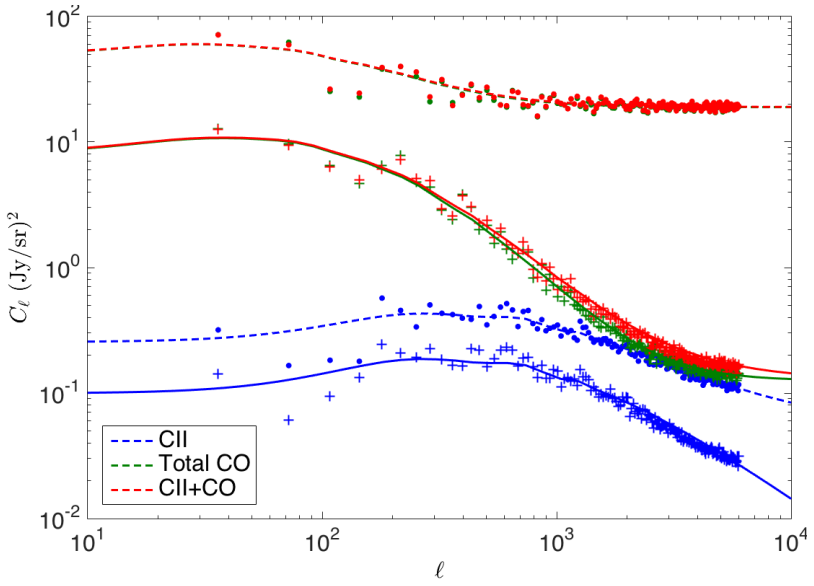

Figure 5.9: Power spectra and best fit curves for simulated maps of CII at $z = 7$ (blue), the total CO foreground (green), and the total emission (red). Dashed curves/dots show the spectra before masking, solid curves/pluses show the results after masking all voxels brighter than $\sim 10^4$ Jy/sr (1% of the voxels in the total map). After masking, the foregrounds still dominate over the signal. The foregrounds continue to dominate no matter what masking percentage is chosen.

of masking 1% of the voxels, the same basic result holds true no matter how many voxels we mask. At this resolution, it is not possible to move the foreground power spectrum below that of the signal no matter how many voxels are removed.

5.5 Discussion

The details of masking the simulated maps explain why the three lines we consider behave somewhat differently under masking. For the CO map, the contamination is shot noise from the few brightest galaxies, which is easily removed. However, both Lyα and CII have foregrounds which include strong clustering components. In

addition, for both of these lines the average intensity of the foregrounds is large compared to that of the signal, as seen in G14 and S15. Despite this, the masking works well for Lyα and poorly for CII. The reason for this has to do with the size of the voxels used in each map. Lyα is a higher frequency line than CII or CO, so it is easier to map with high angular resolution. This means that it is easier to isolate a single bright foreground source without removing as much signal. Our results show that if a CII survey is limited to arcminute scale voxels then foregrounds cannot be easily cleaned through voxel masking, and we must resort to cross-correlations to isolate the target line.

The drop in the amplitudes of the target spectra after masking is an unfortunate side effect of voxel masking. Ideally, an intensity mapping survey would be able to recover the dashed blue curve in Figure 5.7 rather than the solid blue curve. This is because we cannot avoid masking some CO sources when removing HCN. As noted by Visbal et al.[68], some of the information in the power spectrum is lost in this masking process, specifically that encapsulated in the amplitudes of the clustering and shot noise terms. On a brighter note, since we are only masking a percent or so of the voxels, the shape of the clustering term of the power spectrum does not change significantly when masked. This means that masking allows us to recover the *shape* of the galaxy power spectrum even on scales where the foregrounds dominate over the signal.

This is a potentially useful measurement for a variety of *cosmological* purposes.

For example, Brax et al.[87] show that modified gravity models alter the shape of the power spectrum which would be measured by a 21 cm survey. Similar changes could be studied in the masked maps we describe here. However, all of the *astrophysical* information which was contained in the *amplitudes* of the clustering and shot noise components is lost in masking, so we cannot make statements about the luminosity functions of the target galaxies. This means that, while these intensity mapping surveys could still be useful with only this simple foreground cleaning, it may be necessary to use other foreground cleaning methods to reach their full potential.

Nonetheless, the information lost in voxel masking could potentially be recovered in a number of ways. Unfortunately, the estimate of the change in power given in Equation (5.13) differs significantly from our simulations, so this cannot be used to recover the unmasked power. As stated previously, cross-correlations, direct observations, and power spectrum anisotropies all provide potential unbiased means of cleaning foregound lines. One could also make use of the one-point statistics of a map to separate signal from foreground (see Chapter 7).

Note that care must be taken when fitting uncleaned intensity mapping power spectra to determine cosmological or astrophysical parameters. For example, The power spectrum of a CO intensity map can in principle be fit to determine the values of various model parameters[35], and the fitted parameters will depend on how much shot noise is present in the map. However, the foreground line will add an uncertain amount of extra shot noise to the spectrum, which will bias the results of any fit.

We have not included the effects of instrumental noise in any of these calculations. Though we leave for future work a full analysis of the behavior of noisy maps with masked voxels, we can make some basic arguments to predict whether or not our basic premise would hold in a map with noise. If a map has too much noise, then the voxels which are masked will be bright due to random noise fluctuations in no correlation with the brightness of the target or foreground galaxies in those voxels. Therefore, if a simulation has too much noise this voxel masking technique is useless. The survey parameters used in Chapter 3, which were chosen to provide a reasonable chance of detecting CO at $z = 3$, give a noise per voxel $\sigma_N = 1.7$ μK. If we generate noise maps with the value in each voxel drawn from a Gaussian with zero mean and standard deviation of σ_N, we can estimate how many of the masked voxels will be bright due to noise rather than signal. We find that adding noise to our simulated CO+HCN map increases the number of voxels above our 6 μK cutoff by roughly 25%. These spurious bright voxels will reduce the effectiveness of the masking, but it should still be possible to remove most of the bright sources from a map. If the noise is significantly stronger, the masking will be ineffective. However, in this case the signal to noise ratio for detecting the CO line at all becomes considerably smaller as well.

In addition, we assumed the linear form for the underlying matter power spectrum in our calculations. Though we account for some nonlinearity by using a lognormal galaxy density field, a full treatment of the nonlinearities would add more power to

our maps on small scales. However, this would likely not have a significant qualitative effect on our results. There may be some minor differences in the ratio of signal/foreground power after masking, however the scales where nonlinearity becomes most significant are also scales where the power spectra tend to be shot noise dominated, so we do not expect any dramatic effects from a nonlinear calculation.

5.6 Summary

We have presented an exploration of the effectiveness of bright voxel masking on removing foreground lines from intensity maps. Using empirical luminosity function models and simulated intensity maps we have illustrated how masking changes the power spectra of maps for three cases: CO contaminated with HCN, Lyα contaminated with various atomic lines, and CII contaminated with higher order CO lines. For the CO survey, the foreground line was faint enough on average that removing the brightest voxels significantly dropped the amplitude of the foreground spectrum. The high angular resolution possible in the Lyα survey meant that the foreground contamination was limited to a few voxels which could be easily masked. The CII survey, however, had both bright foregrounds and large voxels, so the masking was found to be ineffective.

For all of the lines, masking bright voxels altered the amplitude of the recovered power spectrum away from the desired uncontaminated value. This means that

masking loses some of the information in the spectrum. However, in the two surveys where masking was effective, the masked spectrum had a clustering component with the same shape as the unmasked clustering spectrum. Therefore, though the astrophysical content of the map is lost, the cosmological information contained in the shape of the clustering spectrum can be recovered from a masked map. Thus, voxel masking seems to be a useful technique for obtaining information from even a highly contaminated CO or Lyα map. If we are to obtain the remainder of the information in these surveys, it will be necessary to use some other foreground cleaning technique, such as cross correlation, or to augment it with a $P(D)$ analysis of the progressively masked power spectra. If we are to fully unlock all of the benefits of intensity mapping surveys, it is imperative that we utilize these or other methods to isolate the signal from the foregrounds.

Chapter 6

Observing Molecular Gas with CO Cross-Correlations

In Chapter 5, we discussed how bright foreground lines can contaminate intensity maps, and described one possible method of removing them. Even faint foreground lines, however, still contain information about their emitters. In this chapter, we will discuss how to isolate and access this information. To illustrate this, we will use the example of the first rotational transition of ^{13}CO, an isotopologue of the ^{12}CO molecule we have been discussing so far. In particular, we will show how we can use the ^{13}CO line to study molecular gas properties at high redshift.

This chapter is primarily based on work published in Ref.[48]. In Section 6.1, we will describe the usefulness of ^{13}CO as a tracer of galaxy properties. In Section 6.2, we will explain schematically how to use cross-correlations to access the ^{13}CO line in

a map dominated by the brighter ^{12}CO isotopologue, with a more formal treatment present in Section 6.3. We forecast the ability of a hypothetical CO intensity mapping experiment to constrain molecular gas physics using this cross-spectrum in Section 6.4. A brief discussion of these results can be found in Section 6.5, along with a summary in Section 6.6.

6.1 Molecular Gas and ^{13}CO

Most available observations of molecular gas come from the Milky Way and other nearby galaxies. However, recent years have seen an explosion of information about the high redshift universe, from precise measurements of the cosmic microwave background[2] to large galaxy surveys such as the Sloan Digital Sky Survey[88]. A consensus picture of the cosmic star formation history has emerged, with star formation rates gradually rising from the end of reionization to a peak around $2 \lesssim z \lesssim 3$ before beginning to decline[89]. Unfortunately, our understanding of molecular gas at these redshifts is still quite limited; direct molecular emission has only been measured for a small selection of sources at these redshifts, limited to the most luminous of far-infrared and submillimeter galaxies[90]. A full picture of high-redshift molecular gas and its relationship to star formation requires a more systematic approach.

Though it has seen quite a bit of study, there are drawbacks to using ^{12}CO, the most common CO isotopologue, alone as a molecular gas tracer. Since the rotational

lines of ^{12}CO are strong, they saturate at relatively low column densities compared to other lines, becoming optically thick relatively soon after densities rise sufficiently to support CO formation in the first place[26]. Thus it is a common practice to combine observations of ^{12}CO with those of other, fainter lines which can penetrate further into molecular clouds. ^{13}CO, a stable isotopologue of the usual ^{12}CO molecule, is commonly used to gain information about denser material in environments ranging from protoplanetary disks[91] to giant molecular clouds[92]. ^{13}C has an abundance on the order of a few percent of that of ^{12}C[93], and the relative abundances of the two CO isotopologues is similar, up to corrections due to chemical fractionation and selective photodissociation[94]. Combining measurements of ^{12}CO and ^{13}CO for a single system facilitates the study of molecular gas across a wider density range[95].

At cosmological distances, the ^{12}CO emission lines can only be observed in very bright sources, and ^{13}CO observations are limited only to extreme systems such as the strongly lensed Cloverleaf quasar[96]. The difficulty of detecting these lines severely limits what can be learned from traditional galaxy surveys at these redshifts. Even in these very bright sources, most radio surveys lack the resolution needed to probe sub-galactic scales, preventing the study of gas distributions within any individual galaxy. Intensity mapping, on the other hand, probes the entire galaxy population, including the large numbers of faint sources inaccessible to traditional surveys. Below, we will demonstrate how CO intensity maps can measure the detailed molecular gas composition and distribution at high redshifts. These measurements will help unveil

how star formation is fed in distant galaxies, systems that are all but invisible to current techniques.

6.2 Leveraging Isotopologues

Any ^{12}CO survey will also contain ^{13}CO in an overlapping cosmological volume. Though ^{13}CO is too faint and only marginally separated in frequency in comparison to ^{12}CO to ever be detected by itself, by cross-correlating the proper frequency bands in a CO survey, we can measure the correlation between the intensities of these lines. Since the lines come from the same population of galaxies, this cross-correlation depends on the relative abundances of the two carbon species and the optical depths of the CO lines. Thus by performing this cross-correlation, we can make statements about the distribution of gas densities in our galaxy population, despite having no information about any individual galaxy or molecular cloud. This measurement cannot easily be made using any other method, yet it is available without any extra observation from any CO intensity mapping survey.

Consider a CO intensity mapping survey similar to the planned COMAP experiment described in Table 2 of Li et al.[35] which would cover a frequency range from 30 to 34 GHz. This corresponds to $z = 2.39$ to 2.84 in ^{12}CO and $z = 2.24$ to 2.67 in ^{13}CO. Figure 6.1 shows a schematic view of what the contributions to a survey from these two lines might look like. A given set of galaxies will emit both ^{12}CO and

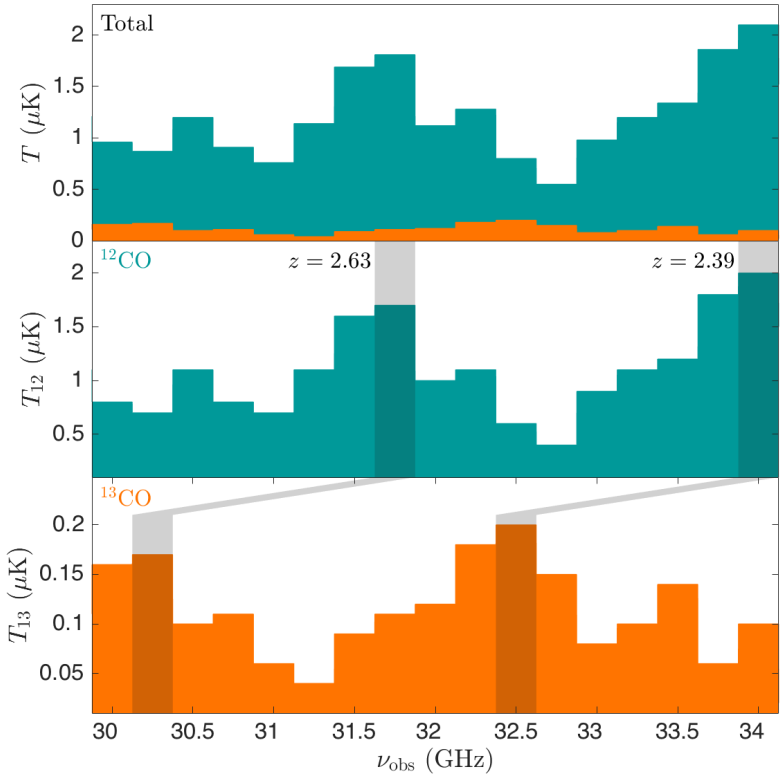

Figure 6.1: A schematic view of the contributions from ^{12}CO and ^{13}CO to a hypothetical intensity mapping survey. The top panel shows the total observed intensity in each frequency bin assuming that the observed ^{13}CO intensity from all galaxies is 10% of the ^{12}CO intensity. The middle and bottom panels show the contribution to the total signal from ^{12}CO and ^{13}CO emission respectively. The shaded regions in these panels highlight emission that comes from the same slice of physical space.

^{13}CO lines at a given position in physical space. The two lines are then redshifted to different bands in frequency space as shown in the bottom two panels, then added together to produce the observed signal in the top panel. For illustration purposes, Figure 6.1 assumes that the observed ^{13}CO intensity from all galaxies is 10% of the ^{12}CO intensity.

Even with a somewhat generous 10% intensity ratio, the contribution from ^{13}CO

to the total intensity is small, and will therefore be virtually impossible to isolate in any individual band. However, looking at the bottom two panels of Figure 6.1 demonstrates that the two lines trace the same structure, just shifted in frequency. At $z = 2.6$, this shift in observed frequency is $\Delta\nu_{12/13} = 1.4$ GHz. The shaded bands show the two peaks which appear at the same location in physical space. This means that we should expect significant correlation between bands separated by $\Delta\nu_{12/13}$. By cross-correlating said bands we can compare the emission in both isotopologues coming from the same set of sources.

What we will observe in this cross-correlation depends on the details of the molecular gas properties in the surveyed volume. Obviously we expect the signal to depend on the relative amounts of ^{12}C and ^{13}C, which depends on the nucleosynthesis history up to the observed redshift. The observed signal will also depend on the density of the emitting molecular clouds. Since ^{13}C is significantly less abundant than ^{12}C, we expect the ^{12}CO line to become optically thick well before the the ^{13}CO line does. If much of the CO molecules in a galaxy are found in very dense regions, this effect can significantly alter the observed intensities.

It is relatively straightforward to qualitatively predict what the observed intensity ratio I_{13}/I_{12} between the lines will be in two extreme cases. We will ignore for now small differences between the emission and absorption properties of the two molecules. In the limit of very diffuse gas, where all of the ^{12}CO emission is optically thin, the only thing that matters is the relative abundance of the two species and we expect

I_{13}/I_{12} to approach the abundance ratio of the two isotopologues. In the opposite limit, where the emission comes from a very dense region and both lines are saturated, we expect I_{13}/I_{12} to approach unity, since each photon will undergo a large enough number of emission and absorption events to wash out any dependence on the relative abundances. These results depend only on basic radiative transfer physics, and are independent of how the gas is modeled.

In real galaxies, we expect there to be contributions from both optically thin and optically thick regions. By cross-correlating frequency bands separated by the frequency difference between the two isotopologues, we can constrain both the abundance ratio of the two species and the amount of emission that comes from dense, optically thick regions. Both of these properties depend on both the instantaneous star formation rate and the cumulative star formation history at a given redshift. Below, we will derive this intensity ratio as a function of these quantities in a more quantitative fashion.

6.3 Formalism

In this section we will present a general formalism for predicting the cross power spectrum of ^{12}CO and ^{13}CO emission in an intensity map. This method should be applicable regardless of the exact model used to predict the molecular gas properties of a galaxy distribution.

6.3.1 ^{12}CO/^{13}CO Intensity Ratios

We derive the intensity ratio of the two CO lines for a given molecular cloud and its dependence on gas properties informed by Pineda et al.[95]. For the purpose of this work, we assume that the different isotopologues of CO are uniformly mixed throughout the medium, and consequently, both molecules will have identical excitation environments at any given line of sight, including densities and temperatures. Consider a generic emission line coming from a line of sight through molecular material with total optical depth τ_r. If there is no background intensity and the line profile is narrow compared to the frequency bandwidth of the observing instrument, then the observed intensity in a given band will be

$$I = \int_0^{\tau_r} \frac{j(\tau)}{\kappa(\tau)} e^{-\tau} d\tau, \tag{6.1}$$

where $j(\tau)$ is the emissivity and $\kappa(\tau) = n\sigma$ is the absorption coefficient for a molecule with number density n and absorption cross section σ. Because our observations come from only a fairly narrow redshift range, it is reasonable to assume that there is no significant effect from any background sources. This also means that the probability of seeing emission from two galaxies in the same frequency band along the same line of sight is negligible.

The emissivity is[36]

$$j = \frac{h\nu}{4\pi} n_u A_{ul}, \tag{6.2}$$

where A_{ul} is the Einstein A coefficient for the transition, ν is the rest frame emission frequency, and n_u is the number density of molecules in the upper state. We can rewrite this in terms of the total number density n as

$$j = \frac{3h^2\nu^2 A_{ul}}{8\pi k_B T_{\rm ex}} n e^{-h\nu/k_B T_{\rm ex}}, \tag{6.3}$$

where $T_{\rm ex}$ is the excitation temperature of the molecule and we have approximated the partition function as $Z \approx 2kT/h\nu$. This approximation is valid when $kT_{\rm ex} \gg 2h\nu$[95], and thus should hold in typical molecular cloud conditions.

If we change the integration variable in Equation (6.1) to column density we have

$$I = \frac{3h^2\nu^2 A_{ul}}{8\pi k T_{\rm ex}\sigma} e^{-h\nu/k_B T_{\rm ex}} \left(1 - e^{-N_r^l \sigma}\right), \tag{6.4}$$

where N_r^l is the total column density of molecules in the lower state, and we have assumed that the cloud properties are constant along the line of sight. The fraction of molecules in the lower state is given by

$$\frac{N_r^l}{N_r} = \frac{1}{Z} e^{-E_l/kT_{\rm ex}}. \tag{6.5}$$

where E_l is the energy of the ground state and N_r is the total column density.

Now consider a population of CO molecules with isotope ratio $R \equiv n_{13}/n_{12}$, which

in turn means that $N_{r13} = RN_{r12}$. The ratio of the line intensities is then

$$\frac{I_{13}}{I_{12}} = \left(\frac{\nu_{13}}{\nu_{12}}\right)^2 \left(\frac{A_{13}}{A_{12}}\right) \left(\frac{\sigma_{12}}{\sigma_{13}}\right) e^{-h(\nu_{13}-\nu_{12})/k_B T_{\rm ex}} \left[\frac{1-\exp\left(-R\sigma_{13}N_{r12}^l\right)}{1-\exp\left(-\sigma_{12}N_{r12}^l\right)}\right], \tag{6.6}$$

The cross sections σ_{12} and σ_{13} corrected for stimulated emission are given by

$$\sigma = \frac{3c^2 A_{ul}}{8\pi\nu^2}\left(1 - e^{-h\nu/k_B T_{\rm ex}}\right)\Delta\nu_{\rm FWHM}, \tag{6.7}$$

where $\Delta\nu_{\rm FWHM}$ is the frequency full width at half maximum of the line[36].

The Einstein A coefficients for the two transitions are[97] $A_{1-0}^{^{12}\rm CO} = 7.21 \times 10^{-8}$ s^{-1} and $A_{1-0}^{^{13}\rm CO} = 6.34 \times 10^{-8}$ s^{-1}. We assume that the molecular material has a constant excitation temperature $T_{\rm ex} = 20$ K, and that the width of the line is dominated by a turbulent velocity distribution with a velocity FWHM of 10 km/s, similar to the velocities of the largest local GMCs[98], which translates to $\Delta\nu_{\rm FWHM} = 3.83$ MHz for the ^{12}CO line and 3.67 MHz for the ^{13}CO line. This yields cross sections $\sigma_{12} = 3.67 \times 10^{-17}$ cm^2 and $\sigma_{13} = 3.55 \times 10^{-17}$ cm^2, and

$$\frac{I_{13}}{I_{12}} = 0.84\frac{1-\exp\left(-0.34R\sigma_{13}N_{r12}\right)}{1-\exp\left(-0.36\sigma_{12}N_{r12}\right)}, \tag{6.8}$$

where the factors of 0.34 and 0.36 in the exponentials come from the conversion between N_r^l and N_r.

