Studying the largest scales in the Universe with the kinetic Sunyaev-Zel'dovich effect

by

Juan Ignacio Cayuso

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Doctor of Philosophy
in
Physics

Waterloo, Ontario, Canada, 2021

Examining Committee Membership

The following served on the Examining Committee for this thesis. The decision of the Examining Committee is by majority vote.

External Examiner:	Gilbert Holder Professor Physics, University of Illinois
Supervisor(s):	Matthew Johnson Adjunct Faculty Physics & Astronomy, University of Waterloo Associate Professor Physics & Astronomy, York University
	Niayesh Afshordi Associate Professor Physics & Astronomy, University of Waterloo
Internal Member:	Michael Hudson Professor Physics & Astronomy, University of Waterloo
Internal Member:	Kendrick Smith Adjunct Assistant Professor Physics & Astronomy, University of Waterloo
Internal-External Member:	Florian Girelli Associate Professor Applied Mathematics, University of Waterloo

Author's Declaration

This thesis consists of material all of which I authored or co-authored: see Statement of Contributions included in the thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Statement of Contributions

This thesis contains edited material from research papers which I authored or co-authored, as well as material from textbooks and other public sources. Chapter 1 serves as an introduction and review of relevant material, which I have collected and written myself. Chapter 2 was co-written with Richard Bloch, Selim C. Hotinli, Matthew C. Johnson, and Fiona McCarthy and published in [1] . Chapter 3 was co-written with Matthew C. Johnson and James B. Mertens and was published in [2]. Chapter 4 was co-written with Matthew C. Johnson and was published in [3].

Abstract

As we progress further into the era of precision cosmology, new avenues to test our fundamental models of the Universe are opening up. This PhD thesis is concerned with the development and understanding of a technique known as kinetic Sunyaev-Zel'dovich (kSZ) velocity reconstruction, which aims to extract information about the Universe on the largest accessible scales using measurements of the Cosmic Microwave Background (CMB) anisotropies sourced by the kSZ effect and data from galaxy redshift surveys. kSZ velocity reconstruction estimates the remote CMB dipole, i.e. the $\ell = 1$ multipole moment of the observed CMB sky as seen by observers on our past lightcone. This observable probes cosmological perturbations on scales of several Gpc, and thus has the potential to be a valuable source of information on the fundamental early Universe phenomena that leaves imprints on such scales. Preliminary forecasts from the foundational literature on kSZ velocity reconstruction indicate that high signal to noise reconstructions of the remote dipole will be possible in the context of next generation CMB experiments and galaxy surveys. The goal of this thesis is to further develop the technical details of the technique and provide motivation for its use as a tool to probe physics on ultra-large scales. Chapter 1 elaborates on the motivation behind this thesis and provides a review of the key material necessary to understand kSZ velocity reconstruction. Chapter 2 presents an extended formalism for kSZ velocity reconstruction, which describes new sources of noise and bias and incorporates more realistic experimental conditions. Forecasts for the reconstruction of the remote dipole are presented in the same chapter, and these show that high signal to noise is still achievable using the new estimators. Chapter 3 presents the first suite of N-body simulations of remote dipole reconstruction on our past lightcone, which implement a novel methodology to treat the wide range of scales involved in kSZ velocity reconstruction (tens of Mpc to tens of Gpc). These simulations were used to test the robustness of the reconstruction technique against the effects of gravitational non-linearities, redshift space distortions, and CMB lensing. Additionally, these simulations were used to demonstrate the relevance of large scale contributions to the remote dipole that are not captured by other approaches to kSZ velocity reconstruction. Chapter 4 presents an analysis of parameter constraints on CMB anomalies models, aimed to demonstrate that the reconstructed remote CMB dipole and the reconstructed remote CMB quadrupole (obtained using a similar technique to kSZ velocity reconstruction) can help us go beyond the constraints achievable with traditional probes of the anomalies like the primary CMB temperature, primary CMB polarization, and large-scale galaxy distribution.

Acknowledgements

I would like to thank all the people who made this thesis possible. On the academic side, I would like to thank Matthew Johnson for his guidance and generosity as my academic advisor. As his student, I felt heard and encouraged to be my best, and I'll always be grateful for that. I would like to thank all the faculty, postdocs and students with whom I collaborated or had productive discussions with surrounding my research. I would like to thank the University of Waterloo and the Perimeter Institute for Theoretical Physics for supporting me in my studies and research, and for providing a vibrant, enthusiastic environment to tackle the greatest questions about physics and our Universe.

On the personal side, I would like to thank my parents, Marisa and Daniel, who have always believed in and been there for me in every big step in my life. I truly was blessed with amazing parents. I would like to thank my siblings, Ramiro, Renata and Laureana, as well as the rest of my family, for their unconditional love and support. I would like to thank Natalie, my partner, for her love and companionship and for helping me become a better person every day. I would like to say thanks to the horde of friends that have been there for me during my studies. To my friends from Cordoba: Alejandra, Julian, Marina, Pedro, Natalia and Barbara, who have brightened my days since I started my journey in physics; to the "kabros of the Polis", Ramiro, Riki, Alex and Eric, for our brotherhood and unconditional support. To Santiago and Belen, for making Tuesday an extra special day. To Natacha Altamirano, the person who motivated me to come to Canada, opened my eyes to the beauty of cosmology and gifted me with her amazing friendship. To Barbara, Lucia, Fiona, Tomas, Richard, Sara and countless other friends who made Perimeter Institute a great place to be at.

Dedication

This is dedicated to my grandmothers, Marta and Martha.

Table of Contents

List of Figures

List of Tables

Chapter 1

Introduction and review

1.1 Motivation

Cosmology, the field of research that studies the evolution and structure of the Universe, has gone though an unprecedented period of growth in the last several decades as a consequence of our ability to observe the cosmos with increasing amount of detail. The collection of observations we have gathered so far, especially those from the *Planck* satellite mission measuring the Cosmic Microwave Background radiation (CMB) [4], has lead us to establish a **baseline** or standard model of Cosmology: Lambda Cold Dark Matter or ΛCDM for short. Armed with only six parameters, ΛCDM offers the most consistent explanation for **current** observational data, including statistical properties of the CMB sky, large-scale distribution of galaxies and expansion history of the Universe. The words *baseline* and *current* are intentionally highlighted; despite its success, ΛCDM is just a model and, if history is any indication, it will face some shortcomings under the scrutiny of more precise data.

There already exist some indications that the standard model might be incomplete or not fully consistent with observations. On large angular scales, the CMB sky presents a series of anomalous features that could indicate deviations from the standard slow roll inflation and violations to statistical isotropy on large scales [5]. Puzzling as well is the tension between measurements of the Hubble constant derived from CMB observations and those derived from distance measurement of standard candle supernovae [6]. The weight and significance of these discrepancies is still under debate. The CMB anomalies are only marginally significant in the statistical sense due to cosmic variance on large angular scales, and current analysis has not been conclusive enough to point towards a particular early

Universe model or mechanism for isotropy breaking that could explain these anomalies [7]. Regarding the Hubble tension, it is possible that unaccounted for large-scale structures or multiple systematic errors could be behind the discrepancy. Still, many proposals for new physics both at late and early time have been put forward as an alternative explanation [8]. The belief that ΛCDM is an incomplete model of the Universe also finds support on the theory front. Two crucial components of the model, dark matter and dark energy in the form of a cosmological constant, are far from being completely understood. It is also a possibility that the introduction of these dark substances is an indication that General Relativity, a central piece of ΛCDM, is not entirely correct and it fails to describe the effects of gravity on cosmological scales.

In order to decidedly identify the shortcomings of the ΛCDM model, larger and more detailed datasets are needed. Fortunately, we find ourselves on the brink of an era of promising experiments in Cosmology. Tens of millions to billions of galaxies will be observed with instruments like the Dark Energy Spectroscopic Instrument (DESI) [9] , the Vera C. Rubin Observatory (Rubin) [10], and the EUCLID satellite [11]. The 3-dimensional maps constructed in these experiments will allow us to further constraint the physics of dark energy and modified gravity, study the nature of massive neutrinos, better understand the clustering properties of dark matter and look for imprints of inflation on large angular scales. The Canadian Hydrogen Intensity Mapping Experiment (CHIME) [12] and the Hydrogen Intensity and Real-time Analysis eXperiment (HIRAX) [13, 14] radio telescopes will measure the properties of large-scale structures and dark energy not by observing individual galaxies, but by mapping the large-scale 21 cm neutral hydrogen emission, a technique known as Hydrogen Intensity Mapping. This technique measures linear perturbation modes without the need of measuring individual galaxy redshifts, which will result in deeper and wider survey volumes. A new frontier of high precision CMB experiments is also in our immediate future. The Simons Observatory collaboration (SO) [15] and the fourth-generation ground-based CMB experiment (CMB-S4) [16] will measure the temperature and polarization of the CMB to new levels of sensitivity and pursue a series of exciting science goals: to search for gravitational wave signatures from inflation in the B-mode polarization of the CMB; to improve the reconstruction of the CMB lensing potential (a powerful probe of the matter distribution between us and the last scattering survey), which will allow for better constraints on primordial gravitational waves from delensing, neutrino masses, dark energy and high redshift astrophysics; to further characterize the primordial spectrum of scalar perturbations and achieve stronger constraints on inflation from the CMB bispectrum; to provide new measurements of the Hubble constant H_0 at higher statistical significance than previous probes and using multiple observables (temperature and E-mode polarization auto and cross-spectra) for a strong check of consistency; and to

measure the CMB anisotropies arising from the thermal and kinetic Sunyaev-Zel'dovich effects to high levels of sensitivity.

Each one of these experiments by themselves will provide rich data to further test the validity of ΛCDM. Still, much more information is to be gained if we adopt a **multi-tracer** approach to data analysis by combining the observations from different experiments. For example, the reconstructed lensing potential from SO and the galaxy maps from Rubin can be combined to determine the galaxy bias on large scales [15] (essentially, the difference in clustering properties of galaxies and dark matter). The scale dependence of this bias can reveal the presence of local primordial non-Gaussianities, a signature of multi-field inflation models. When it comes to studying large-scale properties of the Universe, multi-tracer analysis offers a powerful mechanism to surpass the limitations of cosmic variance present in single-probe analysis: sample variance cancellation [17]. For example, given two biased tracers of the same underlying dark matter perturbation, a multi-tracer analysis yields better constraints on the relative bias by directly comparing modes compared to what is achievable from constraining each bias independently. Apart from the benefits of sample variance cancellation, combining multiple observables that probe perturbations on ultra-large scales can serve as a consistency check for the existence of features like the CMB anomalies; if the anomalies are truly of primordial origin, counterpart imprints should be present in other probes.

Finding new observables that can inform us of the largest accessible scales will certainly increase the benefits of multi-tracer analysis. An interesting approach to do this, which doesn't rely on direct observations of long-range patterns in wide cosmic volumes or angles on the sky, is to **reconstruct large-scale fields using statistical properties of small-scale fields**. The canonical example of this approach is the reconstruction of the CMB lensing potential: the statistical properties of the small-angle temperature anisotropies generated by gravitational lensing can be used to infer the large-angle features of the lensing potential. This thesis is concerned with the study and exploration of a similar technique, here referred to as **kinetic Sunyaev-Zel'dovich velocity reconstruction**, also known in previous literature as **kSZ tomography** or **remote dipole reconstruction** [18, 19, 20, 21]. In simple terms, the idea is to use measurements of the kinetic Sunyaev-Zel'dovich (kSZ) effect together with tracers of the electron distribution on our past lightcone in order to reconstruct information about the peculiar velocity of these electrons on large scales. The principle behind kSZ velocity reconstruction is the following: the kSZ signal originates from the Doppler shift of CMB photons when they are scattered by electrons with peculiar motion respect to the CMB rest frame [22, 23]. This process leads to new temperature anisotropies and the accumulation of these along our past lightcone sums up to become the the dominant contribution to the blackbody spectrum of the CMB for

multipoles $\ell \gtrsim 4000$ (scales of a few arcminutes). One can show that the cross-correlation between the kSZ signal and a tracer of the electron distribution on small scales carries information about the large-scale bulk motion of the electrons projected along the line of sight. In analogy with lensing potential reconstruction, this velocity information can be recovered by constructing quadratic estimators based on the statistical properties of the CMB-LSS correlation. The tracer of the electron distribution (a galaxy survey for example) can be a 3-dimensional probe separated into multiple redshift bins. If this is the case, a tomographic separation of the total kSZ signal can be achieved by cross-correlating the temperature with the electron tracers at multiple redshifts (a technique known as kSZ tomography [24, 25, 26, 27, 28]) and kSZ velocity reconstruction can be applied for each one of these set of correlations. The result is a collection of sky maps carrying the large-angle features of the radial peculiar velocity on our past lightcone. The large-scale velocity field directly traces the underlying dark matter density field and thus can be a powerful additional probe of the large-scale perturbations in the Universe and their evolution.

The new generation of CMB and galaxy survey experiments will bring an exciting opportunity to apply kSZ velocity reconstruction and this has motivated a deeper study into its technical details, as well as its potential. Several approaches have been developed to formalise the reconstruction procedure, test its validity with simulations, and forecast its capabilities in next generation experiments. The research presented in this thesis utilizes the "Lightcone Picture" of kSZ velocity reconstruction, introduced in Refs.[18, 19]. This approach incorporates the geometry of the past lightcone and its projection effects on observables. A detailed discussion about its advantages and disadvantages with respect to other approaches will be provided later in the text. Preliminary forecasts for a Lightcone Picture reconstruction indicate that high signal to noise can be achieved for the large-angle radial velocity for a variety of redshift bin configurations, assuming SO-like measurements of the kSZ effect and Rubin-like galaxy maps [19]. These results, although promising, were obtained under a series of simplifying assumptions both at the level of the experimental conditions and the theoretical modelling of the reconstruction machinery. The research presented in this thesis expands on the baseline results from [19] in multiple ways. More realistic experimental conditions are considered by including frequency dependent temperature foregrounds and instrumental noise, incorporating the effects of photometric redshift errors in the galaxy survey, and evaluating the impact of masking of the sky. On the theory side, an extended formalism for the Lightcone Picture is presented. This formalism brings an improved understanding of sources of bias to the velocity reconstruction as well as the effects of coarse-graining of 3-dimensional information on the lightcone. Moreover, the theoretical tools presented are general enough to be applied to other CMB-LSS cross-correlation pairs that carry information about large scale fields, and have been implemented

in the publicly available code " ReCCO".

In addition to the extended Lightcone Picture formalism, this thesis presents two other studies on kSZ velocity reconstruction. The first one consists on a validation of the principles behind the technique using N-body simulations. A challenging aspect of simulating velocity reconstruction on the lightcone is the wide range of scales that need to be simulated: from several Gpc scales for the large-scale velocity field, down to Mpc scales for the small-scale density field. Evolving such large volumes on the lightcone with enough small-scale resolution is currently a computationally intractable task. To work around this challenge, the simulations presented in this text utilize a "box-in-box" approach: an N-body simulation in a Gpc size volume is embedded in a larger volume containing large-scale modes evolved using linear theory. This approach not only facilitates the construction of properly correlated temperature and density maps that contain the information from large-scale modes, but also allows for the inclusion of a contribution to the kSZ signal coming from the intrinsic CMB dipole, a feature that is easily described in the Lightcone Picture. Overall, these simulations served as a test of robustness against previously neglected effects such as gravitational non-linearities, redshift space distortions and CMB lensing. The second study presented in this thesis explores the potential of using the reconstructed large-scale velocity to further constraint some of the CMB anomalies models. Any physical model constructed to explain the existence of anomalous features in the CMB temperature should produce predictable signatures on other observables, such as the large-angle CMB polarization or the large-angle distribution of distant galaxies [29, 30, 31, 32, 33, 34]. The objective of the study presented here is to determine if constraints on anomaly models coming from joint analysis of CMB temperature, CMB E-mode polarization, and galaxy surveys can be improved by adding the reconstructed velocity. This study also considered a second observable that can be added to the list: the remote CMB quadrupole. When CMB photons are scattered by free electrons, a polarization signal arises and this signal depends on the CMB quadrupole observed by the electrons [35]. This remote CMB quadruple can be reconstructed on the lightcone using an appropriate modification of kSZ velocity reconstruction [36], and has been identified as a promising probe of the largest scales that can be observed [37].

This thesis is organized as follows: the rest of Chapter 1 consists of a review of the key concepts involved in kSZ velocity reconstruction. Chapter 2 presents the extended Lightcone Picture of kSZ velocity reconstruction and an analysis of the new contributions to the formalism. Chapter 3 presents the "box-in-box" simulations of kSZ velocity reconstruction on the lightcone using N-body simulations. Chapter 4 analyses improvement on parameter constraints in CMB anomaly models that can be achieved by using the reconstructed large-scale velocity in combination with standard probes of the anomalies. Chapter 5 sum-

marizes the results of this thesis and discusses future goals for the development of kSZ velocity reconstruction.

1.2 Conventions

1.2.1 Units

Unless stated otherwise, we work in natural units with $\hbar = c = G_N = 1$.

1.2.2 Lightcone parametrization

Our past lightcone defines a 3-dimensional surface on spacetime. We parametrize events on the lightcone using a unit direction vector $\hat{\mathbf{n}}$ and a radial coordinate χ, corresponding to the comoving distance from the origin (us) to the event. With this parametrization, $(\hat{\mathbf{n}}, \chi)$ represents the spacetime event $(\eta(\chi), \mathbf{x})$, where $\eta(\chi)$ is the conformal time of the event and $\mathbf{x} = \chi(\eta)\hat{\mathbf{n}}$ is the spatial comoving position of the event. The value of a field F at the spacetime event $(\eta(\chi), \mathbf{x})$ is expressed as $F(\hat{\mathbf{n}}, \chi)$.

1.2.3 Gauge choice

Where necessary, we present relations in the Newtonian gauge, where at late-times when we can neglect anisotropic stress, the metric is:

$$ds^2 = a(\eta)^2 \left(- \left[1 + 2\Psi(\eta, \mathbf{x})\right] d\eta^2 + \left[1 - 2\Psi(\eta, \mathbf{x})\right] d\mathbf{x}^2 \right) \tag{1.1}$$

1.3 The primary and secondary CMB anisotropies

The Cosmic Microwave Background is the oldest electromagnetic signal that we can observe, a remnant of the hot and dense beginnings of our Universe. Originating around 13.8 billion years ago, it is detected today as an extremely isotropic radiation coming from every direction on the sky, with tiny anisotropies that can be resolved with sensitive enough instruments. The origin of the CMB can be traced back to the epoch of recombination, when the first neutral hydrogen atoms in the Universe started forming and photons decoupled from the baryonic matter. At earlier times, photons would be tightly coupled to free

electrons via Compton scattering and couldn't travel unimpeded for long distances. Once the ratio of free electrons to photons had drastically decreased due to recombination, light was able to propagate freely and carry with itself an imprint of the small perturbations from homogeneity that existed in the early Universe. The process of decoupling occurred relatively fast in cosmic terms, defining a region of space-time known as the last-scattering surface. The perturbations on the photon field shortly after decoupling can be modelled as a combination of: a primordial component with origins in some high energy process in the very early Universe; the effects of gravity captured by the Sachs-Wolfe effect and the early Integrated Sachs-Wolfe (ISW) effect; the Doppler effect arising from peculiar motions of the last-scattering surface sourced by density perturbations; the imprints of acoustic oscillations in the photon-baryon fluid; and small-scale damping of perturbations due to imperfect coupling between baryon and photons during recombination and due to the non-negligible thickness of the last-scattering surface. The anisotropies associated to these perturbations are referred to as the **primary CMB anisotropies**, or primary CMB for abbreviation.

The different sources of primary CMB anisotropies can be more or less relevant depending on the angular scale of observation. For large-angles, anisotropies are dominated by gravitational and Doppler effects, which offer a cleaner window into the primordial perturbations. As smaller angular scales are inspected, the imprints from the acoustic oscillations and damping start dominating and drastically depart from the primordial perturbations. This last fact does not render the small-angle anisotropies useless or uninteresting for Cosmology; to the contrary, the imprints from the photon-baryon physics have been crucial sources of information to determine the relative proportions of mark matter and baryonic matter in the Universe. When we measure the CMB today, part of what we see are these primary anisotropies from the time of decoupling, but we also see anisotropies that are sourced by multiple physical processes the CMB photons go through in their journey from the last-scattering surface to our location. These additional anisotropies are referred to as the **secondary CMB anisotropies**, or secondary CMB for abbreviation. An example of these are the temperature anisotropies sourced by gravitational lensing of the CMB: the path of photons travelling through the Universe can be deflected by intervening matter, leading to modifications of the primary CMB anisotropies on arcminute scales.

The secondary CMB signals provide new ways to study cosmological perturbations through their ability to probe large volumes of our observable Universe. In contrast to the primary CMB, which is mostly sourced in a 2-dimensional surface, the secondary CMB carries 3-dimensional integrated information, thus probing more linear modes. Moreover, these line of sight integrated signals offer an additional avenue to probe the expansion history of the Universe and evolution of large-scale structure. We continue our discussion

of the CMB secondaries by briefly reviewing the different types of signals that arise when primary CMB photons travel across the Universe. We will pay a special amount of detail to the kSZ signal as it is a central piece of this thesis.

1.3.1 The late-time ISW effect

Photons travelling through time evolving gravitational potentials suffer a shift in energy due to the unequal energy modifications from entering and exiting the potentials. We can write this signal as a line of sight integral on our past lightcone:

$$\Theta^{ISW}(\hat{\mathbf{n}}) = -2 \int_0^{\chi_{ls}} d\chi \, \frac{d\Psi}{d\chi}(\hat{\mathbf{n}}, \chi) \tag{1.2}$$

where χ_{ls} is the conformal distance to last scattering. An early component of the ISW effect is sourced around the time of decoupling, when the Universe is not fully matter dominated and the gravitational potentials decay under the influence of radiation. Once the Universe becomes matter dominated, linear perturbations of the density grow as fast as the scale factor, leading to non-evolving gravitational potential and therefore no ISW effect. The situation changes either when dark energy starts dominating the expansion of the Universe or when non-linear perturbations of the gravitational potential start to grow considerably. These conditions lead to the late-time ISW effect. The classical approach to detect the linear ISW effect is to cross-correlate CMB temperature maps with large-scale structure tracers like galaxies, a technique first implemented in [38]. This detection approach, as well as others using correlations with the CMB lensing potential and stacking CMB anisotropies associated with supervoids and superclusters, have been used with more recent data from the *Planck* mission and a series of large scale structure surveys [39]. The linear ISW effect leads to a small additional signal to the CMB power spectrum on the largest angular scales we observe. Understanding and properly modelling the ISW effect allows for a clearer differentiation of this signal from early Universe contributions to the large angle CMB. The linear ISW signal is also interesting for its sensitivity to the properties of dark energy, as it can serve as a complementary source of (not extremely competitive) constraints on dark energy model parameters.

The non-linear extension of the late-time ISW effect is often referred as the Rees-Sciama (RS) effect [40]. When density perturbations evolve non-linearly, the gravitational potentials change differently than they would if they where only under the effect of dark energy. The signatures generated by the RS effect have been studied using N-body simulations in [41]. The authors identify that the origin of the non-linear ISW effect is mostly sourced

by the non-linear relation between the density field and the velocity field:

$$\Theta^{ISW-NL}(\hat{\mathbf{n}}) \approx 2 \int_0^{\chi_{ls}} d\chi \; \nabla_{\mathbf{x}} \Psi(\hat{\mathbf{n}}, \chi) \cdot \mathbf{v}(\hat{\mathbf{n}}, \chi) \tag{1.3}$$

which manifests as ring like features surrounding clusters and voids, and dipole features sourced by the transverse motion lumps of dark matter respect to the line of sight. This last signature is usually referred in the literature as the moving lens effect [42, 43, 44, 45]. The temperature anisotropies coming from the non-linear ISW start becoming important respect to the linear ISW ones for multipoles $\ell > 80$ and are completely dominant for $\ell > 200$. A detection of the RS effect is still to be achieved. According to another study using N-body simulations [46], detecting the RS signal through cross-correlation with LSS tracers is difficult at low redshifts ($z < 4$) due to the small contribution to the signal coming from those redshifts, and a lack of suitable LSS tracers makes the task difficult at higher redshifts ($4 < z < 30$) where most of the signal is sourced. A direct detection of the RS signal in the CMB temperature power spectrum is also difficult due to its small amplitude respect to the primary anisotropies and other secondary anisotropies like CMB lensing and the kSZ effect.

1.3.2 CMB lensing

As CMB photons travel across the Universe, their paths can be deflected by the gravitational influence from intervening matter. These deflections are integrated along the line of sight and translate to distortions of the observed CMB temperature and polarization. For the temperature we have:

$$\Theta^{L}(\hat{\mathbf{n}}) = \Theta^{pCMB}\left(\hat{\mathbf{n}} + \alpha(\hat{\mathbf{n}})\right) \tag{1.4}$$

where the deflection angle can be expressed in terms of a lensing potential ϕ:

$$\alpha(\hat{\mathbf{n}}) = \nabla_{\perp}\phi(\hat{\mathbf{n}}) \quad , \quad \phi(\hat{\mathbf{n}}) = -2 \int_0^{\chi_{ls}} d\chi \; \frac{\chi_{ls} - \chi}{\chi_{ls}\; \chi} \Psi(\hat{\mathbf{n}}, \chi) \tag{1.5}$$

where $\nabla_{\perp}$ is the angular gradient in the unit 2-sphere. In terms of the lensing potential, the lensed CMB can be expanded as:

$$\Theta^{L}(\hat{\mathbf{n}}) \;\approx\; \Theta^{pCMB}(\hat{\mathbf{n}}) + \nabla_{\perp}\Theta^{pCMB}(\hat{\mathbf{n}}) \cdot \nabla_{\perp}\phi(\hat{\mathbf{n}}) \tag{1.6}$$

and a similar expansion can be written for the polarization using appropriate tensor objects. Understanding the signatures from CMB lensing is of importance for cosmology for several

reasons. Lensing leads to a B-mode polarization signal that can dominate over a possible polarization signal originated by primordial gravitational waves from inflation. The effects of lensing can be alleviated by "delensing procedures", which can undo the distortions in the CMB sky if a good enough estimate of the lensing potential is available. Interestingly enough, an estimate of the lensing potential can be obtained using the lensed CMB itself: the non-Gaussianities induced by lensing can be used to construct a quadratic estimator of the lensing potential [47]. This type of estimation procedure, which takes small angular data to reconstruct large-angle features of an underlying field, will be exemplified in high detail later in the text in the context of kSZ velocity reconstruction. Apart from facilitating delensing, the reconstructed lensing potential is also a useful probe of cosmology due to its simple relation to the matter power spectrum. It can be used in cross-correlation with other probes of dark matter to better characterize the matter power spectrum and its evolution, as well as providing stronger constraints on relative large-scale biases between probes through sample variance cancellation [17].

1.3.3 Thermal Sunyaev-Zel'dovich effect

Primary CMB photons can be scattered into our line of sight by free electrons on our past lightcone. The scattered CMB radiation from a cloud of hot electrons suffers a small distortion of its original blackbody energy spectrum. Accumulated along the line of sight, these distortions translate to a change in CMB temperature given by:

$$\Theta^{tSZ,\nu}(\hat{\mathbf{n}}) = f(x)y(\hat{\mathbf{n}}) \tag{1.7}$$

where x is the dimensionless frequency parameter $x = \frac{h\nu}{k_B T_{CMB}}$, the frequency dependence is $f(x) = x\coth(\frac{x}{2}) - 4$, and $y(\hat{\mathbf{n}})$ is the Compton-y parameter, given by:

$$y(\hat{\mathbf{n}}) = \int d\chi\ a(\chi)\ n_e(\hat{\mathbf{n}}, \chi)\frac{k_B T_e(\hat{\mathbf{n}}, \chi)}{m_e}\sigma_T \tag{1.8}$$

where $n_e(\hat{\mathbf{n}}, \chi)$ is the electron number density, $T_e(\hat{\mathbf{n}}, \chi)$ is the electron temperature, m_e is the electron mass, and σ_T is the Thompson cross-section. In contrast to the late-time ISW and lensing signals, the tSZ effect can be separated from the blackbody CMB due to its particular frequency spectrum. The tSZ signal has been identified as an extremely useful probe to find galaxy clusters and characterize their abundances. One of the reasons is that it facilitates the detection of clusters at higher redshifts than other sources: because the tSZ is a fractional change in the CMB temperature, the amplitude of the change is redshifted along with the baseline temperature of the CMB which leads to a surface

brightness that does not fade as $(1+z)^{-4}$. The tSZ signal provides other advantages for the characterization of cluster abundance: the selection function as a function of redshift is generally well behaved and, because the signal probes the intra-cluster medium, a larger portion of the baryonic matter can be detected compared to observations of objects like galaxies. This makes the tSZ signal a better proxy of the total cluster mass. Using methods to calibrate the observable-mass relation, the cluster abundances measured with tSZ can be used to constrain cosmology using forward modelling techniques [48, 49, 50]. Near future measurements of the tSZ effect by the SO collaboration forecast the detection of tens of thousands of clusters, roughly an order of magnitude more than current datasets, as well as measurements of the tSZ power spectrum with total signal to noise ≈ 250, which can be used as an additional route to measure the sum of neutrino masses and other cosmological parameters [15, 51, 52] .

1.3.4 Kinetic Sunyaev-Zel'dovich effect

The scattering of primary CMB photons by free electrons can lead to additional signals apart from the one described by the tSZ effect. Sunyaev and Zel'dovich predicted that the bulk peculiar velocity of a cloud of electrons would lead to new temperature anisotropies in the scattered CMB light [22, 23]:

$$\frac{\Delta T_{kSZ}}{T_{CMB}} = -v_c \tau_c \tag{1.9}$$

where τ_c is the optical depth of the cloud and v_c is the peculiar motion of the cloud, projected along the line of sight. These Doppler shifts of energy preserve the Planckian spectrum of the CMB radiation, and manifest in the temperature map as cold (hot) spots for clusters with positive (negative) line of sight velocity. When integrated along the line of sight, the kSZ signal is given by:

$$\Theta^{kSZ}(\hat{\mathbf{n}}) = -\int d\chi \ \ (\mathbf{v}(\hat{\mathbf{n}}, \chi) \cdot \hat{\mathbf{n}}) \ \ \dot{\tau}(\hat{\mathbf{n}}, \chi) \tag{1.10}$$

where $\dot{\tau}(\hat{\mathbf{n}}, \chi)$ is the differential optical depth:

$$\dot{\tau}(\hat{\mathbf{n}}, \chi) = \sigma_T a(\chi) n_e(\hat{\mathbf{n}}, \chi) \tag{1.11}$$

and $\mathbf{v}(\hat{\mathbf{n}}, \chi)$ is the peculiar velocity field along our past lightcone. A correction to this temperature distortion can come if the CMB sky observed by the electron, which we refer as the **remote CMB**, has contributions to the dipole anisotropy apart from the velocity

induced kinematic dipole. This "non-kinematic" or "stationary" dipole, which we will describe in detail in Sec.1.5, is typically small compared to the largely dominant peculiar velocity sourced term (except on scales of several Gpc), and we will refer to the combination of both sources as an "effective" velocity. The projection of the effective velocity along the line of sight leads to the temperature anisotropy:

$$\Theta^{kSZ}(\hat{\mathbf{n}}) = -\int d\chi \; v_{\rm eff}(\hat{\mathbf{n}}, \chi) \; \dot{\tau}(\hat{\mathbf{n}}, \chi) \tag{1.12}$$

with

$$v_{\rm eff}(\hat{\mathbf{n}}, \chi) = \frac{3}{4\pi} \int d^2\hat{n}_e \; \Theta(\mathbf{x}, \hat{\mathbf{n_e}}) \; (\hat{\mathbf{n}} \cdot \hat{\mathbf{n_e}}). \tag{1.13}$$

where $\Theta(\mathbf{x}, \hat{\mathbf{n_e}})$ is the remote CMB at the spacetime position $(\eta(\chi), \mathbf{x} = \chi\hat{\mathbf{n}})$, containing both stationary and kinematic dipole anisotropy as observed by the electrons. Due to its small contribution to the total integrated signal, the stationary remote dipole contribution is often not accounted for in studies of the kSZ signal. In Sec.1.5, we will argue that taking into consideration the stationary dipole terms is important when using kSZ velocity reconstruction to probe of the homogeneity of the Universe on ultra large scales. For the moment, we concentrate on the dominant peculiar velocity term. The lightcone geometry for the kSZ effect is depicted in Fig.(1.1).

Given that the kSZ effect arises due to the presence of free electrons, this signal is sensitive to the history of reionization. In fact, the kSZ signal generated from a fully reionized Universe has different characteristics that the one arising as the Universe reionizes [53], and this has motivated a separation between the so called reionization kSZ and the late-time kSZ. Measuring and distinguishing these two signals would bring valuable cosmological information: the reionization kSZ can inform us about the patchiness and duration of reionization, while the late-time kSZ can serve as a probe of the large scale velocity field and distribution of electrons in dark matter halos. Detecting the kSZ signal from individual clusters can be challenging for several reasons: for large clusters the tSZ signal can be dominant over the kSZ and removing it requires sensitive multi-frequency measurements in addition to modelling of the electron environment in the cluster; for smaller clusters, where kSZ and tSZ can be comparable, the sensitivity of current experiments is not enough for a clear separation. Beyond individual cluster signals, the kSZ effect could also be detected statistically by measuring its angular power spectrum. For $\ell \gtrsim 4000$, the blackbody CMB is dominated by the kSZ effect, with same order of magnitude contributions from the reionization and late-time signals. Although current measurements of the kSZ power spectrum are of low significance [54], next generation CMB experiments with arcminute resolution will allow for precise measurements of the power spectrum, assuming successful enough foreground removal and degeneracy breaking with the primary CMB [55].

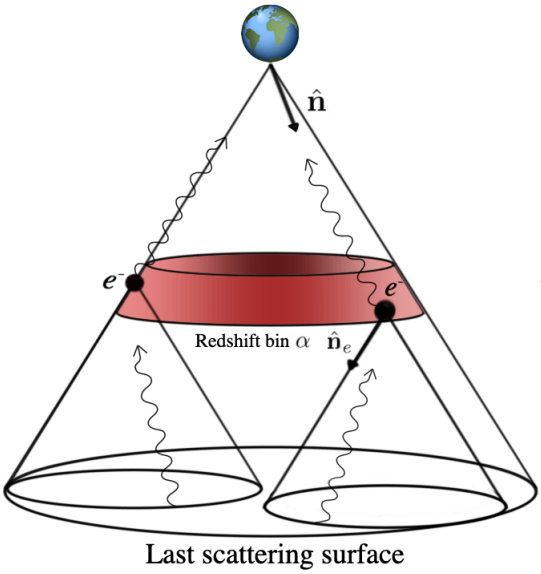

Figure 1.1: Points along the lightcone are parametrized by an angular direction $\hat{\mathbf{n}}$ and the comoving distance χ. Electrons at spacetime position $(\eta(\chi), \mathbf{x} = \chi\hat{\mathbf{n}})$ observe the remote CMB sky $\Theta(\mathbf{x}, \hat{\mathbf{n}_e})$, and will scatter some of the incoming CMB into our line of sight. The temperature distortions induced by the kSZ effect are sourced by the remote CMB dipole observed by the electrons, projected along the line of sight $\hat{\mathbf{n}}$.

A whole new world of possibilities for detecting and using the kSZ signal opens up if we have external measurements of a tracer of the electron distribution on our past lightcone, like a galaxy redshift survey for example. Evidence for the late-time kSZ effect has been obtained using statistics based on the pairwise motion of galaxy clusters [56, 57, 58, 59]: gravity induces correlations on the relative bulk velocities of galaxy clusters as a function of their comoving distance. These relative motions between galaxy clusters lead to patterns in the CMB temperature that can be picked up by specially tailored statistics. A galaxy survey can also be used to create a template for the kSZ signal that can later be cross-correlated with the temperature to pick up the kSZ effect [25, 26, 27, 28]. Using the galaxy distribution as a tracer of dark matter, one can estimate the large-scale velocity using the linear relation between velocity and density perturbations. The estimated velocity and the galaxy distribution, together with some modelling of the relation between electron and galaxy perturbations, can then be used to construct a template of the kSZ signal and cross correlate it with the CMB to yield a detection. A similar statistical approach based on estimated velocities was derived in [60], where the temperature field is filtered in a

particular way that enhances its small correlation with the velocity field at each cluster location. The different statistical approaches described above (collectively referred as to kSZ tomography) were eventually found to be connected in [20], were it is shown that they are equivalent formulations of an optimal bispectrum estimator with two powers of the galaxy field and one power of the temperature field $\langle ggT \rangle$. The technique of interest of this thesis, kSZ velocity reconstruction, is another example of an approach that exploits the statistics between the kSZ signal and a galaxy survey [18, 19, 20, 21]. As we will discuss later in the text, the kSZ effect leads to a statistical anisotropy in the galaxy-temperature cross-correlation $\langle gT \rangle$, and this anisotropy encodes information about the large scale velocity field inside the volume of the survey. Using the structure of the statistical anisotropy, a quadratic estimator $\hat{v}$ for the large scale velocity can be constructed using products of the galaxy and temperature field. kSZ velocity reconstruction is also discussed in [20], as the cross-correlation between the reconstructed velocity (gT) and the galaxy survey used for its reconstruction (g) is also the optimally estimated bispectrum $\langle ggT \rangle$. We leave a more detailed discussion of kSZ velocity reconstruction for Sec.1.4

1.3.5 Polarized Sunyaev-Zel'dovich effect

We conclude our summary of the secondary CMB by considering the polarization anisotropies that arise due to scattering with free electrons on the past lightcone. In the presence of a local CMB quadrupole, the scattered photons are endowed with a polarization [23, 35]. The polarized component of the CMB arising after reionization, primarily from collapsed structures, is known as the polarized Sunyaev-Zel'dovich (pSZ) effect (as opposed to the component sourced near decoupling and at reionization, which is simply CMB polarization). The pSZ effect can be expressed in terms of the contributions to the Stokes parameters $Q \pm iU$ from scattering events along the line of sight:

$$(Q \pm iU)^{\rm pSZ}(\hat{\mathbf{n}}) = -\frac{\sqrt{6}}{10} \int d\chi \sum_{m=-2}^{2} \Theta_2^m(\hat{\mathbf{n}}, \chi) {}_{\pm 2}Y_{2m}(\hat{\mathbf{n}}) \times \dot{\tau}(\hat{\mathbf{n}}, \chi) \tag{1.14}$$

where $\Theta_2^m(\hat{\mathbf{n}}, \chi)$ are the remote $\ell = 2$ CMB multipoles observed by electrons at spacetime position $(\eta(\chi), \mathbf{x} = \chi\hat{\mathbf{n}})$ on our past lightcone:

$$\Theta_2^m(\hat{\mathbf{n}}, \chi) = \int d^2\hat{n}\, \Theta(\mathbf{x}, \hat{\mathbf{n}}) Y^*_{2m}(\hat{\mathbf{n}}) \tag{1.15}$$

and ${}_{\pm 2}Y_{2m}(\hat{\mathbf{n}})$ are spin-2 weighted spherical harmonics. The pSZ signal has been proposed as a gateway to measure the remote CMB quadrupole [61], which probes different sections

of the last-scattering surface than our locally observed quadrupole and thus can offer additional information on large-scale inhomogeneities within our lightcone. For moving clusters, electrons observe additional contributions to their observed quadrupole coming from the Doppler leakage of other remote CMB multipoles, and the frequency dependence of these contributions could be used to separate and measure them [62]. These corrections are not likely to be measured in the near future, given that even the leading order contribution will be challenging to measure with CMB Stage 4 experiments [63], having a forecasted signal to noise of 0.5% for individual clusters and a 3σ detection for cumulative signal, although 9σ detection are possible for observations in the 150 GHz band (corresponding to a 9-meter mirror). It was shown in [19] that these levels of detection still allow for a reconstruction of the remote quadrupole in a similar fashion to kSZ velocity reconstruction, the details of which we will review in the following section. The reconstruction of the remote quadrupole is reviewed separately in Appendix A.1

1.4 kSZ velocity reconstruction in the Lightcone Picture

In this section, we review the details of kSZ velocity reconstruction. The basic principle behind the method was briefly mentioned in Sec.1.3.4: the kSZ effect induces a statistical anisotropy in the cross-correlation between the CMB temperature and tracers of the electron density on our past lightcone:

$$\left\langle \Theta^{kSZ}_{\ell m}\ \delta_{\ell' m'} \right\rangle \neq 0 \quad \text{for} \quad \ell \neq \ell', m \neq m'. \tag{1.16}$$

where $\Theta^{kSZ}_{\ell m}$ are the multipole moments of the kSZ signal on the whole sky and $\delta_{\ell' m'}$ are the multipole moments of the electron density tracer on the full sky. On small angular scales ($\ell, \ell' \gg 1$), this statistical anisotropy is mostly sourced by the large-scale modes of the remote dipole field defined in Eq.(1.13). kSZ velocity reconstruction utilizes the structure of the statistical anisotropy to estimate the remote dipole field.

Several approaches have emerged to describe the reconstruction procedure. Two descriptions of kSZ velocity reconstruction that have received attention are the so called "Box Picture" [20], and "Lightcone Picture" [18, 19] . In the Box Picture, the kSZ effect is estimated from the momentum field along one direction in a 3D box at the median redshift of a galaxy survey. The box formulation is convenient since it sidesteps spherical projection effects and allows one to work in the familiar Fourier domain. This approach is a good approximation for reconstruction on relatively small sky areas and over limited

ranges in redshift, for which the underlying geometry of the lightcone can be ignored. For this same reason, it is cumbersome to accurately incorporate redshift evolution, relativistic contributions to the kSZ effect, and large sky area in the Box Picture.

In the Lightcone Picture , kSZ velocity reconstruction is formulated in terms of observables on our past light cone, which facilitates the incorporation of redshift evolution and relativistic contributions to the radial velocity field (promoting it to the remote dipole field). The Lightcone Picture has its own drawbacks when compared to the Box Picture: computing observables is usually more expensive due to the cumbersome projection integrals on the lightcone, and the incorporation of photometric redshift errors and other systematics is less straightforward. However, in an application to future datasets, the Light Cone Picture has a number of advantages compared to the Box Picture. First, it is formulated in terms of direct observables (e.g. fields on the sphere), making contact between theory and observation precise. Second, the signal to noise of the reconstruction is estimated to be largest on the largest scales, where the redshift evolution and projection effects captured by the Lightcone Picture are most important.

The research presented in this thesis was done in the Lightcone Picture of kSZ velocity reconstruction. In order to provide context to the reader and properly identify the new contributions to the formalism that are presented in this work, we dedicate the rest of this section to review the foundations of the Lightcone Picture presented in Refs. [18, 19]. The exposition pays special attention to the assumptions and approximations made, with the ultimate purpose of highlighting avenues for the improvement of the technique.

1.4.1 Observables

Let us assume that we have full sky measurements of the CMB where the only contributions are the primary CMB anisotropies (lensed) and kSZ effect, and measurements are limited only by instrumental noise. The spherical harmonic moments for the temperature anisotropies are given by:

$$\Theta_{\ell m} = \Theta^{p,lens}_{\ell m} + \Theta^{kSZ}_{\ell m} + n_{\ell m} \tag{1.17}$$

where the multipoles $\Theta^{kSZ}_{\ell m}$ are given by the spherical harmonic transform of Eq.(1.12):

$$
\begin{aligned}
\Theta^{kSZ}_{\ell m} &= \int d^2\hat{n}\, \Theta^{kSZ}(\hat{\mathbf{n}}) Y^*_{\ell m}(\hat{\mathbf{n}}) \\
&= -\int d^2\hat{n} \int d\chi\; v_{\rm eff}(\hat{\mathbf{n}},\chi)\dot{\tau}(\hat{\mathbf{n}},\chi) Y^*_{\ell m}(\hat{\mathbf{n}}) \\
&= -\int d\chi\; v_{\ell_1 m_1}(\chi)\dot{\tau}_{\ell_2 m_2}(\chi) \int d^2\hat{n} Y_{\ell_1 m_1}(\hat{\mathbf{n}}) Y_{\ell_2 m_2}(\hat{\mathbf{n}}) Y^*_{\ell m}(\hat{\mathbf{n}}) \\
&= -\sum_{\ell_1 m_1}\sum_{\ell_2 m_2} (-1)^m W^{\ell_1,\ell_2,\ell}_{m_1,m_2,-m} \int d\chi\; v_{\ell_1 m_1}(\chi)\dot{\tau}_{\ell_2 m_2}(\chi)
\end{aligned}
\tag{1.18}
$$

where

$$
W^{\ell_a,\ell_b,\ell_c}_{m_a,m_b,m_c} = \sqrt{\frac{(2\ell_a+1)\,(2\ell_b+1)\,(2\ell_c+1)}{4\pi}} \begin{pmatrix} \ell_a & \ell_b & \ell_c \\ 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} \ell_a & \ell_b & \ell_c \\ m_a & m_b & m_c \end{pmatrix} \tag{1.19}
$$

is the multipole coupling constructed with 3-j Wigner symbols resulting from the angular integral, $v_{\ell_1 m_1}(\chi)$ are the spherical harmonic moments of the remote dipole $v_{\rm eff}(\hat{\mathbf{n}},\chi)$ along the lightcone (subscript eff is suppressed for ease of notation), and $\dot{\tau}_{\ell_2 m_2}(\chi)$ are the spherical harmonic moments of the differential optical depth $\dot{\tau}(\hat{\mathbf{n}},\chi)$.

We assume we have full-sky measurements of a tracer of the electron distribution on the past lightcone, such as a redshift galaxy survey $g(\hat{\mathbf{n}},\chi(z))$, and we also assume we can subdivide the survey by taking averages inside N top-hat bins of equal size in comoving space:

$$
g^\alpha(\hat{\mathbf{n}}) = \int d\chi\; \Pi^\alpha(\chi) g(\hat{\mathbf{n}},\chi) \tag{1.20}
$$

where

$$
\Pi^\alpha(\chi) = \begin{cases} \frac{1}{\Delta\chi}, & \chi_{\rm min} + \alpha\Delta\chi \le \chi < \chi_{\rm min} + (\alpha+1)\Delta\chi, \\ 0, & \text{otherwise} \end{cases} \tag{1.21}
$$

with α in $(0,\ldots,N-1)$. For this idealized scenario we assume perfect redshift measurements, ignore for the moment the effects of redshift spatial distortions so that we can have a one to one correspondence between the comoving distance χ and the cosmological redshift z of the galaxies, and assume the only source of noise is shot noise.

1.4.2 CMB-LSS statistical anisotropy

The cross-correlation between the multipole moments of the CMB temperature and the binned galaxy survey shows a statistical anisotropy sourced by the kSZ signal:

$$\left\langle \Theta_{\ell m}\, g^{\alpha}_{\ell' m'} \right\rangle = -\sum_{\ell_1 m_1}\sum_{\ell_2 m_2} (-1)^m W^{\ell_1,\ell_2,\ell}_{m_1,m_2,-m} \int d\chi \left\langle v_{\ell_1 m_1}(\chi)\dot{\tau}_{\ell_2 m_2}(\chi) g^{\alpha}_{\ell' m'} \right\rangle. \tag{1.22}$$

The three-point function in Eq.(1.22) can be approximated in the limit that two of the (ℓ_1, ℓ_2, ℓ') multipoles are far greater than the remaining multipole (squeezed limit). Below, we assume that the dominant contribution to the three-point function comes from the squeezed limit where $\ell_1 \ll \ell, \ell_2$. Furthermore, we will consider a fixed realization for the large scale field $v_{\ell_1 m_1}(\chi)$ and take the the statistical average of the three-point function only over small angular scales. With these assumptions we can approximate:

$$\left\langle v_{\ell_1 m_1}(\chi)\, \dot{\tau}_{\ell_2 m_2}(\chi)\, g^{\alpha}_{\ell' m'} \right\rangle \approx\ v_{\ell_1 m_1}(\chi) \left\langle \dot{\tau}_{\ell_2 m_2}(\chi)\, g^{\alpha}_{\ell' m'} \right\rangle. \tag{1.23}$$

Following [19], we approximate the continuous fields $v_{\ell_1 m_1}(\chi)$ and $\dot{\tau}_{\ell_2 m_2}(\chi)$ by their bin averages:

$$v_{\ell_1 m_1}(\chi) \approx \Delta\chi \sum_{\alpha'=0}^{N-1} v^{\alpha'}_{\ell_1 m_1} \Pi^{\alpha'}(\chi) \quad \text{with} \quad v^{\alpha'}_{\ell_1 m_1} = \frac{1}{\Delta\chi}\int_{\chi^{\alpha'}_{min}}^{\chi^{\alpha'}_{max}} d\chi\, v_{\ell_1 m_1}(\chi), \tag{1.24}$$

$$\dot{\tau}_{\ell_2 m_2}(\chi) \approx \Delta\chi \sum_{\alpha'=0}^{N-1} \dot{\tau}^{\alpha'}_{\ell_2 m_2} \Pi^{\alpha'}(\chi) \quad \text{with} \quad \dot{\tau}^{\alpha'}_{\ell_2 m_2} = \frac{1}{\Delta\chi}\int_{\chi^{\alpha'}_{min}}^{\chi^{\alpha'}_{max}} d\chi\, \dot{\tau}_{\ell_2 m_2}(\chi). \tag{1.25}$$

Using $\left\langle \dot{\tau}^{\alpha'}_{\ell_2 m_2}\, g^{\alpha}_{\ell' m'} \right\rangle = C_{\ell'}^{\dot{\tau}^{\alpha} g^{\alpha}}\, \delta_{\alpha\alpha'}\, \delta_{\ell_2 \ell'}\, \delta_{m_2 m'}$, we arrive to the final expression for the CMB-LSS statistical anisotropy :

$$\left\langle \Theta_{\ell m}\, g^{\alpha}_{\ell' m'} \right\rangle = -\sum_{\ell_1 m_1} (-1)^m W^{\ell_1,\ell',\ell}_{m_1,m',-m}\, v^{\alpha}_{\ell_1 m_1}\, C_{\ell'}^{\dot{\tau}^{\alpha} g^{\alpha}}\, \Delta\chi. \tag{1.26}$$

The calculation above shows that the statistical anisotropy on small angular scales ($\ell, \ell' \gg 1$) between the CMB temperature and a galaxy survey averaged inside a comoving bin α is sourced by the low-ℓ_1 moments of the remote dipole field, averaged inside the same comoving bin.

1.4.3 Quadratic estimator for the remote dipole

The structure of Eq.(1.26) can be used to construct a quadratic estimator for the large angle multipoles of the remote dipole inside each comoving bin. Reference [19] considers

the quadratic sum:

$$\hat{v}^{\alpha}_{LM} = A^{\alpha}_{L} \sum_{\ell m;\ell' m'} (-1)^{M} \begin{pmatrix} \ell & \ell' & L \\ m & m' & -M \end{pmatrix} G^{\alpha}_{\ell\ell' L}\, \Theta_{\ell m} g^{\alpha}_{\ell' m'} \tag{1.27}$$

and chooses the weights A^{α}_{L} and $G^{\alpha}_{\ell\ell' L}$ such that the estimator is unbiased $\left\langle \hat{v}^{\alpha}_{LM} \right\rangle = v^{\alpha}_{LM}$ and has minimum variance. The first requirement leads to the following constraint:

$$A^{\alpha}_{L} = (2L+1) \left(\sum_{\ell;\ell'} G^{\alpha}_{\ell\ell' L} f^{\alpha}_{\ell L\ell'} \right)^{-1}. \tag{1.28}$$

where $f^{\alpha}_{\ell L\ell'}$ is given by:

$$f^{\alpha}_{\ell\ell_1\ell'} \equiv \sqrt{\frac{(2\ell+1)(2\ell_1+1)(2\ell'+1)}{4\pi}} \begin{pmatrix} \ell & \ell' & \ell_1 \\ 0 & 0 & 0 \end{pmatrix} C^{\dot{\tau} g}_{\alpha,\ell'}\, \Delta\chi. \tag{1.29}$$

As part of the calculation to find the minimum variance estimator, we need to compute a four-point function of the form $\langle \Theta g \Theta g \rangle$. Reference [19] approximates this 4-point function as $\langle \Theta\Theta \rangle \langle gg \rangle$ and finds the weights that lead to a minimum variance estimator subject to Eq.(1.28):

$$G^{\alpha}_{\ell\ell' L} = \frac{f^{\alpha}_{\ell L\ell'}}{C^{\Theta\Theta}_{\ell} C^{g^{\alpha} g^{\alpha}}_{\ell'}}. \tag{1.30}$$

where $C^{\Theta\Theta}_{\ell}$ is the full CMB power spectrum including instrumental noise, and $C^{g^{\alpha} g^{\alpha}}_{\ell'}$ is the galaxy power spectrum in bin α including shot noise. With this choice of weights, the variance of the estimator gives:

$$\left\langle \hat{v}^{\alpha}_{L'M'} \hat{v}^{\beta *}_{LM} \right\rangle = C^{v^{\alpha} v^{\beta}}_{L} \delta_{LL'} \delta_{MM'} + N^{\alpha\beta}_{L} \delta_{LL'} \delta_{MM'} \tag{1.31}$$

where $C^{v^{\alpha} v^{\beta}}_{L}$ is the bin-bin covariance matrix of the remote dipole multipoles $\hat{v}^{\alpha}_{LM}$ and $N^{\alpha\beta}_{L}$ is the reconstruction noise matrix given by:

$$N^{\alpha\beta}_{L} = A^{\alpha}_{L}\, \delta_{\alpha\beta}. \tag{1.32}$$

For CMB noise levels in the ballpark of next-generation CMB Stage 4 experiments and galaxy surveys with specifications similar to Rubin, the analysis in [19] forecasted the possibility of high signal to noise reconstruction for the multipoles $\hat{v}^{\alpha}_{LM}$ (several thousands for the lowest multipoles) for 6 and 12 bins of equal comoving size in the redshift range (0,6).

1.4.4 Discussion of assumptions

The formalism and results from [19] served as a baseline to investigate the potential of kSZ velocity reconstruction on the lightcone. The promising forecasts that were obtained motivated a more detailed study of the assumptions made and led to the research presented in this thesis. Here are listed the aspects of the Lightcone Picture that required revisions and extensions:

1. **Assumptions about the CMB temperature:** Eq.(1.17) assumes full-sky measurements of the blackbody CMB spectrum limited only by instrumental noise. Incomplete sky coverage is a feature that realistic experiments often have to operate with, and multiple frequency dependent foregrounds are present in measurement of the CMB temperature anisotropies, like the tSZ effect, the Cosmic Infrared Background (CIB), and radio point sources. Instrumental noise also depends on the frequency of observation. A realistic application of kSZ velocity reconstruction will certainly involve some multi-frequency foreground removal procedure which could deteriorate the signal to noise ratio. Eq.(1.17) also ignores additional contributions to the blackbody CMB temperature like the linear and non-linear ISW effects.

2. **Assumptions about the galaxy survey:** Realistic galaxy surveys are subject to redshift errors, redshift spatial distortions and limitations on how many custom redshift bins the survey can be divided into. To see how this effects can impact the reconstruction, consider the cross-correlation between galaxies in two redshift bins:

$$\left\langle g^{\alpha}_{\ell m}\, g^{*\beta}_{\ell m} \right\rangle = \int d\chi_1 \int d\chi_2\, \Pi^{\alpha}(\chi_1)\Pi^{\beta}(\chi_2) C^{gg}_{\ell}(\chi_1,\chi_2), \tag{1.33}$$

with

$$C^{gg}_{\ell}(\chi_1,\chi_2) = \frac{2}{\pi}\int_0^{\infty} \frac{dk}{k} k^3 P_{gg}\left(k,\chi_1,\chi_2\right) j_{\ell}\left(k\chi_1\right) j_{\ell}\left(k\chi_2\right) \tag{1.34}$$

where P_{gg} is the underlying galaxy-galaxy power spectrum in Fourier space and j_ℓ are spherical Bessel functions. Since kSZ reconstruction is concerned with small angular scale galaxy perturbations ($\ell \gg 1$) we can study the equations above using the Limber approximation [64, 65], which picks out the $\chi_1 = \chi_2 = \frac{\ell+\frac{1}{2}}{k}$ part of the radial integrals. In the absence of redshift errors, the window functions $\Pi^{\alpha}(\chi_1)$ and $\Pi^{\beta}(\chi_2)$ don't overlap for $\alpha \neq \beta$ and the correlation vanishes. Redshift errors, like the ones expected in photometric redshift surveys like Rubin, lead to a smearing of the window functions, allowing for non vanishing overlaps for $\alpha \neq \beta$ and thus galaxy correlation between bins. The reconstruction noise for the remote dipole would not be diagonal,

reflecting the mixing of radial information due to redshift errors. Imperfect redshift measurements also lead to biases in the reconstructed remote dipole. The smearing of the window functions introduces $\alpha \neq \alpha'$ correlation in $\left\langle \dot{\tau}^{\alpha'}_{\ell_2 m_2}\, g^{\alpha}_{\ell' m'} \right\rangle$, which changes the final statistical anisotropy Eq.(1.26) to:

$$\left\langle \Theta_{\ell m}\, g^{\alpha}_{\ell' m'} \right\rangle = -\sum_{\alpha'=0}^{N-1} \sum_{\ell_1 m_1} (-1)^m W^{\ell_1,\ell',\ell}_{m_1,m',-m}\, v^{\alpha'}_{\ell_1 m_1}\, C_{\ell'}^{\dot{\tau}^{\alpha'} g^{\alpha}}\, \Delta\chi. \tag{1.35}$$

Thus, we see that redshift errors, by mixing radial information, bias the remote dipole reconstruction by mixing different bins.

Redshift spatial distortions are expected to be less of a problem, due to the fact that their effect modifies the galaxy-galaxy angular power spectrum at multipoles that are far smaller than the ones needed for velocity reconstruction.

3. **Assumptions about the CMB-LSS statistical anisotropy:** Eq.(1.22) assumes that the only source of statistical anisotropy in the temperature-galaxy cross-correlation is sourced by the kSZ signal. In reality, other statistically anisotropic terms will be present, sourced by CMB lensing and the non-linear ISW effect. Both of these terms lead to a three-point function when correlated with the galaxy moments which can be approximated in a squeezed limit similarly to the kSZ-g term. For example, the lensing term leads to the following statistical anisotropy:

$$\left\langle \Theta^{p,lens}_{\ell m}\, g^{\alpha}_{\ell' m'} \right\rangle = \sum_{\ell_1 m_1} (-1)^{m_1} \begin{pmatrix} \ell & \ell' & \ell_1 \\ m & m' & -m_1 \end{pmatrix} f^{\alpha}_{\ell \ell_1 \ell'}\, \Theta^{p}_{\ell_1 m_1} \tag{1.36}$$

where $\Theta^{p}_{\ell_1 m_1}$ are the unlensed primary CMB moments and the couplings $f^{\alpha}_{\ell\ell_1\ell'}$ are given by:

$$f^{\alpha}_{\ell\ell_1\ell'} \equiv [\ell_1(\ell_1+1) + \ell'(\ell'+1) - \ell(\ell+1)] \tag{1.37}$$

$$\times \sqrt{\frac{(2\ell+1)(2\ell_1+1)(2\ell'+1)}{16\pi}} \begin{pmatrix} \ell & \ell' & \ell_1 \\ 0 & 0 & 0 \end{pmatrix} C_{\ell'}^{\phi g^{\alpha}} \tag{1.38}$$

where $C_{\ell'}^{\phi g^{\alpha}}$ is the angular cross-spectrum between the lensing potential Eq.(1.5) and the binned galaxy moments. Systematic errors in the galaxy survey can also lead to a statistical anisotropy in the CMB-LSS cross-correlation, a possibility we will explore in the Chapter 2.

4. **Assumptions about the reconstruction noise:** The variance of the estimator involves a four-point function of the form $\langle \Theta g \Theta g \rangle$. The contribution from the kSZ effect to this four-point function can be expressed as a six-point function of the form $\langle v\dot{\tau} g v\dot{\tau} g \rangle$. Assuming Gaussian fields, we see that the approximation $\langle \Theta g \Theta g \rangle \approx \langle \Theta\Theta \rangle \langle gg \rangle$ used in [19] is considering only a subset of the 15 possible contractions of the 6-point function. These ignored terms could lead to underestimations of the reconstruction noise. For non-Gaussian fields, one must additionally compute the connected part of the six-point function.

5. **Assumptions about the coarse graining procedure:** Eqs.(1.24) and (1.25) expand the radial dependence of the remote dipole field and the differential optical depth along the interval (χ_{min}, χ_{max}) of the line of sight in terms of a finite set of top hat functions $\Pi^{\alpha}(\chi)$. This set of functions does not form a complete set, so the expansion used in [19] does not account for contributions to the kSZ line of sight integral that do not come from the bin-averages of the remote dipole. In the limit of $N \to \infty$ such contributions should be negligible, but for realistic numbers of galaxy redshift bins the neglected terms could introduce a significant bias in the reconstructed average dipoles.

6. **Assumptions about the galaxy-electron cross-correlation**: In order to construct the quadratic estimator, the couplings $f^{\alpha}_{\ell\ell_1\ell'}$ defined in Eq.(1.29]) have to be modelled from theory. The underlying quantity that needs to be modelled is the Fourier-space power spectrum P_{eg} between electrons and galaxies. Reference [19] use the model $P_{eg} = bP_{mm}$ where P_{mm} is the nonlinear matter power spectrum and b is the scale independent the galaxy bias on large scales. As the authors rightly pointed out, this approximation ignores baryonic feedback effects that make the electron distribution differ from the dark matter one on small scales, and also neglects the scale dependence of the galaxy bias. The uncertainty in the P_{eg} modelling leads to what is referred in the kSZ velocity reconstruction literature as the optical depth bias [20, 21, 66].

The list above points towards many aspects of the Lightcone Picture of kSZ velocity reconstruction that require revisions. Chapter 2 addresses these issues and develops an extended formalism to account for the additional sources of noise, sources of bias, and modelling of more realistic experimental conditions. An issue that is not treated in this thesis is that of the non-Gaussian contributions to the estimator variance. The extra contributions to the variance beyond the $\langle \Theta\Theta \rangle \langle gg \rangle$ term were explored in the context of the Box Picture formalism in Ref.[21] using a suite of N-body simulations. The terms

coming from the disconnected 6-point function were referred to as the $N^{(1)}$ bias and the terms coming from the connected part were referred as the $N^{(\frac{3}{2})}$ bias. It was shown that the $N^{(\frac{3}{2})}$ is far larger than the $N^{(1)}$, and can even become comparable to the $\langle\Theta\Theta\rangle\langle gg\rangle$ term is signal to noise ratio is high enough. As we will discuss in Chapter 2, we expect the Lightcone Picture equivalent of $N^{(\frac{3}{2})}$ to be small for the fiducial experimental conditions we explored. However, as the authors of [21] point out, futuristic experiments might be sensitive enough to make the $N^{(\frac{3}{2})}$ term important, and its calculation in the Lightcone Picture will be necessary.

1.5 The remote dipole as a probe of the large-scale Universe

In the sections above, we have discussed how kSZ velocity reconstruction allows us to estimate the large angular features of the remote CMB dipole $v_{\rm eff}(\hat{\mathbf{n}}, \chi)$ defined in Eq(1.13). We conclude this introductory chapter by discussing the different contributions to the remote dipole and how they inform us about cosmological perturbations on the largest accessible scales. We follow the description of the remote dipole presented in [18]. Freely falling electrons located along our past lightcone observe a CMB dipole that is sourced by the Sachs-Wolfe effect, the integrated Sachs-Wolfe effect and the Doppler effect. With these contributions, the line of sight projected remote dipole is given by

$$v_{\rm eff}(\hat{\mathbf{n}}, \chi) = \frac{3}{4\pi}\int d^2\hat{n}_e \ \left[\ \Theta^{SW}(\mathbf{x}, \hat{\mathbf{n_e}}) + \Theta^{ISW}(\mathbf{x}, \hat{\mathbf{n_e}}) + \Theta^{\mathcal{D}}(\mathbf{x}, \hat{\mathbf{n_e}})\ \right]\ (\hat{\mathbf{n}}\cdot\hat{\mathbf{n_e}}) \tag{1.39}$$

We can express each one of these contributions in terms of the Newtonian gauge gravitational potential Ψ (neglecting anisotropic stress is a good approximation for the large scale modes of interest [67]). The Fourier modes of the gravitational potential can be expressed in linear theory in terms of their primordial values at $a \to 0$ using appropriate transfer and growth functions. For modes that enter the horizon well after matter radiation equality, the gravitational potential $\Psi(a, \mathbf{k})$ can be related to its primordial value $\Psi_p(\mathbf{k})$ by:

$$\frac{\Psi(a,\mathbf{k})}{\Psi_p(\mathbf{k})} = D_{\bar{\Psi}}(a) = \frac{16\sqrt{1+y} + 9y^3 + 2y^2 - 8y - 16}{10y^3} \times \left[\frac{5}{2}\Omega_m \frac{E(a)}{a}\int_0^a \frac{da}{E^3(a)a^3}\right] \tag{1.40}$$

with $y = a(1+z_{eq})$ and $E(a) = \sqrt{\Omega_m a^{-3} + \Omega_\Lambda}$ is the normalized Hubble parameter. The first factor is takes care of the matter-radiation equality crossing and the second factor

takes care of late time evolution. In order to accommodate for evolution on sub-horizon scales, we can introduce a transfer function:

$$\Psi(a, \mathbf{k}) = \Psi_{\rm p}(\mathbf{k}) T(k) D_\Psi(a) \tag{1.41}$$

where

$$T(k) = \frac{\ln[1 + 0.171x]}{0.171x} \left[1 + 0.284x + (1.18x)^2 + (0.399x)^3 + (0.49x)^4\right]^{-0.25} \tag{1.42}$$

with $x = k/k_{eq}$ and $k_{\rm eq} = a_{\rm eq} H\left(a_{\rm eq}\right) \simeq 82.5 H_0$. In terms of this quantities, the remote dipole can be written as an integral on Fourier space:

$$v_{\rm eff}(\hat{\mathbf{n}}, \chi) = i \int \frac{d^3k}{(2\pi)^3} T(k) \Psi_p(\mathbf{k}) \mathcal{K}^v\left(k, \chi\right) \mathcal{P}_1\left(\hat{\mathbf{k}} \cdot \hat{\mathbf{n}}\right) e^{i\chi \mathbf{k}\cdot\hat{\mathbf{n}}} \tag{1.43}$$

where $\mathcal{P}_1$ is the $\ell = 1$ Legendre polynomial and $\mathcal{K}^v$ is an integral kernel with contributions from the SW, ISW, and Doppler effects:

$$\mathcal{K}^v\left(k, \chi\right) = \mathcal{K}^{SW}\left(k, \chi\right) + \mathcal{K}^{ISW}\left(k, \chi\right) + \mathcal{K}^{\mathcal{D}}\left(k, \chi\right) \tag{1.44}$$

given by

$$\begin{aligned}
\mathcal{K}^{SW}\left(k, \chi\right) &\equiv 3\left(2 D_\Psi\left(a_{\rm ls}\right) - \frac{3}{2}\right) j_1\left(k(\chi_{\rm ls} - \chi)\right) \\
\mathcal{K}^{ISW}\left(k, \chi\right) &\equiv 6 \int_{a_{\rm ls}}^{a(\chi)} da' \frac{dD_\Psi}{da'} j_1(k(\chi(a') - \chi)) \\
\mathcal{K}^{\mathcal{D}}\left(k, \chi\right) &\equiv k D_v\left(a_{\rm ls}\right) j_0\left(k(\chi_{\rm ls} - \chi)\right) - 2k D_v\left(a_{\rm ls}\right) j_2\left(k(\chi_{\rm ls} - \chi)\right) - k D_v\left(a(\chi)\right)
\end{aligned} \tag{1.45}$$

where $D_v(a)$ is a velocity growth function defined as:

$$D_v(a) \equiv \frac{2a^2 H(a)}{H_0^2 \Omega_m} \frac{y}{4 + 3y} \left[D_\Psi + \frac{dD_\Psi}{d \ln a}\right]. \tag{1.46}$$

The last term of the Doppler integral kernel, $-k D_v\left(a(\chi)\right)$, is the only term that is typically considered when modelling the kSZ signal. We refer to this term as the "local Doppler" term and denote it by $\mathcal{K}^{\mathcal{LD}}$. In what follows, we discuss how the other contributions to the remote dipole become important when a kSZ based observable probes large enough volumes of the Universe. We refer to these additional contributions, which source

the entire CMB dipole for observers with no peculiar velocities, as the "stationary dipole" and denote it by $\mathcal{K}^{\mathcal{S}t}$. The left panel of Fig.(1.2) depicts the behaviour of the full integral kernel as a function of the scale and for several redshift values; the right panel shows the relative size of the local Doppler term to stationary dipole term. We can see that for scales $k \gtrsim 1[\mathrm{Gpc}^{-1}]$ the linear behaviour of the local Doppler term is dominant and the stationary dipole contributions to the kernel manifest as oscillations that become negligible as we go to smaller scales. For scales $k \lesssim 1[\mathrm{Gpc}^{-1}]$ the situation is more interesting: the contributions beyond the local Doppler term become more relevant and, as shown in [18], precisely cancel the local Doppler term in the limit $k \to 0$ and leave $\mathcal{O}(k^3)$ residuals[1]. It is clear that when large enough volumes of the Universe are studied, the local Doppler term becomes insufficient to describe the remote CMB dipole.