If we expand this quantity in the limit of small optical depth, we find that I_{13}/I_{12}

goes close R at zero order. In the opposite limit, where all of the emission is optically thick, we find that I_{13}/I_{12} approaches 0.84. Up to the differences in atomic-scale physics between the two CO species, this agrees with the prediction from Section 2. Figure 6.2 shows the full behavior of this intensity ratio as a function of column density for different isotope ratios. The two extremes can clearly be seen, along with a transition region where the ^{12}CO line is optically thick but the ^{13}CO line remains optically thin. Again, this result is independent of the detailed distribution of molecular material. We note, however, the precise values of the intensity ratio will be dependent on the excitation temperatures of the medium along the line of sight, as well as the details of selective chemistry and photodissociation. For the purpose of forecasting, however, we neglect these effects for simplicity.

6.3.2 Power spectra

Each galaxy contributes a brightness temperature T_b to the map, where the value of T_b is drawn from a distribution $dn_{\rm gal}/dT_b$, which in turn depends on the galaxy luminosity function and the parameters of the intensity mapping instrument. The total number density of emitting sources is then $n_{\rm gal} = \int dn_{\rm gal}/dT_b dT_b$.

In addition to the usual auto-spectrum of each line, we can compute a cross power spectrum for pairs of frequency bands separated by $\Delta\nu_{12/13}$, the difference in observed frequencies of the ^{12}CO and ^{13}CO lines at a desired redshift. The first term in Equation 2.1 comes from correlations between pairs of sources, each weighted by

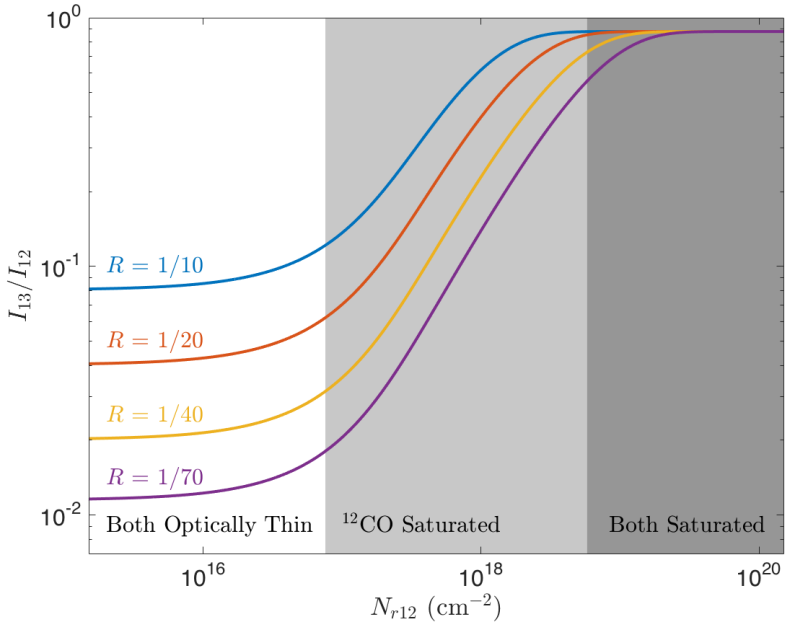

Figure 6.2: Intensity ratios between the two different CO lines as a function of ^{12}CO column density for different carbon isotope ratios. The black dashed line shows the column density where the ^{12}CO optical depth is unity, the colored dashed lines show where the ^{13}CO optical depth is unity for the four shown values of R.

the amount of intensity they produce. The shot noise term comes from the zero-lag correlation function, and thus essentially gives the correlation of each source with itself. For the case of ^{12}CO crossed with ^{13}CO, the two lines are each tracing the same population of sources. To compute the large-scale clustering term of a cross spectrum, we then simply need to weight one of each pair of sources by ^{12}CO intensity and the other by ^{13}CO intensity. The cross spectrum is then

$$P_{12\times13}(k,z) = \overline{T}_{b,12}(z)\overline{T}_{b,13}(z)\overline{b_{12}}(z)\overline{b_{13}}(z)P_m(k,z) + P_{\rm shot}^{12\times13}. \tag{6.9}$$

If we assume the ^{13}CO intensity is a function $T_{b,13}(T_{b,12})$ of the ^{12}CO intensity, as

implied by Figure 2, the new shot-noise term is given by

$$P_{\rm shot}^{12\times13} = \int_0^\infty T_{b,12}T_{b,13}(T_{b,12})\frac{dn_{\rm gal}}{dT_{b,12}}dT_{b,12}. \tag{6.10}$$

Since we are considering cross-correlations between different frequency bands of an intensity map, it may be more intuitive to again project down to two-dimensional power spectra, either within an individual band in the case of an auto-spectrum or between two bands in the case of a cross spectrum. As shown in Chapter 2, the angular power spectrum C_ℓ at a given multipole ℓ is given by

$$C_\ell = \frac{2}{\pi}\int k^2 P(k)\left[\int f_1(r_1)j_\ell(kr_1)dr_1\right]\left[\int f_2(r_2)j_\ell(kr_2)dr_2\right]dk, \tag{6.11}$$

The selection functions f_1 and f_2 are now set by the shape of the two frequency channels being correlated. Equation (6.11) is somewhat troublesome to evaluate exactly, so we will follow the approximations in Section 2.2 to simplify our computations.

For the purposes of this work, we will assume that the selection functions f_1 and f_2 take the form of top-hat functions with widths set by the instrumental frequency bandwidths. For a cross-correlation, these selection functions should be chosen so that ^{12}CO emitters at redshift z are included in f_1 and ^{13}CO emitters at redshift z are included in f_2, i.e. that f_1 and f_2 correspond to frequency bands separated by $\Delta\nu_{12/13}$. From here on we will assume for simplicity that the redshift ranges of f_1 and f_2 overlap exactly. This may not be the case for a real experiment, but this will

not alter our final results. Note also that f_1 will contain ^{13}CO emitters and f_2 will contain ^{12}CO emitters from outside the target redshift range. This should not be an issue, as the spatial separation between the ^{12}CO emitters in the two bands is large enough that they should be uncorrelated on the scales we care about. See Appendix A for a quantitative proof of this result.

6.4 Forecast

In the previous section we outlined the basic physical process which controls the emission of these two lines within a cloud and the final observables which will be computed from an intensity map. What we need now to forecast constraints on these observables is a model connecting the cloud-scale physics to the cosmological-scale power spectra. Given the large astrophysical uncertainties at these high redshifts, we will model this connection in a fairly simple manner. However, the advantage of the formalism presented above is that it holds regardless of exactly how the molecular gas is modeled. One could easily replace the calculations presented below with a more sophisticated model if one were made available.

At the redshifts we are considering, virtually all CO emission comes from within galaxies. Since our beam is large compared to an individual galaxy, we can model each galaxy as a point source with a given L_{12} and L_{13}. We can predict the ratio of

these luminosities by slightly modifying Equation (6.8):

$$\frac{L_{13}}{L_{12}} = 0.84 \frac{1 - \exp\left(-0.34 R \sigma_{13} \overline{N}_{r12}\right)}{1 - \exp\left(-0.36 \sigma_{12} \overline{N}_{r12}\right)}, \tag{6.12}$$

where $\overline{N}_{r12}$ is now the ^{12}CO column density averaged over all lines of sight through the galaxy. These luminosities can then be converted to brightness temperature following Equation (4.4).

Equation (6.12) relates the luminosity of a galaxy in the two CO lines to the properties of its molecular clouds. In order to compute power spectra from a population of such galaxies, we will use empirical results to compute a distribution of ^{12}CO luminosities, as well as a relation between ^{12}CO luminosity and molecular cloud density. With these tools in hand, we can predict a distribution of ^{13}CO luminosities and forecast a power spectrum for an intensity mapping experiment.

6.4.1 ^{12}CO Luminosity Function

The literature contains a wide variety of methods for modeling CO luminosity functions at high redshift, ranging from simple scaling arguments[62,31] to sophisticated semianalytic calculations[99,100]. We are more interested here in the relation between ^{12}CO and ^{13}CO rather than the exact behavior of the ^{12}CO line, so we will err on the

side of simplicity. We assume that the ^{12}CO luminosity function takes the form

$$\frac{dn}{dL_{12}} = \phi_* \left(\frac{L_{12}}{L_*}\right)^{\alpha} \exp\left(-\frac{L_{12}}{L_*} - \frac{L_{\rm min}}{L_{12}}\right), \tag{6.13}$$

which is simply a Schechter function with an exponential cutoff added to the low-luminosity end. This choice has the advantage that a wide variety of CO emission models can be expected to produce luminosity functions with similar shapes.

We choose values for the free parameters $(\phi_*, \alpha, L_*, L_{\rm min})$ by comparing to the results of Li et al.[35]. This model uses a simulated relationship between halo mass and SFR[101], then uses scaling relationships between SFR and far infrared luminosity and between FIR luminosity and CO luminosity to get L_{12} as a function of halo mass. Finally, they add lognormal scatter around this relation to obtain a luminosity function. If we fit our Schechter function to these results, we get best fit parameters $\phi_* = 2.8 \times 10^{-10}$ $({\rm Mpc/h})^{-3}$ $L_\odot^{-1}$, $\alpha = -1.87$, and $L_* = 2.1 \times 10^6$ $L_\odot$. The Li et al. model assumes a hard cutoff in ^{12}CO luminosity for halos smaller than 10^{10} $M_\odot$. We set the location of our low-luminosity cutoff at $L_{\rm min} = 500$ $L_\odot$, which is the luminosity in their model which corresponds to 10^{10} $M_\odot$.

In order to compute power spectra, we also need a prediction for the luminosity-weighted bias $\bar{b}$. In principle, the luminosity weighting means that the bias should take into account the full luminosity function model. However, in the interest of keeping things simple, we will follow literature models[77] in assuming for the purpose

of the bias calculation that luminosity is proportional to halo mass. The bias is then given by Equation (3.11). We again use the Tinker et al.[102] mass function dn/dM along with the corresponding mass-dependent bias $b(M)$[103]. When calculating the cross spectrum between ^{12}CO and ^{13}CO we will use the same bias factor for both lines. One could just as easily assume a more sophisticated weighting in Equation (3.11) and adjust it slightly to take into account a relation between L_{12} and L_{13}. Since the two lines trace the same population of galaxies, this would cause at most an order unity change in the final amplitudes of our power spectra, and likely much smaller. Using the Li et al. $L(M)$ relation to weight the bias, for example, only increases $\overline{b_{12}}$ by $\sim 5\%$. This is a very small effect, especially since the astrophysical uncertainties in the rest of the model are so very large.

6.4.2 Relating Luminosity to Column Density

In order to determine the distribution of ^{13}CO intensities, we would like to rewrite the ratio L_{13}/L_{12} from Equation (6.12) in terms of the ^{12}CO luminosity L_{12}. Doing so requires a relationship between average column density $\overline{N}_{r12}$ and L_{12}, which we will derive here using a similar set of scaling relations to those used to get the Li et al. model described above.

We start with the assumption that the average ratio $Z_{\rm CO}$ of ^{12}CO and H_2 column densities is roughly 10^{-4}, following Bolatto et al[26] based on the observations of Sofia et al.[104]. We can then convert to a mean surface star formation rate density Σ_{SFR}

using the Schmidt-Kennicutt law

$$\frac{\Sigma_{SFR}}{M_\odot \text{ kpc}^{-2} \text{ yr}^{-1}} = 2.5 \times 10^{-4} \left(\frac{\Sigma_{\text{H}_2}}{M_\odot \text{ pc}^2} \right)^{1.4} = 4.8 \times 10^{-32} \left(\frac{N_{\text{H}_2}}{\text{cm}^{-2}} \right)^{1.4}, \quad (6.14)$$

where we have assumed that the mass of the gas is dominated by H_2.

The star formation rate is then simply

$$SFR = \pi r_{\text{gal}}^2 \Sigma_{SFR}, \quad (6.15)$$

where we assume a representative galaxy radius $r_{\text{gal}} = 30$ kpc for all galaxies. We relate star formation rate to L_{12} using Model A of Pullen et al.[31], which uses the relation between SFR and far infrared luminosity from Kennicutt[105] and the relation between FIR luminosity and CO luminosity from Wang et al.[80] to obtain

$$\frac{L_{12}}{L_\odot} = 3.2 \times 10^4 \left(\frac{SFR}{M_\odot \text{ yr}^{-1}} \right)^{3/5}. \quad (6.16)$$

Combining all of this together yields

$$\frac{\overline{N}_{12}}{\text{cm}^{-2}} = X_L \left(\frac{L_{12}}{L_\odot} \right)^{1.2}, \quad (6.17)$$

Table 6.1: Survey parameters used for Fisher analysis.

Parameter	Pathfinder	Full	Future
Frequency range (GHz)	30-34	30-34	30-34
Patch Area Ω_s (deg^2)	2.5	6.25	30
Beam Size $\theta_{\rm FWHM}$ (arcmin)	6	3	2
Observing time/Patch $t_{\rm obs}$ (hr)	1500	2250	2250×5
Number of patches $N_{\rm patch}$	4	4	4
Sensitivity s (μK s$^{1/2}$)	1026	783	585
Channel width $\Delta\nu$ (MHz)	40	10	10
Number of channels $N_{\rm ch}$	100	400	400

where $X_L \equiv 3.5 \times 10^{10}$. The luminosity ratio of our two CO lines is then

$$\frac{L_{13}}{L_{12}} = 0.84 \frac{1 - \exp\left[-0.34 R \sigma_{13} X_L \left(L_{12}/L_\odot\right)^{1.2}\right]}{1 - \exp\left[-0.36 \sigma_{12} X_L \left(L_{12}/L_\odot\right)^{1.2}\right]}, \tag{6.18}$$

These luminosities then can be converted to intensities using Equation (4.4).

6.4.3 Experimental Parameters

When forecasting the constraining power of the measurements discussed here, we will consider three different surveys: the COMAP "Pathfinder" and "Full" surveys described by Li et al.[35] and a hypothetical "Future" survey which has improved sensitivity, resolution, and additional observing time. The parameters we use for these three instruments are given in Table 6.1. Construction on the "Pathfinder" survey is currently underway, the "Full" experiment is planned as a next step. The "Future" experiment assumes a modest increase in sensitivity and angular resolution, and either additional observing time or an array of several dishes observing simultaneously.

6.4.4 Fisher analysis

Combining the results from previous sections assuming the frequency bandwidth of the "Full" or "Future" experiments yields the four power spectra shown in Figure 6.3. We assume a fiducial CO isotopologue ratio $R = 1/70$, which is representative of the many local and extragalactic measurements[94,106,107].

The blue and orange dashed curves show the auto-spectra for the ^{12}CO and ^{13}CO lines respectively. Neither of these are observable on their own, as correlating any given band with itself will simply give the sum of the two, which is shown in green. This total auto-spectrum is essentially just the ^{12}CO spectrum, with a correction of order a few percent from the ^{13}CO emitters. The red curve shows the cross spectrum between two bands separated by $\Delta\nu_{12/13}$. Note that the quantity plotted on the y-axis of Figure 6.3 is C_ℓ rather than the commonly seen $\ell(\ell+1)C_\ell/(2\pi)$. The error σ_ℓ at a given multipole is given by Equation (3.18). The shaded region in Figure 6.3 shows the error on the cross-spectrum assuming the parameters of the "Full" experiment with multipoles binned in sets of 10.

We stated previously that the cross-spectrum of these two CO lines is interesting because it depends on both the relative abundances of the two carbon species and the distribution of molecular gas densities at a given redshift. The abundance information can be trivially parameterized using the isotopologue ratio R, which depends on the nucleosynthesis history. We parameterize the gas density distribution by defining the

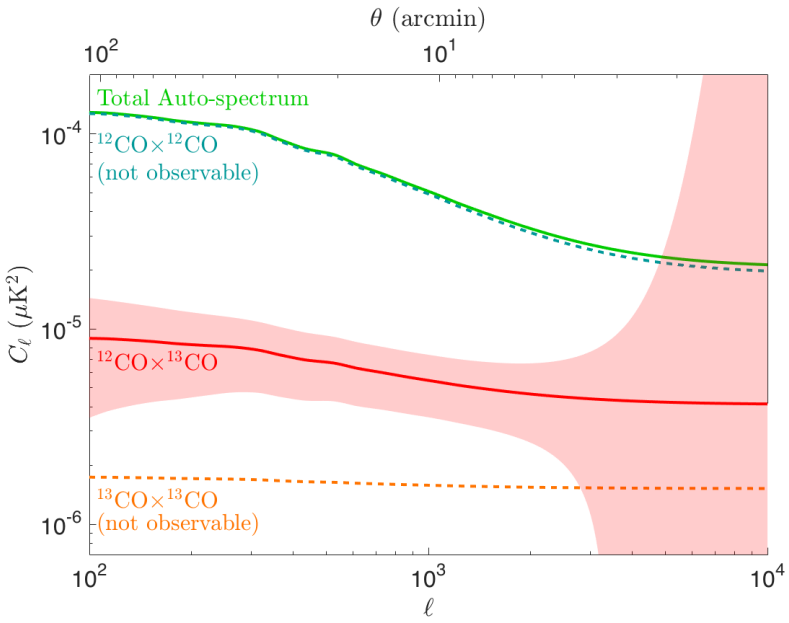

Figure 6.3: Angular power spectra for different combinations of the two CO lines. Dashed lines show auto-spectra of ^{12}CO (blue) and ^{13}CO (orange), which are not independently observable. Solid lines show the two observable spectra, the auto-spectrum of a single frequency band (green) and the cross-spectrum of two bands separated by $\Delta\nu_{12/13}$ (red). The red shaded region shows the instrumental error on the cross-spectrum assuming the parameters of the "Full" experiment and multipoles binned in sets of 10.

quantity

$$f_s \equiv \frac{\int_{T_\tau}^{\infty} T_{b,12} dn_{\rm gal}/dT_{b,12} dT_{b,12}}{\int_0^{\infty} T_{b,12} dn_{\rm gal}/dT_{b,12} dT_{b,12}}, \tag{6.19}$$

which is the fraction of the measured ^{12}CO emission coming from optically thick lines of sight. In order to constrain f_s, we treat the quantity X_L in Equation (6.17) as a free parameter. The ^{12}CO line becomes optically thick when $\overline{N}_{12}\sigma_{12} = 1$, or when $\overline{N}_{12} = 6 \times 10^{16}$ cm^{-2}. From a given value of X_L we can compute an intensity T_τ from Equations (6.17) and (4.4) which corresponds to this column density, then use Equation (6.19) to compute a value of f_s. Our fiducial value for X_L corresponds to a saturated fraction $f_s = 0.36$.

No relevant constraints on these quantities currently exist at these redshifts. The parameters R and f_s as we have defined them here are values averaged over all of the molecular gas in a given redshift slice. Since only an intensity mapping survey can access the vast majority of the galaxy population, these quantities cannot be realistically constrained by existing data. This also means that even weak constraints from an intensity map hold great scientific value.

We estimate constraints on R and f_s using the Fisher matrix formalism[108,109,110]. This assumes that the likelihood distribution of the parameter values is a multivariate Gaussian centered on the fiducial values. The Fisher matrix $F_{\mu\nu}$ for a model with parameters p_μ is given by

$$F_{\mu\nu} = \sum_\ell \frac{1}{\sigma_\ell^2} \frac{\partial C_\ell}{\partial p_\mu} \frac{\partial C_\ell}{\partial p_\nu}, \tag{6.20}$$

where C_ℓ is the power spectrum we are analyzing. For simplicity, we use the power spectrum computed at the central redshift $z = 2.6$ of our survey for all frequency bins and stack the contribution from each. The Fisher matrix computed from Equation (6.20) can then be inverted to obtain the covariance matrix of the model parameters.

Figure 6.4 shows the results of our Fisher analysis for our three experiments. Dark colors show 1-σ constraints, light colors show 2-σ constraints. The "Pathfinder" survey only yields a signal-to-noise ratio of ~ 1 for the cross-spectrum, so it effectively provides an upper limit on the cross-correlation signal. The "Full" survey does better, however the constraints on the two parameters are quite degenerate, leading to order unity 1-σ fractional errors on both parameters. Despite this degeneracy, the volume of parameter space is still dramatically reduced compared to the complete lack of constraints currently available. Fractional errors on our two parameters fall to $\sim 30\%$ for the "Future" experiment.

For these constraints, we have assumed that the ^{12}CO luminosity function is known exactly. This may seem like a poor assumption, since the cross-spectrum depends equally on the properties of both lines. However, the auto-spectrum, which is essentially the ^{12}CO spectrum, will be measured as well. The auto-spectrum has an amplitude roughly 10 times greater than the cross-spectrum, so the remaining uncertainty in the cross-spectrum from ^{12}CO emission should be very small compared to that from ^{13}CO emission, especially if things like the one-point PDF of the map (see Chapter 7) are used to further constrain dn/dL_{12}. For more detail, see Appendix B.

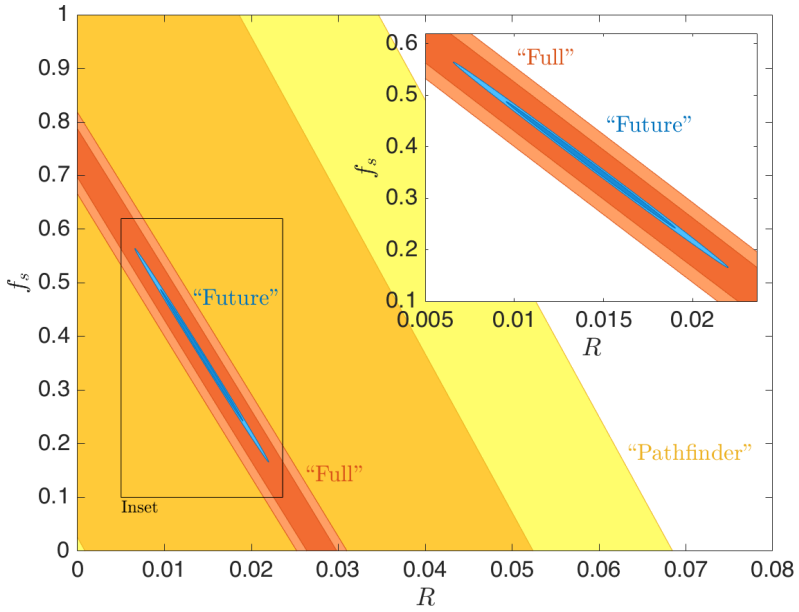

Figure 6.4: Forecasted 1- and 2-sigma constraints (dark and light respectively) on isotopologue ratio R and saturated fraction f_s for the "Pathfinder" (yellow), "Full" (red), and "Future" (blue) experiments.

6.5 Discussion

Using intensity mapping, we have demonstrated that it is possible to constrain both carbon isotopologue ratios and molecular gas density distributions at cosmological distances. These measurements cannot be made using any other method, as even the deepest targeted survey would only be able to study a handful of the brightest CO emitters[35]. Intensity mapping surveys will allow measurement of these quantities from the entire galaxy population. Using the cross-correlation techniques discussed here also allows significant improvement beyond what is possible using intensity mapping of only a single spectral line, allowing for a much more in-depth probe of the molecular gas which feeds star formation.

As seen in Figure 6.4, constraints on the parameters R and f_s are quite degenerate for all three of our surveys. This degeneracy arises because, as shown in Figure 6.3, most of the signal-to-noise comes from the clustering component of the signal, which depends only on the mean intensities of the two lines. Using the mean intensity alone cannot distinguish between a higher ^{13}CO abundance or more saturated ^{12}CO emission. This degeneracy can be broken somewhat by a measurement of the shot noise. Indeed, a large part of why the "Future" survey results in a weaker degeneracy is because it has much more constraining power in the shot-noise component of the spectrum. It may thus be possible to reproduce some of the extra constraining power of the "Future" experiment by combining the "Full" experiment with an interferometric intensity mapping experiment. Interferometers such as the one used in the COPSS[34] survey are sensitive to considerably smaller scales than single-dish COMAP-like experiments. This makes them well suited to studying the shot-noise dominated regime.

One could also potentially further improve these constraints by combining an intensity map with traditional galaxy survey data. Even if only the brightest end of the luminosity function were constrained through these surveys, that could still help to break this degeneracy. An in-depth study of quasar lines could accomplish the same purpose. Neither of these measurements can hope to do as well as our intensity mapping forecasts on its own, but they could provide valuable complements. If the degeneracy between R and f_s were completely broken, the fractional uncertainties on

these quantities in our forecast for the "Full" survey would drop from order unity to a few percent. It is probably overly optimistic to assume perfect degeneracy breaking, but the potential for dramatic improvement is clear.

For the above forecasts, we constrained the properties of the CO emission averaged over the entire survey volume. However, for a sufficiently sensitive survey, one could break the volume down into multiple redshift bins and constrain the redshift variation of these quantities. One could then compare to other redshift-dependent measurements, such as the star-formation history[89]. We used a "Future" survey targeting the same redshift range as the planned COMAP "Pathfinder", but there is no reason why such a futuristic experiment could not be extended to a broader redshift range. In addition, some have proposed conducting CO intensity mapping surveys at $z \sim 6 - 10$ as a probe of reionization physics[30]. At these redshifts, we begin to run into the limits of what is known about molecular gas, the formation of carbon and consequently CO, and the formation of the earliest generations of stars where metallicity is limited. With sufficient sensitivity, one could apply the cross-correlation methods applied here to understanding this mysterious era in cosmic history.

It is important to note here that the forecasts given in this work are based on deliberately simplified modeling. We have neglected a number of effects here which deserve more consideration in future work. For example, our model assigns a ^{13}CO luminosity to a given galaxy based entirely on its average ^{12}CO column density, which in turn is determined solely by its ^{12}CO luminosity. In reality, a galaxy will contain

a large number of heterogeneous molecular clouds with different densities, leading to a more complicated relation between the two CO lines. In the short term, this could perhaps be taken into account by simply assuming a scatter in the luminosity relations[35]. The basic framework we have outlined here can be straightforwardly generalized to more sophisticated models as we gain a better understanding of the environments in these distant galaxies.

The interpretation of this type of measurement may also be complicated by the fact that the metallicity of high-redshift systems could deviate significantly from solar. It is well-established that the molecular gas to CO ratio increases in lower-metallicity systems[26]. This fact must be taken into account when attempting to constrain global molecular gas properties from these observations, and will add additional uncertainty. In addition, low metallicity may alter the correlation between molecular gas and star formation. In such systems, star formation can proceed in regions which remain dominated by atomic gas. These regions will produce much less CO emission than would be expected from star formation at solar metallicity[111,112].

We have also neglected the effects of various foregrounds in this analysis. However, we do not expect this to have a substantial impact on the results in this case. There are two types of foregrounds which affect intensity mapping surveys. The first are those with continuum spectra, such as Milky Way dust and synchrotron emission. These foregrounds have been extensively studied for the case of 21 cm reionization experiments[12], which have a considerably higher foreground-to-signal ratio than CO

experiments. Since these foregrounds have little spectral structure, they only contribute substantially to Fourier modes which fall close to the line-of-sight, and can thus be cleaned simply by subtracting these modes from an analysis[33]. The second type of foreground are spectral lines from other redshifts which fall into our target frequency bands. As mentioned previously, ^{13}CO can be thought of as a line foreground to ^{12}CO, though not an important one. A sufficiently bright line can in principle contaminate the CO auto-spectrum (see Chapter 5), but unless there exists a pair of foreground lines separated by exactly $\Delta\nu_{12/13}$ there should be no contamination to the cross-spectrum we looked at here.

As mentioned previously, intensity mapping surveys are planned in many other lines besides CO, such as [CII], Lyα, and the 21 cm HI line. The broad strokes of the work we have presented here could easily be applied to cross-correlation between CO and these other lines. Such cross-correlations would require additional planning compared to our CO isotope correlation, since both lines would not appear in the same survey. However, if experiments were planned well to target the same volumes, one could potentially learn about high-redshift galaxies in even greater detail, allowing us to study the complex ecology of gas dynamics and star formation across the entire history of the universe.

6.6 Summary

We have demonstrated here a method whereby we can dramatically improve our understanding of molecular gas at high redshifts by combining information from ^{13}CO in intensity maps with the usual ^{12}CO. This is a direct extension of similar techniques used when studying local molecular clouds. By cross-correlating properly chosen slices of a CO intensity map it is possible to determine the total amount of ^{13}CO emission and how it varies with ^{12}CO emission. We showed how the cross-spectrum of these two lines can be used to constrain not only the abundance ratio of these two species but also the density distribution of molecular gas in the mapped galaxy population, quantities which are extremely difficult if not impossible to measure with any other method. This technique will allow us to gain deep insights into the processes that feed star formation throughout cosmic history. By branching out to more detailed models, additional spectral lines, and broader redshift ranges we can study the complex ecology of star formation and galaxy evolution at a level of detail unimaginable with traditional methods. We have likely only scratched the surface of what intensity mapping can teach us about the distant universe.

Chapter 7

One-Point Statistics of Intensity Maps

Up to this point, we have considered intensity maps entirely in terms of their power spectra. However, as described in Chapter 4, the brightness fluctuations in these maps are highly non-Gaussian. For non-Gaussian fields, the power spectrum, a two-point statistic, leaves out a significant amount of information about the higher-order moments of an intensity distribution. This limits the ability to constrain line luminosity functions with power spectra alone. In the CMB, non-Gaussianities are typically probed using higher-point statistics such as the bispectrum and trispectrum[113]. Unfortunately, these statistics are rather difficult to work with, both from a theoretical and observational perspective. We propose in this chapter instead to study the one-point statistics of intensity maps using a quantity we will refer to as

the voxel intensity distribution, or VID. This VID statistic is the probability distribution function of observed intensities in individual voxels. It can be predicted in a straightforward manner from a model luminosity function, and it can be estimated from a map simply by making a histogram of the observed intensity values.

Our VID method is an extension of a technique known as probability of deflection, or $P(D)$ analysis, which is a general method for predicting observed intensities from confusion-limited populations. $P(D)$ analysis was originally developed for radio astronomy[114], but has since been applied to observations ranging from gamma rays[115,116] to X-rays[117] to the optical[118] to the submillimeter[119,120]. Since intensity maps provide deliberately confused observations of galaxy populations, they are good candidates for $P(D)$ analysis. Here, we study this method in detail with the goal of creating a procedure that can be readily applied to many different intensity-mapping surveys targeting different lines. We will again demonstrate this method using a hypothetical CO intensity mapping survey.

This chapter is primarily based on work published in Ref.[49]. Section 7.1 contains a discussion of the power spectrum and its limitations along with the presentation of our VID formalism. Section 7.2 describes the CO emission model we will use in this chapter, which we apply in Section 7.3 to demonstrate the constraining power of the VID. Section 7.4 investigates how contamination from continuum emission, interloper lines, and gravitational lensing effects our constraints. We discuss our results in detail in Section 7.5 and conclude in Section 7.6.

7.1 $P(D)$ Formalism

The power spectrum is a useful statistic when studying cosmological density fields, but it suffers from one key limitation. All of the information about a random field is contained within its power spectrum if and only if the field is perfectly Gaussian. However, we expect the small-scale fluctuations in an intensity map to be highly non-Gaussian, as the measured intensity is the product of highly nonlinear processes within the galaxy population. Thus, the power spectrum alone misses out on much of the information content of a map.

This can be easily seen by looking at Equations (2.5) and (2.6). The power spectrum depends only on the first two moments of the luminosity function $\Phi(L)$. No higher moments can be measured from this statistic. As with CMB measurements, higher moments may be measurable using higher order n-point statistics such as the bispectrum, but these are computationally difficult[121]. One could also obtain additional information through cross-correlations of different lines (see Chapter 6), but while this allows the study of different properties of the galaxy distribution, it adds little information about the initial target line. A much more straightforward method would be to consider instead the one-point statistics of an intensity map through the VID statistic, which we will now derive. Our derivation is based on the $P(D)$ computation presented by Lee et al.[115], modified somewhat to include the effects of large-scale clustering.

Consider a volume of space at redshift z containing a population of point sources

emitting a line with rest frequency $\nu_{\rm em}$ and luminosity function $\Phi(L)$. If we divide our space into voxels with volume $V_{\rm vox}$, then the observed intensity in a given voxel is given by Equation 4.4. We neglect beam effects here, i.e. a source contributes all of its intensity to the voxel it is contained within, and there is no smoothing effect spreading the intensity over multiple voxels.

If we consider only voxels that contain exactly one emitter, the probability[1] $\mathcal{P}_1(T)$ of observing intensity T is given by

$$\mathcal{P}_1(T) = \frac{V_{\rm vox}}{\overline{n} X_{LT}} \Phi\left(T V_{\rm vox}/X_{LT}\right), \tag{7.1}$$

(compare with Equation 4.4) where

$$\overline{n} = \int_0^\infty \Phi(L) dL, \tag{7.2}$$

is the average comoving number density of sources. The probability of observing T in a voxel with two sources is then

$$\mathcal{P}_2(T) = \int\int \mathcal{P}_1(T')\mathcal{P}_1(T'')\delta_D(T - T' - T'')dT'dT'', \tag{7.3}$$

[1]For clarity, we will use the symbol P when referring to power spectra and $\mathcal{P}$ when referring to probability distributions.

where δ_D is the Dirac delta function. This simplifies to the convolution

$$\mathcal{P}_2(T) = \int \mathcal{P}_1(T')\mathcal{P}_1(T - T')dT' = (\mathcal{P}_1 * \mathcal{P}_1)(T). \tag{7.4}$$

From this it is clear that the probability of observing T in a voxel with N sources is simply

$$\mathcal{P}_N(T) = (\mathcal{P}_{N-1} * \mathcal{P}_1)(T). \tag{7.5}$$

An empty voxel obviously will always give zero intensity, so $\mathcal{P}_0(T) = \delta_D(T)$.

With Equations (7.1) and (7.5) we can recursively compute these probability distributions for voxels containing any arbitrary number of sources. The full VID is then given by

$$\mathcal{P}(T) = \sum_{N=0}^{\infty} \mathcal{P}_N(T)\mathcal{P}(N), \tag{7.6}$$

where $\mathcal{P}(N)$ is the probability of observing a voxel that contains N sources. It is possible[115] to replace the series of convolutions used to calculate $\mathcal{P}_N(T)$ with products in Fourier space. However, as luminosity functions often span many orders of magnitude, it may be computationally easier to compute $\mathcal{P}_N(T)$ using convolutions.

If the sources are unclustered, then $\mathcal{P}(N)$ is the Poisson distribution $\mathcal{P}_{\rm Poiss}(N, \overline{N})$ with mean $\overline{N} = \overline{n}V_{\rm vox}$. Equation (7.6) can then be simplified further, as shown by Lee et al.[115]. However, intensity maps typically cover large enough volumes that large-scale clustering will cause their number count distribution to differ significantly

from Poisson. We must therefore modify our computation to take this clustering into account. One method for $P(D)$ analysis of clustered sources is given by Equation (21) of Barcons[122] (see also Takeuchi & Ishii[123]). However, this method requires knowing all of the higher N-point statistics of the intensity distribution, which makes computing the VID intractable.

We instead make use of the fact that the galaxy number-count distribution is known to be approximately lognormal[72]. Using this fact, we can follow the same procedure we used in Chapter 4 and assume that for each voxel there is an expectation value μ for the number of galaxies contained within it which depends on the value of the lognormal cosmic density field at that point. The observed number of galaxies within that voxel will then be a Poisson draw from a distribution with mean μ. We can then write $\mathcal{P}(N)$ as

$$\mathcal{P}(N) = \int_0^\infty \mathcal{P}_{LN}(\mu)\mathcal{P}_{\rm Poiss}(N,\mu)d\mu, \tag{7.7}$$

where $\mathcal{P}_{LN}$ is lognormal probability of finding a voxel with expectation value μ.

The lognormal probability $\mathcal{P}_{LN}$ is computed assuming that the galaxy density field has density contrast in $\delta_{LN}(\vec{x}) = [\mu(\vec{x}) - \overline{N}]/\overline{N}$ in a voxel located at $\vec{x}$. We can write δ_{LN} in terms of a Gaussian random variable δ_G as

$$1 + \delta_{LN} = \exp\left(\delta_G - \frac{\sigma_G^2}{2}\right), \tag{7.8}$$

where σ_G^2 is the variance of the Gaussian random field. We can then write[124] $\mathcal{P}_{LN}$ as

$$\mathcal{P}_{LN}(\mu) = \frac{1}{\mu\sqrt{2\pi\sigma_G^2}} \exp\left\{-\frac{1}{2\sigma_G^2}\left[\ln\left(\frac{\mu}{\bar{N}}\right) + \frac{\sigma_G^2}{2}\right]^2\right\}. \tag{7.9}$$

The quantity σ_G sets the overall "strength" of the clustering, i.e. fields with a larger σ_G have comparatively more voxels containing very many or very few sources, and comparatively fewer "mid-range" voxels. We can compute it from the power spectrum $P_G(k)$ of δ_G using

$$\sigma_G^2 = \int P_G(k) \left|W(\vec{k})\right|^2 \frac{d^3\vec{k}}{(2\pi)^3}, \tag{7.10}$$

where $W(\vec{k})$ is the Fourier transform of the voxel window function. The spectrum $P_G(k)$ can be calculated using the fact that the real-space correlation functions $\xi(r)$ of δ_{LN} and δ_G are related by[72]

$$\xi_G = \ln\left[1 + \xi_{LN}(r)\right], \tag{7.11}$$

and the power spectra and correlation functions are related in the usual way. We thus only need to assume a power spectrum for our lognormal field to compute σ_G. Here, we assume that this spectrum is given by $P_{\rm LN}(k) = \bar{b}^2 P_m(k)$ calculated using Equations (2.1) and (2.7). Note that since the bias $\bar{b}$ used here is luminosity weighted, this is slightly different from the power spectrum typically used for galaxy surveys.

In a realistic experiment, many of the fluctuations observed in a map will be caused by instrumental noise. The measured intensity in a given voxel will then be the sum of the signal and noise contributions. The noise will have its own VID $\mathcal{P}_{\rm Noise}(T)$ which is determined by the instrumental properties. For example, in the case of simple Gaussian noise the noise VID is

$$\mathcal{P}_{\rm Noise}(T) = \frac{1}{\sqrt{2\pi\sigma_N^2}} \exp\left(-\frac{T^2}{2\sigma_N^2}\right), \tag{7.12}$$

where the variance σ_N is set by the survey sensitivity. By the same arguments used in Equation (7.4), the VID for the sum of signal and noise is

$$\mathcal{P}_{\rm Total}(T) = \left(\mathcal{P}_{\rm Signal} * \mathcal{P}_{\rm Noise}\right)(T). \tag{7.13}$$

Contributions from other sources of contamination, such as line or continuum foreground emission, can be added to the VID in a similar fashion.

To summarize the above formalism, the steps to compute a VID for a given model are as follows:

- Assume a luminosity function $\Phi(L)$, a voxel shape, and a cosmological model.
- Compute the probability $\mathcal{P}_N(T)$ of observing T in a voxel containing N sources from $\Phi(L)$ using Equation (7.5).
- Determine the mean number $\overline{N}$ of galaxies/voxel using the assumed $\Phi(L)$ and

voxel dimensions.

- Compute a power spectrum for the lognormal galaxy field from the assumed cosmological model.
- Calculate σ_G from this power spectrum using Equations (7.10) and (7.11).
- Use Equation (7.7) to compute the probability $\mathcal{P}(N)$ of observing a voxel with N galaxies using the calculated $\overline{N}$ and σ_G.
- Sum $\mathcal{P}_N(T)\mathcal{P}(N)$ over all values of N as in Equation (7.6).
- Convolve the resulting VID with VIDs computed for instrumental noise and any foreground contamination.

Note that this method as presented here makes a subtle approximation about the halo bias. By including the bias in our chosen $P_{LN}(k)$ spectrum, we take into account the fact that galaxies are more strongly clustered than the underlying dark matter. However, putting it into the model in this manner effectively assigns each galaxy the average bias value, when in reality the brightest galaxies should be more strongly clustered. The model therefore underestimates the number of very bright voxels. Fortunately, we expect the effect of this to be small, especially given the immense uncertainties that currently exist in the modeling of $\Phi(L)$ for most intensity mapping lines. For more discussion of this effect, see Appendix C.

7.2 Fiducial CO Model

The formalism described above can be readily applied to a wide variety of different lines and models. In order to demonstrate its effectiveness, we will now without loss of generality apply the VID to a model of a CO intensity map. Literature models for CO emission span a wide range of signal amplitudes (see Chapter 3), but most share several essential features. Luminosity functions tend to have a roughly power-law shape for moderate luminosities. Reduced star formation efficiency in high-mass halos introduces a cutoff in $\Phi(L)$ at high luminosities, and the difficulty of forming galaxies in very low-mass halos creates a cutoff at low luminosities. To capture this behavior in the simplest possible form we again model the CO luminosity function as a slightly modified Schechter function[65],

$$\frac{\Phi(L)}{(\mathrm{Mpc}/h)^{-3}\ L_\odot} = \phi_* \left(\frac{L}{L_*}\right)^{\alpha} \exp\left(-\frac{L}{L_*} - \frac{L_{\mathrm{min}}}{L}\right), \tag{7.14}$$

where $p_i \equiv (\phi_*, \alpha, L_*, L_{\mathrm{min}})$ are free parameters. The quantity L_{min} here serves the purpose of the hard minimum mass or luminosity seen in most literature models, but replaced with an exponential cutoff both for added realism and to prevent numerical issues which can arise around hard cutoffs.

We choose values for our four parameters by fitting this Schechter model to the luminosity function plotted in Figure 8 of Li et al.[35], in a similar fashion to the procedure in Chapter 6. This luminosity function is calculated from a suite of N-

body simulations using the relation between halo mass and star formation rate[101]. Star formation rates are then connected to CO luminosity through a series of empirical scaling relations[105,27]. The best fit Schechter function has parameters $\phi_* = 2.8 \times 10^{-10}$, $\alpha = -1.87$, and $L_* = 2.1 \times 10^6\ L_\odot$. The Schechter function fits the Li et al. results reasonably well, though it does produce a steeper high-luminosity cutoff. Since essentially nothing is known currently about CO emission from very faint galaxies, we somewhat arbitrarily choose $L_{\rm min} = 5000\ L_{\rm sun}$.

As for the bias, If we assume the linear relation between halo mass and CO luminosity as inChapter 3, this $L_{\rm min}$ corresponds to a halo mass of 2.5×10^9 solar masses. This is comparable to literature values that usually place the CO luminosity cutoff around $10^9 - 10^{10}\ M_\odot$. When computing the mean bias $\bar{b}$, we use this same linear mass-luminosity relation along with the Tinker form of the mass function and $b(M)$[102,103].

7.3 Constraining Power

We will now forecast the constraining power of a COMAP-like experiment for our CO emission model. The parameters of the full COMAP experiment can be found in Table 6.1. The planned survey would target the 115 GHz CO(1-0) line in 400 bands between $z = 2.4$ and $z = 2.8$, giving each band a width $\delta\nu = 10$ MHz. The planned instrument has an angular resolution of 3 arcminutes with a total survey area of 6.25

deg^2. Based on these parameters, we assume a model voxel which is a $3' \times 3'$ square in the plane of the sky with a comoving depth set by the 10 MHz frequency bandwidth. Our choice to set the voxel angular scale at scale of the beam FWHM matches the optimization performed by Vernstrom et al.[125], who found that larger voxels than this increasingly smooth out small-scale structure while smaller voxels induce excess correlations between voxels due to beam effects. The COMAP survey aims for a noise/voxel of 5.8 μK, so we adopt a Gaussian noise VID with $\sigma_N = 5.8$ μK.

Note that, given the large uncertainty on the amplitude of the CO signal, the constraints presented here are model dependent, with brighter-on-average models yielding better constraints, and vice versa. However, because our model captures the essential shape, if not necessarily the amplitude, of the CO luminosity function, the qualitative behaviors presented below should hold for a wide variety of CO models.