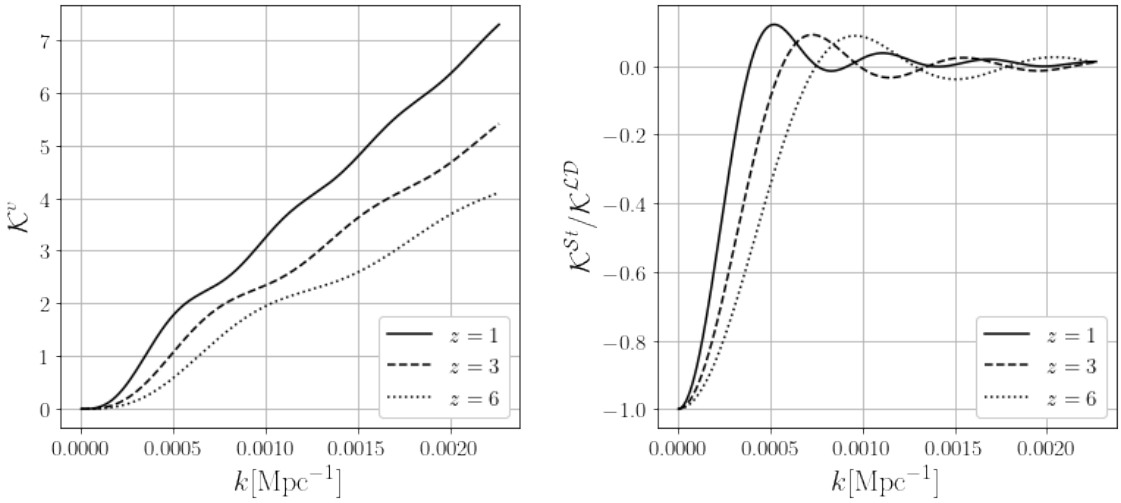

Figure 1.2: **Left panel:** Remote dipole integral kernel as a function of scale, for redshifts $z = 1, 3, 6$. As k increases, the linear behaviour of the local term Doppler dominates. For $k \to 0$, the integral kernel goes as $\mathcal{O}(k^3)$ due to the precise cancellation of the local Doppler by the $\mathcal{O}(k)$ pieces of other contributions to the kernel. **Right panel:** Relative size of local Doppler term respect to the other contributions to $\mathcal{K}^v$. The dominance of the local Doppler terms falls as we go to larger scales.

For most of the kSZ statistics discussed in Sec.(1.3.4), the scales that are probed are largely dominated by the local Doppler contribution to the remote dipole. The Lightcone

[1]This cancellation occurs because a pure gradient in the gravitational potential, which can be removed by a gauge transformation, should not contribute to observable quantities.

Picture of kSZ velocity reconstruction is the only approach with an adequate infrastructure to describe the signal sourced by the stationary dipole, which depends strongly on the lightcone geometry. Moreover, the formulation in the Ligthcone Picture allows for estimations of the remote dipole on larger portions of the sky compared to other methods, allowing us to access the scales at which the corrections coming from the stationary dipole terms become more relevant. We can inspect the typical size of these corrections by looking at the harmonic moments of the remote dipole:

$$v_{\rm eff;LM}(\chi) = i\int d^2\hat{n}\int \frac{d^3k}{(2\pi)^3}T(k)\Psi_p(\mathbf{k})\mathcal{K}^v\left(k,\chi\right)\mathcal{P}_1\left(\hat{\mathbf{k}}\cdot\hat{\mathbf{n}}\right)e^{i\chi\mathbf{k}\cdot\hat{\mathbf{n}}}Y^*_{LM}\left(\hat{\mathbf{n}}\right) \quad (1.47)$$

$$= \int \frac{d^3k}{(2\pi)^3}\Delta^v_L\left(k,\chi\right)\Psi_p(\mathbf{k})Y^*_{LM}(\hat{\mathbf{k}}) \quad (1.48)$$

where we have defined the remote dipole transfer function:

$$\Delta^v_L\left(k,\chi\right) \equiv \frac{4\pi i^L}{2L+1}\mathcal{K}^v\left(k,\chi\right)\left[Lj_{L-1}\left(k\chi\right)-(L+1)j_{L+1}\left(k\chi\right)\right]T(k) \quad (1.49)$$

In Fig.(1.3) we compare two versions of the transfer function: in solid lines we plot the transfer function when only the local Doppler term sources the remote dipole and in dashed lines we plot the full transfer function which contains local Doppler and stationary contributions to the remote dipole. We do the comparison at redshifts $z = 1, 3, 6$ and for multipoles $L = 1, 3, 5$. As expected, the large angular scales of the remote dipole are the ones more sensitive to the stationary dipole, and this sensitivity increases with redshift as larger volumes of the Universe are being probed. Fig.(1.3) tells us two important things about kSZ velocity reconstruction in the Lightcone Picture: first, the lowest reconstructed multipoles (which are forecasted to have the highest signal to noise [19]) can inform us about cosmological perturbations on scales of several Gpc; second, proper modelling of the stationary dipole contribution plays an important role in interpreting measurements on these scales. If the signal is modelled only with the local Doppler term, one would infer incorrectly the features of the underlying primordial perturbations. Due to its sensitivity to Gpc-scale physics, the remote dipole could provide additional statistical power to test the nature of the CMB anomalies and determine if the anomalies are the result of deviations from the ΛCDM primordial spectrum of perturbations on ultra-large scales. We discuss more about this potential use of the remote CMB dipole in Chapter 4.

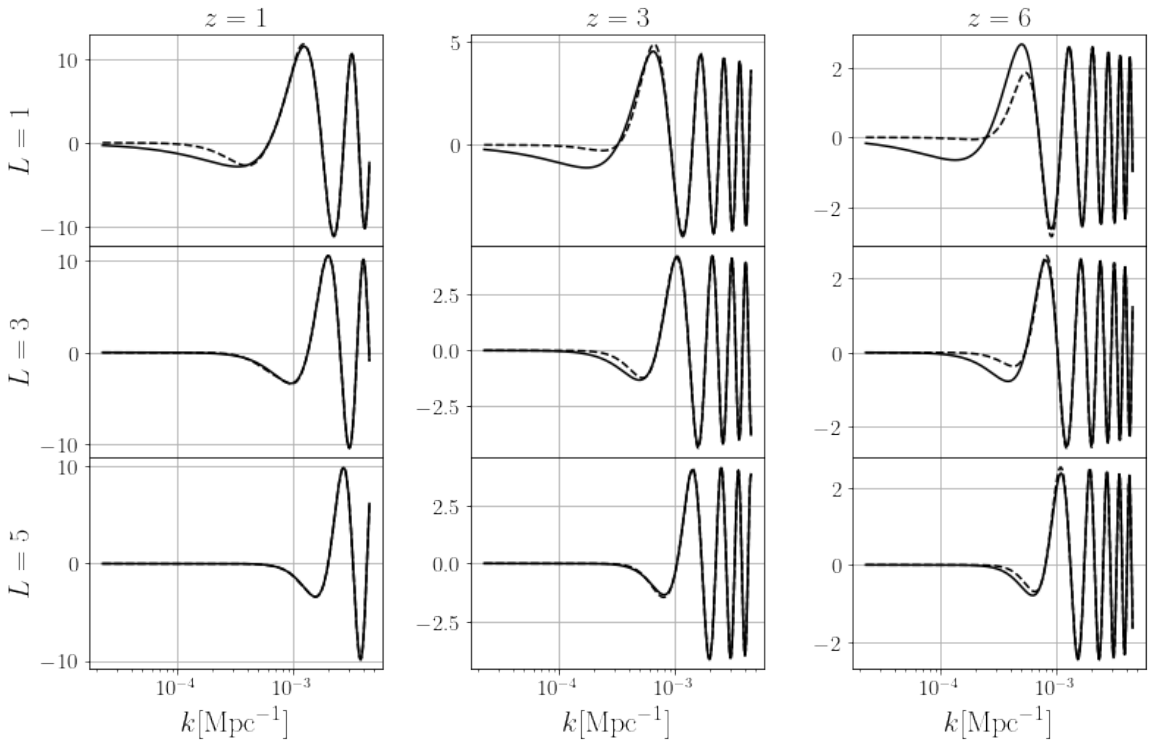

Figure 1.3: Transfer function $\Delta_L^v(k, \chi)$ for the remote dipole spherical harmonic moments as a function of the scale k, for redshifts $z = 1, 3, 6$ and for multipoles $L = 1, 3, 5$. Solid lines are obtained when only the local Doppler term $\mathcal{K}^{\mathcal{LD}}$ is considered as a source for the remote dipole. Dashed lines corresponds to the full transfer function, including local Doppler and stationary sources to the remote dipole.

Chapter 2

The extended Lightcone Picture formalism

2.1 Chapter summary

The kSZ effect carries significant cosmological information through its dependence on the large-scale peculiar velocity field. Previous work identified promising means of extracting this cosmological information using a set of quadratic estimators for the radial component of the velocity field [18, 19, 20]. These estimators are based on the statistically anisotropic components of the cross-correlation between the CMB and a tracer of large-scale structure, such as a galaxy redshift survey. In this chapter, we assess the challenges posed by various foregrounds and systematics in the CMB and galaxy surveys, as well as biases in the estimator. We work in the context of the Lightcone Picture approach for kSZ velocity reconstruction described in Sec.(1.4) of Chapter 1.

A primary goal of this chapter is to further develop the formalism for the Lightcone Picture. We begin by developing a self-consistent theoretical framework based on the halo model [68] for predicting the auto and cross-spectra of fields on the lightcone, including dark matter density, electron density, velocities, galaxy number counts, the Newtonian potential and its time derivative, as well as the frequency-dependent contribution to the CMB from the thermal Sunyaev Zel'dovich (tSZ) effect and the Cosmic Infrared Background (CIB). We employ a coarse-graining scheme in radial distance along the lightcone based on Haar wavelets and use this complete basis to perform line of sight integrals for the kSZ, ISW, lensing, tSZ, CIB, and binned galaxy number counts. The contributions we include for the CMB represent the most important (in terms of the amplitude of power spectra) blackbody

and frequency dependent components. A challenging aspect of kSZ velocity reconstruction is that large-scale fields are reconstructed from small angular scale anisotropies, and so we develop tools to accurately compute spectra from the dipole down to sub-arcminute angular scales. We quantify the level of coarse graining given a fiducial CMB experiment and galaxy survey that will be necessary to capture the relevant cosmological information accessible using kSZ velocity reconstruction.

Another goal of this chapter is to assess the impact of various foregrounds and systematics on kSZ velocity reconstruction. Previous work has largely neglected these effects in forecasts. We assess the impact of extragalactic foregrounds by forecasting the level of residuals in auto and cross-spectra given a fiducial CMB experiment and the resulting effect on the variance of the quadratic estimators. We develop a formalism, analogous to bias hardening in CMB lensing [69], to remove biases associated with photometric redshift errors, and compute the variance of the resulting unbiased estimators in both a redshift-binned and principal component basis. The radial velocity quadratic estimators are biased by other sources of statistical anisotropy in the CMB-galaxy cross-power such as CMB lensing and the non-linear ISW signal. We confirm that these biases are small enough to be neglected in near-term experiments. A related systematic arises from redshift calibration errors on large angular scales, or any other effect that modulates the amplitude of the underlying statistically isotropic CMB anisotropies or galaxy number counts. This leads to a statistically anisotropic modulation of the CMB-galaxy cross-power induced by foreground residuals that can bias the radial velocity quadratic estimators. While this is the dominant source of estimator bias, we determine that it is below the estimator variance for the fiducial CMB experiment we consider.

To assess the impact of partial sky coverage for the CMB experiment and galaxy survey, we work with simulations in map space. We develop a numerical framework to produce sets of properly correlated CMB maps and redshift-binned galaxy number counts, assuming the underlying fields are Gaussian. We derive and implement a set of real space quadratic estimators and an associated pipeline to reconstruct the radial velocity field from an ensemble of simulated CMB maps (including both blackbody components and foreground residuals) and correlated binned galaxy number count maps. We find that no significant bias is introduced by masking, and confirm that the reconstructed power spectra are essentially scaled by the fraction of the sky that remains unmasked.

Our results suggest that the main limitations on kSZ velocity reconstruction will be various modelling errors that give rise to a biased reconstruction of the velocity field. The largest among these are biases introduced by photometric redshift errors and mismodelling of the galaxy-electron cross-power spectra, highlighting areas for future work. With a fixed experimental setup, improvements in the fidelity of the reconstruction can

be made through better foreground removal techniques on small angular scales. However, our investigations have not found any effect that seriously impacts the performance of kSZ velocity reconstruction as presented in previous literature, which is good news for this technique.

The plan of this chapter is as follows: Sec.2.2 presents the extended Ligthcone Picture formalism. Sec.2.3 describes our modelling of the many observables involved in kSZ velocity reconstruction. Sec.2.4 presents a forecast for the reconstruction of the radial velocity and an analysis of the new sources of bias and noise described in the extended formalism. Sec.2.5 presents our pipeline for generating realizations of properly correlated CMB maps and redshift-binned galaxy number counts, assuming the underlying fields are Gaussian, and use this to validate a velocity reconstruction pipeline and assess map-based systematics such as masking. We summarize our conclusions in Sec.2.6.

2.2 Formalism

In this section, we describe the extended formalism for kSZ velocity reconstruction in the Lightcone Picture. We begin by reviewing the projection of cosmological fields onto our past lightcone. Coarse-grained cosmological fields on the lightcone constitute the inputs to our estimator formalism, and we describe a coarse graining scheme as well as the statistics of the coarse grained fields. We outline the quadratic estimator formalism, proceeding from the simplest to the most realistic scenario.

2.2.1 Continuous fields on the lightcone

Continuous fields defined on our past lightcone constitute some of the basic building blocks of our formalism. A simple way of constructing a field on the lightcone is to take the projection of an underlying 4-dimensional space time field $\mathrm{U}(\eta, \mathbf{x})$, where η is the conformal time coordinate and $\mathbf{x}$ are the comoving spatial coordinates. If one parametrizes the lightcone with a unit direction vector $\hat{\mathbf{n}}$ and a comoving distance χ, then the projected field is defined by:

$$\mathrm{F}(\hat{\mathbf{n}}, \chi) \equiv \mathrm{U}(\eta(\chi), \mathbf{x} = \chi\hat{\mathbf{n}}). \tag{2.1}$$

In many cases, it is convenient to express $\mathrm{F}(\hat{\mathbf{n}}, \chi)$ in terms of the spatial Fourier moments of the field U, defined by:

$$\mathrm{U}(\eta, \mathbf{x}) = \int \frac{d^3\mathbf{k}}{(2\pi)^3}\, \tilde{\mathrm{U}}(\eta, \mathbf{k})\; e^{i\mathbf{k}\cdot\mathbf{x}}, \tag{2.2}$$

which gives

$$\mathrm{F}(\hat{\mathbf{n}},\chi) = \int \frac{d^3\mathbf{k}}{(2\pi)^3}\,\tilde{\mathrm{U}}(\eta(\chi),\mathbf{k})\; e^{i\mathbf{k}\cdot\chi\hat{\mathbf{n}}}. \tag{2.3}$$

The direction dependence of the lightcone field can be expanded in terms of spherical harmonics:

$$\mathrm{F}(\hat{\mathbf{n}},\chi) = \sum_{\ell m} \mathrm{F}_{\ell m}(\chi) Y_{\ell m}(\hat{\mathbf{n}}) \tag{2.4}$$

and we will refer to the coefficients $\mathrm{F}_{\ell m}(\chi)$ as the *lightcone moments* (LC moments for short) of the field F. The LC moments can be expressed as:

$$\mathrm{F}_{\ell m}(\chi) = \int \frac{d^2\hat{n}}{(2\pi)^3} Y^*_{\ell m}(\hat{\mathbf{n}}) \int \frac{d^3\mathbf{k}}{(2\pi)^3}\,\tilde{\mathrm{U}}(\eta(\chi),\mathbf{k})\; e^{i\mathbf{k}\cdot\chi\hat{\mathbf{n}}}, \tag{2.5}$$

and, using properties of the spherical harmonics and $e^{i\mathbf{k}\cdot\chi\hat{\mathbf{n}}}$, we can further simplify the expression to:

$$\mathrm{F}_{\ell m}(\chi) = \int \frac{d^3\mathbf{k}}{(2\pi)^3}\, 4\pi(i)^\ell\, j_\ell(k\chi)\,\tilde{\mathrm{U}}(\eta(\chi),\mathbf{k})\, Y^*_{\ell m}(\hat{k}), \tag{2.6}$$

where $j_\ell(k\chi)$ is a spherical Bessel function. It is possible to define fields on the lightcone using a more complex projection of the underlying field U than the one used in Eq.(2.1). For example, the projection could depend on the direction $\hat{\mathbf{n}}$ or introduce weights depending on the conformal time η. A more general expression for the LC moments is then:

$$\mathrm{F}_{\ell m}(\chi) = \int \frac{d^3\mathbf{k}}{(2\pi)^3}\, \mathcal{K}_\ell(\chi,k)\,\tilde{\mathrm{U}}(\eta(\chi),\mathbf{k})\, Y^*_{\ell m}(\hat{k}), \tag{2.7}$$

where $\mathcal{K}_\ell(\chi,k)$ is an integral kernel determined by the particular observable and typically containing linear combinations of spherical Bessel functions.

2.2.2 Integrated and coarse grained fields on the lightcone

A second type of building block of our formalism are line-of-sight integrals of continuous fields on the lightcone. Given a generic window function $W(\chi)$ we define the windowed F field:

$$\mathrm{F}^W(\hat{\mathbf{n}}) \equiv \int d\chi\, W(\chi)\,\mathrm{F}(\hat{\mathbf{n}},\chi), \tag{2.8}$$

and its spherical harmonic moments:

$$\mathrm{F}^W_{\ell m} \equiv \int d\chi\; W(\chi)\,\mathrm{F}_{\ell m}(\chi). \tag{2.9}$$

where $\mathrm{F}_{\ell m}(\chi)$ are the LC moments defined in the previous section. Given a finite portion of the lightcone determined by an interval $[\chi_{min}, \chi_{max}]$, we can consider a complete set of normalized functions $\mu^i(\chi)$ and expand the LC moments of the field F:

$$\mathrm{F}_{\ell m}(\chi) = \sum_i \mathrm{F}^i_{\ell m}\, \mu^i(\chi) \tag{2.10}$$

where the coefficients $\mathrm{F}^i_{\ell m}$ are obtained using Eq.(2.9) with $W(\chi) = \mu^i(\chi)$. From now on, we refer to these coefficients as the *μ-binned LC moments* of the field F. In this work, we choose to expand the LC moments in the radial direction using the Haar basis. Haar wavelets are defined on the interval $\chi_{\rm min} \leq \chi \leq \chi_{\rm max}$ by:

$$h^s(\chi) = \frac{1}{\sqrt{\chi_{\rm max} - \chi_{\rm min}}} \begin{cases} 2^{p/2}, & (q-1)\ 2^{-p} \leq (\frac{\chi}{\chi_{\rm max}-\chi_{\rm min}}) < (q-1/2)\ 2^{-p} \\ -2^{p/2}, & (q-1/2)\ 2^{-p} \leq (\frac{\chi}{\chi_{\rm max}-\chi_{\rm min}}) < q\ 2^{-p} \\ 0, & \text{otherwise} \end{cases} \tag{2.11}$$

with $s = 2^p + q - 1$ for integer p, q for $s > 0$; for $s = 0$, the Haar wavelet is $h^0(\chi) = \frac{1}{\sqrt{\chi_{\rm max}-\chi_{\rm min}}}$. Visually, Haar wavelets consist on a sequence of re-scaled top-hat functions (see Fig.2.1). The scale is determined by p and the location is determined by q; for each value of s there is a unique choice of p, q. The Haar basis functions are orthonormal over the interval $\chi_{\rm min} \leq \chi \leq \chi_{\rm max}$:

$$\int_{\chi_{\rm min}}^{\chi_{\rm max}} d\chi\ h^s(\chi) h^{s'}(\chi) = \delta_{ss'} \tag{2.12}$$

We choose the the Haar basis to expand the LC moments because conveniently the truncated Haar expansion up to $s = N - 1$ is equivalent to representing the LC moments by their average values in comoving bins of equal size $\Delta\chi = \frac{\chi_{\rm max}-\chi_{\rm min}}{N}$:

$$\sum_{s=0}^{N-1} \mathrm{F}^s_{\ell m}\ h^s(\chi) = \sum_{\alpha=0}^{N-1} \mathrm{F}^\alpha_{\ell m}\ \Pi^\alpha(\chi) \tag{2.13}$$

where

$$\Pi^\alpha(\chi) = \begin{cases} \frac{1}{\Delta\chi}, & \chi_{\rm min} + j\Delta\chi \leq \chi < \chi_{\rm min} + (j+1)\Delta\chi, \\ 0, & \text{otherwise} \end{cases} \tag{2.14}$$

and $\mathrm{F}^\alpha_{\ell m}$ are the Π-binned LC moments (we reserve Greek letters to index the Π-binned LC moments and Latin letters to index the Haar-binned LC moments). This property allows

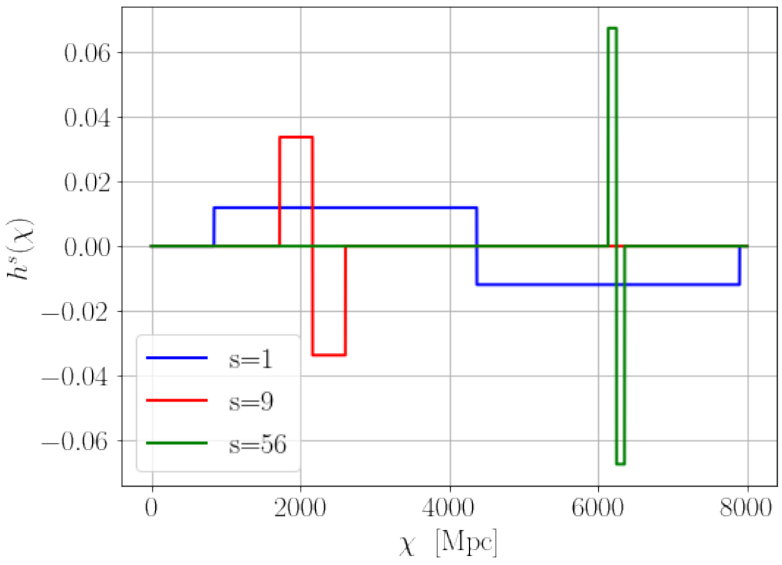

Figure 2.1: $s = 1, 9, 56$ Haar wavelets $h^s(\chi)$ for the interval $\chi \in [840, 7900]$ Mpc.

us to express the LC moments in a way that coarse-graining in the radial direction is clear:

$$\mathrm{F}_{\ell m}(\chi) = \sum_{\alpha=0}^{N-1} \mathrm{F}^{\alpha}_{\ell m}\, \Pi^{\alpha}(\chi) + \sum_{s=N}^{\infty} \mathrm{F}^{s}_{\ell m}\, h^{s}(\chi) \tag{2.15}$$

where the first sum represents the "coarse" or "bulk" radial modes and the second sum, orthogonal to the first, represents the "fine" modes that don't contribute to the bulk averages. We note that the spherical Fourier-Bessel decomposition (see e.g. [70]) could have been chosen instead of the Haar basis used here. In the context of galaxy redshift surveys, a comparison between the spherical Fourier-Bessel decomposition and the redshift-binned approach employed here can be found in Ref. [71]. Exploring the advantages of various choices of basis is deferred to future work.

2.2.3 Statistically isotropic correlations

The statistically isotropic correlations between μ-binned or integrated lightcone moments can conveniently be expressed in terms of a set of angular auto and cross-spectra that depend on ℓ and the window labels only. Let's consider two fields F and G on the lightcone, constructed from underlying 4-dimensional fields U^{F} and U^{G} as described in 2.2.1, and

integrated on the line of sight with windows W and W' respectively. The cross-spectra is:

$$\begin{aligned} C_\ell^{F^W G^{W'}} &\equiv \left\langle \mathrm{F}^W_{\ell m} \mathrm{G}^{W'}_{\ell m} \right\rangle \\ &= \int d\chi_1 d\chi_2 \; W(\chi_1) W'(\chi_2) \left\langle \mathrm{F}_{\ell m}(\chi_1) \mathrm{G}_{\ell m}(\chi_2) \right\rangle \\ &= \int d\chi_1 d\chi_2 \; W\left(\chi_1\right) W'\left(\chi_2\right) \\ &\times \int \frac{k^2 dk}{(2\pi)^3} \; \mathcal{K}^F_\ell(\chi_1, k) \, \mathcal{K}^G_\ell(\chi_2, k) \; P_{FG}(\chi_1, \chi_2, k), \end{aligned} \tag{2.16}$$

where we have assumed a statistically isotropic cross-correlation power spectrum between the underlying 4-dimensional space time fields:

$$\left\langle \tilde{\mathrm{U}}^F(\eta(\chi_1), \mathbf{k})^* \tilde{\mathrm{U}}^G(\eta(\chi_2), \mathbf{k}') \right\rangle = (2\pi)^3 P_{FG}(\chi_1, \chi_2, k) \delta^{(3)}(\mathbf{k} - \mathbf{k}'). \tag{2.17}$$

Although a brute force computation of the integrals in Eqs.(2.16) is feasible for certain values of ℓ and certain χ ranges, the oscillatory behaviour of the integral kernels makes a brute force approach a rather cumbersome one if accuracy across a wide range of multipole moments and redshifts is desired. This is exactly our case, as we aim to have consistent modelling of large-angle and small-angle observables across a large redshift range. The Limber approximation (see e.g. [72]) can be used to simplify the oscillatory integrals and provide accurate spectra under certain circumstances. For our purposes, an implementation of the Limber approximation is challenged by several factors: first, part of our calculations require narrow window functions, which can drive the Limber approximation beyond its regime of validity if the multipole ℓ is not high enough. Second, the Limber approximation only picks up the equal-time contribution to the cross-correlation power spectra ($\chi_1 = \chi_2$), and does not capture non-negligible contributions from unequal-time correlations [73]. Here we adopt the "Beyond Limber approximation" method from [74], which separates (2.16) into a term suitable for the Limber approximation and a term with separable structure that allows for fast Bessel integrations. We briefly summarize this method in Appendix B.1.

2.2.4 Statistically anisotropic cross-correlations

We now discuss our modelling for anisotropic cross-correlations between the temperature field and a windowed density tracer on the lightcone. We write the observed temperature

field as the sum of two contributions:

$$\Theta(\hat{\mathbf{n}}) = I(\hat{\mathbf{n}}) + \int d\chi\ M(\chi,\hat{\mathbf{n}})B(\chi,\hat{\mathbf{n}}) \tag{2.18}$$

where the first term $I(\hat{\mathbf{n}})$ represents all the contributions to the temperature coming from integrated lightcone fields:

$$I(\hat{\mathbf{n}}) = \sum_a \int d\chi\, W_a(\chi)\, \mathrm{F}_a(\hat{\mathbf{n}},\chi), \tag{2.19}$$

and the second term consists on the line of sight integration of the product of two lightcone fields. Consider as well a large scale structure tracer obtained as a line of sight integration of a density field $\delta(\chi,\hat{\mathbf{n}})$ on the lightcone:

$$\delta^W(\hat{\mathbf{n}}) = \int d\chi\, W(\chi)\, \delta(\hat{\mathbf{n}},\chi), \tag{2.20}$$

and we assume that $I(\hat{\mathbf{n}})$, $M(\chi,\hat{\mathbf{n}})$, $G(\chi,\hat{\mathbf{n}})$ and $\delta^W(\hat{\mathbf{n}})$ are isotropically correlated among each other as described in Sec. 2.2.3.

The second term in Eq.(2.18) leads to a statistical anisotropy when the temperature harmonics are correlated with the harmonics of the density field:

$$\begin{aligned}\left\langle \Theta_{\ell m}\, \delta^W_{\ell' m'} \right\rangle &= (-1)^m C_\ell^{I\delta^W} \delta_{\ell\ell'}\delta_{mm'} \\ &+ \sum_{\ell_1 m_1}\sum_{\ell_2 m_2} (-1)^m W^{\ell_1,\ell_2,\ell}_{m_1,m_2,-m} \int d\chi\, \left\langle M_{\ell_1 m_1}(\chi)\, B_{\ell_2 m_2}(\chi) \delta^W_{\ell' m'} \right\rangle\end{aligned} \tag{2.21}$$

where $W^{\ell_a,\ell_b,\ell_c}_{m_a,m_b,m_c}$ is an angular mode coupling matrix containing Wigner 3-j symbols; $M_{\ell_1 m_1}(\chi)$ and $B_{\ell_1 m_1}(\chi)$ are the LC moments of the fields $M(\chi,\hat{\mathbf{n}})$ and $B(\chi,\hat{\mathbf{n}})$ as defined in Sec.2.2.1. The type of mode coupling matrices that we will use in this chapter have the following structure:

$$\begin{aligned} W^{\ell_a,\ell_b,\ell_c}_{m_a,m_b,m_c} &= \omega(\ell_a,\ell_b,\ell_c)\sqrt{\frac{(2\ell_a+1)\,(2\ell_b+1)\,(2\ell_c+1)}{4\pi}} \\ &\times \begin{pmatrix} \ell_a & \ell_b & \ell_c \\ 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} \ell_a & \ell_b & \ell_c \\ m_a & m_b & m_c \end{pmatrix} \end{aligned} \tag{2.22}$$

were $\omega(\ell_a,\ell_b,\ell_c)$ is an extra scale dependence that appears only for some of the signals we consider. The three-point function in Eq.(2.21) can be approximated in the limit that two

of the (ℓ_1, ℓ_2, ℓ) multipoles are far greater than the remaining multipole (squeezed limit). In the squeezed limit where $\ell_1 \ll \ell, \ell_2$, we consider the field $M_{\ell_1 m_1}(\chi)$ to be deterministic, and therefore approximate:

$$\left\langle M_{\ell_1 m_1}(\chi)\, B_{\ell_2 m_2}(\chi)\, \delta^W_{\ell' m'} \right\rangle \approx\ M_{\ell_1 m_1}(\chi)\, \left\langle B_{\ell_2 m_2}(\chi)\, \delta^W_{\ell' m'} \right\rangle. \tag{2.23}$$

Inserting the expression above into Eq.(2.21) and expanding $M_{\ell_1 m_1}(\chi)$ and $B_{\ell_2 m_2}(\chi)$ using the Haar basis gives:

$$\begin{aligned} \left\langle \Theta_{\ell m}\, \delta^W_{\ell' m'} \right\rangle &= (-1)^m C_\ell^{I\delta^W} \delta_{\ell\ell'} \delta_{mm'} \\ &+ \sum_{s=0}^{\infty} \sum_{\ell_1 m_1} (-1)^{m_1} \begin{pmatrix} \ell & \ell' & \ell_1 \\ m & m' & -m_1 \end{pmatrix} f^{M^s W}_{\ell\ell_1\ell'}\, M^s_{\ell_1 m_1} \end{aligned} \tag{2.24}$$

where we have the defined the coupling:

$$f^{M^s W}_{\ell\ell_1\ell'} \equiv \omega(\ell, \ell_1, \ell') \sqrt{\frac{(2\ell+1)(2\ell_1+1)(2\ell'+1)}{4\pi}} \begin{pmatrix} \ell & \ell' & \ell_1 \\ 0 & 0 & 0 \end{pmatrix} C_{\ell'}^{B^s \delta^W}. \tag{2.25}$$

Eq.(2.24) tells us that the statistical anisotropy in the small angular scale temperature-density cross-correlation $\ell, \ell' \gg 1$ is modulated by the large angular scale field M. More precisely, each separate Haar-binned LC moment $M^s_{\ell_1 m_1}$ sources an independent statistically anisotropic term.

As discussed in Sec. 2.2.2, it is possible to use a mixed Π-binned and Haar-binned expansion by choosing a truncation value $s = N - 1$:

$$\begin{aligned} \left\langle \Theta_{\ell m}\, \delta^W_{\ell' m'} \right\rangle &= (-1)^m C_\ell^{I\delta^W} \delta_{\ell\ell'} \delta_{mm'} \\ &+ \sum_{\alpha=0}^{N-1} \sum_{\ell_1 m_1} (-1)^{m_1} \begin{pmatrix} \ell & \ell' & \ell_1 \\ m & m' & -m_1 \end{pmatrix} f^{M^\alpha W}_{\ell\ell_1\ell'}\, M^\alpha_{\ell_1 m_1} \\ &+ \sum_{s=N}^{\infty} \sum_{\ell_1 m_1} (-1)^{m_1} \begin{pmatrix} \ell & \ell' & \ell_1 \\ m & m' & -m_1 \end{pmatrix} f^{M^s W}_{\ell\ell_1\ell'}\, M^s_{\ell_1 m_1} \end{aligned} \tag{2.26}$$

where

$$f^{M^\alpha W}_{\ell\ell_1\ell'} \equiv \omega(\ell, \ell_1, \ell') \sqrt{\frac{(2\ell+1)(2\ell_1+1)(2\ell'+1)}{4\pi}} \begin{pmatrix} \ell & \ell' & \ell_1 \\ 0 & 0 & 0 \end{pmatrix} C_{\ell'}^{B^\alpha \delta^W} \Delta\chi. \tag{2.27}$$

Eq.(2.26) tells us that some of the statistical anisotropy comes from the 'coarse', 'bulk' radial modes $M^{\alpha}_{\ell_1 m_1}$ and the rest of the statistical anisotropy comes from the 'fine', small-scale modes $M^{s}_{\ell_1 m_1}$ that don't contribute to the bulk averages. Our formalism builds upon previous literature [19] in which the contribution from fine modes on the lightcone has not been considered and the statistical anisotropy is approximated as only sourced by the bulk modes. This can be a good approximation if N is high enough (how high depends on the radial profile of $M_{\ell_1 m_1}(\chi)$); in this work we will keep these terms and quantify their relevance in the modelling of the temperature-density statistical anisotropy.

We can add more generality to our modelling of the temperature-density cross-correlation by adding additional contributions to the temperature signal:

$$\Theta(\hat{\mathbf{n}}) = I(\hat{\mathbf{n}}) + \int d\chi\; M(\chi,\hat{\mathbf{n}})B(\chi,\hat{\mathbf{n}}) + \int d\chi\; Q(\chi,\hat{\mathbf{n}})D(\chi,\hat{\mathbf{n}}) + \dots \tag{2.28}$$

which translate to:

$$\begin{aligned}
\left\langle \Theta_{\ell m}\, \delta^{W}_{\ell' m'} \right\rangle \;=\;& (-1)^m C_\ell^{I\delta^W}\, \delta_{\ell\ell'}\delta_{mm'} \\
&+ \sum_{\alpha=0}^{N-1}\sum_{\ell_1 m_1} (-1)^{m_1} \begin{pmatrix} \ell & \ell' & \ell_1 \\ m & m' & -m_1 \end{pmatrix} f^{M^\alpha W}_{\ell\ell_1\ell'}\, M^{\alpha}_{\ell_1 m_1} \\
&+ \sum_{s=N}^{\infty}\sum_{\ell_1 m_1} (-1)^{m_1} \begin{pmatrix} \ell & \ell' & \ell_1 \\ m & m' & -m_1 \end{pmatrix} f^{M^s W}_{\ell\ell_1\ell'}\, M^{s}_{\ell_1 m_1} \\
&+ \sum_{\alpha=0}^{N-1}\sum_{\ell_1 m_1} (-1)^{m_1} \begin{pmatrix} \ell & \ell' & \ell_1 \\ m & m' & -m_1 \end{pmatrix} f^{Q^\alpha W}_{\ell\ell_1\ell'}\, Q^{\alpha}_{\ell_1 m_1} \\
&+ \sum_{s=N}^{\infty}\sum_{\ell_1 m_1} (-1)^{m_1} \begin{pmatrix} \ell & \ell' & \ell_1 \\ m & m' & -m_1 \end{pmatrix} f^{Q^s W}_{\ell\ell_1\ell'}\, Q^{s}_{\ell_1 m_1} \\
&+ \dots
\end{aligned} \tag{2.29}$$

2.2.5 Quadratic estimator

In this subsection we discuss how the Π-binned modes appearing in the statistical anisotropy Eq.(2.29) can be estimated by constructing appropriately weighted sums of products of temperature and density multipoles. The most general case discussed in the previous subsection included statistical anisotropies sourced by the Π-binned and Haar-binned LC

moments of a series of modulating fields ($M(\hat{\mathbf{n}},\chi)$, $Q(\hat{\mathbf{n}},\chi)$, ...). We will first consider the simplified case in which there is only one modulating field $M(\hat{\mathbf{n}},\chi)$ and a single Π mode α and show how the estimation works. After that, we progressively add layers of complexity until we reach the most general case.

Case 1: Single modulating field and single Π mode

Let's start by considering the simple case in which there is only one modulating bulk mode $M^{\alpha}_{\ell_1 m_1}$ sourcing the statistical anisotropy in the temperature-density cross-correlation. Starting from Eq.(2.18), we write the temperature multipoles as:

$$\Theta_{\ell m} = I_{\ell m} + \sum_{\ell_1 m_1}\sum_{\ell_2 m_2}(-1)^m W^{\ell_1,\ell_2,\ell}_{m_1,m_2,-m} M^{\alpha}_{\ell_1 m_1}\, B^{\alpha}_{\ell_2 m_2}\Delta\chi \tag{2.30}$$

and write the temperature cross-correlation as:

$$\begin{aligned}\left\langle \Theta_{\ell m}\, \delta^{W}_{\ell' m'} \right\rangle &= (-1)^m C^{I\delta^W}_{\ell}\delta_{\ell\ell'}\delta_{mm'} \\ &+ \sum_{\ell_1 m_1}(-1)^{m_1}\begin{pmatrix}\ell & \ell' & \ell_1 \\ m & m' & -m_1\end{pmatrix} f^{M^{\alpha}W}_{\ell\ell_1\ell'}\, M^{\alpha}_{\ell_1 m_1}\end{aligned} \tag{2.31}$$

Let's construct a quadratic sum of temperature and density multipoles with the following structure

$$\hat{M}^{\alpha}_{LM} = A^{M^{\alpha}}_{L}\sum_{\ell m;\ell' m'}(-1)^M\begin{pmatrix}\ell & \ell' & L \\ m & m' & -M\end{pmatrix} G^{M^{\alpha}W}_{\ell\ell' L}\ \Theta_{\ell m}\delta^{W}_{\ell' m'} \tag{2.32}$$

and choose weights $G^{M^{\alpha}W}_{\ell\ell' L}$ such that the estimator is unbiased:

$$\left\langle \hat{M}^{\alpha}_{LM} \right\rangle = M^{\alpha}_{LM} \tag{2.33}$$

and has minimum variance. The first condition translates to

$$A^{M^{\alpha}}_{L} = (2L+1)\left(\sum_{\ell;\ell'} G^{M^{\alpha}W}_{\ell\ell' L} f^{M^{\alpha}W}_{\ell L\ell'}\right)^{-1}. \tag{2.34}$$

The minimum variance estimator can be found using the Lagrange multiplier method subject to the constraint Eq.(2.34), which gives:

$$G^{M^{\alpha}W}_{\ell\ell' L} \equiv \frac{C^{\Theta\Theta}_{\ell'} C^{\delta^W\delta^W}_{\ell} f^{M^{\alpha}W}_{\ell L\ell'} - (-1)^{\ell+\ell'+L} C^{I\delta^W}_{\ell} C^{I\delta^W}_{\ell'} f^{M^{\alpha}W}_{\ell' L\ell}}{C^{\Theta\Theta}_{\ell} C^{\Theta\Theta}_{\ell'} C^{\delta^W\delta^W}_{\ell} C^{\delta^W\delta^W}_{\ell'} - \left(C^{I\delta^W}_{\ell}\right)^2\left(C^{I\delta^W}_{\ell'}\right)^2} \tag{2.35}$$

where $C_\ell^{\Theta\Theta}$ is the full temperature power spectrum. A complete derivation can be found in Appendix B.2. In the computation of the estimator variance used to derive $G_{\ell\ell' L}^{M^\alpha W}$, we have only included the disconnected part of the temperature-galaxy-temperature-galaxy four-point function. Under this approximation, the estimator variance is given by:

$$\left\langle \hat{M}^\alpha_{L'M'} \hat{M}^{\alpha *}_{LM} \right\rangle = C_L^{M^\alpha M^\alpha} + A_L^{M^\alpha} \tag{2.36}$$

However, looking in more detail at the estimator variance:

$$\begin{aligned}
\left\langle \hat{M}^\alpha_{LM} \hat{M}^{\alpha *}_{L'M'} \right\rangle &= A_L^{M^\alpha} A_{L'}^{M^\alpha} \sum_{\ell_a m_a; \ell_b m_b} \sum_{\ell_c m_c; \ell_d m_d} (-1)^{M+M'} \begin{pmatrix} \ell_a & \ell_b & L \\ m_a & m_b & -M \end{pmatrix} \begin{pmatrix} \ell_c & \ell_d & L' \\ m_c & m_d & -M' \end{pmatrix} \\
&\times G^{M^\alpha W}_{\ell_a \ell_b L} \; G^{M^\alpha W}_{\ell_c \ell_d L'} \left\langle \Theta_{\ell_a m_a} \delta^W_{\ell_b m_b} \Theta^*_{\ell_c m_c} \delta^{W*}_{\ell_d m_d} \right\rangle \\
&= A_L^{M^\alpha} A_{L'}^{M^\alpha} \sum_{\ell_a m_a; \ell_b m_b} \sum_{\ell_c m_c; \ell_d m_d} (-1)^{M+M'} \begin{pmatrix} \ell_a & \ell_b & L \\ m_a & m_b & -M \end{pmatrix} \begin{pmatrix} \ell_c & \ell_d & L' \\ m_c & m_d & -M' \end{pmatrix} \\
&\times G^{M^\alpha W}_{\ell_a \ell_b L} \; G^{M^\alpha W}_{\ell_c \ell_d L'} \Bigg[\left\langle I_{\ell_a m_a} \delta^W_{\ell_b m_b} I^*_{\ell_c m_c} \delta^{W*}_{\ell_d m_d} \right\rangle + (\Delta\chi)^2 \sum_{\ell_1 m_1} \sum_{\ell_2 m_2} \sum_{\ell_1' m_1'} \sum_{\ell_2' m_2'} \\
&\times (-1)^{m_a + m_c} W^{\ell_1, \ell_2, \ell_a}_{m_1, m_2, -m_a} W^{\ell_1', \ell_2', \ell_c}_{m_1', m_2', -m_c} \left\langle M^\alpha_{\ell_1 m_1} \; B^\alpha_{\ell_2 m_2} \delta^W_{\ell_b m_b} M^{\alpha *}_{\ell_1' m_1'} \; B^{\alpha *}_{\ell_2' m_2'} \delta^{W*}_{\ell_d m_d} \right\rangle \Bigg]
\end{aligned} \tag{2.37}$$

we see that the estimator variance depends on a six-point function of the underlying fields. Therefore, even if all the fields are Gaussian, the disconnected four-point function is not a complete description of the estimator variance – one must in principle include the 15 terms that contribute to the disconnected six-point function. Fortunately, as we describe in Appendix B.3, for the observables considered in this work the relevant components of the six-point function do not yield any significant additional variance beyond the terms in Eq. 2.36. This additional contribution to the variance was computed for kSZ tomography in the box formalism in Ref. [21], where in analogy with lensing reconstruction, it was referred to as the $N^{(1)}$ bias. For non-Gaussian fields, one must additionally compute the connected part of the six-point function. This was also computed in Ref. [21], where it was shown that this "$N^{(3/2)}$ bias" is far larger than the $N^{(1)}$ bias, and can even become comparable to $A_L^{M^\alpha}$ at sufficiently high SNR. A full computation of these additional contributions to the variance within the Lightcone Picture will appear in future work. Moving forward, we will only consider the contribution from $C_L^{M^\alpha M^\alpha}$ and $A_L^{M^\alpha}$ in our estimator variance.

Case 2: Single modulating field and multiple Π modes

The next step to add more realism is to have multiple Π-binned LC moments from a single modulating field. In this case we have:

$$\Theta_{\ell m} = I_{\ell m} + \sum_{\alpha=0}^{N-1} \sum_{\ell_1 m_1} \sum_{\ell_2 m_2} (-1)^m W^{\ell_1,\ell_2,\ell}_{m_1,m_2,-m} M^{\alpha}_{\ell_1 m_1} \, B^{\alpha}_{\ell_2 m_2} \Delta\chi \tag{2.38}$$

and

$$\begin{aligned} \left\langle \Theta_{\ell m} \, \delta^W_{\ell' m'} \right\rangle &= (-1)^m C_\ell^{I\delta^W} \delta_{\ell\ell'} \delta_{mm'} \\ &+ \sum_{\alpha=0}^{N-1} \sum_{\ell_1 m_1} (-1)^{m_1} \begin{pmatrix} \ell & \ell' & \ell_1 \\ m & m' & -m_1 \end{pmatrix} f^{M^\alpha W}_{\ell\ell_1\ell'} \, M^{\alpha}_{\ell_1 m_1} \end{aligned} \tag{2.39}$$

We want to construct N unbiased quadratic estimators, one for each modulating source $M^{\alpha}_{\ell_1 m_1}$. The strategy we choose is to first construct N *biased* estimators, by taking N versions of the single field estimator Eq.(2.32) described in Case 1:

$$\begin{aligned} \hat{M}^0_{LM} &= A_L^{M^0} \sum_{\ell m;\ell' m'} (-1)^M \begin{pmatrix} \ell & \ell' & L \\ m & m' & -M \end{pmatrix} G^{M^0 W}_{\ell\ell' L} \, \Theta_{\ell m} \delta^W_{\ell' m'} \\ &\vdots \\ \hat{M}^{N-1}_{LM} &= A_L^{M^{N-1}} \sum_{\ell m;\ell' m'} (-1)^M \begin{pmatrix} \ell & \ell' & L \\ m & m' & -M \end{pmatrix} G^{M^{N-1} W}_{\ell\ell' L} \, \Theta_{\ell m} \delta^W_{\ell' m'} \end{aligned} \tag{2.40}$$

where the weights $A_L^{M^\alpha}, G^{M^\alpha,W}_{\ell\ell' L}$ are chosen exactly as if $M^{\alpha}_{\ell_1 m_1}$ was the only source of statistical anisotropy. These estimators will be biased:

$$\left\langle \hat{M}^0_{LM} \right\rangle = M^0_{LM} + \sum_{\alpha\neq 0} \hat{M}^{\alpha}_{LM} \, \frac{A_L^{M^0}}{2L+1} \sum_{\ell\ell'} G^{M^0 W}_{\ell\ell' L} \, f^{M^\alpha W}_{\ell L\ell'} \tag{2.41}$$

$$\vdots$$

$$\left\langle \hat{M}^{N-1}_{LM} \right\rangle = M^{N-1}_{LM} + \sum_{\alpha\neq N-1} \hat{M}^{\alpha}_{LM} \, \frac{A_L^{M^{N-1}}}{2L+1} \sum_{\ell\ell'} G^{M^{N-1} W}_{\ell\ell' L} \, f^{M^\alpha W}_{\ell L\ell'} \tag{2.42}$$

We can define a "rotation matrix":

$$(\mathbf{R}_L)^{XY} \equiv \frac{\sum_{\ell\ell'} G^{XW}_{\ell\ell' L} f^{YW}_{\ell L\ell'}}{\sum_{\ell\ell'} G^{XW}_{\ell\ell' L} f^{XW}_{\ell L\ell'}} \tag{2.43}$$

with indices X, Y in $(M^0, \ldots, M^{N-1})$. With this matrix we write the system Eqs.(2.41) as:

$$\left\langle \hat{\mathbf{d}}_{LM} \right\rangle = \mathbf{R}_L \cdot \mathbf{d}_{LM}, \tag{2.44}$$

where $\hat{\mathbf{d}}_{LM} = (\hat{M}^0_{LM}, \ldots, \hat{M}^{N-1}_{LM})$ and $\mathbf{d}_{LM} = (M^\alpha_{LM}, \ldots, M^{N-1}_{LM})$. If the rotation matrix is invertible, we can now define unbiased estimators $\tilde{\mathbf{d}}_{LM}$ for $\mathbf{d}_{LM}$:

$$\tilde{\mathbf{d}}_{LM} \equiv \mathbf{R}_L^{-1} \cdot \hat{\mathbf{d}}_{LM} \tag{2.45}$$

The procedure above serves as an example of *bias hardening* the quadratic estimators [69] in the presence of multiple sources of statistical anisotropy in the temperature-density cross-correlation[1].

The two point function for the unbiased estimator is:

$$\left\langle \tilde{\mathbf{d}}_{LM} \tilde{\mathbf{d}}^\dagger_{LM} \right\rangle = \mathbf{R^{-1}}_L \cdot \left\langle \hat{\mathbf{d}}_{LM} \hat{\mathbf{d}}^\dagger_{LM} \right\rangle \cdot (\mathbf{R^{-1}}_L)^\dagger \tag{2.46}$$

where $\left\langle \hat{\mathbf{d}}_{LM} \hat{\mathbf{d}}^\dagger_{LM} \right\rangle$ is the two point function of the biased estimator; note that this is a matrix containing all auto and cross-spectra. Similarly to Case 1, the biased two point function can be written in terms of 4-point and 6-point functions and broken down into

[1] It is important to highlight that these unbiased estimators won't necessarily be the minimum variance estimators for $\mathbf{M}_{LM}$. It is possible to construct an unbiased and minimum variance estimator by correlating linear combinations of δ^W maps with Θ. However, as we demonstrate below, the simpler approach taken here of deriving separate estimators and then rotating yields very good results in practice. Exploring such new estimators is deferred to future work.

signal and noise terms:

$$\begin{aligned}
&\left\langle (\hat{\mathbf{d}}_{LM})^X (\hat{\mathbf{d}}^*_{L'M'})^Y \right\rangle \\
= \;& A_L^X A_{L'}^Y \sum_{\ell_a m_a;\ell_b m_b} \sum_{\ell_c m_c;\ell_d m_d} (-1)^{M+M'} \begin{pmatrix} \ell_a & \ell_b & L \\ m_a & m_b & -M \end{pmatrix} \begin{pmatrix} \ell_c & \ell_d & L' \\ m_c & m_d & -M' \end{pmatrix} \\
\times \;& G^{XW}_{\ell_a \ell_b L}\, G^{YW}_{\ell_c \ell_d L'} \left\langle \Theta_{\ell_a m_a} \delta^W_{\ell_b m_b} \Theta^*_{\ell_c m_c} \delta^{W*}_{\ell_d m_d} \right\rangle \\
= \;& A_L^X A_{L'}^Y \sum_{\ell_a m_a;\ell_b m_b} \sum_{\ell_c m_c;\ell_d m_d} (-1)^{M+M'} \begin{pmatrix} \ell_a & \ell_b & L \\ m_a & m_b & -M \end{pmatrix} \begin{pmatrix} \ell_c & \ell_d & L' \\ m_c & m_d & -M' \end{pmatrix} \\
\times \;& G^{XW}_{\ell_a \ell_b L}\, G^{YW}_{\ell_c \ell_d L'} \Bigg[\left\langle I_{\ell_a m_a} \delta^W_{\ell_b m_b} I^*_{\ell_c m_c} \delta^{W*}_{\ell_d m_d} \right\rangle + \sum_{\ell_1 m_1} \sum_{\ell_2 m_2} \sum_{\ell'_1 m'_1} \sum_{\ell'_2 m'_2} \\
\times \;& (-1)^{m_a + m_c} W^{\ell_1, \ell_2, \ell_a}_{m_1, m_2, -m_a} W^{\ell'_1, \ell'_2, \ell_c}_{m'_1, m'_2, -m_c} \\
\times \;& \left\langle (\mathbf{X}_{\ell_1 m_1} \cdot \mathbf{Y}_{\ell_2 m_2})\, \delta^W_{\ell_b m_b} \left(\mathbf{X}^*_{\ell'_1 m'_1} \cdot \mathbf{Y}^*_{\ell'_2 m'_2} \right) \delta^{W*}_{\ell_d m_d} \right\rangle \Bigg]
\end{aligned} \tag{2.47}$$

where X, Y are indices in $(M^0, \ldots, M^{N-1})$ and $\mathbf{X}_{\ell m}, \mathbf{Y}_{\ell m}$ are the vectors:

$$\mathbf{X}_{\ell m} = (M^0_{\ell m}, \ldots, M^{N-1}_{\ell m}) \tag{2.48}$$

$$\mathbf{Y}_{\ell m} = (\Delta\chi B^0_{\ell m}, \ldots, \Delta\chi B^{N-1}_{\ell m}) \tag{2.49}$$

The dominant terms of the 2-point function of the bias estimator are:

$$\left\langle \hat{\mathbf{d}}_{LM} \hat{\mathbf{d}}^\dagger_{LM} \right\rangle = \mathbf{R}_L \mathbf{C}_L^{MM} (\mathbf{R}_L)^\dagger + \mathbf{N}^0_L + \ldots \tag{2.50}$$

where $\mathbf{C}_L^{MM}$ is the modulating field covariance matrix and the elements of the $\mathbf{N}^0_L$ matrix are given by:

$$(\mathbf{N}^0_L)^{\alpha\beta} = \frac{A_L^{M^\alpha} A_L^{M^\beta}}{2L+1} \sum_{\ell_1 \ell_2} G^{M^\alpha W}_{\ell_1 \ell_2 L} \Big[G^{M^\beta W}_{\ell_1 \ell_2 L} C_{\ell_1}^{\Theta\Theta} C_{\ell_2}^{\delta^W \delta^W} + (-1)^{\ell_1+\ell_2+L} G^{M^\beta W}_{\ell_2 \ell_1 L} C_{\ell_1}^{\Theta\delta^W} C_{\ell_2}^{\Theta\delta^W} \Big] \tag{2.51}$$

It is easy to check that the diagonal elements satisfy $(\mathbf{N}^0_L)^{\alpha\alpha} = A^{M^\alpha}$. This is to be expected because we constructed the Case 2 estimators as collections of the Case 1 estimator. Eq.(2.50) has further contributions from the 6-point function: some relatively simple terms proportional to the covariances $\mathbf{C}_L^{BB}$,$\mathbf{C}_L^{\delta^W \delta^W}$,$\mathbf{C}_L^{\delta^W B}$, etc., and more complicated terms coming from various contractions of the 6-point function. Again, these contributions are expected to be small enough to be neglected.

Case 3: Single modulating field, multiple Π and Haar modes

The estimators constructed in Case 2 ignore the contributions to the temperature multipoles and temperature-density statistical anisotropy coming from the Haar-binned LC moments of the field $M(\hat{\mathbf{n}}, \chi)$:

$$\Theta_{\ell m} = I_{\ell m} + \sum_{\alpha=0}^{N-1} \sum_{\ell_1 m_1} \sum_{\ell_2 m_2} (-1)^m W^{\ell_1,\ell_2,\ell}_{m_1,m_2,-m} M^{\alpha}_{\ell_1 m_1} B^{\alpha}_{\ell_2 m_2} \Delta\chi \tag{2.52}$$

$$+ \sum_{s=0}^{s_{max}} \sum_{\ell_1 m_1} \sum_{\ell_2 m_2} (-1)^m W^{\ell_1,\ell_2,\ell}_{m_1,m_2,-m} M^{s}_{\ell_1 m_1} B^{s}_{\ell_2 m_2} \tag{2.53}$$

$$\begin{aligned} \left\langle \Theta_{\ell m}\, \delta^W_{\ell' m'} \right\rangle &= (-1)^m C_\ell^{I\delta^W} \delta_{\ell\ell'} \delta_{mm'} \\ &+ \sum_{\alpha=0}^{N-1} \sum_{\ell_1 m_1} (-1)^{m_1} \begin{pmatrix} \ell & \ell' & \ell_1 \\ m & m' & -m_1 \end{pmatrix} f^{M^\alpha W}_{\ell\ell_1\ell'}\, M^{\alpha}_{\ell_1 m_1} \\ &+ \sum_{s=N}^{s_{max}} \sum_{\ell_1 m_1} (-1)^{m_1} \begin{pmatrix} \ell & \ell' & \ell_1 \\ m & m' & -m_1 \end{pmatrix} f^{M^s W}_{\ell\ell_1\ell'}\, M^{s}_{\ell_1 m_1} \end{aligned} \tag{2.54}$$

.

where s_{max} is chosen such that Haar expansion is mostly converged. If we follow the same steps as in Case 2 to construct N biased estimators, the Haar modes lead to an additional bias:

$$\begin{aligned} \left\langle \hat{M}^0_{LM} \right\rangle\Big|_{Case\ 3} &\to \left\langle \hat{M}^0_{LM} \right\rangle\Big|_{Case\ 2} + \sum_{s=N}^{s_{max}} M^s_{LM}\, \frac{A_L^{M^0}}{2L+1} \sum_{\ell\ell'} G^{M^0 W}_{\ell\ell' L}\, f^{M^s W}_{\ell\ell_2\ell'} \\ &\vdots \\ \left\langle \hat{M}^{N-1}_{LM} \right\rangle\Big|_{Case\ 3} &\to \left\langle \hat{M}^{N-1}_{LM} \right\rangle\Big|_{Case\ 2} + \sum_{s=N}^{s_{max}} M^s_{LM}\, \frac{A_L^{M^{N-1}}}{2L+1} \sum_{\ell\ell'} G^{M^{N-1} W}_{\ell\ell' L}\, f^{M^s W}_{\ell\ell_2\ell'} \end{aligned} \tag{2.55}$$

The relevance of this bias depends on the truncation number N, which in principle can be chosen to be high enough such that the contribution from Haar modes can be ignored.

Thus, quantifying the size of these terms as a function of N is a useful way of determining the level of coarse graining that we need to model our observables with.

The 2-point function of the estimator can be computed using the same expression Eq.(2.47) as in Case 2 just by expanding the vectors $\mathbf{X}, \mathbf{Y}$ defined in Eqs.(2.48-2.49):

$$\mathbf{X}_{\ell m} = (M^0_{\ell m}, \ldots, M^{N-1}_{\ell m}, M^{s=N}_{\ell m}, \ldots, M^{s=s_{max}}_{\ell m}). \tag{2.56}$$

$$\mathbf{Y}_{\ell m} = (\Delta\chi B^0_{\ell m}, \ldots, \Delta\chi B^{N-1}_{\ell m}, B^{s=N}_{\ell m}, \ldots, B^{s=s_{max}}_{\ell m}). \tag{2.57}$$

The resulting 2-point function contains three terms which we identify as dominant:

$$\left\langle \hat{\mathbf{d}}_{LM} \hat{\mathbf{d}}^\dagger_{LM} \right\rangle = \mathbf{R}_L \mathbf{C}^{MM}_L (\mathbf{R}_L)^\dagger + \mathbf{N}^0_L + \mathbf{N}^{MM\ fine}_L + \ldots \tag{2.58}$$

where $\mathbf{R}_L$ is the rotation matrix defined in Eq.(2.43), $\mathbf{C}^{MM}_L$ is the modulating field covariance matrix, $\mathbf{N}^0_L$ computed exactly as in Case 2, and $\mathbf{N}^{MM\ fine}_L$ is given by:

$$\begin{aligned} \left(\mathbf{N}^{MM\ fine}_L\right)^{\alpha\beta} &= \sum_{s,s'=N}^{k_{max}} \left\langle M^s_{LM} M^{s'*}_{LM} \right\rangle \\ &\times \left(\sum_{\ell_1\ell_2} \frac{A^{M^\alpha}_L}{2L+1} G^{M^\alpha W}_{\ell_1\ell_2 L} f^{M^s W}_{\ell_1 L \ell_2} \right) \left(\sum_{\ell_1\ell_2} \frac{A^{M^\beta}_L}{2L+1} G^{M^\beta W}_{\ell_1\ell_2 L} f^{M^{s'} W}_{\ell_1 L \ell_2} \right) \end{aligned} \tag{2.59}$$

We call this term the fine mode noise, as it is sourced by the Haar modes of the modulating field above the truncation number N. $\mathbf{N}^{MM\ fine}_L$ can become comparable to $\mathbf{N}^0_L$ for low enough N. Conversely, one can find a high enough truncation number N such that the fine mode noise can be neglected. In a realistic scenario, the truncation number is limited by the details of the 3-dimensional large-scale structure survey that is being used for the reconstruction. In further sections we will show the size of the fine mode noise in the estimation of the radial velocity Π-binned LC moments.

Case 4: Multiple modulating fields, multiple Π and Haar modes

Generalizing the results from the previous cases to the multiple field case is straightforward. The temperature and the statistical anisotropy are:

$$
\begin{aligned}
\Theta_{\ell m} = I_{\ell m} \;&+\; \sum_{\alpha=0}^{N-1}\sum_{\ell_1 m_1}\sum_{\ell_2 m_2}(-1)^m W^{\ell_1,\ell_2,\ell}_{m_1,m_2,-m} M^{\alpha}_{\ell_1 m_1}\; B^{\alpha}_{\ell_2 m_2}\Delta\chi \\
&+\; \sum_{s=N}^{s_{max}}\sum_{\ell_1 m_1}\sum_{\ell_2 m_2}(-1)^m W^{\ell_1,\ell_2,\ell}_{m_1,m_2,-m} M^{s}_{\ell_1 m_1}\; B^{s}_{\ell_2 m_2} \\
&+\; \sum_{\alpha=0}^{N-1}\sum_{\ell_1 m_1}\sum_{\ell_2 m_2}(-1)^m W^{\ell_1,\ell_2,\ell}_{m_1,m_2,-m} Q^{\alpha}_{\ell_1 m_1}\; D^{\alpha}_{\ell_2 m_2}\Delta\chi \\
&+\; \sum_{s=N}^{s_{max}}\sum_{\ell_1 m_1}\sum_{\ell_2 m_2}(-1)^m W^{\ell_1,\ell_2,\ell}_{m_1,m_2,-m} Q^{s}_{\ell_1 m_1}\; D^{s}_{\ell_2 m_2} \\
&\;\vdots
\end{aligned}
\tag{2.60}
$$

$$
\begin{aligned}
\left\langle \Theta_{\ell m}\, \delta^W_{\ell' m'} \right\rangle \;&=\; (-1)^m C^{I\delta^W}_{\ell}\, \delta_{\ell\ell'}\delta_{mm'} \\
&+\; \sum_{\alpha=0}^{N-1}\sum_{\ell_1 m_1}(-1)^{m_1}\begin{pmatrix}\ell & \ell' & \ell_1 \\ m & m' & -m_1\end{pmatrix} f^{M^\alpha W}_{\ell\ell_1\ell'}\; M^{\alpha}_{\ell_1 m_1} \\
&+\; \sum_{s=0}^{s_{max}}\sum_{\ell_1 m_1}(-1)^{m_1}\begin{pmatrix}\ell & \ell' & \ell_1 \\ m & m' & -m_1\end{pmatrix} f^{M^s W}_{\ell\ell_1\ell'}\; M^{s}_{\ell_1 m_1} \\
&+\; \sum_{\alpha=0}^{N-1}\sum_{\ell_1 m_1}(-1)^{m_1}\begin{pmatrix}\ell & \ell' & \ell_1 \\ m & m' & -m_1\end{pmatrix} f^{Q^\alpha W}_{\ell\ell_1\ell'}\; Q^{\alpha}_{\ell_1 m_1} \\
&+\; \sum_{s=0}^{s_{max}}\sum_{\ell_1 m_1}(-1)^{m_1}\begin{pmatrix}\ell & \ell' & \ell_1 \\ m & m' & -m_1\end{pmatrix} f^{Q^s W}_{\ell\ell_1\ell'}\; Q^{s}_{\ell_1 m_1} \\
&\;\vdots
\end{aligned}
\tag{2.61}
$$

Similarly to the previous cases, we can construct N biased estimators for the Π-binned LC moments of the M field. The 2-point function is calculated using Eq.(2.47) as in Case 2 just by expanding the vectors $\mathbf{X}_{\ell m}, \mathbf{Y}_{\ell m}$ space to:

$$
\mathbf{X}_{\ell m} = (M^0_{\ell m}, \dots, M^{N-1}_{\ell m}, M^{k=N}_{\ell m}, \dots, M^{k=k_{max}}_{\ell m}, Q^0_{\ell m}, \dots, Q^{N-1}_{\ell m}, Q^{k=N}_{\ell m}, \dots, Q^{k=k_{max}}_{\ell m}, \dots) \tag{2.62}
$$

$$
\mathbf{Y}_{\ell m} = (\Delta\chi B^0_{\ell m}, \dots, \Delta\chi B^{N-1}_{\ell m}, B^{k=N}_{\ell m}, \dots, B^{k=k_{max}}_{\ell m}, \Delta\chi D^0_{\ell m}, \dots, \Delta\chi D^{N-1}_{\ell m}, D^{k=N}_{\ell m}, \dots, D^{k=k_{max}}_{\ell m}, \dots) \tag{2.63}
$$

The number of terms contributing to the 2-point function will clearly increase with the introduction of new modulating fields. In addition to the terms from Case 3, each new modulating field Q will introduce a term $\mathbf{N}_L^{QQ}$ given by:

$$\begin{aligned}\left(\mathbf{N}_L^{QQ}\right)^{\alpha\beta} &= \sum_{\alpha',\beta'=0}^{N-1} \left\langle Q_{LM}^{\alpha'} Q_{LM}^{\beta'*} \right\rangle \\ &\times \left(\sum_{\ell_1\ell_2} \frac{A_L^{M^\alpha}}{2L+1} G_{\ell_1\ell_2 L}^{M^\alpha W} f_{\ell_1 L\ell_2}^{Q^{\alpha'} W} \right) \left(\sum_{\ell_1\ell_2} \frac{A_L^{M^\beta}}{2L+1} G_{\ell_1\ell_2 L}^{M^\beta W} f_{\ell_1 L\ell_2}^{Q^{\beta'} W} \right), \end{aligned} \tag{2.64}$$

as well as other less significant terms with a similar structure. In principle, there is no immediate way to determine if $\mathbf{N}_L^{QQ}$ are negligible respect to $\mathbf{N}_L^0$ and $\mathbf{N}_L^{MM\ fine}$, as this depends on the specific details of the modulating fields giving rise to the temperature signal. We show examples comparison in Sec.2.4, where we compute the reconstruction noise for the remote CMB dipole.