Figure 7.1 shows the VIDs for our fiducial model both with and without noise. The dashed curve shows $\mathcal{P}_1(T)$, which is a simple rescaling of our fiducial Schechter function. The amplitude of the VID is reduced from that of $\mathcal{P}_1(T)$ because the expected number of galaxies per voxel is only ~ 0.3 in this setup, leading to a significant number of voxels that contain zero galaxies. This creates a delta function in $\mathcal{P}(T)$ at $T = 0$ (not shown). This suppression is less for brighter intensities though because some of the difference is made up by voxels that contain several sources. These multiple-source voxels cause the VID to deviate a modest amount from the Schechter power law in the middle of the distribution. The effect of the instrumental noise is

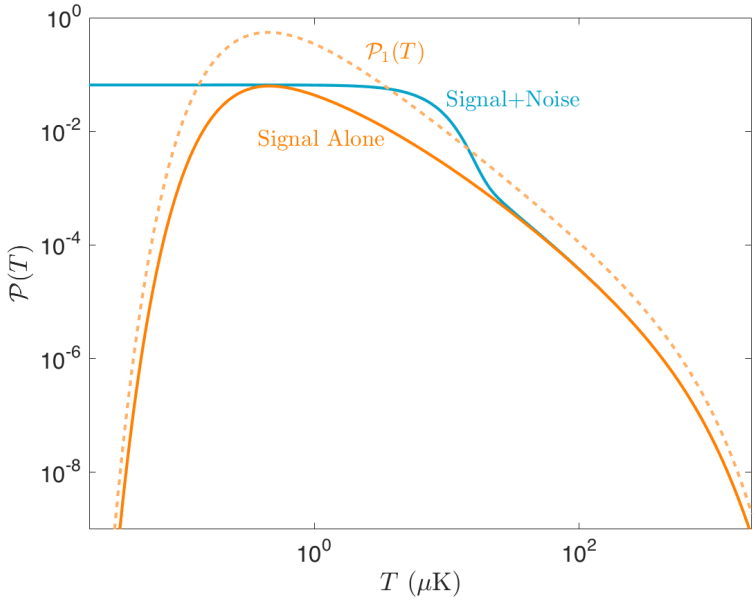

Figure 7.1: VIDs for our fiducial model and experiment with (orange) and without (blue) instrumental noise. Also plotted is the probability distribution $\mathcal{P}_1(T)$ for voxels that are known to contain exactly one source (dashed orange). This distribution is simply a rescaling of our Schechter luminosity function.

to remap all of the faint voxels into a Gaussian distribution, with the signal VID dominating the bright end of the distribution. These calculated VIDs agree well with simulated VIDs prepared using the method from Chapter 4, as shown in detail in Appendix D.

In order to study the constraining power of our example CO model, we will again make use of the Fisher matrix, which for a VID parameterized by p_μ takes the form

$$F_{\mu\nu} = \sum_i \frac{1}{\sigma_i^2} \frac{\partial B_i}{\partial p_\mu} \frac{\partial B_i}{\partial p_\nu}, \tag{7.15}$$

where

$$B_i = N_{\rm vox} \int_{T_{{\rm min},i}}^{T_{{\rm max},i}} \mathcal{P}(T) dT \tag{7.16}$$

is the number of voxels in an intensity bin with edges $[T_{{\rm min},i}, T_{{\rm max},i}]$ for a survey containing $N_{\rm vox}$ voxels. The expected variance σ_i^2 on B_i is assumed to be equal to B_i, i.e. we assume the bins obey Poisson statistics. This should be a reasonable assumption for bins that contain many voxels, which are the bins that will contribute most of the signal-to-noise. For simplicity, we neglect the evolution of the VID across the redshift range of the survey.

We compute the Fisher matrix using five parameters: the four parameters of our model luminosity function and σ_G. Adding σ_G as a free parameter takes into account our uncertainty on both the average bias and the clustering behavior of the target galaxies. Since our goal is to measure the luminosity function, we will report constraints on the four Schechter parameters marginalized over σ_G. The result of this analysis for our fiducial model and experiment are shown in Figure 7.2. The smaller orange ellipses show the constraints obtained for an ideal measurement with zero instrumental noise, with the only errors due to sample variance. For an instrument with infinite sensitivity, these errors could be further reduced by observing a larger area of the sky. The larger blue ellipses show the effect of the Gaussian COMAP-like instrumental noise. Adding noise obviously makes the constraints somewhat worse, but there is still a substantial amount of constraining power. The constraint on the low-luminosity cutoff $L_{\rm min}$ is hardest hit, as the cutoff occurs well below the intensity

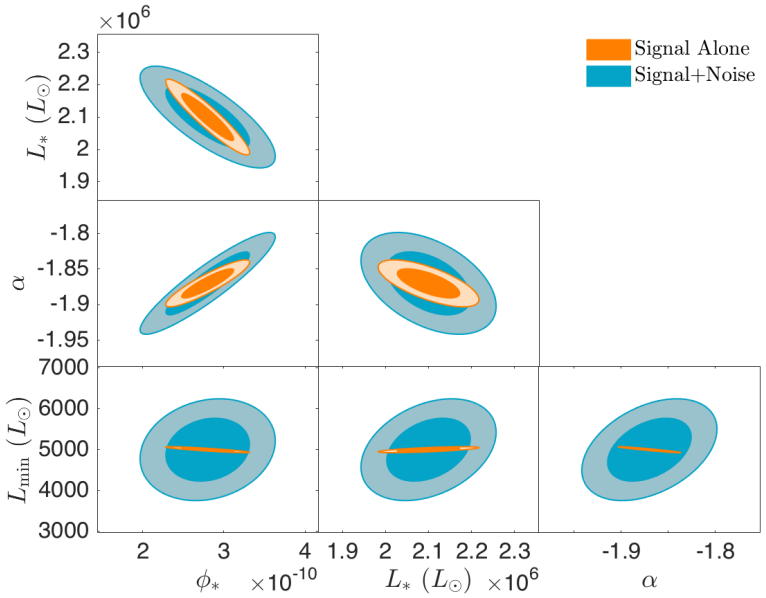

Figure 7.2: Fisher constraints on the four parameters of our fiducial CO luminosity function both with (blue) and without (orange) Gaussian instrumental noise. Dark ellipses show 1-σ errors, light ellipses show 2-σ errors.

where noise begins to dominate.

Using equation (7.14), we can convert these constraints on the model parameters into constraints on $\Phi(L)$. The resulting 95% confidence region for the case with COMAP Full instrumental noise are shown in blue in Figure 7.3. We can compare these uncertainties to those forecasted using the COMAP power spectrum[35]. If we apply the same fractional uncertainties plotted in Figure 8 of Li et al.[35] to our fiducial $\Phi(L)$, we get the 95% confidence region plotted in grey. Because the VID statistic is much better suited to measuring the luminosity function, it leads to substantially better constraints than the power spectrum despite using the same instrumental setup. It should be noted that these constraints only hold if the real luminosity function has

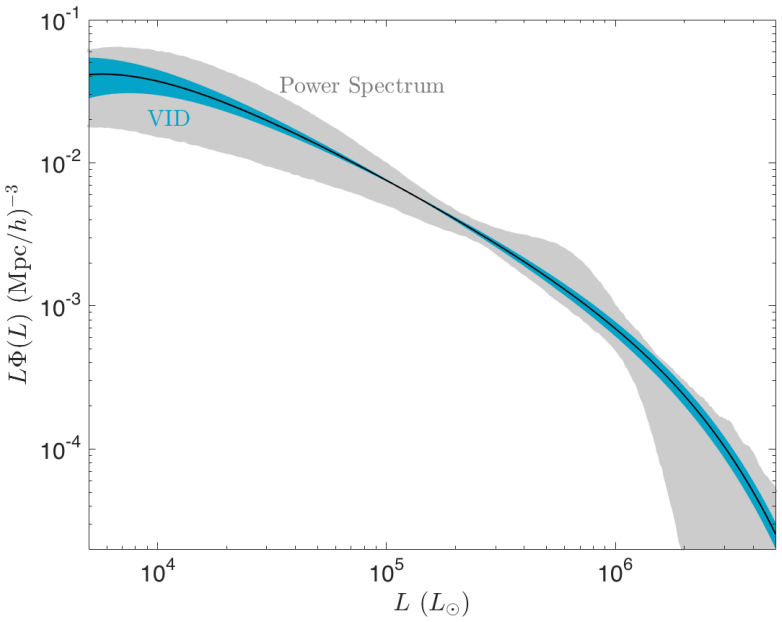

Figure 7.3: 95% confidence regions around our fiducial $\Phi(L)$ obtained from the above VID constraints including COMAP instrumental noise (blue) and from the COMAP power spectrum analysis[35] (gray).

exactly the form given in equation (7.14). For this reason, computing uncertainties on $\Phi(L)$ in this manner likely underestimates the true error, especially at the faint end where the VID is noise dominated.

7.4 Foreground Effects

The results shown above assume that the measured fluctuations in an intensity map are caused only by the source galaxies and instrumental noise. However, as when studying the power spectrum, there are a number of effects that can alter the observed intensity fluctuations and reduce the constraining power of a given experiment. Below we provide examples of how three of these effects, continuum foregrounds, line

foregrounds, and gravitational lensing, can affect the constraints obtained from the VID.

7.4.1 Continuum Foregrounds

One of the most substantial difficulties facing any intensity mapping experiment is the presence of bright continuum foreground emission. This problem is most evident in 21 cm experiments, where foreground contamination can be ~ 5 orders of magnitude brighter than the signal[12]. Other lines face the same problems to a somewhat lesser degree. For example, a COMAP like survey would observe at roughly 30 GHz, an area in frequency space familiar to CMB observers for containing large amounts of galactic synchrotron and free-free emission[126], as well as contamination from radio point sources[33]. Indeed, the CMB itself creates a substantial amount of "foreground" contamination at these frequencies.

Despite the overall strength of continuum foregrounds compared to the signal, they can be cleaned out of a map by taking advantage of their smooth frequency spectra. The intensity mapping signal will have a significant amount of structure in frequency space, as it maps the distribution of galaxies along the line of sight. Foregrounds such as synchrotron emission, however, are expected to be quite spectrally smooth. This means that continuum contamination is confined to Fourier modes oriented near the plane of the sky, and it can be effectively removed by subtracting out these modes, as was demonstrated by the COPSS experiment[33].

The effects of this mode subtraction are straightforward to model in power spectrum space, as the only effect is to reduce the number of available Fourier modes. Unfortunately, there is no clear way within our formalism to exactly replicate the effects of subtracting out only line-of-sight modes. Studying this procedure accurately may require the use of simulated maps. As an approximation to the correct foreground cleaning procedure, we consider a case where all of the $k = 0$ modes are removed from a map, both along the line of sight and in the plane of the sky. This will not exactly duplicate the true effect of continuum foregrounds, but it provides a rough estimate of the significance of the effect. For simplicity, we will also consider a spatially uniform foreground

Consider then a map made up of three contributions: the usual CO signal, Gaussian, zero-mean instrumental noise, and a spatially and spectrally smooth foreground with intensity $\overline{T}_{CF}$. The observed intensity in any given voxel will then be

$$T_{\rm obs} = \overline{T}_{\rm CO} + \Delta T_{\rm CO} + \Delta T_{\rm Noise} + \overline{T}_{CF}, \tag{7.17}$$

where we have divided the CO signal into mean and fluctuation parts. Subtracting out the $k = 0$ modes means that we can only observe

$$\Delta T_{\rm obs} = T_{\rm obs} - \overline{T} = \Delta T_{\rm CO} + \Delta T_{\rm Noise}, \tag{7.18}$$

where $\overline{T} = \overline{T}_{\rm CO} + \overline{T}_{CF}$. The probability of observing a voxel with fluctuation ΔT is

then

$$\mathcal{P}_\Delta(\Delta T) = \mathcal{P}(\Delta T + \overline{T}), \tag{7.19}$$

where $\mathcal{P}(T)$ is the original signal+noise VID.

When computing the Fisher matrix for this new model, the number of voxels in a fluctuation-space bin is

$$B_i = N_{\rm vox} \int_{\Delta T_{{\rm min},i}}^{\Delta T_{{\rm max},i}} \mathcal{P}_\Delta(\Delta T) d\Delta T. \tag{7.20}$$

Using Equation (7.19) we can rewrite this as

$$B_i = N_{\rm vox} \int_{\Delta T_{{\rm min},i}+\overline{T}}^{\Delta T_{{\rm max},i}+\overline{T}} P(T) dT. \tag{7.21}$$

We have now introduced a sixth unknown parameter $\overline{T}$ into our calculation. In the presence of our simple foreground, it is no longer known *a priori* which absolute intensity T corresponds to a given fluctuation ΔT. In other words, we can only measure the VID up to an additive constant in every voxel. Fortunately, we can easily add this extra parameter to our Fisher analysis to determine how it affects our constraints.

Figure 7.4 shows the effect of this added uncertainty, with the signal+COMAP noise constraints from Section 5 shown in blue and the new continuum-subtracted constraints shown in purple. The continuum-subtracted constraints are marginalized

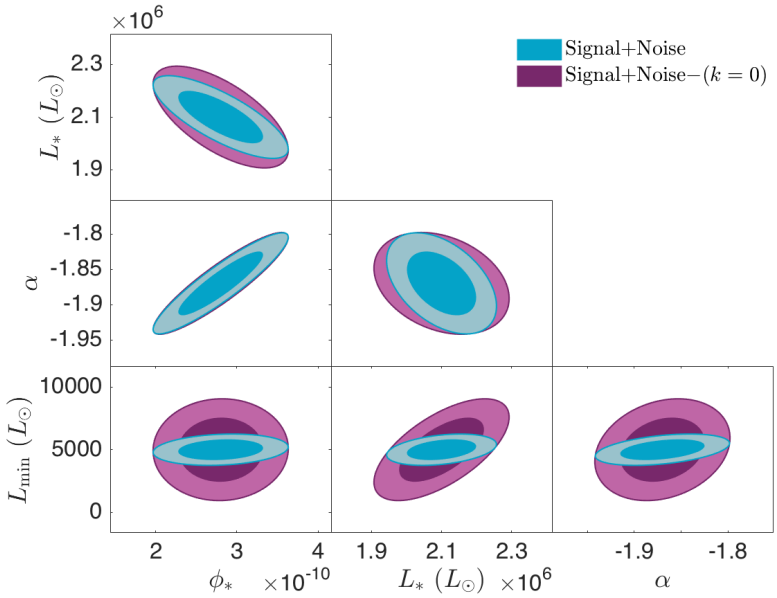

Figure 7.4: Fisher constraints for a CO map where all $k = 0$ modes have been subtracted to remove continuum (purple) compared with the signal+noise constraints from Figure 7.2.

over the unknown mean intensity $\overline{T}$ in addition to σ_G. The foreground subtraction procedure worsens the constraints slightly, but the effect is not dramatic. The effect of this subtraction on the $\Phi(L)$ constraints shown in Figure 7.3 are given in Appendix E. Simply subtracting out the mean of the map, therefore, does not lead to a substantial loss of constraining power. Note that a true spatially-varying foreground may leave some residual contamination which will affect these constraints. This effect will be highly model-dependent, so we leave it for future work.

7.4.2 Line Foregrounds

As discussed in Chapter 5, intensity maps can also suffer from interloper emission from other spectral lines emitted by galaxy populations at different redshifts. 21 cm surveys are not expected to see significant line contamination, as there are few other bright lines at such low frequencies[75]. Surveys aiming for other lines, however, could see interloper lines bright enough to rival or dominate over the target lines[77,38]. In such cases, the ability to study the original target line can be significantly degraded.

The literature contains several proposed means of recovering the power spectrum of a target line from a foreground-contaminated map. If one has access to other data in the target volume, either traditional data like a galaxy map or an intensity map of a different line, it is possible to cross-correlate the two data sets together to isolate emission from a single redshift[68]. It may also be possible to use anisotropies in the power spectrum to separate signal from foreground[78,79]. One could also seek simply to mask out voxels where the foreground contribution is brightest, either blindly as we demonstrated in Chapter 5 or by using galaxy surveys to locate bright foreground emitters[39].

We can divide line contamination into two broad categories based on how it affects the power spectrum. The most obviously problematic lines are those for which the mean foreground intensity $\overline{T}_F$ is greater than $\overline{T}_{\rm CO}$. We will refer to these lines as "clustering foregrounds", as the foreground dominates over the signal in the clustering component of the power spectrum. Examples of clustering foregrounds are Hα, OIII,

and OII in Lyα surveys[77] and higher-order CO rotational lines in CII surveys[38]. It is also possible for foreground lines that are substantially fainter than the signal on average to produce a small population of very bright sources that contribute a disproportionate amount of shot noise to a map, leading us to refer to them as "shot-noise foregrounds". The model considered in Chapter 5 for HCN contamination in CO surveys is an example of such a foreground.

Here we seek to understand how contamination from foreground lines affects the VID statistic. For the sake of simplicity, we will continue using our fiducial CO signal model and invoke hypothetical foregrounds to test their effects. We leave for future work a detailed exploration of signal and foreground models for surveys targeting CII, Lyα, and other lines. We choose the luminosity functions of our fiducial foreground lines to best demonstrate the behavior of the two types of foregrounds described above.

The first line, which we will name FG1, we choose to be a shot-noise foreground with $\overline{T}_{\rm FG1} = 0.1\overline{T}_{\rm CO}$ and $P_{\rm shot,FG1} = 2P_{\rm shot,CO}$. This roughly duplicates the behavior of the HCN model in Chapter 5, and is something of a worst-case scenario for a CO intensity mapping survey. The second line, which we will name FG2, we choose to be a clustering foreground with $\overline{T}_{\rm FG2} = 5\overline{T}_{\rm CO}$ and $P_{\rm shot,FG2} = 25P_{\rm shot,CO}$. Though no lines of this type are expected to appear in CO surveys, this is the type of line expected to cause issues in high-redshift CII and Lyα surveys. We assign both FG1 and FG2 an emission frequency of 88 GHz, which is the rest frequency of HCN(1-0).

Table 7.1: Luminosity function parameters for fiducial CO model and hypothetical foreground lines

Line	ϕ_*	L_* ($L_\odot$)	α	$L_{\rm min}$ ($L_\odot$)
CO(1-0)	2.8×10^{-10}	2.1×10^6	-1.87	5000
FG1	4.1×10^{-18}	6.5×10^8	-2.26	500
FG2	5.1×10^{-10}	3.4×10^6	-1.6	5000

We then choose values for the four Schechter parameters of each line to reproduce the desired power spectra. The chosen parameter values can be found in Table 7.1.

Figure 7.5 shows the luminosity functions, VIDs, and power spectra of the fiducial CO signal compared with those of the FG1 and FG2 models. Power spectra are computed using matter power spectra from CAMB[55]. The different behaviors of the two types of foregrounds can be clearly seen. The high-luminosity tail of FG1 leads to a shot-noise dominated power spectrum, which is strong enough to compete with the signal despite the overall weakness of the FG1 line. The FG2 luminosity function is very similar to that of the CO signal, but contains somewhat more bright sources. Since these sources are also closer to the observer, this leads to a VID that dominates over that of CO at most observed intensities, and a power spectrum that dominates on all scales. Note that the power spectra plotted here are unprojected, which means that the foreground spectra would be amplified even more relative to the signal in a true measurement.

It should be noted that the parameters in Table 7.1 do not uniquely determine the power spectra shown in the right-hand panel of Figure 7.5. Because we have four free luminosity function parameters from which to determine the two terms of

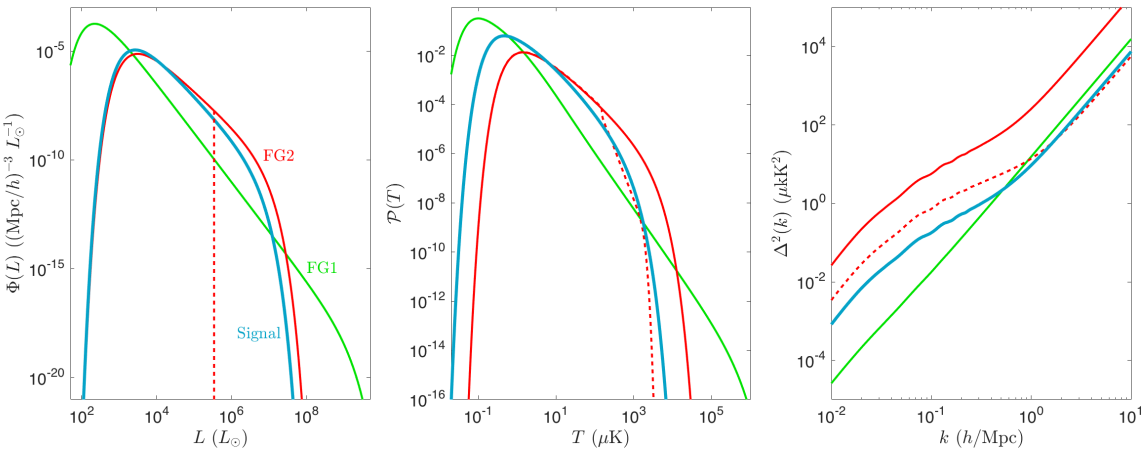

Figure 7.5: Luminosity functions (left), VIDs (center), and power spectra (right) for the fiducial CO signal (blue), shot noise foreground FG1 (green), and clustering foreground FG2 (red). Dashed red curves show the effect of masking FG2 emitters brighter than $L_{*,\mathrm{FG2}}/10$.

the power spectrum, we could choose an infinite number of different parameters and generate the exact same power spectra. The particular four parameters shown here have thus been chosen somewhat arbitrarily. As we are only using these fictional lines for a proof of concept, the exact parameter choices are not particularly important. However, this does serve to again illustrate the limitations of power spectra, as the same measured spectrum could be the result of many different luminosity functions.