Multiple density windows

Our formalism so far has discussed the construction of estimators for Π-binned LC moments of a lightcone field $M(\hat{\mathbf{n}}, \chi)$ using only one window window function $W(\chi)$ for the density tracer. We remind the reader that the window function shows up in the coupling Eq.(2.27) through the cross-spectra $C_{\ell'}^{B^\alpha \delta^W} \equiv \left\langle \mathrm{B}_{\ell m}^\alpha \delta_{\ell m}^W \right\rangle$, where $B(\hat{\mathbf{n}}, \chi)$ is the field integrated together with $M(\hat{\mathbf{n}}, \chi)$ to form a temperature signal. Since the estimator utilizes large multipoles ℓ, ℓ', the cross-spectra will be non-negligible only if the window function $W(\chi)$ overlaps with $\Pi^\beta(\chi)$. A density window function with a wide support on the lightcone will lead to well defined couplings and estimators for all the Π-bins, but at the same time leads to a increased mix in the 3-dimensional information we are trying to reconstruct. In contrast, more localized density window functions will be better at isolating contributions coming from different redshifts, but can lead to ill-defined estimators for Π-bins with zero overlap with the density window. One can remediate this last issue by constructing estimators with a changing density window function:

$$
\begin{aligned}
\hat{M}^0_{LM} &= A_L^{M^0W^0} \sum_{\ell m;\ell' m'} (-1)^M \begin{pmatrix} \ell & \ell' & L \\ m & m' & -M \end{pmatrix} G^{M^0W^0}_{\ell\ell' L}\ \Theta_{\ell m}\delta^{W^0}_{\ell' m'} \\
&\vdots \\
\hat{M}^{N-1}_{LM} &= A_L^{M^{N-1}W^{N-1}} \sum_{\ell m;\ell' m'} (-1)^M \begin{pmatrix} \ell & \ell' & L \\ m & m' & -M \end{pmatrix} G^{M^{N-1}W^{N-1}}_{\ell\ell' L}\ \Theta_{\ell m}\delta^{W^{N-1}}_{\ell' m'}
\end{aligned} \tag{2.65}
$$

such that $\Pi^\alpha(\chi)$ and $W^\alpha(\chi)$ are overlapped. One intuitive choice is to take $W^\alpha(\chi) \equiv \Pi^\alpha(\chi)$, which could be a possibility if one has 3 dimensional measurements of the large scale structure that can be separated into custom redshift bins [2]. The rotation matrix is defined similarly as in Eq.(2.43), where the only difference comes from adequately changing the window functions involved in each component to match the one used in the estimators:

$$
(\mathbf{R}_L)^{XY} \equiv \frac{\sum_{\ell\ell'} G^{XW^X}_{\ell\ell' L} f^{YW^Y}_{\ell L\ell'}}{\sum_{\ell\ell'} G^{XW^X}_{\ell\ell' L} f^{XW^X}_{\ell L\ell'}} \tag{2.66}
$$

where W^X is the density window function associated to the observable X. The 2-point function introduced in Case 2 can be easily generalized to include the varying density window functions:

$$
\begin{aligned}
&\left\langle \hat{\mathbf{d}}^X_{LM} \hat{\mathbf{d}}^{Y*}_{L'M'} \right\rangle \\
&= A_L^{XW^X} A_{L'}^{YW^Y} \sum_{\ell_a m_a;\ell_b m_b} \sum_{\ell_c m_c;\ell_d m_d} (-1)^{M+M'} \begin{pmatrix} \ell_a & \ell_b & L \\ m_a & m_b & -M \end{pmatrix} \begin{pmatrix} \ell_c & \ell_d & L' \\ m_c & m_d & -M' \end{pmatrix} \\
&\times\ G^{XW^X}_{\ell_a\ell_b L}\ G^{YW^Y}_{\ell_c\ell_d L'} \left\langle \Theta_{\ell_a m_a}\delta^{W^X}_{\ell_b m_b} \Theta^*_{\ell_c m_c}\delta^{W^{Y}*}_{\ell_d m_d} \right\rangle \\
&= A_L^{XW^X} A_{L'}^{YW^Y} \sum_{\ell_a m_a;\ell_b m_b} \sum_{\ell_c m_c;\ell_d m_d} (-1)^{M+M'} \begin{pmatrix} \ell_a & \ell_b & L \\ m_a & m_b & -M \end{pmatrix} \begin{pmatrix} \ell_c & \ell_d & L' \\ m_c & m_d & -M' \end{pmatrix} \\
&\times\ G^{XW^X}_{\ell_a\ell_b L}\ G^{YW^Y}_{\ell_c\ell_d L'} \Bigg[\left\langle I_{\ell_a m_a}\delta^{W^X}_{\ell_b m_b} I^*_{\ell_c m_c}\delta^{W^{Y}*}_{\ell_d m_d} \right\rangle + \sum_{\ell_1 m_1} \sum_{\ell_2 m_2} \sum_{\ell'_1 m'_1} \sum_{\ell'_2 m'_2} \\
&\times\ (-1)^{m_a+m_c} W^{\ell_1,\ell_2,\ell_a}_{m_1,m_2,-m_a} W^{\ell'_1,\ell'_2,\ell_c}_{m'_1,m'_2,-m_c} \left\langle (\mathbf{X}_{\ell_1 m_1} \cdot \mathbf{Y}_{\ell_2 m_2})\, \delta^{W^X}_{\ell_b m_b} \left(\mathbf{X}^*_{\ell'_1 m'_1} \cdot \mathbf{Y}^*_{\ell'_2 m'_2} \right) \delta^{W^{Y}*}_{\ell_d m_d} \right\rangle \Bigg]
\end{aligned} \tag{2.67}
$$

[2] Surveys with big redshift errors can make this separation more difficult

The advantage of using multiple localized density window functions over fewer, wider ones, is clear: more 3-dimensional information of the lightcone fields is retained. However, realistic large-scale structure surveys might be limited to the second scenario. In this work we will explore both strategies for density windows.

Principal component analysis

Consider a set of estimators for the Π-binned LC moments of a field $M(\hat{\mathbf{n}}, \chi)$, constructed using the methods described above. The variance of the estimator is given by:

$$\left\langle \hat{\mathbf{d}}_{LM} \hat{\mathbf{d}}^*_{LM} \right\rangle = \mathbf{R}_L \mathbf{C}_L^{MM} (\mathbf{R}_L)^\dagger + \mathbf{N} \tag{2.68}$$

where $\mathbf{N}$ is the sum of all sources of noise. Although the Π basis is useful when it comes to separation of scales and localization on the lightcone, it can be less useful when it comes to separate the independent information contained in the 2-point function of the estimator. Using a principal component analysis, we can find the uncorrelated linear combinations of bins that yield the highest signal to noise. We use the following:

- Transform to a basis in which the noise matrix is diagonal.
- Perform a second transformation to a basis in which the noise matrix is the Identity.
- Perform a third transformation to a basis in which the signal matrix is diagonal. The noise matrix, due to being equal to the identity, is unchanged by the third transformation. The resulting signal matrix $\mathbf{C}_L^{pp}$ is diagonal and contains the signal to noise for the different uncorrelated principal components.

The linear combinations of bins associated to the principal components can be found using the transformation matrices $\mathbf{T_1}, \mathbf{T_2}, \mathbf{T_3}$ and the rotation matrix $\mathbf{R}_L$:

$$\mathbf{X^P} = \mathbf{T_3} \cdot \mathbf{T_2} \cdot \mathbf{T_1} \cdot \mathbf{R_L} \cdot \mathbf{X} \tag{2.69}$$

where $\mathbf{X} = (M^0_{LM}, \ldots, M^N_{LM})$. The j-th principal components then is characterized by a set of N coefficients $c_L^{j\beta}$ such that:

$$(\mathbf{X^P})^j_{LM} = \sum_\beta c_L^{j\beta} M^\beta_{LM} \tag{2.70}$$

The signal to noise per mode for the j-th principal component is simply given by the diagonal element $(\mathbf{C}_L^{pp})^{jj}$. We define a signal to noise per harmonic LM mode as

$$SN_{LM} = \sum_j (\mathbf{C}^{pp})^{jj}_L. \tag{2.71}$$

We define the total signal to noise as a sum over all principal components and harmonic modes:

$$SN_{tot} = \sum_j \sum_L (2L+1)(\mathbf{C}^{pp})^{jj}_L \tag{2.72}$$

Multiplicative bias from theory modelling

In order to construct the quadratic estimators, one has to assume a model for the couplings Eq.(2.27), which depend on $C_{\ell'}^{B^\alpha \delta^W}(\chi)$. If an incorrect $\tilde{C}_{\ell'}^{B^\alpha \delta^W}$ is used instead of the true physical $C_{\ell'}^{B^\alpha \delta^W}$, a multiplicative bias will be introduced:

$$\left\langle \hat{\mathbf{d}}^X_{LM} \right\rangle = \sum_Y \Gamma^{\mathrm{XY}}_{\mathrm{L}} \left(\tilde{\mathbf{R}}_L \right)^{XY} \mathbf{d}^Y_{LM} + \tilde{\Delta} \tag{2.73}$$

where $\tilde{\mathbf{R}}_L$ is the rotation Eq.(2.66), $\tilde{\Delta}$ is the reconstruction noise, and $\Gamma^{\mathrm{XY}}_{\mathrm{L}}$ is the multiplicative bias:

$$\Gamma^{\mathrm{XY}}_{\mathrm{L}} = \frac{\sum_{\ell\ell'} \tilde{G}^{XW^X}_{\ell\ell' L} f^{YW^Y}_{\ell L \ell'}}{\sum_{\ell\ell'} \tilde{G}^{XW^X}_{\ell\ell' L} \tilde{f}^{YW^Y}_{\ell L \ell'}} \tag{2.74}$$

where $f^{YW^Y}_{\ell L \ell'}$ is the true physical coupling. The bias Eq. 2.74 is not symmetric in the indices X, Y so in principle there are N^2 bias parameters at each scale L. In the context of kSZ velocity reconstruction, where the B-field is the differential optical depth, this multiplicative factor is commonly referred as the optical depth bias; see e.g. [75, 20, 21, 66]. As discussed in more detail in Sec. 2.4.1, the bias does not depend on the scale L over the relevant range for reconstruction, leaving a total of N^2 bias parameters to account for. Note that in the absence of off-diagonal terms in the rotation matrix Eq. 2.66 (or if these terms are very small), there would only be N bias parameters. This is the assumption that has been made in previous literature utilizing the Lightcone Picture to forecast cosmological constraints, e.g. Refs. [76, 77] (and Chapter 4 of this thesis). We comment on this assumption and the general problem of mitigating the optical depth bias in Sec. 2.4.1.

2.3 Modeling of Observables

In this work, we will be interested in statistically anisotropic correlations between various contributions to the observed CMB and a tracer of large scale structure (LSS). Our goal is to use such statistical anisotropies to reconstruct (on large angular scales) a set of modulating fields – here, our focus is on the radial velocity field. Our prototype tracer is a photometric galaxy survey, as considered in e.g. Refs. [18, 19, 20, 78]; other tracers such as spectroscopic surveys [20, 78], the Cosmic Infrared Background (CIB) [79], or line-intensity maps [13, 80] are other interesting candidates. Throughout the this chapter, we assume a fiducial cosmological model consistent with Planck 2018 [4]; in particular, we set: $\{10^9 A_s = 2.2, n_s = 0.965, \Omega_m = 0.31, \Omega_b = 0.049, H_0 = 68 \text{ km s}^{-1} \text{ Mpc}^{-1}, \tau = 0.06\}$. There is no strong dependence on cosmological parameters for any of our conclusions.

Where necessary relations are presented in the Newtonian gauge, where at late-times when we can neglect anisotropic stress, the metric is:

$$ds^2 = a(\eta)^2 \left(-\left[1 + 2\Psi(\eta, \mathbf{x})\right] d\eta^2 + \left[1 - 2\Psi(\eta, \mathbf{x})\right] d\mathbf{x}^2 \right) \tag{2.75}$$

Because our halo model code calculates perturbations in the synchronous gauge defined by the cold dark matter frame, we need to relate Newtonian gauge quantities to synchronous gauge ones. For late-times, simple relations can be written for the Newtonian gauge gravitational potential Ψ and the peculiar velocity field $\mathbf{v}$ of dark matter in terms of the synchronous gauge dark matter perturbations in Fourier space:

$$\Psi(\eta, \mathbf{k}) = -\frac{3\Omega_m H_0^2}{a^2(\eta) k^2} \delta_m^{(sync)}(\eta, \mathbf{k}) \tag{2.76}$$

$$\mathbf{v}(\eta, \mathbf{k}) = i \frac{\mathbf{k}}{k^2} f(\eta) H(\eta) a(\eta) \delta_m^{(sync)}(\eta, \mathbf{k}) \tag{2.77}$$

where H is the Hubble rate and f is the growth rate, defined as $\frac{a}{D}\frac{dD}{da}$, with $D(a)$ the linear theory growth factor of dark matter perturbations. Given that perturbations of the galaxy and electron fields are only needed on small scales for the reconstruction procedure, we approximate $\delta_g^{(Newt)} \approx \delta_g^{(sync)}$ and $\delta_e^{(Newt)} \approx \delta_e^{(sync)}$.

2.3.1 Constructing observables

All of the observables presented below are constructed from a set of fundamental cosmological fields, which we compute using linear cosmological perturbation theory and the halo

model for large-scale structure. Combining Eqs.(2.7) and (2.9), a field on the lightcone can be characterized by specifying a window function $W(\chi)$, an integral kernel $\mathcal{K}_\ell^F(\chi, k)$ and an underlying perturbation field in Fourier space $\tilde{\mathrm{U}}^F(\eta(\chi), \mathbf{k})$:

$$\mathrm{F}_{\ell m}^W \equiv \int d\chi\ W(\chi) \int \frac{d^3\mathbf{k}}{(2\pi)^3}\, \mathcal{K}_\ell(\chi, k)\, \tilde{\mathrm{U}}^F(\eta(\chi), \mathbf{k})\, Y_{\ell m}^*(\hat{k}). \tag{2.78}$$

Below, we refer to these function as the "building functions" and we specify them for each observable we construct.

2.3.2 The CMB

The de-beamed CMB temperature measured in a frequency band ν through an instrument with an isotropic beam $(B_\ell^\Theta)^\nu$ and noise $n_{\ell m}^\nu$ has contributions from a variety of sources. As a baseline model, we take:

$$\Theta_{\ell m}^\nu = \Theta_{\ell m}^{pCMB} + \Theta_{\ell m}^{ISW,lin} + \Theta_{\ell m}^{ISW,nlin} + \Theta_{\ell m}^{kSZ} + \Theta_{\ell m}^{ReI} + \Theta_{\ell m}^{L} + (\Theta_{\ell m}^{XG})^\nu + (\Theta_{\ell m}^{G})^\nu + n_{\ell m}^\nu/(B_\ell^\Theta)^\nu \tag{2.79}$$

There are blackbody contributions including: $\Theta_{\ell m}^{pCMB}$ which contains the SW, Doppler, and early ISW contributions to the primary CMB, $\Theta_{\ell m}^{ISW,lin}$ the linear contribution to the late-time ISW component, $\Theta_{\ell m}^{ISW,nlin}$ the non-linear contribution to the late-time ISW component (also referred to as the moving lens effect, Rees-Sciama effect), $\Theta_{\ell m}^{kSZ}$ the late-time kSZ component, $\Theta_{\ell m}^{ReI}$ the reionization kSZ, and $\Theta_{\ell m}^{L}$ the lensing contribution to the primary CMB. There are frequency-dependent extragalactic components $(\Theta_{\ell m}^{XG})^\nu$, whose dominant components for the experimental configurations considered below include the CIB and the thermal Sunyaev-Zel'dovich effect (tSZ). Finally, there is a frequency-dependent galactic component $(\Theta_{\ell m}^{G})^\nu$.

Below we describe in detail the components which have significant cross-correlation with late-time tracers of LSS, since such components must be computed in a self-consistent way. The SW, Doppler, and early ISW contributions $\Theta_{\ell m}^{pCMB}$ to the primary CMB do not contribute to the cross-correlation with tracers of LSS; we compute their power spectra using CAMB [81]. The reionization kSZ component $\Theta_{\ell m}^{ReI}$ is modelled as a Gaussian field with power spectrum $\ell^2 C_\ell^{rei}/2\pi = 1\mu\mathrm{K}^2$. If higher redshift tracers of LSS are considered, then the reionization kSZ can be used to reconstruct the radial velocity field as described in Ref. [82]; in this work, we focus on tracers of LSS that do not have any significant cross-correlation with reionization kSZ. For the fiducial CMB experiments considered below, including reionization kSZ does not affect any of our results below, and we therefore neglect this contribution in our analysis. Galactic foregrounds on the small angular scales

Freq. [GHz]	FWHM [arcmin]	$N_{\rm white}$ [μK-arcmin]	$N_{\rm red}$ [μK^2 s]
27	7.4	71	100
39	5.1	36	39
93	2.2	8	230
145	1.4	10	1,500
225	1.0	22	17,000
280	0.9	54	31,000

Table 2.1: CMB experimental noise parameters consistent with Simons Observatory Large Area Telescope.

relevant to the reconstruction considered below are generally sub-dominant to extragalactic foregrounds on a line of sight away from the galactic plane (see e.g. Ref. [83]). We assume that regions with significant contamination can be masked. Other than considering the effect of a mask, we therefore neglect galactic foregrounds. We model instrumental noise $n^{\nu}_{\ell m}$ as a frequency-dependent constant and the beam $(B^{\Theta}_{\ell})^{\nu}$ as a Gaussian with a frequency-dependent Full Width at Half-Maximum (FWHM) $\theta^{\nu}_{\rm FWHM}$. Our fiducial CMB experiment is consistent with the properties of the Simons Observatory Large Area Telescope [15], with:

$$n^{\nu}_{\ell m} = N_{\rm red}\left(\frac{\ell}{1000}\right)^{-3.5} + N_{\rm white} \tag{2.80}$$

where $N_{\rm red}$ describes the level of $1/f$ atmospheric noise and $N_{\rm white}$ describes the sensitivity of the frequency band. The frequencies, beam, and noise levels are collected in Table 2.1. We assume an observation time of 5 years when computing the level of $1/f$ noise, and choose the "baseline" values for noise found in [15].

In Fig. 2.2, we summarize the blackbody components of our CMB model: the primary CMB, late-time ISW (linear and non linear), lensed CMB, and kSZ. At low-ℓ, the dominant components are the primary CMB and late-time linear ISW effects. Crucially, at high-ℓ ($\ell \gtrsim 4000$), kSZ is the dominant blackbody component of the CMB. Each of these contributions is discussed in more detail in the following sub-sections. In the left panel of Fig. 2.3, we show the frequency-dependent components of our CMB model, including extragalactic foregrounds and instrumental noise – as mentioned above, we do not consider a galactic component because it will not be an important contaminant for velocity reconstruction. In the right panel of Fig. 2.3, we compare the effective noise obtained by using multi-frequency information with the blackbody component of the CMB and with the noise and foregrounds in the "cleanest" channel of our fiducial experiment (for veloc-

ity reconstruction), at 145 GHz. Note that the blackbody CMB dominates the noise and foregrounds below $\ell \gtrsim 3000$. We assume that multi-frequency information can be used to clean foregrounds using a standard harmonic space internal linear combination (ILC) procedure, described in more detail below. Such a procedure can reduce the level of noise and foregrounds by roughly a factor of 2 at high-ℓ when compared to the 145 GHz channel. Unless otherwise specified, in the analyses to follow we will use the ILC-cleaned CMB generated with the specifications in Table 2.1; we consider a maximum value of $\ell_{\rm max} = 6000$ which roughly corresponds to maps with a Healpix resolution NSIDE of 2048.

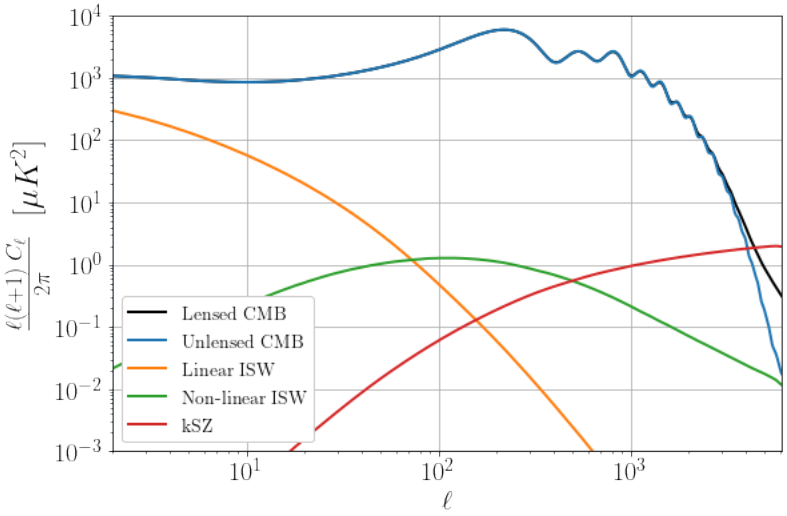

Figure 2.2: Contributions to blackbody CMB.

We now describe in more detail how we model the various CMB components listed above.

Late-time ISW (linear)

Gravitational potentials that evolve in time induce a temperature anisotropy known as the integrated Sachs-Wolfe (ISW) effect. The late-time ISW contribution to the CMB is given by

$$\Theta_{\ell m}^{ISW} = -2 \int_0^{\chi_{\rm max}} d\chi \, \frac{\partial \Psi_{\ell m}}{\partial \chi}(\chi) \tag{2.81}$$

where $\chi_{\rm max}$ is a fiducial maximum range in comoving distance large enough to capture the majority of the late-time decay of the potential due to the presence of a cosmological

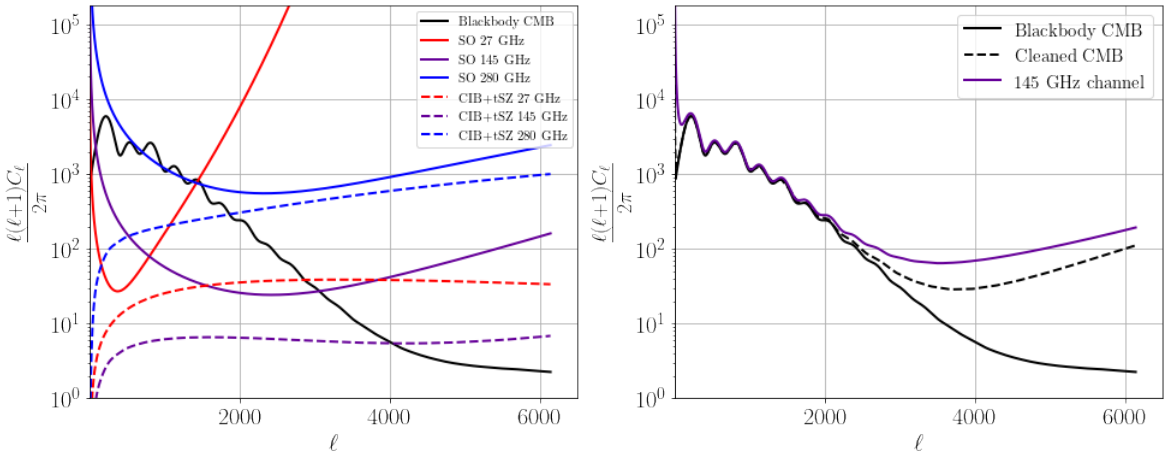

Figure 2.3: **Left panel:** Frequency dependent components of CMB compared to the blackbody component. Colored solid lines correspond to the de-beamed instrumental noise and dashed lines correspond to the CIB+tSZ contributions (including their cross-spectra). **Right panel:** the ILC-cleaned power spectrum compared to the blackbody component and the full 145 Ghz channel.

constant. The building functions (see Eq. 2.78) for the linear late-time ISW effect are:

$$W^{ISW,lin}(\chi) = \begin{cases} 1, & 0 \leq \chi < \chi_{\rm max}, \\ 0, & \text{otherwise} \end{cases} \tag{2.82}$$

$$\mathcal{K}_\ell^{ISW,lin}(\chi,k) = 4\pi i^\ell \frac{j_\ell(k\chi)}{k^2} \frac{3\Omega_m H_0^2}{a^2(\chi)} \left(-\frac{da}{d\chi}(\chi) + a(\chi)\frac{d}{d\chi} \right) \tag{2.83}$$

$$\tilde{\rm U}^{ISW,lin}(\eta(\chi),\mathbf{k}) = \delta_m^{(lin)}(\eta(\chi),\mathbf{k}) \tag{2.84}$$

where $\delta_m^{(lin)}$ are the Fourier modes of the linear dark matter perturbations. The power spectrum of the linear late-time ISW is calculated using Eq.(2.16).

Non-linear ISW

At the non-linear order, the ISW effect can be sourced on small-scales by the long-wavelength peculiar velocities of gravitational potentials. Taking the limit $k' \ll k$, the non-linear evolution of the gravitational potential due to a long wavelength velocity modes can be approximated as:

$$\dot{\Psi}_{\rm NL}(\boldsymbol{k}) \sim i\boldsymbol{k}\Psi(\boldsymbol{k}) \cdot \int \frac{{\rm d}^3k'}{(2\pi)^3}\mathbf{v}(\boldsymbol{k'}) \tag{2.85}$$

which translates in real space to :

$$\dot{\Psi}_{\rm NL}(\eta, \mathbf{x}) \sim \nabla\Psi_{\rm NL}(\eta, \mathbf{x}) \cdot \mathbf{v}(\eta, \mathbf{x}) \tag{2.86}$$

Simulations of the non-linear ISW effect indicate that the line of sight integral will be mostly sourced by motions transverse to the line of sight (see Sec. 4.1 of [41] for more details). With this approximation, the non-linear ISW effect contribution is given by

$$\Theta^{\rm ISW,nlin}(\hat{\mathbf{n}}) \approx 2\int_0^{\chi_{ls}} d\chi\ \frac{(\nabla_\perp\Psi)}{\chi}(\hat{\mathbf{n}}, \chi) \cdot \mathbf{v}_\perp(\hat{\mathbf{n}}, \chi), \tag{2.87}$$

where $\mathbf{v}_\perp$ is the peculiar (comoving) transverse velocity and $\nabla_\perp$ is the gradient on 2-sphere. In this work, we assume that the large-scale velocity is pure-gradient, and therefore the transverse velocity component can be expressed as $\mathbf{v}_\perp = \nabla_\perp\Upsilon(\hat{\mathbf{n}}, \chi)$. We refer to Υ as the *transverse velocity potential*. In spherical harmonics, the effect on the CMB temperature takes the form

$$\Theta^{\rm ISW,nlin}_{\ell,m} = \int_0^{\chi_{ls}} d\chi \sum_{\ell\ell'\,mm'} \Upsilon_{\ell'm'}(\chi)\psi_{\ell''m''}(\chi) \int {\rm d}^2\hat{n}\, Y^*_{\ell m}\nabla_i Y_{\ell'm'}\nabla^i Y_{\ell''m''},, \tag{2.88}$$

where

$$\psi_{\ell m}(\chi) \equiv 2\frac{\Psi_{\ell m}(\chi)}{\chi}\,, \tag{2.89}$$

and we will refer to this quantity as the *moving lens potential* (in reference to the manifestation of the non-linear ISW effect through the moving lens effect). We can expand the signal in terms of the Haar-binned LC moments of Υ and ψ:

$$\Theta^{\rm ISW,nlin}_{\ell,m} = \sum_{s=0}^{\infty}\sum_{\ell\ell'\,mm'} \Upsilon^s_{\ell'm'}\psi^s_{\ell''m''} \int {\rm d}^2\hat{n}\, Y^*_{\ell m}\nabla_i Y_{\ell'm'}\nabla^i Y_{\ell''m''}. \tag{2.90}$$

For the transverse velocity potential we use the following building functions:

$$W^{\Upsilon^s}(\chi) = h^s(\chi) \quad \text{from Eq.(2.11)} \tag{2.91}$$

$$\mathcal{K}^{\Upsilon^s}_\ell(\chi, k) = 4\pi i^\ell \frac{j_\ell(k\chi)}{k^2}\frac{f(\chi)H(\chi)a(\chi)}{\chi} \tag{2.92}$$

$$\tilde{\rm U}^{\Upsilon^s}(\eta(\chi), \mathbf{k}) = \delta_m(\eta(\chi), \mathbf{k}) \tag{2.93}$$

For the moving lens potential we use the following building functions:

$$W^{\psi^s}(\chi) = h^s(\chi) \quad \text{from Eq.(2.11)} \tag{2.94}$$

$$\mathcal{K}_\ell^{\psi^s}(\chi,k) = -4\pi i^\ell \frac{j_\ell(k\chi)}{k^2}\frac{3\Omega_m H_0^2}{a(\chi)\chi} \tag{2.95}$$

$$\tilde{\mathrm{U}}^{\psi^s}(\eta(\chi),\mathbf{k}) = \delta_m(\eta(\chi),\mathbf{k}) \tag{2.96}$$

The angular integral in Eq.(2.90) is given by:

$$\begin{aligned}\int \mathrm{d}^2\hat{n}\, Y^*_{\ell m}\nabla_i Y_{\ell' m'}\nabla^i Y_{\ell'' m''} &= \frac{1}{2}[\ell(\ell+1)+\ell'(\ell'+1)-\ell''(\ell''+1)] \\ &\times \sqrt{\frac{(2\ell+1)(2\ell'+1)(2\ell''+1)}{4\pi}} \\ &\times \begin{pmatrix}\ell & \ell' & \ell''\\ m & m' & m''\end{pmatrix}\begin{pmatrix}\ell & \ell' & \ell''\\ 0 & 0 & 0\end{pmatrix}\end{aligned} \tag{2.97}$$

The non-linear ISW power spectrum can be calculated in terms of the auto and cross-spectra of the Haar-binned moments of Υ and ϕ.

$$\begin{aligned}C_\ell^{\mathrm{ISW,nlin}} &= \sum_{\ell',\ell''}\frac{(2\ell'+1)(2\ell''+1)}{4\pi}\frac{1}{4}[\ell'(\ell'+1)+\ell''(\ell''+1)-\ell(\ell+1)]^2 \\ &\times \begin{pmatrix}\ell & \ell' & \ell''\\ 0 & 0 & 0\end{pmatrix}^2 \sum_{s,s'=0}^{\infty}\left[(C^{\Upsilon\Upsilon})^{ss'}_{\ell'}(C^{\psi\psi})^{ss'}_{\ell''}+(C^{\Upsilon\psi})^{ss'}_{\ell'}(C^{\Upsilon\psi})^{ss'}_{\ell''}\right]\end{aligned} \tag{2.98}$$

where $(C^{\Upsilon\Upsilon})^{ss'}_{\ell'}$, $(C^{\psi\psi})^{ss'}_{\ell''}$, $(C^{\Upsilon\psi})^{ss'}_{\ell'}$, and $(C^{\Upsilon\psi})^{ss'}_{\ell''}$ are calculated using the building functions and Eq.(2.16). In the squeezed limit where $\ell' \ll \ell''$, we can set $\ell \simeq \ell''$ and neglect cross-correlations between the bins s and s' as well as the cross-power between Υ and ψ:

$$C_\ell^{\mathrm{ISW,nlin}} = \frac{1}{4}\sum_{s=1}^{\infty}\left[\sum_{\ell'}\frac{(2\ell'+1)}{4\pi}[\ell'(\ell'+1)]^2(C^{\Upsilon\Upsilon})^{ss}_{\ell'}\right](C^{\psi\psi})^{ss}_{\ell}. \tag{2.99}$$

Defining a potential $\theta_\perp \equiv \frac{1}{2}\nabla^2_\perp\Upsilon$, and taking the continuum limit of the sum, we obtain:

$$C_\ell^{\mathrm{ISW,nlin}} = \int d\chi\, \langle\theta_\perp(0,\chi)^2\rangle P^{\psi\psi}(\frac{\ell}{\chi},\chi). \tag{2.100}$$

We demonstrate the convergence of the non-linear ISW signal with number of bins in Fig. 2.4.

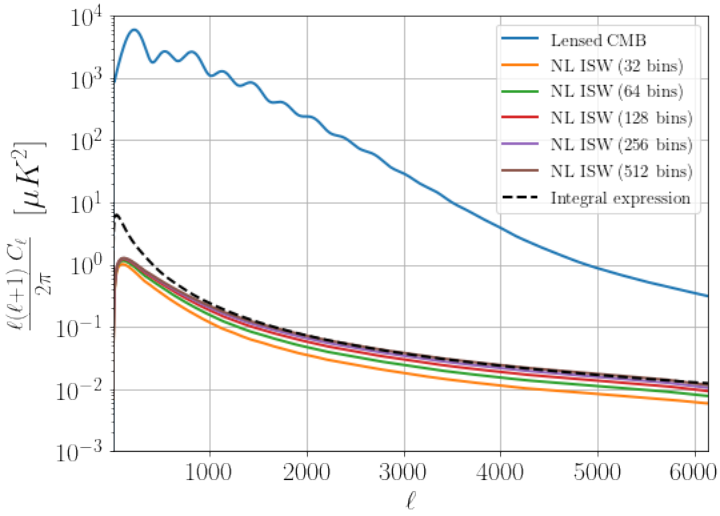

Figure 2.4: Convergence of non-linear ISW power with number of bins.

The kSZ effect

The contribution to the CMB temperature from the late-time kSZ effect is:

$$\Theta^{kSZ}(\hat{\mathbf{n}}) = -\int_0^{\chi_{\max}} d\chi \; v_{\rm eff}(\hat{\mathbf{n}}, \chi)\dot{\tau}(\hat{\mathbf{n}}, \chi) \tag{2.101}$$

where $\dot{\tau}(\hat{\mathbf{n}}, \chi)$ is the differential optical depth and $v_{\rm eff}(\hat{\mathbf{n}}, \chi)$ is the remote dipole field. The differential optical depth is

$$\dot{\tau}(\hat{\mathbf{n}}, \chi) = \sigma_T a(\chi)\bar{n}_e(\chi)(1 + \delta_e(\hat{\mathbf{n}}, \chi)) \tag{2.102}$$

where σ_T is the Thompson cross section, $a(\chi)$ the scale factor, $\bar{n}_e(\chi)$ the average electron density, and $\delta_e(\hat{\mathbf{n}}, \chi)$ the electron overdensity field. The remote dipole field is defined as

$$v_{\rm eff}(\hat{\mathbf{n}}, \chi) = \frac{3}{4\pi}\int d^2\hat{n}_e \; \Theta(\mathbf{x}, \hat{\mathbf{n_e}}) \; (\hat{\mathbf{n}} \cdot \hat{\mathbf{n_e}}). \tag{2.103}$$

where $\Theta(\mathbf{x}, \hat{\mathbf{n_e}})$ is the remote CMB at the spacetime position $(\eta(\chi), \mathbf{x} = \chi\hat{\mathbf{n}})$ along our past lightcone. The dominant contribution to the remote dipole comes from the radial peculiar velocity of electrons, and corrections to the dipole relevant on large scales come from the stationary dipole components discussed in Sec.1.5. In general, the correction coming from the stationary components can be safely neglected, only becoming significant enough when

inspecting the correlations of the remote dipole field on large enough scales. For simplicity, we will only consider the dominant kinetic term in this chapter, and use the terminology radial velocity in lieu of remote dipole.

We will focus on the late-time kSZ effect here, where the limits of integration extend from the origin out to a radial comoving distance $\chi_{\rm max}$ after reionization ended. We assume a fiducial value of $\chi_{\rm max} = 8.4$ Gpc, which corresponds to a redshift $z_{\rm max} = 6$ in our fiducial cosmology. Computing the multipoles of the kSZ temperature anisotropies Eq. 2.101 in terms of the Haar-binned LC moments, we have:

$$\begin{aligned}\Theta_{\ell m}^{kSZ} &= \sum_{\ell_1 m_1;\ell_2 m_2} (-1)^m \sqrt{\frac{(2\ell+1)(2\ell_1+1)(2\ell_2+1)}{4\pi}} \\ &\times \begin{pmatrix}\ell_1 & \ell_2 & \ell \\ 0 & 0 & 0\end{pmatrix} \begin{pmatrix}\ell_1 & \ell_2 & \ell \\ m_1 & m_2 & -m\end{pmatrix} \sum_{s=1}^{\infty} v^s_{\ell_1 m_1} \dot{\tau}^s_{\ell_2 m_2} \end{aligned} \tag{2.104}$$

For the radial velocity moments we use the following building functions:

$$W^{v^s}(\chi) = h^s(\chi) \quad \text{from Eq.(2.11)} \tag{2.105}$$

$$\mathcal{K}_\ell^{v^s}(\chi,k) = 4\pi i^\ell \frac{f(\chi)H(\chi)a(\chi)}{(2\ell+1)k}[\ell j_{\ell-1}(k\chi) - (\ell+1)j_{\ell+1}] \tag{2.106}$$

$$\tilde{\rm U}^{v^s}(\eta(\chi),\mathbf{k}) = \delta_m(\eta(\chi),\mathbf{k}) \tag{2.107}$$

where $H(\chi)$ is the Hubble rate and $f(\chi)$ is the growth rate, defined as $\frac{a}{D}\frac{dD}{da}$, with $D(a(\chi))$ the linear theory growth factor of dark matter perturbations. For the differential optical depth we use the following building functions:

$$W^{\dot{\tau}^s}(\chi) = h^s(\chi) \quad \text{from Eq.(2.11)} \tag{2.108}$$

$$\mathcal{K}_\ell^{\dot{\tau}^s}(\chi,k) = 4\pi i^\ell j_\ell(k\chi)\sigma_T a(\chi)\bar{n}_e(\chi) \tag{2.109}$$

$$\tilde{\rm U}^{\dot{\tau}^s}(\eta(\chi),\mathbf{k}) = \delta_e(\eta(\chi),\mathbf{k}) \tag{2.110}$$

The kSZ temperature power spectrum is

$$C_\ell^{kSZ} = \sum_{\ell_1 m_1;\ell_2 m_2} \sum_{\ell_1' m_1';\ell_2' m_2'} (2\ell+1)\sqrt{\frac{(2\ell_1+1)(2\ell_2+1)}{4\pi}}\sqrt{\frac{(2\ell_1'+1)(2\ell_2'+1)}{4\pi}} \tag{2.111}$$

$$\times \begin{pmatrix}\ell_1 & \ell_2 & \ell \\ 0 & 0 & 0\end{pmatrix} \begin{pmatrix}\ell_1' & \ell_2' & \ell \\ 0 & 0 & 0\end{pmatrix} \begin{pmatrix}\ell_1 & \ell_2 & \ell \\ m_1 & m_2 & -m\end{pmatrix} \begin{pmatrix}\ell_1' & \ell_2' & \ell \\ m_1' & m_2' & -m\end{pmatrix}$$

$$\times \sum_{s,s'=1}^{\infty} \langle (v^s_{\ell_1 m_1})^* (\dot{\tau}^s_{\ell_2 m_2})^* v^{s'}_{\ell_1' m_1'} \dot{\tau}^{s'}_{\ell_2' m_2'} \rangle \tag{2.112}$$

Keeping the disconnected parts of the four-point function only, the power spectrum is:

$$C_\ell^{kSZ} = \sum_{\ell_1;\ell_2} \frac{(2\ell_1+1)(2\ell_2+1)}{4\pi} \begin{pmatrix} \ell_1 & \ell_2 & \ell \\ 0 & 0 & 0 \end{pmatrix}^2 \times \sum_{s,s'=1}^{\infty} \left[(C^{vv})_{\ell_1}^{ss'} (C^{\dot\tau\dot\tau})_{\ell_2}^{ss'} + (C^{v\dot\tau})_{\ell_1}^{ss'} (C^{v\dot\tau})_{\ell_2}^{ss'} \right] \tag{2.113}$$

where $(C^{vv})_{\ell_1}^{ss'}$, $(C^{\dot\tau\dot\tau})_{\ell_2}^{ss'}$, $(C^{v\dot\tau})_{\ell_1}^{ss'}$, and$(C^{v\dot\tau})_{\ell_2}^{ss'}$are calculated using the building functions and Eq.(2.16). Focusing on $\ell \gg 1$, the majority of the power will come from $\ell_1 \ll \ell$ where $\ell_2 \sim \ell$. In this regime, we also expect that there is little bin-bin correlation in the differential optical depth, so we can take $(C^{\dot\tau\dot\tau})_{\ell_2}^{ss'} \simeq (C^{\dot\tau\dot\tau})_{\ell_2}^{ss}\, \delta_{ss'}$. Finally, the first term in parentheses above will dominate the second on small angular scales. In this limit, the kSZ power can be approximated by:

$$C_\ell^{kSZ} \simeq \sum_{s=1}^{s_{\rm max}} \left[\sum_{\ell_1} \frac{(2\ell_1+1)}{4\pi} (C^{vv})_{\ell_1}^{ss} \right] (C^{\dot\tau\dot\tau})_\ell^{ss} \tag{2.114}$$

$$= \sum_{s=1}^{s_{\rm max}} \langle \bar v_s(0)^2 \rangle (C^{\dot\tau\dot\tau})_\ell^{ss} \tag{2.115}$$

Taking $s_{\rm max} \to \infty$, this is equivalent to the expression:

$$C_\ell^{kSZ} = \int d\chi\ \langle v(0,\chi)^2 \rangle\ \dot\tau(\chi)^2\ P^{ee}(\frac{\ell}{\chi}, \chi) \tag{2.116}$$

which is consistent with previous literature [84].

In Fig. 2.5 we show the coarse grained kSZ power spectrum for $\chi_{\rm max} = 8.4$ Gpc (corresponding to a redshift $z_{\rm max} = 6$ in our fiducial cosmology) with $k_{\rm max} = 32, 64, 128, 256$ and 512 bins (corresponding to $\Delta\chi = 263, 131, 66, 33, 17$ Mpc). We compare with the continuum expression Eq. 2.116. It can be seen from this figure that ~ 512 bins, corresponding of a coarse graining scale of ~ 17 Mpc is sufficient to capture the majority of the kSZ power. Based on this, below we take 512 bins to correspond to the continuum limit.

Lensing of the primary CMB

We approximate the lensing of the primary CMB by the first order term:

$$\Theta^L(\hat{\mathbf{n}}) = (\nabla_\perp \phi) \cdot (\nabla_\perp \Theta^{pCMB}) \tag{2.117}$$

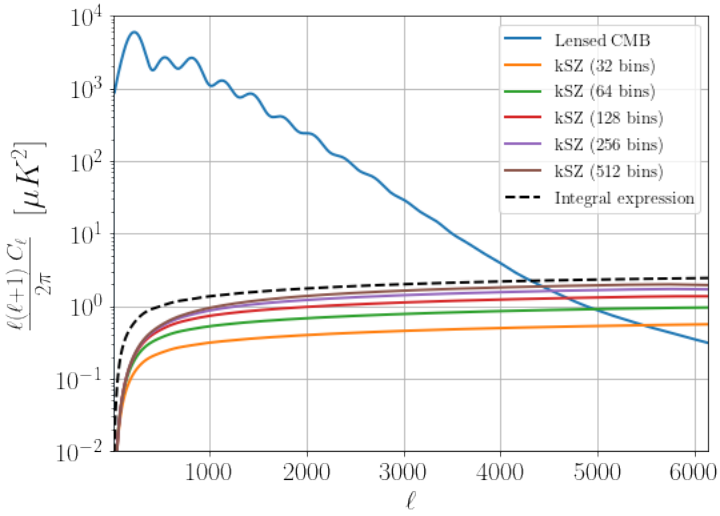

Figure 2.5: Convergence of kSZ power with number of bins.

where $\nabla_\perp$ is the angular gradient in the unit 2-sphere and $\phi(\hat{\mathbf{n}})$ is the lensing potential, defined as:

$$\phi(\hat{\mathbf{n}}) = -2\int_0^{\chi_{ls}} d\chi \frac{\chi_{ls}-\chi}{\chi_{ls}\chi}\Psi(\hat{\mathbf{n}},\chi). \tag{2.118}$$

In terms of the multipole moments of the primary CMB and the lensing potential, the lensed CMB contribution is written as:

$$\Theta^{\rm L}_{\ell m} = \sum_{\ell\ell' mm'} \phi_{\ell' m'}\Theta^{pCMB}_{\ell'' m''}\int {\rm d}^2\hat{n}\, Y^*_{\ell m}\nabla_i Y_{\ell' m'}\nabla^i Y_{\ell'' m''} \tag{2.119}$$

where the primary CMB is computed using CAMB and the lensing potential moments are computed using the building functions:

$$W^\phi(\chi) = \begin{cases} 1, & 0\le\chi<\chi_{ls}, \\ 0, & \text{otherwise} \end{cases} \tag{2.120}$$

$$\mathcal{K}^\phi_\ell(\chi,k) = 4\pi i^\ell \frac{j_\ell(k\chi)}{k^2}\frac{3\Omega_m H_0^2}{a(\chi)}\frac{\chi_{ls}-\chi}{\chi_{ls}\chi} \tag{2.121}$$

$$\tilde{\rm U}^\phi(\eta(\chi),\mathbf{k}) = \delta_m(\eta(\chi),\mathbf{k}) \tag{2.122}$$

The power spectrum calculation is similar to the one for the non-linear ISW effect, and

yields:

$$C_\ell^{\rm L} = \sum_{\ell',\ell''} \frac{(2\ell'+1)(2\ell''+1)}{4\pi}\frac{1}{4}[\ell'(\ell'+1)+\ell''(\ell''+1)-\ell(\ell+1)]^2 \times \begin{pmatrix} \ell & \ell' & \ell'' \\ 0 & 0 & 0 \end{pmatrix}^2 \left[(C^{\phi\phi})_{\ell'}(C^{\Theta\Theta})_{\ell''} + (C^{\phi\Theta})_{\ell'}(C^{\phi\Theta})_{\ell''}\right] \quad (2.123)$$

Extragalactic foregrounds

There are a number of extragalactic foregrounds that contribute to the CMB, whose relative importance depends on the frequency and scale being observed. At low frequencies ($\lesssim$ 150 GHz) on arcminute scales , the thermal Sunyaev Zel'dovich (tSZ) effect and radio point sources dominate. At high frequencies ($\gtrsim$ 150 GHz) on the same scales, the CIB is the dominant extragalactic foreground. Below, we assume that enough radio point sources can be masked to make tSZ the dominant source at low frequencies. With this assumption, we include the tSZ and CIB only in our extragalactic foreground model.

We model the CIB and tSZ using the Halo Model for large scale structure, combining elements of the models described in Refs. [20, 79, 85, 86, 87, 88]. Our assumptions are outlined in detail in Appendix B.4. In Fig. 2.3, we show angular power spectra of the CIB and tSZ at several frequencies for our fiducial CMB experiment. Since all observables are computed within the same halo model, it is possible to capture the correlations between the CIB, tSZ, and galaxy number counts – e.g. the spectra in Fig. 2.3 include the CIB-tSZ cross-power. We discuss the detailed properties of the galaxy-foreground cross-spectra below in Sec. 2.3.6.

2.3.3 Foreground Cleaning of the CMB

To access the blackbody components of the CMB necessary for velocity reconstruction, we can estimate how well one can use the multi-frequency information in the CMB to clean the extragalactic foregrounds in our model. Here, we use the harmonic Internal Linear Combination (ILC) algorithm [89].

We write the covariance between the de-beamed CMB at different frequencies as a matrix:

$$\mathbf{C}_\ell = C_\ell^{TT}\mathbf{e}\mathbf{e}^\dagger + \mathbf{C}_\ell^{XG} + (\mathbf{B^{-1}N})_\ell \quad (2.124)$$

where C_ℓ^{TT} contains the blackbody components of the CMB (primary CMB, kSZ, ISW, etc.), $\mathbf{e} = \{1,1,1,\ldots\}$, $\mathbf{C}_\ell^{XG}$ contains the CIB, tSZ, and the cross-correlations between

these various components at the measured frequencies, and $(\mathbf{B^{-1}N})_\ell$ is the de-beamed instrumental noise covariance (assumed diagonal). Following the ILC method in harmonic space [89], we estimate the blackbody component as:

$$\hat{\Theta}_{\ell m} = \mathbf{w}_\ell^\dagger \, \mathbf{\Theta}_{\ell m} \tag{2.125}$$

where the weights $\mathbf{w}_\ell$ that minimize the variance of the resulting multipole moments $\hat{\Theta}_{\ell m}$ are given by:

$$\mathbf{w}_\ell = \frac{(\mathbf{C}_\ell)^{-1}\mathbf{e}}{\mathbf{e}^\dagger(\mathbf{C}_\ell)^{-1}\mathbf{e}} \tag{2.126}$$

The ensemble averaged power spectrum of the cleaned map is:

$$C_\ell^{\Theta\Theta;\text{clean}} = C_\ell^{TT} + \mathbf{w}_\ell^\dagger \left(\mathbf{C}_\ell^{XG} + (\mathbf{B^{-1}N})_\ell\right) \mathbf{w}_\ell \tag{2.127}$$

To the extent that the second term is small, we have successfully isolated the blackbody component of the CMB in the resulting map. Note that the residuals represented by the second term include both foreground residuals as well as an effective noise for the linear combination of maps. In the right panel of Fig. 2.3, we show C_ℓ^{clean} for our fiducial CMB experiment. For the fiducial experimental parameters we choose, from the left panel of Fig. 2.3, we see that the experimental noise is somewhat larger than the extragalactic foregrounds. Therefore, much of the improvement of the cleaned CMB over the 145 GHz channel comes from a lower effective noise rather than the removal of extragalactic foregrounds.

We can also estimate the cleaned galaxy-Temperature cross-power:

$$C_\ell^{\Theta\delta^{W^\alpha};clean} = C_\ell^{ISW,lin\ \delta^{W^\alpha}} + \mathbf{w}_\ell^\dagger \mathbf{C}_\ell^{XG\ \delta^{W^\alpha}} \tag{2.128}$$

Here, because the CMB noise is uncorrelated with the galaxy field, there is no effective noise term. The ILC algorithm in this case reduces the variance of the cross-power due to the removal of extragalactic foregrounds.

2.3.4 Galaxy number counts

We now consider a tracer of the electron overdensity field, which for the purposes of the present work we take to be the galaxy overdensity field, measured using a photometric redshift survey. Other tracers such as the redshifted 21cm Hydrogen line (or transitions such as CII) measured by line intensity mapping surveys [13, 80], the CIB [79], or the

dispersion measure of Fast Radio Bursts [66] have been considered as well. Spectroscopic surveys were considered in Refs. [78, 20, 21], which may be more computationally feasible to analyze in the Box Picture; we defer a discussion of spectroscopic surveys in the Lightcone Picture to future work.

For the purposes of velocity reconstruction, the three dimensional information in a galaxy redshift survey is used to construct a series of 2-dimensional fields that are later cross-correlated with CMB temperature data. In harmonic space, these 2-dimensional fields can be expressed as integrals over redshift space:

$$g_{\ell m}^{W^\alpha} = \int dz_o W^\alpha(z_o) g_{\ell m}(z_o) \tag{2.129}$$

where z_o denotes the observed redshift for the galaxies in the survey, $g_{\ell m}(z_o)$ are spherical harmonic coefficients of the measured 3-dimensional galaxy overdensity field, and $W^\alpha(z_o)$ is the window function used to construct the average. The equation above is not immediately related to the comoving space integral Eq.(2.9) of a lightcone field introduced in section 2.2. First, the observed redshift z_o of a galaxy may be subject to instrumental errors and therefore different from the actual redshift z'. Second, due to redshift space distortions, the redshift z' can be different from the background cosmological redshift z of the galaxy (which is simply related to the comoving distance χ). The second issue can be safely ignored for high enough multipoles, where the RSD correction to the power spectrum is unimportant [20, 90]. Since only small angular scale galaxy data is necessary for velocity reconstruction, we don't include RSD in our modelling (for the impact of RSD on correlations between velocity reconstruction and number counts see [76]) and will treat the actual redshift z' as the cosmological redshift z. The issue of measurement errors is discussed below in the context of a redshift galaxy survey subject to photometric redshift errors. In our analysis below, we consider two prototype galaxy surveys: a Rubin-like survey with many photometric redshift bins and a WISE-like survey with a single wide photometric redshift bin. For velocity reconstruction, these two surveys will be used as prototypes for the 'multiple density window' and 'single density window' cases for the quadratic estimators described in Sec. 2.2.5.

Rubin-like survey

For the Rubin-like survey, we consider Gaussian errors on photometric redshifts, with the probability of assigning redshift z_o to a galaxy with true redshift z (following Ref.[10])

given by:

$$P(z, z_o) = \frac{\exp\left[-\frac{(z_o - z)^2}{2\sigma_z^2}\right]}{\int_0^\infty d\tilde{z} \exp\left[-\frac{(\tilde{z} - z)^2}{2\sigma_z^2}\right]} \tag{2.130}$$

where $\sigma_z = \sigma_0(1+z)$ with σ_0 parametrizing the size of the photometric errors. We assume a fiducial value of $\sigma_0 = 0.05$. With this probability distribution, the galaxy average for the window $W^\alpha(z_o)$ can be expressed as an integral over the actual redshifts z:

$$g_{\ell m}^{W^\alpha} = \int dz \left[\int dz_o W^\alpha(z_o) P(z, z_o)\right] g_{\ell m}(z), \tag{2.131}$$

and in terms of the comoving distance

$$g_{\ell m}^{W^\alpha} = \int d\chi \; W_{\rm eff}^\alpha(\chi) \; g_{\ell m}(\chi), \tag{2.132}$$

where we have defined the effective window function

$$W_{\rm eff}^\alpha(\chi) = H(z(\chi)) \int dz_o W^\alpha(z_o) P(z(\chi), z_o), \tag{2.133}$$

and $g_{\ell m}(\chi)$ are the lightcone moments of the underlying galaxy overdensity field. The angular power spectrum between two galaxy redshift bins coming from a photometric survey can then be expressed using Eq.(2.16) plus a shot noise term:

$$\begin{aligned} C_\ell^{g^{W^\alpha} g^{W^\beta}} &= \int d\chi_1 d\chi_2 \; W_{\rm eff}^\alpha(\chi_1) \, W_{\rm eff}^\beta(\chi_2) \\ &\times \int \frac{k^2 dk}{(2\pi)^3} \; \mathcal{K}_\ell^g(\chi_1, k) \, \mathcal{K}_\ell^g(\chi_2, k) \; P_{gg}(\chi_1, \chi_2, k) + \delta_{\alpha\beta} \frac{1}{\bar{n_g}^\alpha} \end{aligned} \tag{2.134}$$

where $P_{gg}(\chi_1, \chi_2, k)$ is the galaxy-galaxy power spectrum computed using the halo model (consistent with Refs. [78, 20]; see Appendix B.4 for a summary), $\mathcal{K}_\ell^g(\chi, k) = 4\pi i^\ell j_\ell(k\chi)$ is the galaxy projection kernel from three dimensional Fourier space onto the sky, and $\bar{n_g}^\alpha$ is the number of galaxies per steradian in redshift bin α. We assume shot noise that is uncorrelated between redshift bins, and compute the number density per bin assuming the galaxy number density $n(z)$ per square arcmin is [10]

$$n(z) = \frac{n_{\rm g}}{2z_0} \left(\frac{z}{z_0}\right)^2 \exp\left(\frac{z}{z_0}\right) \; , \tag{2.135}$$

with $z_0 = 0.3$ and $n_g = 40/\text{arcmin}^2$. We construct the effective window functions using:

$$W^\alpha(z_o) = \frac{\Pi^\alpha(\chi(z_o))}{H(z_o)} \tag{2.136}$$

with $\Pi^\alpha(\chi)$ defined as in Eq.(2.14). In the limit of $\sigma_0 \to 0$, where photometric redshift errors can be neglected, these window functions correspond to normalized top-hat windows in comoving space. We show the effects of the photometric errors in the galaxy-galaxy covariance matrix in Fig. 2.6. Bin-bin correlations are enhanced as expected and the auto-power at a particular bin is reduced due to the contamination from distant bins. The principal components of the galaxy survey can be founds using the procedure described in Sec.2.2.5 just by appropriately replacing the signal and noise matrices. Fig. 2.7 compares the effect of different photometric redshift error levels on the total signal to noise Eq.(2.72) of the galaxy survey as a function of the number of bins N. As expected, we observe that the photometric errors put a limit on how much radial resolution our galaxy survey can have. For our fiducial value of $\sigma_0 = 0.05$, the signal to noise is mostly saturated for more than 32 redshift bins.

unWISE-like survey

For the unWISE-like survey considered in this chapter, we model the "blue" sample used in Refs. [91, 92, 93] from the unWISE catalogue [94], which is based on data from the WISE mission [95]. This sample is characterized by a median redshift of $\bar{z} \sim 0.6$ and is reasonably uniform over a redshift range of $\delta z \sim 0.3$. The number density of the resulting map is $\bar{n} \simeq 0.95/\text{arcmin}^2$.

Following Ref. [91], we model the unWISE blue sample as a linearly biased tracer of dark matter plus shot noise. In particular, we model the galaxy-galaxy angular power spectrum as

$$\begin{aligned} C_\ell^{g^W g^W} &= \int d\chi_1 d\chi_2 \; W_{\text{eff}}\left(\chi_1\right) W_{\text{eff}}\left(\chi_2\right) \\ &\times \int \frac{k^2 dk}{(2\pi)^3} \; \mathcal{K}_\ell^g(\chi_1, k)\, \mathcal{K}_\ell^g(\chi_2, k) \; P_{gg}(\chi_1, \chi_2, k) + \frac{1}{\bar{n}_g} \end{aligned} \tag{2.137}$$

where the power spectrum is:

$$P_{gg}(\chi_1, \chi_2, k) = b(\chi_1) b(\chi_2) P_{mm}(\chi_1, \chi_2, k), \quad b\left[z(\chi)\right] = 1.2 z(\chi) + 0.8 \tag{2.138}$$

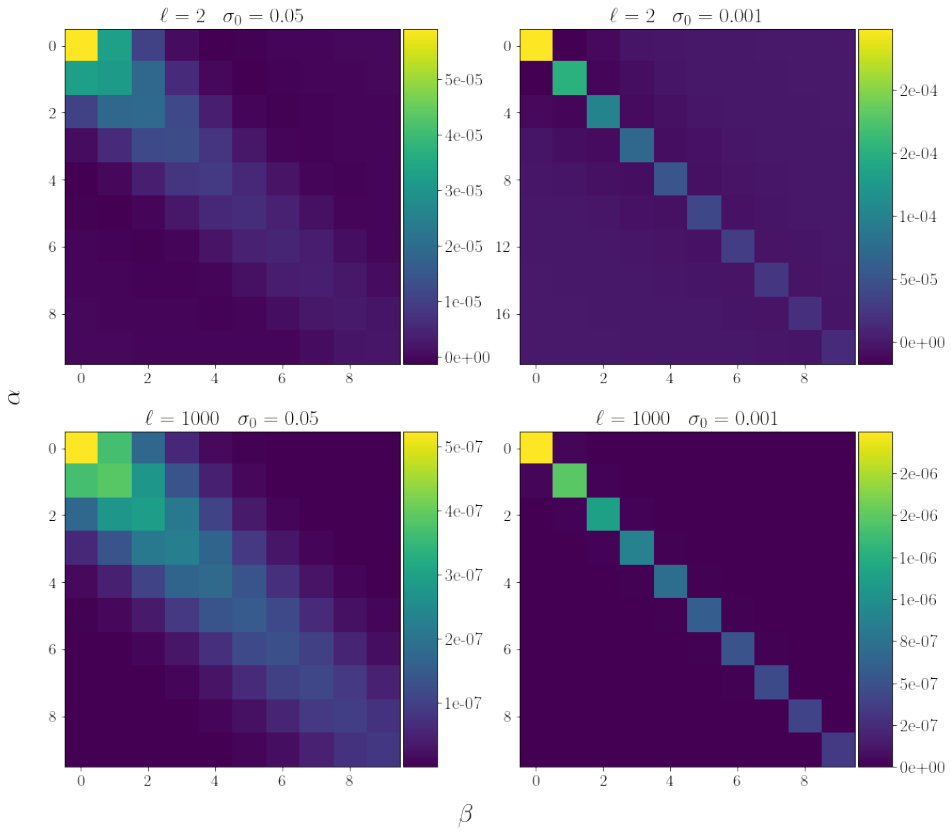

Figure 2.6: Section of the 32×32 redshift bins galaxy-galaxy covariance matrix, for different values of σ_0 and multipole ℓ.

Here, $P_{mm}(\chi_1, \chi_2, k)$ is computed in the halo model as described in Appendix B.4. The galaxy window function W_{eff} is simply the normalized comoving galaxy density

$$W_{\text{eff}}(\chi) = H(z(\chi)) \frac{dN}{dz} \tag{2.139}$$

where the redshift distribution of galaxies dN/dz is defined to be normalized as $1 = \int dz \; dN/dz$ and is reasonably uniform within a range $\delta z \sim 0.3$ of the median redshift $\bar{z} \sim 0.6$; the redshift distribution is shown in Fig. 2.16. The total number of galaxies in the survey is $\sim 1.4 \times 10^8$, yielding a shot noise of $1/\bar{n}_g = 9.2 \times 10^{-8}$. When performing velocity reconstruction, we must also compute the cross-power with the Π-binned optical depth and potential. In these cases, it is convenient to expand the observed moments of

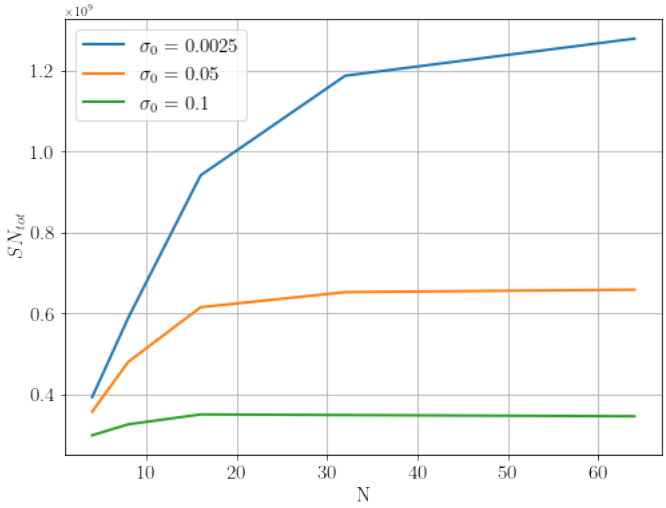

Figure 2.7: Total signal to noise for a photometric galaxy survey as a function on the number of redshift bins and for different error levels.

the galaxy overdensity as:

$$g^W_{\ell m} = \int d\chi\ H(\chi)\frac{dN}{dz}g_{\ell m}(\chi) \tag{2.140}$$

$$= \sum_\alpha \int d\chi\ \left[H\frac{dN}{dz}\ \Delta\chi\ \Pi^\alpha\right]\ g_{\ell m}(\chi) \tag{2.141}$$

and define a set of window functions

$$W^\alpha_{\rm eff}\,(\chi) = H\frac{dN}{dz}\ \Delta\chi\ \Pi^\alpha \tag{2.142}$$

We then define a set of binned galaxy moments as in Eq. 2.132 using the window functions Eq. 2.142. These binned galaxy moments are used to compute the cross-power with other Π-binned LC moments.

2.3.5 Galaxy survey systematics

Aside from the photometric redshift errors described above, one must consider a wide variety of systematics associated with a galaxy survey, many of which manifest on large angular scales (see e.g. [96, 97, 98, 99, 100]). Systematics that modulate the observed

number counts of galaxies are most problematic for velocity reconstruction, as they lead to a statistically anisotropic cross-power between the galaxy overdensity and CMB temperature that mimics the signal of interest. Additive effects that are uncorrelated with extragalactic sources, e.g. mis-identified stars included in the sample, are less problematic, adding only noise to the estimators but not bias. Starting from the observed number counts (following Refs. [98, 99, 100]), we model systematics effects as:

$$N^{W^\alpha}_{\rm obs}(\hat{\mathbf{n}}) = (1 + c(\hat{\mathbf{n}}))\, N^{W^\alpha}(\hat{\mathbf{n}}) \tag{2.143}$$

where $N^{W^\alpha}_g(\hat{\mathbf{n}})$ are the number counts of galaxies in a bin defined by the window function W^α. The modulating field $c(\hat{\mathbf{n}})$ encodes calibration errors which we might expand as a sum of effects associated with the instrument/observation strategy, extinction due to galactic dust, etc. Defining the underlying galaxy overdensity field $g^{W^\alpha}(\hat{\mathbf{n}})$ by $N^{W^\alpha}(\hat{\mathbf{n}}) = \bar{N}^{W^\alpha}(1 + g^{W^\alpha}(\hat{\mathbf{n}}))$, with $\bar{N}^{W^\alpha}$ the mean number of objects on the sky, the moments of the observed galaxy overdensity field are:

$$\begin{aligned}(g^{W^\alpha}_{\ell m})_{\rm obs} &= g^{W^\alpha}_{\ell m} + c_{\ell m} + \sum_{\ell_1,m_1;\ell_2,m_2} (-1)^m \sqrt{\frac{(2\ell+1)(2\ell_1+1)(2\ell_2+1)}{4\pi}} \\ &\times \begin{pmatrix} \ell_1 & \ell_2 & \ell \\ 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} \ell_1 & \ell_2 & \ell \\ m_1 & m_2 & -m \end{pmatrix} c_{\ell_1 m_1} g^{W^\alpha}_{\ell_2 m_2} + \mathcal{O}(\epsilon g^{W^\alpha}_{\ell m})\end{aligned} \tag{2.144}$$

where $g^{W^\alpha}_{\ell m}$ is defined as above in Eq. 2.132 and

$$\epsilon = \frac{1}{\bar{N}^{W^\alpha}} \sum_{\ell m} c_{\ell m} \langle N^{W^\alpha}(\hat{\mathbf{n}}) Y_{\ell m}(\hat{\mathbf{n}}) \rangle_{\rm sky} \tag{2.145}$$

is the correction to the mean number counts from each moment $c_{\ell m}$. Below, we neglect this correction to the mean. To model the form of the large-angular scale systematics, we assume that the modulating field $c(\hat{\mathbf{n}})$ is a Gaussian random field with power spectrum:

$$C^c_\ell = A^c e^{-(\ell/10)^2} \tag{2.146}$$

where we set the fiducial value for the amplitude A_c such that the variance of $c(\hat{\mathbf{n}})$ satisfies :

$$\sum \frac{(2\ell+1)}{4\pi} A^c e^{-(\ell/10)^2} = 10^{-4} \tag{2.147}$$

which corresponds to a level of calibration errors somewhere between the best current datasets and futuristic datasets (see Fig. 3 of [99])

2.3.6 CMB temperature-galaxy cross-power

As discussed above, there are a number of components of the CMB temperature that are correlated with tracers of large scale structure, such as the galaxy surveys considered above. Some of these contributions, such as the late-time linear ISW and extragalactic foregrounds, have a statistically isotropic cross-power. On the other hand, secondary components of the CMB such as lensing, kSZ, and late-time non-linear ISW will have a statistically anisotropic cross-power with the galaxy survey. Indeed, this statistical anisotropy is the basis for velocity reconstruction. We now consider these two cases in tern.

Statistically isotropic cross-correlations

The observed CMB anisotropies have contributions that are isotropically correlated with galaxies, including: extragalactic foregrounds (CIB, tSZ) and the late time linear ISW effect. To calculate isotropic cross-correlations we use Eq. 2.16:

$$\begin{aligned} C_\ell^{F^W G^{W'}} &= \int d\chi_1 d\chi_2 \; W\left(\chi_1\right) W'\left(\chi_2\right) \\ &\times \int \frac{k^2 dk}{(2\pi)^3} \; \mathcal{K}_\ell^F(\chi_1, k)\, \mathcal{K}_\ell^G(\chi_2, k) \; P_{FG}(\chi_1, \chi_2, k), \end{aligned} \tag{2.148}$$

and therefore we need to specify the window functions, integral kernels, and underlying power spectra for each of the temperature-galaxy signals. For the galaxies, we use the window functions introduced in Sec. 2.3.4 and the integral kernel $\mathcal{K}_\ell^g(\chi, k) = 4\pi i^\ell j_\ell(k\chi)$.

Extragalactic foregrounds are themselves tracers of large scale structure, and therefore are well-correlated with binned galaxy number density. We assume that extragalactic foregrounds can be described by random Gaussian fields. For this signals we use trivial window functions $W\left(\chi_1\right) = 1$, kernels $\mathcal{K}_\ell^{extra}(\chi, k) = 4\pi i^\ell j_\ell(k\chi)$ and underlying spectra $P^\nu_{CIBg}(\chi, k)$ and $P^\nu_{tSZg}(\chi, k)$ computed at each frequency ν using the halo model. In the left panel of Fig. 2.8, we show the cross-correlation between the extragalactic foregrounds at different frequencies and a Rubin-like galaxy survey in the redshift bin $z = (0.20, 0.26)$. We show as well the cross-power between the ILC cleaned temperature discussed in Sec. 2.3.3 and the galaxy survey in that same bin. The right panel shows the cross-power between cleaned temperature and galaxies at different redshift bins together with the cross-power between the linear late-time ISW signal and galaxies. The linear ISW-g correlation is calculated using Eq.(2.16) with the corresponding building functions.

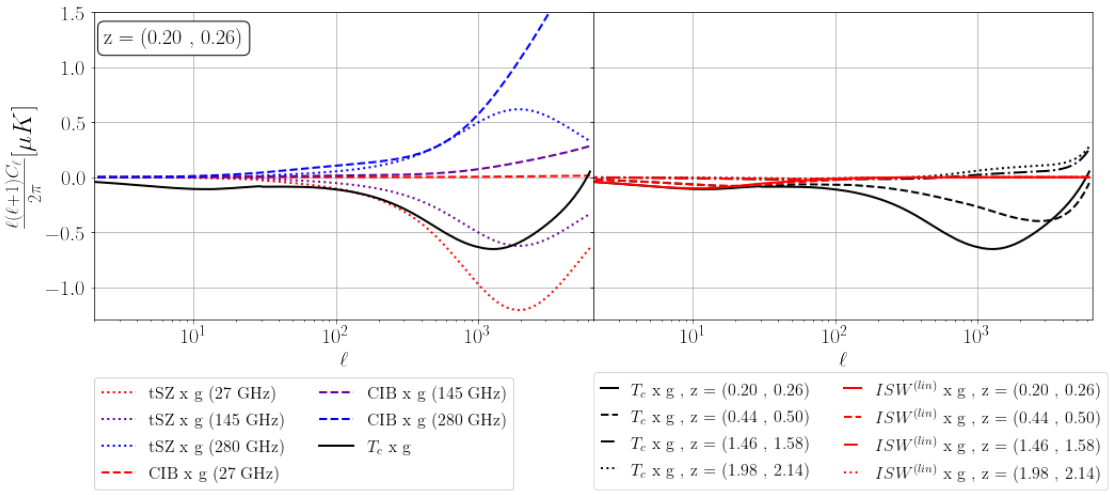

Figure 2.8: **Left panel:** Extragalactic foregrounds cross galaxies at redshift bin $z = (0.20, 0.26)$ and ILC cleaned temperature cross galaxies (solid line). **Right panel:** ILC cleaned temperature cross galaxies for several redshift bins compared to the linear-ISW cross galaxies.

Anisotropic cross-correlations

The main focus of this chapter are the statistically anisotropic cross-correlations between the CMB and galaxy surveys, as these are what allow us to perform radial velocity reconstruction. We work in the basis introduced in Sec. 2.2.5, expanding in terms of Π and Haar-binned LC moments to define the 'bulk' and 'fine' modes, respectively. Apart from the galaxy fields, the statistical anisotropies are expressed in terms of a series of fields described in Sec.2.3, namely: the radial velocity v, the differential optical depth $\dot{\tau}$, the transverse velocity potential Υ, the moving lens potential ψ, the lensing potential ϕ, the primary CMB anisotropy $\Theta^{\rm pCMB}$ and the photometric calibration c.

For the kSZ-galaxy cross-power we have:

$$\begin{aligned}\left\langle \Theta^{kSZ}_{\ell m}\, g^{W}_{\ell' m'} \right\rangle &= \sum_{\alpha=0}^{N-1} \sum_{\ell_1 m_1} (-1)^{m_1} \begin{pmatrix} \ell & \ell' & \ell_1 \\ m & m' & -m_1 \end{pmatrix} f^{v^\alpha W}_{\ell \ell_1 \ell'}\, v^{\alpha}_{\ell_1 m_1} \\ &+ \sum_{k=N}^{\infty} \sum_{\ell_1 m_1} (-1)^{m_1} \begin{pmatrix} \ell & \ell' & \ell_1 \\ m & m' & -m_1 \end{pmatrix} f^{v^k W}_{\ell \ell_1 \ell'}\, v^{k}_{\ell_1 m_1}\end{aligned}$$

where the bulk mode couplings $f^{v^\alpha W}_{\ell\ell_1\ell'}$ and the fine mode couplings $f^{v^k W}_{\ell\ell_1\ell'}$ are given by:

$$f^{v^\alpha W}_{\ell\ell_1\ell'} \equiv \sqrt{\frac{(2\ell+1)(2\ell_1+1)(2\ell'+1)}{4\pi}} \begin{pmatrix} \ell & \ell' & \ell_1 \\ 0 & 0 & 0 \end{pmatrix} C_{\ell'}^{\dot{\tau}^\alpha g^W} \Delta\chi. \tag{2.149}$$

and

$$f^{v^k W}_{\ell\ell_1\ell'} \equiv \sqrt{\frac{(2\ell+1)(2\ell_1+1)(2\ell'+1)}{4\pi}} \begin{pmatrix} \ell & \ell' & \ell_1 \\ 0 & 0 & 0 \end{pmatrix} C_{\ell'}^{\dot{\tau}^k g^W}. \tag{2.150}$$

The kSZ cross-power with galaxies forms the basis of the estimators used for radial velocity reconstruction. However, there are additional sources of statistical anisotropy in the cross-power that potentially introduce biases on the reconstructed velocity fields. Here, we focus on non-linear ISW, CMB lensing, and large angular scale calibration error in the galaxy survey. For the non-linear ISW-galaxy cross-power we have:

$$\begin{aligned} \left\langle \Theta^{ISW,nlin}_{\ell m}\, g^W_{\ell' m'} \right\rangle &= \sum_{\alpha=0}^{N-1} \sum_{\ell_1 m_1} (-1)^{m_1} \begin{pmatrix} \ell & \ell' & \ell_1 \\ m & m' & -m_1 \end{pmatrix} f^{\Upsilon^\alpha W}_{\ell\ell_1\ell'}\, \Upsilon^\alpha_{\ell_1 m_1} \\ &+ \sum_{k=N}^{\infty} \sum_{\ell_1 m_1} (-1)^{m_1} \begin{pmatrix} \ell & \ell' & \ell_1 \\ m & m' & -m_1 \end{pmatrix} f^{\Upsilon^k W}_{\ell\ell_1\ell'}\, \Upsilon^k_{\ell_1 m_1} \end{aligned}$$

where the bulk mode couplings $f^{\Upsilon_\alpha W}_{\ell\ell_1\ell'}$ and the fine mode couplings $f^{\Upsilon^k W}_{\ell\ell_1\ell'}$ are given by:

$$\begin{aligned} f^{\Upsilon^\alpha W}_{\ell\ell_1\ell'} &\equiv [\ell_1(\ell_1+1) + \ell'(\ell'+1) - \ell(\ell+1)] \\ &\times \sqrt{\frac{(2\ell+1)(2\ell_1+1)(2\ell'+1)}{16\pi}} \begin{pmatrix} \ell & \ell' & \ell_1 \\ 0 & 0 & 0 \end{pmatrix} C_{\ell'}^{\psi^\alpha g^W} \Delta\chi \end{aligned} \tag{2.151}$$

and

$$\begin{aligned} f^{\Upsilon^k W}_{\ell\ell_1\ell'} &\equiv [\ell_1(\ell_1+1) + \ell'(\ell'+1) - \ell(\ell+1)] \\ &\times \sqrt{\frac{(2\ell+1)(2\ell_1+1)(2\ell'+1)}{16\pi}} \begin{pmatrix} \ell & \ell' & \ell_1 \\ 0 & 0 & 0 \end{pmatrix} C_{\ell'}^{\psi^k g^W}. \end{aligned} \tag{2.152}$$

For the CMB lensing-galaxy cross-power we have:

$$\left\langle \Theta^L_{\ell m}\, g^W_{\ell' m'} \right\rangle = \sum_{\alpha=0}^{N-1} \sum_{\ell_1 m_1} (-1)^{m_1} \begin{pmatrix} \ell & \ell' & \ell_1 \\ m & m' & -m_1 \end{pmatrix} f^{\Theta W}_{\ell\ell_1\ell'}\, \Theta^{pCMB}_{\ell_1 m_1} \tag{2.153}$$

where the couplings $f^{\Theta W}_{\ell\ell_1\ell'}$ are given by:

$$\begin{aligned} f^{\Theta W}_{\ell\ell_1\ell'} &\equiv [\ell_1(\ell_1+1)+\ell'(\ell'+1)-\ell(\ell+1)] \\ &\times \sqrt{\frac{(2\ell+1)(2\ell_1+1)(2\ell'+1)}{16\pi}} \begin{pmatrix} \ell & \ell' & \ell_1 \\ 0 & 0 & 0 \end{pmatrix} C^{\phi g^W}_{\ell'} \end{aligned} \tag{2.154}$$

For the calibration error contribution we have:

$$\left\langle \Theta^{\nu}_{\ell m}\, g^{W}_{\ell' m'} \right\rangle\Big|_{cal} = \sum_{\alpha=0}^{N-1} \sum_{\ell_1 m_1} (-1)^{m_1} \begin{pmatrix} \ell & \ell' & \ell_1 \\ m & m' & -m_1 \end{pmatrix} f^{cW}_{\ell\ell_1\ell'}\, c_{\ell_1 m_1} \tag{2.155}$$

where the couplings $f^{cW}_{\ell\ell_1\ell'}$ are given by:

$$f^{cW}_{\ell\ell_1\ell'} \equiv [\ell_1(\ell_1+1)+\ell'(\ell'+1)-\ell(\ell+1)] \begin{pmatrix} \ell & \ell' & \ell_1 \\ 0 & 0 & 0 \end{pmatrix} C^{Ig^W}_{\ell'} \tag{2.156}$$

where I stands for the CMB components isotropically correlated with the galaxies. Other effects leading to a statistically anisotropic cross-power which we anticipate will be less important, and which we do not compute here, include: relativistic aberration of the CMB [101] (similar effect as calibration error, but with a smaller magnitude), SZ effects at higher order in velocity and temperature (see e.g. Refs. [102, 103, 104, 105, 62, 106]), anisotropic/ill-characterized beam patterns in the CMB experiment (see e.g. Ref. [107] for an assessment of the impact on lensing reconstruction), and perhaps others. In the case of CMB lensing, note that the modulating field is the primary CMB temperature. Although we do not explore it further here, we note that a quadratic estimator for the low-ℓ primary CMB can be formulated from the CMB-galaxy cross-power using the formalism introduced in Sec. 2.2.5. A similar estimator was introduced in Ref. [108] as a means to reconstruct the primary CMB dipole, which is not directly measurable due to the contribution from our local peculiar velocity.

2.4 Reconstruction analysis

In Section 2.2.5, we discussed the details involved in constructing quadratic estimators for fields sourcing a statistical anisotropy in the CMB-LSS cross-correlation. We showed that information about these fields can be reconstructed up to a series of noise terms. The purpose of this section is to analyse the relations between signal and noise for the reconstruction of the Π-binned LC moments of the radial velocity, which act as sources for the kSZ-LSS statistical anisotropy. We estimate the signal and noise for a reconstruction using the modelling for the CMB, LSS, and their correlation presented in Section 2.3.

2.4.1 Radial velocity reconstruction for SO x Rubin

Applying the formalism of Section 2.2.5 to the reconstruction of $(v_{\rm eff}^{\alpha})_{LM}$ leads to a collection of estimators with the following two point function:

$$\begin{aligned}\left\langle (\hat{v_{\rm eff}^{\alpha}})_{LM} (\hat{v_{\rm eff}^{\beta}})^{*}_{LM} \right\rangle &= \left(\mathbf{R}_L \mathbf{C}_L^{vv} (\mathbf{R}_L)^{\dagger}\right)^{\alpha\beta} + \left(\mathbf{N}_L^0\right)^{\alpha\beta} + \left(\mathbf{N}_L^{fine}\right)^{\alpha\beta} \\ &+ \left(\mathbf{N}_L^{cal}\right)^{\alpha\beta} + \left(\mathbf{N}_L^{\Upsilon}\right)^{\alpha\beta} + \left(\mathbf{N}_L^{lens}\right)^{\alpha\beta} \end{aligned} \tag{2.157}$$

where the various terms are defined by:

- $\mathbf{R}_L \mathbf{C}_L^{vv} (\mathbf{R}_L)^{\dagger}$: the covariance matrix of Π-binned LC moments of the radial velocity field. The rotation matrix $\mathbf{R}_L$ encodes the bin-bin mixing of the signal covariance due to the redshift error in the galaxy survey.
- $\mathbf{N}_L^0$: the Gaussian reconstruction noise Eq. 2.51, with coupling functions defined by Eq. 2.149. This term comes from the disconnected contractions in Eq.(2.67) (e.g. $\langle\Theta\Theta\rangle\langle\delta\delta\rangle$ and $\langle\Theta\delta\rangle\langle\Theta\delta\rangle$). Note that we do not include the non-Gaussian contributions to the estimator noise in the present analysis (e.g. the $N^{3/2}$ and N^1 noise terms, in the terminology of Ref. [21]); see Appendix B.3 for discussion.
- $\mathbf{N}_L^{fine}$: the estimator variance coming from the fine mode bias Eq. 2.59, with coupling functions for the bulk and fine modes of the radial velocity field defined by Eq. 2.149 and 2.150, respectively. The relative importance of this term decreases with an increasing number of bins; we explore this in detail below.
- $\mathbf{N}_L^{\Upsilon}$: the estimator variance due to the non-linear ISW effect, defined by Eq. 2.64 using the coupling function for the transverse velocity potential Eq. 2.149.
- $\mathbf{N}_L^{cal}$: the estimator variance due to galaxy survey calibration error systematics, defined by Eq. 2.64 using the coupling function for the calibration error Eq. 2.156.
- $\mathbf{N}_L^{lens}$: the estimator variance due the lensing of the primary CMB, defined by Eq. 2.64 using the coupling function for the lensing potential Eq. 2.154.

Note that we refer to the contribution $\mathbf{R}_L \mathbf{C}_L^{vv} (\mathbf{R}_L)^{\dagger}$ as the "signal" and all other terms as the "noise" in the discussion that follows.

In Fig.2.9 we show a few diagonal elements (e.g. $\alpha = \beta$) of Eq.(2.157) for a near bin at $z \sim 0.5$ and far bin at $z \sim 1.5$. The dominant source of reconstruction noise is the

$\mathbf{N}_L^0$ term, followed by the fine mode and calibration error contributions to the variance. The variance arising from the transverse velocity potential and lensing are negligibly small compared to the Gaussian estimator noise; we therefore neglect these terms in our analysis below.

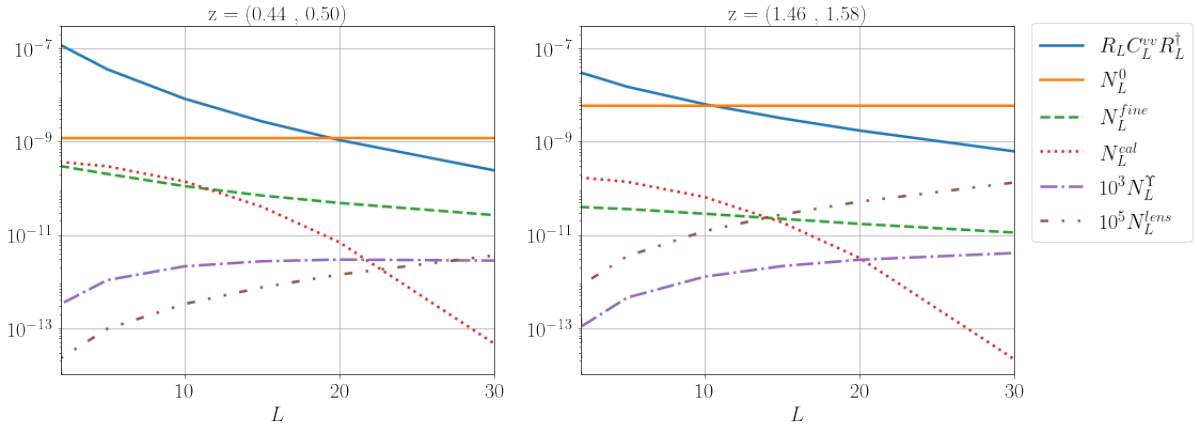

Figure 2.9: Radial velocity signal and noise sources at redshift bins $z = (0.44\ ,\ 0.50)$ and $z = (1.46\ ,\ 1.58)$, corresponding to bins 4 and 16 of 32.

There are significant bin-bin correlations in the estimator variance both due to the signal and the various noise terms. Photometric errors in the galaxy surveys lead to mixing of radial information that contributes to the bin-bin correlation, and this radial mixing is captured by the rotation matrix $\mathbf{R}_L$. In Fig. 2.10 we show, for fixed L, the radial mixing for a set of redshift bins and illustrate how the mixing decreases when the photometric errors are less severe. The rotation matrix is found to be largely independent of the multipole L for $L \lesssim 200$.

In Fig. 2.11 we show, for fixed L, the contributions to the bin-bin covariance from the various noise terms. The Gaussian reconstruction noise is correlated between bins, mainly due to the the correlation between structures in nearby bins induced by the redshift error in the galaxy survey. This is the largest contribution to the bin-bin covariance in nearby bins, independent of L. There is a less significant, but non-negligible, short-range correlation induced by the fine-mode noise which is most important at low-L. We observe that the bias from the calibration error induces long-range bin-bin correlations in the estimator variance, as expected due to our assumption that the calibration error is the same for each bin. Had we assumed a different calibration error in each bin, there would be no such correlation.

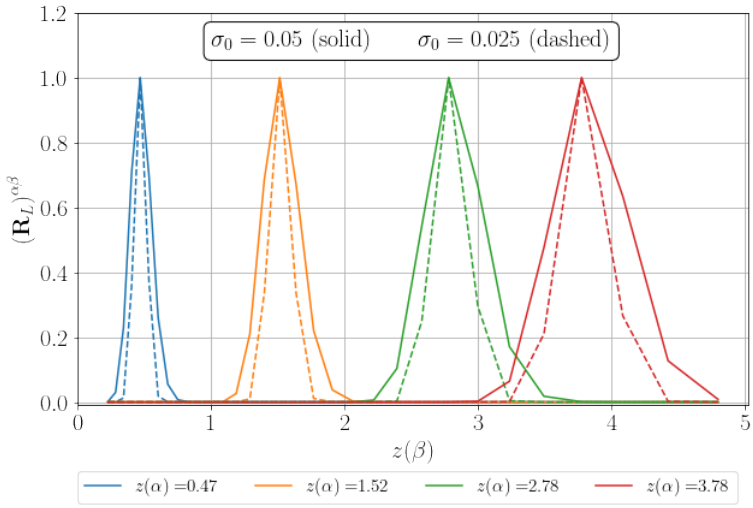

Figure 2.10: Rows 4, 16, 24, and 28 of the rotation matrix, corresponding to redshift bins with central redshifts 0.47, 1.52, 2.78 and 3.78. Solid (dashed) lines are obtained using the photometric redshift error model described in Sec.2.3.4 with $\sigma_0 = 0.05$ ($\sigma_0 = 0.025$). As expected, the mixing of radial information reduces for less severe redshift errors.

Principal components

In light of the significant bin-bin covariance present at all scales L in both the signal and the noise terms in the bin basis, it is instructive to consider the principal component basis where there is no covariance. The transformation to the principal component basis was outlined above in Sec. 2.2.5 and defined by:

$$(\hat{v}_{\rm eff}^{Pj})_{LM} = \sum_{\beta} c_L^{j\beta} (\hat{v}_{\rm eff}^{\beta})_{LM} \tag{2.158}$$

Note that we employ the full signal covariance and all noise terms in Sec. 2.157 to define the principal components. In Fig. 2.12 we show the $j = 1, 2, 3$ principal component coefficients $c_L^{j\beta}$ as a function of bin β for $N = 64$ bins at $L = 1, 2, 5, 10$. Note that at each scale L, the weight for the most significant principal components receives support primarily from lowest redshifts. This is where the galaxy density is relatively high (hence the reconstruction noise is minimized) and the amplitude of velocities is relatively large (e.g. due to linear growth). In addition, the number of nodes along the radial direction increases with L and for the lower signal to noise principal components at fixed L.

In the principal component basis, we can define a measure of the total signal to noise

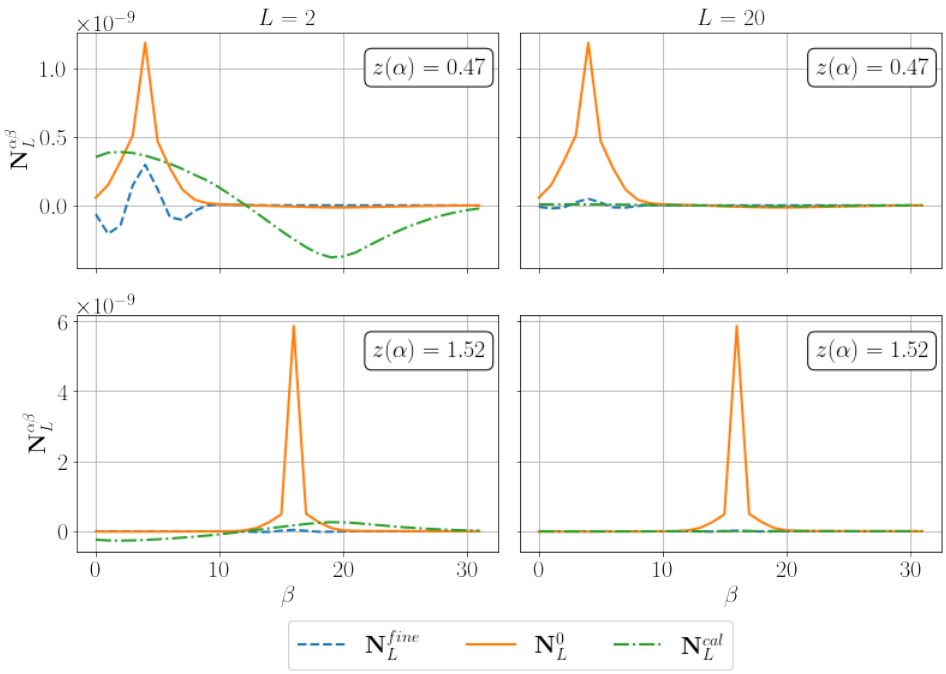

Figure 2.11: Contributions to the noise covariance matrix from the Gaussian reconstruction noise $\mathbf{N}_L^0$, the fine mode noise $\mathbf{N}_L^{fine}$, and the calibration error noise $\mathbf{N}_L^{cal}$. We show rows 4 (top panels) and 16 (bottom panels) of the noise matrices for multipoles $L = 2, 20$.

per mode LM by

$$SN_{LM} = \sum_{j=1}^{N} (\langle \hat{v}_{\rm eff}^{Pj})_{LM} (\hat{v}_{\rm eff}^{Pj})_{LM}^{\dagger} \rangle \tag{2.159}$$

We evaluate this quantity in Fig. 2.13. Each panel of that figure compares, for a reconstruction with N redshift bins, the effect of adding the $\mathbf{N}_L^{fine}$ and $\mathbf{N}_L^{cal}$ compared with the Gaussian reconstruction noise $\mathbf{N}_L^0$. As expected, with an increasing number redshift bins, the fine mode contribution becomes less important and the calibration error becomes the leading correction to the Gaussian reconstruction noise. The reconstruction of the Π-binned LC moments of the radial velocity suffer a considerable loss of signal to noise per mode as we reduce the number of bins even when only including the Gaussian reconstruction noise. For our binning scheme, $N = 64$ corresponds to redshift bins of equal coming size of approximately 110 Mpc. The coherence length of the velocity field is around 70 Mpc, and therefore it makes sense that the fine modes become more relevant for $N = 32$ and smaller, as the size of the bins are considerably larger than 70 Mpc. Even with $N = 64$

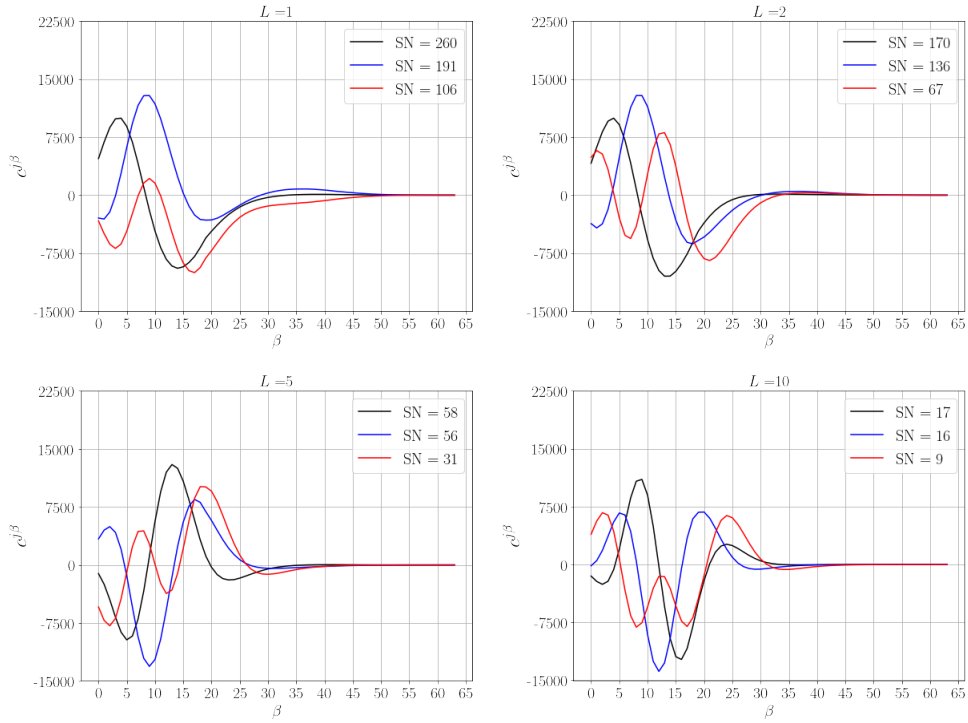

Figure 2.12: Coefficients $c_L^{j\beta}$ for the three largest signal to noise principal components $\sum_\beta c_L^{j\beta}(\hat{v}_{\rm eff}^\beta)_{LM}$ for a reconstruction with 64 redshift bins, at various L multipoles.

bins, comparing the orange and green curves, we see that calibration error leads to a significant degradation in SN_{LM} of greater than 10%. Efforts to mitigate systematics in galaxy surveys on large angular scales can therefore meaningfully impact the fidelity of the reconstruction. Regardless, we see that velocity reconstruction with SO x Rubin will have exceedingly high SNR on large angular scales, with $SN_{LM} > 1$ for $L < 30$ with the most significant principal component.

Optical depth bias

As discussed in Sec.2.2.5, incorrect modelling of the correlation between the electron and the galaxy distribution leads to a multiplicative bias in the reconstructed radial velocity, commonly referred to as the optical depth bias. We illustrate how the bias shows up in our formalism by considering a one parameter toy model for the electron-galaxy correlation function, based on the halo model. If we fix the model parameters determining how galaxies

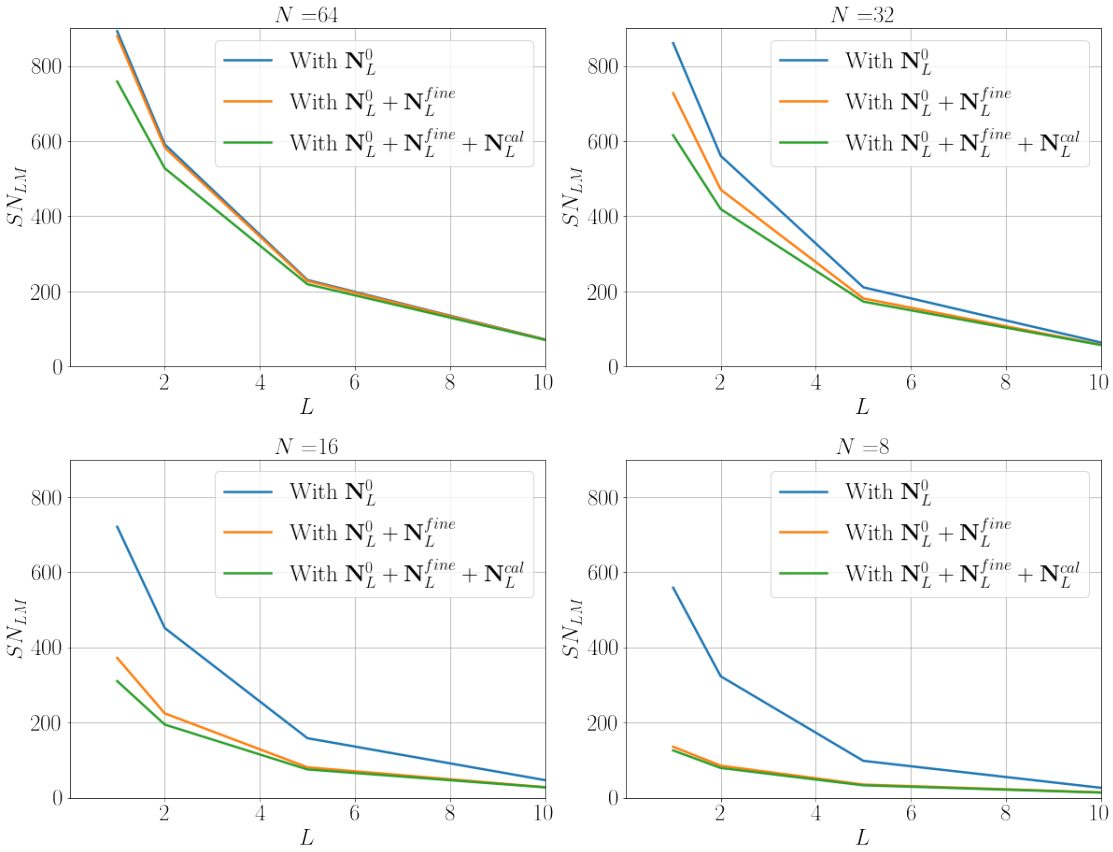

Figure 2.13: Signal to noise per mode as defined in Eq.(2.71), as a function of L, for different binning and sources of noise. The loss of signal to noise due to the fine mode noise accentuates for wider redshift bins.

inhabit dark matter halos, the electron-galaxy cross-power is determined by the model for the electron density profile inside dark matter halos (see Appendix B.4 for more details). Due to physical processes such as AGN feedback, baryonic matter does not track dark matter inside of halos. In Fourier space, this translates into an electron density profile $u_e(k, M, z)$ that is different from the dark matter density profile $u(k, M, z)$. To explore a one-parameter family of models, we construct the following toy model for the electron density profile :

$$u_e^{(toy)}(k, M, z) = \frac{u_e(Ak, M, z)}{u(Ak, M, z)} u(k, M, z) \tag{2.160}$$

where A is a continuous parameter that interpolates between the dark matter profile ($A = 0$) and our fiducial model ($A = 1$) that incorporates feedback. If we take our fiducial model to be the true model for the electron density, we can explore the optical depth bias when 'incorrect' $A \neq 1$ profiles are used for velocity reconstruction.

Fig. 2.14 shows the behaviour of the elements of the bias matrix Eq. 2.74 as a function of the angular scale (left panel) and redshift (right panel). We find that the bias is practically independent of L for large angular scales $L \lesssim 200$ (in accordance to what was found in [21]) and that it is less significant at higher redshifts, were the difference between electron and dark matter perturbations is less pronounced. The off-diagonal elements of the bias matrix are similarly scale-independent and approach one at high redshift. The L independence on large angular scales is a feature that we expect to be robust independently of the models under consideration. Note that the bias is always less than one for this family of models. This is because feedback in the fiducial model causes the electron halo profiles to be more diffuse than their host dark matter halos, leading to a power suppression at high-k, and therefore on small angular scales where the sums in Eq. 2.74 receive the most weight. Even in the extreme case where baryons are assumed to trace dark matter, the magnitude of the bias over the entire range of redshifts lies within a reasonably small range $0.6 \lesssim \Gamma_L^{\mathrm{XY}} < 1$.

To obtain cosmological constraints from velocity reconstruction using future datasets, it will be necessary to incorporate the optical depth bias into the analysis. For example, if we wish to obtain constraints on a set of cosmological parameters $\mathbf{m}$ appearing in the radial velocity power spectrum $\mathbf{C}_L^{vv}(\mathbf{m})$, it is necessary to compare (via e.g. a likelihood function) the measured velocity spectra to the model:

$$(C^{\mathrm{recon}})_L^{\alpha\beta} = \left[(\mathbf{\Gamma R})\,\mathbf{C}_L^{vv}(\mathbf{\Gamma R})^\dagger\right]^{\alpha\beta} + (\mathbf{N}_L)^{\alpha\beta} \tag{2.161}$$

where $(\mathbf{N}_L)^{\alpha\beta}$ includes the most relevant noise terms (e.g. $N^{(0)}$, fine-mode, calibration) and $(\mathbf{\Gamma R})^{ij} = \Gamma^{ij}(u_e)\mathrm{R}^{ij}(\sigma_z)$, where we assume the optical depth bias and rotation matrix are independent of L (a good approximation, as shown above) and indicate explicitly the dependence on the electron profile u_e and redshift error σ_z. To get access to the cosmological information contained in $\mathbf{C}_L^{vv}(\mathbf{m})$ it is necessary to encapsulate the redshift errors and electron profile into a set of nuisance parameters that can be marginalized over. In the absence of any modelling, there are N_{bin}^2 nuisance parameters. This is the same number of independent entries in $\mathbf{C}_L^{vv}(\mathbf{m})$ that determine redshift-redshift correlations, implying that the only residual cosmological information is in the shape of the velocity power spectrum. This impedes one's ability to learn about e.g. the growth function using the reconstructed velocity field. But this scenario is far too pessimistic, as it does not incorporate information from other sources, or physical constraints present in the modelling.

Optimistically, it may be sufficient to characterize redshift errors and the electron profile by a small number of model parameters. For example, the fiducial model of Gaussian redshift errors considered above contains a single parameter σ_0. Assuming for the moment that this is an accurate model for Rubin redshift errors, there is a single model parameter associated with the rotation matrix. In addition, one can put a prior on the ranges this parameter might take by using other available information: the galaxy-galaxy power spectrum itself, simulations, comparing with a spectroscopic survey, etc. Likewise, if Eq. 2.160 is a reasonable description of the range of possible electron profiles, then a single model parameter would determine the optical depth bias. Again, one could incorporate additional measurements to provide a prior on A, for example by independently measuring the galaxy-electron cross-power using Fast Radio Bursts [66] or by correlating the reconstructed velocity field with the galaxy survey [109, 20, 76] or, if available, the transverse velocity field [110]. In reality, there are likely more than two model parameters to consider to fully characterize redshift errors and the electron profiles. But by evaluating a range of physical models and finding complementary observations, one can likely put an informative prior on the $N_{\rm bin}^2$ degrees of freedom in the rotation matrix and optical depth bias.

In a number of previous analyses, e.g. Refs. [77, 76, 111] it was assumed that the rotation matrix was diagonal, and therefore that the optical depth bias consisted of $N_{\rm bin}$ parameters to be marginalized over. In the presence of photometric redshift errors, we have seen that this assumption does not hold. – the off-diagonal nature of the rotation matrix gives rise to greater than $N_{\rm bin}$ parameters. How many additional parameters need to be incorporated depends on how dominant the diagonal terms in $(\mathbf{N}_L)^{\alpha\beta}$ are compared with the off-diagonal terms, since small off-diagonal terms can be neglected. This depends primarily on the magnitude of redshift errors, so more accurate photometric redshifts, or spectroscopic redshifts can simplify the analysis of the reconstructed velocity field. In future work, cosmological forecasts and analyses should take into account the off-diagonal terms in the optical depth bias, either through a physical model or by marginalizing over a sufficient number of degrees of freedom.

2.4.2 "Double" SO and pre-reconstruction vs post-reconstruction cleaning

In this section, we investigate two scenarios related to the effect of foregrounds on the reconstruction. First, we investigate whether or not additional frequency channels can help with mitigating the effect of foregrounds on the reconstruction. To do so, we define a hypothetical experiment we refer to as "Double" SO, which has a set of channels (in

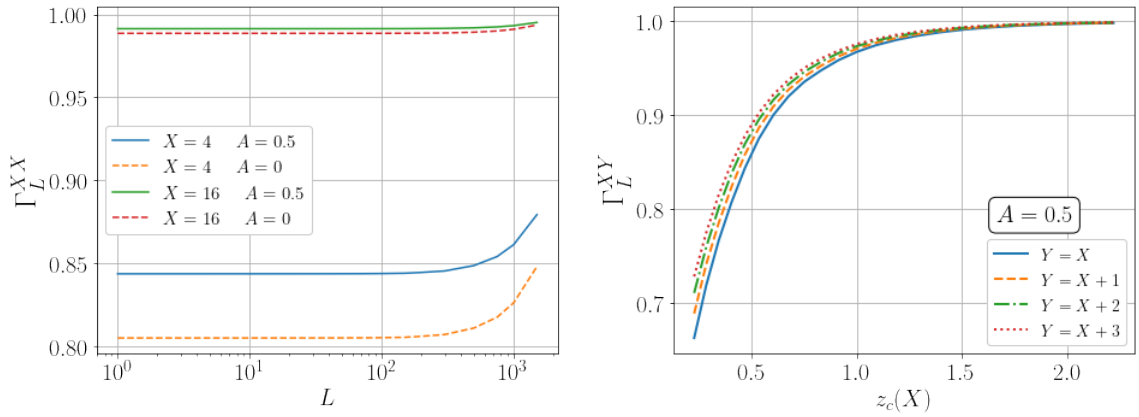

Figure 2.14: **Left panel:** Diagonal elements of the optical depth bias matrix as a function of the multipole L. Bin 4 and bin 16 correspond redshifts ranges $(0.44, 0.50)$ and $(1.46, 1.58)$ respectively. Solid (dashed) lines correspond to $A = 0.5$ ($A = 0$). **Right panel:** Diagonal and some off-diagonal elements of the bias matrix as a function of redshift. Bias tends to 1 for higher redshifts as electrons trace dark matter more closely at earlier times.

GHz) at: 47, 52, 63, 75, 91, 109, 131, 158, 190, 228, 275, 330, 397, 478, 575, 691, 831, 1000. The boundaries and spacing of this selection were chosen to minimize residuals in the cleaned CMB temperature spectrum for our foreground model. Including frequencies below ∼50 GHz and above ∼1000 GHz provides no improvement for removing extragalactic foregrounds. Our choice of 12 frequency channels in the relevant range is somewhat arbitrary, and is simply meant to be representative of a reasonable number of detectors as compared to SO. To define the noise properties of Double SO, we first take the SO LAT TT noise model [15] assumed above and define a linear interpolating function on the three free parameters in the noise model, extrapolating when necessary to higher frequencies that are not in the SO selection. We then analyzed the reconstruction noise for the fiducial $N = 64$ bin case assumed for SO x Rubin above. In Fig.(2.15) we show the signal to noise for the first two principal components over a range of scales for SO and Double SO and Rubin. It can be seen that the Double SO experiment (true to its name) yields a signal to noise that is about twice as good as SO. This is due to a combination of a lower effective noise in the auto-power as well as a reduction of foreground residuals in the cross-power.

Above, we considered the scenario where a linear combination of CMB maps was used to remove foregrounds before velocity reconstruction. It is also possible to perform velocity reconstruction on each frequency map, and then find the linear combination of reconstructions that minimizes the variance of the reconstruction. To do so, we describe the power

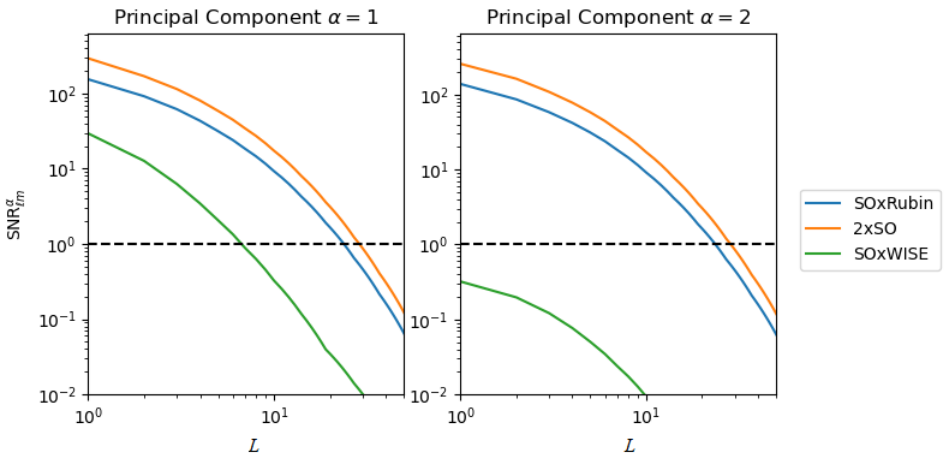

Figure 2.15: SNR of principal components 1 and 2 of velocity reconstruction.

spectrum of the reconstruction as

$$\mathbf{C}_\ell = C_\ell^{vv}\mathbf{e}\mathbf{e}^t + \mathbf{N}_\ell \tag{2.162}$$

We can then apply the same harmonic space ILC method defined above to find a map that minimizes the variance due to reconstruction noise:

$$C_\ell^{\mathrm{clean}} = \mathbf{w}_\ell^\dagger \mathbf{N}_\ell \mathbf{w}_\ell + C_\ell^{vv} \tag{2.163}$$

where the ILC weights $\mathbf{w}_\ell$ are defined using the reconstructed spectra:

$$\mathbf{w}_\ell = \frac{\mathbf{C}_\ell^{-1}\mathbf{e}}{\mathbf{e}^t\mathbf{C}_\ell^{-1}\mathbf{e}} \tag{2.164}$$

Since we know the signal C_ℓ^{vv} we can subtract this from the reconstruction in the pre-reconstruction cleaning scenario to arrive at the residual noise, which we compare directly to $\mathbf{w}_\ell^\dagger \mathbf{N}_\ell \mathbf{w}_\ell$ from the post-reconstruction cleaning scenario. We find that the residuals for pre-reconstruction cleaning are smaller than the residuals for post-reconstruction cleaning. Therefore, we focus on the scenario where foregrounds are mitigated before reconstruction is performed.

2.4.3 Radial velocity reconstruction for SO x unWISE

We now turn to the second scenario we consider, where velocity reconstruction is performed with SO and data currently available from the unWISE blue sample. Here, there is a

single galaxy window function, which is plotted in Fig. 2.16. We consider a reconstruction using 8 bins in the redshift range between $0.2 < z < 1.5$ corresponding to a comoving bin width of $\Delta\chi \simeq 450$ Mpc. To compute the fine mode noise, we use 512 bins in the same redshift range. We increase the calibration error from our Rubin framework by a factor of 10^2 to account for the difference in precision of redshift measurement between the two experiments. Because there is a single galaxy window function, the reconstructed velocity field and reconstruction noise will be highly correlated among the 8 bins in which we perform the reconstruction. Therefore, it is crucial in this case to use the principal component basis. Fig. 2.16 shows the $\alpha = 1$ principal component coefficients $c_L^{\alpha\beta}$ both with and without the inclusion of fine-mode noise and calibration errors for $L = 1$ and $L = 5$. Note that for $L = 1$, the first principal component roughly traces out the unWISE window function dN/dz when the fine-mode noise and calibration errors are neglected. However, the first principal component becomes oscillatory once the additional noise is included. This is due to the redshift-redshift correlations of the noise terms obscuring the redshift correlations in the signal. The signal to noise of the first principal component at $L = 1$ drops from 25 to 16.4 as the additional noise terms are added. At $L = 5$, the first principal component has an oscillatory structure in redshift both with and without the additional sources of noise. The signal to noise of the first principal component at $L = 5$ is 1.7 and 1.3 with and without the additional noise terms. Therefore, most of the signal lies at the lowest L. Analyzing the higher principle components, they make an insignificant contribution to the signal to noise at all scales. We therefore can focus on the first principle component only.

We explore the effect of changing the number of bins used in the analysis by computing the signal to noise SN_{LM} defined in Eq. 2.71, summing over principal components at fixed L. We find that it is numerically difficult to consider greater than 8 bins. Large bin-bin correlations in the signal covariance and Gaussian reconstruction noise, especially in bins where the redshift distribution is small, lead to poorly conditioned rotation matrices (Eq.2.43) that spoil the construction of the principal component basis. Therefore, we consider scenarios with 4 and 8 bins. The result for the signal to noise per mode, summed over principal components, is shown in Fig. 2.17. Here, the dependence of the signal to noise on the number of bins is less dramatic than for SO x Rubin. This is to be expected, since not much information is gained by finer sampling in redshift due to the fact that there is a single wide galaxy window function. In this figure, we also demonstrate the effect of fine mode noise. For 4 bins, there is a significant correction beyond the Gaussian reconstruction noise. However, we see that for 8 bins, we are able to improve on the signal to noise in the presence of fine mode noise.

Finally, in Fig. 2.15 we compare the SN_{LM}^{α} attainable for SO x unWISE compared

with SO x Rubin. The signal to noise per mode for the first principle component for SO x unWISE is roughly an order of magnitude lower than for SO x Rubin; for the second principle component the difference is three orders of magnitude. Although there is a significant galaxy density in the unWISE sample, yielding a small Gaussian reconstruction noise (at least over some range in L for the first principal component), there is little redshift information. We therefore can only expect to obtain coarse-grained knowledge of the velocity field from such an analysis. Nevertheless, this is in principle important information, and the reconstruction of the first principle component at signal to noise greater than unity can be obtained for $L \lesssim 10$. This represents a modest, but non-trivial, number of well measured modes.

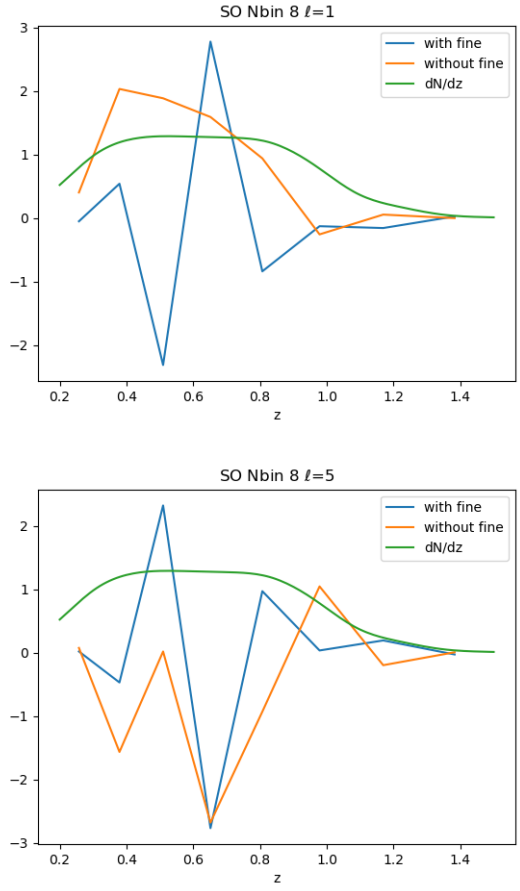

Figure 2.16: First principal component of the velocity reconstruction for SO x unWISE.

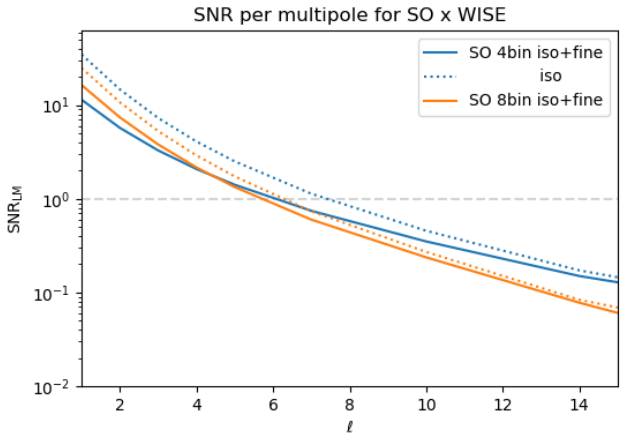

Figure 2.17: SNR per multipole for SO x unWISE for 4 and 8 coarse bins and 512 fine bins. SNR with and without fine mode contributions are shown.

2.5 Velocity reconstruction pipeline

In this section we assess the performance of radial velocity reconstruction with future datasets using a suite of simulations and a reconstruction pipeline based on the quadratic estimators described in previous sections. Previous work has demonstrated the effectiveness of the quadratic estimator for reconstruction of the radial velocity field using N-body simulations in the box geometry [21], and in Chapter 3 we present simulations for the lightcone geometry. Here, we focus on simulated data that consists of properly correlated random Gaussian fields including: the radial velocity field, galaxy number counts with photometric redshift errors, the electron density field, the primary CMB, the kSZ contribution to the CMB, and extragalactic foreground contributions to the CMB. We develop a reconstruction pipeline for the radial velocity field using fast real-space versions of the quadratic estimators described above. Theoretical modelling is an important component of velocity reconstruction, since it appears in the estimator for the Π-binned moments of the velocity fields and also in the rotation matrices required to de-bias the estimators. This makes a combined pipeline including both the simulation of the maps and the application of the estimators essential.

The benefit of using a Gaussian simulation framework is that the ensemble-average properties of the estimator are well-understood on the full sky using the results of previous sections, which allows us to validate the analysis pipeline. Another benefit is that we can isolate and investigate the effect of map-based systematics such as masking on the

reconstruction to compare with results on the full sky. Since the generation of correlated random Gaussian fields is far less computationally intensive than running a suite of N-body simulations, it is possible to explore ensemble averages, and quantify underlying numerical inaccuracies or biases. A disadvantage of this approach is that we miss important non-linear contributions to the reconstruction. As shown in Ref. [21], for radial velocity reconstruction this includes a contribution to the reconstruction noise analogous to the $N^{(3/2)}$ bias in lensing reconstruction [112]. At the resolutions considered in Ref. [21], this was in fact larger than the Gaussian contributions to the reconstruction noise by a factor of ~ 2. At the somewhat lower resolution and higher instrumental noise we consider, we expect this contribution to be smaller, and sub-dominant to the Gaussian contributions. Another non-linear effect included in the simulations presented in Chapter 3 is redshift space distortions, which were found to have minimal impact on the reconstruction at the resolutions simulated. All previous work has relied on dark matter-only N-body simulations, making the approximation that baryons follow the N-body particles. This assumption will fail on the small scales relevant for velocity reconstruction. Under the assumption of statistical isotropy made above, this mis-modelling of baryons shows up as a multiplicative bias (see Sec. 2.4.1 for a discussion), but at the non-linear level there may be additional effects. Future work with simulations should certainly include baryonic effects to explore the impact on simulations at the non-linear level.

2.5.1 Simulations

In this work, we approximate the primary CMB, galaxy number counts, components of the velocity field, electron density, and extragalactic foreground contributions to the CMB as correlated random Gaussian fields. Using the complete set of spectra and cross-spectra between these fields, we can construct a multivariate Gaussian distribution from which to drawn properly correlated realizations. These realizations can be used to compute the kSZ contribution to the CMB. Signals constructed this way will show the expected statistical anisotropy when correlated with the galaxy density. For each set of realizations, the quadratic estimator for the underlying radial velocity can be applied, allowing us to validate the statistics of the estimators by averaging over many realizations. Here is the list of steps we take to generate a suite of Gaussian simulations for radial velocity reconstruction:

1. **Determine fields for simulation:** The first step we take is to determine which fields need to be simulated "simultaneously", that is, from a single multivariate Gaussian distribution capturing all the crucial correlations. Ideally, all of the cosmological

fields we consider in this work should be simulated simultaneously. This can become a difficult computational task if we want to simulate the Π-binned moments of various fields in many redshift bins, which translates into large covariance matrices with non-vanishing off-diagonal terms and many high-resolution maps. The smaller the covariance matrix, the more likely it is that numerical errors can be avoided, so there is good motivation to be as economical as possible. We can ask ourselves, for example, which fields are necessary for a simulation of the kSZ signal. Certainly, joint simulations of the Π-binned moments of the radial velocity v, the differential optical depth $\dot{\tau}$, and the galaxy fields g are necessary if we want to ensure that the kSZ-g cross-correlation has the correct statistical anisotropy. Having identified these completely necessary fields for velocity reconstruction, we can ask ourselves if the fields that source other forms of temperature-galaxy statistical anisotropy should also be considered in the multivariate Gaussian distribution. The analysis of Sec.2.4 answers this question for us: the bias introduced by the non-linear ISW effect and the CMB lensing are negligible. This means that, for simulating radial velocity reconstruction, the non-linear ISW and the CMB lensing signals can be simply treated as "effective" sources of CMB anisotropies with no statistical correlation with the galaxy distribution. Finally, we ask ourselves if the fields that are isotropically correlated with the galaxy distribution need to be simulated together with v, $\dot{\tau}$ and g. These are the linear late-time ISW signal and the frequency cleaned extragalactic temperature foregrounds. The isotropic correlation between temperature and galaxies appears in the estimator weights Eq.(2.35). A quick inspection shows that, for our fiducial experimental noise levels, the relative difference in these weights when the small-angle ($\ell > 200$) temperature-galaxy cross-power is ignored is at most 3%. Thus, we consider it to be safe to ignore all isotropic correlations between temperature and galaxies in the reconstruction pipeline. Summarizing, we only need to generate simultaneous simulations of the Π-binned moments of v, $\dot{\tau}$, and g in order to capture all the important correlations for radial velocity reconstruction. The non-kSZ CMB anisotropies can be simulated separately from a single temperature spectra and later combined with the kSZ map from the joint simulations.

2. **Simulate the fields:** Once we determine which fields have to be simulated from a single multivariate Gaussian distribution, we construct a joint covariance matrix $\mathbf{C}_\ell$ for $0 < \ell \leq \ell_{\rm max}$ including all spectra and cross-spectra. Our fiducial resolution is $N_{\rm bin} = 32$ and $\ell_{\rm max} = 6144$ (corresponding to the band-limited multipole for a Healpix map of $N_{\rm side} = 2048$). The choice of number of bins and angular resolution are fixed by computational resources; it would be desirable to include more bins when possible to incorporate the effects of fine-mode noise. At each ℓ we find the Cholesky

decomposition $\mathbf{L}_\ell \mathbf{L}_\ell^\dagger = \mathbf{C}_\ell$ and generate data vectors $\mathbf{a}_{\ell m}$ corresponding to the $2\ell+1$ spherical harmonic coefficients at fixed ℓ of for all maps using the relation:

$$\mathbf{a}_{\ell m} = \mathbf{L}_\ell \cdot \mathbf{X}_{\ell m} \tag{2.165}$$

where $\mathbf{X}$ is a vector of random Gaussian numbers with zero mean and unit variance. In general we find good agreement between the ensemble-average spectra and cross-spectra from simulations and the input spectra. At $\ell \lesssim 30$ the built-in routine for generating Gaussian maps in Healpix (which generates realizations using a Cholesky decomposition as above) performs somewhat better than our algorithm. For large ℓ, the Healpix algorithm performs worse than ours. We therefore employ a hybrid method, generating low-ℓ moments using Healpix and high-ℓ moments using our algorithm. In either case, it is necessary to compute spectra and cross-spectra at sufficiently high accuracy to ensure that $\mathbf{C}_\ell$ is numerically positive definite and therefore the Cholesky decomposition well-defined. In this respect, our code for computing spectra is sufficiently accurate at the resolutions we have explored, but we expect increasingly accurate spectra are necessary for larger numbers of radial bins.

3. **Construct the temperature signals:** For radial velocity reconstruction, the kSZ signal is constructed from products of the simulated maps as:

$$\Theta^{kSZ}(\hat{\mathbf{n}}) = -\sum_\alpha \Delta\chi \; v^\alpha(\hat{\mathbf{n}}) \dot{\tau}^\alpha(\hat{\mathbf{n}}) \tag{2.166}$$

The part of the CMB temperature that we approximate as uncorrelated with galaxies is simulated as above, using a single temperature spectrum including: primary CMB, lensing contribution, linear and non-linear ISW contribution, and ILC cleaned extragalactic + instrumental noise components. The total temperature map is the sum of the uncorrelated map and the kSZ signal.

We apply a mask corresponding to an SO-like experiment, consisting of a cut between a declination of -70 degrees and +20 degrees and a Galactic mask that removes $\sim 30\%$ of the sky. The total sky fraction covered by the joint mask is $f_{\rm sky} = 0.45$.

4. **Run the estimator pipeline:** The galaxy maps and temperature maps generated this way are processed using the real space estimators described below.

2.5.2 Real-space estimators

The harmonic-space quadratic estimators for the radial velocity field cannot be implemented efficiently at the resolutions we wish to explore. We therefore derive mathematically

equivalent real-space estimators that take advantage of fast forward- and inverse-spherical harmonic transforms, which can be efficiently implemented (with a few caveats explored below).

Radial velocity estimator

To derive an efficient real-space estimator for the radial velocity, we first re-write $G^{\alpha}_{\ell\ell' L}$ as:

$$\begin{aligned} G^{\alpha}_{\ell\ell' L} &= \left(1 - \frac{(C_{\ell}^{g^{W^\alpha}\Theta})^2}{C_{\ell}^{\Theta\Theta} C_{\ell}^{g^{W^\alpha} g^{W^\alpha}}} \frac{(C_{\ell'}^{g^{W^\alpha}\Theta})^2}{C_{\ell'}^{\Theta\Theta} C_{\ell'}^{g^{W^\alpha} g^{W^\alpha}}}\right)^{-1} \\ &\times \left[\frac{f^{v^\alpha W^\alpha}_{\ell L \ell'}}{C_{\ell}^{\Theta\Theta} C_{\ell'}^{g^{W^\alpha} g^{W^\alpha}}} - \frac{(-1)^{\ell+\ell'+L} f^{v^\alpha W^\alpha}_{\ell' L \ell}}{C_{\ell'}^{\Theta\Theta} C_{\ell}^{g^{W^\alpha} g^{W^\alpha}}} \frac{C_{\ell}^{g^{W^\alpha}\Theta} C_{\ell'}^{g^{W^\alpha}\Theta}}{C_{\ell}^{\Theta\Theta} C_{\ell'}^{g^{W^\alpha} g^{W^\alpha}}}\right] \\ &= \left[\sum_{n=0}^{\infty} \left(\frac{(C_{\ell}^{g^{W^\alpha}\Theta})^2}{C_{\ell}^{\Theta\Theta} C_{\ell}^{g^{W^\alpha} g^{W^\alpha}}}\right)^n \left(\frac{(C_{\ell'}^{g^{W^\alpha}\Theta})^2}{C_{\ell'}^{\Theta\Theta} C_{\ell'}^{g^{W^\alpha} g^{W^\alpha}}}\right)^n\right] \\ &\times \left[\frac{f^{v^\alpha W^\alpha}_{\ell L \ell'}}{C_{\ell}^{\Theta\Theta} C_{\ell'}^{g^{W^\alpha} g^{W^\alpha}}} - \frac{(-1)^{\ell+\ell'+L} f^{v^\alpha W^\alpha}_{\ell' L \ell}}{C_{\ell'}^{\Theta\Theta} C_{\ell}^{g^{W^\alpha} g^{W^\alpha}}} \frac{C_{\ell}^{g^{W^\alpha}\Theta} C_{\ell'}^{g^{W^\alpha}\Theta}}{C_{\ell}^{\Theta\Theta} C_{\ell'}^{g^{W^\alpha} g^{W^\alpha}}}\right] \end{aligned} \tag{2.167}$$

Next, we use the relation

$$\begin{aligned} \int d^2\hat{n}\; Y_{\ell m} Y_{\ell' m'} Y^*_{LM} &= \sqrt{\frac{(2\ell+1)(2\ell_1+1)(2\ell'+1)}{4\pi}} (-1)^M \\ &\times \begin{pmatrix} \ell & \ell' & L \\ m & m' & -M \end{pmatrix} \begin{pmatrix} \ell & \ell' & L \\ 0 & 0 & 0 \end{pmatrix} \end{aligned} \tag{2.168}$$

and the definitions:

$$\begin{aligned} \xi^{\alpha}_n(\hat{\mathbf{n}}) &= \sum_{\ell m} \left(\frac{(C_{\ell}^{g^{W^\alpha}\Theta})^2}{C_{\ell}^{\Theta\Theta} C_{\ell}^{g^{W^\alpha} g^{W^\alpha}}}\right)^n \frac{\Theta_{\ell m}}{C_{\ell}^{\Theta\Theta}}\; Y_{\ell m}(\hat{\mathbf{n}}) \\ \zeta^{\alpha}_n(\hat{\mathbf{n}}) &= \sum_{\ell m} \left(\frac{(C_{\ell}^{g^{W^\alpha}\Theta})^2}{C_{\ell}^{\Theta\Theta} C_{\ell}^{g^{W^\alpha} g^{W^\alpha}}}\right)^n \frac{C_{\ell}^{\dot{\tau}^\alpha g^{W^\alpha}} g^{W^\alpha}_{\ell m}}{C_{\ell}^{g^{W^\alpha} g^{W^\alpha}}}\; Y_{\ell m}(\hat{\mathbf{n}}) \end{aligned} \tag{2.169}$$

$$\begin{aligned} \bar{\xi}^{\alpha}_n(\hat{\mathbf{n}}) &= \sum_{\ell m} \left(\frac{(C_{\ell}^{g^{W^\alpha}\Theta})^2}{C_{\ell}^{\Theta\Theta} C_{\ell}^{g^{W^\alpha} g^{W^\alpha}}}\right)^{n+1} \frac{C_{\ell}^{\dot{\tau}^\alpha g^{W^\alpha}} \Theta_{\ell m}}{C_{\ell}^{g^{W^\alpha}\Theta}}\; Y_{\ell m}(\hat{\mathbf{n}}) \\ \bar{\zeta}^{\alpha}_n(\hat{\mathbf{n}}) &= \sum_{\ell m} \left(\frac{(C_{\ell}^{g^{W^\alpha}\Theta})^2}{C_{\ell}^{\Theta\Theta} C_{\ell}^{g^{W^\alpha} g^{W^\alpha}}}\right)^{n+1} \frac{g^{W^\alpha}_{\ell m}}{C_{\ell}^{g^{W^\alpha}\Theta}}\; Y_{\ell m}(\hat{\mathbf{n}}). \end{aligned} \tag{2.170}$$

The real-space estimator is given by:

$$d^{\alpha}_{LM} = N^{\alpha}_{L} \sum_{n=0}^{\infty} \int d^2\hat{n}\; Y^{*}_{LM}(\hat{\mathbf{n}}) \left[\xi^{\alpha}_{n}(\hat{\mathbf{n}})\zeta^{\alpha}_{n}(\hat{\mathbf{n}}) - \bar{\xi}^{\alpha}_{n}(\hat{\mathbf{n}})\bar{\zeta}^{\alpha}_{n}(\hat{\mathbf{n}})\right] \tag{2.171}$$

For the simulations presented below, where we effectively set $C_{\ell}^{g^{W^{\alpha}}\Theta} = 0$ by not including the statistically isotropic correlations between the galaxy and temperature fields (as argued above, these contributions are insignificant for our fiducial CMB experiment), and therefore we can work at the $n = 0$ level. For different experimental configurations, it may become necessary to consider higher order terms in the sum over a number of n. So long as $(C_{\ell}^{g^{W^{\alpha}}\Theta})^2 \ll C_{\ell}^{\Theta\Theta} C_{\ell}^{g^{W^{\alpha}} g^{W^{\alpha}}}$, it is sufficient to consider the $n = 0$ and $n = 1$ terms in the sum only.

Note that the fields $\xi^{\alpha}_{n}(\hat{\mathbf{n}})$, $\zeta^{\alpha}_{n}(\hat{\mathbf{n}})$, $\bar{\xi}^{\alpha}_{n}(\hat{\mathbf{n}})$ and $\bar{\zeta}^{\alpha}_{n}(\hat{\mathbf{n}})$ are convolutions of an azimuthally symmetric function and the moments of the CMB and galaxy density maps. In pixel-space, we can therefore write e.g.:

$$\xi^{\alpha}_{n}(\hat{\mathbf{n}}) = \int d^2\hat{n}'\; B^{\xi^{\alpha}_{n}}(|\hat{\mathbf{n}} - \hat{\mathbf{n}}'|)\Theta(\hat{\mathbf{n}}') \tag{2.172}$$

where the "beam" $B^{\xi^{\alpha}_{n}}(|\hat{\mathbf{n}} - \hat{\mathbf{n}}'|)$ for the field $\xi^{\alpha}_{n}(\hat{\mathbf{n}})$, and the beams for the other filtered fields, are given by:

$$B^{\xi^{\alpha}_{n}}(\theta) = \sum_{\ell} \sqrt{\frac{2\ell+1}{4\pi}} \left(\frac{(C_{\ell}^{g^{W^{\alpha}}\Theta})^2}{C_{\ell}^{\Theta\Theta} C_{\ell}^{g^{W^{\alpha}} g^{W^{\alpha}}}}\right)^{n} \frac{1}{C_{\ell}^{\Theta\Theta}} P_{\ell}(\cos\theta)\,, \tag{2.173}$$

$$B^{\zeta^{\alpha}_{n}}(\theta) = \sum_{\ell} \sqrt{\frac{2\ell+1}{4\pi}} \left(\frac{(C_{\ell}^{g^{W^{\alpha}}\Theta})^2}{C_{\ell}^{\Theta\Theta} C_{\ell}^{g^{W^{\alpha}} g^{W^{\alpha}}}}\right)^{n} \frac{C_{\ell}^{\dot{\tau}^{\alpha} g^{W^{\alpha}}}}{C_{\ell}^{g^{W^{\alpha}} g^{W^{\alpha}}}} P_{\ell}(\cos\theta)\,, \tag{2.174}$$

and

$$B^{\bar{\xi}^{\alpha}_{n}}(\theta) = \sum_{\ell} \sqrt{\frac{2\ell+1}{4\pi}} \left(\frac{(C_{\ell}^{g^{W^{\alpha}}\Theta})^2}{C_{\ell}^{\Theta\Theta} C_{\ell}^{g^{W^{\alpha}} g^{W^{\alpha}}}}\right)^{n+1} \frac{C_{\ell}^{\dot{\tau}^{\alpha} g^{W^{\alpha}}}}{C_{\ell}^{g^{W^{\alpha}}\Theta}} P_{\ell}(\cos\theta)\,, \tag{2.175}$$

$$B^{\bar{\zeta}^{\alpha}_{n}}(\theta) = \sum_{\ell} \sqrt{\frac{2\ell+1}{4\pi}} \left(\frac{(C_{\ell}^{g^{W^{\alpha}}\Theta})^2}{C_{\ell}^{\Theta\Theta} C_{\ell}^{g^{W^{\alpha}} g^{W^{\alpha}}}}\right)^{n+1} \frac{1}{C_{\ell}^{g^{W^{\alpha}}\Theta}} P_{\ell}(\cos\theta)\,. \tag{2.176}$$

Some insight into the map-based properties of the estimator can be gained by examining the shape of these functions, which we plot in Fig. 2.18 for the redshift bin $z = (0.44, 0.50)$

and in the $n = 0$ case. The beams receive support only over a scale of ~ 4 arcmin for the experimental parameters considered here, corresponding to ~ 3 pixels at Healpix resolution $N_{\rm side} = 2048$. This implies that the quadratic estimator is highly local, and that systematic errors due mixing information from masked or contaminated regions of the sky will be minimal. Unlike the case of CMB lensing (see e.g. [47]), we therefore expect that there is only a very small bias from the mask on the reconstructed velocity field.

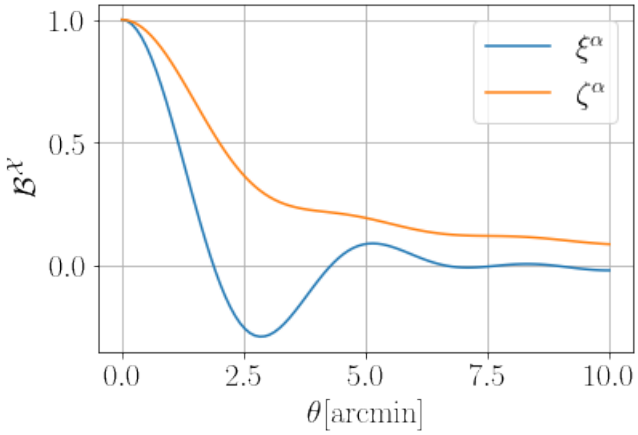

Figure 2.18: Different beams defined in Sections 2.5.2, normalised at $\theta \ll 1$, for bin α corresponding to redshift range $z = (0.44, 0.50)$. The beams receive support only over a scale of few arcmins and highly localised with enhanced scaling (equivalently blue spectrum) on smallest scales corresponding to $\lesssim 2$ arcmin, similar to the the pixel resolution ~ 1.7 armin at Healpix resolution $N_{\rm side} = 2048$.