Figure 7.6 shows the effect of including the shot noise foreground model in our Fisher analysis. We allow the parameters of the foreground model to vary along with those of the CO signal and marginalize over the foreground luminosity function parameters as well as the σ_G values for both lines. The result is that the constraining power diminishes, but not prohibitively. With the exception of the low-luminosity cutoffs, which mostly affect the noise-dominated portion of the VID, the signal and

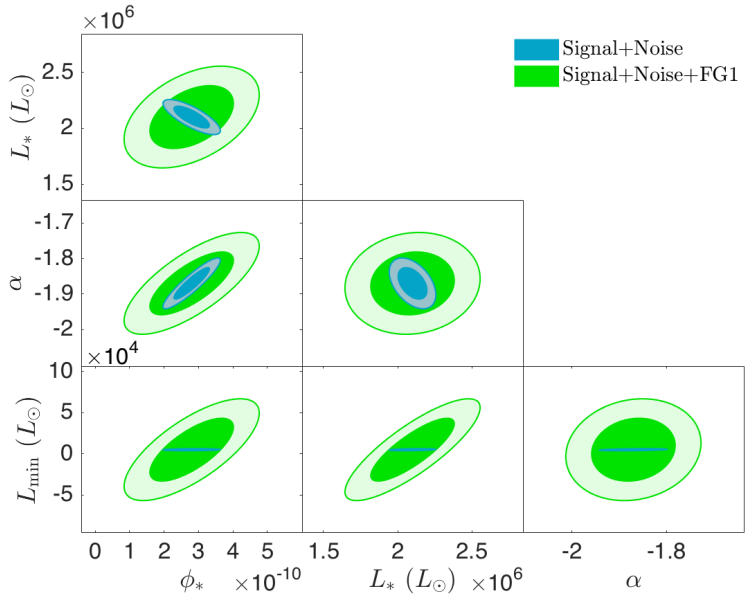

Figure 7.6: Fisher constraints on the CO signal parameters marginalized over the parameters of the shot-noise foreground FG1 (green), compared with the original CO constraints (blue).

foreground parameters are only weakly degenerate, and the constraints are fairly good. This is a marked improvement over what is seen in power spectrum space, where the shot noise components of the signal and foreground spectra are exactly degenerate. The constraints on $\Phi(L)$ with this foreground component included are given in Appendix E.

Figure 7.7 demonstrates that, as one might expect, the effect of the clustering foreground FG2 is much more dramatic. Marginalizing over the foreground parameters and the two σ_G values leaves significantly worse constraints on the CO model. In this case, all of the parameters except α have fractional uncertainties greater than unity. This is caused by the fact that the combined VID cannot easily distinguish the

two similar luminosity functions, leading to significant degeneracies. Because these constraints are so poor, we can do little more than set upper limits on the target luminosity function.

However, we can improve this result significantly if we have access to an external data set which we can use to isolate and mask voxels which are dominated by bright foreground emitters[44,38]. This effectively cuts off the foreground luminosity function at some $L_{\rm mask}$, as shown by the dashed curve in the left panel of Figure 7.5. While one could attempt to directly detect line emission in individual foreground sources, it is often easier to locate bright foreground emitters using some proxy observable which is correlated with line luminosity. For example, Silva et al.[38] find that foreground emission in a CII survey can be effectively reduced by masking sources with bright K-band magnitudes. Alternatitively, in Chapter 5 we demonstrated that, when measuring the power spectrum, it can also be helpful to simply mask out the brightest voxels in a map (with no need for external data) regardless of whether they are dominated by signal or foreground emission. In a one-point analysis however, this merely removes information about both luminosity functions without adding any extra means of distinguishing between them.

To study the effects of targeted foreground masking, we will assume that all voxels containing a source with FG2 luminosity greater than $L_{\rm mask} = L_{*,\rm FG2}/10$ are masked out of the map. This corresponds to a few percent of the total number of voxels. The effects of this masking on the luminosity function, VID, and power spectrum are

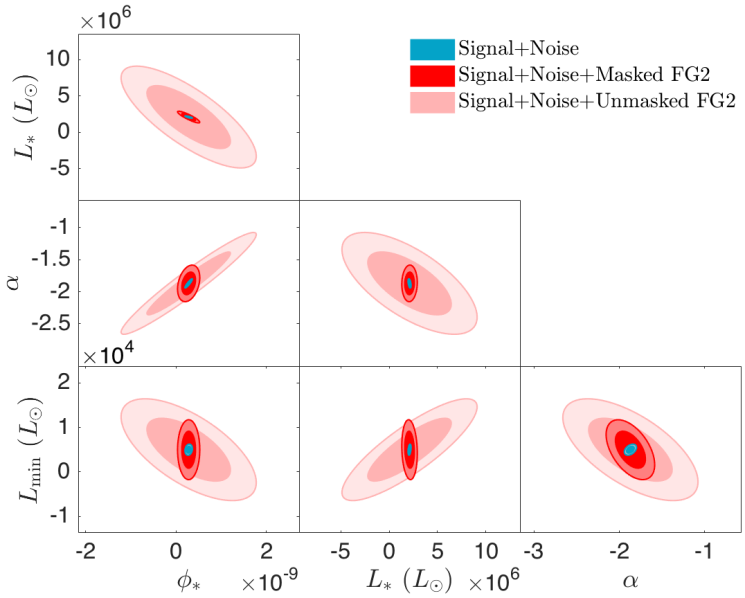

Figure 7.7: Fisher constraints on the CO signal parameters marginalized over the parameters of the clustering foreground FG2 (light red), compared with the original CO constraints (blue). Constraints with FG2 emitters brighter than $L_{*,\rm FG2}$ masked out are shown in red.

shown by the dashed red curves in Figure 7.5. Though this amount of masking still leaves a significant FG2 power spectrum, we can see from the dark red ellipses in Figure 7.7 that the constraints from the VID have been substantially improved. The resulting luminosity function constraints can be seen in Appendix E.

The exact value of the masking luminosity in a real survey will depend on which observables are used to find foreground emitters. Our choice of $L_{*,\rm FG2}/10$ is somewhat arbitrary given that the CO line is not expected to suffer from issues with clustering foregrounds. However, the results given here should be taken as fairly pessimistic, as the mask we have used is insufficient to reduce the FG2 power spectrum below that of CO.

Finally, note that the difference between "foreground" and "signal" in this context is simply a matter of experimental goals. The approach described here models each line in the same way, with the same number of parameters. If one desired to measure the luminosity function of the foreground line, one would only have to marginalize over the parameters of the original target line. The full Fisher constraints on the luminosity functions of our CO and FG2 models can be found in Appendix F.

7.4.3 Gravitational Lensing

In addition to emission from other astrophysical sources, structure between the observer and the emitters can affect the VID through gravitational lensing. Lensed galaxies have their positions on the sky altered slightly, their shapes distorted, and their intensities magnified. Since the voxels in an intensity map are large compared to any single galaxy, only the latter effect is important to consider when computing the VID. Lensing magnification, both weak and strong, will change the apparent luminosity function of the source population[127,128,129], which in turn will alter the VID in ways that could systematically affect luminosity function constraints. Lensing effects are expected to increase substantially with redshift, so in addition to our usual $z \sim 3$ CO model we will consider a case where we use the same luminosity function but take the emission redshift out to $z = 7$.

To estimate the effect of lensing on the VID, we need to compute the probability $\mathcal{P}_{\rm mag}(m)$ of a given galaxy to have a magnification between m and $m + dm$. The

observed luminosity after magnification is $L' = mL$. $\mathcal{P}_{\rm mag}(m)$ contains contributions from both the large-scale matter distribution as well as compact, virialised halos. To estimate the former, we adopt the method of Das & Ostriker[130]. First, we divide the intervening mass distribution up to the target redshift z into N uncorrelated, thin mass sheets. Each sheet i spans a comoving radius between r_i and $r_i + 1$ and has central redshift z_i. For each sheet, we consider fluctuations in the projected surface mass density Σ_i, where we have defined the surface density contrast

$$x(\theta, z_i) \equiv \frac{\Sigma(\theta, z_i) - \overline{\Sigma}(z_i)}{\overline{\Sigma}(z_i)}, \tag{7.22}$$

as a function of the sky position θ. In the Limber approximation[57,58], the two-dimensional power spectrum $P_2(\ell, z_i)$ for $x(\theta, z_i)$ is given by

$$P_2(\ell, z_i) = \frac{1}{\ell(r_{i+1} - r_i)^2} \int_{\ell/r_{i+1}}^{\ell/r_i} P_{NL}(k, z_i) dk, \tag{7.23}$$

where ℓ is the magnitude of the two-dimensional Fourier wavenumber and $P_{NL}(k, z_i)$ is the nonlinear three-dimensional matter power spectrum evaluated using `Halofit`[131]. The rms fluctuation σ_2^2 smoothed over an angular scale θ_0 is then

$$\sigma_2^2(\theta_0, z_i) = \frac{1}{2\pi} \int \ell P_2(\ell, z_i) e^{-\ell^2 \theta_0^2} d\ell. \tag{7.24}$$

The following parametric non-Gaussian PDF $\mathcal{P}_x(x)$ for the density fluctuation x

has been found[130] found to be a good fit to numerical simulations:

$$\mathcal{P}_x(x) = \frac{N}{x} \exp\left[-\frac{(\ln(x) + \omega^2/2)^2(1 + A/x)}{2\omega^2}\right]. \tag{7.25}$$

This is essentially a one-parameter family, as the three parameters N, A, and ω^2 are fixed by σ_2^2 through the requirements that $\langle x\rangle$=1, $\langle x^2\rangle - \langle x\rangle^2 = \sigma_2^2$, and that the PDF is normalized to unity.

In the weak lensing regime, the convergence κ receives a contribution

$$\kappa_i(\theta) = \frac{\Sigma(\theta, z_i) - \overline{\Sigma}(z_i)}{\Sigma_c(z_i)}, \tag{7.26}$$

from each mass sheet. We have introduced the critical surface density

$$\Sigma_c(z_i) = \frac{c^2}{4\pi G} \frac{D_s}{D_l(z_i)\,[D_s - D_l(z_i)]}, \tag{7.27}$$

where G is Newton's gravitational constant, D_s is the comoving radial distance to the target redshift, and D_l is that to the lens-plane redshift z_i. An overdense patch therefore contributes $\kappa > 0$ and an underdense patch contributes $\kappa_i < 0$.

A distribution for the cumulative convergence $\kappa = \sum_i \kappa_i$ can be estimated numerically by randomly drawing x for each mass sheet according to the corresponding PDF $\mathcal{P}_x(x)$. This gives the convergence PDF $\mathcal{P}_\kappa(\kappa)$. Using the weak-lensing relation

$m = (1-\kappa)^{-2}$, we can then compute the source-plane magnification PDF

$$\mathcal{P}_{\rm mag}(m) = \frac{(1-\kappa)^5}{2}\mathcal{P}_\kappa(1-m^{-1/2}), \tag{7.28}$$

which accounts for the non-Gaussian statistics of weak (de-)magnification fairly well.

However, this method underestimates the effect of strong magnification from virialised lenses. In particular, the large-m tail[132] is expected to decay roughly as m^{-3}. This can be remedied by manually adding a power law tail to Equation (7.28),

$$\mathcal{P}_{\rm mag}(m) \rightarrow \mathcal{P}_{\rm mag}(m) + \Theta(m-1)\exp\left[\frac{1}{4(m-1)^4}\right]\frac{(1-\kappa_0)^3}{2}\mathcal{P}_\kappa(\kappa_0)\left(\frac{m}{m_0}\right)^{-3}, \tag{7.29}$$

where a typical matching point is $\mu_0 = (1-\kappa_0)^{-2} \sim 3$ and Θ is the Heaviside function. The resultant semi-analytic model is found to agree reasonably well with ray tracing of N-body simulations. Figure 7.8 shows magnification PDFs both with and without strong lensing for a given draw from $\mathcal{P}_x$ with source redshift $z = 7$. The two PDFs are similar in the low-magnification regime, but differ substantially at higher magnifications.

The effect of magnification is to alter the apparent luminosity distribution $\mathcal{P}(L) = \Phi(L)/\overline{n}$ of the source galaxies, replacing it with

$$\mathcal{P}'(L) = \int \frac{\mathcal{P}_{\rm mag}(m)}{m}\mathcal{P}\left(\frac{L}{m}\right)dm, \tag{7.30}$$

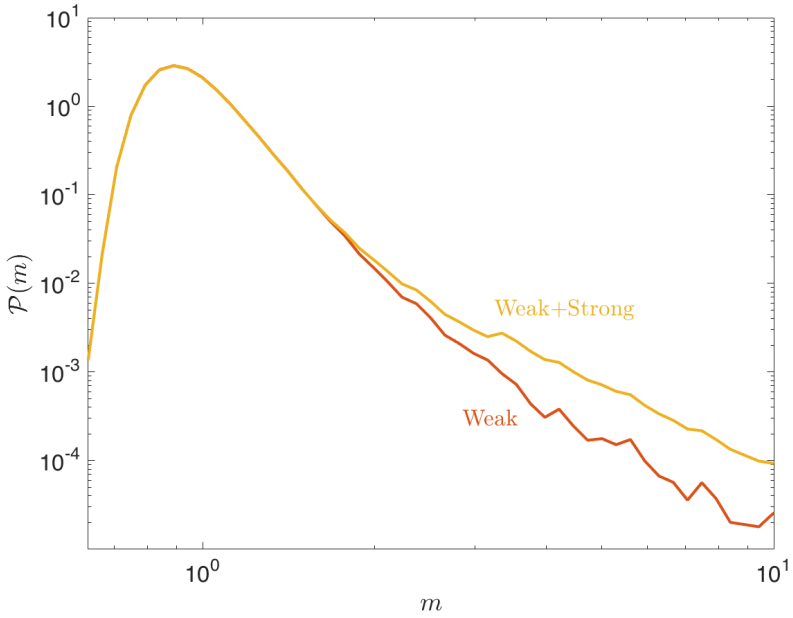

Figure 7.8: Magnification PDFs drawn both with (yellow) and without (red) the power-law strong-lensing tail.

as galaxies appear brighter or fainter due to lensing. If we compute the VID from $\mathcal{P}'(L)$, we see that the distribution is altered somewhat, as shown in Figure 7.9. We have plotted both the VIDs for the full $\mathcal{P}_{\rm mag}$ as well as the weak-lensing-only $\mathcal{P}_{\rm mag}$ at both $z = 3$ and $z = 7$. Note that in order to better understand the size of the effect, we have plotted these VIDs as bin counts rather than $\mathcal{P}(T)$. If we divide the difference between the lensed and unlensed bin counts by the square root of the unlensed counts, we can get an idea of how strong the lensing effect is relative to the Poisson error in a given bin. The effect is visible at both redshifts, and as expected it increases as redshift increases. The primary impact is to make the cutoffs at both ends less sharp. This will thus likely not be a hugely significant effect for near-future experiments, as the low-luminosity cutoff is below the noise limit and there are

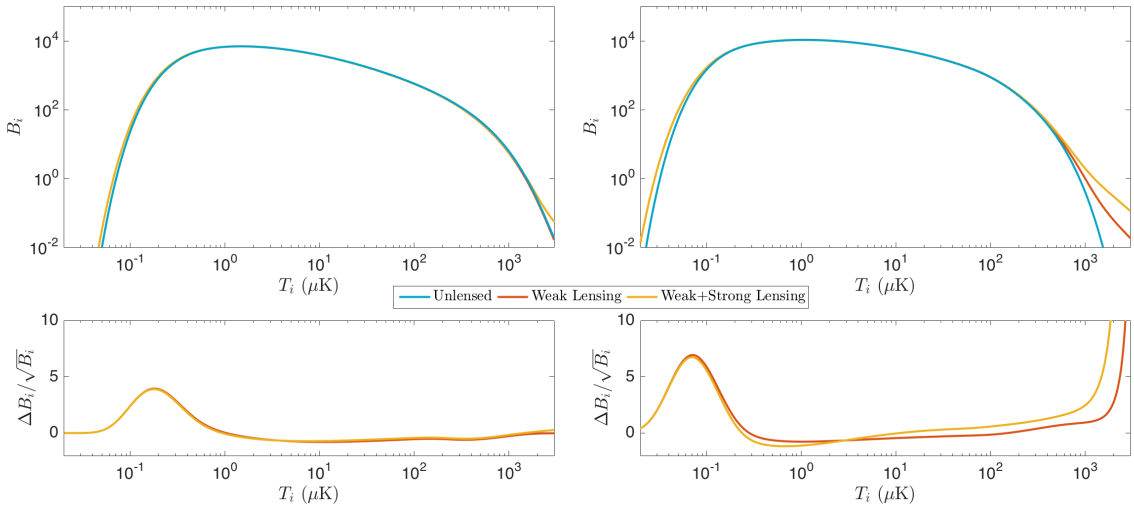

Figure 7.9: Top row: Number of voxels in bins of width $\Delta \log T = 0.012$ dex at redshift $z = 3$ (left column) and $z = 7$ (right column) for our unlensed fiducial CO model (blue), a weakly lensed model (red), and a full weak+strong lensed model (yellow). Bottom row: Difference between the lensed and unlensed voxel counts divided by the Poisson errors used in our Fisher analysis.

relatively few voxels above the high-luminosity cutoff. However, future experiments with sufficient sensitivity and area may need to take into account magnification effects when attempting to constrain luminosity functions, especially when targeting high redshifts.

In the above lensing computations, we neglected the fact that lensing also alters the volume subtended by a given solid angle[129]. Voxels with greater magnifications also subtend a smaller volume at the same solid angle, so they will, on average, contain fewer galaxies. When computing the source luminosity function in Equation (7.30), this effect should average out over the full survey volume, as the total magnification must average to unity on large scales. However, there is an additional effect when moving from $\mathcal{P}'(L)$ to the VID. This magnification/volume correlation means that

a magnified voxel will contain fewer-than-average sources with higher-than-average apparently luminosities, introducing a correlation between $\mathcal{P}_N(T)$ and $\mathcal{P}(N)$ in Equation (7.6). This will likely tend on average to reduce the strength of the lensing effect on the VID, further weakening the significance of lensing as a confounding factor in our analysis. The results presented above are therefore a conservative estimate of the effect of lensing on the VID.

7.5 Discussion

The constraining power of the VID statistic is clear from the above results. These constraints cannot be obtained using the power spectrum alone, as the power spectrum only measures the quantities $P_{\rm shot}$ and $\overline{T}^2\overline{b}^2$. This means that the power spectrum can only constrain at most two of the four model parameters of our luminosity function model, resulting in significant degeneracies. The only way to obtain useful results for a model with more than two parameters is then to apply priors. However, as shown in Figure 2, the VID allows us to constrain our four-parameter Schechter model to order $\sim 10\%$ even without prior data. Even the low-end exponential cutoff $L_{\rm min}$, which falls well below our noise limit, can be constrained by the VID measurement.

This effect can be clearly seen in Figure 7.3, where the VID produces much better constraints on $\Phi(L)$ than power spectrum methods. The constraints are improved

significantly despite the fact that both forecasts use the same experimental setup and very similar luminosity functions. As mentioned above, however, the constraints plotted in Figure 7.3 likely somewhat underestimate the luminosity function errors, as both the VID and power spectrum constraints assume that the CO emitters have a luminosity function exactly described by a single model. More realistic constraints could be obtained, for example, by using a spline model[119,120], or by using the values of $\Phi(L)$ in different luminosity bins as the model parameters.

Though the VID provides considerably more information about the luminosity function, the power spectrum likewise contains information that is not present in the VID. Just as the power spectrum contains only integrals over the luminosity function, the VID contains only integrals over the galaxy power spectrum. The VID thus leaves out almost all of the information about the spatial distribution of galaxies. This is intuitively obvious, since a one-point statistic by definition does not take into account voxel locations. Because of this, the VID will not be nearly as effective as the power spectrum when attempting to measure, for example, baryon acoustic oscillations. The VID and the power spectrum are highly complementary statistics, and using both to study a map will yield substantially more information than using only one or the other.

There are some subtleties involved when attempting to directly combine these two statistics, however. Though they contain different information, there is a substantial amount of covariance between the VID and the power spectrum of a map. For ex-

ample, observing a region of space with a greater-than-average number of very bright sources will obviously yield more very bright voxels, but will also yield an excess amount of shot noise power. Since shot noise is scale-independent, the net effect of this will be to create covariance between the size of the brightest voxel bins and the amplitude of the power spectrum on all scales. Accurately treating covariances like this will like require detailed numerical simulations, so we leave a full study of this issue for future work.

As expected, the various forms of foreground contamination we studied do add extra uncertainty to our forecasts. However, the VID retains significant constraining power even in the face of this contamination. Our model of a continuum-subtracted map, though simplified, suggests that the process of cleaning foregrounds like Galactic dust and synchrotron will not destroy our ability to constrain the luminosity function from the VID. This is due to the fact that the continuum foreground in our model merely adds a constant to every voxel, and this additive shift is not degenerate with any of our Schechter parameters. The real situation, with foregrounds that vary in amplitude along different lines of sight, will be somewhat more complicated, but we do not expect the end result to be radically different than that plotted in Figure 7.4.

The VID statistic also performs fairly well in the presence of line foregrounds. In the case of a shot-noise foreground like our FG1, the parameters L_*, α, and ϕ_* are well constrained even without any attempts to clean out the interloper line. Degeneracies with FG1 weaken our ability to constrain $L_{\rm min}$, but overall our constraints remain

reasonably strong. FG1 is something of a worst-case-scenario contaminant for a CO map, so it is encouraging that this method remains viable. The unmasked clustering foreground is less optimistic looking, though we are still able to obtain useful constraints on the power-law slope of the luminosity function. However, surveys that suffer from this type of foreground, such as those targeting CII and Lyα, already plan to use some type of masking routine to clean out interloper lines. When we apply a mask to our FG2 luminosity function, out VID formalism produces strong constraints on all of the signal parameters except for $L_{\rm min}$. These results would improve even further with more aggressive masking or the application of other foreground-cleaning methods.

Foreground lines illustrate another key advantage of the VID over the power spectrum. Since a shot-noise dominated spectrum is purely scale independent, $P_{\rm shot}$ for the signal line is exactly degenerate with $P_{\rm shot}$ of the foreground. This means that, even in the best case scenario, the power spectrum in the presence of a foreground line can only reliably measure the quantity $\overline{T}^2\bar{b}^2$ for the signal line. The interloper removes one of the only two quantities the power spectrum can access. As stated before however, the VID can still put strong constraints on most signal parameters even in the presence of foregrounds.

As for gravitational lensing, which can be thought of as contamination from foreground mass, the effect is very weak for our fiducial $z \sim 3$ model and only slightly stronger when we take our emission out to $z \sim 7$. As shown in Figure 7.9, the biggest

effects are on the two exponential tails of the VID. On the faint end, we see an excess due to weak lensing along underdense lines of sight. Bins in this region are shifted by several sigma, but this effect will be negligible in a real observation because these bins will be strongly noise dominated. As for the bright end, there is an excess, mostly caused by strongly lensed sources. This could yield underestimates of the bright end slope of the luminosity function, but there are few enough voxels in this part of the distribution that the effect will likely not be large, at least for early surveys. However, as surveys become more sensitive and target larger volumes, this effect will only become more important. It is important to note though that gravitational lensing is somewhat different than the other types of contamination in that the observed luminosity function can in principle be deconvolved after the fact. This could potentially reduce the importance of lensing effects compared to other systematics. One may also eventually be able to treat this lensing effect as a signal in its own right, and use it to trace the distribution of dark matter between the observer and the emitters.

The formalism we have designed here is applicable to many different intensity-mapping surveys. Specifically, it can be applied as is to any survey where all of the emitting sources are small compared to the voxel size. For most target lines, all of the emission comes from within galaxies, and most intensity maps will have resolutions much larger than any galaxy, so this assumption holds. However, this is not the case for measurements of the 21 cm line at the epoch of reionization, as this emission comes from large volumes of diffuse intergalactic gas. Lyα surveys may have a similar IGM

component as well[40], though the intensity of this component is a subject of debate. For diffuse emission, the concept of a luminosity function is not well defined, and our formalism breaks down. The one-point statistics of such lines must be treated in a very different manner[133,134].

Our results here provide an excellent proof-of-concept of the VID method. However, there are a number of effects that require deeper study in future work. One such effect is due to the clustering of the source galaxies. We have treated the different bins of the voxel histogram as entirely independent, which may not be the case in a real, clustered map. If a map contains a large overdensity, for example, we would expect to see an excess of voxels in several bright bins. The opposite is true for a large underdensity.

Similarly, we have neglected the detailed effects of beam smoothing. Our choice to set voxel size equal to the beam FWHM may not be optimal for every survey, so depending on the exact beam characteristics it may be necessary to use larger voxels in order to minimize systematic errors. Larger voxels will smooth out more details of the galaxy distribution, leading to weaker constraints. For example, increasing our voxel size from $3'$ to $6'$ increases the 1-σ error on ϕ_* from Figure 7.2 by a factor of roughly 1.5. Even after choosing the optimal voxel size, there may still be excess voxel-to-voxel correlations due to beam effects that need to be taken into account. Many other instrumental systematics, such as ground contamination and pointing errors may also degrade our constraints.

One way to test these and other effects in greater detail would be to make use of maps simulated from large N-body codes. Many components of our VID calculation, such as the halo bias and the galaxy number count distribution could be easily studied based on such simulations. It would also be substantially easier to test clustering, beam smoothing, and line-of-sight mode subtraction in a simulated map than in our analytic work. We used the simulations from Chapter 4 in Appendix D to test the numerical stability of our VID calculations, but these simulations do not include many of the effects described here. Though there are many aspects of realistic galaxy distribution which will be still be difficult to capture, full-scale numerical simulations will be an important tool as we prepare to apply this formalism to real data.

7.6 Summary

We have presented in this chapter a powerful new method for measuring line luminosity functions from intensity maps using the probability distribution of voxel intensities. This voxel intensity distribution can be calculated using $P(D)$ analysis techniques and measured from a map by making a histogram of voxel intensities. Because intensity maps are extremely non-Gaussian, this one-point statistic contains a substantial amount of information that cannot be obtained from usual power spectrum analyses.

We tested our formalism on a four-parameter model of CO emission observed by an

experiment similar to the planned COMAP survey. We found that the VID statistic was able to constrain these four parameters with an average error of order $\sim 10\%$, despite not including any prior information. Incorporating various forms of foreground contamination such as continuum emission, interloper lines, and gravitational lensing weakens these constraints by varying degrees. However, the VID statistic still provides useful information despite these contaminants, even in very pessimistic cases where the power spectrum would be completely swamped by foregrounds. Our results here serve as an excellent proof of the VID concept. Though more work is necessary to fine tune the various subtleties of this method, this work suggests that the VID will make a powerful addition to the intensity mapping toolbox as more and more experiments come online in the coming years.

Chapter 8

Constraining the Cosmic Star-Formation History with CO Intensity Maps

In this chapter, we will demonstrate how the VID method we created in Chapter 7 can be used to constrain astrophysically interesting quantities, specifically the cosmic star formation history. Our current measurements of star formation as a function of redshift indicate that the star formation rate (SFR) increased with time up to a maximum at around $z \sim 2-3$ before declining until the present day[89]. However, as with all high-redshift measurements, our knowledge of star formation in the distant universe comes from a relatively small population of high-redshift galaxies. Here we discuss how CO intensity mapping can be used to improve our understanding of how

the star formation rate density (SFRD) evolved with redshift.

Our P(D)-based VID statistic is ideally suited to this type of measurement. One could make use of the CO power spectrum, but both $\overline{T}_{\rm CO}$ and $P_{\rm shot}$ in Equation 2.1 depend only on integrals over the CO luminosity function, which limits their ability to constrain its shape. On top of that, $\overline{T}_{\rm CO}$ is degenerate with the galaxy bias and the matter growth function, and $P_{\rm shot}$ may be degenerate with shot noise contributions from foreground lines such as HCN(1-0) (see Chapter 5). As the CO luminosity and star formation rate (SFR) of a galaxy may be related nonlinearly, these integral constraints are suboptimal for constraining SFRD. The VID provides a full measurement of the CO luminosity function, which would allow us to make much better SFR estimates. We will show that future CO surveys could provide competitive constraints on SFRD at high redshift, while directly observing the large numbers of faint sources which are invisible to current measurements.

This chapter is primarily based on work published in Ref.[50]. We describe the model we use to relate CO luminosity and SFR in Section 8.1, and forecast constraints on the cosmic SFRD in Section 8.2. We discuss these results in Section 8.3, and summarize in Section 8.4.

8.1 CO-to-SFR Model

In Chapters 6 and 7, we modeled CO emission as a Schechter function. Here we need a means of connecting CO luminosity to star formation physics, so we will return to the type of model we used in Chapter 5, where the CO luminosity L of a halo is related to its mass M by a power law with amplitude A and index b. We again set $A_{\rm CO} = 2 \times 10^{-6}$ and $b_{\rm CO} = 1$[31], and assume that only a fraction $f_{\rm duty} = t_{\rm age}/10^8$ yr of halos with masses greater than $M_{\rm min} = 10^9\ M_\odot$ emit CO at any given time, where $t_{\rm age}$ is the age of the universe. The number density of halos with a given mass is then determined by a mass function dn/dM[102]. If $\bar{n}$ is the mean number density of halos with masses greater than $M_{\rm min}$ calculated by integrating this mass function, then we can write $\mathcal{P}_1(M) = (1/\bar{n})dn/dM$. We can then convert this to $\mathcal{P}_1(L)$ using our mass-luminosity relation, and from there to a brightness temperature following Equation 7.1. The resulting PDF can then be used to compute the VID using the method described in Chapter 7

It should be noted that the proof-of-concept model considered here is simplified in several respects for ease of computation. For example, the duty cycle $f_{\rm duty}$ we use here is independent of halo mass, and we have a hard cutoff on SFR below $M_{\rm min}$. A more physical treatment would be to use a duty cycle that varies smoothly with mass and goes to zero at low masses[135]. In addition, a realistic CO luminosity function likely has a "knee" at a fainter luminosity than our model predicts[136], meaning that our calculation may over-predict the number counts of the brightest sources. It may be

preferable to use a luminosity function model with more than two parameters which can accurately capture this "knee", such as a Schechter function.

A real CO map will include contributions from instrumental noise as well as foreground emission. For this chapter we ignore the effects of continuum foregrounds such as dust and synchrotron emission, but we do consider foregrounds with line spectra as well as instrumental noise. Specifically, we consider contamination from the 88 GHz HCN(1-0) line, which we model in the same way as CO, with $A_{HCN} = 1.7 \times 10^{-15}$ and $b_{HCN} = 5/3$. We assume the instrumental noise has a Gaussian PDF with zero mean and variance σ_N. These contaminants can be added to the VID by convolving $P(T)$ for the original map with that of the contaminant.

We again base our model survey on the "Full" CO mapping experiment described in Table 6.1. This experiment surveys 6.25 square degrees between $z = 2.4$ and 2.8, with $\sigma_N = 5.8$ μK. For simplicity, we assume that neither the signal nor the foregrounds evolve significantly across the observed frequency range. We use information from all of the frequency channels to study the CO properties averaged over the full survey volume.

Figure 8.1 shows the predicted VIDs for the CO signal and the two contaminants. Below $T \sim 1$ μK, unphysical "ringing" effects come into the VID due to the hard cutoff at $M_{\rm min}$. Since this regime would be noise dominated in a realistic experiment, the effect on the CO constraints should be small. As expected, the behavior we see here is similar to the CO line with shot-noise foreground seen in Section 7.4. Noise

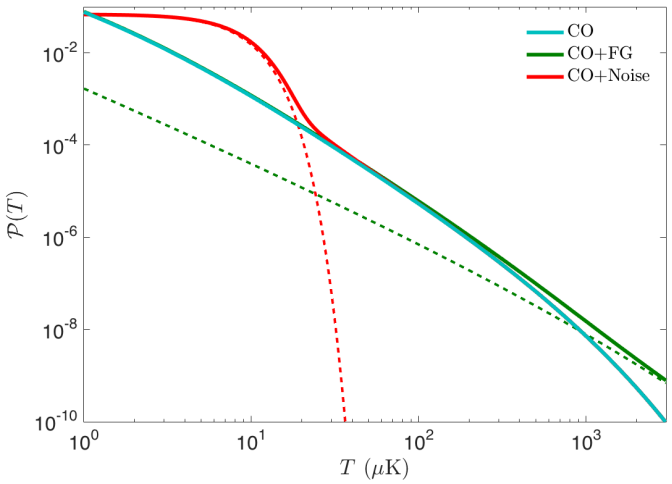

Figure 8.1: Predicted VIDs for intensity maps of CO, CO with HCN, and CO with instrumental noise (*solid curves*: the full VID; *dashed curves*: the contaminant alone).

dominates the VID at small T, and the foreground HCN dominates at large T.

We can use the above VID to obtain constraints on A and b, and we can then compute constraints on other astrophysically interesting quantities. The focus of this chapter is on star formation, so we need a model relating CO luminosity to SFR. Our mass-luminosity parameters $A_{\rm CO}$ and $b_{\rm CO}$ in Chapter 5 were originally derived from a set of empirical scaling relations[31,30].

$$\frac{L_{\rm FIR}}{L_\odot} = C_{\rm FIR}\left(\frac{L'_{\rm CO}}{\rm K\ km\ s^{-1}\ pc^2}\right)^{X_{\rm FIR}}, \tag{8.1}$$

where $C_{\rm FIR}$ and $X_{\rm FIR}$ are constants and $L_{\rm CO}/L_\odot = 4.6\times10^{-5}(L'_{\rm CO}/{\rm K\ km\ s^{-1}\ pc^2})$[80,27]. The FIR luminosity of a galaxy can then be related to its SFR through the Kennicutt relation

$$\frac{\rm SFR}{M_\odot/{\rm yr}} = C_{\rm SFR}\frac{L_{\rm FIR}}{L_\odot}, \tag{8.2}$$

for some constant $C_{\rm SFR}$ [105] (Note that in the notation of Li et al.[35], $\alpha = X_{\rm FIR}$, $\beta = \log(C_{\rm FIR})$, and $\delta_{\rm MF} = 10^{10} C_{\rm SFR}$). The values for the above constants used in Chapter 5 are $C_{\rm FIR} = 1.35 \times 10^{-5}$, $X_{\rm FIR} = 5/3$, and $C_{\rm SFR} = 1.5 \times 10^{-10}$. The weakest part of this process is the relation between SFR and halo mass[31]; here we assume they are related by a power law and use the above relations to write it in terms of A and b:

$$\mathrm{SFR(M)} = 9.8 \times 10^{-18} \left(\frac{\mathrm{A_{CO}}}{2 \times 10^{-6}} \right) \mathrm{M}^{5\mathrm{b_{CO}}/3}. \tag{8.3}$$

We can then integrate this over the mass function to obtain the mean SFRD ψ in our survey, which turns out to be 0.12 $M_\odot/\mathrm{yr/Mpc^3}$ for our fiducial model.

The VID can also be used to predict constraints on other astrophysically interesting quantities, such as the mean CO intensity $\overline{T}_{\rm CO}$. This quantity is an important measure of the quantity of molecular gas at high redshift, and is important for understanding the power spectrum of an intensity map. The mean volume emissivity of CO emitters in this model is

$$j_{\rm CO} = A_{\rm CO} \int_{M_{\rm min}}^{\infty} M^{b_{\rm CO}} \frac{dn}{dM} dM. \tag{8.4}$$

We can use the procedure outlined in Section 2.1 to convert this to $\overline{T}_{\rm CO}$.

8.2 SFRD Constraints

In order to determine the constraining power of such an experiment, we perform a Fisher analysis on our calculated $P(T)$ curves. We compute a Fisher matrix (see Equation 7.15) over the parameters $X_i = (A_{\rm CO},\ b_{\rm CO},\ A_{HCN},\ b_{HCN})$. We again assume Poisson errors on the number of voxels in a given bin.

We invert the Fisher matrix to obtain the covariance matrix, and marginalize over the foreground parameters to plot confidence regions for our CO parameters. Figure 8.2 shows the 95% confidence ellipses for a map containing just CO, as well as for maps contaminated by HCN and instrumental noise. With no foreground lines, the parameters $A_{\rm CO}$ and $b_{\rm CO}$ are fairly degenerate, since to linear order increasing one or the other simply makes every halo brighter. Even with both foregrounds and noise included, the uncertainties are around the $\sim 20\%$ level in $A_{\rm CO}$ and the $\sim 1\%$ level in $b_{\rm CO}$, so relying on one-point statistics, these intensity maps can provide excellent constraints on the CO luminosity function. For comparison, note that in Chapter 3 we estimated an uncertainty on $A_{\rm CO}$ of roughly an order of magnitude.

Constraints on our original parameters X_i are straightforward to convert to constraints on other parameters Y_i. We simply need to multiply the Fisher matrix calculated in Equation (7.15) on both sides by the Jacobian matrix $J_{ij} = \partial X_i/\partial Y_j$. If $Y_i = (\overline{T}_{\rm CO}, \psi)$, we get the confidence regions plotted in Figure 8.3. The uncertainty on SFRD with noise and foregrounds included is on the order of $\sim 10\%$. It should be noted that this is an optimistic calculation. When converting our Fisher matrix to

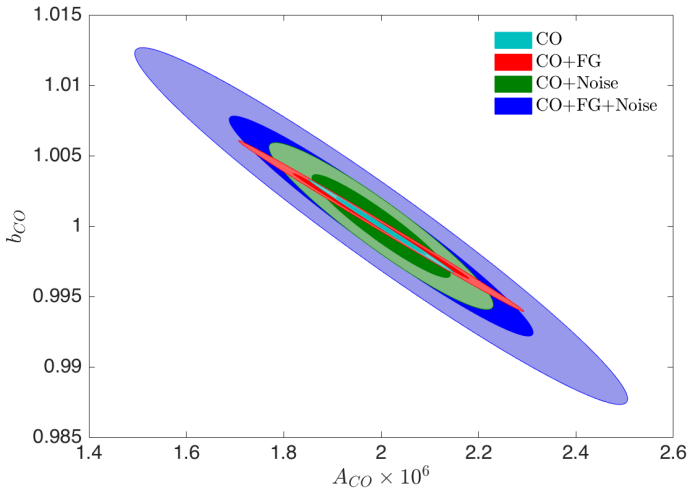

Figure 8.2: Predicted 95% confidence regions for CO mass-luminosity parameters $A_{\rm CO}$ and $b_{\rm CO}$ for four cases: only CO with no noise or foregrounds (light blue), CO with foreground HCN (green), CO with instrumental noise (red) and CO with both noise and foreground (blue).

SFRD and $\overline{T}_{\rm CO}$, we assumed that the scaling relations used to connect $L_{\rm CO}$ to SFR are well constrained. In practice, the uncertainties in these relations will reduce the constraining power of such a measurement. More galaxy observations will be required to reach the constraints shown here. However, it is clear from these calculations that the potential of a CO intensity mapping experiment to constrain SFRD is quite high.

Up until now we have been considering a survey centered at $z = 2.6$. It is interesting to consider how the SFRD constraint varies with redshift, especially as the reliability of other methods diminishes with redshift. Figures 4 and 5 illustrate this redshift dependence. For this calculation, we hold all of the instrumental parameters constant except the central observing frequency. Figure 8.4 shows the fractional uncertainty in SFRD as a function of redshift. Figure 8.5 places the result with noise and foregrounds in context with results from the far ultraviolet (FUV) observations

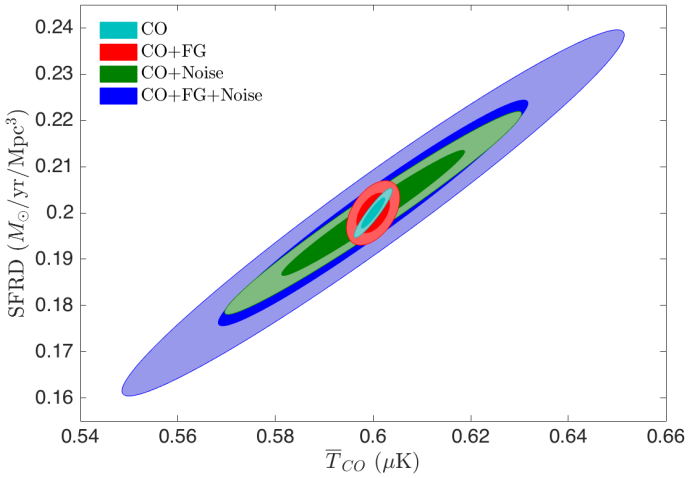

Figure 8.3: Predicted 95% confidence regions for mean CO brightness temperature $\overline{T}_{\rm CO}$ and SFRD for the same four cases shown in Figure 5.3.

described by Madau & Dickinson[89] (Hereafter M&D) and gamma-ray burst (GRB) observations described by Kistler et al.[137]. For Figure 8.5, we use a somewhat different form for $\psi(z)$ than we have used thus far. Our modeling, which is quite simplistic and only intended to serve as a proof-of-concept, gives a form for $\psi(z)$ (dotted black line) which differs from the M&D fit (solid black line). To facilitate comparison with the data (grey points), the blue curves in Figure 8.5 show the 1-σ uncertainty on SFRD calculated assuming the M&D form for $\psi(z)$.

As mentioned above, there is theoretical uncertainty regarding the relations in Equations (8.1) and (8.2). The magenta curves in Figures 8.4 and 8.5 show the effect of taking it into account. The dashed magenta curves in both figures assume a 10% uncertainty on $C_{\rm FIR}$ and $C_{\rm SFR}$. Current results[27] tend to give a constraint on $X_{\rm FIR}$ which is roughly an order of magnitude better than the constraint on $C_{\rm FIR}$, similar to how in Figure 2 we obtain a much better fractional constraint on $b_{\rm CO}$ than

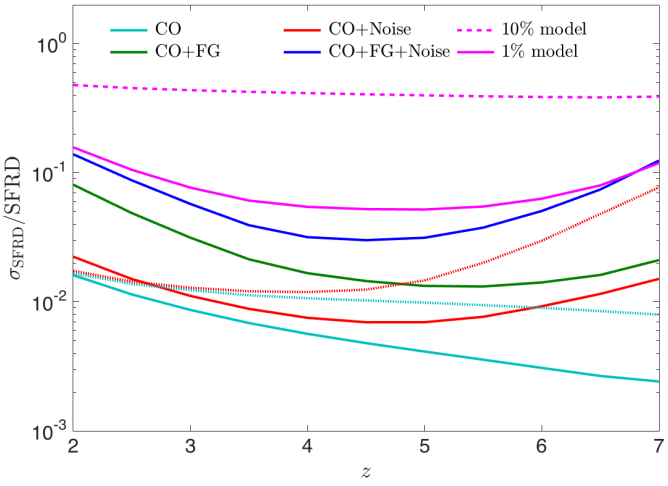

Figure 8.4: Predicted fractional 1-σ uncertainties on SFRD for different target redshifts. The cyan, green, red, and blue curves show the same four cases as Figures 2 and 3. The dashed magenta curve shows the effect of adding a pessimistic 10% uncertainty on $C_{\rm FIR}$ and $C_{\rm SFR}$ and 1% uncertainty on $X_{\rm FIR}$. The solid magenta curve shows the effect of reducing these model uncertainties to a more optimistic 1% and 0.1% respectively. The dotted cyan and red curves show the effect of adding an exponential cutoff and scatter to the power law $L(M)$ model.

$A_{\rm CO}$. These curves therefore assume a 1% uncertainty on $X_{\rm FIR}$. These constraints are somewhat better than current values, but they are likely pessimistic compared to what will be available once intensity mapping data are available. If there are resolved CO emitters in the same volume, one could calibrate the CO-SFR relation and dramatically reduce this uncertainty. The solid magenta curve in Figure 8.4 shows a more optimistic scenario where the errors on $C_{\rm FIR}$ and $C_{\rm SFR}$ are 1% and that on $X_{\rm FIR}$ is 0.1%.