2.5.3 Reconstruction on simulated maps

We now present the results obtained by applying the real-space estimators derived above to the simulated maps for the fiducial data combination of SO x Rubin employed in previous sections. We generated a set of 30 realizations, each with 32 bins in the redshift range $0.2 \leq z \leq 5$ and output resolutions of NSIDE 2048. The resolution and number of simulations were dictated by available computational resources. Note that our reconstructions will not include fine-mode noise, since the simulations are constructed using a limited number of bins. We show examples of the reconstruction in Fig.(2.19), where we compare the rotated true velocities $\mathbf{R}_L \cdot \mathbf{v_{LM}}$ to the output of the estimator $\hat{\mathbf{v}}_{\mathbf{LM}}$ at two representative redshift bins located at $z \sim 0.5$ and $z \sim 1.5$. These maps have been filtered to show only the largest

angular scales, where the reconstruction is signal-dominated. A visual inspection of these maps indicates that a successful reconstruction of the radial velocity has been achieved. A comparison of the angular power spectra for the reconstruction and the masked actual velocity field indicate good quantitative agreement. Before undertaking a full quantitative analysis of the full set of realizations, and comparing to theoretical expectations, we take a brief digression to discuss the effects of the mask on the reconstruction.

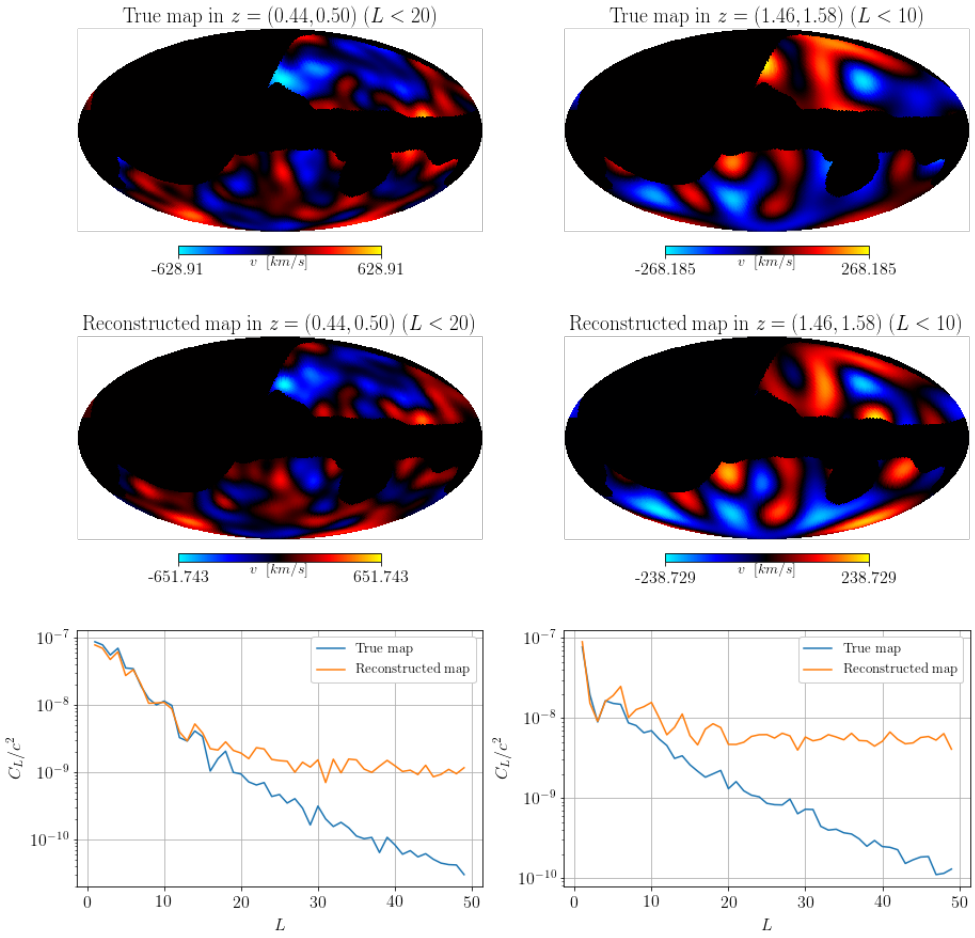

Figure 2.19: **Top panels**: Low-multipole filtered maps of the true rotated velocities. **Middle panels**: Low-multipole filtered reconstructed maps. **Bottom panels**: power spectra comparison between true and reconstructed maps.

As we discussed in Sec. 2.5.2, we expect contamination from masked regions to extend only a few pixels from the mask boundary due to the local nature of the radial velocity

estimator. We can corroborate this by fixing the realization and subtracting the full-sky reconstruction noise (defined as the reconstruction minus the actual radial velocity) from the masked-sky reconstruction noise and studying the residuals, which we refer to as the mask bias. We can see from Fig. 2.20 that the dominant effect of mask in concentrated on the edge of the unmasked region as expected, and that this contamination can be removed post-reconstruction by extending the mask by a few pixels. As can be seen by comparing the maps in the top and middle panels, extending the mask by a single pixel (at $N_{\rm side} = 64$) removes the majority of contaminated pixels. In the bottom panel of Fig. 2.20, we see that the mask bias is always below the reconstruction noise, and that extending the mask by one pixel decreases the mask bias by orders of magnitude. We conclude that one need not worry about mask bias for reconstruction of the radial velocity. Note that the story presented here will become more complicated for apodized maps, since apodization will introduce a statistical anisotropy that may be picked up by the estimator and which must be accounted for in the reconstruction.

Returning to our ensemble of simulations, we now confirm that the statistics of the ensemble are as expected from the analytic estimates presented in Sec.2.4. On the full sky this can be thought of as a validation exercise for our simulations and reconstruction pipeline, since in the absence of numerical errors, the agreement should be perfect. On the masked sky, we can determine what the effect of the mask is on the reconstructed power spectrum. To mitigate the mask bias, we extend the mask post-reconstruction by one pixel in the results presented below.

The top panel of Fig. 2.21 compares the ensemble average reconstruction signal and noise to the theoretical expectations. Comparing the theory signal (blue) to the simulated signal (green dashed), there is excellent agreement in both redshift bins. Comparing the power spectrum of the reconstruction (red dot-dashed) to the theory signal and the theory noise (orange), we see excellent agreement in the signal-dominated and noise-dominated regimes as well. We also compute the power spectrum of the reconstruction minus the actual signal (purple dots) and the result of applying the estimator to a temperature map whose kSZ component is uncorrelated with the galaxy density (brown squares). In both cases, the agreement with the theory reconstruction noise is excellent, aside from some excess in the power spectrum of the reconstruction minus the theory signal at low-L, at the level of about one percent of the signal. This result is a powerful validation of our simulation and reconstruction pipeline.

The bottom panel of Fig. 2.21 shows the comparison between the theory and ensemble-averaged reconstruction on the masked sky. We make no attempt here to de-project the mask, and simply find the power spectra of the masked maps. The theory curves (signal and noise) have been multiplied by $f_{\rm sky} = 0.44$ to account for the loss of variance from masking.

Comparing the theory signal multiplied by $f_{\rm sky}$ (blue) to the simulated masked signal (green dashed) the scale-dependent effect of the mode-coupling with the mask is evident, especially at low redshift. Nevertheless the factor of $f_{\rm sky}$ gives a reasonable estimate of the power spectrum of the velocity field on the masked sky. Comparing the reconstruction on the masked sky to the masked actual velocity and the theory noise reduced by $f_{\rm sky}$, we see good agreement at both high- and low-L. Checking this in more detail by finding the difference between the reconstruction and the actual masked signal (purple dots), we again find good agreement with the expected reconstruction noise aside from a few percent excess at low-L and low redshift.

The collection of reconstructed maps are not statistically independent of each other due to velocity correlations along the lightcone. On the full-sky, harmonic moments with different multipoles L, M are independent from each other and only correlated in the radial direction. Using the principal component decomposition described in Sec. 2.2.5, we can construct N_{bin} linear combinations of Π-binned harmonic multipoles $a^j_{LM} = \sum_\alpha c_L^{j\beta} a^\beta_{LM}$ such that a^j_{LM} and a^i_{LM} are uncorrelated for all $i \neq j$. Note that the transformation coefficients depend on the multipole L. The maps constructed using these rotated harmonic moments at each L constitute the principal components of the radial velocity on the lightcone. The power spectrum of the principal components at each multipole L is diagonal and is obtained by rotating the bin basis power spectrum $\mathbf{c}_L \mathbf{C}_L (\mathbf{c}_L)^\dagger$, where $\mathbf{c}_L$ is the matrix defined by the coefficients $c_L^{j\beta}$. Fig.(2.22) shows the true maps, reconstructed maps and spectra for the 2 highest signal to noise principal components from a single realization. The 'unit' of the principle component power spectra and maps is signal to noise, since the noise has been normalized to unity. Clearly, the fidelity of the principle component reconstruction is far higher than for the single bins presented in Fig. 2.19 (although the information in the set of principle component maps is equivalent to the information in the set of bin maps), making the principle component basis desirable for a visual representation of the results. In the top panel of Fig. 2.23, we show the ensemble-averaged power spectra of the first two principle components without masking. Performing the comparisons described above in the bin basis between theoretical expectations and data from the reconstructions, we again find excellent agreement, aside from sub-percent level effects at the lowest L.

The principal component analysis discussed here is subject to some complications when a mask is introduced. Statistical isotropy is broken for the masked map, which introduces $L \neq L'$ correlations at each bin and between bins. Using the $c_L^{j\beta}$ coefficients defined for the full-sky scenario would not lead to statistically independent maps. A more rigorous procedure to find the uncorrelated combinations of the data could be done in pixel space rather than harmonic space. Such a procedure would involve the construction and diagonalization of a huge matrix containing the covariance between every pair of pixels at

every pair of redshift bins, a task that is computationally demanding. Here, we consider a "pseudo" principal component transformation at the level of the power spectra of the reconstructions with mask. As we discussed in the previous subsection, the power spectrum of reconstructed maps with a mask traces the underlying full-sky spectra up to a factor of f_{sky} and some scale-dependent corrections due to convolution with the mask. For multipoles where the scale dependent correction is small at all redshifts, the masked C_L is approximately proportional to the unmasked one and therefore can be diagonalized using the full-sky transformation matrix $\mathbf{c}_L$. In the bottom panel of Fig. 2.23 we compare the theory signal and noise (in this case, the noise is the identity matrix) reduced by a factor of $f_{\rm sky}$ to the actual masked velocity and reconstructed velocity for the first two principal components. Despite the complications from mode coupling with the mask, there is reasonable agreement with the theory curves. Finally, we can explicitly check that the rotation associated with the principal components $\mathbf{c}_L$ take the reconstructed spectra $\left\langle \hat{\mathbf{v}}_{\mathbf{LM}} \hat{\mathbf{v}}^{\dagger}_{\mathbf{LM}} \right\rangle$ to a close-to diagonal form. This is shown in Fig. 2.24. where we plot $\left\langle \hat{\mathbf{v}}_{\mathbf{LM}} \hat{\mathbf{v}}^{\dagger}_{\mathbf{LM}} \right\rangle$ averaged over 30 masked reconstructions. The rotation associated with the pseudo principal component basis does indeed result in a more diagonal signal covariance matrix even in the presence of a mask.

2.6 Conclusions

This chapter has outlined the formalism for velocity reconstruction in the Lightcone Picture using CMB experiments and galaxy surveys. One of the main goals of developing this formalism has been to explore some of the challenges posed by systematics and foregrounds for velocity reconstruction. The range of effects we have explored include: properly correlated extragalactic foregrounds, large-angular scale systematics in the galaxy survey, photometric redshift errors, masking of regions contaminated by galactic emission, modelling errors in the galaxy-electron correlation function (the optical depth degeneracy), biases introduced due to additional physical effects that lead to a statistically anisotropic CMB-galaxy correlation (e.g. lensing of the primary CMB), biases introduced by coarse graining on the lightcone (e.g. the 'fine mode' noise), CMB instrumental noise and beam, choice of frequency channels for cleaning extragalactic foregrounds in the CMB, and the effect of performing foreground cleaning on reconstructed maps. We have developed a numerical pipeline to compute the properly correlated auto and cross-spectra necessary to assess this range of effects. We have also developed a real-space reconstruction pipeline that we have validated using Gaussian simulations. This pipeline was used to assess the

impact of systematics in real-space such as masking. The good news is that none of the systematic effects we have explored seriously degrade the fidelity of the reconstruction, indicating the promise of velocity reconstruction for extracting new cosmological information from future datasets.

Our fiducial datasets were a Rubin-like galaxy survey and an SO-like CMB experiment. We also considered the data combination of the existing unWISE galaxy catalogue with SO. These choices determine factors such as: redshift error, depth of the survey, galaxy shot noise, the level of large-angular scale systematics, frequency channels assumed for the CMB experiment, the associated level of instrumental noise and resolution, and sky coverage. For these datasets, some of the take-away points of our analysis include:

- The total information available in the reconstructed velocity fields (quantified by the total signal to noise) is limited mainly by the redshift error, sky coverage, and factors contributing to the Gaussian reconstruction noise (CMB instrumental noise and beam, foreground residuals, and the level of galaxy shot noise).

- It is essential to incorporate the "fine mode" noise associated with coarse-graining fields on the lightcone into the estimator formalism for velocity reconstruction. This source of bias can be mitigated by ensuring that velocity reconstruction is performed in a sufficient number of bins along the radial direction. For radial velocity reconstruction in the fiducial SO x Rubin scenario, it is necessary to use around 64 bins to mitigate this bias and include most of the signal to noise in the reconstruction.

- For radial velocity reconstruction, large-scale systematics in the galaxy survey have a significant ($\sim$ 10%-level) effect on the total signal to noise due to the additional bin-bin correlations it introduces.

- The biases induced by CMB lensing and non-linear ISW on radial velocity reconstruction are negligible.

- From the principle components of the radial velocity reconstructions (see Figs. 2.16), most of the signal to noise in the reconstruction comes from large angular scales $L \lesssim 30$ and redshifts $z \lesssim 1$.

- For SO x unWISE, it is possible to reconstruct a single principal component on the very largest angular scales that roughly follows the unWISE redshift distribution. This demonstrates that even for a single broad photometric redshift bin it is possible to reconstruct the large-scale radial velocity field.

- The real space estimators for velocity reconstruction are highly local, and contamination from masking is restricted to the region in close proximity to the mask. It is therefore possible to remove the mask bias by extending the mask by a few pixels post-reconstruction.

- We have demonstrated a pipeline for velocity reconstruction which performs as expected for Gaussian simulations. This pipeline can serve as a prototype for analysis of future datasets.

A fundamental assumption of the quadratic estimator formalism explored in this chapter is that the underlying fields are Gaussian. On the small-scales that contribute the most to the quadratic estimators, this assumption is far from accurate. One consequence is the presence of the '$N^{3/2}$ bias', explored by Ref. [21] in the Box Picture. We have not performed an analysis of this contribution in the Lightcone Picture, although a similar computation based on the halo model could in principle be performed (albeit with complex projection integrals to contend with). This is a necessary ingredient for using velocity reconstruction in the Lightcone Picture to measure and constrain cosmological parameters. Aside from this bias Refs. [21] demonstrated that velocity reconstruction in the Box Picture essentially works as advertised even for non-linear N-body simulations. Chapter 3 addresses the impact of non-linearities in the Ligthcone Picture using N-body simulations. An important future analysis will be to directly compare the importance of the various systematic effects discussed in this chapter between Gaussian and non-linear N-body simulated datasets.

In the future, our framework can be extended to assess velocity reconstruction using different tracers such as the CIB [79] or intensity maps [80]. Other extensions include velocity reconstruction using reionization kSZ and 21cm maps [82] or reconstruction of the remote quadrupole field [36, 19]. Having a unifying framework, or at least a unifying basis, allows one to combine the cosmological information from these various probes. Examples where this may be important include constraints on modified gravity [77] and various early-Universe scenarios (see Chapter 4). The reconstruction tools described in this chapter can be found in the publicly available code ReCCO. This code calculates angular spectra for a variety of fields on the lightcone and uses them to compute the numerous signals, noises and biases described in this chapter. We expect this code to become a useful tool for the future of kSZ velocity reconstruction and other methods involving CMB secondaries.

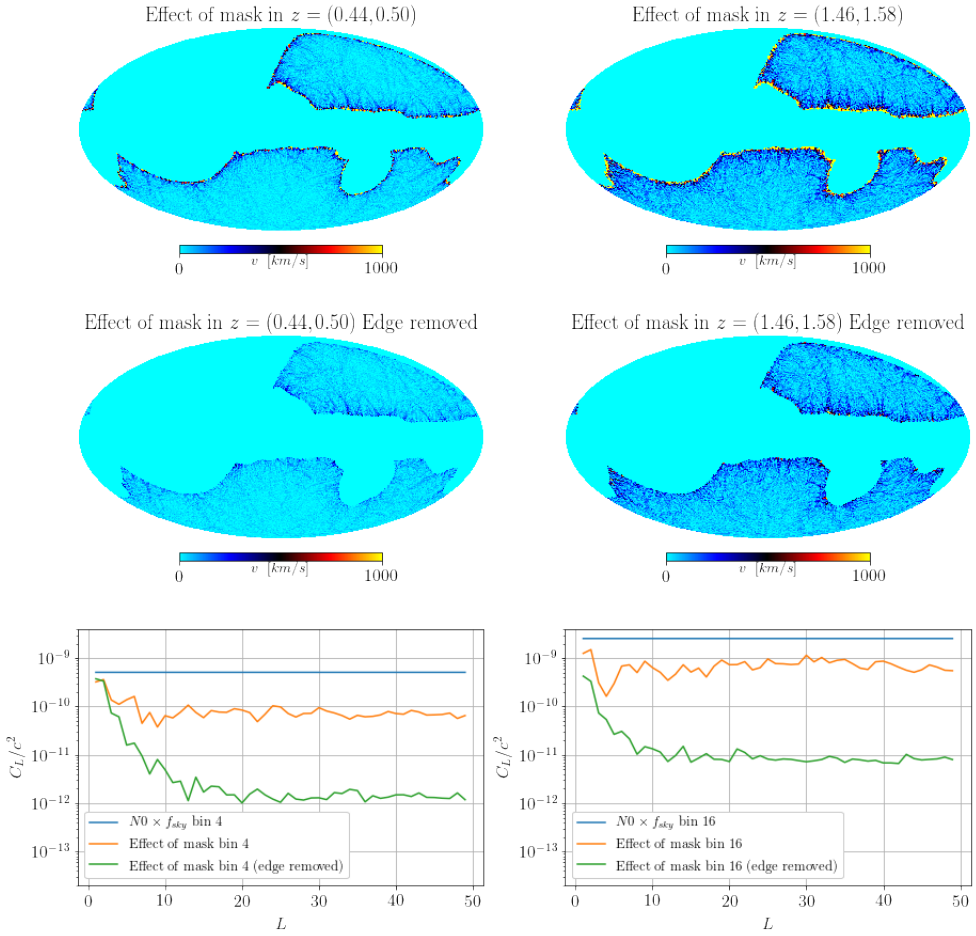

Figure 2.20: **Top panels**: Mask bias at Nside 64 when no edge pixels are removed. **Middle panels**: mask bias after removing a 1-pixel thick border at Nside = 64, which leads to a reduction of f_{sky} from 0.45 to 0.44. **Bottom panels**: power spectrum of the mask bias with and without an extended mask compared to the N^0 noise for radial velocity reconstruction.

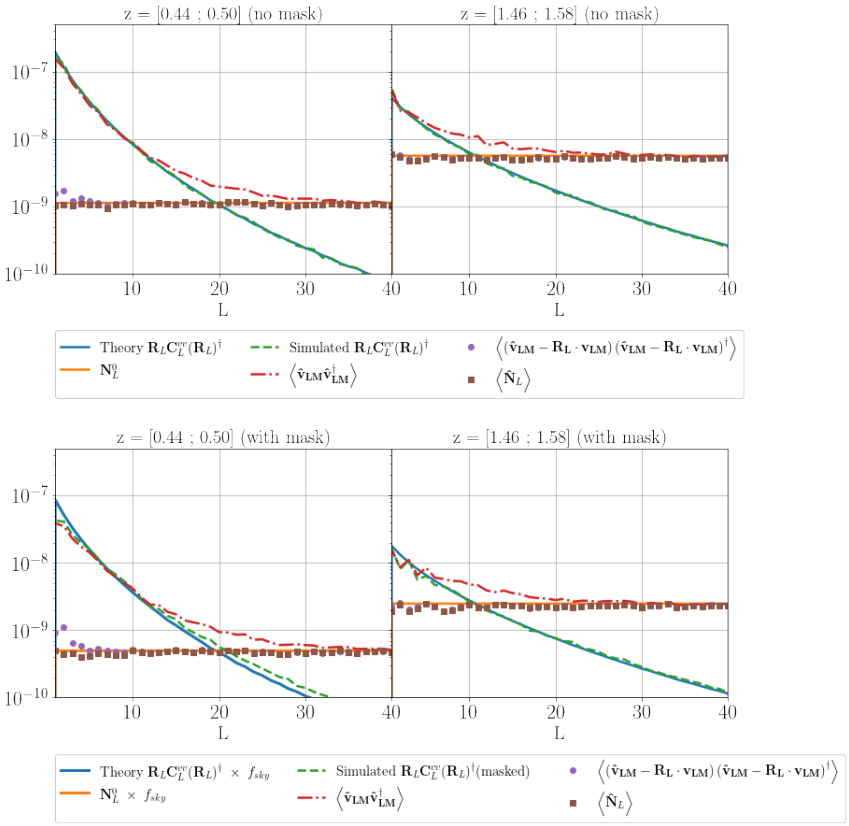

Figure 2.21: Average power spectrum from 20 simulated reconstructions on the full-sky (top panels) and on masked-sky (bottom panels). Solid lines correspond to signal and noise from theory, dashed lines show simulated radial velocities (rotated by $\mathbf{R}_L$), dot-dashed lines shows the estimator output, circles show the difference between the estimator output and the simulated rotated velocities, and squares shows the estimator noise.

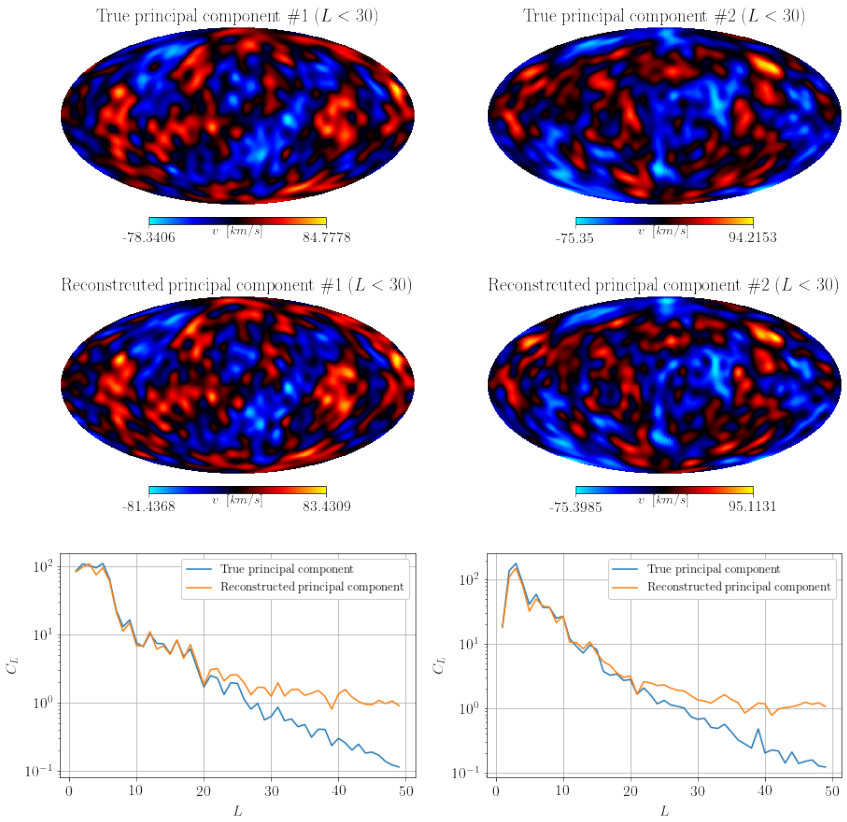

Figure 2.22: **Top panels**: Low-multipole filtered maps of the 2 highest signal to noise true principal components on the full sky. **Middle panels**: Low-multipole filtered reconstructed principal components. **Bottom panels**: power spectra comparison between true and reconstructed maps.

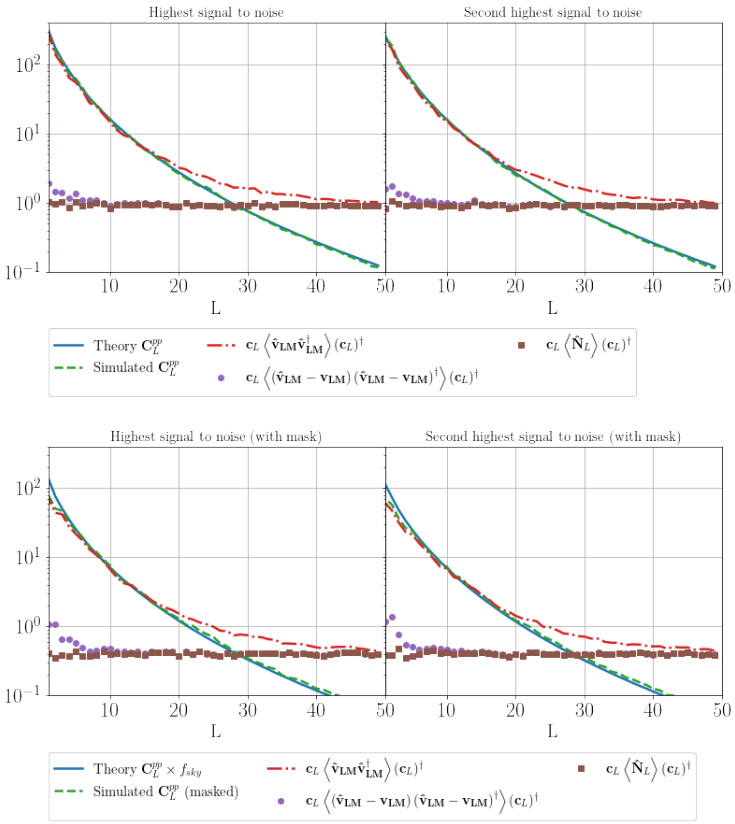

Figure 2.23: Signal and noise averages for the pseudo principal component transformation of reconstructed spectra on the full-sky (top panels) and masked sky (bottom panels).

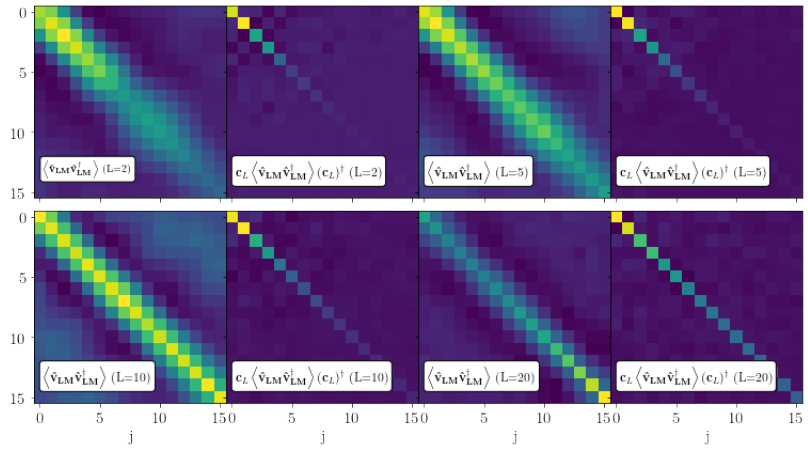

Figure 2.24: Pseudo principal component transformation of the reconstructed spectra $\left\langle \hat{\mathbf{v}}_{\mathbf{LM}} \hat{\mathbf{v}}^{\dagger}_{\mathbf{LM}} \right\rangle$ at different L multipoles, averaged over 30 simulations.

Chapter 3

Simulated reconstruction of the remote dipole

3.1 Chapter summary

In Chapter 2, a series of extensions and additional layers of realism were incorporated into the Ligthcone Picture formalism of kSZ velocity reconstruction. The quadratic estimator for the radial velocity field was validated using a numerical framework to produce sets of properly correlated CMB maps and redshift binned galaxy number counts, assuming Gaussian underlying fields. Despite being practical for testing many aspects of the reconstruction machinery, these mock datasets do not inform us about the robustness of kSZ velocity reconstruction in the presence of non-linear evolution of structure. Moreover, the simulation pipeline from Chapter 2 is somewhat agnostic of the remote dipole power spectrum on large angular scales: $C_L^{v^\alpha v^\beta}$ is simply an input that we aim to reconstruct using small-angle CMB and galaxy power spectra; any reasonable model for the low-L $C_L^{v^\alpha v^\beta}$ would lead to a successful reconstruction as long as the CMB-galaxy correlation shows a kSZ induced statistical anisotropy consistent with a remote dipole field sampled from the input model. Indeed, the reconstruction presented in Chapter 2 assumed for simplicity a remote dipole sourced only by the local Doppler term and ignored the stationary dipole contributions which become important at the lowest multipoles (see Sec.1.5 for a discussion of the many contributions to the remote CMB dipole).

This chapter, which corresponds to a study done before the development of the extended formalism of Chapter 2, explores the previously neglected effects of gravitational non-linearities and demonstrates the relevance of the stationary dipole contributions to the

remote CMB dipole using a set of mock CMB and LSS maps generated from a suite of N-body simulations. Additionally, such simulations allow us to explore the effects of redshift space distortions and CMB lensing (predicted to be negligible in Chapter 2). Simulating kSZ velocity reconstruction on the light cone while properly treating all of the relevant physics is intrinsically challenging. One must model both the dipole field, which receives important contributions from scales of order the size of the observable Universe, as well as the density fields, which depend on small scales and include baryonic physics. Evolving a standard N-body simulation incorporating such a large range of scales is currently computationally intractable. To overcome this limitation, we develop a novel box-in-box simulation framework which consistently embeds a ∼Gpc-sized N-body simulation inside of a box whose volume is larger than the observable Universe, and which contains large-scale modes evolved using linear theory. This box-in-box procedure allows us to use the data from both of these simulations to produce properly correlated maps of the lensed primary CMB temperature anisotropies, kSZ temperature anisotropies, the dipole field, and the dark matter overdensity field. We do not model baryonic physics in the present analysis, and therefore use the dark matter density as a proxy for the electron density. N-body simulations including baryonic physics will be incorporated in the future.

We find that the quadratic estimator efficiently reconstructs the remote dipole field over a range of angular scales and redshift bins, indicating that kSZ tomography is generally robust. However, we present evidence for a bias due to non-linear structure at low redshifts. We demonstrate the ability of kSZ tomography to reconstruct the fundamental component of the observed CMB dipole, supporting the suggestion in previous work [19] that this could be an early application of kSZ tomography on large angular scales. The results we present here are intended primarily as a proof-of-principle both of the simulation framework and remote dipole reconstruction. To lay the groundwork for the analysis of near-term datasets, various layers of realism will be added to our simulation framework in future work, including the construction of mock galaxy catalogs, improved resolution, inclusion of baryonic physics, correlated foregrounds such as thermal SZ, and partial sky data. Some of these aspects have already been explored in Chapter 2.

3.2 A brief review

In this section we briefly review the expressions for the kSZ signal, the remote dipole field, its theoretical power spectrum, and the real-space quadratic estimator that we use to reconstruct the dipole field. The kSZ effect can be expressed as the line of sight integral:

$$\Theta^{kSZ}(\hat{\mathbf{n}}) = -\int d\chi\; v_{\rm eff}(\hat{\mathbf{n}},\chi)\;\dot{\tau}(\hat{\mathbf{n}},\chi) \tag{3.1}$$

where $\dot{\tau}(\hat{\mathbf{n}},\chi)$ is the differential optical depth defined in Eq.(1.11) and $v_{\rm eff}(\hat{\mathbf{n}},\chi)$ is the remote CMB dipole field defined as

$$v_{\rm eff}(\hat{\mathbf{n}},\chi) = \frac{3}{4\pi}\int d^2\hat{n}_e\;\Theta(\mathbf{x},\hat{\mathbf{n}_e})\;(\hat{\mathbf{n}}\cdot\hat{\mathbf{n}_e}). \tag{3.2}$$

where $\Theta(\mathbf{x},\hat{\mathbf{n}_e})$ is the remote CMB at the spacetime position $(\eta(\chi),\mathbf{x}=\chi\hat{\mathbf{n}})$, which receives contributions from the Sachs-Wolfe effect, the integrated Sachs-Wolfe effect due to the evolution of the gravitational potential along the line of sight, and the Doppler effect:

$$\Theta(\mathbf{x},\hat{\mathbf{n}_e}) = \Theta^{SW}(\mathbf{x},\hat{\mathbf{n}_e}) + \Theta^{ISW}(\mathbf{x},\hat{\mathbf{n}_e}) + \Theta^{\mathcal{D}}(\mathbf{x},\hat{\mathbf{n}_e}) \tag{3.3}$$

In order to work with a binned version of Eq. 3.1, we consider a bin-averaged (Π-binned as defined in Sec.2.2.2) remote dipole field:

$$v^{\alpha}_{\rm eff}(\hat{\mathbf{n}}) = \frac{3}{4\pi}\frac{1}{\Delta\chi_\alpha}\int_{\chi^\alpha_{min}}^{\chi^\alpha_{\rm max}} d\chi\int d^2\hat{n}_e\;\Theta(\mathbf{x},\hat{\mathbf{n}_e})\;(\hat{\mathbf{n}}\cdot\hat{\mathbf{n}_e}), \tag{3.4}$$

Here, the index α labels each bin, which extend over the range in comoving distance $\chi^\alpha_{min}<\chi<\chi^\alpha_{max}$, and where $\Delta\chi_\alpha = \chi^\alpha_{max}-\chi^\alpha_{min}$. The binned power spectrum is given by

$$C_L^{v^\alpha v^\beta} = \int \frac{d^3k}{(2\pi)^3} P_{\Psi_p}(k)\Delta_L^{v^\alpha *}(k)\Delta_L^{v^\beta}(k), \tag{3.5}$$

where Greek indices denote redshift bins, $P_{\Psi_p}(k)$ is the power spectrum of the Newtonian gauge primordial gravitational potential Ψ_p, and $\Delta_L^{v^\alpha}(k)$ is bin averaged remote dipole transfer function (non-averaged function defined in Eq.(1.49)). As shown in [19], the presence of a large scale dipole will manifest in the cross correlation between the kSZ contribution to the CMB temperature and the moments of a redshift binned density distribution δ^α defined by

$$\delta^\alpha(\hat{\mathbf{n}}) = \frac{1}{\Delta\chi_\alpha}\int_{\chi^\alpha_{min}}^{\chi^\alpha_{\rm max}} d\chi\;\delta(\hat{\mathbf{n}},\chi). \tag{3.6}$$

We use the real-space optimal quadratic estimator for the Π-binned LC moments of the remote dipole field defined in Sec.2.5.2. At leading order, the real space estimator is given by:

$$\widehat{v}^\alpha_{{\rm eff},LM} = \left(\mathbf{N}^0_L\right)^{\alpha\alpha}\int d^2\hat{n}\;Y^*_{\ell m}(\hat{n})\xi(\hat{n})\zeta^\alpha(\hat{n}). \tag{3.7}$$

$$\xi(\hat{n}) = \sum_{\ell m} \frac{\Theta_{\ell m}}{C_\ell^{\Theta\Theta}} Y_{\ell m}(\hat{n}) \tag{3.8}$$

$$\zeta^\alpha(\hat{n}) = \sum_{\ell m} \frac{C_\ell^{\dot{\tau}^\alpha \delta^\alpha} \delta^\alpha_{\ell m}}{C_\ell^{\delta^\alpha \delta^\alpha}} Y_{\ell m}(\hat{n}) \tag{3.9}$$

where $(\mathbf{N}_L^0)^{\alpha\alpha}$ are the diagonal elements of the Gaussian reconstruction noise given by Eq.(2.34), and $\dot{\tau}^\alpha$ denotes the Π-binned differential optical depth. We remind the reader that this form for the real space quadratic estimators is obtained when the statistically isotropic correlations between the density tracer and temperature fields $C_\ell^{\delta^\alpha \Theta}$ is not included. This approximation is valid as long as $(C_\ell^{\delta^\alpha \Theta})^2 \ll C_\ell^{\Theta\Theta} C_\ell^{\delta^\alpha \delta^\alpha}$ on small angular scales, which we found to be true to the percent level in Chapter 2 in the presence of cleaned extragalactic foregrounds. For the simulations below, which do not include such foregrounds, the approximation is expected to be even more accurate. With this approximation, the diagonal elements of the reconstruction noise are given by:

$$\left(\mathbf{N}_L^0\right)^{\alpha\alpha} = (2L+1) \left(\sum_{l_1 l_2} \frac{f^{v^\alpha}_{\ell_1 L \ell_2} f^{v^\alpha}_{\ell_1 L \ell_2}}{C_{\ell_1}^{\Theta\Theta} C_{\ell_2}^{\delta^\alpha \delta^\alpha}} \right)^{-1}. \tag{3.10}$$

where

$$f^{v^\alpha}_{\ell \ell_1 \ell'} \equiv \sqrt{\frac{(2\ell+1)(2\ell_1+1)(2\ell'+1)}{4\pi}} \begin{pmatrix} \ell & \ell' & \ell_1 \\ 0 & 0 & 0 \end{pmatrix} C_{\ell'}^{\dot{\tau}^\alpha g^W} \Delta\chi_\alpha. \tag{3.11}$$

are the f-couplings defined in Eq.(2.27).

Because the simulations presented below do not contain baryons, we assume that the electron density field traces the dark matter density field.

3.3 Simulations

Our simulation framework includes two components: a small-scale N-body simulation and a large-scale random field evolved using linear perturbation theory. We explore the idea of "sewing" these simulations together in order to accurately model physics on both large and small scales, thereby obtaining consistent realizations of both the primary CMB and angular, projected matter fields.

In order to obtain lightcone data on small scales, we use the publicly available L-PICOLA code [113]. L-PICOLA is a "Lightcone-enabled Parallel Implementation of the COLA" method, providing an efficient means for generating both data on an observer's

past lightcone and data on spatial hypersurfaces. The COmoving Lagrangian Acceleration ("COLA") method [114, 115] works by solving the second-order Lagrangian perturbation theory (2LPT) equations in order to generate an initial guess for the motion of particles in the simulation, and subsequently solves a set of equations describing the difference between the 2LPT solution and the full N-body equations in order to improve the accuracy of the 2LPT solution. This method allows L-PICOLA to obtain results with an accuracy similar to full N-body simulations on the scales we are interested in, but with a substantially larger simulation timestep, and therefore at a substantially reduced computational cost. In the limit of many timesteps, the output from L-PICOLA should be equivalent to a traditional N-body simulation.

Although these N-body simulations are able to provide us with particular realizations of physics on small-scales, we are interested in modeling both the primary CMB and kSZ temperature fields. In order to obtain contributions to the kSZ temperature from the full dipole field (Eq. 3.4) in a manner consistent with the small-scale L-PICOLA data, as well as to generate the primary CMB, we utilize a novel "box-in-box" technique. This technique is similar in spirit to the mode-adding procedure (MAP) described in [116, 117], in that information about large scales is added to a small-scale simulation. However, the technique we utilize differs in several important regards. Similar to [118], we add information at the level of the density and peculiar velocity fields directly in Eulerian or configuration space, rather than in either Fourier space or Lagrangian space; additionally, no information is removed from the small-scale simulation.

We utilize N-body simulations with a number of particles $N_p = 1280^3$ in a comoving volume $(2\mathrm{Gpc/h})^3$, corresponding to a maximum simulation redshift of $z \sim 0.37$ and particle mass $2 \times 10^{12}\,\mathrm{M}_\odot$. While this coarse resolution does not allow us to resolve the structure of small mass halos, and also does not necessarily result in high-fidelity simulation data on the associated length scales, we find that the data we do obtain is sufficient for use in producing maps at angular resolutions of interest to us. We require the large-scale random field to encompass a volume containing the CMB (and ideally larger modes), so we utilize a large-scale "box" with volume $(32\mathrm{Gpc/h})^3$, resolved by 320^3 grid points.

The "box-in-box" method should be valid in a regime similar to the MAP method, which itself has been shown to perform well when linear theory provides a good description of the field content. This is a somewhat stronger condition than requiring mode amplitudes or the power spectrum to be well-described by linear theory. While the linear and nonlinear matter power spectra agree to within a few percent down to scales of order 10 Mpc, mode coupling can exist – nonlinear terms of order $\delta\rho^2$ can constitute percent or larger corrections to evolution of the density field on scales of order a few hundred Mpc. So long as we remain in a regime where the field configuration is sufficiently well-described by

(only) linear theory, we can expect the box-in-box technique to work. For the Gpc-scale N-body box sizes we employ here, this is the case.

3.3.1 Simulating small scales using L-PICOLA

We make use of both the lightcone output from L-PICOLA as well as data from spatial slices. The particle data from spatial slices is used to compute both primordial and large-scale components of the kSZ and primary CMB, and will be discussed in Section 3.3.2. The lightcone data is used to construct lightcone-projected sky maps of the density contrast field, velocity, and momentum fields, as well as convergence maps.

We generate radially binned maps of various fields, both in order to examine the underlying physics of the simulations, as well as to test reconstruction techniques at various redshifts. We divide the lightcone data into a number of radial bins between us and the largest redshift probed by the simulation. These radial bins can then be selectively integrated over to construct the contributions to a given field, such as density or kSZ temperature, from a given redshift range.

In order to produce density maps in both radial and angular bins, we bin particle data by noting that

$$\delta_{\rm bin} = \frac{\rho_{\rm bin} - \bar{\rho}_{\rm bin}}{\bar{\rho}_{\rm bin}} = \frac{n_{\rm bin}}{\bar{n}_{\rm bin}} - 1\,, \tag{3.12}$$

where ρ is the physical density inside a radial-angular-bin on the lightcone with comoving volume $V_{\rm bin} = \frac{\Omega_{\rm bin}}{3}(\chi_B^3 - \chi_A^3)$, where the bin has radial boundaries at χ_A and χ_B, and subtends a solid angle $\Omega_{\rm bin}$. The number of simulated particles of mass m in a bin is $n = \rho/m$, and the average/expected/background number of particles in a pixel is

$$\bar{n}_{\rm bin} = N_{\rm sim}\frac{V_{\rm bin}}{V_{\rm sim}} = \frac{N_{\rm sim}}{V_{\rm sim}}\frac{\Omega_{\rm bin}}{3}(\chi_B^3 - \chi_A^3)$$

in the case of discrete bins, or

$$\bar{n}_{\rm bin} = \frac{N_{\rm sim}}{V_{\rm sim}}\frac{\Omega_{\rm bin}}{3}3\chi^2 d\chi$$

in the continuum limit, with $N_{\rm sim}$ the total number of particles in a simulation of comoving volume $V_{\rm sim}$. The overdensity is then given by taking $n_{\rm bin}$ to be the number of particles in a given bin, so explicitly,

$$\delta_{\rm bin} = -1 + \sum_{\rm particles\in bin}\frac{1}{\bar{n}_{\rm bin}}\,. \tag{3.13}$$

This expression is similar in spirit to that of [119], although not identical. We also integrate the density contrast along a line of sight – or in a pixel subtending some solid angle on the sky; this can be written as a sum over the densities of all bins along the line of sight of the pixel,

$$\delta = \int d\chi \delta(\chi) = \sum_{\text{bin}\in\text{pix}} \delta_{\text{bin}} d\chi_{\text{bin}} \,, \tag{3.14}$$

where the bins that lie along the direction of the pixel on the sky are summed over. For different choices of radial binning, the sum will agree up to terms $\mathcal{O}(d\chi^2)$.

We are additionally interested in accounting for redshift-space distortions (RSDs) within this framework, requiring a small modification to the density field used in the reconstruction, Eq. 3.6. In order to take RSDs into account, we perturb particle positions by a small amount corresponding to the mis-inferred distance. Quantitatively, we compute

$$\chi_{\text{RSD}} = \chi_{\text{FRW}}(z_{\text{FRW}}(\chi) + v_{\text{Doppler}}) \,, \tag{3.15}$$

where functions with the FRW subscripts indicate the background FRW cosmology has been used, and where we then bin particles using Eq. 3.13 but according to their position χ_{RSD}. During the later discussion of reconstruction in this work, the density field used in reconstruction is the one that accounts for RSDs.

The convergence, formally written as

$$\kappa = \frac{3}{2} H_0^2 \Omega_{m,0} \int_0^{\chi_{\text{ls}}} d\chi \frac{\chi(\chi_{\text{ls}} - \chi)}{\chi_{\text{ls}}} \frac{\delta(\chi)}{a(\chi)} \,, \tag{3.16}$$

can similarly be binned. An expression for convergence binned in discrete angular pixels that is independent of radial binning is used [120, 121], allowing contributions to be placed into radial bins that can be summed over later to examine the convergence contribution from a given radial bin or range of radial bins,

$$\kappa_{\text{bin}} = \frac{3}{2} H_0^2 \Omega_{m,0} \frac{V_{\text{sim}}/N_{\text{sim}}}{\Omega_{bin}} \sum_{\text{particles},\, p,\in\text{bin}} \frac{1}{\chi_p a(\chi_p)} \frac{\chi_{\text{ls}} - \chi_p}{\chi_{\text{ls}}} \,, \tag{3.17}$$

so that for each angular pixel on the sky the total convergence will be

$$\kappa_{\text{pix}} = \sum_{\text{bin}\in\text{pix}} \kappa_{\text{bin}} \,. \tag{3.18}$$

We use this convergence map to lens the primary CMB.

There are several ways to compute the kSZ temperature fluctuations from particle data. The kSZ temperature fluctuations given by Eq. 3.1 can be evaluated by binning the components of the fields v and δ separately. However, the peculiar velocity field can be severely undersampled in simulated data, with nonzero velocities determined by only a single particle, or not at all in some pixels. In the case of the density field the issue is not as severe, as a lack of particles is merely indicative of an underdense region, where the density should be small anyways. A standard practice is therefore to write the integral in terms of a sum over peculiar particle momenta [122],

$$\Theta^{kSZ} = -\frac{\sigma_T f_b \mu}{\Omega_{\rm pix}} \sum_{\text{particles } p \in \text{bin}} \frac{m_p v}{D_{A,p}^2} \,. \tag{3.19}$$

In standard techniques used to construct kSZ temperature maps, the only contribution to the temperature field considered is the peculiar velocity of matter in Newtonian gauge projected along the line of sight, $v = v_{\rm Doppler,N-body}$. Thus, important contributions to the observed kSZ temperature perturbations on large angular scales from ISW, SW, or large-scale velocity modes (modes larger than the simulation volume) have not been modeled, each of which will contribute to the kSZ temperature fluctuations as described by Eq. 3.3.

In Section 3.3.2 we discuss more precisely how we model these additional contributions, however at the level of binning, we have two options. We can include these fields at the level of the already-binned lightcone data, replacing v with

$$v \to v_{\rm eff} = v_{\rm Dopp,\,N-Body} + v_{\rm Dopp,\,LS} + v_{\rm ISW} + v_{\rm SW} \tag{3.20}$$

where the ISW and SW components are given by the respective contributions of the effects (Eq. 3.3) to the temperature perturbation (Eq. 3.1), and where the Doppler contributions are from both the N-body simulation and large-scale (LS) modes not included in the N-body simulation. Alternatively, we can compute the overdensity $\delta_{\rm bin}$ and the velocity $v_{\rm eff}^{\rm bin}$ in each bin, and evaluate Eq. 3.1 directly. We find that both methods result in nearly identical kSZ temperature maps and power spectra for the angular resolutions we are interested in, although for the final maps we use Eq. 3.19 and 3.20.

The final quantity we compute using lightcone data is the peculiar velocity field, taking the velocity in each bin to simply be the average velocity of particles within each bin. At low angular resolutions, which for our simulations means a HEALPIX [123] resolution of Nside=1024, the narrowest redshift-angular bins we consider will typically contain at least one particle. At higher resolutions, artifacts become apparent in velocity maps due to undefined velocities in cells without particles [124]. However, for the dipole field in

particular, it is sufficient to compare low-Nside maps to our reconstructed velocity maps as we are interested in reconstructing the dipole field on large angular scales ($\ell \lesssim 20$).

As a final point of note, and as a check that the temperature maps and especially the dipole field reconstruction is insensitive to the precise binning method used, we employ binning using both a "nearest gridpoint" assignment scheme, and a "cloud-in-cell"-type assignment scheme where contributions from individual particles of fields are distributed to a weighted average of nearby cells, both radially and in an angular direction. The latter of these methods introduces additional smoothing, or aliasing, on bin-sized scales; this suppresses power on these scales, but also suppresses the effects of shot noise. Despite this difference, we find that the performance of the dipole field reconstruction presented below is largely insensitive to this detail.

3.3.2 Large scales: "box-in-box"

We formally describe the process of sewing the N-body data and the large-scale modes together using a "coloring" operator $\mathcal{C}_{P_c}(f)$ that rescales a stochastic field f (with its own power spectrum P_f) by a power spectrum P_c,

$$\mathcal{C}_{P_c}(f) = \int \frac{d^3k}{(2\pi)^3} e^{i\vec{k}\vec{x}} f(\vec{k}) P_c^{1/2}(k) \tag{3.21}$$

so the power spectrum of the resulting field is given by $P_f P_c$. For a coloring spectrum $P_c = P_f^{-1}$, the field will be whitened. We additionally make use of an "inlay" operator, $\mathcal{I}(f_1, f_2)$, which acts in configuration space to replace values in the interior of a (large-scale) field by values of a second (coarsened, small-scale) field. The procedure of sewing a small field into another larger field then consists of the following operations:

$$f_{\rm sewn} = \mathcal{C}_{P_f}(\mathcal{I}(\mathcal{C}_{P_f^{-1}}(f_{\rm lg}), \mathcal{C}_{P_f^{-1}}(f_{\rm sm}))) \tag{3.22}$$

Evaluating the above expression entails taking Gaussian random fields $f_{\rm sm}$ and $f_{\rm lg}$, both with statistical properties described by P_f, whitening these fields, replacing values of the lg field by ones from the sm field, and finally de-whitening the fields.

The result of this procedure on the lg field is that the small-scale modes in the region of replacement are now given by modes from the sm field, while large-scale modes have been preserved and superimposed upon the small-scale field.

Although L-PICOLA provides us with information about the density field, we are ultimately interested in obtaining the primordial potential, from which we can compute

corrections to the velocity field using linear theory. In order to obtain the primordial potential on large scales, we extract the potential on the initial slice using the density field and Poisson equation[1],

$$\nabla^2\Phi = 4\pi G a^2 \delta_\rho^{\rm sim}\,. \tag{3.23}$$

The potential can then be evolved back in time using the transfer function for the potential, $T(\phi_{\rm sim} \to \phi_{\rm prim})$, and corrections to the velocity field then found using the velocity growth function defined in [18].

We are therefore interested in computing

$$\phi_{\rm sewn} = T(\phi_{\rm sim} \to \phi_{\rm prim}) \frac{4\pi G a^2}{\nabla^2} \mathcal{C}_{P_\delta}(\mathcal{I}(\mathcal{C}_{P_\delta^{-1}}(\delta_{\rm Box}), \mathcal{C}_{P_\delta^{-1}}(\delta_{\rm LP}))) \tag{3.24}$$

where the L-PICOLA density field is noted by the LP subscript, and the Box subscript refers to a random realization of a density field with power spectrum P_δ. Written in Fourier space, the outermost coloring operation, transfer function operation, and the inverse Laplacian operation can all be combined into an operation equivalent to coloring by the primordial spectrum. Coloring the large-scale Box modes with its inverse spectrum is also equivalent to simply generating a field of white noise, $N_{\rm white}$. Thus the final operation we perform in order to obtain a large-scale primordial potential consistent with the density field from the L-Picola simulation is

$$\phi_{\rm sewn} = \mathcal{C}_{P_{\phi,\rm prim}}(\mathcal{I}(N_{\rm white}, \mathcal{C}_{P_\delta^{-1}}(\delta_{\rm LP}))\,. \tag{3.25}$$

The power spectra for the comoving density field is obtained using the CLASS code [128], and the primordial Newtonian potential is chosen to be the usual scale-invariant one. We show snapshots of various steps of this procedure in Fig. 3.1

Once we have the primordial potential, we use the CMB radiation transfer functions to obtain the primary CMB, and velocity transfer functions to compute the contributions to the dipole field due to large-scale modes[2]. From the large-scale modes, we can then compute the contributions to $v_{\rm eff}$ from Eq. 3.20, and thus their contribution to the observed

[1] As a technical note, we can safely interpret output from the L-PICOLA simulations in a standard way without worrying about relativistic effects given our accuracy requirements for the scales we are interested in [125, 126, 127]: the density field is interpreted as the comoving synchronous gauge one, while the velocity and metric fields are interpreted as the Newtonian ones. The evolution of large-scale modes is then determined using linear cosmological perturbation theory in the appropriate gauge, which takes into account relativistic effects at linear order in metric and density perturbations.

[2] We could also modulate large-scale modes in the density field, however long-wavelength density perturbations contribute negligibly to the cross correlation between the kSZ temperature and density field [18], so we do not include this modulation.

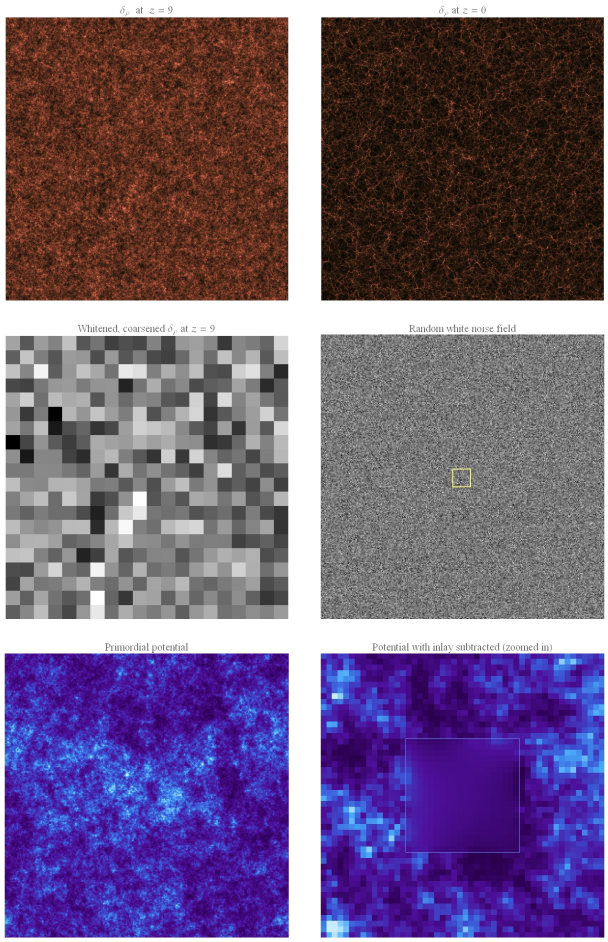

Figure 3.1: Slices of spatial hypersurfaces of simulations during various parts of the sewing-in procedure described by Eq. 3.25. Top left: the initial L-Picola density contrast field at $z = 9$, with comoving box size $L = 2\,\mathrm{Gpc}/h$. Top right: the density field at $z = 0$. Middle left: the initial $z = 9$ density field, whitened using the matter power spectrum, averaged over (coarsened) so the resolution is the same as that of the box containing large-scale modes. Middle right: A random realization of white noise for large-scale modes, with $L = 32\,\mathrm{Gpc}/h$. The central $2\,\mathrm{Gpc}/h$ region that will be replaced has been outlined with a yellow border. Bottom left: The primordial potential with white-noise values in the large box replaced using the whitened L-PICOLA field, then colored using the primordial power spectrum. Bottom right: The central $5\,\mathrm{Gpc}/h$, with the colored small-scale box values directly subtracted. Small residual large-scale modes can be seen in the center. The region where the subtraction has been performed is outlined.

kSZ temperature fluctuations. When computing large-scale contributions to $v_{\rm eff}$, we also need to ensure we do not double-count modes already accounted for by the N-body simulation. Therefore, when computing the large-scale Doppler contribution to $v_{\rm eff}$, we only integrate over modes with wavelengths larger than the N-body simulation volume.

The CMB multipoles are then computed using the large-volume simulated primordial potential up to $\ell = 28$. In principle we could generate additional CMB modes using simulated data, however they will not be correlated with the remote dipole field or the density field. We therefore use a random realization of the primary CMB a_{lm}s at $\ell > 28$, based on the theoretical power spectrum obtained from CLASS.

We also include lensing of the primary CMB, utilizing the convergence maps generated from the lightcone data. From the maps we can compute the lensing potential ϕ in harmonic space as

$$\phi_{\ell m} = \frac{2\kappa_{\ell m}}{\ell(\ell+1)}. \tag{3.26}$$

The lensed CMB temperature is then given by

$$T(\hat{n}) \to T(\hat{n} + \nabla\phi) \simeq T(\hat{n}) + \nabla\phi\nabla T(\hat{n})\,. \tag{3.27}$$

Although the CMB is lensed, the kSZ temperature is not. In principle, there could be lensing of the kSZ temperature fluctuations due to any structures between kSZ sources and an observer, however we do not model this. The lensing we compute is also derived from only the N-body volume we simulate, thus in a more realistic treatment, structure at higher redshifts and on large scales would need to be included. However, the small-scale density-temperature correlations induced by lensing from the density field we use for reconstruction are accounted for. In future work, we would nevertheless prefer to include lensing (and kSZ) contributions from additional redshifts.

To model kSZ temperature anisotropies sourced at redshifts beyond our N-body simulation, we include Gaussian random noise on angular scales $\ell \gtrsim 1000$ with amplitude $\sim 2\ \mu$K.

In Figure 3.2, we show various outputs of our simulation pipeline. Of particular note is that the that the projected density is properly correlated with the convergence field and the kSZ temperature anisotropies, and the dipole field is properly correlated with the primary CMB and the kSZ temperature anisotropies. Focusing on the kSZ map, both the large-scale contributions and large-scale modulation of power from the dipole field are visible. In addition, the primordial components of the dipole field (i.e. contributions from modes in the big box) are visible as responsible for the structure of the dipole field on large angular scales. On these scales, the primordial contributions can dominate the power, resulting in

a temperature asymmetry that can be seen "by eye". The particular north-south direction of the asymmetry in Figure 3.2 is peculiar to the realization.

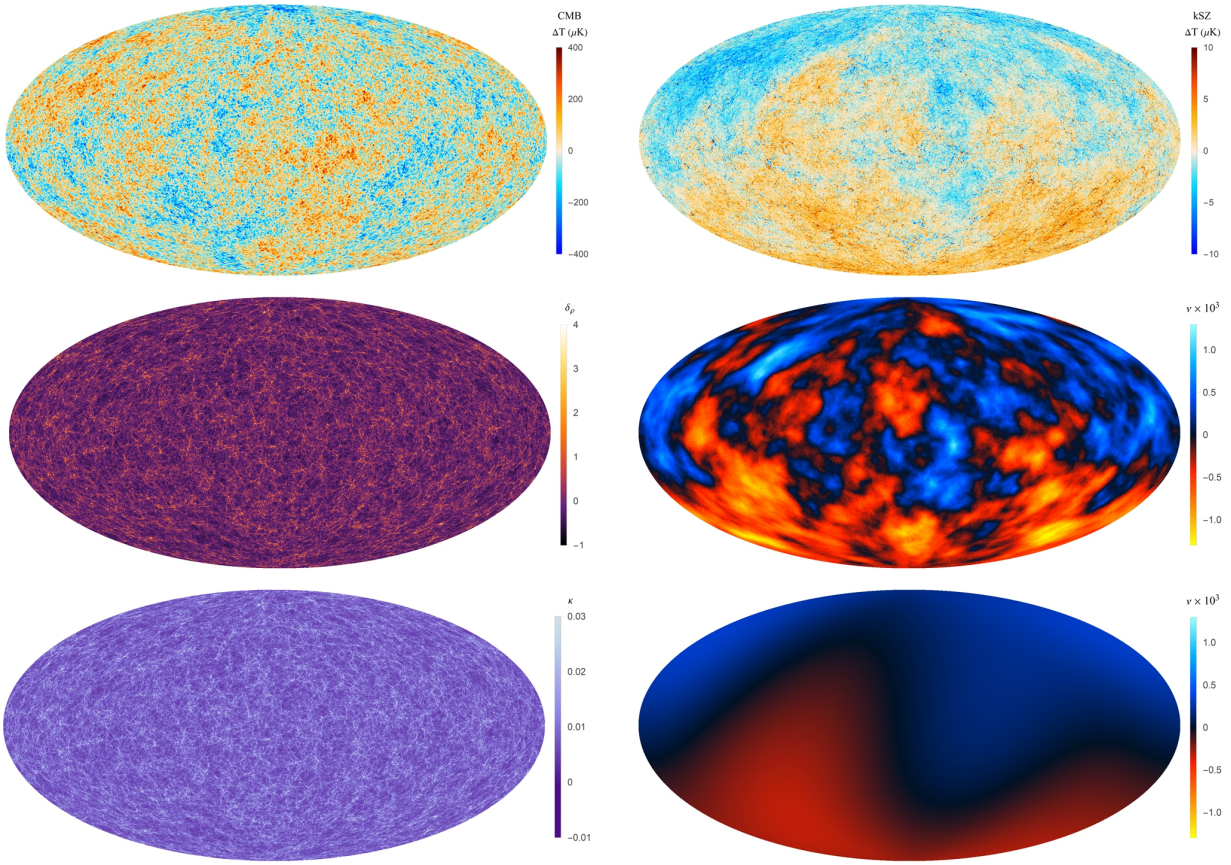

Figure 3.2: Hammer-Aitoff projections of different fields on the sky from the box-in-box simulations; all fields are properly correlated. Top left: the total CMB temperature fluctuations, including kSZ contributions. The CMB dipole is not included. Top right: the contribution of the kSZ effect to temperature fluctuations. Middle left: the binned, average density field (Eq. 3.6); middle right: the binned, average dipole field; bottom left: the binned convergence field; bottom right: the contribution to the remote dipole field from the big-box modes. Binning is performed over a redshift range $z = 0.18$ to $z = 0.27$.

3.4 Results

3.4.1 Reconstruction using a quadratic estimator

We now analyze data from an ensemble of ten simulations to assess the performance of the quadratic estimator, Eq. 3.7. We utilize two radial binning schemes, with the density field on the light cone of each simulation arranged into either a single bin or eight bins of equal radial comoving width. For each simulation and bin we construct maps of ξ defined in Eq. 3.8 and ζ^α defined in Eq. 3.9. The power spectra $C_\ell^{\Theta\Theta}$, $C_\ell^{\delta^\alpha\delta^\alpha}$, and $C_\ell^{\dot{\tau}^\alpha\delta^\alpha}$ used in Eqs. 3.8, 3.9, 3.10 to generate the ξ, ζ^α fields and reconstruction noise are the sample variances from each realization. We then obtain the estimated moments of the binned dipole field from Eq. 3.7, and generate a map of the reconstructed average dipole field in each bin.

In Figure 3.3 we compare the reconstructed and actual bin-averaged dipole fields for a single bin and for the 8th bin of the eight bin configuration. All maps are filtered to contain only multipoles $\ell < 28$. "By-eye," the reconstruction performs well on large angular scales. We quantify the agreement between the reconstructed and actual dipole field in two ways.

First, we make a comparison at the level of the power spectra in Fig. 3.4. We compute the mean and standard deviation of the reconstructed dipole field power (with the noise bias removed) and the actual dipole field power (total, and separate contributions from the small and big box modes), as well as the prediction from linear theory using Eq. 3.5. In this figure, we plot these quantities for the single bin (top left) and bins 2, 4, and 8 of the eight bin configuration. In general, the agreement between the mean reconstructed and the mean actual power is quite good at low multipoles, within a single standard deviation. For higher multipoles, the reconstruction is poor and there is an excess of power due to the reconstruction noise. In addition, there appears to be a systematic bias towards extra power in the reconstructed field at low multipoles, especially in the single-bin configuration and the lowest redshift bins of the eight bin configuration; the agreement with linear theory becomes better at higher redshift. This is consistent with a bias due to gravitational non-linearities, which we expect to be more important at low redshift. A similar bias exists in CMB lensing [112], and we hope to investigate this possibility in future work.

As an additional diagnostic of the performance of the reconstruction, we compute the correlation coefficient between the reconstructed and the original field, which we refer to

as the reconstruction efficiency

$$r_L \equiv \frac{\hat{C}_L^{\hat{v}\bar{v}}}{\left(\hat{C}_L^{\hat{v}\hat{v}}\hat{C}_L^{\bar{v}\bar{v}}\right)^{1/2}}, \tag{3.28}$$

where $\hat{v}$ denotes the reconstructed field and $\bar{v}$ the actual field. The efficiency is not sensitive to an overall change in normalization, but instead provides us with a measure of how strongly correlated reconstructed and simulated modes are. In general, we find that the reconstructed modes agree well with the simulated modes on the largest angular scales. The reconstruction efficiency is found to be better at higher redshift, again we expect this due to a lack of nonlinear effects. Reconstruction is also found to perform better in smaller bins, an effect we can at least partially attribute to the increased information content: information from small-scale modes has not been so heavily averaged away. However, in larger redshift bins, the correlation with primordial modes is larger, as discussed in the next section.

3.4.2 CMB-kSZ Dipole Correlation

We now consider how well we can determine the "intrinsic" CMB dipole using information from the reconstructed large-scale velocity field, as suggested in Ref. [19]. This idea is not without ambiguity – because one can arbitrarily change the CMB dipole by performing a boost, there is no unique definition of the intrinsic dipole. Instead, one must settle on a definition universal and specific enough to facilitate a meaningful comparison. We can make progress by noting that the local CMB dipole should, to an extent depending on one's definition of the intrinsic CMB dipole, be correlated with the $\ell = 1$ moments of the remote dipole field. The contributions to our measured CMB dipole and the remote dipole field of a nearby observer are determined primarily by small-scale modes which source local peculiar velocities. However, there are also subdominant contributions to the CMB dipole from larger-scale (but still local) velocity modes and other effects both along our past lightcone and at the CMB last scattering surface. In Sec.1.5 of Chapter 1, we referred to the dipole sources beyond the local Doppler as the "stationary dipole" sources.

A standard definition of the fundamental CMB dipole is obtained by boosting to a reference frame in which the relativistic aberration of the CMB vanishes (see e.g. [129]). In Newtonian gauge, this aberration-free dipole is calculated in the frame where an observer has vanishing local peculiar velocity, altering the Doppler term in Eq. 3.3. This definition of the fundamental CMB corresponds to the stationary dipole from Sec.1.5. A more general

definition of the fundamental CMB dipole is obtained by applying a low-pass filter to the Fourier modes contributing to local peculiar velocities. The aberration-free dipole is a special case, where all modes contributing to the local Doppler term are filtered out. This more general definition is also more closely related to the dipole field obtained in kSZ tomography, since the bin-averaging effectively imposes a low-pass filter on radial peculiar velocities. We will refer to this as the large-scale Doppler dipole.

We can quantitatively express the correlation between the remote dipole field and the various definitions of the CMB dipole in terms of transfer functions, with the CMB transfer function filtered below a given scale $k_{\rm cut}$,

$$C_1^{Tv^\alpha} = \int \frac{d^3k}{(2\pi)^3} P_\Psi(k) \Delta_1^{v^\alpha *}(k) \Delta^T_{\rm filt,1}(k)\,, \tag{3.29}$$

where as before, α labels a redshift bin in which the remote dipole field v is averaged. The filtered CMB transfer function for the dipole is given by

$$\Delta^T_{\rm filt,1} = \Theta(k_{\rm cut} - k)\Delta^T_{\rm dopp,\,local,1}(k) + \Delta^T_{\rm dopp,\,CMB,1}(k) + \Delta^T_{\rm ISW,1}(k) + \Delta^T_{\rm SW,1}(k) \tag{3.30}$$

where Θ is the Heaviside step function, and the individual contributions to the radiation transfer function include ISW, SW, and both local and last-scattering-surface (CMB) Doppler contributions. For the large-scale Doppler dipole, we choose a filtering scale equal to the N-body simulation volume ($L_{\rm cut} \sim 2\pi/k_{\rm cut} \sim 3\,{\rm Gpc}$). For the aberration-free dipole, $k_{\rm cut} \to 0$.