In order to test the model dependence of our results, we consider two aspects of more sophisticated CO emission models which have been neglected thus far in our deliberately simplified analysis. First, the true CO luminosity function likely cuts off

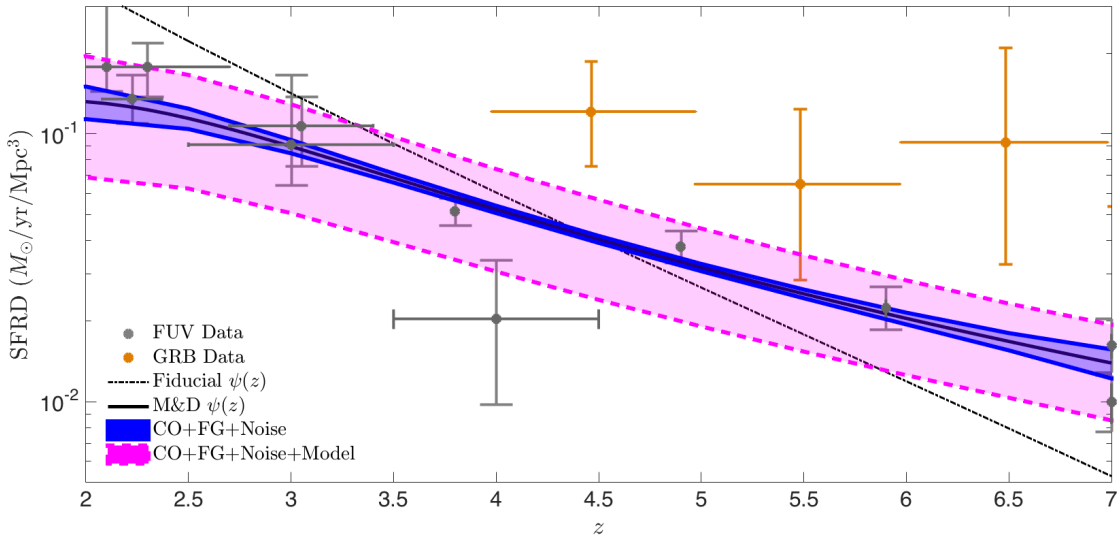

Figure 8.5: Comparison between our predicted SFRD constraints and existing FUV data from[89] (grey points) and GRB data from[137] (orange points). The dash-dotted black line shows our simplistic fiducial $\psi(z)$ model, and the solid black curve shows the[89] fit to the data. The blue curves show the $\pm 1\sigma$ uncertainty from CO intensity mapping calculated assuming the fitted $\psi(z)$, including foregrounds and noise, but excluding modeling uncertainty. The dashed magenta curves show the effect of adding the pessimistic 10% model uncertainty from the dashed magenta curve in Figure 4.

significantly earlier than the halo mass function[136]. In other words, the CO $L(M)$ relation has a turnover at some halo mass. We take this into account by adding an exponential cutoff to our $L(M)$ at a halo mass of $M_* = 2 \times 10^{12}$ solar masses, a value consistent with other calculations[35]. Secondly, semianalytic models[99,138] show that the CO luminosity of a galaxy depends on many parameters besides its mass. Again following Li et al.[35], we add lognormal scatter with $\sigma = 0.3$ dex to our $L(M)$ model to account for this, preserving some dependence on mass while allowing for fluctuations of other galaxy properties.

The dotted curves in Figure 8.4 show the effects of these changes for the cases with and without instrumental noise. We can see that increasing the model complexity and suppressing the bright end of the VID has relatively little effect at lower redshifts but noticeably worsens the constraints on SFR at high redshift, up to roughly a factor of 5 worse at $z = 7$. However, even at these high redshifts the overall uncertainty would still be dominated by the error on the CO-SFR conversion seen in the magenta curves.

8.3 Discussion

From the above results we can get a good sense of the potential of this technique for constraining SFRD. In an ideal world with no foregrounds, noise, or modeling uncertainty, this 6.25 deg^2 survey could constrain SFRD to $\sim 1\%$ at $z \sim 3$. This

scenario is obviously unrealistic, but we could reduce the modeling uncertainty with more observations and reduce the noise with more sensitive instruments. It may be possible to reduce the foreground contamination as well. The masking procedure described in Chapter 5 will have little effect on the one-point statistics since it would simply involve ignoring all pixels greater than some $T_{\rm max}$. However, as described in Chapter 7 it may be possible to reduce the foreground levels by using other data from the same area of the sky.

Counterintuitively, the constraint on SFRD in Figure 8.4 seems to improve with redshift when foregrounds and noise are neglected. This is because at higher redshifts, the same survey parameters correspond to a larger volume of space, which will include more sources. In our simple model, $A_{\rm CO}$ and $b_{\rm CO}$ do not evolve with redshift, so the only effect of moving to more distant sources comes from the evolution of the mass function. Eventually there will be few enough sources above $M_{\rm min}$ that the constraint will worsen, but for our redshift range the volume increase outweighs this effect. When noise is included, the volume effect dominates at low redshift, but the amplitude of the noise relative to $\overline{T}_{\rm CO}$ takes over at $z \sim 5$ or so. With foregrounds, the relative amplitude of CO and HCN decreases with redshift, starting to dominate over volume effects at around $z \sim 6$. The model uncertainties as we have considered them here are nearly redshift independent, so the constraint with these included depends weakly on redshift.

As seen in the dotted curves in Figure 8.4, our constraints do not change drasti-

cally when we consider a more sophisticated model. The model used for the dotted curves contains significant scatter around the mean $L(M)$ relation and adds a sharp cutoff to the luminosity of high mass halos, yet the change in fractional uncertainty on SFR is only order unity. Though even this more intricate model is still relatively simplistic, our qualitative conclusions should still hold. We leave a detailed quantitative prediction for future work.

Currently there is substantial disagreement between high redshift SFRD's measured using different methods, as illustrated by the roughly order-of-magnitude disagreement between the FUV and GRB data plotted in Figure 8.5. This discrepancy can possibly be explained if the FUV analysis underpredicts the quantity of low-luminosity galaxies[137]. Since intensity mapping is far more sensitive to the fainter population than other methods it could provide a powerful means to resolve this discrepancy.

An effect we have not taken into account here is contamination from continuum foregrounds. Though in principle these should be easy to clean out, it is unclear what form the residual contamination would take and how it would alter our results. Another effect worth considering is the dependence of these results on the voxel size. As discussed in Chapter 7, larger beams will wash out additional information about the signal, but may also reduce covariance between nearby voxels due to beam smoothing.

It may also be possible to constrain other quantities related to star formation with

these intensity maps besides global SFRD. For example, because the convolutions in the VID computation depend on the full range of possible luminosities, the VID at all intensities depends on the chosen value of $M_{\rm min}$. Thus, even though the faintest pixels will be noise-dominated, it could be possible to determine $M_{\rm min}$ from our VID, in an analogous way to the determination of $L_{\rm min}$ in Chapter 7. Another common unknown for high-redshift galaxies is what is the relative contribution to the overall intensity of high-mass vs. low-mass galaxies. This question is difficult to answer with galaxy surveys, but it should be more tractable with intensity mapping. The methods described in this chapter should be readily applicable to these and many other problems of high-redshift star formation.

From Figure 5, it is clear that CO intensity mapping has great potential for constraining the high-redshift SFRD. Though we are currently limited by modeling uncertainty, this should improve with better models and more observations. The potential bound shown by the blue curves, though idealized, is remarkably strong despite using only a small survey area. In addition, the results shown above do not include any additional constraints from the two-point statistics of a map. Though there will likely be substantial covariance between the VID and power spectrum, the power spectrum does contain spatial information which may help refine this type of measurement. Combining intensity mapping with existing multiwavelength galaxy surveys should further improve the constraints. More work is necessary before CO intensity mapping can be used to accurately determine SFRs, but these results clearly

demonstrate the value of this effort.

8.4 Summary

We have demonstrated the usefulness of our VID statistic by using it to forecast the ability of a CO intensity mapping survey to constrain the cosmic star formation history. Our constraints on SFRD are limited by model uncertainty to a fractional uncertainty of 50% between redshifts 2 and 7. However, if this model uncertainty could be reduced, our constraints improve to roughly 5-10%, which is highly competitive with current measurements. Because intensity maps directly measure the emission from very faint galaxies, even the pessimistic constraints can be used to resolve important conflicts in current measurements which cannot resolve the faint population.

Chapter 9

Conclusions

Intensity mapping offers a powerful tool to study physics in the distant universe, on scales ranging from individual molecular clouds to the large-scale structure. In this thesis, we have studied several aspects of intensity mapping surveys, and created useful tools for extracting information from them. We have also highlighted the many ways in which intensity maps allow science above and beyond what is possible using conventional observations.

In Chapter 2, we described how to compute the power spectrum of an intensity map. We performed this computation in two ways, one using three-dimensional spatial coordinates, and another using two-dimensional angular coordinates. The three-dimensional form captures the full behavior of the map along the line of sight, while the two-dimensional form uses coordinates which adapt more readily to the frame of observation, and is easier to apply to measurements at different redshifts. In the

future, it may prove valuable to use a third formalism which combines the advantages of both of these by decomposing the full three-dimensional intensity fluctuations into basis functions in spherical coordinates[139]. Such a procedure would be computationally intensive, but would most accurately capture the behavior of observed maps.

In Chapter 3, we used two-dimensional power spectra to study the prospects for conducting a CO intensity mapping survey at $z \sim 3$. Using a single-parameter approximation of four literature models, we found that there is little agreement between different models on the amplitude of the CO signal. The brightest and faintest models produced spectra which differed by roughly two orders of magnitude. We also found, however, that all but the faintest models predicted detectable signal-to-noise ratios in a hypothetical survey. While this result is encouraging, it is also important to note that the high SNR we obtained is highly dependent on survey design, so upcoming experiments must plan their observing strategies carefully to ensure detection of the CO signal.

Because existing N-body methods of simulating intensity maps are slow and unwieldy when dealing with large volumes, we created in Chapter 4 a method for quickly simulating maps with known galaxy power spectra. By drawing lognormal random fields which approximate the true galaxy number distribution and assigning luminosities from a model mass or luminosity function we can create two-dimensional slices of intensity maps for a wide range of line models. Though these simulations have proved themselves quite useful, in the future we will look to expand them to three

dimensions. This extension will be necessary to study important problems such as continuum foreground contamination.

We applied these simulations in Chapter 5 to the issue of foreground line emission. By adding together simulations of signals and foregrounds for surveys targeting CO, Lyα, and CII we were able to study the effects of voxel masking on combined power spectra. Because bright voxels tend to be dominated by foreground emission, we were able to show that masking them out can in some cases allow the measurement of a "clean" signal power spectrum. Unfortunately, this masking also changes the amplitude of the signal spectrum, removing in the process most of the astrophysical information content. The shape of the spectrum is preserved however, which could allow the recovery of useful cosmological information. We found that this method worked successfully for our models of CO and Lyα, but failed for CII owing to a very high foreground-to-signal ratio and a low angular resolution. Compared to other foreground cleaning methods, this has the advantage of requiring neither high signal-to-noise ratio nor external data, but it would be useful to consider how voxel masking could be combined with other techniques to maximize the amount of recovered signal.

In Chapter 6 we expanded our discussion of foreground lines to include faint lines which nevertheless provide valuable probes of astrophysics. We focused on the ^{13}CO (1-0) line in surveys targeting the usual ^{12}CO transition. In the local universe ^{13}CO observations compliment ^{12}CO measurements by allowing the measurement of the density of molecular gas clouds, which provide the nurseries for forming stars. We

find that by cross-correlating pairs of frequency bands, it is possible to access the ^{13}CO transition at high redshift in an intensity map. Furthermore, we showed that by doing so it is possible to measure the global properties of molecular gas in the distant universe, offering a potential new window into star formation and galaxy formation at high redshift. In the future, we intend to expand this type of analysis to other pairs of lines, which could dramatically improve our understanding of the conditions within these galaxies.

Though power-spectrum methods have proven their worth as a tool to study intensity maps, they are incapable of determining the full details of an emitting population. This stems from the fact that any non-Gaussian map, like the ones we study here, contains information which is not accessible in the power spectrum. In Chapter 7, we applied $P(D)$ analysis to intensity mapping to create a new statistic, the VID, which allows for deeper study of the luminosity function of a target population. We showed that, even when contaminated by continuum emission, line foregrounds, and gravitational lensing, the VID allows considerably stronger constraints on line luminosity functions than the power spectrum. Though there is more work needed before the VID is fully ready to be applied to real data, most notably on the problems of realistic continuum foregrounds and covariance with the power spectrum, it is clear that it represents a powerful new tool for studying intensity maps.

Using the VID, we demonstrated in Chapter 8 that a CO intensity map could be used to make a competitive measurement of the star formation rate at high red-

shift, a quantity of significant interest when studying galaxy formation and evolution. Because intensity maps directly observe emission from faint populations invisible to traditional measurements, they suffer from very different systematic effects and can therefore resolve important tensions in existing data.

The work presented in this thesis represents a significant step in the maturation of intensity mapping as a field. As the first data begin to become available, the ideas presented here will be crucial as we attempt to realize the potential of this new observable.

Appendix A

Spurious ^{12}CO Auto Power in Cross-Spectrum

When we take the cross-spectrum between two chosen bands, our goal is to get at the correlation between ^{12}CO and ^{13}CO emission from galaxies at a single redshift. However, the cross-spectrum also contains power from pairs of ^{12}CO emitters and pairs of ^{13}CO emitters at different redshifts. For example, cross-correlating bands centered at 32 and 30.6 GHz would correlate emission from ^{12}CO and ^{13}CO at $z = 2.6$. It will also correlate ^{12}CO emission from $z = 2.6$ and 2.8 and ^{13}CO emission from $z = 2.4$ and 2.6. If the spatial separation between these pairs of redshifts is large compared to the scale set by a given multipole, then this spurious correlation should be small. Here we demonstrate this fact quantitatively.

The 3D auto spectrum $P_{\rm CO}(k, z)$ of a CO line is given by Equation (2.1). If we

wish to compute the angular cross-spectrum between two different redshift bands we simply apply Equation (2.9) with the ^{12}CO 3D spectrum and selection functions f_1 and f_2 centered on our two chosen bins. As stated in Section 2.2, we can evaluate this expression in two limits. If our bands have a comoving spatial width δr which satisfies $\ell\delta r/r \gg 1$ (i.e. if we consider fluctuations on length scales small compared to the width of the redshift slice), we can use the well-known Limber approximation[57,58]

$$C_\ell^{12\times12,s} \approx \int \frac{H(z)}{c}\frac{f_1(z)f_2(z)}{r^2(z)}P_{\rm CO}\left[k=\ell/r(z),z\right]dz. \tag{A.1}$$

In this limit, it is clear that if f_1 and f_2 do not overlap for any value of z, this integral vanishes and we will see no spurious correlation in our cross-spectrum. This agrees with our previous intuition, as the Limber approximation is valid for large ℓ's where we expected our signal to be small.

In the opposite limit, where $\ell\delta r/r \ll 1$, we can replace both selection functions with Dirac delta functions centered at the two redshifts z_1 and z_2. Evaluating Equation (2.9) in this limit gives

$$C_\ell^{12\times12,s} \approx \frac{2}{\pi}\int k^2 P_{\rm CO}(k) j_\ell\left[kr(z_1)\right] j_\ell\left[kr(z_2)\right] dk. \tag{A.2}$$

Since the extra ^{12}CO correlation vanishes in the Limber approximation, we can take this narrow-band approximation as an upper limit on the amount of spurious power. If we evaluate this integral numerically at $\ell = 100$, we find a value for $C_\ell^{12\times12,s}$ which

is approximately 1% of the cross-spectrum $C^{12\times13}$. The ratio falls to $\sim$ a few parts in 10^5 at $\ell = 500$. The largest scales accessible to the instruments we consider here are around $\ell \sim 100$, so we can safely ignore this extra power in our analysis. Since the ^{13}CO line is so much fainter than the ^{12}CO line, spurious ^{13}CO correlation will be even less significant.

It should be noted that this argument would not hold if we were to consider an experiment with a substantially larger survey area. At $\ell = 10$, contamination from this extra ^{12}CO in the narrow-band approximation rises to $\sim 60\%$ of the cross-spectrum. This is still an overestimate, but it would still likely need to be taken into account if performing this analysis on very large scales. It may also needed to be taken into account for measurements at higher redshifts, such as those which would target the Epoch of Reionization[30]. As the target redshift increases, the comoving separation of the two ^{12}CO populations in the two bands decreases, leading to additional correlation on smaller scales. This is particularly significant because these band pairs in a survey at $z = 7$ will cover populations of ^{12}CO emitters separated by nearly the baryon acoustic oscillation scale, which will significantly boost the correlation.

Appendix B

Marginalizing over ^{12}CO in Cross-Power

For the molecular gas constraints shown in Figure 6.4, we assumed that the ^{12}CO luminosity function was known exactly. Here we will relax that assumption and show that the effects on our final constraints is small. Unfortunately, we cannot simply perform a full Fisher analysis combining all four ^{12}CO parameters from Equation (6.13) with the molecular gas parameters R and f_s. Even if we use both the auto- and cross-spectra, the only degrees of freedom we can use are the amplitudes of the clustering and shot noise components of each spectrum. This leaves us with four degrees of freedom and six parameters, ensuring that several parameter constraints will be perfectly degenerate.

This issue arises due to the fact that the power spectrum only contains all of

the information about a map if that map is purely Gaussian[140]. Since a galaxy's CO luminosity is determined by a variety of nonlinear astrophysical processes, the intensity distribution in an intensity map will be very non-Gaussian. In order to fully constrain the luminosity functions which give rise to these non-Gaussian maps, we need some prior information from another source. One promising source of extra information is the one-point pixel intensity distribution. We show in Chapter 7 that this statistic could be used with an experiment similar to the COMAP "Full" survey to constrain the ^{12}CO luminosity function to within $\sim 10 - 20\%$.

Consider then a Fisher matrix computed from both power spectra over all six parameters. We can add to this a prior matrix assuming uncorrelated 10% errors on the four ^{12}CO parameters. The new Fisher matrix will then be invertible, allowing us to forecast parameter constraints. If we perform this analysis, the fractional 1-σ error on R increases to $\sim$35% from the $\sim$30% value quoted above. As stated above, this increase is small because the ^{12}CO emission is substantially brighter that the ^{13}CO emission. Uncertainties on R and f_s will therefore always be dominated by uncertainty in the ^{13}CO measurements.

Appendix C

Effects of Luminosity-Dependent Bias on $P(D)$ Computations

As mentioned in Section 7.1, though our P(D) formalism takes into account the average halo bias when computing $\mathcal{P}(N)$, we do not correctly include the full luminosity dependence of the bias. Here we will describe an approximate method for computing the full biased VID and attempt to get an idea of how important the effect is. Since our fiducial model is not derived from any sort of mass-luminosity relation, for the purposes of this discussion we will look at the probability distribution $\mathcal{P}(M)$ of total halo mass contained within a voxel rather than total intensity. We will assume that halo masses are drawn from the Tinker mass function[102] and we will use the corresponding $b(M)$[103].

In our normal formalism, we would compute σ_G using the average bias $\bar{b}$, which

when calculating $\mathcal{P}(M)$ would be

$$\overline{b} = \frac{\int b(M) dn/dM dM}{\int dn/dM dM}. \tag{C.1}$$

Now, instead of computing one average bias for all galaxies, assume instead that we split our population in half around some mass value $M_{\rm edge}$. We can then compute two average biases, two σ_G's, and two mass PDFs $\mathcal{P}_i(M)$ for the low- and high-mass populations. By the same logic used in Equation (7.4), the full $\mathcal{P}(M)$ will be the convolution of those of the two subsets. We can then divide our mass range into smaller and smaller subsets to more accurately model a full continuous $b(M)$.

Figure C.1 shows the fractional change in $\mathcal{P}(M)$ from a single average bias bin to many narrow bins. If the bins are too wide, there are significant discontinuities at the bin edges, but as the bins become smaller we can see that these edge effects become small. At the low mass end, the effect of this bias is of order $\sim 10\%$, however these low-mass voxels would likely fall into the noise-dominated regime of a full VID calculation. This effect is small enough ($\lesssim 5\%$ outside the noise-dominated region) that we do not expect a full Fisher analysis including the effects of bias to produce significantly different constraints. However, this does imply that leaving this effect out of the analysis of future measurements could lead to non-negligible systematic errors. It is therefore important for future models to accurately take into account the luminosity-dependent bias.

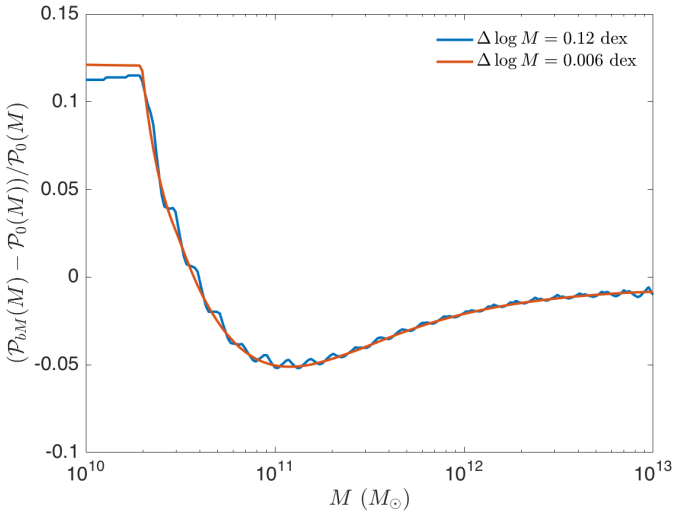

Figure C.1: Fractional change in $\mathcal{P}(M)$ when including mass-dependent bias in wider bins (blue) ad narrower bins (red).

Appendix D

Comparison between numerical and analytic VIDs

We demonstrated in Chapter 4 a method for simulating 2D slices of intensity maps with given power spectra. Our VID formalism is based in part on these simulations, and the two methods are based on the same set of assumptions. Both assume galaxies have randomly assigned luminosities drawn from a luminosity function and are distributed on the sky according to a lognormal random field. If we estimate a VID from a simulation prepared using our routine, we can then verify that the calculations presented here produce reliable results.

We generate 400 slices of a CO intensity map with galaxy luminosities drawn from our fiducial $\Phi(L)$ with voxel sizes set by the COMAP parameters. We then bin the resulting map in bins of width $\Delta \log T = 0.12$ dex, and compare to the bin values

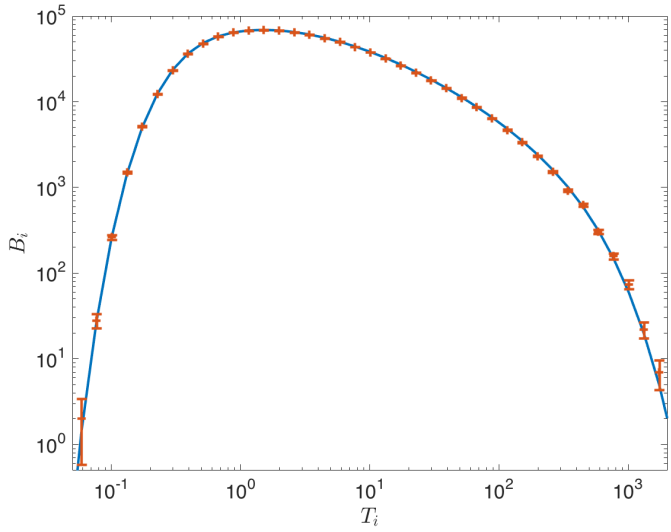

Figure D.1: Number of voxels in intensity bins B_i simulated using the method from Chapter 4 (red points) compared with the predicted values from our fiducial CO VID (blue). Error bars on the simulated bins are Poisson.

predicted by our fiducial VID. Figure D.1 shows the results. The plotted error bars on the simulated bins are the Poisson errors we use when computing Fisher matrices. This result clearly shows that our simulations and VID calculations are in good agreement, and that our numerical computations do not introduce a significant amount of error into our final VIDs. The simulations still use the same set of approximations we used when deriving the VIDs though (see discussions in Chapter 7), so it would be useful in the future to test our formalism against more in-depth N-body simulations which would take into account, for example, the full mass-dependent halo bias.

Appendix E

Luminosity function constraints including foreground effects

Just as we did in Figure 7.3, we can use our Fisher matrix constraints on the fiducial luminosity function parameters under various forms of contamination to estimate uncertainties on the luminosity function $\Phi(L)$. The results of this procedure are shown in Figure E.1. The region shown in blue give the same constraints shown previously for the case including only signal and instrumental noise. The purple region shows the effect of subtracting out the mean of the map to remove continuum foregrounds. As described previously, the constraints worsen somewhat but on the whole do not change substantially.

The green region shows the constraints with the shot-noise FG1 line added. The constraints get notably worse when this contamination is added, but the VID remains

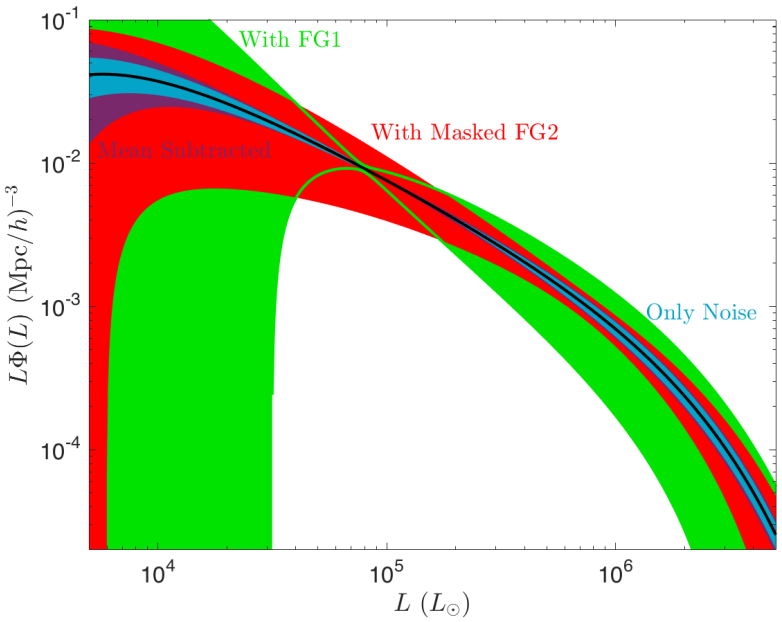

Figure E.1: 95% confidence regions around our fiducial luminosity function for cases with only instrumental noise (blue), noise plus mean subtraction (purple), noise plus shot-noise foreground FG1 (green), and noise plus masked clustering foreground FG2 (red).

reasonably constrained above a few times 10^4 $L_\odot$ despite the fact that no attempt has been made to clean out the interloper line. The errors reach a minimum at around $L \sim 10^5$ $L_\odot$. This is roughly the luminosity where the difference between the signal and FG1 luminosity functions is greatest. It is also just above the point where the instrumental noise falls off, so the signal-to-noise is maximized. Errors with a masked clustering foreground FG2 are shown in red. However, once we mask out the brightest foreground sources, we see that we get significant constraining power over nearly the entire luminosity range, with the best constraints coming at high L where the foreground is masked out.

Appendix F

Full VID constraints on CO with clustering foreground

Figure F.1 shows the full 8-parameter Fisher matrix constraints on the parameters of our fiducial CO model and those of the clustering foreground FG2. Light red ellipses show the constraints with the unmasked foreground, dark ellipses show those with the masking from Figure 7.5 applied. The parameter constraints change in some counter-intuitive ways when the masking is applied, due to the substantial degeneracies between the signal and foreground parameters in the unmasked case. For example, the constraints in the $[L_*, L_{*,2}]$ parameter space rotate by nearly 90 degrees when masked. This occurs because the masking removes nearly all of the information about the foreground cutoff, while simultaneously leaving the signal cutoff free of contamination.

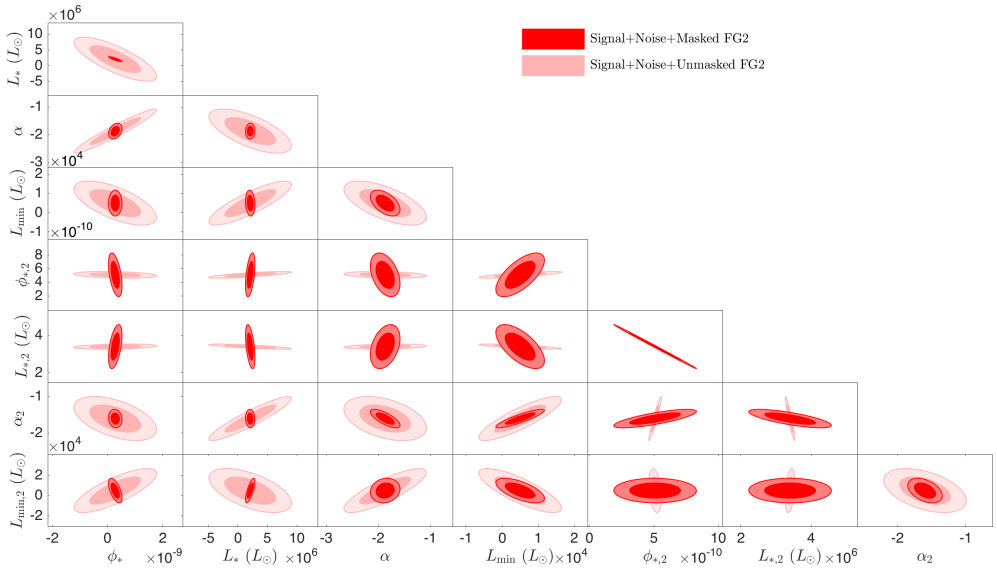

Figure F.1: Full Fisher matrix constraints for a model including the CO signal, the COMAP instrumental noise, and clustering foreground FG2 both masked (dark) and unmasked (light). Parameters with subscript "2" denote parameters of the FG2 luminosity function.

Bibliography

[1] Hinshaw, G., Larson, D., Komatsu, E., et al. 2013, ApJS, 208, 19

[2] Planck Collaboration, Adam, R., Ade, P. A. R., et al. 2016, A&A, 594, A1

[3] The Polarbear Collaboration, Errard, J., Ade, P. A. R., et al. 2010, ArXiv e-prints, arXiv:1011.0763

[4] Benson, B. A., Ade, P. A. R., Ahmed, Z., et al. 2014, in Proc. SPIE, Vol. 9153, Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VII, 91531P

[5] Louis, T., Grace, E., Hasselfield, M., et al. 2016, ArXiv e-prints, arXiv:1610.02360

[6] Harrington, K., Marriage, T., Ali, A., et al. 2016, in Proc. SPIE, Vol. 9914, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 99141K

[7] Blanton, M. R., Bershady, M. A., Abolfathi, B., et al. 2017, ArXiv e-prints, arXiv:1703.00052

[8] Guth, A. H. 1981, Phys. Rev. D, 23, 347

BIBLIOGRAPHY

[9] Eisenstein, D. J., Zehavi, I., Hogg, D. W., et al. 2005, ApJ, 633, 560

[10] Tegmark, M., Hamilton, A. J. S., Strauss, M. A., Vogeley, M. S., & Szalay, A. S. 1998, ApJ, 499, 555

[11] Planck Collaboration, Aghanim, N., Arnaud, M., et al. 2016, A&A, 594, A11

[12] Morales, M. F., & Wyithe, J. S. B. 2010, ARA&A, 48, 127

[13] Ali, Z. S., Parsons, A. R., Zheng, H., et al. 2015, ApJ, 809, 61

[14] Tingay, S. J., Goeke, R., Bowman, J. D., et al. 2013, PASA, 30, e007

[15] van Haarlem, M. P., Wise, M. W., Gunst, A. W., et al. 2013, A&A, 556, A2

[16] DeBoer, D. R., Parsons, A. R., Aguirre, J. E., et al. 2016, ArXiv e-prints, arXiv:1606.07473

[17] Santos, M., Bull, P., Alonso, D., et al. 2015, Advancing Astrophysics with the Square Kilometre Array (AASKA14), 19

[18] Cohen, A., Fialkov, A., Barkana, R., & Lotem, M. 2016, ArXiv e-prints, arXiv:1609.02312

[19] Bandura, K., Addison, G. E., Amiri, M., et al. 2014, in Proc. SPIE, Vol. 9145, Ground-based and Airborne Telescopes V, 914522

[20] Newburgh, L. B., Bandura, K., Bucher, M. A., et al. 2016, in Proc. SPIE, Vol. 9906, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 99065X

[21] Suginohara, M., Suginohara, T., & Spergel, D. N. 1999, ApJ, 512, 547

[22] Bigiel, F., Leroy, A., Walter, F., et al. 2008, AJ, 136, 2846

[23] Leroy, A. K., Walter, F., Bigiel, F., et al. 2009, AJ, 137, 4670

[24] Murray, N. 2011, ApJ, 729, 133

[25] Kennicutt, R. C., & Evans, N. J. 2012, ARA&A, 50, 531

[26] Bolatto, A. D., Wolfire, M., & Leroy, A. K. 2013, ARA&A, 51, 207

[27] Carilli, C. L., & Walter, F. 2013, ARA&A, 51, 105

[28] Heyer, M., & Dame, T. M. 2015, ARA&A, 53, 583

[29] Righi, M., Hernández-Monteagudo, C., & Sunyaev, R. A. 2008, A&A, 489, 489

[30] Lidz, A., Furlanetto, S. R., Oh, S. P., et al. 2011, ApJ, 741, 70

[31] Pullen, A. R., Chang, T.-C., Doré, O., & Lidz, A. 2013, ApJ, 768, 15

[32] Mashian, N., Sternberg, A., & Loeb, A. 2015, JCAP, 11, 028

[33] Keating, G. K., Bower, G. C., Marrone, D. P., et al. 2015, ApJ, 814, 140

[34] Keating, G. K., Marrone, D. P., Bower, G. C., et al. 2016, ApJ, 830, 34

[35] Li, T. Y., Wechsler, R. H., Devaraj, K., & Church, S. E. 2016, ApJ, 817, 169

[36] Spitzer, L. 1978, Physical processes in the interstellar medium, doi:10.1002/9783527617722

[37] Yue, B., Ferrara, A., Pallottini, A., Gallerani, S., & Vallini, L. 2015, MNRAS, 450, 3829

[38] Silva, M., Santos, M. G., Cooray, A., & Gong, Y. 2015, ApJ, 806, 209

[39] Crites, A. T., Bock, J. J., Bradford, C. M., et al. 2014, in Proc. SPIE, Vol. 9153, Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VII, 91531W

[40] Pullen, A. R., Doré, O., & Bock, J. 2014, ApJ, 786, 111

[41] Doré, O., Bock, J., Ashby, M., et al. 2014, ArXiv e-prints, arXiv:1412.4872

[42] Fonseca, J., Silva, M. B., Santos, M. G., & Cooray, A. 2017, MNRAS, 464, 1948

[43] Glover, S. 2013, in Astrophysics and Space Science Library, Vol. 396, The First Galaxies, ed. T. Wiklind, B. Mobasher, & V. Bromm, 103

[44] Gong, Y., Cooray, A., & Santos, M. G. 2013, ApJ, 768, 130

[45] Visbal, E., Haiman, Z., & Bryan, G. L. 2015, MNRAS, 450, 2506

[46] Breysse, P. C., Kovetz, E. D., & Kamionkowski, M. 2014, MNRAS, 443, 3506

[47] —. 2015, MNRAS, 452, 3408

[48] Breysse, P. C., & Rahman, M. 2016, ArXiv e-prints, arXiv:1606.07820

[49] Breysse, P. C., Kovetz, E. D., Behroozi, P. S., Dai, L., & Kamionkowski, M. 2017, MNRAS, arXiv:1609.01728

[50] Breysse, P. C., Kovetz, E. D., & Kamionkowski, M. 2016, MNRAS, 457, L127

[51] Cooray, A., & Sheth, R. 2002, Physics Reports, 372, 1

[52] Gong, Y., Cooray, A., Silva, M., et al. 2012, ApJ, 745, 49

[53] Schulz, A. E., & White, M. 2006, Astroparticle Physics, 25, 172

[54] Jeong, D., & Komatsu, E. 2009, ApJ, 703, 1230

[55] Lewis, A., Challinor, A., & Lasenby, A. 2000, ApJ, 538, 473

[56] Jeong, D. 2010, PhD thesis, University of Texas at Austin

[57] Limber, D. N. 1953, ApJ, 117, 134

[58] Rubin, V. C. 1954, PNAS, 40, 541

[59] Zaldarriaga, M., Furlanetto, S. R., & Hernquist, L. 2004, ApJ, 608, 622

[60] Hopkins, A. M., & Beacom, J. F. 2006, ApJ, 651, 142

[61] Muñoz, J. A., & Furlanetto, S. R. 2013, MNRAS, 435, 2676

[62] Visbal, E., & Loeb, A. 2010, JCAP, 11, 016

[63] Weiß, A., Walter, F., & Scoville, N. Z. 2005, A&A, 438, 533

[64] Carilli, C. L. 2011, ApJL, 730, L30

[65] Schechter, P. 1976, ApJ, 203, 297

[66] Mo, H. J., & White, S. D. M. 1996, MNRAS, 282, 347

[67] —. 2002, MNRAS, 336, 112

[68] Visbal, E., Trac, H., & Loeb, A. 2011, JCAP, 8, 010

[69] Tegmark, M. 1997, Physical Review Letters, 79, 3806

[70] Jaffe, A. H., Kamionkowski, M., & Wang, L. 2000, Phys. Rev. D, 61, 083501

[71] Chiang, L.-Y., & Chen, F.-F. 2012, ApJ, 751, 43

[72] Coles, P., & Jones, B. 1991, MNRAS, 248, 1

[73] Di Matteo, T., Perna, R., Abel, T., & Rees, M. J. 2002, ApJ, 564, 576

[74] Wang, X., Tegmark, M., Santos, M. G., & Knox, L. 2006, ApJ, 650, 529

[75] Gong, Y., Chen, X., Silva, M., Cooray, A., & Santos, M. G. 2011, ApJ, 740, L20

[76] Comaschi, P., & Ferrara, A. 2016, MNRAS, 463, 3078

[77] Gong, Y., Silva, M., Cooray, A., & Santos, M. G. 2014, ApJ, 785, 72

[78] Lidz, A., & Taylor, J. 2016, ApJ, 825, 143

[79] Cheng, Y.-T., Chang, T.-C., Bock, J., Bradford, C. M., & Cooray, A. 2016, ApJ, 832, 165

[80] Wang, R., Carilli, C. L., Neri, R., et al. 2010, ApJ, 714, 699

[81] Gao, Y., & Solomon, P. M. 2004, ApJ, 606, 271

[82] Riechers, D. A., Walter, F., Cox, P., et al. 2007, ApJ, 666, 778

[83] Zhang, Z.-Y., Gao, Y., Henkel, C., et al. 2014, ApJ, 784, L31

[84] Riechers, D. A., Walter, F., Carilli, C. L., et al. 2006, ApJ, 645, L13

[85] Silva, M. B., Santos, M. G., Gong, Y., Cooray, A., & Bock, J. 2013, ApJ, 763, 132

[86] Ly, C., Malkan, M. A., Kashikawa, N., et al. 2007, ApJ, 657, 738

[87] Brax, P., Clesse, S., & Davis, A.-C. 2013, JCAP, 1, 003

[88] Eisenstein, D. J., Weinberg, D. H., Agol, E., et al. 2011, AJ, 142, 72

[89] Madau, P., & Dickinson, M. 2014, ARA&A, 52, 415

[90] Solomon, P. M., & Vanden Bout, P. A. 2005, ARA&A, 43, 677

[91] Miotello, A., Bruderer, S., & van Dishoeck, E. F. 2014, A&A, 572, A96

[92] Heyer, M. H., Carpenter, J. M., & Ladd, E. F. 1996, ApJ, 463, 630

[93] Binney, J., & Merrifield, M. 1998, Galactic Astronomy

[94] Wilson, T. L. 1999, Reports on Progress in Physics, 62, 143

[95] Pineda, J. L., Goldsmith, P. F., Chapman, N., et al. 2010, ApJ, 721, 686

[96] Henkel, C., Downes, D., Weiß, A., Riechers, D., & Walter, F. 2010, A&A, 516, A111

[97] Müller, H. S. P., Thorwirth, S., Roth, D. A., & Winnewisser, G. 2001, A&A, 370, L49

[98] Solomon, P. M., Rivolo, A. R., Barrett, J., & Yahil, A. 1987, ApJ, 319, 730

[99] Lagos, C. d. P., Bayet, E., Baugh, C. M., et al. 2012, MNRAS, 426, 2142

[100] Popping, G., van Kampen, E., Decarli, R., et al. 2016, MNRAS, 461, 93

[101] Behroozi, P. S., Wechsler, R. H., & Conroy, C. 2013, ApJ, 770, 57

[102] Tinker, J., Kravtsov, A. V., Klypin, A., et al. 2008, ApJ, 688, 709

[103] Tinker, J. L., Robertson, B. E., Kravtsov, A. V., et al. 2010, ApJ, 724, 878

[104] Sofia, U. J., Lauroesch, J. T., Meyer, D. M., & Cartledge, S. I. B. 2004, ApJ, 605, 272

[105] Kennicutt, Jr., R. C. 1998, ApJ, 498, 541

[106] Davis, T. A. 2014, MNRAS, 445, 2378

[107] Smith, R. L., Pontoppidan, K. M., Young, E. D., & Morris, M. R. 2015, ApJ, 813, 120

[108] Fisher, R. A. 1935, Journal of the Royal Statistical Society, 98, 39

[109] Jungman, G., Kamionkowski, M., Kosowsky, A., & Spergel, D. N. 1996, Physical Review Letters, 76, 1007

[110] —. 1996, Phys. Rev. D, 54, 1332

[111] Krumholz, M. R. 2012, ApJ, 759, 9

[112] Glover, S. C. O., & Clark, P. C. 2012, MNRAS, 426, 377

[113] Bartolo, N., Matarrese, S., & Riotto, A. 2010, Advances in Astronomy, 2010, 157079

[114] Scheuer, P. A. G. 1957, Proceedings of the Cambridge Philosophical Society, 53, 764

[115] Lee, S. K., Ando, S., & Kamionkowski, M. 2009, JCAP, 7, 007

[116] Lee, S. K., Lisanti, M., & Safdi, B. R. 2015, JCAP, 5, 056

[117] Barcons, X., Raymont, G. B., Warwick, R. S., et al. 1994, MNRAS, 268, 833

[118] Windridge, D., & Phillipps, S. 2000, MNRAS, 319, 591

[119] Patanchon, G., Ade, P. A. R., Bock, J. J., et al. 2009, ApJ, 707, 1750

[120] Glenn, J., Conley, A., Béthermin, M., et al. 2010, MNRAS, 409, 109

[121] Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2016, A&A, 594, A17

[122] Barcons, X. 1992, ApJ, 396, 460

[123] Takeuchi, T. T., & Ishii, T. T. 2004, ApJ, 604, 40

[124] Kayo, I., Taruya, A., & Suto, Y. 2001, ApJ, 561, 22

[125] Vernstrom, T., Scott, D., Wall, J. V., et al. 2014, MNRAS, 440, 2791

[126] Planck Collaboration, Adam, R., Ade, P. A. R., et al. 2016, A&A, 594, A10

[127] Negrello, M., Perrotta, F., González-Nuevo, J., et al. 2007, MNRAS, 377, 1557

[128] Paciga, G., Scott, D., & Chapin, E. L. 2009, MNRAS, 395, 1153

[129] Lima, M., Jain, B., & Devlin, M. 2010, MNRAS, 406, 2352

[130] Das, S., & Ostriker, J. P. 2006, ApJ, 645, 1

[131] Smith, R. E., Peacock, J. A., Jenkins, A., et al. 2003, MNRAS, 341, 1311

[132] Takahashi, R., Oguri, M., Sato, M., & Hamana, T. 2011, ApJ, 742, 15

[133] Barkana, R., & Loeb, A. 2008, MNRAS, 384, 1069

[134] Shimabukuro, H., Yoshiura, S., Takahashi, K., Yokoyama, S., & Ichiki, K. 2015, MNRAS, 451, 467

[135] Jaacks, J., Nagamine, K., & Choi, J. H. 2012, MNRAS, 427, 403

[136] Obreschkow, D., Heywood, I., Klöckner, H.-R., & Rawlings, S. 2009, ApJ, 702, 1321

[137] Kistler, M. D., Yüksel, H., Beacom, J. F., Hopkins, A. M., & Wyithe, J. S. B. 2009, ApJ, 705, L104

[138] Popping, G., Pérez-Beaupuits, J. P., Spaans, M., Trager, S. C., & Somerville, R. S. 2014, MNRAS, 444, 1301

[139] Liu, A., Zhang, Y., & Parsons, A. R. 2016, ApJ, 833, 242

[140] Peebles, P. J. E. 1980, The large-scale structure of the universe

Vita

Patrick C. Breysse was born in January 1991. He received B.S. degrees in Physics and Astronomy and Astrophysics from the Schreyer Honors College at the Pennsylvania State University in 2012. He subsequently enrolled in the Physics and Astronomy Ph.D. program at Johns Hopkins University, where he was awarded the Max Bodmer Graduate Research Fellowship. He won an Astronaut Scholarship in 2011 and was named the Astronomy and Astrophysics Student Marshall at Penn State in 2012. His research covers a wide variety of topics in theoretical cosmology with a particular focus on the theory behind intensity mapping observations. He was invited to present his work on intensity mapping at the 66th Lindau Nobel Laureate Meeting.

In September 2017, Patrick will start a postdoctoral fellowship at the Canadian Institute for Theoretical Physics.