In Figure 3.6, we plot the theoretical prediction for the correlation coefficient (e.g. Eq. 3.28) using linear theory between the $\ell = 1$ moment of the bin-averaged remote dipole field and three definitions of the CMB dipole: the observed CMB dipole ("all Doppler"), the aberration-free or stationary dipole, and the large-scale Doppler dipole. We plot the theory prediction for a single redshift bin $(0, z)$ of varying radial extent. In addition, we show the mean and standard deviation of the correlation coefficient calculated from ten simulations for redshift bins of two different size using the simulated CMB large-scale-filtered dipole and the reconstructed dipole field. As expected from the discussion above, the correlation between the observed CMB dipole and the bin-averaged dipole field is small for all but the smallest bins. Because they are composed primarily of large-scale modes, the correlation between the bin-averaged dipole field and the aberration-free and large-scale Doppler dipoles improves with bin width. However, the dipole field has a finite correlation length, and therefore the correlation coefficient eventually goes down. We find that the large-scale Doppler dipole can in principle be determined with a maximum correlation coefficient of $r \sim 0.9$ while the aberration-free dipole can be determined with a maximum

correlation coefficient of $r \sim 0.65$. The optimal reconstruction bin width corresponds to a redshift of $z \sim 0.4$. In conclusion, our simulations indicate that constraints on the intrinsic CMB dipole should reasonably be attainable in individual realizations.

3.5 Discussion and Conclusions

kSZ tomography is a useful tool for probing the largest observable scales in our Universe, providing information in addition to what the primary CMB and large-scale density surveys alone can tell us. In this chapter we have explored the ability of a quadratic estimator to reconstruct the remote dipole field using simulated maps of the CMB and density field. We have found that the reconstruction process is able to capture highly significant information about large scales, even in the presence of physical effects with the potential to contaminate our ability to reconstruct, including nonlinear growth of structure, RSDs, lensing, and contributions to the kSZ temperature from structures outside the range of redshifts considered for reconstruction.

We have accomplished this using a novel simulation technique, in which a small-scale N-body simulation is sewn into a large-scale volume evolved with linear theory, allowing us to generate self-consistent maps of kSZ temperature fluctuations, the primary CMB, CMB lensing, density, and dipole fields. In turn, the consistency of these components allows us to explore the ability of reconstruction techniques to probe fundamental physics such as determining the intrinsic CMB dipole. More generally, the reconstructed remote dipole can be used as a general-purpose cosmological observable, and is in principle a powerful probe of primordial non-Gaussianity [78], relativistic effects in galaxy surveys [76], modified gravity [77], isocurvature perturbations [111, 80], and CMB anomalies, which we discuss in Chapter 4.

While this work furthers our confidence in the ability of the reconstruction procedure to work in practice, it will be important that future simulations of remote dipole reconstruction on the lightcone account for the many additional physics discussed in Chapter 2 in order to make future predictions as realistic and robust as possible. In the big picture, Chapters 2 and 3 have laid the groundwork for simulations of kSZ velocity reconstruction that incorporate the effects of multi-frequency CMB foregrounds, the impact of the modelling of the electron distribution on small scales through the optical depth bias, the effects of photometric redshift errors and calibration errors in galaxy surveys. and the impact of partial coverage of the sky.

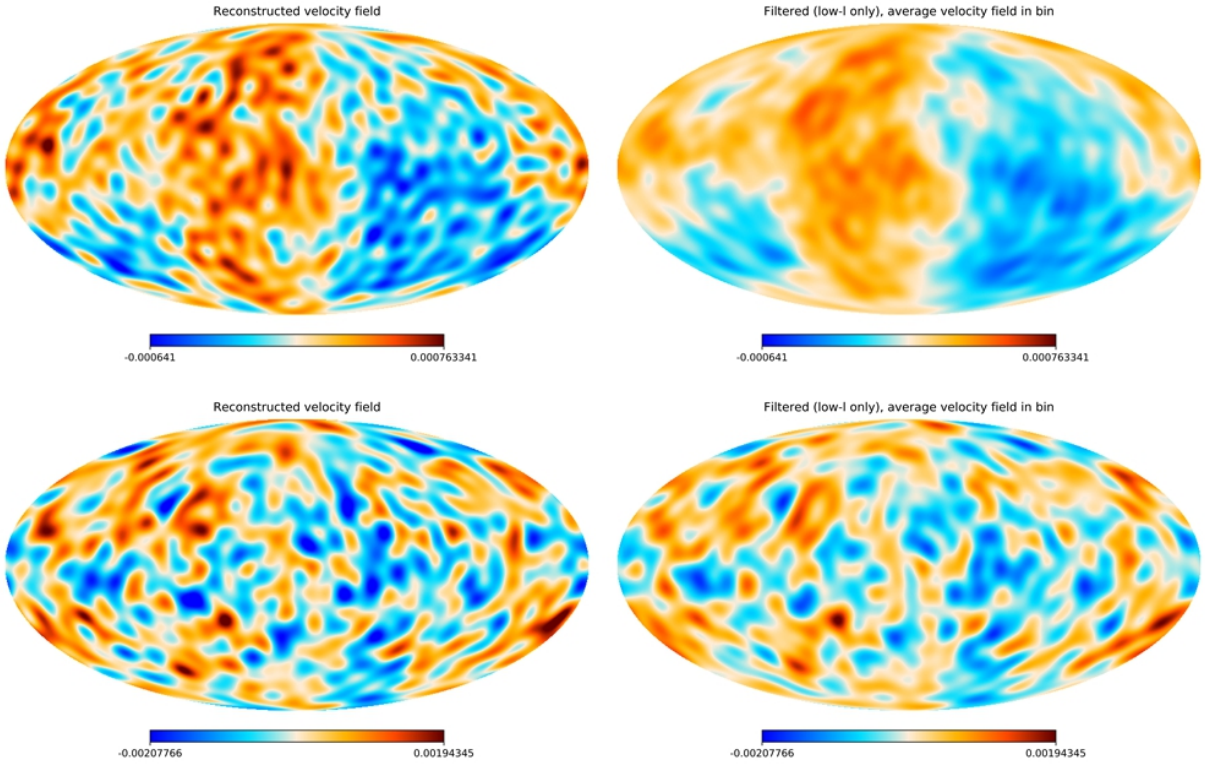

Figure 3.3: The remote dipole field obtained from simulations compared to the reconstructed remote dipole field. The maps do not include modes higher than $\ell > 28$. The reconstruction of the top two plots was done using a single redshift bin from $z = 0.086$ to $z = 0.37$, while the bottom plots are a redshift bin from $z = 0.33$ to $z = 0.37$. By eye, it is noticeable that large angular modes between the two maps tend to agree, while smaller-scale modes only do to a moderate extent. The reconstruction of smaller scales is also found to be better in the smaller, higher-redshift bin. This is in agreement with results obtained by looking at the reconstruction efficiency, shown in the top left panel of Figure 3.5. Excess power can also be seen on small scales, consistent with the spectra found in Figure 3.4.

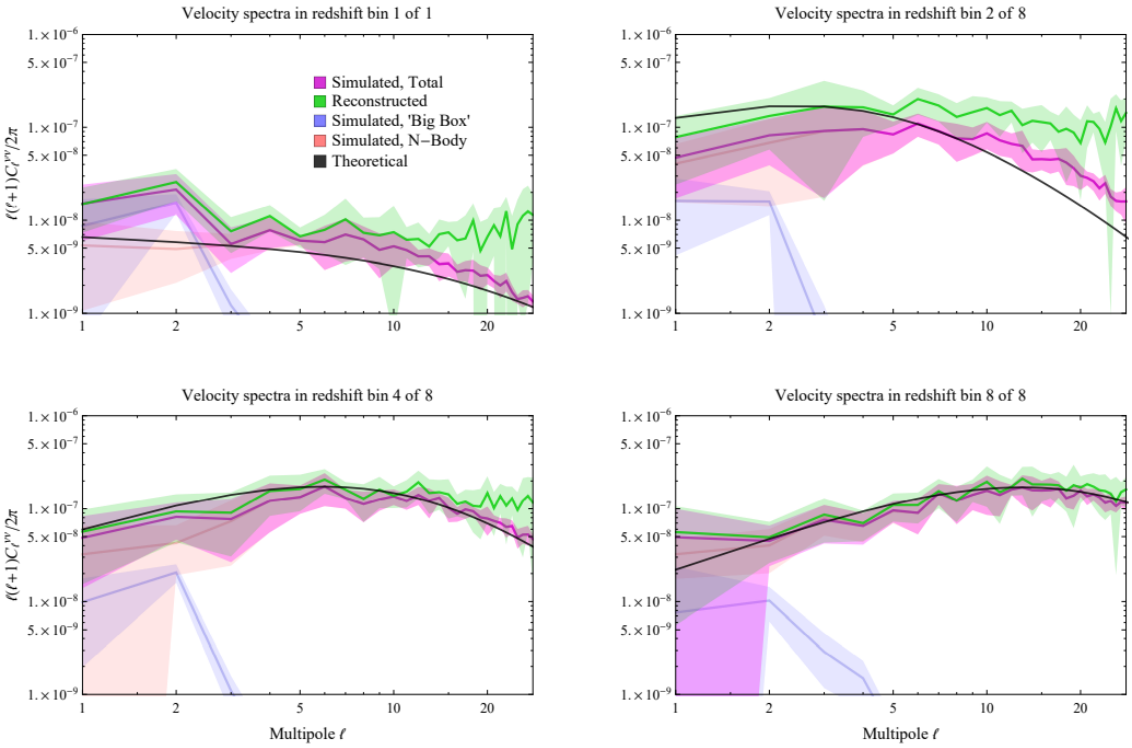

Figure 3.4: The velocity power spectra from simulated data, compared to the theoretical and reconstructed spectra. Contributions to the spectra from the N-body simulation are shown in red, contributions from the large-scale box modes in blue, and the total in purple. The reconstructed spectra with noise subtracted is in green, and linear theory prediction in black. Lines indicate the mean spectrum from our simulations, while solid bands indicate the variance. The reconstruction is performed using redshift data in bins over a redshift range of $z = 0.086$ to $z = 0.37$, subdivided into one or eight bins of equal comoving distance. Reconstruction efficiencies are shown in Figure 3.5.

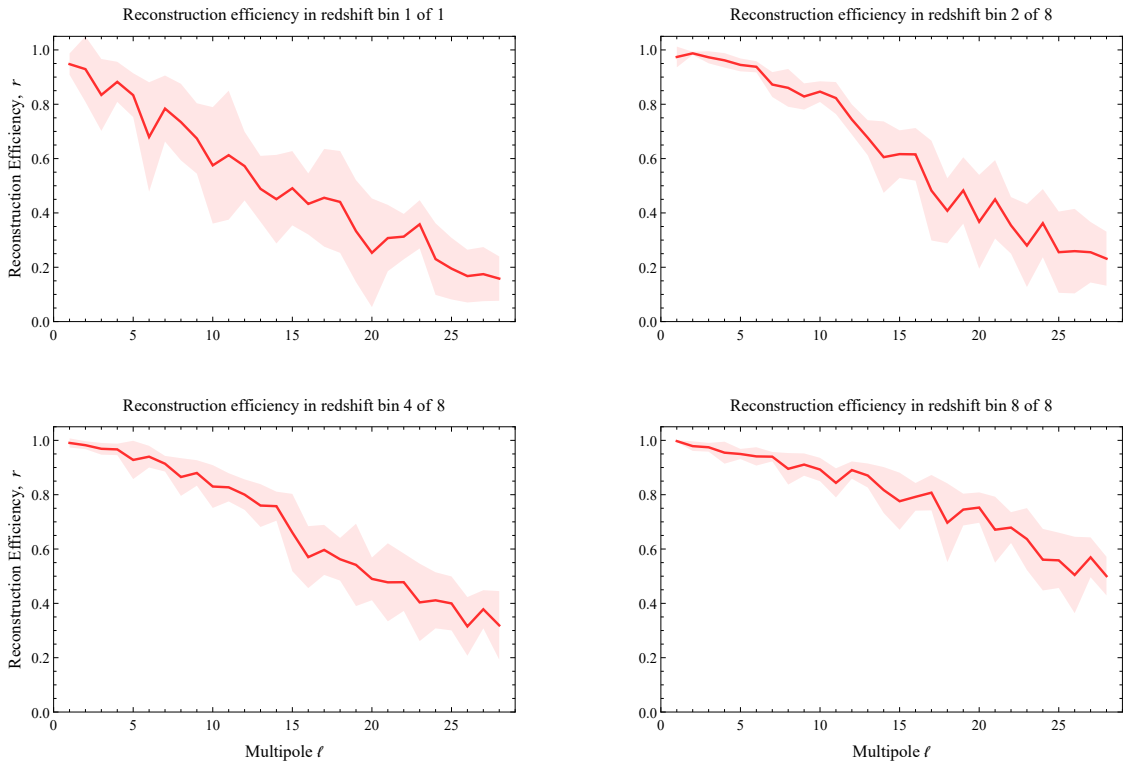

Figure 3.5: The reconstruction efficiency for the runs in Figure 3.4. The solid line indicates the mean reconstruction efficiency in each bin for each simulation realization, and the solid band the standard deviation.

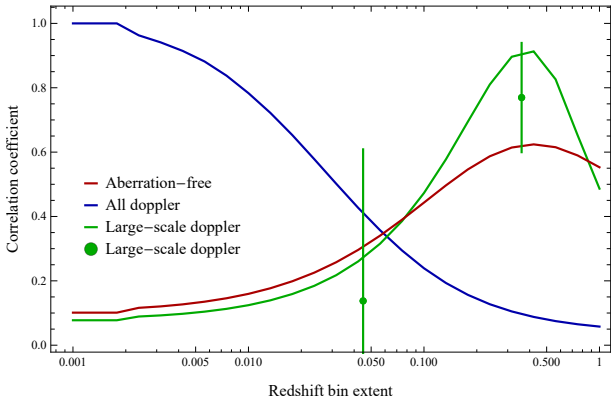

Figure 3.6: The reconstructed velocity field and CMB temperature dipole correlation coefficient, $C_1^{T\bar{v}}/\sqrt{C_1^{TT}C_1^{\bar{v}\bar{v}}}$, computed using different CMB dipoles. The theoretical correlation using the full CMB transfer function is shown in blue, correlation with the aberration-free dipole in red, and the correlation with "filtered" CMB dipole shown in green. Data point show the correlation of the simulated CMB dipole filtered on 3 Gpc (box-sized) scales for two redshift bin sizes. The points are the mean correlation from all simulations we perform, and error bars denote the standard deviation.

Chapter 4

Testing anomalies

4.1 Chapter summary

Anisotropies in the Cosmic Microwave Background (CMB) are a powerful probe of early Universe physics. On large angular scales, these anisotropies encode primordial density fluctuations, which may ultimately have been produced at energy scales far beyond the reach of any terrestrial particle accelerator. Interestingly, a series of anomalous large angular scale features in the microwave sky have been reported by the WMAP and Planck [130, 131] satellite missions, offering what could be hints of physics beyond the standard model of cosmology, ΛCDM. Several notable anomalies include: a hemispherical power asymmetry, a lack of correlations on large angular scales, features in the angular power spectrum, and an alignment of multipoles; see [5] for a recent review. As the statistical significance of each of these anomalies is rather modest, the most conservative position is to attribute them to statistical flukes, given the a posteriori nature of their discovery, systematics or foregrounds. Unfortunately, as a stand-alone probe, the CMB temperature has already reached the limit imposed by cosmic variance on large angular scales, so new information can only come from alternative or complementary probes of the largest scales in the Universe.

Several observables have been identified as potential probes of physical models of the CMB anomalies, including: CMB polarization (see e.g. Refs. [132, 133, 134, 135, 136, 137, 138, 139]), CMB lensing (see e.g. [140, 141]), the integrated Sachs-Wolfe (ISW) effect (see e.g. [100, 142, 143]), and probes of large scale structure in the late Universe (see e.g. [144, 145, 146, 147]). Each of these observables has both advantages and disadvantages. CMB polarization can access scales comparable to those in the CMB temperature. On the largest

angular scales, however, the mapping between the observed polarization anisotropies and physical scales is dependent on the (relatively poorly constrained) history of reionization. In addition, on large scales galactic foregrounds are challenging (though not impossible) to remove [4]. The lensing potential can be reconstructed with high fidelity using future CMB datasets (e.g. Simons Observatory [15] or CMB-S4 [16]), however there is limited support from the physical scales associated with many of the CMB anomalies (see [141]). If the (late time) ISW contribution to the CMB temperature can be isolated (e.g. using the technique of [143]), this could contribute a modest number of modes probing large scales. Finally, future galaxy surveys (e.g. Rubin [10], Euclid [11], Spherex [148]) or 21cm experiments (e.g. CHIME [12], HIRAX [14]; see also [13]) can reach large enough volumes to offer new information on some of the CMB anomalies. While promising the measurement of a huge number of modes on linear scales, there will be limited support on physical scales responsible for the lowest multipoles of the CMB temperature, and measurement of the largest modes will be noisy and plagued by various systematics (see e.g. [149]).

The goal of this chapter is to explore a new set of observables that may become important tools in the study of the physical nature of CMB anomalies: the remote dipole and quadrupole fields, i.e. the $\ell = 1, 2$ moments of the microwave sky measured throughout our observable Universe (see Secs.1.3.4 and 1.3.5). We remind the reader that the remote dipole manifests itself through the kinetic Sunyaev Zel'dovich (kSZ) effect [23]: the dominant blackbody temperature contribution to *our* CMB sky on angular scales corresponding to multipoles $\ell \gtrsim 4000$ originates from free electrons on our past light-cone scattering *their* locally observed CMB dipole. Similarly, in the presence of a local CMB quadrupole, the scattered photons are endowed with a polarization. The polarized component of the CMB arising after reionization, primarily from collapsed structures, is known as the polarized Sunyaev Zel'dovich (pSZ) effect (as opposed to the component sourced near decoupling and at reionization, which is simply CMB polarization). Information of the remote dipole and quadrupole on our past lightcone can be reconstructed using the techniques of kSZ velocity reconstruction (thoroughly discussed throughout Chapters 1,2 and 3) and pSZ quadrupole reconstruction (reviewed in Appendix A.1). Below, we refer to these two techniques more generally as SZ tomography. Previous forecasts and analyses [18, 19, 20, 21, 36], as well as the results presented so far in this thesis, have established the feasibility of the reconstruction techniques with future datasets.

Being primarily sensitive to inhomogeneities on large physical scales, these new observables that will become accessible with the next generation of CMB and galaxy surveys stand as potential candidates to further extend our understanding of the CMB anomalies. The remote quadrupole field receives support from the same scales contributing to the low-ℓ moments of the CMB temperature. Although at low redshift and on large angular

scales the remote quadrupole field is strongly correlated with the primary CMB temperature quadrupole [37, 150], there is significant new information on moderate angular scales and at high redshift [36, 151, 19, 63]. The remote dipole field is dominated by the coarse-grained line-of-sight peculiar velocity field, and is therefore sensitive to somewhat smaller scales than the remote quadrupole. However, it can be reconstructed at far higher signal to noise, and carries a significant amount of information on scales relevant to a variety of CMB anomalies. As general probes of physics on large scales, the remote dipole and quadrupole fields can yield improved constraints on: primordial non-Gaussianity [78], primordial gravitational waves [152, 151], and pre-inflationary relics [153].

The central question we wish to address in this work is whether or not the remote dipole and quadrupole fields could serve as alternative and complementary probes of the CMB anomalies to more traditionally considered tracers: the large scale E-mode CMB polarization and 3-dimensional galaxy maps on our past light-cone. An important fact to highlight is that measurements of these observables significantly differ in their nature: a key feature of SZ tomography is that it reconstructs large-scale inhomogeneities from anisotropies on the smallest angular scales. The fidelity of the reconstruction improves with the sensitivity and resolution of the CMB experiment and the depth and redshift errors of the galaxy survey. Therefore, the information accessible using SZ tomography will improve greatly with time, while direct probes of the largest scales are already close to the cosmic variance limit. Apart from offering an alternative way of measuring the large scale properties of the Universe, the remote dipole and quadrupole fields may also capture new independent information to that already available through the traditional probes.

Our methodology to estimate how informative the remote dipole and quadrupole fields could be for the study of the CMB anomalies consists of forecasting parameter constraints for physical models of the anomalies under the experimental conditions of idealized next-generation CMB experiments and galaxy surveys. We compare the performance of different combinations of the primary CMB temperature, CMB polarization, galaxy clustering and the remote dipole and quadrupole fields. We assume that foreground-cleaned data is accessible for the CMB polarization on large scales, for the galaxy survey and also for the kSZ and pSZ effects on small angular scales. The limiting experimental factors that we do take into account are CMB noise levels as expected for Stage-4 CMB experiments and Rubin-like specifications for the galaxy survey, as well as calibration errors on large scales for the latter. We perform our forecast using a Fisher matrix formalism on a series of physical models for three representative CMB anomalies: the power asymmetry, the lack of power on large angular scales, and a feature in the temperature power spectrum at multipoles around $\ell \sim 20 - 30$. We find that the remote CMB dipole and quadrupole can yield constraints on anomalies models significantly beyond what can be done with CMB

temperature and polarization on large scales, and comparable to what is possible with a large photometric galaxy survey. The best constraints are obtained from a multi-tracer analysis including the primary CMB, remote fields, and galaxy survey. It is important to mention that the work presented in this chapter was done prior to the extended formalism presented in Chapter 2, and thus follows the treatment for the quadratic estimators from [19] (reviewed in Sec.1.4 and Appendix A.1) plus some additions like the optical depth bias marginalization and the impact of photometric calibration errors on the galaxy redshift survey. Still, the goal of this study was to provide evidence of the sensitivity of the reconstructed remote dipole and quadrupole to the physics involved in CMB anomalies models, and identify them as valuable probes to include in a multi-tracer analysis. The extended extended formalism will serve as a new platform for future forecasts and the eventual analysis of real data.

The plan of the chapter is as follows. In Sec. 4.2, we review SZ tomography and describe the properties of the remote dipole and quadrupole fields. In Sec. 4.3, we describe the details of our forecast and introduce a figure of merit which is used to quantify the potential constraining power with different combinations of observables. In Sec. 4.4 we present the results of our forecast, and we conclude in Sec. 4.5.

4.2 The remote dipole and quadrupole fields

In this section, we review the details of the remote dipole and quadrupole fields and argue for their utility as probes for the CMB anomalies. Contributions to the CMB temperature and polarization generated via the kinetic and polarized SZ effects can be expressed through the line of sight integrals

$$\Theta^{kSZ}(\hat{\mathbf{n}}) = -\int d\chi\, v_{\rm eff}(\hat{\mathbf{n}},\chi)\,\dot{\tau}(\hat{\mathbf{n}},\chi), \tag{4.1}$$

$$v_{\rm eff}(\hat{\mathbf{n}},\chi) \equiv \frac{3}{4\pi}\int d^2\hat{n}_e\,\Theta(\mathbf{x},\hat{\mathbf{n}_\mathbf{e}})\,(\hat{\mathbf{n}}\cdot\hat{\mathbf{n}_\mathbf{e}}) = \sum_{m=-1}^{1}\Theta_1^m(\hat{\mathbf{n}},\chi)\,Y_{1m}(\hat{\mathbf{n}}) \tag{4.2}$$

$$(Q\pm iU)^{\rm pSZ}(\hat{\mathbf{n}}) = -\frac{\sqrt{6}}{10}\int d\chi\, q^{\pm}_{\rm eff}(\hat{\mathbf{n}},\chi)\,\dot{\tau}(\hat{\mathbf{n}},\chi), \tag{4.3}$$

$$q^{\pm}_{\rm eff}(\hat{\mathbf{n}},\chi) \equiv \sum_{m=-2}^{2}\Theta_2^m(\hat{\mathbf{n}},\chi)\,{}_{\mp2}Y_{2m}(\hat{\mathbf{n}}) \tag{4.4}$$

where $\hat{\mathbf{n}}$ denotes the line of sight direction, χ the comoving distance to the scatterer, $\Theta_\ell^m(\hat{\mathbf{n}}, \chi)$ are the moments of the CMB temperature at the position of the scatterer, and $\dot{\tau}(\hat{\mathbf{n}}, \chi)$ is the differential optical depth defined as

$$\dot{\tau}(\hat{\mathbf{n}}, \chi) \equiv \sigma_T a(\chi) n_e(\hat{\mathbf{n}}, \chi) = \sigma_T a(\chi) \bar{n}_e(\chi) [1 + \delta_e(\hat{\mathbf{n}}, \chi)] \tag{4.5}$$

with $a(\chi)$ the scale factor, σ_T the Thompson cross-section and $\delta_e(\hat{\mathbf{n}}, \chi)$ the perturbations about the average electron number density $\bar{n}_e(\chi)$.

Figure 4.1 depicts the basic spacetime geometry of the SZ effect. The remote dipole field $v_{\rm eff}(\hat{\mathbf{n}}, \chi_e)$ is a projection of the CMB dipole $\Theta_1^m(\hat{\mathbf{n}}, \chi)$ as observed along the past light cone. The dominant contribution is from the line-of-sight component of the peculiar velocity field (as it is for our own observed CMB dipole), although there are subdominant contributions that come from the Sachs-Wolfe (SW), Integrated Sachs-Wolfe (ISW), and primordial Doppler (velocities of the plasma at last-scattering). In this thesis, these dominant and subdominant contributions are referred to as the kinematic or local Doppler dipole, and stationary or fundamental CMB dipole, respectively (see Sec.1.5 for details). The remote quadrupole field $q_{\rm eff}^{\pm}(\hat{\mathbf{n}}, \chi)$ is a projection of the CMB quadrupole $\Theta_2^m(\hat{\mathbf{n}}, \chi)$ as observed along the past light cone. The remote quadrupole receives contributions from both scalar and tensor fluctuations, although we consider only scalar modes in the present context. In this case, $q_{\rm eff}^{+} = q_{\rm eff}^{-}$, and the remote quadrupole is curl-free. As such, we will denote the pure scalar remote quadrupole field as "q_E". The remote quadrupole is sourced by the SW, ISW, and primordial Doppler effects.

SZ tomography works by first inferring the fluctuations in the optical depth in a set of redshift bins labeled by α from a tracer of structure such as a galaxy survey. A quadratic estimator for the bin-averaged dipole and quadrupole fields is then constructed from the CMB temperature or polarization and each redshift bin of the galaxy survey. We work in harmonic space for the reconstructed fields, denoting the moments of the dipole or quadrupole fields in each bin as v_{LM}^{α} and $q_{E;LM}^{\alpha}$ respectively. The reconstruction noise on the remote dipole and quadrupole fields depends on the specifications of the CMB experiment and the volume and shot noise of the galaxy survey. We discuss our assumptions for the reconstruction noise in detail in the following section, which correspond to the choices made in Ref. [19] (see Sec.1.4 and Appendix. A.1 for a review). An additional consideration is the so-called "optical depth degeneracy" (see e.g. [154, 75] and our discussion from Sec.2.4.1), which is a consequence of the necessarily imperfect inference of the fluctuations in the optical depth from the galaxy survey. This manifests itself as an overall multiplicative bias on the remote dipole and quadrupole fields in each redshift bin that must be marginalized over [20]. Direct measurements of the distribution of free electrons,

for example using fast radio bursts as proposed in Ref. [66], can mitigate the optical depth degeneracy.

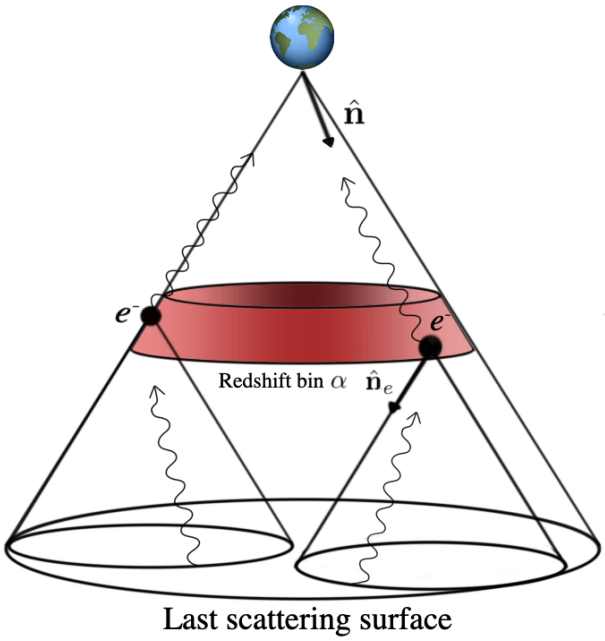

Figure 4.1: Photons travelling from the last scattering surface can be re-scattered by free electrons once the Universe is reionized. The small scale CMB signal generated through this process can be combined with a redshift dependent tracer of the electron density to reconstruct the moments v^{α}_{LM} and $q^{\alpha}_{E;LM}$ of the dipole and quadrupole field.

The remote dipole and quadrupole fields provide new information about the Universe on large scales. The primary CMB photons, travelling to us directly from the last scattering surface, probe the largest accessible scales. The information they provide, however, is somewhat obscured due to the fact that we observe the projection of 3-dimensional inhomogeneities onto a 2-dimensional surface. As illustrated in Fig. 4.1, the remote dipole and quadrupole fields accessed through SZ tomography provide additional information in a number of ways. First, due to the tomographic nature of the reconstruction, we obtain coarse-grained three-dimensional information. Furthermore, the remote dipole and quadrupole fields are sensitive to inhomogeneities *inside our past light cone*, implying that they can access *different* information than what is encoded in the primary CMB temperature. In the case of the remote dipole field, which is dominated by the local peculiar velocity, it is possible to study bulk motion on scales comparable to the size of the observ-

able Universe using long-range correlations.

To go beyond these qualitative remarks, we inspect the scales probed by the remote dipole and quadrupole fields using linear theory, which is a good approximation for the scales under consideration. The various observables under consideration can be related to primordial gravitational potential in Newtonian gauge $\Psi_i(\boldsymbol{k})$ using a set of (bin-averaged) transfer functions $\Delta_\ell^{X,\alpha}(k)$

$$a_{\ell m}^{X,\alpha} = \int \frac{d^3\boldsymbol{k}}{(2\pi)^3} \Delta_\ell^{X,\alpha}(k) \Psi_i(\boldsymbol{k}) Y_{lm}^*(\hat{\boldsymbol{k}}), \tag{4.6}$$

where $X = T, E, v, qE, G$ for the observables we consider in this work (primary CMB temperature, E-mode polarization, remote dipole, remote quadrupole and redshift galaxy distribution, respectively); for $X = T, E$ the index α is superfluous. Expressions for the remote dipole and quadrupole transfer functions, which capture the contributions coming from the SW, ISW and Doppler effects, can be found in Sec.1.5 and Appendix A.2.

The transfer functions reveal to us which scales the remote dipole and quadrupole are sensitive to. As an example, we show in Figure 4.2 the ΛCDM transfer functions (e.g. using parameters from Planck 2018 [4]) in the (ℓ, k) plane for the primary CMB temperature, E-mode polarization and the remote fields at a few different redshifts. For the CMB temperature and remote dipole field, we plot the range $1 \leq \ell \leq 30$, which roughly encompasses the range of scales relevant to the CMB anomalies we consider. For the CMB E-mode polarization and remote quadrupole, we restrict the range to $1 \leq \ell \leq 10$, as this is the range over which the remote quadrupole receives significant support. There are a few things to note from this figure. Comparing with the CMB temperature transfer function, we see that the remote dipole and quadrupole fields have good support over a comparable range of wavenumbers. Because it is sourced mostly by fluctuations near the time of last scattering, the remote quadrupole is relatively more sensitive to large scales than the remote dipole. However, the amplitude of the remote quadrupole falls sharply with ℓ, implying (correctly) that there will be a limited number of measurable modes. It can also be noted that the remote dipole field probes larger scales at higher redshift; this is due to the larger physical distances in the peculiar velocity field which are sampled.

Based on these observations, the observables returned by SZ tomography appear to have the potential to add statistical power into the analysis of CMB anomalies due to their sensitivity to large scale inhomogeneities. However, the amount of *new* information that can be added will depend on the correlations that exist amongst all the observables we consider. Indeed, some correlations are expected to be there by construction. For example, at low redshift the $\ell = 2$ moments of the remote quadrupole field are perfectly correlated with the CMB temperature quadrupole [37, 150, 36].

In Figure 4.3 we plot the correlation coefficient between the remote fields and the primary CMB, defined by

$$r^{X,Y}_{\alpha\beta,\ell\ell',mm'} = \frac{C^{X,Y}_{\alpha\beta,\ell\ell',mm'}}{\sqrt{C^{X,X}_{\alpha\alpha,\ell\ell,mm} C^{Y,Y}_{\beta\beta,\ell'\ell',m'm'}}}, \tag{4.7}$$

where

$$C^{X,Y}_{\alpha\beta,\ell\ell',mm'} = \int \frac{d^3\boldsymbol{k}}{(2\pi)^3} \int \frac{d^3\boldsymbol{k'}}{(2\pi)^3} \, \Delta^{*X,\alpha}_{\ell}(k) \, \Delta^{Y,\beta}_{\ell'}(k') \left\langle \Psi^*_i(\vec{\boldsymbol{k}}) \Psi_i(\vec{\boldsymbol{k'}}) \right\rangle Y_{\ell m}(\hat{\boldsymbol{k}}) Y^*_{\ell' m'}(\hat{\boldsymbol{k}}') \tag{4.8}$$

are the elements of the covariance matrix. In the top panel, we show the correlations between the CMB temperature and remote fields at a few values of ℓ. For $\ell = 1$, the CMB temperature is the aberration-free dipole (see e.g. Refs. [108, 129, 155, 156] for a summary of the various frames for the CMB dipole), not the dipole observed in the Earth's rest frame. There is $\lesssim 10 - 20\%$ correlation between the CMB temperature and the remote dipole field over a range of redshifts and multipoles. There is a far higher correlation between the CMB temperature and the remote quadrupole field. As expected, there is nearly perfect correlation between the CMB quadrupole and the $\ell = 2$ moment of the remote quadrupole field, except at the highest redshifts. The remote dipole field has little correlation with the CMB E-mode polarization. However, at the highest redshift, there is a near-perfect correlation between the E-mode polarization and the remote quadrupole field. This is expected, since at high redshift the CMB polarization is sourced by the same remote quadrupole field that is being reconstructed by SZ tomography (see Ref. [151] for further discussion). Significant correlations will also exists between the galaxy redshift distribution and the reconstructed dipole field, as the latter is dominated by the line of sight peculiar velocity field, which certainly has a tight relationship with the density field itself.

In conclusion, including the full covariance between the various observables can be important in a joint analysis, such as the one we present below. This is particularly important at low-multipoles/low-redshift for the CMB temperature, at high-redshift for the CMB E-mode polarization and in general for the galaxy-velocity correlations. Conversely, we see that over a wide range of multipoles and redshifts, the remote dipole and quadrupole fields carry significant independent information beyond the primary CMB temperature and polarization. For the case of the remote dipole, velocities are sensitive to the gradient of the gravitational potential and thus can probe inhomogeneities on slightly larger scales than the directly observed densities.

4.3 Forecast setup

The reconstruction of the the remote dipole and quadrupole fields using SZ tomography opens the possibility of bringing new information into the study of the large scale CMB anomalies. Determining how informative a data set is will depend, of course, on the type of questions we are trying to answer. From the Bayesian perspective, we might strive for model-selection: How would adding a new observable change the odds-ratio between the anomaly model and ΛCDM in a future experiment (see e.g. [157])? Such an approach requires a motivated set of theoretical priors, as well as an understanding of the full likelihood function over model parameters. Due to the lack of strongly motivated models and the computational complexity of evaluating the full likelihood function, we do not pursue this approach here. Another possibility (less computationally expensive than model selection) for predicting how informative a data set can be is to determine its constraining power on the parameters of a model using a Fisher matrix-based approach. In general, such results are not sufficient to decide if a future experiment could choose among competing theoretical models. However, this approach does offer a way of quantifying the additional constraining power a new observable might add. We will adopt this methodology in order to study the information content of the remote dipole and quadrupole field on a series of models for CMB anomalies, and compare it to what is achievable using the primary CMB temperature, E-mode polarization and galaxy clustering on large angular scales. We now proceed to describe our definitions for information content and modelling for the signal and noise covariance.

4.3.1 Fisher analysis, figure of merit and parameter space

Given a cosmological model with parameters $\{\lambda_i\}$, one can forecast how well these parameters can be constrained using different set of observables by implementing a Fisher matrix analysis. The elements of the Fisher matrix are given by the following expression:

$$\mathbf{F}_{ij} = \frac{1}{2}\mathrm{Tr}[(\mathbf{C}+\mathbf{N})^{-\mathbf{1}}\mathbf{C}_{,\lambda_\mathbf{i}}(\mathbf{C}+\mathbf{N})^{-\mathbf{1}}\mathbf{C}_{,\lambda_\mathbf{j}}] \tag{4.9}$$

where $\mathbf{C}$ is the signal covariance matrix (whose elements are defined by Eq. 4.8), $\mathbf{C}_{,\lambda_\mathbf{j}}$ denotes its derivative respect to λ_j, and $\mathbf{N}$ is the noise covariance matrix. The Fisher matrix encodes information about the curvature of the likelihood function around its maximum in parameter space, and this information can be turned into fully marginalized constraints on the model parameters:

$$\sigma_{\lambda i} = \sqrt{\left(\mathbf{F}^{-1}\right)_{ii}}. \tag{4.10}$$

A good way of comparing the results of the Fisher matrix analysis among models with different number of parameters is to define a single quantity that condenses the information of how well the model parameters are constrained. For a forecast using a set $\{X\}$ of observables, we define the figure of merit FoM for a subset of N parameters:

$$\mathrm{FoM}(X) = \left(\frac{1}{\sqrt{\det(\mathbf{F_N^{-1}})}} \right)^{\frac{1}{N}}. \tag{4.11}$$

where $\mathbf{F_N^{-1}}$ is the part of the inverse Fisher matrix containing information about the fully marginalized constraints on the subset of parameters. Furthermore, since we want to highlight the relative performance respect to the primary CMB, we will express our results in terms of a figure of merit ratio, defined by :

$$\mathrm{FoMr}(X) = \frac{\mathrm{FoM}(X)}{\mathrm{FoM}(T)} \tag{4.12}$$

The figure of merit ratio encodes the geometrical mean improvement on model parameter constraints. Similar figure of merit ratios have been used in previous literature, e.g. as a measure of improvement in constraints on CMB anomalies parameters when comparing current to future missions [158].

Apart from the standard ΛCDM parameters and the extra parameters a_i present in any of its extensions, it is important to marginalize over other "nuisance" parameters that can be introduced depending on our set of observables $\{X\}$. For the case of the remote fields, we need to include a bias parameter b_α^d multiplying the multipole coefficients $a_{LM}^{v,\alpha}$ and $a_{LM}^{qE,\alpha}$ in each bin α due to the optical depth degeneracy in kSZ/pSZ : having incomplete inference of the electron-galaxy cross spectrum, we can only reconstruct the remote dipole and quadrupole inside each redshift bin up to an overall amplitude. We refer the reader to [20] for a more detailed discussion of the optical depth degeneracy. Modelling of galaxy clustering also involves the introduction of other nuisance parameters: galaxies are biased tracers of the dark matter distribution, so its necessary to include a galaxy bias $b^G(z)$ which will depend on the details of the galaxy survey and can also be marginalized over. On large scales, there are as well contributions to the observed galaxy number counts which come from redshift spatial distortions, lensing and GR effects on very large scales. These contributions have their own internal biases (intrinsic alignment bias, evolution bias, magnification bias) for which we provide references in the following section. The fiducial values for the ΛCDM cosmology we take are $\Omega_b = 0.049$, $\Omega_c = 0.263$, $h = 0.675$, $\tau = 0.054$, $10^9 A_s = 2.096$ and $n_s = 0.965$. Fiducial values for each of the anomalies model parameters

are presented together with the models in the next section and Appendix C.1. The bias parameters b^d_α are assigned fiducial values of unity. Priors on b^d_α can come from other astrophysical probes [66], but for the anomaly models under consideration we find that the constraints are relatively insensitive to the addition of such a prior. Parameters b^G_α are obtained by averaging the galaxy bias $b^G(z)$ over the survey redshift bins. The fiducial values we use for this bias are $b^G(z) \simeq 0.95 + 0.67z$ as quoted in the LSST science book [10].

4.3.2 Modelling of the signal

The signal covariance matrix $\mathbf{C}$ we construct is split into two pieces: $\mathbf{C}^{\rm low}$ and $\mathbf{C}^{\rm high}$. For multipoles $\ell \leq 60$, corresponding to $\mathbf{C}^{\rm low}$, we investigate different combinations of all observables under consideration, accounting for the auto and cross-correlations between the primary CMB temperature and E-mode polarization, remote dipole/quadrupole fields and galaxy number counts. For multipoles $60 < \ell < \ell_{\rm high}$, where the reconstruction of the remote dipole and quadrupole fields is poor, and where correlations with the primary CMB are vanishingly small, we include only the CMB temperature and polarization (and their covariance) in $\mathbf{C}^{\rm high}$. We choose $\ell_{\rm high} = 3000$ since for higher multipoles the primary CMB becomes a subdominant contribution to the measured microwave sky. With these assumptions, the Fisher matrix factorizes into the sum of a low-ℓ and high-ℓ piece, $\mathbf{F}^{\rm low}$ and $\mathbf{F}^{\rm high}$, respectively. The main effect of $\mathbf{F}^{\rm high}$ is to constrain the ΛCDM parameters. We further assume that $\mathbf{F}^{\rm high}$ is zero for the entries corresponding to the anomaly model parameters a_i, since the anomaly models under consideration will have little to no effect on these scales. The elements of the covariance matrix are computed using Eq. 4.8. We obtain the primary CMB temperature and polarization transfer functions using the publicly available Code for Anisotropies in the Microwave Background (CAMB) [81]. We use the same conventions and definitions for the remote fields transfer functions as in [19] (Sec.1.5 and Appendix A.2), and for the galaxy number counts we follow what is done in related work [76] (Appendix A) where dipole reconstruction is put together with galaxy number counts to determine if large scale general relativistic effects could be detected with near future surveys.

4.3.3 Modelling of the noise

The noise covariance matrices for the primary CMB fields and for galaxy number counts are constructed under the assumption of idealized next generation CMB experiments and

galaxy surveys. We assume CMB temperature and polarization data on the full sky with noise

$$\mathbf{N}_\ell^{\rm CMB} = \eta^2 \exp\left(\frac{\theta^2_{\rm FWHM}}{8{\rm log}2}\ell(\ell+1)\right). \tag{4.13}$$

We choose fiducial values of $\eta = 1\mu K$-arcmin and $\theta_{\rm FWHM} = 1$ arcmin, representative of Stage-4 CMB-like experiments, and explore how constraints vary for larger ($5\mu K$-arcmin), smaller ($0.1\mu K$-arcmin) noise and increased beam size ($\theta_{\rm FWHM} = 5$ arcmin). To determine the galaxy survey shot-noise, we assume an Rubin-like number density per arc-minute square given by :

$$n(z) = n_{\rm gal}\frac{1}{2z_0}\left(\frac{z}{z_0}\right)^2 \exp\left(-z/z_0\right). \tag{4.14}$$

where $z_0 = 0.3$ and $n_{\rm gal} = 40$ arcmin^{-2}. In addition to the shot noise we also account for photometric calibration errors, which introduce additional biases to the large galaxy power spectrum, and we base our parametrization on previous studies on the impact of these systematics in the reconstruction of the ISW effect from galaxy surveys [100, 99]. Calibration errors introduce extra power $C_\ell^{\rm cal}$ which can be modelled by:

$$C_\ell^{\rm cal} = A_{\rm cal} e^{-(\ell/10)^2} \tag{4.15}$$

for multipoles $\ell < 30$ and zero otherwise, where $A_{\rm cal}$ is a normalization constant chosen such that the variance of the calibration error field defined on the sky is equal to a desired value. We choose this variance to be equal to 10^{-4} which corresponds to a level of calibration of ≈ 0.01 magnitudes and is a rather conservative value for what is expected from future surveys. We do not marginalize over the calibration error parameters.

For the remote fields, the reconstruction noise is computed as in Ref. [19]. The instrumental CMB noise discussed above is one of the necessary pieces to calculate the reconstruction noise for the dipole and quadrupole correlations appearing in $\mathbf{C}^{\rm low}$. Multipoles up to $\ell = 9000$ (assumed to be accessible with next generation experiments) are used to calculate the reconstruction noise, and thus different choices of the CMB noise level η and beam size $\theta_{\rm FWHM}$ will have an impact on the signal to noise for the low-ℓ dipole and quadrupole fields. The galaxy shot noise enters into the calculation in a similar way, but we have kept the parameters appearing in Eq.4.14 fixed in order to focus on improvements on constraints due to the reconstructed fields and their dependence on CMB noise parameters.

The construction of both the signal and noise covariance matrices also involves a choice of redshift binning for the galaxy survey, which determines how coarse-grained the reconstructed remote dipole and quadrupole fields are. The thinner the binning is, the more

information can be collected. Clearly all independent information would be captured in the limit of having infinitely small redshift bins, but the redshift resolution of the (photometric redshift) surveys used in the reconstruction process imposes a limit on how many redshift bins can be used. We use 45 redshift bins of equal comoving radial width between $Z = [0.1, 3]$, which, translating to redshift, is roughly consistent with the expected photometric redshift errors for Rubin [10].

4.4 Information content forecast

In what follows we present the Fisher forecast for constraints on models of the large scale CMB anomalies using different subsets of observables: $X = (T, E)$, $X = (T, E, R)$, $X = (T, E, G)$ and $X = (T, E, R, G)$. We first introduce two general classes of physical models: models that break statistical isotropy, which could be responsible for the power asymmetry, and models that deviate from a nearly scale-invariant primordial power spectrum, which could be responsible for a lack of power on large scales and a feature in the power spectrum at $\ell \sim 20 - 30$. While this is a small subset of physical models considered to explain only a subset of the CMB anomalies, we hope that the cases we do consider are illustrative of the potential utility of SZ tomography for providing further insight into the nature of the CMB anomalies. Following this we present the figure of merit ratio FoMr given by Eq. (4.12) for each model, which provides a quantitative measure of the overall improvement on parameter constraints relative to what is achievable with measurements of the primary CMB temperature only. The fully marginalized parameter constraints for each model can be found in Appendix C.1.

4.4.1 Statistical isotropy breaking

A subset of the observed CMB anomalies suggest the existence of statistical anisotropies [159]: unexpected alignment between the low multipole moments, a hemispherical power asymmetry, parity asymmetry of the CMB etc. It is still not known whether or not these features are due to foregrounds, local cosmic structure or possible statistical flukes present our observed realization of ΛCDM. However, if due to true physical departures from ΛCDM, the underlying model must break statistical isotropy.

We consider phenomenological models of spontaneous isotropy breaking [160] (see also e.g. [161, 162, 163, 164]), in which local observers would detect statistical anisotropy, while the Universe as a whole is globally statistically homogeneous and isotropic. Following

Ref. [132], we include a field $h(\vec{\boldsymbol{x}})$ with super horizon fluctuations that modulates the potential $g_1(\vec{\boldsymbol{x}})$ only on large scales, leaving small scale fluctuations described by $g_2(\vec{\boldsymbol{x}})$ unaffected:

$$\Psi_i(\vec{\boldsymbol{x}}) = g_1(\vec{\boldsymbol{x}})(1 + h(\vec{\boldsymbol{x}})) + g_2(\vec{\boldsymbol{x}}), \tag{4.16}$$

Here, $g_1(\vec{\boldsymbol{x}})$ and $g_2(\vec{\boldsymbol{x}})$ are random Gaussian fields, while $h(\vec{\boldsymbol{x}})$ is deterministic within our Hubble volume. It is the slow variation of $h(\vec{\boldsymbol{x}})$ inside our Hubble volume that is responsible for the existence of statistical anisotropy in the CMB. Such a modulation can occur, for example, in inflation models with more than one field contributing to the primordial curvature perturbations. Here, rather than focusing in a particular early Universe mechanism for generating the modulation, we are interested in determining how the imprints of a preferred direction on the remote dipole and quadrupole fields can help to constrain the amplitude of the modulation.

The effect on the primordial power spectrum is given by:

$$\begin{aligned}\left\langle \Psi_i^*(\vec{\boldsymbol{k}})\Psi_i(\vec{\boldsymbol{k'}})\right\rangle &= (2\pi)^3\delta(\vec{\boldsymbol{k}} - \vec{\boldsymbol{k'}})(P_{g_1}(k) + P_{g_2}(k')) \\ &+ (P_{g_1}(k) + P_{g_1}(k'))h(\vec{\boldsymbol{k}} - \vec{\boldsymbol{k'}}) \\ &+ \int \frac{d^3\tilde{k}}{(2\pi)^3} P_{g_1}(\tilde{k})h(\vec{\boldsymbol{k}} - \vec{\tilde{\boldsymbol{k}}})h(\vec{\boldsymbol{k'}} - \vec{\tilde{\boldsymbol{k}}}),\end{aligned} \tag{4.17}$$

where $P_{g_1}(k)$ and $P_{g_2}(k)$ are the power spectra for $g_1(\vec{\boldsymbol{x}})$ and $g_2(\vec{\boldsymbol{x}})$. We will fix these power spectra to that $P_{g_1}(k) + P_{g_2}(k) = P_{\Lambda\mathrm{CDM}}(k)$ when $h(\vec{\boldsymbol{x}}) = 0$. The second and third term will induce couplings between different (ℓ, m) multipoles and this manifests the breaking of statistical isotropy for local observers.

For this work consider a dipolar modulation given by a super-horizon scale mode varying in the direction of the Z axis:

$$h(\vec{\boldsymbol{x}}) = A\frac{\sin(\vec{\boldsymbol{k}}_0 \cdot \vec{\boldsymbol{x}})}{k_0\,\chi_{dec}} \approx A\frac{Z}{\chi_{dec}} \tag{4.18}$$

This physical model could explain the observed power asymmetry [132] (see also e.g. [165, 141, 137]); we do not consider other modulation models here, which could be responsible for the observed alignment of low-ℓ multipoles (see e.g. [132]) or other observed CMB anomalies. For the modulating field Eq. 4.18, expressions for elements of the covariance matrix up to second order in A can be obtained analytically and are presented in Appendix C.2. Analysis of temperature data by Planck suggests a phenomenological dipole modulation up to $\ell \sim 60$ with a value for the amplitude parameter of approximately [159] $A = 0.07 \pm 0.02$. An open question is to what scales the asymmetry might persist, e.g.

where the cross-over occurs from the observed fluctuations being sourced by g_1 to being sourced by g_2. There is a hard upper bound implied by the low hemispherical asymmetry of the distribution of high redshift quasars [166] of $k \lesssim 1$ Mpc^{-1}. In the following we treat the cross-over to statistical homogeneity as in Ref. [141], were a new set of parameters describing the cross-over is introduced:

$$P_{g_1}(k) = \frac{1}{2} P_{\Lambda\mathrm{CDM}}(k) \left[1 - \tanh\left(\frac{\ln k - \ln k_c}{\Delta \ln k} \right) \right], \tag{4.19}$$

with k_c a cutoff scale with fiducial value $7.83 \times 10^{-3}\mathrm{Mpc}^{-1}$ and $\Delta \ln k = 0.5$ a parameter controlling the steepness of the cross-over.

4.4.2 Deviations from ΛCDM power law

The other class of models we consider involves possible deviations from the ΛCDM power law primordial power spectrum. On large angular scales it has been observed by WMAP and Planck that the CMB temperature shows an unexpected lack of variance compared to ΛCDM. Features in the temperature power spectrum have also been identified, remarkable ones being a low quadrupole and a lack of power at multipoles $\ell \sim 20-30$. One simple and theoretically interesting possibility is that these CMB anomalies are due to corresponding features in the primordial power spectrum of curvature fluctuations. Such features can arise as a signature of: the onset of inflation (e.g. [167, 168, 169]), oscillations [170, 171] or sharp steps [172] in the inflaton potential, steps in the sound speed [173] or DBI inflation warp factor [174], among other scenarios. In this section, we determine the additional constraining power offered by SZ tomography for a subset of these feature models, choosing a few representative examples that have previously been investigated by Planck [175].

Phenomenological models for large scale power suppression

Following Ref. [175], we consider a set of two-parameter phenomenological models for the suppression of power on large angular scales in the primary CMB temperature. The first model we consider [169] implements an exponential suppression of power below a wavenumber k_c:

$$P_{s_1}(k) = P_{\Lambda\mathrm{CDM}}(k) \left(1 - \exp\left[- \left(\frac{k}{k_c} \right)^{\lambda} \right] \right) \tag{4.20}$$

where $P_{\Lambda\mathrm{CDM}}(k) = A_s\left(\frac{k}{k_*}\right)^{n_s-1}$ and the best-fit model parameters from Planck 2015 [175] are $k_c = 3.74 \times 10^{-4}$ Mpc^{-1} and $\lambda = 0.53$. The second model has a break in the power law at a scale k_b:

$$P_{s_2}(k) = \begin{cases} A_s\left(\frac{k_b}{k_*}\right)^{-\delta}\left(\frac{k}{k_*}\right)^{n_s-1+\delta} & \text{if } k \leq k_b, \\ A_s\left(\frac{k}{k_*}\right)^{n_s-1} & \text{if } k \geq k_b, \end{cases} \tag{4.21}$$

The best-fit model parameters from Planck 2015 [175] are $k_b = 5.26 \times 10^{-4}$ Mpc^{-1} and $\delta = 1.14$. In both cases, we choose the central values for the model parameters as the best-fit Planck values and we fix the pivot scale to $k_* = 0.05$ Mpc^{-1}. We plot the two fiducial models in Fig. 4.4.

Features in the power spectrum

We now review two physical scenarios that give rise to features in the primordial power spectrum. In the first model, we consider a period of kinetic domination preceding slow-roll inflation. This gives rise to a suppression of power on large scales, as well as oscillations in the power spectrum on intermediate scales [169]. The one parameter in this model is a scale k_c, roughly corresponding to the comoving size of the horizon when slow-roll begins. Clearly, we are able to constrain this model only when there are a minimal number of e-folds of inflation, in which case k_c is on observable scales. The full form of the power spectrum is given by:

$$\ln P_c(k) = \ln P_0(k) + \ln\left(\frac{\pi}{16}\frac{k}{k_\mathrm{c}}\left|C_\mathrm{c} - D_\mathrm{c}\right|^2\right), \tag{4.22}$$

where

$$\begin{aligned} C_\mathrm{c} &= \exp\left(\tfrac{-ik}{k_\mathrm{c}}\right)\left[H_0^{(2)}\left(\tfrac{k}{2k_\mathrm{c}}\right) - \left(\tfrac{k_\mathrm{c}}{k} + i\right)H_1^{(2)}\left(\tfrac{k}{2k_\mathrm{c}}\right)\right], \\ D_\mathrm{c} &= \exp\left(\tfrac{ik}{k_\mathrm{c}}\right)\left[H_0^{(2)}\left(\tfrac{k}{2k_\mathrm{c}}\right) - \left(\tfrac{k_\mathrm{c}}{k} - i\right)H_1^{(2)}\left(\tfrac{k}{2k_\mathrm{c}}\right)\right], \end{aligned} \tag{4.23}$$

and $H_n^{(2)}$ denotes the Hankel function of the second kind. We assume the best-fit value from Planck 2015 [175] of $k_c = 3.63 \times 10^{-4}$ Mpc^{-1} as the central values in our analysis below. The second model we consider arises when there is a tanh-shaped step in the inflaton potential as in Ref. [176], which gives rise to oscillations in the primordial power spectrum. This is a three-parameter model, which, at the level of the inflaton potential, correspond to the location, height, and width of the step. The resulting power spectrum is given by

$$\ln P_s(k) = \exp\left[\ln P_0(k) + I_0(k) + \ln\left(1 + I_1^2(k)\right)\right] \tag{4.24}$$

where

$$I_0(k) = \frac{\mathcal{A}_{\rm s}}{2x^3}\left[\left(18x - 6x^3\right)\cos 2x + \left(15x^2 - 9\right)\sin 2x\right]\Big|_{x=(k/k_{\rm s})} \times \mathcal{D}\left(\frac{k/k_{\rm s}}{x_{\rm s}}\right) \tag{4.25}$$

$$I_1(k) = \frac{1}{\sqrt{2}}\Bigg[\frac{\pi}{2}\left(1 - n_{\rm s}\right) - \frac{\mathcal{A}_{\rm s}}{x^3} \tag{4.26}$$

$$\times \left\{3(x\cos x - \sin x)\left[3x\cos x + \left(2x^2 - 3\right)\sin x\right]\right\}\Big|_{x=(k/k_{\rm s})} \times \mathcal{D}\left(\frac{k/k_{\rm s}}{x_{\rm s}}\right)\Bigg]. \tag{4.27}$$

$$\mathcal{D}(x) = \frac{x}{\sinh x} \tag{4.28}$$

$$\tag{4.29}$$

We assume the best-fit value from Planck 2015 [175] of $\mathcal{A}_s = 0.374$, $k_s = 7.94\times10^{-4}$ Mpc^{-1}, and $x_s = 1.41$. We plot the power spectra for the two models in Fig. 4.5. As commented on in Ref. [175], for these choices of parameters both of these models give rise to a deficit in the CMB temperature power spectrum at $\ell \sim 20 - 30$, similar to what is observed.

4.4.3 Results

In Table 4.1 we present the results of our multi-tracer analysis. For each model we show the FoMr that captures how constraints on model parameters improve when the observable set is extended beyond the primary CMB temperature T. For those cells that correspond to a set including the remote fields R we have included two values: the FoMr when the small scale CMB noise has beam size $\theta_{\rm FWHM}$ equal to 5 arcminutes and 1 arcminute. This further emphasizes the dependence of the reconstruction signal to noise ratio on the instrumental noise of the high resolution CMB experiment.

The results displayed on Table 4.1 tell us a consistent story across the different models. The inclusion of the E mode polarization into the analysis of CMB anomalies brings substantial constraining power, here represented by a FoMr ranging from 1.5 to 3.0. Beyond this, the following two questions are important to address: How do constraints improve if we add the remote dipole and quadrupole fields, and how does this improvement relate to what can be done by instead choosing galaxy clustering on large scales as the additional observable? We can see that depending on the noise conditions for the small scale CMB experiment, the remote fields compete with the large scale galaxy clustering in improving the FoMr beyond the (T, E) combination, with a tendency of the galaxies to be superior than the dipole and quadrupole except for very sensitive CMB experiments. However, the remote fields contain a significant amount of new information, as evidenced by the

Model	Noise [μ-Karcmin]	FoMr(T,E)	FoMr(T,E,R)	FoMr(T,E,G)	FoMr(T,E,R,G)
Dipolar modulation	5.0	1.48	1.60 - 2.37	2.90	2.90 - 3.02
	1.0	1.48	1.87 - 2.67	2.90	2.93 - 3.13
	0.1	1.48	2.27 - 2.72	2.90	3.00 - 3.17
Exponential suppression	5.0	1.46	1.48 - 1.63	1.69	1.69 - 1.75
	1.0	1.46	1.55 - 1.70	1.69	1.72 - 1.81
	0.1	1.46	1.63 - 1.78	1.69	1.76 - 1.86
Broken power law	5.0	1.69	1.76 - 1.79	1.90	1.94 - 1.98
	1.0	1.69	1.78 - 1.83	1.90	1.97 - 1.99
	0.1	1.69	1.82 - 1.92	1.90	1.99 - 2.02
Cutoff	5.0	2.05	2.14 - 2.46	2.60	2.63 - 2.74
	1.0	2.05	2.27 - 2.63	2.60	2.68 - 2.82
	0.1	2.05	2.46 - 2.77	2.60	2.75 - 2.92
Step in inflaton potential	5.0	3.09	3.25 - 3.81	3.88	3.89 - 4.08
	1.0	3.09	3.50 - 4.07	3.88	3.98 - 4.24
	0.1	3.09	3.78 - 4.18	3.88	4.09 - 4.31

Table 4.1: Figure of merit ratio for different combinations of observables. We present results for different levels of CMB noise. Columns that involve the remote fields (R) show two values corresponding to CMB experiments with $\theta_{\rm FWHM} = 5 - 1$ arcmin. Columns involving only combinations of T, E, and G do not depend on the details of the small scale CMB noise levels and beams and therefore show unchanged values.

last column of Table 4.1. Adding the dipole and quadrupole fields on top of (T, E, G) can increase the FoMr by $\sim 10\%$ percent. Given that observations of the large scale polarization or galaxy clustering can be limited by different systematics and foregrounds than the high angular resolution measurements necessary for SZ tomography and vice versa,the remote dipole and quadrupole are an "extra handle" to work with when other tracers fall short, and which has prospects of becoming increasingly informative as detector technology keeps developing.

The FoMr gives us a general idea of how constraints improve for different datasets, but its also useful to look directly at parameter constraints (Appendix C.1) to see how helpful these combinations of observables can be in the task of determining the physical nature of the CMB anomalies. Keeping in mind that ultimately only a full likelihood analysis will reveal if departures of ΛCDM are favoured by data, we can see that the full set of observables here considered can push marginal 1σ-3σ constraints to tighter ones. For the dipolar modulation model, constraints on the amplitude of modulation can be pushed from

3σ to 8σ and for the cross-over scale constraints can be pushed from the 1σ to the 3σ level. For the generic suppression models constraints on the characteristic scale of suppression can be pushed from the 1σ to 2σ level, and in the case of the step model the amplitude $\mathcal{A}_s$ of the step and the parameter x_s controlling its shape see constraints jumping from marginality to above the 4σ level. These results indicate a much more optimistic possibility of studying the CMB anomalies beyond what the primary CMB temperature allows for. The reconstructed dipole and quadrupole will therefore form part of a set of observables that can provide stronger evidence to favour or rule out new early Universe physics.

4.5 Conclusions

Determining whether or not the observed large angular scale anomalies in the CMB are indications of physics beyond ΛCDM is a matter of great interest and intense debate. Faced with the obstacle imposed by cosmic variance on our study of the largest scales in the Universe, we are driven to analyze datasets that incorporate additional observables on top of the primary CMB temperature in order to favour or rule out the different hypotheses for the origin of the anomalies.

In this chapter we explored the constraining power on CMB anomalies models provided by a new set of observables: the remote CMB dipole and quadrupole fields. These fields, which correspond to the projected $\ell = 1, 2$ moments of the microwave sky as measured at different locations in the Universe, can be reconstructed using SZ tomography. The remote dipole and quadrupole fields carry 3-dimensional information about large scale fluctuations in the Universe, and a significant number of independent modes can in principle be reconstructed from next-generation CMB and galaxy surveys. This additional information is largely independent of the primary CMB, and can therefore offer more statistical power for the analysis of physical models of large scale CMB anomalies.

Our methodology consisted of deriving forecasted constraints on a series of anomalies models using different combinations of observables probing the largest scales, including the primary CMB temperature (T), E-model polarization (E), galaxy clustering on large-scales (G) and the remote dipole and quadrupole fields (R). The improvement on constraints relative to what is achievable with the primary CMB temperature serves as a measure of how informative additional observables can be, and this was expressed in terms of appropriate figure of merit ratios (FoMr); see Eq. 4.12. We assumed access to data on the full sky, with no systematics (aside from CMB instrumental noise, photometric redshift errors, and photometric calibration errors) or foregrounds, in each of our forecasts.

We considered a series of representative models for CMB anomalies, capturing departures from statistical isotropy and the standard inflationary ΛCDM primordial power spectrum. These modifications have been considered as possible explanations for the statistically marginal power asymmetry on the CMB sky, as well as features and lack of power in the temperature power spectrum on large angular scales. Based on our analysis, we can make a number of general statements about the utility of SZ tomography for addressing the possible physical nature of the CMB anomalies. As many previous analyses have shown [132, 133, 134, 135, 136, 137, 138, 139], E-mode polarization has been identified as a powerful discriminator for physical models of CMB anomalies. We found that one can go beyond the (T, E) combination by adding the reconstructed remote dipole and quadrupole field, as it was expected due to their 3-dimensional nature. A comparable amount of information can also be accessed using a more commonly considered 3-dimensional probe: galaxy clustering on our past lightcone. Because of this, we explored the ability of the remote fields to improve over the (T, E, G) constraints, finding that typically one can achieve higher FoMrs up to ten percent, with better results for more sensitive CMB experiments. This means that there is new information on the large scale Universe that becomes available through SZ tomography. Overall, our results suggest that the remote dipole/quadrupole fields could play the role of an alternative and complementary probe of the CMB anomalies, affected by different systematics or foregrounds that make difficult a cosmic variance limited measurement of large scale E-mode polarization or galaxy clustering. SZ tomography also offers the possibility to systematically improve constraints on CMB anomaly models in the current era of rapidly evolving high-resolution, low-noise CMB experiments. All together, the observable set that we considered here was able to push several marginal parameter constraints on anomaly models to above the 3σ level and even higher. This indicates that next-generation CMB experiments and galaxy surveys will allow for an enhanced testing of the nature of the large scale CMB anomalies.

Our analysis has a number of shortcomings. First, our Fisher-based analysis is insensitive to the shape of the likelihood function, which can deviate significantly from a Gaussian for many of the models considered here. A future investigation could improve upon this by sampling the full likelihood function; however given the size of the covariance matrix including all observables and the dimensionality of the parameter space, there will be computational challenges for doing so. Future analyses should also incorporate realistic foregrounds and systematics in the CMB and galaxy surveys, and investigate their impact on the reconstruction of the remote dipole and quadrupole fields. In addition, the effects of masking should be taken into account, which will degrade the information available on the largest angular scales. Despite these limitations, our analysis highlights there is useful information on the physical nature of the observed CMB anomalies that is *in principle*

accessible using SZ tomography. This provides a useful target for future analyses and observations.

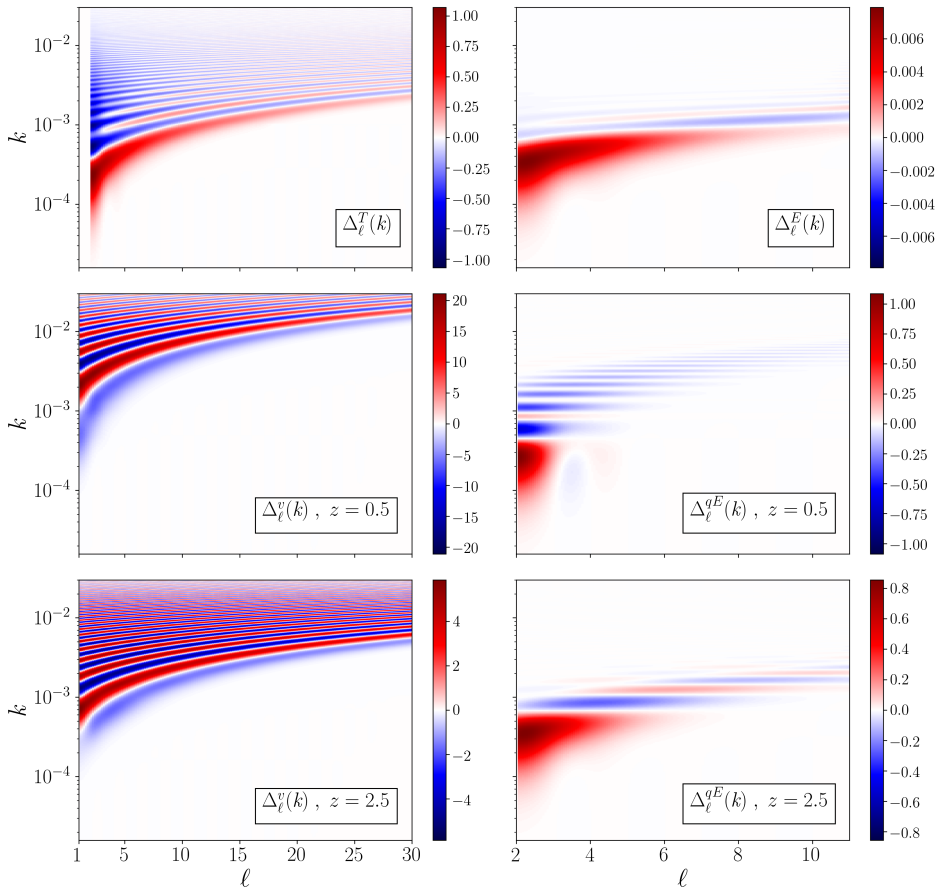

Figure 4.2: On the top panels, the transfer functions for the primary CMB temperature ($\ell = 1$ is not plotted here) and E-mode polarization. On the middle and bottom, the bin averaged transfer functions for the remote dipole (left) and quadrupole (right) for bins centered on redshifts $z = 0.1$ and $z = 2.5$. The binning scheme used for this figure consisted on 60 bins of equal comoving size between $0.1 \leq z \leq 6$.

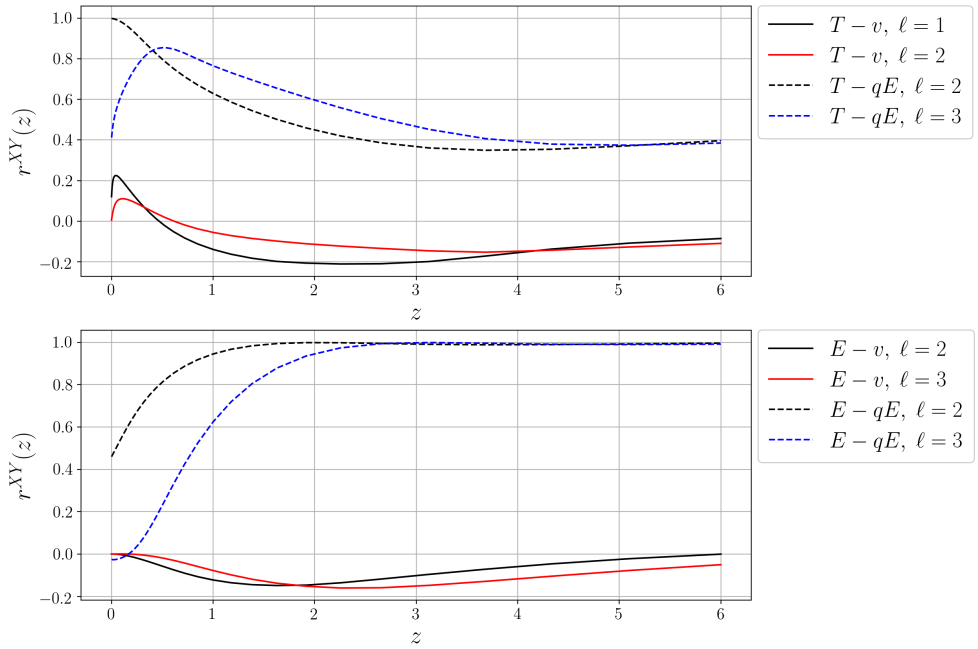

Figure 4.3: Correlation coefficient between the primary CMB fields and the remote dipole an quadrupole fields. As expected, the $\ell = 2$ moment of the primary CMB temperature is perfectly correlated with the very low redshift remote quadrupole (top panel, black dashed line) and the remote dipole captures the primordial contributions to the $\ell = 1$ aberration-free CMB dipole measured at $z = 0$ (top panel, black solid line). The remote quadrupole exhibits longer range correlations with the primary CMB than the remote dipole does (bottom panel).

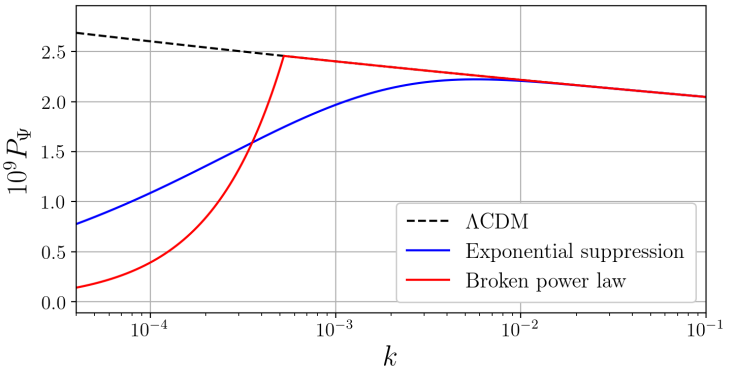

Figure 4.4: Primordial power spectrum for the exponential suppression model and the broken power law model together with the standard ΛCDM spectrum. The first model shows suppression starting at scales of several hundred Mpc while the second one deviates from the standard power law on scales of several Gpc.

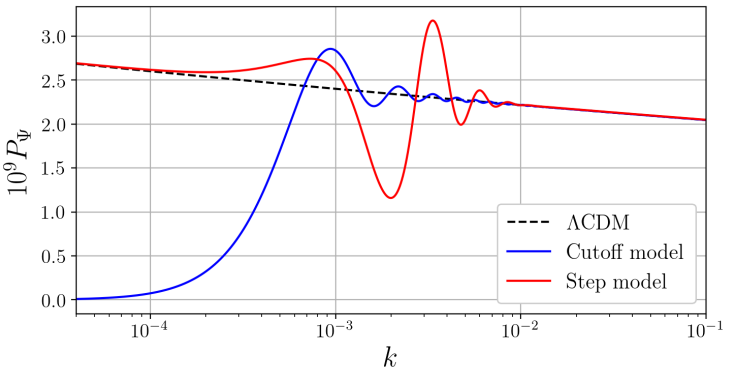

Figure 4.5: Primordial power spectrum for the cutoff model and the step model together with the standard ΛCDM spectrum.

Chapter 5

Conclusions

In this thesis, we have presented advances in the development and use of kSZ velocity reconstruction, a technique that proposes using next-generation measurements of the CMB anisotropies from the kinetic Sunyaev-Zel'dovich effect in combination with tracers of the large-scale structure of the Universe to reconstruct a new observable to do cosmology with: the remote CMB dipole. The research discussed in this thesis stands as a continuation of the foundational work presented in [18, 19], which introduced an approach to kSZ velocity reconstruction referred to as the "Lightcone Picture". In Chapter 1 we discussed the basic features of this approach and pointed out a series of avenues to improve the modelling of the technique and account for more realistic experimental scenarios. Later in the same chapter, we took a closer look at the remote CMB dipole and argued for its use as a probe of cosmological perturbations on scales of several Gpc.

Chapter 2 presented an extended formalism for the Lightcone Picture of kSZ velocity reconstruction which expands on the results from [19] in several ways:

1. First, additional sources of bias and noise for the reconstructed dipole were identified. It was found that photometric redshift errors in the galaxy survey, which cause a mixing of the radial information along the lightcone, lead to biases in the reconstructed remote dipole; for a given photometric redshift bin, the estimated remote dipole moments carry information about the dipole in neighbouring bins. This mixing of radial information can be undone if a good enough model for the photometric redshift errors is available. The multiplicative bias coming from the improper modelling of the galaxy-electron correlation on small scales, known as the optical depth bias, was also described for the extended Lightcone Picture formalism using a simple toy model. We found that this bias is largely scale independent for the reconstruction

of the remote dipole and we discussed how to model the uncertainty introduced by it in cosmological parameter constraints analyses. Additional sources of bias coming from calibration errors for the galaxy surveys on large angular scales, CMB lensing, and the non-linear ISW signal were also studied and it was determined that only the first of these can lead to noticeable reductions in the total signal to noise for the reconstructed dipole. We also found that the coarse graining of radial information can lead to a significant error in the remote dipole reconstruction if wide enough redshift bins are used, a feature we refer to as the "fine mode" noise. We characterized the size of this new noise term and determined the level of radial resolution that is needed to safely ignore it.

2. Second, temperature foregrounds from the CIB and tSZ effect were included at multiple frequency bands, as well as a more realistic instrumental noise model consistent with the properties of the Simons Observatory Large Area Telescope [15]. The information from multiple frequency bands can be used to clean the temperature foregrounds before implementing the quadratic estimator. We found that the cleaning procedure is good enough to still produce high signal to noise reconstructions of the remote dipole, but lower than those predicted in [19].

3. Third, we tested the validity of the estimator using a suite of Gaussian simulations. We implemented a method to construct properly correlated Gaussian maps of all the fields involved in kSZ reconstruction and demonstrated good agreement with theory expectations. The suite of Gaussian simulations also allowed us to test the reconstruction in the presence of a mask. Using appropriate real space versions of the quadratic estimator we showed that, apart from some effects at the edge of the mask (which can be easily removed), the reconstruction performs well in unmasked regions of the sky.

In Summary, our analysis shows that forecasts for kSZ velocity reconstruction in the Lightcone Picture are still promising in the presence of more realistic experimental conditions. The extended formalism has been developed with the knowledge that it will soon play an important role for achieving the SO collaboration goal to constrain the local primordial non-Gaussianity parameter f_{NL} to the $\sigma(f_{NL}) \approx 1$ level using the reconstructed radial velocity field [15]. The extended formalism presented in this chapter is general enough to be applied to the reconstruction of other observables, such as the transverse velocity field which sources the moving lens effect. The tools presented in this chapter have been implemented in ReCCO, a publicly available code which will enable the broader research community to further explore the potential uses of velocity reconstruction and other similar methods.

Chapter 3 presented the results from N-body simulations of kSZ velocity reconstruction on the lightcone. In order to construct properly correlated mock data on scales ranging from a few Mpc to several Gpc, we developed a "Box in box" method to combine small-scale perturbation modes evolved with an N-body code and large-scale perturbation modes evolved using linear theory. One of the objectives of this study was to test the robustness of the remote dipole estimator against the effects of redshift space distortions, gravitational lensing and non-linear evolution of structure. To concentrate on these features, we assumed electrons trace dark matter, ignored the effect of photometric redshift errors, and did not include CMB temperature foregrounds. Additionally, the suite of simulations allowed us to study the stationary contributions to the remote dipole described in Sec.(1.5). We found that high fidelity reconstructions of the low-ℓ multipole moments of the remote dipole are achievable for redshifts $z \gtrsim 0.2$, and that the reconstruction improves as we approach our redshift limit of $z = 0.37$. For redshifts $z \lesssim 0.2$ the fidelity decreases, which we attribute to non-linear effects. We also demonstrated that the reconstructed dipole contains information from the stationary dipole defined in Sec.1.5.

In Chapter 4 we explored the possibility of using the reconstructed remote dipole and reconstructed remote quadrupole to improve constraints on CMB anomalies models beyond what is achievable with traditional probes of the anomalies, like the primary CMB temperature, E-mode CMB polarization and large-angle distribution of galaxies. For a series of representative anomaly models, we constructed a figure of merit which quantifies improvements on parameter constraints and found that the remote dipole and quadrupole can lead to a $\approx 10\%$ increase in the figure of merit in a scenario consistent with next-generation CMB and galaxy survey experiments. Moreover, since the large-angle information of the remote dipole and quadrupole is inferred from data on small angular scales, these observables are not subject to the same systematics as traditional probes of anomalies and thus play an important complementary role in the analysis of the CMB anomalies. The results from this chapter validate the idea that the remote dipole reconstructed with kSZ velocity reconstruction is sensitive to physics on ultra-large scales and justifies the need for the appropriate modelling of all the contributions to the CMB dipole.

Overall, the research progress on kSZ velocity reconstruction presented in this thesis will play an important role in realizing the full potential of the technique. With regards to the road ahead, there are several avenues for future research and further development of the technique. Non-Gaussian contributions to the reconstruction noise, which have been identified as significant in the Box Picture approach for the high signal to noise limit [21], still need a description in the Lightcone Picture. A deeper understanding of the optical depth bias could be achieved using N-body simulations that include the physics of electrons. With respect to the CMB anomalies, a more exhaustive Bayesian model

comparison approach could be used to determine if the remote dipole and quadrupole fields can favor or disprove the physical reality of the anomalies.

We foresee a bright future for the application of kSZ velocity reconstruction. Regarding the near future, this technique is a clear example of extracting "free" or extra information from datasets that will come from already planned CMB experiments and galaxy surveys. Eventually, the potential of using kSZ velocity reconstruction to test fundamental physics may motivate an even stronger push towards the low-noise, high-resolution frontier of CMB experiments. In the big picture, we expect techniques like kSZ velocity reconstruction, pSZ quadrupole reconstruction and lensing potential reconstruction to become the flagships of a new paradigm in cosmology, one based on harnessing the wealth of information contained in CMB secondaries and use it to test our fundamental models of the Universe in new ways.

References

[1] Juan Cayuso, Richard Bloch, Selim C. Hotinli, Matthew C. Johnson, and Fiona McCarthy. Velocity reconstruction with the cosmic microwave background and galaxy surveys. *arXiv e-prints*, page arXiv:2111.11526, November 2021.

[2] Juan I. Cayuso, Matthew C. Johnson, and James B. Mertens. Simulated reconstruction of the remote dipole field using the kinetic Sunyaev Zel'dovich effect. *Phys. Rev.*, D98(6):063502, 2018.

[3] Juan I. Cayuso and Matthew C. Johnson. Towards testing CMB anomalies using the kinetic and polarized Sunyaev-Zel'dovich effects. *Phys. Rev. D*, 101(12):123508, 2020.

[4] N. Aghanim et al. Planck 2018 results. VI. Cosmological parameters. 2018.

[5] Dominik J. Schwarz, Craig J. Copi, Dragan Huterer, and Glenn D. Starkman. CMB Anomalies after Planck. *Class. Quant. Grav.*, 33(18):184001, 2016.

[6] Adam G. Riess et al. A 2.4% Determination of the Local Value of the Hubble Constant. *Astrophys. J.*, 826(1):56, 2016.

[7] Y. Akrami et al. Planck 2018 results. X. Constraints on inflation. *Astron. Astrophys.*, 641:A10, 2020.

[8] Eleonora Di Valentino, Olga Mena, Supriya Pan, Luca Visinelli, Weiqiang Yang, Alessandro Melchiorri, David F. Mota, Adam G. Riess, and Joseph Silk. In the realm of the Hubble tension—a review of solutions. *Class. Quant. Grav.*, 38(15):153001, 2021.

[9] Michael E. Levi et al. The Dark Energy Spectroscopic Instrument (DESI). 7 2019.

[10] Paul A. Abell et al. LSST Science Book, Version 2.0. 2009.

[11] Euclid definition study report, 2011.

[12] J. Richard Shaw, Kris Sigurdson, Ue-Li Pen, Albert Stebbins, and Michael Sitwell. All-Sky Interferometry with Spherical Harmonic Transit Telescopes. *Astrophys. J.*, 781:57, 2014.

[13] Réza Ansari et al. Inflation and Early Dark Energy with a Stage II Hydrogen Intensity Mapping experiment. 2018.

[14] L. B. Newburgh et al. HIRAX: A Probe of Dark Energy and Radio Transients. *Proc. SPIE Int. Soc. Opt. Eng.*, 9906:99065X, 2016.

[15] Peter Ade et al. The Simons Observatory: Science goals and forecasts. *JCAP*, 02:056, 2019.

[16] Kevork N. Abazajian et al. CMB-S4 Science Book, First Edition. 10 2016.

[17] Uros Seljak. Extracting primordial non-gaussianity without cosmic variance. *Phys. Rev. Lett.*, 102:021302, 2009.

[18] Alexandra Terrana, Mary-Jean Harris, and Matthew C. Johnson. Analyzing the cosmic variance limit of remote dipole measurements of the cosmic microwave background using the large-scale kinetic sunyaev zel'dovich effect. *Journal of Cosmology and Astroparticle Physics*, 2017(02):040–040, Feb 2017.

[19] Anne-Sylvie Deutsch, Emanuela Dimastrogiovanni, Matthew C. Johnson, Moritz Münchmeyer, and Alexandra Terrana. Reconstruction of the remote dipole and quadrupole fields from the kinetic Sunyaev Zel'dovich and polarized Sunyaev Zel'dovich effects. *Phys. Rev. D*, 98(12):123501, 2018.

[20] Kendrick M. Smith, Mathew S. Madhavacheril, Moritz Münchmeyer, Simone Ferraro, Utkarsh Giri, and Matthew C. Johnson. KSZ tomography and the bispectrum. 10 2018.

[21] Utkarsh Giri and Kendrick M. Smith. Exploring KSZ velocity reconstruction with N-body simulations and the halo model. 10 2020.

[22] R. A. Sunyaev and Ya. B. Zeldovich. The Observations of relic radiation as a test of the nature of X-Ray radiation from the clusters of galaxies. *Comments Astrophys. Space Phys.*, 4:173–178, 1972.

[23] R. A. Sunyaev and Ya. B. Zeldovich. The velocity of clusters of galaxies relative to the microwave background - The possibility of its measurement. , 190:413–420, February 1980.

[24] Guilhem Lavaux, Niayesh Afshordi, and Michael J. Hudson. First measurement of the bulk flow of nearby galaxies using the cosmic microwave background. *Mon. Not. Roy. Astron. Soc.*, 430:1617–1635, 2013.

[25] Shirley Ho, Simon Dedeo, and David Spergel. Finding the Missing Baryons Using CMB as a Backlight. *arXiv e-prints*, page arXiv:0903.2845, March 2009.

[26] Jiawei Shao, Pengjie Zhang, Weipeng Lin, Yipeng Jing, and Jun Pan. Kinetic sunyaev-zel'dovich tomography with spectroscopic redshift surveys. *Monthly Notices of the Royal Astronomical Society*, 413(1):628–642, Mar 2011.

[27] D. Munshi, I. T. Iliev, K. L. Dixon, and P. Coles. Extracting the late-time kinetic sunyaev–zel'dovich effect. *Monthly Notices of the Royal Astronomical Society*, 463(3):2425–2442, Aug 2016.

[28] Simone Ferraro, J. Colin Hill, Nick Battaglia, Jia Liu, and David N. Spergel. Kinematic sunyaev-zel'dovich effect with projected fields. ii. prospects, challenges, and comparison with simulations. *Physical Review D*, 94(12), Dec 2016.

[29] Cora Dvorkin, Hiranya V. Peiris, and Wayne Hu. Testable polarization predictions for models of cmb isotropy anomalies. *Physical Review D*, 77(6), Mar 2008.

[30] Craig J. Copi, Dragan Huterer, Dominik J. Schwarz, and Glenn D. Starkman. Large-angle cosmic microwave background suppression and polarization predictions. *Monthly Notices of the Royal Astronomical Society*, 434(4):3590–3596, Aug 2013.

[31] Amanda Yoho, Simone Aiola, Craig J. Copi, Arthur Kosowsky, and Glenn D. Starkman. Microwave background polarization as a probe of large-angle correlations. *Physical Review D*, 91(12), Jun 2015.

[32] Hu Zhan, Lloyd Knox, J. Anthony Tyson, and Vera Margoniner. Exploring large-scale structure with billions of galaxies. *The Astrophysical Journal*, 640(1):8–17, Mar 2006.

[33] Cameron Gibelyou, Dragan Huterer, and Wenjuan Fang. Detectability of large-scale power suppression in the galaxy distribution. *Physical Review D*, 82(12), Dec 2010.

[34] Andrew P Hearin, Cameron Gibelyou, and Andrew R Zentner. Testing the origin of the cmb large-angle correlation deficit with a galaxy imaging survey. *Journal of Cosmology and Astroparticle Physics*, 2011(10):012–012, Oct 2011.

[35] S. Y. Sazonov and R. A. Sunyaev. Microwave polarization in the direction of galaxy clusters induced by the CMB quadrupole anisotropy. *Monthly Notices of the Royal Astronomical Society*, 310(3):765–772, 12 1999.

[36] Anne-Sylvie Deutsch, Matthew C. Johnson, Moritz Münchmeyer, and Alexandra Terrana. Polarized sunyaev zel'dovich tomography. *Journal of Cosmology and Astroparticle Physics*, 2018(04):034–034, Apr 2018.

[37] Marc Kamionkowski and Abraham Loeb. Getting around cosmic variance. *Phys. Rev. D*, 56:4511–4513, Oct 1997.

[38] Stephen Boughn and Robert Crittenden. A correlation between the cosmic microwave background and large-scale structure in the universe. *Nature*, 427(6969):45–47, Jan 2004.

[39] P. A. R. Ade et al. Planck 2015 results. XXI. The integrated Sachs-Wolfe effect. *Astron. Astrophys.*, 594:A21, 2016.

[40] M. J. Rees and D. W. Sciama. Large-scale Density Inhomogeneities in the Universe. , 217(5128):511–516, February 1968.

[41] Yan-Chuan Cai, Shaun Cole, Adrian Jenkins, and Carlos S. Frenk. Full-sky map of the isw and rees-sciama effect from gpc simulations, 2010.

[42] M. Birkinshaw and S. F. Gull. A test for transverse motions of clusters of galaxies. , 302(5906):315–317, March 1983.

[43] N. Aghanim, S. Prunet, O. Forni, and F. R. Bouchet. Moving gravitational lenses: Imprints on the CMB. *Astron. Astrophys.*, 334:409, 1998.

[44] Robin Tuluie and Pablo Laguna. The Imprint of Proper Motion of Nonlinear Structures on the Cosmic Microwave Background. , 445:L73, June 1995.

[45] Robin Tuluie, Pablo Laguna, and Peter Anninos. Cosmic Microwave Background Anisotropies from the Rees-Sciama Effect in $\Omega_0 \leq 1$ Universes. , 463:15, May 1996.

[46] M. J. Fullana, J. V. Arnau, R. J. Thacker, H. M. P. Couchman, and D. Sáez. On the estimation and detection of the Rees-Sciama effect. , 464(4):3784–3795, February 2017.

[47] Takemi Okamoto and Wayne Hu. CMB lensing reconstruction on the full sky. *Phys. Rev. D*, 67:083002, 2003.

[48] William H. Press and Paul Schechter. Formation of Galaxies and Clusters of Galaxies by Self-Similar Gravitational Condensation. , 187:425–438, February 1974.

[49] Jeremy Tinker, Andrey V. Kravtsov, Anatoly Klypin, Kevork Abazajian, Michael Warren, Gustavo Yepes, Stefan Gottlöber, and Daniel E. Holz. Toward a halo mass function for precision cosmology: The limits of universality. *The Astrophysical Journal*, 688(2):709–728, Dec 2008.

[50] Mathew S. Madhavacheril, Nicholas Battaglia, and Hironao Miyatake. Fundamental physics from future weak-lensing calibrated sunyaev-zel'dovich galaxy cluster counts. *Physical Review D*, 96(10), Nov 2017.

[51] G. Mark Voit. Tracing cosmic evolution with clusters of galaxies. *Reviews of Modern Physics*, 77(1):207–258, April 2005.

[52] Steven W. Allen, August E. Evrard, and Adam B. Mantz. Cosmological Parameters from Observations of Galaxy Clusters. , 49(1):409–470, September 2011.

[53] Marcelo A. Alvarez. The kinetic sunyaev–zel'dovich effect from reionization: Simulated full-sky maps at arcminute resolution. *The Astrophysical Journal*, 824(2):118, Jun 2016.

[54] E. M. George, C. L. Reichardt, K. A. Aird, B. A. Benson, L. E. Bleem, J. E. Carlstrom, C. L. Chang, H-M. Cho, T. M. Crawford, A. T. Crites, and et al. A measurement of secondary cosmic microwave background anisotropies from the 2500 square-degree spt-sz survey. *The Astrophysical Journal*, 799(2):177, Jan 2015.

[55] Erminia Calabrese, Renée Hložek, Nick Battaglia, J. Richard Bond, Francesco de Bernardis, Mark J. Devlin, Amir Hajian, Shawn Henderson, J. Colin Hil, Arthur Kosowsky, and et al. Precision epoch of reionization studies with next-generation cmb experiments. *Journal of Cosmology and Astroparticle Physics*, 2014(08):010–010, Aug 2014.

[56] Nick Hand, Graeme E. Addison, Eric Aubourg, Nick Battaglia, Elia S. Battistelli, Dmitry Bizyaev, J. Richard Bond, Howard Brewington, Jon Brinkmann, Benjamin R. Brown, and et al. Evidence of galaxy cluster motions with the kinematic sunyaev-zel'dovich effect. *Physical Review Letters*, 109(4), Jul 2012.

[57] F. De Bernardis, S. Aiola, E.M. Vavagiakis, N. Battaglia, M.D. Niemack, J. Beall, D.T. Becker, J.R. Bond, E. Calabrese, H. Cho, and et al. Detection of the pairwise kinematic sunyaev-zel'dovich effect with boss dr11 and the atacama cosmology telescope. *Journal of Cosmology and Astroparticle Physics*, 2017(03):008–008, Mar 2017.

[58] P. A. R. Ade, N. Aghanim, M. Arnaud, M. Ashdown, E. Aubourg, J. Aumont, C. Baccigalupi, A. J. Banday, R. B. Barreiro, and et al. Planckintermediate results. *Astronomy Astrophysics*, 586:A140, Feb 2016.

[59] B. Soergel, S. Flender, K. T. Story, L. Bleem, T. Giannantonio, G. Efstathiou, E. Rykoff, B. A. Benson, T. Crawford, S. Dodelson, and et al. Detection of the kinematic sunyaev–zel'dovich effect with des year 1 and spt. *Monthly Notices of the Royal Astronomical Society*, 461(3):3172–3193, Jun 2016.

[60] Ming Li, R. E. Angulo, S. D. M. White, and J. Jasche. Matched filter optimization of ksz measurements with a reconstructed cosmological flow field. *Monthly Notices of the Royal Astronomical Society*, 443(3):2311–2326, Jul 2014.

[61] Marc Kamionkowski and Abraham Loeb. Getting around cosmic variance. , 56(8):4511–4513, October 1997.

[62] Siavash Yasini and Elena Pierpaoli. Kinetic sunyaev-zeldovich effect in an anisotropic cmb model: Measuring low multipoles of the cmb at higher redshifts using intensity and polarization spectral distortions. *Physical Review D*, 94(2), Jul 2016.

[63] Thibaut Louis, Emory F. Bunn, Benjamin Wandelt, and Joseph Silk. Measuring polarized emission in clusters in the cmb s4 era. *Physical Review D*, 96(12), Dec 2017.

[64] D. Nelson Limber. The Analysis of Counts of the Extragalactic Nebulae in Terms of a Fluctuating Density Field. , 117:134, January 1953.

[65] D. Nelson Limber. The Analysis of Counts of the Extragalactic Nebulae in Terms of a Fluctuating Density Field. II. , 119:655, May 1954.

[66] Mathew S. Madhavacheril, Nicholas Battaglia, Kendrick M. Smith, and Jonathan L. Sievers. Cosmology with kSZ: breaking the optical depth degeneracy with Fast Radio Bursts. 2019.

[67] Adrienne L. Erickcek, Sean M. Carroll, and Marc Kamionkowski. Superhorizon perturbations and the cosmic microwave background. *Physical Review D*, 78(8), Oct 2008.

[68] Asantha Cooray and Ravi K. Sheth. Halo Models of Large Scale Structure. *Phys. Rept.*, 372:1–129, 2002.

[69] Toshiya Namikawa, Duncan Hanson, and Ryuichi Takahashi. Bias-hardened cmb lensing. *Monthly Notices of the Royal Astronomical Society*, 431(1):609–620, Feb 2013.

[70] Karl Fisher, Ofer Lahav, Yehuda Hoffman, Donald Lynden-Bell, and Saleem Zaroubi. Wiener reconstruction of density, velocity, and potential fields from all-sky galaxy redshift surveys. 6 1994.

[71] F. Lanusse, A. Rassat, and J. L. Starck. 3D galaxy clustering with future wide-field surveys: Advantages of a spherical Fourier-Bessel analysis. *Astron. Astrophys.*, 578:A10, 2015.

[72] Marilena LoVerde and Niayesh Afshordi. Extended Limber Approximation. *Phys. Rev. D*, 78:123506, 2008.

[73] T. D. Kitching and A. F. Heavens. Unequal-Time Correlators for Cosmology. *Phys. Rev. D*, 95(6):063522, 2017.

[74] Xiao Fang, Elisabeth Krause, Tim Eifler, and Niall MacCrann. Beyond Limber: Efficient computation of angular power spectra for galaxy clustering and weak lensing. *JCAP*, 05:010, 2020.

[75] Nicholas Battaglia. The Tau of Galaxy Clusters. *JCAP*, 1608(08):058, 2016.

[76] Dagoberto Contreras, Matthew C. Johnson, and James B. Mertens. Towards detection of relativistic effects in galaxy number counts using kSZ tomography. *Journal of Cosmology and Astroparticle Physics*, 2019(10):024–024, oct 2019.

[77] Zhen Pan and Matthew C. Johnson. Forecasted constraints on modified gravity from Sunyaev-Zel'dovich tomography. *Phys. Rev. D*, 100(8):083522, 2019.

[78] Moritz Münchmeyer, Mathew S. Madhavacheril, Simone Ferraro, Matthew C. Johnson, and Kendrick M. Smith. Constraining local non-Gaussianities with kSZ tomography. 2018.

[79] Fiona McCarthy and Matthew C. Johnson. Remote dipole field reconstruction with dusty galaxies. *Phys. Rev. D*, 102(4):043520, 2020.

[80] Gabriela Sato-Polito, José Luis Bernal, Kimberly K. Boddy, and Marc Kamionkowski. Kinetic Sunyaev-Zel'dovich tomography with line-intensity mapping. 11 2020.

[81] Antony Lewis, Anthony Challinor, and Anthony Lasenby. Efficient computation of CMB anisotropies in closed FRW models. , 538:473–476, 2000.

[82] Selim C. Hotinli and Matthew C. Johnson. Reconstructing large scales at cosmic dawn. 12 2020.

[83] Steve K. Choi et al. The Atacama Cosmology Telescope: a measurement of the Cosmic Microwave Background power spectra at 98 and 150 GHz. *JCAP*, 12:045, 2020.

[84] Chung-Pei Ma and James N. Fry. Nonlinear kinetic Sunyaev-Zeldovich effect. *Phys. Rev. Lett.*, 88:211301, 2002.

[85] J. Colin Hill and Enrico Pajer. Cosmology from the thermal Sunyaev-Zel'dovich power spectrum: Primordial non-Gaussianity and massive neutrinos. *Phys Rev. D*, 88(6):063526, September 2013.

[86] Cien Shang, Zoltán. Haiman, Lloyd Knox, and S. Peng Oh. Improved models for cosmic infrared background anisotropies: new constraints on the infrared galaxy population. , 421(4):2832–2845, April 2012.

[87] Planck Collaboration. Planck 2013 results. XXX. Cosmic infrared background measurements and implications for star formation. , 571:A30, November 2014.

[88] Fiona McCarthy and Mathew S. Madhavacheril. Improving models of the cosmic infrared background using CMB lensing mass maps. *arXiv e-prints*, page arXiv:2010.16405, October 2020.

[89] Max Tegmark, Angelica de Oliveira-Costa, and Andrew Hamilton. A high resolution foreground cleaned CMB map from WMAP. *Phys. Rev. D*, 68:123523, 2003.

[90] Anthony Challinor and Antony Lewis. Linear power spectrum of observed source number counts. , 84(4):043516, August 2011.

[91] Alex Krolewski, Simone Ferraro, Edward F. Schlafly, and Martin White. unWISE tomography of Planck CMB lensing. *JCAP*, 05:047, 2020.

[92] Alex Krolewski, Simone Ferraro, and Martin White. Cosmological constraints from unWISE and Planck CMB lensing tomography. 5 2021.

[93] Aleksandra Kusiak, Boris Bolliet, Simone Ferraro, J. Colin Hill, and Alex Krolewski. Constraining the baryon abundance with the kinematic Sunyaev-Zel'dovich effect: Projected-field detection using Planck, WMAP, and unWISE. *Phys. Rev. D*, 104(4):043518, 2021.

[94] Edward F. Schlafly, Aaron M. Meisner, and Gregory M. Green. The unwise catalog: Two billion infrared sources from five years of wise imaging. *The Astrophysical Journal Supplement Series*, 240(2):30, Feb 2019.

[95] Edward L. Wright, Peter R. M. Eisenhardt, Amy K. Mainzer, Michael E. Ressler, Roc M. Cutri, Thomas Jarrett, J. Davy Kirkpatrick, Deborah Padgett, Robert S. McMillan, Michael Skrutskie, and et al. The wide-field infrared survey explorer (wise): Mission description and initial on-orbit performance. *The Astronomical Journal*, 140(6):1868–1881, Nov 2010.

[96] Boris Leistedt, Hiranya V. Peiris, Daniel J. Mortlock, Aurélien Benoit-Lévy, and Andrew Pontzen. Estimating the large-scale angular power spectrum in the presence of systematics: a case study of Sloan Digital Sky Survey quasars. *Mon. Not. Roy. Astron. Soc.*, 435:1857, 2013.

[97] Franz Elsner, Boris Leistedt, and Hiranya V. Peiris. Unbiased methods for removing systematics from galaxy clustering measurements. *Monthly Notices of the Royal Astronomical Society*, 456(2):2095–2104, Dec 2015.

[98] Dragan Huterer, Carlos E. Cunha, and Wenjuan Fang. Calibration errors unleashed: effects on cosmological parameters and requirements for large-scale structure surveys. *Monthly Notices of the Royal Astronomical Society*, 432(4):2945–2961, May 2013.

[99] Noah Weaverdyck, Jessica Muir, and Dragan Huterer. Integrated sachs-wolfe map reconstruction in the presence of systematic errors. *Phys. Rev. D*, 97:043515, Feb 2018.

[100] Jessica Muir and Dragan Huterer. Reconstructing the integrated sachs-wolfe map with galaxy surveys. *Phys. Rev. D*, 94:043503, Aug 2016.

[101] Anthony Challinor and Floor van Leeuwen. Peculiar velocity effects in high-resolution microwave background experiments. *Phys. Rev. D*, 65:103001, Apr 2002.

[102] Anthony Challinor and Anthony Lasenby. Relativistic corrections to the sunyaev-zeldovich effect. *The Astrophysical Journal*, 499(1):1–6, May 1998.

[103] Anthony Challinor and Anthony Lasenby. Comptonization of an isotropic distribution in moving media: Higher order effects. *The Astrophysical Journal*, 510(2):930–933, Jan 1999.

[104] Naoki Itoh, Yasuharu Kohyama, and Satoshi Nozawa. Relativistic corrections to the sunyaev-zeldovich effect for clusters of galaxies. *The Astrophysical Journal*, 502(1):7–15, Jul 1998.

[105] Satoshi Nozawa, Naoki Itoh, and Yasuharu Kohyama. Relativistic corrections to the sunyaev-zeldovich effect for clusters of galaxies. ii. inclusion of peculiar velocities. *The Astrophysical Journal*, 508(1):17–24, Nov 1998.

[106] William Coulton, Atsuhisa Ota, and Alexander van Engelen. Cosmology with the thermal-kinetic sunyaev-zel'dovich effect. *Physical Review Letters*, 125(11), Sep 2020.

[107] N. J. Miller, M. Shimon, and B. G. Keating. Cmb beam systematics: Impact on lensing parameter estimation. *Phys. Rev. D*, 79:063008, Mar 2009.

[108] P. Daniel Meerburg, Joel Meyers, and Alexander van Engelen. Reconstructing the Primary CMB Dipole. *Phys. Rev.*, D96(8):083519, 2017.

[109] Naonori S. Sugiyama, Teppei Okumura, and David N. Spergel. Will kinematic sunyaev-zel'dovich measurements enhance the science return from galaxy redshift surveys?, 2017.

[110] Selim C. Hotinli, Kendrick M. Smith, Mathew S. Madhavacheril, and Marc Kamionkowski. Cosmology with the moving lens effect. 8 2021.

[111] Selim C. Hotinli, James B. Mertens, Matthew C. Johnson, and Marc Kamionkowski. Probing correlated compensated isocurvature perturbations using scale-dependent galaxy bias. *Phys. Rev. D*, 100(10):103528, 2019.

[112] Vanessa Böhm, Marcel Schmittfull, and Blake D. Sherwin. Bias to CMB lensing measurements from the bispectrum of large-scale structure. *Phys. Rev. D*, 94(4):043519, 2016.

[113] Cullan Howlett, Marc Manera, and Will J. Percival. L-PICOLA: A parallel code for fast dark matter simulation. *Astron. Comput.*, 12:109–126, 2015.

[114] Roman Scoccimarro, Lam Hui, Marc Manera, and Kwan Chuen Chan. Large-scale Bias and Efficient Generation of Initial Conditions for Non-Local Primordial Non-Gaussianity. *Phys. Rev.*, D85:083002, 2012.

[115] Svetlin Tassev, Matias Zaldarriaga, and Daniel Eisenstein. Solving Large Scale Structure in Ten Easy Steps with COLA. *JCAP*, 1306:036, 2013.

[116] Giuseppe Tormen and Edmund Bertschinger. Adding long wavelength modes to an n-body simulation. *Astrophys. J.*, 472:14, 1996.

[117] Shaun Cole. Adding long-wavelength power to n-body simulations. *Mon. Not. Roy. Astron. Soc.*, 286:38, 1997.

[118] Edmund Bertschinger. Multiscale Gaussian random fields for cosmological simulations. *Astrophys. J. Suppl.*, 137:1, 2001.

[119] Raphael Sgier, Alexandre Réfrégier, Adam Amara, and Andrina Nicola. Fast Generation of Covariance Matrices for Weak Lensing. 2018.

[120] Romain Teyssier, Sandrine Pires, Simon Prunet, Dominique Aubert, Christophe Pichon, Adam Amara, Karim Benabed, Stephane Colombi, Alexandre Refregier, and Jean-Luc Starck. Full-Sky Weak Lensing Simulation with 70 Billion Particles. *Astron. Astrophys.*, 497:335, 2009.

[121] A. Kiessling, A. F. Heavens, and A. N. Taylor. SUNGLASS: A new weak lensing simulation pipeline. *Mon. Not. Roy. Astron. Soc.*, 414:2235, 2011.

[122] Samuel Flender, Lindsey Bleem, Hal Finkel, Salman Habib, Katrin Heitmann, and Gilbert Holder. Simulations of the Pairwise Kinematic Sunyaev–Zel'dovich Signal. *Astrophys. J.*, 823(2):98, 2016.

[123] K. M. Gorski, Eric Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelman. HEALPix - A Framework for high resolution discretization, and fast analysis of data distributed on the sphere. *Astrophys. J.*, 622:759–771, 2005.

[124] Jun Koda, Chris Blake, Tamara Davis, Christina Magoulas, Christopher M. Springob, Morag Scrimgeour, Andrew Johnson, Gregory B. Poole, and Lister Staveley-Smith. Are peculiar velocity surveys competitive as a cosmological probe? *Mon. Not. Roy. Astron. Soc.*, 445(4):4267–4286, 2014.

[125] Nora Elisa Chisari and Matias Zaldarriaga. Connection between Newtonian simulations and general relativity. *Phys. Rev.*, D83:123505, 2011. [Erratum: Phys. Rev.D84,089901(2011)].

[126] Christian Fidler, Thomas Tram, Cornelius Rampf, Robert Crittenden, Kazuya Koyama, and David Wands. Relativistic Interpretation of Newtonian Simulations for Cosmic Structure Formation. *JCAP*, 1609(09):031, 2016.

[127] William E. East, Radosław Wojtak, and Tom Abel. Comparing Fully General Relativistic and Newtonian Calculations of Structure Formation. *Phys. Rev.*, D97(4):043509, 2018.

[128] Diego Blas, Julien Lesgourgues, and Thomas Tram. The Cosmic Linear Anisotropy Solving System (CLASS) II: Approximation schemes. *JCAP*, 1107:034, 2011.

[129] Antony Lewis and Anthony Challinor. Weak gravitational lensing of the cmb. *Phys. Rept.*, 429:1–65, 2006.

[130] C. L. Bennett et al. The Microwave Anisotropy Probe (MAP) Mission. *Astrophys. J.*, 583:1–23, 2003.

[131] J. A. Tauber, N. Mandolesi, J.-L. Puget, et al. Planck pre-launch status: The Planck mission. *A&A*, 520, September 2010.

[132] Cora Dvorkin, Hiranya V. Peiris, and Wayne Hu. Testable polarization predictions for models of cmb isotropy anomalies. *Phys. Rev. D*, 77:063008, Mar 2008.

[133] C. J. Copi, D. Huterer, D. J. Schwarz, and G. D. Starkman. Large-Angle CMB Suppression and Polarization Predictions. *Mon. Not. Roy. Astron. Soc.*, 434:3590–3596, 2013.

[134] Amanda Yoho, Simone Aiola, Craig J. Copi, Arthur Kosowsky, and Glenn D. Starkman. Microwave Background Polarization as a Probe of Large-Angle Correlations. *Phys. Rev.*, D91(12):123504, 2015.

[135] Marcio O'Dwyer, Craig J. Copi, Lloyd Knox, and Glenn D. Starkman. CMB-S4 and the Hemispherical Variance Anomaly. *Mon. Not. Roy. Astron. Soc.*, 470(1):372–378, 2017.

[136] Emory F. Bunn, Qingyang Xue, and Haoxuan Zheng. Testing cosmological models with large-scale power modulation using microwave background polarization observations. *Phys. Rev.*, D94(10):103512, 2016.

[137] D. Contreras, J. P. Zibin, D. Scott, A. J. Banday, and K. M. Górski. Testing physical models for dipolar asymmetry with CMB polarization. *Phys. Rev.*, D96(12):123522, 2017.

[138] Georges Obied, Cora Dvorkin, Chen Heinrich, Wayne Hu, and V. Miranda. Inflationary versus reionization features from *Planck* 2015 data. *Phys. Rev.*, D98(4):043518, 2018.

[139] M. Billi, A. Gruppuso, N. Mandolesi, L. Moscardini, and P. Natoli. Polarisation as a tracer of CMB anomalies: Planck results and future forecasts. 2019.

[140] A. Yoho, C. J. Copi, G. D. Starkman, and A. Kosowsky. Probing Large-Angle Correlations with the Microwave Background Temperature and Lensing Cross Correlation. *Mon. Not. Roy. Astron. Soc.*, 442(3):2392–2397, 2014.

[141] J. P. Zibin and D. Contreras. Testing physical models for dipolar asymmetry: from temperature to k space to lensing. *Phys. Rev.*, D95(6):063011, 2017.

[142] Craig J. Copi, Márcio O'Dwyer, and Glenn D. Starkman. The ISW effect and the lack of large-angle CMB temperature correlations. *Mon. Not. Roy. Astron. Soc.*, 463(3):3305–3310, 2016.

[143] Simon Foreman, P. Daniel Meerburg, Joel Meyers, and Alexander van Engelen. Cosmic variance mitigation in measurements of the integrated Sachs-Wolfe effect. 2018.

[144] Xingang Chen, Cora Dvorkin, Zhiqi Huang, Mohammad Hossein Namjoo, and Licia Verde. The Future of Primordial Features with Large-Scale Structure Surveys. *JCAP*, 1611(11):014, 2016.

[145] Xingang Chen, P. Daniel Meerburg, and Moritz Münchmeyer. The Future of Primordial Features with 21 cm Tomography. *JCAP*, 1609(09):023, 2016.

[146] Mario Ballardini, Fabio Finelli, Roy Maartens, and Lauro Moscardini. Probing primordial features with next-generation photometric and radio surveys. *JCAP*, 1804(04):044, 2018.

[147] Gonzalo A. Palma, Domenico Sapone, and Spyros Sypsas. Constraints on inflation with LSS surveys: features in the primordial power spectrum. *JCAP*, 1806(06):004, 2018.

[148] Olivier Doré et al. Cosmology with the SPHEREX All-Sky Spectral Survey. 2014.

[149] Franz Elsner, Boris Leistedt, and Hiranya V. Peiris. Unbiased methods for removing systematics from galaxy clustering measurements. *Mon. Not. Roy. Astron. Soc.*, 456(2):2095–2104, 2016.

[150] Emory F. Bunn. Probing the Universe on gigaparsec scales with remote cosmic microwave background quadrupole measurements. *Physical Review D*, 73(12):123517, jun 2006.

[151] Anne-Sylvie Deutsch, Emanuela Dimastrogiovanni, Matteo Fasiello, Matthew C. Johnson, and Moritz Münchmeyer. Primordial gravitational wave phenomenology with polarized Sunyaev Zel'dovich tomography. 2018.

[152] E. Alizadeh and C. M. Hirata. How to detect gravitational waves through the cross correlation of the galaxy distribution with the CMB polarization. *Physical Review D*, 85(12):123540, 2012.

[153] Pengjie Zhang and Matthew C. Johnson. Testing eternal inflation with the kinetic Sunyaev Zel'dovich effect. *JCAP*, 1506(06):046, 2015.

[154] Alex Hall and Anthony Challinor. Detecting the polarization induced by scattering of the microwave background quadrupole in galaxy clusters. *Physical Review D*, 90(6):063518, sep 2014.

[155] J. P. Zibin and Douglas Scott. Gauging the cosmic microwave background. *Phys. Rev. D*, 78:123529, Dec 2008.

[156] Siavash Yasini and Elena Pierpaoli. Generalized Doppler and aberration kernel for frequency-dependent cosmological observables. *Phys. Rev.*, D96(10):103502, 2017.

[157] Roberto Trotta. Forecasting the Bayes factor of a future observation. *Mon. Not. Roy. Astron. Soc.*, 378:819–824, 2007.

[158] Andreas Albrecht, Luca Amendola, Gary Bernstein, Douglas Clowe, Daniel Eisenstein, Luigi Guzzo, Christopher Hirata, Dragan Huterer, Robert Kirshner, Edward Kolb, and Robert Nichol. Findings of the Joint Dark Energy Mission Figure of Merit Science Working Group. *arXiv e-prints*, page arXiv:0901.0721, Jan 2009.

[159] Planck Collaboration. Planck 2015 results - xvi. isotropy and statistics of the cmb. *A&A*, 594:A16, 2016.

[160] Christopher Gordon, Wayne Hu, Dragan Huterer, and Tom Crawford. Spontaneous isotropy breaking: A mechanism for cmb multipole alignments. *Phys. Rev. D*, 72:103002, Nov 2005.

[161] Adrienne L. Erickcek, Marc Kamionkowski, and Sean M. Carroll. A Hemispherical Power Asymmetry from Inflation. *Phys. Rev.*, D78:123520, 2008.

[162] Andrew R. Liddle and Marina Cortês. Cosmic microwave background anomalies in an open universe. *Phys. Rev. Lett.*, 111(11):111302, 2013.

[163] Takeshi Kobayashi, Marina Cortês, and Andrew R. Liddle. A separate universe view of the asymmetric sky. *JCAP*, 1505(05):029, 2015.

[164] Saroj Adhikari, Sarah Shandera, and Adrienne L. Erickcek. Large-scale anomalies in the cosmic microwave background as signatures of non-Gaussianity. *Phys. Rev.*, D93(2):023524, 2016.

[165] D. Contreras, J. Hutchinson, A. Moss, D. Scott, and J. P. Zibin. Closing in on the large-scale CMB power asymmetry. *Phys. Rev.*, D97(6):063504, 2018.

[166] Christopher M Hirata. Constraints on cosmic hemispherical power anomalies from quasars. *Journal of Cosmology and Astroparticle Physics*, 2009(09):011–011, sep 2009.

[167] Andrei D. Linde. A Toy model for open inflation. *Phys. Rev.*, D59:023503, 1999.

[168] Daisuke Yamauchi, Andrei Linde, Atsushi Naruko, Misao Sasaki, and Takahiro Tanaka. Open inflation in the landscape. *Phys. Rev.*, D84:043513, 2011.

[169] Carlo R Contaldi, Marco Peloso, Lev Kofman, and Andrei Linde. Suppressing the lower multipoles in the CMB anisotropies. *Journal of Cosmology and Astroparticle Physics*, 2003(07):002–002, jul 2003.

[170] Xingang Chen, Richard Easther, and Eugene A. Lim. Large Non-Gaussianities in Single Field Inflation. *JCAP*, 0706:023, 2007.

[171] Eva Silverstein and David Tong. Scalar speed limits and cosmology: Acceleration from D-cceleration. *Phys. Rev.*, D70:103505, 2004.

[172] Jennifer A. Adams, Bevan Cresswell, and Richard Easther. Inflationary perturbations from a potential with a step. *Phys. Rev.*, D64:123514, 2001.

[173] Ana Achucarro, Jinn-Ouk Gong, Sjoerd Hardeman, Gonzalo A. Palma, and Subodh P. Patil. Features of heavy physics in the CMB power spectrum. *JCAP*, 1101:030, 2011.

[174] Rachel Bean, Xingang Chen, Girma Hailu, S. H. Henry Tye, and Jiajun Xu. Duality Cascade in Brane Inflation. *JCAP*, 0803:026, 2008.

[175] Planck Collaboration. Planck 2015 results - xx. constraints on inflation. *A&A*, 594:A20, 2016.

[176] Vinícius Miranda and Wayne Hu. Inflationary steps in the planck data. *Phys. Rev. D*, 89:083529, Apr 2014.

[177] LSST Science Collaboration. Lsst science book, version 2.0, 2009.

[178] Luca Amendola et al. Cosmology and fundamental physics with the Euclid satellite. *Living Rev. Rel.*, 21(1):2, 2018.

[179] Kevin Bandura, Graeme E. Addison, Mandana Amiri, J. Richard Bond, Duncan Campbell-Wilson, Liam Connor, Jean-François Cliche, Greg Davis, Meiling Deng, Nolan Denman, and et al. Canadian hydrogen intensity mapping experiment (chime) pathfinder. *Ground-based and Airborne Telescopes V*, Jul 2014.

[180] R. A. Sunyaev and Ya. B. Zeldovich. Distortions of the Background Radiation Spectrum. , 223(5207):721–722, August 1969.

[181] C. Hernández-Monteagudo and R. A. Sunyaev. Galaxy clusters as mirrors of the distant universe. *Astronomy and Astrophysics*, 509:A82, Jan 2010.

[182] Arthur Kosowsky. The Atacama Cosmology Telescope. *New Astron. Rev.*, 47:939–943, 2003.

[183] Fabian Schmidt, Enrico Pajer, and Matias Zaldarriaga. Large-scale structure and gravitational waves. iii. tidal effects. *Physical Review D*, 89(8), Apr 2014.

[184] Hong-Ming Zhu, Ue-Li Pen, Yu Yu, Xinzhong Er, and Xuelei Chen. Cosmic tidal reconstruction. *Physical Review D*, 93(10), May 2016.

[185] Christopher M. Hirata. Tidal alignments as a contaminant of redshift space distortions. *Mon. Not. Roy. Astron. Soc.*, 399:1074, 2009.

[186] Roman Scoccimarro, Matias Zaldarriaga, and Lam Hui. Power spectrum correlations induced by nonlinear clustering. *Astrophys. J.*, 527:1, 1999.

[187] Emmanuel Schaan, Simone Ferraro, and David N. Spergel. Weak Lensing of Intensity Mapping: the Cosmic Infrared Background. *Phys. Rev. D*, 97(12):123539, 2018.

[188] Hyunbae Park, Marcelo A. Alvarez, and J. Richard Bond. The Impact of Baryonic Physics on the Kinetic Sunyaev–Zel'dovich Effect. *Astrophys. J.*, 853(2):121, 2018.

[189] Jeremy L. Tinker, Brant E. Robertson, Andrey V. Kravtsov, Anatoly Klypin, Michael S. Warren, Gustavo Yepes, and Stefan Gottlöber. The Large-scale Bias of Dark Matter Halos: Numerical Calibration and Model Tests. *apj*, 724(2):878–886, Dec 2010.

[190] Julio F. Navarro, Carlos S. Frenk, and Simon D. M. White. The Structure of Cold Dark Matter Halos. , 462:563, May 1996.

[191] Alan R. Duffy, Joop Schaye, Scott T. Kay, and Claudio Dalla Vecchia. Dark matter halo concentrations in the Wilkinson Microwave Anisotropy Probe year 5 cosmology. , 390(1):L64–L68, October 2008.

[192] Jeremy Tinker, Andrey V. Kravtsov, Anatoly Klypin, Kevork Abazajian, Michael Warren, Gustavo Yepes, Stefan Gottlöber, and Daniel E. Holz. Toward a Halo Mass Function for Precision Cosmology: The Limits of Universality. *apj*, 688(2):709–728, December 2008.

[193] Rahul Kothari, Shamik Ghosh, Pranati K. Rath, Gopal Kashyap, and Pankaj Jain. Imprint of Inhomogeneous and Anisotropic Primordial Power Spectrum on CMB Polarization. *Mon. Not. Roy. Astron. Soc.*, 460(2):1577–1587, 2016.

APPENDICES

Appendix A

A.1 Quadratic estimator for the remote quadrupole

In this appendix, we review the main details of the remote quadrupole field reconstruction described in [36, 19]. The polarized Sunyaev-Zel'dovich signal (discussed in Sec.1.3.5), expressed in terms of the contributions to the Stokes parameters $Q \pm iU$, is related to the remote CMB quadrupole by:

$$(Q \pm iU)^{\rm pSZ}(\hat{\mathbf{n}}) = -\frac{\sqrt{6}}{10}\int d\chi\; q^{\pm}_{\rm eff}\left(\hat{\mathbf{n}},\chi\right)\dot{\tau}(\hat{\mathbf{n}},\chi), \tag{A.1}$$

$$q^{\pm}_{\rm eff}\left(\hat{\mathbf{n}},\chi\right) \equiv \sum_{m=-2}^{2}\Theta^{m}_{2}\left(\hat{\mathbf{n}},\chi\right){}_{\mp 2}Y_{2m}\left(\hat{\mathbf{n}}\right) \tag{A.2}$$

where $\Theta^{m}_{2}(\hat{\mathbf{n}},\chi)$ are the remote $\ell = 2$ CMB multipoles observed by electrons at spacetime position $(\eta(\chi), \mathbf{x} = \chi\hat{\mathbf{n}})$ on our past lightcone:

$$\Theta^{m}_{2}(\hat{\mathbf{n}},\chi) = \int d^{2}\hat{n}\;\Theta(\mathbf{x},\hat{\mathbf{n}})Y^{*}_{2m}\left(\hat{\mathbf{n}}\right) \tag{A.3}$$

and ${}_{\pm 2}Y_{2m}\left(\hat{\mathbf{n}}\right)$ are spin-2 weighted spherical harmonics. The remote quadrupole receives contributions from both scalar and tensor fluctuations. Restricting only to scalar modes, $q^{+}_{\rm eff} = q^{-}_{\rm eff}$, and the remote quadrupole is curl-free. In this thesis, we only will consider the scalar sourced remote quadrupole, which we denote by q_E. However, one should notice that a pure scalar remote quadrupole still leads to a B-mode for the pSZ signal due to the spatial variation of the optical depth [152]. The E-mode and B-polarization arising from a

scalar sourced remote quadrupole are given in harmonic space by:

$$a_{\ell m}^{E,pSZ} = -\frac{\sqrt{6}}{10}\int d\chi \sum_{L_1,M_1,L_2,M_2}(-1)^m \begin{pmatrix} \ell & L_1 & L_2 \\ -m & M_1 & M_2 \end{pmatrix} F_{\ell,L_1,L_2} \times \alpha_{\ell,L_1,L_2} q_{E;L_1M_1}(\chi)\dot{\tau}_{L_2M_2}(\chi) \tag{A.4}$$

$$a_{\ell m}^{B,pSZ} = -\frac{\sqrt{6}}{10}\int d\chi \sum_{L_1,M_1,L_2,M_2}(-1)^m \begin{pmatrix} \ell & L_1 & L_2 \\ -m & M_1 & M_2 \end{pmatrix} F_{\ell,L_1,L_2} \times \gamma_{\ell,L_1,L_2} q_{E;L_1M_1}(\chi)\dot{\tau}_{L_2M_2}(\chi) \tag{A.5}$$

where

$$F_{\ell,L_1,L_2} = \sqrt{\frac{(2\ell+1)\,(2L_1+1)\,(2L_2+1)}{4\pi}} \begin{pmatrix} \ell & L_1 & L_2 \\ 2 & -2 & 0 \end{pmatrix} \tag{A.6}$$

$$\alpha_{\ell,L_1,L_2} = \frac{1}{2}\left(1+(-1)^{\ell+L_1+L_2}\right) \tag{A.7}$$

$$\gamma_{\ell,L_1,L_2} = \frac{1}{2i}\left(1-(-1)^{\ell+L_1+L_2}\right) \tag{A.8}$$

and $q_{E;\ell m}(\chi)$ and $\dot{\tau}_{\ell m}(\chi)$ are the spherical transforms of the fields $q_E(\hat{\mathbf{n}},\chi)$ and $\dot{\tau}(\hat{\mathbf{n}},\chi)$. These continuous fields can be approximated by their bin averages:

$$q_{E;\ell m}(\chi) \approx \Delta\chi \sum_{\alpha'=0}^{N-1} q_{E;\ell m}^{\alpha'}\Pi^{\alpha'}(\chi) \quad \text{with} \quad q_{E;\ell m}^{\alpha'} = \frac{1}{\Delta\chi}\int_{\chi_{min}^{\alpha'}}^{\chi_{max}^{\alpha'}} d\chi\; q_{E;\ell m}(\chi), \tag{A.9}$$

$$\dot{\tau}_{\ell m}(\chi) \approx \Delta\chi \sum_{\alpha'=0}^{N-1} \dot{\tau}_{\ell m}^{\alpha'}\Pi^{\alpha'}(\chi) \quad \text{with} \quad \dot{\tau}_{\ell m}^{\alpha'} = \frac{1}{\Delta\chi}\int_{\chi_{min}^{\alpha'}}^{\chi_{max}^{\alpha'}} d\chi\; \dot{\tau}_{\ell m}(\chi). \tag{A.10}$$

where Π is the top-hat window defined in Eq.(2.14). The pSZ signal can be expressed as a sum over bin averages:

$$a_{\ell m}^{E,pSZ} = -\frac{\sqrt{6}}{10} \sum_{\alpha,L_1,M_1,L_2,M_2}(-1)^m \begin{pmatrix} \ell & L_1 & L_2 \\ -m & M_1 & M_2 \end{pmatrix} F_{\ell,L_1,L_2} \times \alpha_{\ell,L_1,L_2}\; q_{E;L_1M_1}^{\alpha}\dot{\tau}_{L_2M_2}^{\alpha}\Delta\chi \tag{A.11}$$

$$a_{\ell m}^{B,pSZ} = -\frac{\sqrt{6}}{10} \sum_{\alpha,L_1,M_1,L_2,M_2}(-1)^m \begin{pmatrix} \ell & L_1 & L_2 \\ -m & M_1 & M_2 \end{pmatrix} F_{\ell,L_1,L_2} \times \gamma_{\ell,L_1,L_2}\; q_{E;L_1M_1}^{\alpha}\dot{\tau}_{L_2M_2}^{\alpha}\Delta\chi \tag{A.12}$$

It was shown in [36, 19] that, given a tracer $\delta^{\beta}_{\ell_2 m_2}$ of the electron distribution in bin β (a galaxy redshift survey for example), the remote quadrupole leads to a statistical anisotropy in the pSZ-density correlation:

$$\begin{aligned}
\left\langle a^{E,pSZ}_{\ell_1 m_1} \delta^{\beta}_{\ell_2 m_2} \right\rangle &= -\frac{\sqrt{6}}{10} \sum_{\alpha, L_1, M_1, L_2, M_2} (-1)^{m_1} \begin{pmatrix} \ell_1 & L_1 & L_2 \\ -m_1 & M_1 & M_2 \end{pmatrix} F_{\ell_1, L_1, L_2} \\
&\times \alpha_{\ell_1, L_1, L_2}\, q^{\alpha}_{E; L_1 M_1} \left\langle \dot{\tau}^{\alpha}_{L_2 M_2} \delta^{\beta}_{\ell_2 m_2} \right\rangle \qquad \text{(A.13)} \\
\left\langle a^{B,pSZ}_{\ell_1 m_1} \delta^{\beta}_{\ell_2 m_2} \right\rangle &= -\frac{\sqrt{6}}{10} \sum_{\alpha, L_1, M_1, L_2, M_2} (-1)^{m_1} \begin{pmatrix} \ell_1 & L_1 & L_2 \\ -m_1 & M_1 & M_2 \end{pmatrix} F_{\ell_1, L_1, L_2} \\
&\times \gamma_{\ell_1, L_1, L_2}\, q^{\alpha}_{E; L_1 M_1} \left\langle \dot{\tau}^{\alpha}_{L_2 M_2} \delta^{\beta}_{\ell_2 m_2} \right\rangle \qquad \text{(A.14)}
\end{aligned}$$

Assuming $\left\langle \dot{\tau}^{\alpha}_{L_2 M_2} \delta^{\beta}_{\ell_2 m_2} \right\rangle = C^{\dot{\tau}^{\alpha}\delta^{\alpha}}_{\ell_2} \delta_{\alpha\beta} \delta_{L_2 \ell_2} \delta_{M_2 m_2}$, one can use this statistical anisotropy to construct quadratic estimators for the bin averaged multipoles: $q^{\alpha}_{E;LM}$:

$$\begin{aligned}
\widehat{q}^{\alpha}_{\ell m} &= N^{E\alpha}_{\ell} \sum_{\ell_1 m_1 \ell_2 m_2} (-1)^{m} f^{q^{\alpha}_E}_{\ell_1 \ell \ell_2} \begin{pmatrix} \ell_1 & \ell_2 & \ell \\ m_1 & m_2 & -m \end{pmatrix} \\
&\times \frac{\left(\alpha_{\ell,\ell_1,\ell_2} a^{E}_{\ell_1 m_1} - \gamma_{\ell,\ell_1,\ell_2} a^{B}_{\ell_1 m_1}\right) \delta^{\alpha}_{\ell_2 m_2}}{\left(\left|\alpha_{\ell,\ell_1,\ell_2}\right|^2 \tilde{C}^{EE}_{\ell_1} + \left|\gamma_{\ell,\ell_1,\ell_2}\right|^2 \tilde{C}^{BB}_{\ell_1}\right) \tilde{C}^{\delta^{\alpha}\delta^{\alpha}}_{\ell_2}} \qquad \text{(A.15)}
\end{aligned}$$

were $a^{E}_{\ell_1 m_1}$ and $a^{B}_{\ell_1 m_1}$ are the measured CMB polarization E and B modes with power spectra $\tilde{C}^{EE}_{\ell_1}$ and $\tilde{C}^{BB}_{\ell_1}$, including noise and any other source of CMB polarization. The f-coupling is given by:

$$f^{q^{\alpha}_E}_{\ell_1 \ell \ell_2} = -\frac{\sqrt{6}}{10} \sqrt{\frac{(2\ell+1)(2\ell_1+1)(2\ell_2+1)}{4\pi}} \begin{pmatrix} \ell_1 & \ell & \ell_2 \\ 2 & -2 & 0 \end{pmatrix} C^{\dot{\tau}^{\alpha}\delta^{\alpha}}_{\ell_2} \qquad \text{(A.16)}$$

and the reconstruction noise $N^{E\alpha}_{\ell}$ is given by:

$$N^{E\alpha}_{\ell} = (2\ell+1) \left(\sum_{\ell_1 \ell_2} \frac{f^{q^{\alpha}_E}_{\ell_1 \ell \ell_2} f^{q^{\alpha}_E}_{\ell_1 \ell \ell_2}}{\left(\left|\alpha_{\ell,\ell_1,\ell_2}\right|^2 \tilde{C}^{EE}_{\ell_1} + \left|\gamma_{\ell,\ell_1,\ell_2}\right|^2 \tilde{C}^{BB}_{\ell_1}\right) \tilde{C}^{\delta^{\alpha}\delta^{\alpha}}_{\alpha \ell_2}} \right)^{-1} \qquad \text{(A.17)}$$

A.2 Transfer function for the remote quadrupole

Similarly to the the remote dipole field, the remote quadrupole field can be expressed as a Fourier space integral:

$$q_{E;LM}(\chi) = \int \frac{d^3k}{(2\pi)^3} \Delta_L^{q_E}(k,\chi) \Psi_p(\mathbf{k}) Y^*_{LM}(\widehat{\mathbf{k}}) \tag{A.18}$$

where the remote quadrupole transfer function $\Delta_L^{q_E}(k,\chi)$ is given by :

$$\Delta_L^{q_E}(k,\chi) = -5i^L \sqrt{\frac{3}{8}} \sqrt{\frac{(L+2)!}{(L-2)!}} \frac{j_L(k\chi)}{(k\chi)^2} T(k) \left[\mathcal{G}^{SW}(k,\chi) + \mathcal{G}^{\rm ISW}(k,\chi) + \mathcal{G}^{\mathcal{D}}(k,\chi)\right] \tag{A.19}$$

where $T(k)$ is defined by Eq.(1.42) in Sec.1.5 and

$$\begin{aligned}
\mathcal{G}^{SW}(k,\chi) &\equiv -4\pi \left(2D_\Psi(a_{\rm ls}) - \frac{3}{2}\right) j_2\left(k(\chi_{\rm ls}-\chi)\right) \\
\mathcal{G}^{\rm ISW}(k,\chi) &\equiv -8\pi \int_{a_{\rm ls}}^{a(\chi)} da' \frac{dD_\Psi}{da'} j_2(k(\chi(a')-\chi)) \\
\mathcal{G}^{\mathcal{D}}(k,\chi) &\equiv \frac{4\pi}{5} k D_v(a_{\rm ls}) \left[3j_3\left(k(\chi_{\rm ls}-\chi)\right) - 2j_1\left(k(\chi_{\rm ls}-\chi)\right)\right]
\end{aligned} \tag{A.20}$$

are the remote quadrupole integral kernels (equivalent to those for the remote dipole, defined in Eqs.(1.45)) which capture the contributions from the Sachs-Wolfe effect (SW), integrated Sachs-Wolfe effect (ISW) and Doppler effect ($\mathcal{D}$) due to the peculiar velocity of the last scattering surface.

Appendix B

B.1 Beyond Limber approximation

B.1.1 The general picture

We review the "Beyond Limber approximation" method from [74], which we use to evaluate angular power spectra that take the form Eq.(2.16):

$$C_\ell^{F^W G^{W'}} = \int d\chi_1 d\chi_2 \; W\left(\chi_1\right) W'\left(\chi_2\right) \int \frac{k^2 dk}{(2\pi)^3} \; \mathcal{K}_\ell^F(\chi_1, k)\, \mathcal{K}_\ell^G(\chi_2, k) \; P_{FG}(\chi_1, \chi_2, k),$$

The method aims to separate the integral above into a piece suitable for the Limber approximation, and a piece that can be expressed as a simple Hankel transform. The separation occurs at the level of the power spectrum $P_{FG}(\chi_1, \chi_2, k)$ by defining a "non-linear" power spectrum

$$P_{FG}^{(nlin)}(\chi_1, \chi_2, k) = P_{FG}(\chi_1, \chi_2, k) - P_{FG}^{(lin)}(\chi_1, \chi_2, k). \tag{B.1}$$

where $P_{FG}(\chi_1, \chi_2, k)$ is the full power spectrum we calculate using the halo model described below and $P_{FG}^{(lin)}(\chi_1, \chi_2, k)$ is the linear theory power spectrum. The non-linear power spectrum defined this way is negligible on large scales and starts becoming important for scales and redshifts at which non-linearity kicks in. It is argued in [74] that the Limber approximation of the nonlinear correction term $(P_{FG} - P_{FG}^{(lin)})$ is sufficiently accurate in

realistic cases and therefore the angular power spectrum integral can be rearranged as :

$$C_\ell^{F^W G^{W'}} = \text{Limber}\Bigg[\int d\chi_1 d\chi_2 W(\chi_1) W'(\chi_2) \times \int \frac{k^2 dk}{(2\pi)^3} \mathcal{K}_\ell^F(\chi_1, k) \mathcal{K}_\ell^G(\chi_2, k) P_{FG}^{nlin}(\chi_1, \chi_2, k)\Bigg] \quad \text{(B.2)}$$

$$+ \int d\chi_1 d\chi_2 W(\chi_1) W'(\chi_2) \times \int \frac{k^2 dk}{(2\pi)^3} \mathcal{K}_\ell^F(\chi_1, k) \mathcal{K}_\ell^G(\chi_2, k) P_{FG}^{lin}(\chi_1, \chi_2, k) \quad \text{(B.3)}$$

The linear power spectrum can be related to its value at redshift zero using a growth factor. Ignoring any scale dependent growth for the moment, the linear power spectrum can be expressed as:

$$P_{FG}^{lin}(\chi_1, \chi_2, k) = P_{FG}^{lin}(0, 0, k) g_F(\chi_1) g_G(\chi_2), \quad \text{(B.4)}$$

which allows us to separate the χ_1 and χ_2 dependence in the second term of Eq.(B.2):

$$C_\ell^{F^W G^{W'}} = \text{Limber}\Bigg[\int d\chi_1 d\chi_2 W(\chi_1) W'(\chi_2) \times \int \frac{k^2 dk}{(2\pi)^3} \mathcal{K}_\ell^F(\chi_1, k) \mathcal{K}_\ell^G(\chi_2, k) P_{FG}^{nlin}(\chi_1, \chi_2, k)\Bigg]$$

$$+ \int \frac{k^2 dk}{(2\pi)^3} P_{FG}^{lin}(0, 0, k) \times \left[\int d\chi_1 W(\chi_1) g(\chi_1) \mathcal{K}_\ell^F(\chi_1, k)\right] \left[\int d\chi_2 W'(\chi_2) g(\chi_2) \mathcal{K}_\ell^G(\chi_2, k)\right] \quad \text{(B.5)}$$

For kernels $\mathcal{K}_\ell(\chi, k)$ of the form $f_1(\chi) f_2(k) f_3(\ell) j_{\ell'}(k\chi)$, where $f_1, f_2, f3$ are arbitrary functions and $j_{\ell'}(k\chi)$ is a spherical Bessel function[1], the χ space integrals between brackets can be expressed in terms of Hankel transforms, which can be calculated much faster and with more accuracy than brute force integrations of spherical Bessel functions. If the growth factors g_F, g_G are scale dependent,

$$P_{FG}^{lin}(\chi_1, \chi_2, k) = P_{FG}^{lin}(0, 0, k) g_F(\chi_1, k) g_G(\chi_2, k) \quad \text{(B.6)}$$

then the terms in brackets in the second line of Eq.(B.5) cannot be expressed as Hankel transforms. The authors of [74] work around this problem by splitting the χ space integrations into narrow enough bins such that the evolution of the scale dependence inside each

[1] All the integral kernels we consider in this work can be expressed as sums of terms with this form.

bin can be ignored. Inside each bin, the growth factor can be approximated as :

$$g(\chi,k) = g(\bar{\chi},k)\frac{g(\chi,k)}{g(\bar{\chi},k)} \approx g(\bar{\chi},k)g^{\rm eff}(\bar{\chi},\chi) \tag{B.7}$$

where $\bar{\chi}$ is the mean χ in the bin and the approximation comes from ignoring the evolution of the k-dependence. With this, we can approximate the linear power spectrum for χ_1 and χ_2 inside bins with mean $\bar{\chi}_1$ and $\bar{\chi}_2$ as:

$$\begin{aligned} P^{lin}_{FG}(\chi_1,\chi_2,k) &= P^{lin}_{FG}(0,0,k)g_F(\chi_1,k)g_G(\chi_2,k) \\ &\approx P^{lin}_{FG}(\bar{\chi}_1,\bar{\chi}_1,k)g^{\rm eff}_F(\bar{\chi_1},\chi_1)g^{\rm eff}_G(\bar{\chi}_2,\chi_2) \end{aligned} \tag{B.8}$$

and this allows us to restore the separability necessary to construct the Hankel transforms:

$$\begin{aligned} C_\ell^{F^W G^{W'}} &= \text{Limber}\Big[\int d\chi_1 d\chi_2\; W(\chi_1)\, W'(\chi_2) \\ &\times \int \frac{k^2 dk}{(2\pi)^3}\; \mathcal{K}^F_\ell(\chi_1,k)\,\mathcal{K}^G_\ell(\chi_2,k)\; P^{nlin}_{FG}(\chi_1,\chi_2,k)\Big] \\ &+ \int \frac{k^2 dk}{(2\pi)^3} \sum_i \sum_j P^{lin}_{FG}(\bar{\chi_i},\bar{\chi_j},k) \\ &\times \Big[\int d\chi_i W(\chi_i)\, g^{\rm eff}_F(\bar{\chi_i},\chi_i)\mathcal{K}^F_\ell(\chi_i,k)\Big]\Big[\int d\chi_j W'(\chi_j)\, g^{\rm eff}_G(\bar{\chi_j},\chi_j)\mathcal{K}^G_\ell(\chi_j,k)\Big] \end{aligned} \tag{B.9}$$

where the sums are over the auxiliary bins constructed to do the approximation. Since there is no limitation on how small these auxiliary bins can be, this approximation can be made as accurate as necessary.

B.1.2 Our implementation

In our implementation of the Beyond Limber method, we define the non-linear piece of the power spectrum in the following ways depending on the particular observables involved:

- For power spectra involving only dark matter, electrons or galaxies, we define the non-linear spectrum as:

$$P^{(nlin)}_{FG}(\chi_1,\chi_2,k) = P^{1h+2h}_{FG}(\chi_1,\chi_2,k) - b^F(\chi_1,k)b^G(\chi_2,k)P^{lin}_{mm}(\chi_1,\chi_2,k) \tag{B.10}$$

where $P_{FG}^{1h+2h}(\chi_1, \chi_2, k)$ is the full power spectrum computed using the halo model containing the 1-halo term and 2-halo term (see Appendix B.4), $P_{mm}^{lin}(\chi_1, \chi_2, k)$ is the linear dark matter power spectrum from CAMB, and $b^X(\chi, k)$ is the large scale linear bias function computed with the halo model, only different from 1 for galaxies (we assume electron perfectly trace dark matter for linear modes).

- For power spectra involving at least one power of the CIB or tSZ effect,we define the non-linear spectrum as:

$$P_{FG}^{(nlin)}(\chi_1, \chi_2, k) = P_{FG}^{1h+2h}(\chi_1, \chi_2, k) - P_{FG}^{2h}(\chi_1, \chi_2, k) \tag{B.11}$$

where $P_{FG}^{2h}(\chi_1, \chi_2, k)$ is the 2-halo term computed using the halo model. Effectively, we are treating the 1-halo term as the non-linear piece and the 2-halo term as the linear piece. This is not entirely correct, because the 2-halo term does account for part of the non-linearities on small scales and it is not strictly separable as in Eq.(B.6). However, a detailed inspection reveals that the 2-halo term is separable on the scales for which the second term of Eq.(B.9) finds most of its support, and that the 1-halo term dominates the regime for which the Limber approximation is adequate.

B.2 Quadratic estimators

In this appendix, we derive the unbiased and minimum variance quadratic estimator Eq. 2.32 for a modulating field $M^{\alpha}_{\ell_1 m_1}$. The starting point is the statistically anisotropic cross-power Eq. 2.31:

$$\left\langle \Theta_{\ell m}\, \delta^W_{\ell' m'} \right\rangle = (-1)^m C_\ell^{I\delta^W} \delta_{\ell\ell'} \delta_{m-m'} + \sum_{\ell_1 m_1} (-1)^{m_1} \begin{pmatrix} \ell & \ell' & \ell_1 \\ m & m' & -m_1 \end{pmatrix} f^{M^\alpha W}_{\ell\ell_1\ell'}\, M^{\alpha}_{\ell_1 m_1} \tag{B.12}$$

where the form of $f^{M^\alpha W}_{\ell\ell_1\ell'}$ depends on the observable. The quadratic estimator is of the form:

$$\hat{M}^{\alpha}_{LM} = A_L^{M^\alpha} \sum_{\ell m; \ell' m'} (-1)^M \begin{pmatrix} \ell & \ell' & L \\ m & m' & -M \end{pmatrix} G^{M^\alpha W}_{\ell\ell' L}\, \Theta_{\ell m} \delta^W_{\ell' m'} \tag{B.13}$$

Our goal is to find the appropriate weights $G^{M^\alpha W}_{\ell\ell' L}$ such that we minimize $\langle \hat{M}^{\alpha}_{LM} \hat{M}^{\alpha}_{LM} \rangle$ subject to the constraint $\langle \hat{M}^{\alpha}_{LM} \rangle = M^{\alpha}_{LM}$.

First, we find the mean of the estimator:

$$\begin{aligned}\langle \hat{M}^{\alpha}_{LM}\rangle &= A^{M^\alpha}_L \sum_{\ell m;\ell' m'} (-1)^M \begin{pmatrix} \ell & \ell' & L \\ m & m' & -M \end{pmatrix} G^{M^\alpha W}_{\ell\ell' L}\ \langle \Theta_{\ell m}\delta^W_{\ell' m'}\rangle \\ &= A^{M^\alpha}_L \sum_{\ell' m'} (-1)^{M-m'} \begin{pmatrix} \ell' & \ell' & L \\ -m' & m' & -M \end{pmatrix} G^{M^\alpha W}_{\ell\ell' L} C^{I\delta^W}_\ell \\ &+ A^{M^\alpha}_L \sum_{\ell m;\ell' m';L' M'} (-1)^{M+M'} \begin{pmatrix} \ell & \ell' & L \\ m & m' & -M \end{pmatrix} \\ &\times \begin{pmatrix} \ell & \ell' & L' \\ m & m' & -M' \end{pmatrix} G^{M^\alpha W}_{\ell\ell' L}\ f^{M^\alpha W}_{\ell L'\ell'}\ M^\alpha_{L' M'} \end{aligned} \tag{B.14}$$

We now use:

$$(-1)^{-m'} \begin{pmatrix} \ell' & \ell' & L \\ -m' & m' & -M \end{pmatrix} = (-1)^{-m'} \begin{pmatrix} \ell' & \ell' & L \\ -m' & m' & 0 \end{pmatrix} \delta_{M0} \tag{B.15}$$

as well as the following properties of 3j symbols:

$$\sum_{m'} (-1)^{-m'} \begin{pmatrix} \ell' & \ell' & L \\ -m' & m' & 0 \end{pmatrix} = (-1)^{\ell'} \sqrt{2\ell'+1}\delta_{L0} \tag{B.16}$$

and

$$\sum_{m,m'} \begin{pmatrix} \ell & \ell' & L \\ m & m' & -M \end{pmatrix} \begin{pmatrix} \ell & \ell' & L' \\ m & m' & -M' \end{pmatrix} = \frac{\delta_{LL'}\delta_{MM'}}{2L+1} \tag{B.17}$$

Substituting these relations into the estimator, we obtain:

$$\begin{aligned}\langle \hat{M}^{\alpha}_{LM}\rangle &= A^{M^\alpha}_L \sum_{\ell} (-1)^\ell \sqrt{2\ell+1}\ G^{M^\alpha W}_{\ell\ell' L} C^{I\delta^W}_\ell\ \delta_{L0}\delta_{M0} \\ &+ M^\alpha_{LM} \frac{A^{M^\alpha}_L}{2L+1} \sum_{\ell;\ell'} G^{M^\alpha W}_{\ell\ell' L} f^{M^\alpha W}_{\ell L\ell'} \end{aligned} \tag{B.18}$$

Aside from the monopole, we can make the estimator unbiased so long as:

$$A_L = (2L+1) \left(\sum_{\ell;\ell'} G^{M^\alpha W}_{\ell\ell' L} f^{M^\alpha W}_{\ell L\ell'} \right)^{-1} \tag{B.19}$$

We can now fix $G^{M^\alpha W}_{\ell\ell' L}$ by minimizing the variance of the estimator. We compute:

$$\begin{aligned}\langle \hat{M}^{\alpha *}_{LM} \hat{M}^{\alpha}_{LM} \rangle &= A_L^2 \sum_{\ell_1 m_1; \ell_2 m_2} \sum_{\ell_1' m_1'; \ell_2' m_2'} \begin{pmatrix} \ell_1 & \ell_2 & L \\ m_1 & m_2 & -M \end{pmatrix} \begin{pmatrix} \ell_1' & \ell_2' & L \\ m_1' & m_2' & -M \end{pmatrix} \\ &\times G^{M^\alpha W}_{\ell_1 \ell_2 L} G^{M^\alpha W}_{\ell_1' \ell_2' L} \langle \Theta^*_{\ell_1 m_1} \delta^{W*}_{\ell_2 m_2} \Theta_{\ell_1' m_1'} \delta^W_{\ell_2' m_2'} \rangle \end{aligned} \tag{B.20}$$

The four-point function can be decomposed into a connected and disconnected piece:

$$\begin{aligned}\langle \Theta^*_{\ell_1 m_1} \delta^{W*}_{\ell_2 m_2} \Theta_{\ell_1' m_1'} \delta^W_{\ell_2' m_2'} \rangle &= \langle \Theta^*_{\ell_1 m_1} \delta^{W*}_{\ell_2 m_2} \Theta_{\ell_1' m_1'} \delta^W_{\ell_2' m_2'} \rangle_{\rm con} \\ &+ \langle \Theta^*_{\ell_1 m_1} \delta^{W*}_{\ell_2 m_2} \Theta_{\ell_1' m_1'} \delta^W_{\ell_2' m_2'} \rangle_{\rm discon} \end{aligned} \tag{B.21}$$

Here, we minimize the variance considering the disconnected contribution only:

$$\begin{aligned}\langle \Theta^*_{\ell_1 m_1} \delta^{W*}_{\ell_2 m_2} \Theta_{\ell_1' m_1'} \delta^W_{\ell_2' m_2'} \rangle_{\rm discon} &= \langle \Theta^*_{\ell_1 m_1} \delta^{W*}_{\ell_2 m_2} \rangle \langle \Theta_{\ell_1' m_1'} \delta^W_{\ell_2' m_2'} \rangle \\ &+ \langle \Theta^*_{\ell_1 m_1} \Theta_{\ell_1' m_1'} \rangle \langle \delta^{W*}_{\ell_2 m_2} \delta^W_{\ell_2' m_2'} \rangle \\ &+ \langle \Theta^*_{\ell_1 m_1} \delta^W_{\ell_2' m_2'} \rangle \langle \Theta_{\ell_1' m_1'} \delta^{W*}_{\ell_2 m_2} \rangle \\ &= (-1)^{m_2} \delta_{\ell_1 \ell_2} \delta_{m_1 - m_2} C^{\Theta \delta^W}_{\ell_1} (-1)^{m_1'} \delta_{\ell_1' \ell_2'} \delta_{-m_1' m_2'} C^{\Theta \delta^W}_{\ell_1'} \\ &+ \delta_{\ell_1 \ell_1'} \delta_{m_1 m_1'} C^{\Theta\Theta}_{\ell_1} \delta_{\ell_2 \ell_2'} \delta_{m_2 m_2'} C^{\delta^W \delta^W}_{\ell_2} \\ &+ \delta_{\ell_1 \ell_2'} \delta_{m_1 m_2'} C^{\Theta \delta^W}_{\ell_1} (-1)^{m_2 + m_1'} \delta_{-m_2 - m_1'} \delta_{\ell_2 \ell_1'} C^{\Theta \delta^W}_{\ell_1'} \end{aligned} \tag{B.22}$$

Contributions to the connected piece are discussed in Sec.B.3.2. Plugging this expression into the variance:

$$\begin{aligned}\langle \hat{M}^{\alpha *}_{LM} \hat{M}^{\alpha}_{LM} \rangle &= A_L^2 \sum \Bigg[\begin{pmatrix} \ell_1 & \ell_2 & L \\ m_1 & m_2 & -M \end{pmatrix}^2 (G^{M^\alpha W}_{\ell_1 \ell_2 L})^2 C^{\Theta\Theta}_{\ell_1} C^{\delta^W \delta^W}_{\ell_2} \\ &+ (-1)^{m_2 + m_1'} \begin{pmatrix} \ell_2 & \ell_2 & L \\ -m_2 & m_2 & -M \end{pmatrix} \begin{pmatrix} \ell_1' & \ell_1' & L \\ m_1' & -m_1' & -M \end{pmatrix} G^{M^\alpha W}_{\ell_2 \ell_2 L} G^{M^\alpha W}_{\ell_1' \ell_1' L} C^{\Theta \delta^W}_{\ell_2} C^{\Theta \delta^W}_{\ell_1'} \\ &+ \begin{pmatrix} \ell_1 & \ell_1' & L \\ m_1 & m_1' & -M \end{pmatrix} \begin{pmatrix} \ell_1' & \ell_1 & L \\ m_1' & m_1 & -M \end{pmatrix} G^{M^\alpha W}_{\ell_1 \ell_1' L} G^{M^\alpha W}_{\ell_1' \ell_1 L} C^{\Theta \delta^W}_{\ell_1} C^{\Theta \delta^W}_{\ell_1'} \Bigg] \end{aligned} \tag{B.23}$$

We now perform the sums over m. Using Eq. B.16, the second term in parentheses contributes only to the monopole. We neglect this term in the following. To evaluate the third term, we use

$$\begin{pmatrix} \ell_1' & \ell_1 & L \\ m_1' & m_1 & -M \end{pmatrix} = (-1)^{\ell_1 + \ell_1' + L} \begin{pmatrix} \ell_1 & \ell_1' & L \\ m_1 & m_1' & -M \end{pmatrix} \tag{B.24}$$

Changing dummy indices, the variance is

$$\begin{aligned}\langle \hat{M}^{\alpha *}_{LM} \hat{M}^{\alpha}_{LM} \rangle &= A_L^2 \sum_{\ell_1 m_1; \ell_2 m_2} \left[\begin{pmatrix} \ell_1 & \ell_2 & L \\ m_1 & m_2 & -M \end{pmatrix}^2 (G^{M^\alpha W}_{\ell_1 \ell_2 L})^2 C^{\Theta\Theta}_{\ell_1} C^{\delta^W \delta^W}_{\ell_2} \right. \\ &+ \left. (-1)^{\ell_1 + \ell_1' + L} \begin{pmatrix} \ell_1 & \ell_2 & L \\ m_1 & m_2 & -M \end{pmatrix}^2 G^{M^\alpha W}_{\ell_1 \ell_2 L} G^{M^\alpha W}_{\ell_2 \ell_1 L} C^{\Theta \delta^W}_{\ell_1} C^{\Theta \delta^W}_{\ell_2} \right] \end{aligned} \tag{B.25}$$

Using Eq. B.17, we can perform the sums over m_1, m_2 to obtain:

$$\begin{aligned}\langle \hat{M}^{\alpha *}_{LM} \hat{M}^{\alpha}_{LM} \rangle &= \frac{1}{2L+1} \sum_{\ell_1; \ell_2} A_L G^{M^\alpha W}_{\ell_1 \ell_2 L} \left[A_L G^{M^\alpha W}_{\ell_1 \ell_2 L} C^{\Theta\Theta}_{\ell_1} C^{\delta^W \delta^W}_{\ell_2} \right. \\ &+ \left. (-1)^{\ell_1 + \ell_2 + L} A_L G^{M^\alpha W}_{\ell_2 \ell_1 L} C^{\Theta \delta^W}_{\ell_1} C^{\Theta \delta^W}_{\ell_2} \right] \end{aligned} \tag{B.26}$$

To minimize the variance, we can use the Lagrange Multiplier method. First, let's define

$$F_{\ell_1 \ell_2 L} \equiv A_L G^{M^\alpha W}_{\ell_1 \ell_2 L} \tag{B.27}$$

The variance can therefore be written as:

$$\langle \hat{M}^{\alpha *}_{LM} \hat{M}^{\alpha}_{LM} \rangle = \frac{1}{2L+1} \sum_{\ell_1; \ell_2} F_{\ell_1 \ell_2 L} \left[F_{\ell_1 \ell_2 L} C^{\Theta\Theta}_{\ell_1} C^{\delta^W \delta^W}_{\ell_2} + (-1)^{\ell_1 + \ell_2 + L} F_{\ell_2 \ell_1 L} C^{\Theta \delta^W}_{\ell_1} C^{\Theta \delta^W}_{\ell_2} \right] \tag{B.28}$$

We want to minimize the variance subject to the constraint that the estimator is unbiased, which is enforced by Eq. B.19. This condition translates to:

$$\frac{1}{2L+1} \sum_{\ell_1; \ell_2} F_{\ell_1 \ell_2 L} f^{M^\alpha W}_{\ell_1 L \ell_2} - 1 = 0 \tag{B.29}$$

We therefore want to evaluate:

$$\begin{aligned} 0 &= \frac{\delta}{\delta F_{\ell_1 \ell_2 L}} \left(\frac{1}{2L+1} \sum_{\ell_1; \ell_2} F_{\ell_1 \ell_2 L} \left[F_{\ell_1 \ell_2 L} C^{\Theta\Theta}_{\ell_1} C^{\delta^W \delta^W}_{\ell_2} + (-1)^{\ell_1 + \ell_2 + L} F_{\ell_2 \ell_1 L} C^{\Theta \delta^W}_{\ell_1} C^{\Theta \delta^W}_{\ell_2} \right] \right. \\ &+ \left. \lambda \left[\frac{1}{2L+1} \sum_{\ell_1; \ell_2} F_{\ell_1 \ell_2 L} f^{M^\alpha W}_{\ell_1 L \ell_2} - 1 \right] \right) \end{aligned} \tag{B.30}$$

where λ is the Lagrange multiplier. Evaluating the derivative yields

$$F_{\ell_1 \ell_2 L} C^{\Theta\Theta}_{\ell_1} C^{\delta^W \delta^W}_{\ell_2} + (-1)^{\ell_1 + \ell_2 + L} F_{\ell_2 \ell_1 L} C^{\Theta \delta^W}_{\ell_1} C^{\Theta \delta^W}_{\ell_2} + \lambda f^{M^\alpha W}_{\ell_1 L \ell_2} = 0 \tag{B.31}$$

Multiplying by $C_{\ell_2}^{\Theta\Theta} C_{\ell_1}^{\delta^W\delta^W}$ and subtracting $(-1)^{\ell_1+\ell_2+L} C_{\ell_1}^{\Theta\delta^W} C_{\ell_2}^{\Theta\delta^W}$ times Eq. B.31 with permuted indices $\ell_1 \leftrightarrow \ell_2$ we obtain:

$$F_{\ell_1\ell_2 L} + \lambda \frac{C_{\ell_2}^{\Theta\Theta} C_{\ell_1}^{\delta^W\delta^W} f_{\ell_1 L\ell_2}^{M^\alpha W} - (-1)^{\ell_1+\ell_2+L} C_{\ell_1}^{\Theta\delta^W} C_{\ell_2}^{\Theta\delta^W} f_{\ell_2 L\ell_1}^{M^\alpha W}}{C_{\ell_1}^{\Theta\Theta} C_{\ell_2}^{\Theta\Theta} C_{\ell_1}^{\delta^W\delta^W} C_{\ell_2}^{\delta^W\delta^W} - (C_{\ell_1}^{\Theta\delta^W})^2 (C_{\ell_2}^{\Theta\delta^W})^2} = 0 \tag{B.32}$$

For now, let's define a new function:

$$h_{\ell_1\ell_2 L} \equiv \frac{C_{\ell_2}^{\Theta\Theta} C_{\ell_1}^{\delta^W\delta^W} f_{\ell_1 L\ell_2}^{M^\alpha W} - (-1)^{\ell_1+\ell_2+L} C_{\ell_1}^{\Theta\delta^W} C_{\ell_2}^{\Theta\delta^W} f_{\ell_2 L\ell_1}^{M^\alpha W}}{C_{\ell_1}^{\Theta\Theta} C_{\ell_2}^{\Theta\Theta} C_{\ell_1}^{\delta^W\delta^W} C_{\ell_2}^{\delta^W\delta^W} - (C_{\ell_1}^{\Theta\delta^W})^2 (C_{\ell_2}^{\Theta\delta^W})^2} \tag{B.33}$$

so that:

$$F_{\ell_1\ell_2 L} = -\lambda h_{\ell_1\ell_2 L}. \tag{B.34}$$

Multiplying by $f_{\ell_1 L\ell_2}^{M^\alpha W}$ and using the no bias condition Eq. B.29, we can solve for the Lagrange multiplier:

$$\lambda = -(2L+1) \left[\sum_{\ell_1\ell_2} h_{\ell_1\ell_2 L} f_{\ell_1 L\ell_2}^{M^\alpha W} \right]^{-1} \tag{B.35}$$

Substituting this into Eq. B.34, we have

$$F_{\ell_1\ell_2 L} = (2L+1) h_{\ell_1\ell_2 L} \left[\sum_{\ell\ell'} h_{\ell\ell' L} f_{\ell L\ell'}^{M^\alpha W} \right]^{-1} = (2L+1) G_{\ell_1\ell_2 L}^{M^\alpha W} \left[\sum_{\ell\ell'} G_{\ell\ell' L}^{M^\alpha W} f_{\ell L\ell'}^{M^\alpha W} \right]^{-1} \tag{B.36}$$

and so we can identify $G_{\ell_1\ell_2 L}^{M^\alpha W} = h_{\ell_1\ell_2 L}$ as the choice that minimizes the variance:

$$G_{\ell_1\ell_2 L}^{M^\alpha W} \equiv \frac{C_{\ell_2}^{\Theta\Theta} C_{\ell_1}^{\delta^W\delta^W} f_{\ell_1 L\ell_2}^{M^\alpha W} - (-1)^{\ell_1+\ell_2+L} C_{\ell_1}^{\Theta\delta^W} C_{\ell_2}^{\Theta\delta^W} f_{\ell_2 L\ell_1}^{M^\alpha W}}{C_{\ell_1}^{\Theta\Theta} C_{\ell_2}^{\Theta\Theta} C_{\ell_1}^{\delta^W\delta^W} C_{\ell_2}^{\delta^W\delta^W} - (C_{\ell_1}^{\Theta\delta^W})^2 (C_{\ell_2}^{\Theta\delta^W})^2} \tag{B.37}$$

B.3 Additional contributions to the estimator mean and variance

In this appendix we discuss the contributions to the mean and variance of the quadratic estimator for the radial velocity field that arise in the presence of additional non-Gaussian contributions to correlation functions between the CMB temperature and density field. The various contributions to the CMB temperature listed in Eq. 2.79 generate non-trivial 3- and 4-point functions between the measured CMB temperature and density tracer. There are

two distinct types of non-Gaussian contributions that we must consider. The first type is what our quadratic estimators are based on: the non-Gaussianity associated with the fact that CMB secondaries are line-of-sight integrals over products of fields. The second type is the intrinsic non-Gaussianity of the density and velocity fields due to gravitational collapse (or primordial non-Gaussianity, though we expect this to be negligibly small). A complete assessment of the magnitude of the many contributions to the mean and variance of the quadratic estimators is beyond the scope of the this thesis, and will be evaluated in future work. Here, we only attempt to enumerate the contributions that must be considered, and in some cases, estimate their magnitude.

B.3.1 Estimator mean

The mean of the estimators $\langle \hat{M}^{\alpha}_{LM} \rangle$ depend on the two-point function $\langle \Theta_{\ell m} \delta^{W}_{\ell' m'} \rangle$. Quadratic estimators for the radial velocity fields are based on non-Gaussian contributions to the correlation functions $\langle \Theta^{kSZ}_{\ell m} \delta^{W}_{\ell' m'} \rangle$. Because the kSZ temperature anisotropies depend on the product of density contrast and velocity, these are in fact three-point functions. Above, we considered the squeezed limit of these correlators, where the velocity mode is of much larger wavelength than the density modes. Additionally, we treated the velocity and density fields as Gaussian. When the velocity mode is of comparable wavelength to the density modes, there will be a contribution to the three-point function due to gravitational collapse. We expect this to be important at high L, beyond the regime of interest for velocity reconstruction. Likewise, the contributions from lensing $\langle \Theta^{L}_{\ell m} \delta^{W}_{\ell' m'} \rangle$ and non-linear ISW $\langle \Theta^{ISW,nlin}_{\ell m} \delta^{W}_{\ell' m'} \rangle$ treated in Sec. 2.3.6 will receive contributions from non-linearities on small scales, but since the leading order bias these signals is so tiny, we expect these additional contributions to be completely negligible on scales of interest. Also considered in Sec. 2.3.6 was the contribution to the estimator mean from systematics that modulate the observed density field. Similar systematics in the galaxy survey or CMB experiment will lead to similar effects.

Another contribution to the mean of the estimators, which was not considered above, arises from non-linear terms in $\langle \Theta^{XG}_{\ell m} \delta^{W}_{\ell' m'} \rangle$. On scales $\ell, \ell' \gg 1$ where the estimator receives most of it's weight, we must include non-linear contributions to the galaxy density field as well as the extragalactic foregrounds (here, the CIB and tSZ). At second order in perturbation theory, schematically we must consider correlators of the form $\langle t\delta\delta \rangle$ where t is the large-scale tidal field (see e.g. Ref. [183, 184]). It is difficult to imagine this term being larger than the bias induced by calibration error, which takes a similar form, and which is likely far larger in magnitude than the large-scale tidal field. Related to the tidal field, systematics associated with the intrinsic alignment of galaxies lead to a large-scale

statistical anisotropy in the galaxy number counts [185]; again, it is difficult to imagine that the amplitude of this effect is large enough to cause a significant bias. We defer a detailed estimate of these and other effects to future work.

B.3.2 Estimator variance

Above, we considered only the disconnected contributions to the estimator variance $\langle \hat{M}^{\alpha}_{LM} \hat{M}^{\beta}_{LM} \rangle$. There are a number of additional contributions to the variance, arising from the non-Gaussian nature of the kSZ effect as well as other non-Gaussian contributions to the CMB temperature and galaxy survey. Concentrating on non-kSZ, non-Gaussian contributions to the estimator variance Eq. B.20, we conjecture that the most important terms arise from:

$$\langle \Theta^{XG*}_{\ell_1 m_1} \delta^{W*}_{\ell_2 m_2} \Theta^{XG}_{\ell_1' m_1'} \delta^{W}_{\ell_2' m_2'} \rangle_{\rm con}, \quad \langle \Theta^{L*}_{\ell_1 m_1} \delta^{W*}_{\ell_2 m_2} \Theta^{L}_{\ell_1' m_1'} \delta^{W}_{\ell_2' m_2'} \rangle_{\rm con}, \tag{B.38}$$

Note that the relevant shape of the four point function for the estimator variance is the "collapsed" configuration where $\ell_1 \sim \ell_2$ and $\ell_1' \sim \ell_2'$, since the relevant scales are $L \ll \ell_1, \ell_2, \ell_1', \ell_2'$. The terms in Eq. B.38 should be calculable analytically within the halo model since the collapsed four-point function typically has a simple form [186]. For example, similar computations have been performed in the context of the CIB have been performed [187]. Roughly speaking, we expect the disconnected four-point function to dominate the connected four-point function by a power of the matter power spectrum. Therefore, including the connected four-point function will most likely not make a large contribution to the estimator variance. We leave a detailed computation to future work.

Because the kSZ temperature anisotropies arise do to the product of the optical depth and radial velocity, evaluating kSZ contributions to the estimator variance involves computing a six-point function. In terms of the Haar-binned LC moments of v and $\dot{\tau}$ we have:

$$\begin{aligned} \langle \Theta^{kSZ*}_{\ell_1 m_1} \delta^{W*}_{\ell_2 m_2} \Theta^{kSZ}_{\ell_1' m_1'} \delta^{W}_{\ell_2' m_2'} \rangle &= \sum_{\bar{\ell}_1 \bar{m}_1; \bar{\ell}_2 \bar{m}_2} \sum_{\tilde{\ell}_1 \tilde{m}_1; \tilde{\ell}_2 \tilde{m}_2} R^{\bar{\ell}_1 \bar{\ell}_2 \ell_1}_{\bar{m}_1 \bar{m}_2 - m_1} R^{\tilde{\ell}_1 \tilde{\ell}_2 \ell_1'}_{\tilde{m}_1 \tilde{m}_2 - m_1'} \\ &\times \sum_{ss'} \langle v^{s*}_{\bar{\ell}_1 \bar{m}_1} \dot{\tau}^{s*}_{\bar{\ell}_2 \bar{m}_2} \delta^{W*}_{\ell_2 m_2} v^{s'}_{\tilde{\ell}_1 \tilde{m}_1} \dot{\tau}^{s'}_{\tilde{\ell}_2 \tilde{m}_2} \delta^{W}_{\ell_2' m_2'} \rangle \end{aligned} \tag{B.39}$$

where

$$R^{\ell_1 \ell_2 \ell}_{m_1 m_2 - m} = (-1)^m \sqrt{\frac{(2\ell+1)(2\ell_1+1)(2\ell_2+1)}{4\pi}} \begin{pmatrix} \ell_1 & \ell_2 & \ell \\ 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} \ell_1 & \ell_2 & \ell \\ m_1 & m_2 & -m \end{pmatrix} \tag{B.40}$$

To compute the six point function we must consider both connected and disconnected components. There are a total of 15 terms in the disconnected six point function. We can use the fact that the four-point function Eq. B.39 takes the collapsed configuration, together with the property that the velocity power spectrum falls rapidly with ℓ to argue that the relevant scales are $\bar{\ell}_1 \ll \ell_1$, $\tilde{\ell}_1 \ll \ell_1'$, $\ell_1 \sim \ell_2$, and $\ell_1' \sim \ell_2'$. From the 3j symbols in the coupling functions $R^{\ell_1\ell_2\ell}_{m_1m_2-m}$, this in turn implies that $\bar{\ell}_2 \sim \ell_2$ and $\tilde{\ell}_1 \sim \ell_2'$. Therefore, correlators involving the velocity (which is relevant at low-ℓ) and either $\dot{\tau}$ or δ^W (which are relevant at high-ℓ) will not make a significant contribution to the disconnected six-point function. We can therefore make the approximation:

$$\begin{aligned}\langle v^{s*}_{\bar{\ell}_1\bar{m}_1}\dot{\tau}^{s*}_{\bar{\ell}_2\bar{m}_2}\delta^{W*}_{\ell_2m_2}v^{s'}_{\tilde{\ell}_1\tilde{m}_1}\dot{\tau}^{s'}_{\tilde{\ell}_2\tilde{m}_2}\delta^{W}_{\ell_2'm_2'}\rangle_{\rm discon} &\simeq \langle v^{s*}_{\bar{\ell}_1\bar{m}_1}v^{s'}_{\tilde{\ell}_1\tilde{m}_1}\rangle\Big[\langle\dot{\tau}^{s*}_{\bar{\ell}_2\bar{m}_2}\delta^{W*}_{\ell_2m_2}\rangle\langle\dot{\tau}^{s'}_{\tilde{\ell}_2\tilde{m}_2}\delta^{W}_{\ell_2'm_2'}\rangle \\ &+ \langle\dot{\tau}^{s*}_{\bar{\ell}_2\bar{m}_2}\dot{\tau}^{s'}_{\tilde{\ell}_2\tilde{m}_2}\rangle\langle\delta^{W*}_{\ell_2m_2}\delta^{W}_{\ell_2'm_2'}\rangle \\ &+ \langle\dot{\tau}^{s*}_{\bar{\ell}_2\bar{m}_2}\delta^{W}_{\ell_2'm_2'}\rangle\langle\dot{\tau}^{s'}_{\tilde{\ell}_2\tilde{m}_2}\delta^{W*}_{\ell_2m_2}\rangle\Big] \end{aligned} \tag{B.41}$$

As we now show, the first term gives rise to the signal covariance, the second term reproduces the Gaussian estimator variance, and the third term yields the "$N^{(1)}$ bias" from Ref. [21].

Substituting the first term into the estimator variance Eq. B.20, we obtain:

$$\begin{aligned}&\langle\hat{M}^{\alpha}_{LM}\hat{M}^{\beta}_{LM}\rangle_{kSZ,1} \\ =\;& \sum_{s,s'=1}^{\infty}(C^{vv})^{ss'}_L \\ &\left(\sum_{\ell_1\ell_2}\frac{A^{M^\alpha}_L}{2L+1}G^{M^\alpha W}_{\ell_1\ell_2L}\sqrt{\frac{(2\ell_1+1)(2\ell_2+1)(2L+1)}{4\pi}}\begin{pmatrix}\ell_1 & \ell_2 & L\\ 0 & 0 & 0\end{pmatrix}C^{\dot{\tau}^sW}_{\ell_2}\right) \\ &\left(\sum_{\ell_1\ell_2}\frac{A^{M^\beta}_L}{2L+1}G^{M^\beta W}_{\ell_1\ell_2L}\sqrt{\frac{(2\ell_1+1)(2\ell_2+1)(2L+1)}{4\pi}}\begin{pmatrix}\ell_1 & \ell_2 & L\\ 0 & 0 & 0\end{pmatrix}C^{\dot{\tau}^{s'}W}_{\ell_2}\right) \\ =\;& \sum_{s,s'=1}^{\infty}(C^{vv})^{ss'}_L\left(\sum_{\ell_1\ell_2}\frac{A^{M^\alpha}_L}{2L+1}G^{M^\alpha W}_{\ell_1\ell_2L}f^{v^sW}_{\ell_1L\ell_2}\right)\left(\sum_{\ell_1\ell_2}\frac{A^{M^\beta}_L}{2L+1}G^{M^\beta W}_{\ell_1\ell_2L}f^{v^{s'}W}_{\ell_1L\ell_2}\right) \\ =\;& \sum_{s,s'=1}^{\infty}(C^{vv})^{ss'}_L R^{v^sM^\alpha}_L R^{v^{s'}M^\beta}_L \end{aligned} \tag{B.42}$$

This is the signal covariance rotated into the basis defined by the estimators.

Moving to the second term:

$$\begin{aligned}\langle \hat{M}^{\alpha}_{LM}\hat{M}^{\beta}_{LM}\rangle_{kSZ,2} &= A_L^{M^\alpha}A_L^{M^\beta}\sum_{\ell_1\ell_2}\frac{(-1)^{\ell_1+\ell_2+L}}{2L+1}G^{M^\alpha W}_{\ell_1\ell_2 L}G^{M^\beta W}_{\ell_1\ell_2 L}C^{WW}_{\ell_2}\\ &\quad\sum_{\bar{\ell}_1\bar{\ell}_2}\frac{(2\bar{\ell}_1+1)(2\bar{\ell}_2+1)}{4\pi}\begin{pmatrix}\bar{\ell}_1 & \bar{\ell}_2 & \ell_1\\ 0 & 0 & 0\end{pmatrix}^2\sum_{s,s'=1}^{\infty}(C^{vv})^{ss'}_{\bar{\ell}_1}(C^{\dot{\tau}\dot{\tau}})^{ss'}_{\bar{\ell}_2}\\ &= A_L^{M^\alpha}A_L^{M^\beta}\sum_{\ell_1\ell_2}\frac{(-1)^{\ell_1+\ell_2+L}}{2L+1}G^{M^\alpha W}_{\ell_1\ell_2 L}G^{M^\beta W}_{\ell_1\ell_2 L}C^{WW}_{\ell_2}C^{\rm kSZ}_{\ell_1} \qquad \text{(B.43)}\end{aligned}$$

This term combines with the non-kSZ disconnected components of the temperature galaxy four-point function to yield the estimator noise.

The third term is somewhat more complicated,

$$\begin{aligned}\langle \hat{M}^{\alpha}_{LM}\hat{M}^{\beta}_{LM}\rangle_{kSZ,3} &= A_L^{M^\alpha}A_L^{M^\beta}\sum_{\ell_1;\ell_2}\sum_{\ell_1';\ell_2'}\sum_{\bar{\ell}_1}G^{M^\alpha W}_{\ell_1\ell_2 L}G^{M^\beta W}_{\ell_1'\ell_2' L'}\frac{2\bar{\ell}_1+1}{4\pi}\\ &\times \sum_{s,s'=1}^{\infty}(C^{vv})^{ss'}_{\bar{\ell}_1}C^{\dot{\tau}^s W}_{\ell_2'}C^{\dot{\tau}^{s'} W}_{\ell_2}\sqrt{(2\ell_1+1)(2\ell_2+1)(2\ell_1'+1)(2\ell_2'+1)}\\ &\times \begin{pmatrix}\bar{\ell}_1 & \ell_2' & \ell_1\\ 0 & 0 & 0\end{pmatrix}\begin{pmatrix}\bar{\ell}_1 & \ell_2 & \ell_1'\\ 0 & 0 & 0\end{pmatrix}\frac{(-1)^{L+\bar{\ell}_1}}{2L+1}\begin{Bmatrix}\ell_1' & \ell_2' & L\\ \ell_1 & \ell_2 & \bar{\ell}_1\end{Bmatrix} \qquad \text{(B.44)}\end{aligned}$$

When L is much smaller than the other factors in the 6j symbol, we can simplify using:

$$\begin{Bmatrix}\ell_1' & \ell_2' & 0\\ \ell_1 & \ell_2 & \bar{\ell}_1\end{Bmatrix} = \frac{\delta_{\ell_1,\ell_2}\delta_{\ell_1',\ell_2'}}{\sqrt{(2\ell_1'+1)(2\ell_2+1)}}(-1)^{\ell_1'+\ell_2+\bar{\ell}_1}\left\{\ell_1',\ell_2,\bar{\ell}_1\right\} \qquad \text{(B.45)}$$

This gives

$$\begin{aligned}\langle \hat{M}^{\alpha}_{LM}\hat{M}^{\beta}_{LM}\rangle_{kSZ,3} &= A_L^{M^\alpha}A_L^{M^\beta}\sum_{\ell_1';\ell_2}\sum_{\bar{\ell}_1}G^{M^\alpha W}_{\ell_2\ell_2 L}G^{M^\beta W}_{\ell_1'\ell_1' L}\sum_{s,s'=1}^{\infty}\frac{2\bar{\ell}_1+1}{4\pi}(C^{vv})^{ss'}_{\bar{\ell}_1}C^{\dot{\tau}^s W}_{\ell_1'}C^{\dot{\tau}^{s'} W}_{\ell_2}\\ &\times \sqrt{(2\ell_2+1)(2\ell_1'+1)}\begin{pmatrix}\bar{\ell}_1 & \ell_1' & \ell_2\\ 0 & 0 & 0\end{pmatrix}^2 \qquad \text{(B.46)}\end{aligned}$$

This is the $N^{(1)}$ bias first computed in [21]. Evaluating it, we find, in agreement with [21], that this term is negligible compared to the Gaussian estimator noise.

Another contribution to the estimator variance arises due to the connected six-point function, the "$N^{(3/2)}$ bias", which was found in Ref. [21] to be even larger than the Gaussian estimator noise in the high signal-to-noise regime. A full computation of this term is beyond the scope of this thesis, but will be necessary for a complete analysis in the future.

B.4 Halo Model

In this Appendix, we describe the assumptions made in our halo model description of various tracers of large scale structure including: dark matter density, galaxy number counts, electron density, cosmic infrared background (CIB), and the thermal Sunyaev Zel'dovich (tSZ) effect. For a general review of the halo model of large scale structure, see e.g. Ref. [68]. The final product of the numerical computations is a set of auto- and cross-power spectra at a set of redshifts (dark matter density, galaxy number counts, and electron density) and/or frequencies (CIB and tSZ). These auto- and cross-power spectra are then converted to angular spectra as a function of redshift bin and/or frequency using the techniques described in Sec. 2.2.3.

B.4.1 Halo mass function, halo bias, and the matter power spectrum

Halo mass function

Within the halo model, all matter is distributed in discrete halos of different sizes. The halo mass function $\frac{dN}{dM}$ describes the distribution of the halos: the number density of halos $n^h(z)$ between masses M_1 and M_2 at z is given by

$$n^h(z) = \int_{M_1}^{M_2} \frac{dN}{dM}(M, z) dM. \tag{B.47}$$

In our halo model, we use the halo mass function of [189], which parametrizes the halo multiplicity function $f(\upsilon)$ as

$$f(\upsilon) = \alpha(1 + (\beta\upsilon)^{-2\phi})\upsilon^{2\eta} e^{-\eta\upsilon^2/2}. \tag{B.48}$$

$f(\upsilon)$ is related to $\frac{dN}{dM}$ as (see Eq. (2) of [192])

$$\frac{dN}{dM} = \upsilon f(\upsilon) \frac{\rho_m}{M} \frac{d \ln \sigma^{-1}}{dM} \tag{B.49}$$

where ρ_m is the present day cosmological matter density and the peak height υ is

$$\upsilon \equiv \frac{\delta_c}{\sigma} \tag{B.50}$$

with $\delta_c = 1.686$ the critical density required for collapse. σ is the linear matter variance smoothed with a top-hat function over the radius of the halo $R = \left(\frac{3M}{4\pi\rho_m}\right)^{1/3}$

$$\sigma^2(R,z) = \frac{1}{2\pi^2}\int P(k,z)\hat{W}(k,R)k^2 dk \tag{B.51}$$

where $\hat{W}(k,R)$ is the Fourier transform (in k) of a top-hat function with radial extent R.

The values of the parameters $\{\beta,\gamma,\phi,\eta\}$ are listed in Table 4 of [189] with a mild redshift dependence given in Eqs (9)-(12) of [189]. The value of α results from applying the z-dependent normalization condition to be discussed below in Section B.4.1.

Halo bias

Halos are biased with respect to the underlying dark matter power spectrum; in particular, the power spectrum of halos of masses M at redshift z P_{hh} can be written (on large scales, where the bias is scale-independent) as

$$P_{hh}(k,M,z) = b_h(M,z)P_{\rm lin}(k,z), \tag{B.52}$$

where $P_{\rm lin}(k,z)$ is the linear dark matter power spectrum and $b_h(M,z)$ is the halo bias. In our halo model, we use the halo bias of [189], which is parametrized as

$$b(\upsilon) = 1 - A\frac{\upsilon^a}{\upsilon^a + \delta_c^a} + B\upsilon^b + C\upsilon^c. \tag{B.53}$$

The values of the parameters $\{A,a,B,b,C,c\}$ are listed in Table 2 of [189].

Halo density profile

The halo densty profile $\rho(\vec{r},M,z)$ gives the density at a displacement $\vec{r}$ from the centre of a halo and thus governs the distribution of dark matter within a halo. For spherically symmetric halos, $\rho(\vec{r}) = \rho(r)$. We take $\rho(r)$ to be Navarro–Frenk–White (NFW) [190], ie

$$\rho(r) = \frac{\rho_S}{\frac{r}{r_S}\left(1+\frac{r}{r_S}\right)^2} \tag{B.54}$$

with r_S the scale radius, a parameter which is related to the halo radius r_M by the concentration parameter $c = \frac{r_M}{r_S}$; the scale density ρ_S defines the density of the halo, and

can be eliminated in favour of the virial radius and mass by using the definition of mass $M = \int_0^{r_M} 4\pi r^3 \rho(r) dr$. We use the halo concentration parametrization found in [191], which parametrizes the concentration of halos as

$$c = A \left(\frac{M}{M_{\rm pivot}} \right)^B (1+z)^C. \tag{B.55}$$

The values of A, B, C depend on the definition of the halo mass one is using and can be found in Table 1 of [191] (in particular we take the Sample-F redshift 0-2 row). Note that Ref. [191] provides different values for the parameters depending on the definition of the halo mass considered; we take M to be the mass within the radius R_{200m} for which the mean density of the halo is 200 times the mean matter density (labeled $M_{\rm mean}$ in Ref. [191]).

In power spectra, the normalized Fourier transform of $\rho(r)$

$$u(k, M, z) \equiv \frac{\int dr 4\pi r^2 \frac{\sin(kr)}{kr} \rho(r, M, z)}{\int dr 4\pi r^2 \rho(r, M, z)} \tag{B.56}$$

is used.

Dark matter power spectrum

Within the halo model, power spectra are split into a term sourced by correlations in different halos (inter-halo correlations), and correlations within a single halo (intra-halo correlations). These terms are known as the 2-halo and 1-halo power spectra respectively, so we have

$$P_{mm}(k, z) = P_{mm}^{2h}(k, z) + P_{mm}^{1h}(k, z) \tag{B.57}$$

where P_{mm}^{2h} and P_{mm}^{1h} denote the 2-halo and 1-halo dark matter power spectra respectively, and P_{mm} is the total dark matter power spectrum. Each term is an integral over all halo masses:

$$P_{mm}^{2h}(k, z) = \left(\int dM \frac{dN}{dM} b_h(M, z) \frac{M}{\rho_m} u(k, M, z) \right)^2 P_{\rm lin}(k, z); \tag{B.58}$$

$$P_{mm}^{1h}(k, z) = \int dM \frac{dN}{dM} \left(\frac{M}{\rho_m} u(k, M, z) \right)^2. \tag{B.59}$$

On large scales, it is a requirement that $P_{mm}^{2h}(k, z) = P_{\rm lin}(k, z)$; this is a consistency condition that ensures that all dark matter resides in halos, and that it is unbiased with respect

to itself. This results in the following normalization condition:

$$\int b(\nu) f(\nu) d\nu = 1. \tag{B.60}$$

This consistency condition results in a z-dependent constraint on the normalization of the halo mass function: it fixes the value of the parameter α in Equation (B.48).

B.4.2 Large-scale structure tracers

The large-scale structure tracers we are interested in are the galaxy density g, electron density e, the CIB flux density I_ν at frequency ν, and the tSZ temperature anisotropy $\Theta^{\mathrm{tSZ},\nu}$ at frequency ν. Below, we summarize for each tracer the essential details necessary for constructing auto-power and cross-power spectra in the halo model.

Galaxy density

Galaxies are distributed in halos according to a *halo occupation distribution* (HOD). In our HOD, we assign one "central" galaxy to the centre of halos in a mass-dependent way, and additional "satellite" galaxies which are distributed throughout the halo according to the dark matter distribution. Thus, the number of galaxies in a halo of mass M at redshift z is

$$N^{\mathrm{gal}}(M, z) = N^{\mathrm{cen}}(M, z) + N^{\mathrm{sat}}(M, z) \tag{B.61}$$

where N^{cen} denotes the number of central galaxies (always 0 or 1) and N^{sat} the number of satellite galaxies. We use the same HOD as Ref. [20]. The mean number density of galaxies at z is then

$$\bar{n}^g(z) = \int dM \frac{dN}{dM} \left(N^{\mathrm{cen}}(M, z) + N^{\mathrm{sat}}(M, z) \right). \tag{B.62}$$

Electron density

Electrons are distributed inside dark matter halos according to a radial density profile $\rho_e(r)$. As a fiducial model, we choose the 'AGN' gas profiles from [75]. The Fourier space density profile for electrons is then:

$$u_e(k, M, z) = \frac{1}{M_{AGN}} \int_0^R 4\pi r^2 \rho_{AGN}(r) \frac{\sin(kr)}{kr} dr \tag{B.63}$$

where M_{AGN} is the AGN "mass" $M_{AGN} = \int_0^R 4\pi r^2 \rho_{AGN}(r) dr$, with R the cutoff radius at which we cut off the NFW profile Eq.(B.54) (R_{200m} for us).

CIB flux density

The CIB flux density at frequency ν I_ν is given by an integral over the CIB emissivity density $j_\nu(\hat{\mathbf{n}}, \chi)$:

$$I_\nu(\hat{\mathbf{n}}) = \int d\chi a(\chi) j_\nu(\hat{\mathbf{n}}, \chi). \tag{B.64}$$

This can be written as an integral over galaxies with different luminosity densities: the mean emissivity density is

$$\bar{j}_\nu(\chi) = \int dL_{(1+z)\nu} \frac{dN}{dL_{(1+z)\nu}} \frac{L_{(1+z)\nu}}{4\pi} \tag{B.65}$$

where $L_{(1+z)\nu}$ is the luminosity density and $\frac{dN}{dL_{(1+z)\nu}}$ is the halo luminosity function defined in analogy with the halo mass function; the factor of $(1+z)$ in the frequency accounts for the fact that the photons that we receive have been redshifted. Neglecting scatter between M and L_ν, this can be written as an integral over the halo mass function

$$\bar{j}_\nu(\chi) = \int dM \frac{dN}{dM} \frac{L_{(1+z)\nu}}{4\pi}. \tag{B.66}$$

As all luminosity is sourced by galaxies, L_ν can be separated into that sourced by the central galaxies and the satellite galaxies:

$$L_\nu(M, z) = L_\nu^{\text{cen}}(M, z) + L_\nu^{\text{sat}}(M, z) \tag{B.67}$$

tSZ temperature

The tSZ temperature anisotropy at frequency ν is given by

$$\frac{\Delta T^{\text{tSZ}}}{T}(\hat{\mathbf{n}}, \chi) = g_\nu y(\hat{\mathbf{n}}, \chi) \tag{B.68}$$

where $y(\hat{\mathbf{n}}, x)$ is the Compton y-parameter and g_ν is the spectral function of the tSZ

$$g_\nu = x \coth \frac{x}{2} - 4 \tag{B.69}$$

with the dimensionless variable x given by $x \equiv \frac{h\nu}{k_B T_{\rm CMB}}$ (where h is Planck's constant; k_B is the Boltzmann constant; and $T_{\rm CMB}$ is the temperature of the black-body CMB). The Compton y-parameter is a line-of-sight integral over electron pressure

$$y(\hat{\mathbf{n}}, \chi) = \frac{\sigma_T}{m_e c^2} \int d\chi a(\chi) P_e(\hat{\mathbf{n}}, \chi) \tag{B.70}$$

where $P_e(\hat{\mathbf{n}}, \chi)$ is the electron pressure at $(\hat{\mathbf{n}}, \chi)$, and where σ_T is the Thompson scattering cross section; m_e is the electron mass; and c is the speed of light.

To calculate the power spectrum of y, we need the three-dimensional Fourier transform of $P_e(r)$; for spherically symmetric halos this allows us to define the profile

$$y(k, M, z) \equiv \frac{4\pi\sigma_T a}{m_e c^2} \int dr r^2 \frac{\sin(kr)}{kr} P_e(r). \tag{B.71}$$

We use the pressure profiles of [75] in our model.

B.4.3 2-halo power spectra

The 2-halo power spectra are all of the form

$$P^{2h}_{XY}(k, z) = D_X(k, z) D_Y(k, z) P_{\rm lin}(k, z) \tag{B.72}$$

where $P_{\rm lin}(k, z)$ is the linear dark matter power spectrum and $D_X(k, z)$ takes the form

$$D_X(k, z) = \int dM \frac{dN}{dM} b_h(M, z) A_X(M, k, z) \tag{B.73}$$

with $A_X(M, k, z)$ the Fourier profile of the observable X.

The profile $A_X(M, k, z)$ in Equation (B.73) is specific to the observable X. For matter, we saw in Sec. B.4.1 that

$$A_m(M, k, z) = \left(\frac{M}{\rho_m}\right) u(k, M, z) \tag{B.74}$$

where $u(k, M, z)$ is the normalized Fourier-transformed dark matter halo density profile.

For the other observables we have

$$A_e(M,k,z) = \left(\frac{M}{\rho_m}\right) u_e(k,M,z); \tag{B.75}$$

$$A_g(M,k,z) = \frac{N^{\rm cen}(M,z) + N^{\rm sat}(M,z)u(k,m,z)}{\bar{n}_g(z)}; \tag{B.76}$$

$$A_{j_\nu}(M,k,z) = \frac{1}{4\pi}\left(L^{\rm cen}_{(1+z)\nu}(M,z) + L^{\rm sat}_{(1+z)\nu}(M,z)u(k,M,z)\right); \tag{B.77}$$

$$A_y(M,k,z) = y(k,M,z), \tag{B.78}$$

with m referring to dark matter; e to the electron density profile; g to the galaxy density; j_ν to the CIB luminosity density at frequency ν; and y to the Compton y parameter. $y(k,M,z)$ is defined in Equation (B.71).Note that central galaxies are always taken to be at the centre of the halo which is why they are not multiplied by a k-dependent factor[2]. The satellite galaxies (and luminosity) are weighted by the dark matter profile $u(k,M,z)$; this is because that the galaxy distribution is modelled as following the dark matter distribution in the halo.

B.4.4 1-halo power spectra

For the 1-halo power spectra, we distinguish between "discrete" observables (galaxies and CIB) and "continuous" observables (everything else). For dark matter, electrons, and Compton y, we have

$$P^{1h}_{mm} = \int dM \frac{dN}{dM}\left(\frac{M}{\rho_m}u(k,M,z)\right)^2 \tag{B.79}$$

$$P^{1h}_{ee} = \int dM \frac{dN}{dM}\left(\frac{M}{\rho_m}u_e(k,M,z)\right)^2 \tag{B.80}$$

$$P^{1h}_{yy} = \int dM \frac{dN}{dM}\left(\frac{4\pi\sigma_T a}{m_e c^2}y(k,M,z)\right)^2, \tag{B.81}$$

These power spectra are all of the form

$$P^{1h}_{XX} = \int dM \frac{dN}{dM} A_X(M,k,z)^2, \tag{B.82}$$

[2]One could also replace $N^{\rm cen} \to N^{\rm cen}u_c(k)$ to take into account central galaxies that are mis-centered. Here, we take $u_c(k)$ to be 1, as in Appendix B of [20]. One could similarly take into account mis-centering in $L^{\rm cen}_{(1+z)\nu}(M,z)$.

and their cross spectra are

$$P_{XY}^{1h} = \int dM \frac{dN}{dM} A_X(M,k,z) A_Y(M,k,z). \tag{B.83}$$

For galaxies we have [20] (in the "maximally correlated" model)

$$P_{gg}^{1h} = \int dM \frac{dN}{dM} \frac{2N^{\rm sat}(M,z)u(k,M,z) + (N^{\rm sat}(M,z)^2/N^{\rm cen})u(k,M,z)^2}{\bar{n}_g(z)^2}. \tag{B.84}$$

The 1-halo power spectrum for CIB at frequencies ν and ν' is

$$\begin{aligned} P_{\nu\nu'}^{1h} &= \int dM \frac{dN}{dM} \frac{1}{(4\pi)^2} \Big(L^{cen}_{\nu(1+z)} L^{sat}_{\nu'(1+z)} u(k,M,z) \\ &+ L^{cen}_{\nu'(1+z)} L^{sat}_{\nu(1+z)} u(k,M,z) + L^{sat}_{\nu(1+z)} L^{sat}_{\nu'(1+z)} u(k,M,z)^2 \Big) \end{aligned} \tag{B.85}$$

Within this paradigm, the cross power spectrum between the CIB at frequency ν and galaxies is

$$\begin{aligned} P_{g\nu}^{1h} &= \int dM \frac{dN}{dM} \frac{1}{4\pi \bar{n}_g(z)} \Big(L^{cen}_{\nu(1+z)} N^{sat}(M,z) u(k,M,z) & \text{(B.86)} \\ &+ L^{sat}_{\nu(1+z)} N^{cen}(M,z) u(k,M,z) + L^{sat}_{\nu(1+z)} N^{sat}(M,z) u(k,M,z)^2 \Big). & \text{(B.87)} \end{aligned}$$

The cross power-spectra of the "continuous" and the "discrete" observables is:

$$P_{XY}^{1h} = \int dM \frac{dN}{dM} A_X(M,k,z) A_Y(M,k,z). \tag{B.88}$$

B.4.5 Poissonian noise

In all galaxy-galaxy, CIB-CIB, and galaxy-CIB power spectra, we must also include the scale-independent Poissonian noise (or shot noise).

$$C_\ell^{gg,{\rm shot}}(z) = \frac{1}{\bar{n}(z)} \tag{B.89}$$

where $\bar{n}(z)$ is the total galaxy number density in the map in a redshift bin. For the CIB, the shot noise is

$$C_\ell^{\nu\nu,{\rm shot}} = \int dS_\nu \frac{dN}{dS_\nu} S_\nu^2 \tag{B.90}$$

where S_ν represents flux measured at frequency ν.

Appendix C

C.1 Constraints on model parameters

We present here the constraints on model parameters derived using Fisher analysis of different combinations of primary CMB and remote dipole and quadrupole fields. We include the constraint using only the primary CMB temperature as well. Characteristic scale parameters k_c, k_b and k_s are in units of Mpc^{-1}.

Dipolar modulation model

Parameter	Noise [μ-Karcmin]	$\sigma(T)$	$\sigma(T,E)$	$\sigma(T,E,R)$	$\sigma(T,E,G)$	σ(All)
$10^2A = 7.00$	5.0	1.93	1.04	1.03 - 0.92	0.85	0.85 - 0.85
	1.0	1.93	1.04	0.98 - 0.87	0.85	0.84 - 0.83
	0.1	1.93	1.04	0.92 - 0.86	0.85	0.84 - 0.83
$10^3k_c = 7.83$	5.0	6.36	4.50	4.40 - 3.46	2.76	2.76 - 2.73
	1.0	6.36	4.50	4.08 - 3.05	2.76	2.74 - 2.69
	0.1	6.36	4.50	3.58 - 2.97	2.76	2.73 - 2.67
$10\Delta \ln k = 5.00$	5.0	0.17	0.12	0.10 - 0.05	0.04	0.04 - 0.04
	1.0	0.17	0.12	0.08 - 0.04	0.04	0.04 - 0.03
	0.1	0.17	0.12	0.05 - 0.04	0.04	0.04 - 0.03

Table C.1: Parameter constraints for the dipolar modulation model 4.4.1.

Exponential suppression model

Parameter	Noise [μ-Karcmin]	$\sigma(T)$	$\sigma(T,E)$	$\sigma(T,E,R)$	$\sigma(T,E,G)$	σ(All)
$10^4 k_c = 3.74$	5.0	3.05	2.33	2.23 - 2.21	2.04	2.03 - 1.94
	1.0	3.05	2.33	2.19 - 2.03	2.04	1.98 - 1.88
	0.1	3.05	2.33	2.09 - 1.90	2.04	1.92 - 1.80
$\lambda = 0.53$	5.0	0.25	0.18	0.17 - 0.15	0.15	0.14 - 0.14
	1.0	0.25	0.18	0.16 - 0.15	0.15	0.14 - 0.13
	0.1	0,25	0.18	0.15 - 0.14	0.15	0.14 - 0.13

Table C.2: Parameter constraints for the exponential suppression model 4.4.2.

Broken power law model

Parameter	Noise [μ-Karcmin]	$\sigma(T)$	$\sigma(T,E)$	$\sigma(T,E,R)$	$\sigma(T,E,G)$	σ(All)
$10^4 k_b = 5.26$	5.0	5.03	2.89	2.79 - 2.74	2.60	2.55 - 2.51
	1.0	5.03	2.89	2.75 - 2.68	2.60	2.52 - 2.47
	0.1	5.03	2.89	2.70 - 2.54	2.60	2.49 - 2.45
$\delta = 1.14$	5.0	2.85	1.59	1.50 - 1.47	1.36	1.32 - 1.29
	1.0	2.85	1.59	1.48 - 1.43	1.36	1.30 - 1.28
	0.1	2.85	1.59	1.44 - 1.34	1.36	1.28 - 1.26

Table C.3: Parameter constraints for the broken power law model 4.4.2.

Cutoff model

Parameter	Noise [μ-Karcmin]	$\sigma(T)$	$\sigma(T,E)$	$\sigma(T,E,R)$	$\sigma(T,E,G)$	σ(All)
$10^4 k_c = 3.63$	5.0	0.78	0.38	0.37 - 0.32	0.30	0.30 - 0.29
	1.0	0.78	0.38	0.34 - 0.30	0.30	0.29 - 0.28
	0.1	0.78	0.38	0.32 - 0.28	0.30	0.28 - 0.27

Table C.4: Parameter constraints for the cutoff model 4.4.2.

Step model

Parameter	Noise [μ-Karcmin]	$\sigma(T)$	$\sigma(T,E)$	$\sigma(T,E,R)$	$\sigma(T,E,G)$	σ(All)
$10\mathcal{A}_s = 3.74$	5.0	2.80	1.11	1.05 - 0.91	0.93	0.91 - 0.85
	1.0	2.80	1.11	0.97 - 0.85	0.93	0.88 - 0.81
	0.1	2.80	1.11	0.90 - 0.82	0.93	0.85 - 0.80
$10^4 k_s = 7.94$	5.0	0.72	0.18	0.17 - 0.14	0.14	0.14 - 0.13
	1.0	0.72	0.18	0.16 - 0.13	0.14	0.14 - 0.13
	0.1	0.72	0.18	0.14 - 0.13	0.14	0.13 - 0.13
$x_s = 1.41$	5.0	0.60	0.25	0.24 - 0.20	0.19	0.19 - 0.18
	1.0	0.60	0.25	0.22 - 0.18	0.19	0.19 - 0.17
	0.1	0.60	0.25	0.20 - 0.18	0.19	0.18 - 0.17

Table C.5: Parameter constraints for the step model 4.4.2.

C.2 Mode coupling

The super-horizon modulating field $h(\vec{\boldsymbol{x}})$ introduced in the spontaneous isotropy breaking mechanism of Sec.4.4.1 leads to couplings between different multipole moments. The modified primordial spectrum Eq. 4.17 is used to compute the covariance matrix, which differs from the ΛCDM covariance matrix by terms linear and quadratic in the modulation amplitude A:

$$C^{X,Y}_{\alpha\beta,\ell\ell',mm'} = C^{X,Y,(\Lambda\mathrm{CDM})}_{\alpha\beta,\ell\ell',mm'} + C^{X,Y,(A)}_{\alpha\beta,\ell\ell',mm'} + C^{X,Y,(A^2)}_{\alpha\beta,\ell\ell',mm'}$$

$$C^{X,Y,(A)}_{\alpha\beta,\ell\ell',mm'} = \delta_{mm'}\sqrt{\frac{4\pi}{3}}\frac{A}{i\chi_{dec}}\int\frac{dk\,k^2}{(2\pi)^3}P_\psi(k)\Big[\Delta^{*X,\alpha}_\ell(k)\,\partial_k\Delta^{Y,\beta}_{\ell'}(k) - \partial_k\Delta^{*X,\alpha}_\ell(k)\,\Delta^{Y,\beta}_{\ell'}(k)$$
$$- \frac{2\,\Delta^{*X,\alpha}_\ell(k)\,\Delta^{Y,\beta}_{\ell'}(k)}{k}\Big(\ell\;\delta_{l',l-1} - (\ell+1)\;\delta_{l',l+1}\Big)\Big]R^{1\ell'}_{\ell m}$$

$$C^{X,Y,(A^2)}_{\alpha\beta,\ell\ell',mm'} = \delta_{mm'}\frac{4\pi}{3}\frac{A^2}{\chi^2_{dec}}\int\frac{dk\,k^2}{(2\pi)^3}P_\psi(k)\sum_L R^{1L}_{\ell m}R^{1L}_{\ell' m}\Big[\partial_k\Delta^{*X,\alpha}_\ell(k)\,\partial_k\Delta^{*Y,\beta}_{\ell'}(k)$$
$$+ \frac{\partial_k\Delta^{*X,\alpha}_\ell(k)\,\Delta^{Y,\beta}_{\ell'}(k)}{k}\Big((1+\ell')\delta_{L,\ell'-1} - \ell'\delta_{L,\ell'+1}\Big)$$
$$+ \frac{\Delta^{*X,\alpha}_\ell(k)\,\partial_k\Delta^{Y,\beta}_{\ell'}(k)}{k}\Big((1+\ell)\delta_{L,\ell-1} - \ell\delta_{L,\ell+1}\Big)$$
$$+ \frac{\Delta^{*X,\alpha}_\ell(k)\,\Delta^{Y,\beta}_{\ell'}(k)}{k^2}\Big((1+\ell)^2\,\delta_{L,\ell-1}\,\delta_{\ell',\ell} + \ell^2\,\delta_{L,\ell+1}\,\delta_{\ell',\ell}$$
$$- (1+\ell)(\ell-2)\,\delta_{L,\ell-1}\,\delta_{\ell',\ell-2} - \ell(\ell+3)\,\delta_{L,\ell+1}\,\delta_{\ell',\ell+2}\Big)\Big]$$

where the couplings $R^{\ell_1\ell_2}_{\ell m}$ are defined through the 3-j Wigner symbols

$$R^{\ell_1\ell_2}_{\ell m} = (-1)^m\sqrt{\frac{(2\ell+1)(2\ell_1+1)(2\ell_2+1)}{4\pi}}\begin{pmatrix}\ell_1 & \ell_2 & \ell\\ 0 & 0 & 0\end{pmatrix}\begin{pmatrix}\ell_1 & \ell_2 & \ell\\ 0 & m & -m\end{pmatrix}. \tag{C.1}$$

The term linear in A induces couplings between multipoles ℓ and $\ell\pm1$, while the quadratic term adds couplings between ℓ and $\ell\pm2$ as well as the same multipole corrections

to the covariance matrix. For the particular case of temperature transfer functions in the Sachs-Wolfe approximation, i.e. $\Delta_\ell^T(k) \propto j_\ell(k\chi_{dec})$, the above expressions can be reduced to those presented in [160, 132] using appropriate recursion relations for the derivatives of the spherical Bessel functions. A similar approach to compute the multipole couplings in terms of derivatives of the transfer functions was taken in reference [193], although the assumptions on the modulation of the primordial spectrum are not the same as ours and the $\mathcal{O}(A^2)$ term was not computed.